
Representation and Inference for Natural Language

A First Course in Computational Semantics

Volume I
Working with First-Order Logic

Patrick Blackburn & Johan Bos

September 3, 1999

Contents

Preface	iii
1 First-Order Logic	1
1.1 First-Order Logic	1
1.2 A Simple Model Checker	14
1.3 Some Refinements	21
1.4 First-Order Logic and Natural Language	28
2 Lambda Calculus	29
2.1 Compositionality	29
2.2 Two Experiments	33
2.3 The Lambda Calculus	40
2.4 Implementing Lambda Calculus	46
2.5 Grammar Engineering	54
3 Underspecified Representations	63
3.1 Scope Ambiguities	63
3.2 Montague's Approach	65
3.3 Storage Methods	70
3.4 Underspecification	83
4 Propositional Inference	95
4.1 Propositional Tableaux	97

4.2	Implementing Propositional Tableaux	108
4.3	Equality Constraints	112
4.4	Theoretical Remarks	112
5	First-Order Inference	117
5.1	First-Order Tableaux	117
5.2	Unification and Free Variable Tableaux	122
5.3	Implementing Free Variable Tableaux	132
5.4	Off-the-shelf Inference Tools	137
6	Putting It All Together	141
6.1	Natural Language Questions	141
6.2	Natural Language Argumentation	144
6.3	Building Ontologies	147
6.4	Logical Redundancies	155
	Afterword: Where to from here?	157
	Bibliography	159
A	Propositional Languages	163
B	Type Theory	165
C	Theorem Provers and Model Builders	169
D	Prolog in a Nutshell	171
E	Listing of Programs	189

Preface

This book introduces a number of fundamental techniques for computing semantic representations for fragments of natural language and performing inference with the result. Both the underlying theory and their implementation in Prolog are discussed. We believe that the reader who masters these techniques will be in a good position to appreciate (and critically assess) ongoing developments in computational semantics.

Computational semantics is a relatively new subject, and trying to define such a lively area (if indeed it is a single area) seems premature, even counterproductive. However, in this book we take ‘semantics’ to mean ‘formal semantics’ (that is, the business of giving model-theoretic interpretations to fragments of natural language, usually with the help of some intermediate level of logical representation) and ‘computational semantics’ to be the business of using a computer to actually build such representations (*semantic construction*) and reason with the result (*inference*). Thus this book is devoted to introducing techniques for tackling the following two questions:

1. *How can we automate the process of associating semantic representations with expressions of natural language?*
2. *How can we use logical representations of natural language expressions to automate the process of drawing inferences?*

Two comments should immediately be made. First, we’ve presented semantic construction and inference as if they were independent, but of course they’re not. Indeed, how semantic construction and inference can best be interleaved is a deep and important problem. We won’t solve it, but by the end of the book we will have seen some nice examples of semantic construction and inference harnessed together in a single architecture.

Second, our working definition of computational semantics isn’t quite as innocent as it looks. Many semanticists claim that intermediate levels of logical representation are essentially redundant. Whether or not this is true, the move to a computational perspective on formal semantics certainly increases the *practical* importance of the representation level. Logical representations—that is, formulas of a logical language—encapsulate meaning in a clean compact way. They make it possible to use well understood proof systems to

perform inference, and we shall learn how to exploit this possibility. Models may be the heart of semantics, but representations are central to its computational cousin. However, as will become clear in the course of the book, we feel that the computational perspective vividly brings out *theoretical* importance of representations. The success of Discourse Representation Theory, the explosion of interest in underspecification, and the exploration of alternative glue languages such as linear logic to drive the process of semantic construction all bear witness to an important lesson: semanticists ignore representations at their peril. Representations are well-defined mathematical entities that (among other things) can be formed into more abstract abstract entities, explored geometrically, and specified indirectly with the aid of constraints; we'll see examples of all these things in this book. Even if representations are eliminable, they need to be taken seriously.

Chapter 1. First-Order Logic. Here we introduce the syntax and semantics of first order logic (the representation language used in the first half of the book), show how formulas and (simple finite) models may be represented in Prolog, build a simple model checker, and discuss the problems that arise when Prolog variables are used to represent first-order variables.

Chapter 2. Lambda Calculus. Here we start studying semantic construction. We outline the methodology underlying our work (namely, *compositionality*) and motivate our use of Prolog's Definite Clause Grammars. We then write two rather naive programs that build semantic representations for a very small fragment of English. These experiments lead us to the lambda calculus, the tool that drives this book's approach to semantic construction. We implement β -conversion, the computational core of the lambda calculus, in two different ways, and then integrate it into the grammatical architecture that will be used throughout the book.

Chapter 3. Underspecified Representations. Here we investigate a fundamental problem for computational semantics: scope ambiguities. These are semantic ambiguities that can arise in syntactically unambiguous expressions, and they pose a problem for compositional approaches to semantic construction. We illustrate the problem, and present four (increasingly more sophisticated) solutions: Montague's use of quantifier raising, Cooper storage, Keller storage, and a more recent underspecification-based method called Hole Semantics. We implement all four approaches.

Chapter 4. Propositional Inference. Here we turn to the second major theme of the book: inference. Our approach to inference will be based on first-order theorem proving and model building, and in this chapter we start developing the necessary tools. We introduce a *signed tableaux* system for

propositional calculus, show how to implement it in Prolog, discuss some extensions, and conclude with an informal discussion of a number of theoretical issues.

Chapter 5. First-Order Inference. Here we extend our propositional signed tableaux system to full-first order logic—and immediately run into a problem. Naive approaches to first-order tableaux are doomed to hopeless inefficiency. So we rethink the problem. Instead of asking the tableaux system to solve the entire inference problem, we reformulate it so that it generates a collection of constraints on possible solutions. We then feed these constraints to a *unification* component, which (hopefully) will compute a useful solution. As we shall see, this new approach lends itself to computational implementation.

Chapter 6. Putting It All Together. In this chapter we combine the techniques we have learned so far. By plugging together our lambda calculus, quantifier storage, model checking, and theorem proving programs we are able (more or less immediately) to define simple question handling and argumentation predicates. Does Vincent give every woman a foot massage? Which robber loves Honey Bunny? Find out in this week’s exciting episode . . .

Each chapter concludes with two sections: *Software Summary* lists the programs developed in the chapter, and *Notes* lists references the reader may find helpful, and discusses more advanced topics.

We have tried to make this book self-contained, and believe that readers with relatively modest backgrounds in linguistics and logic should be able to follow the theoretical discussion. But the fact remains that this is a book in *computational* semantics, and getting the most out of it requires a working knowledge of Prolog. We have included an appendix which outlines the main ideas of Prolog, and we hope that this will be useful—but, bluntly, the reader who has never thought about computational problems from the declarative perspective typical of logic or constraint-based programming will find parts of it tough going. We develop many theoretical ideas by thinking problems through from a declarative perspective. For example, our account of the lambda calculus in Chapter 2 makes no appeal to types, function valued functions, or higher-order logic—rather, lambda calculus is presented as a beautiful piece of data-abstraction that emerges naturally from a declarative analysis of semantic construction. Computational ideas are not an optional extra in computational semantics; they are the heart of the enterprise.

This book developed out of material for a course on Computational Semantics we regularly offer at the Department of Computational Linguistics, University of the Saarland. We also taught a preliminary version (essentially the material that now makes up Part I) as an introductory course at ESSLLI’97, the *European Summer School in Logic, Language and Information* held at Aix-en-Provence, France in August 1997. When designing these

courses, we found no single source which contained all the material we wanted to present. At that time, the only notes solely devoted to computational semantics we knew of were Cooper et al. 1993. These notes, which we recommend to our readers, were developed at the University of Edinburgh, and are probably the first systematic introduction to modern computational semantics. Like the present book they are Prolog based, and cover some of the same ground using interestingly different tools and techniques. However we wanted to teach the subject in a way that emphasized such ideas as inference, underspecification, and architectural issues. This led us to a first version of the book, which was heavily influenced by Pereira and Shieber 1987 for semantic construction, and Fitting 1996 and Smullyan 1995 for tableaux systems. Since then, the project has taken on a life of its own, and grown in a variety of (often unexpected) directions. Both the code and the text has been extensively rewritten and we are now (we hope!) in the final stretch of producing the kind of introduction to computational semantics that we wanted all along.

Acknowledgments

We would like to thank Manfred Pinkal and all our colleagues at the Department of Computational Linguistics, University of Saarland, Saarbrücken, Germany. Thanks to Manfred, the department has become an extremely stimulating place for working on computational semantics; it was the ideal place to write this book.

Johan Bos would like to thank his colleagues in the *Verbmobil* Project (grant 01 IV 701 R4 and 01 IV 701 N3) at IMS-Stuttgart, TU-Berlin, CSLI Stanford, the DFKI, and the Department of Computational Linguistics at Saarbrücken, and those in the Trindi Project (Task Oriented Instructional Dialogue; LE4-8314) at the University of Gothenburg, SRI Cambridge, University of Edinburgh, and Xerox Grenoble.

Conversations with Paul Dekker, Andreas Franke, Martin Kay, Hans Kamp, Michael Kohlhase, Karsten Konrad, Christof Monz, Reinhard Muskens, Maarten de Rijke, and Henk Zeevat were helpful in clarifying our goals. We're grateful to David Beaver, Judith Baur, Björn Gambäck, Ewan Klein, Emiel Kraemer, Rob van der Sandt, and Frank Schilder for their comments on an early draft of this book. We're also very grateful to all our students, who have provided us with invaluable feedback. In particular we would like to thank Aljoscha Burchardt, Gerd Fliedner, Malte Gabsdil, Kristina Striegnitz, Stefan Thater, and Stephan Walter. Patrick Blackburn would like to thank Aravind Joshi and the staff of Institute for Research in Cognitive Science at the University of Pennsylvania for their hospitality while the Montreal version was being prepared.

Finally, we are deeply grateful to Robin Cooper, who taught a course based on the previous draft of this book, and provided us with extremely detailed feedback on what worked and what didn't. His comments have greatly improved the book.

*Patrick Blackburn and Johan Bos,
Computerlinguistik, Universität des Saarlandes,
September 1999.*

Chapter 1

First-Order Logic

First-order logic is the formalism used in the first half of this book to represent the meaning of natural language sentences. In this chapter we introduce first-order logic, write some simple Prolog programs that work with it, and discuss the role of first-order logic in natural language semantics.

Here's what we'll do. First, we'll discuss the syntax and semantics of first-order logic. Next, we'll write a simple first-order model checker (or semantic evaluator) in Prolog. The model checker takes a first-order formula and a (special kind of) first-order model as input and checks whether the formula is satisfied in the model. As we shall see, such a model checker is extremely simple to implement if we use Prolog variables to simulate first-order variables, but this strategy leads to problems if carried out naively; we discuss these problems and show how to deal with them. We conclude with a discussion of the relevance of first-order logic to natural language semantics.

1.1 First-Order Logic

In this section we review the syntax and semantics of first-order logic. We discuss *vocabularies*, *first-order models* and *first-order languages*, tie them together via the crucial *satisfaction* definition, define the key logical concepts of *validity* and *logical consequence*, and conclude with a discussion of *function symbols* and *equality*.

Vocabularies

Our ultimate goal is to define how first-order formulas (that is, certain kinds of descriptions) are evaluated in first-order models (that is, certain kinds of situations). Simplifying somewhat, the purpose of the evaluation process is to tell us whether a description is true or false in a situation.

We shall shortly do this—but we need to exercise a little care. Intuitively, it doesn't make much sense to ask whether or not an arbitrary description is true in an arbitrary situation. Some descriptions and situations simply don't belong together. For example, if we are given a formula (that is, a description) from a first-order language intended for talking about the various relationships and properties (such as *loving*, *being a robber*, and *being a customer*) that hold between Mia, Honey Bunny, Vincent, and Pumpkin, and we are given a model (that is, a situation) which records information about something completely different (for example, which household cleaning products are best at getting rid of particularly nasty stains) then it doesn't really make much sense to evaluate this particular formula in that particular model. *Vocabularies* allow us to avoid such problems: they tell us which first-order languages and models belong together.

Here is our first vocabulary:

$$\{ \begin{array}{l} (\text{LOVE},2), \\ (\text{CUSTOMER},1), \\ (\text{ROBBER},1), \\ (\text{MIA},0), \\ (\text{VINCENT},0), \\ (\text{HONEY-BUNNY},0), \\ (\text{PUMPKIN},0) \end{array} \}$$

Intuitively, this vocabulary is telling us two important things: the topic of conversation, and the language the conversation is going to be conducted in. Let's spell out this a little.

First, the vocabulary tells us *what* we're going to be talking about. In the present case, we're going to be talking about *loving* (the 2 indicates that loving is taken to be a two place relation) and the properties (or 1-place relations) of *being a customer* and *being a robber*. In addition to these relations we're going to be talking about four special entities named *Mia*, *Honey Bunny*, *Vincent*, and *Pumpkin* (the 0s indicate that these symbols are constants, or names).

Second, the vocabulary also tells us *how* we can talk about these things. In the above case it tells us that we will be using a symbol LOVE of arity 2 (that is, a 2-place symbol) for talking about loving, two symbols of arity 1 (CUSTOMER and ROBBER) for talking about customers and robbers, and four constant symbols (or names), namely MIA, VINCENT, HONEY-BUNNY, and PUMPKIN for naming certain entities of special interest.

In short, a vocabulary gives us all the information needed to define the class of models of interest (that is, the kinds of situations we want to describe) and the relevant first-order language (that is, the kinds of descriptions we can use). So let's now look at what first-order models and languages actually are.

First-Order Models

Suppose we've fixed some vocabulary. What should a first-order *model* for this vocabulary be?

Actually, our previous discussion has pretty much given the answer. Intuitively a model is a situation. That is, it is a *semantic* entity: it contains the kinds of things we want to talk about. Thus a model for a given vocabulary gives us two pieces of information. First, it tells us which collection of entities we are talking about; this collection is usually called the *domain*. Second, for each symbol in the vocabulary, it gives us an appropriate semantic entity, built from the items in D . This task is carried out by a function F which specifies, for each symbol in the vocabulary, an appropriate semantic value; we call such functions *interpretation functions*. Thus, in set theoretic terms, a model \mathbf{M} is an ordered pair (D, F) consisting of a domain D and an interpretation function F specifying semantic values in D .

What are appropriate semantic values? There's no mystery here. As constants are names, each constant should be interpreted as an element of D . (That is, for each constant symbol c in the vocabulary, $F(c) \in D$.) As n -place relation symbols are intended to denote n -place relations, each n -place relation symbol R should be interpreted as an n -place relation on D . (That is, $F(R)$ should be a set of n -tuples of elements of D .)

Let's consider an example. We shall define a simple model for the vocabulary given above. Let D be $\{d_1, d_2, d_3, d_4\}$. That is, this four element set is the domain of our little model.

Next, we must specify an interpretation function F . Here's one possibility:

$$\begin{aligned} F(\text{MIA}) &= d_1 \\ F(\text{HONEY-BUNNY}) &= d_2 \\ F(\text{VINCENT}) &= d_3 \\ F(\text{PUMPKIN}) &= d_4 \\ F(\text{CUSTOMER}) &= \{d_1, d_3\} \\ F(\text{ROBBER}) &= \{d_2, d_4\} \\ F(\text{LOVE}) &= \{(d_4, d_2), (d_3, d_1)\} \end{aligned}$$

Note that every symbol in the vocabulary does indeed correspond to an appropriate semantic entity. The four names pick out individuals, the two arity 1 symbols pick out subsets of D (that is, properties, or 1-place relations on D) and the arity 2 symbol picks out a 2-place relation on D . Intuitively, in this model d_1 is called Mia, d_2 is called Honey Bunny, d_3 is called Vincent and d_4 is called Pumpkin. Both Honey Bunny and Pumpkin are robbers, while both Vincent and Mia are customers. Pumpkin loves Honey Bunny and Vincent loves Mia. Sadly, Honey Bunny does not love Pumpkin, Mia does not love Vincent, and nobody loves themselves.

Here's a second model for the same vocabulary. We'll use the same domain (that is,

$D = \{d_1, d_2, d_3, d_4\}$) but change the interpretation function. To emphasize that the interpretation function has changed, we'll use a different symbol (namely F_2) for it.

$$\begin{aligned} F_2(\text{MIA}) &= d_2 \\ F_2(\text{HONEY-BUNNY}) &= d_1 \\ F_2(\text{VINCENT}) &= d_4 \\ F_2(\text{PUMPKIN}) &= d_3 \\ F_2(\text{CUSTOMER}) &= \{d_1, d_2, d_4\} \\ F_2(\text{ROBBER}) &= \{d_3\} \\ F_2(\text{LOVE}) &= \{(d_3, d_4)\} \end{aligned}$$

In this model, three of the individuals are customers, only one is a robber, and Pumpkin loves Vincent.

One point is worth emphasizing. Both models just defined are rather special, for they have the following property: *every entity in D is named by exactly one constant*. We shall call any model with this property an *exact* model. Now, exact models are particularly easy to work with, but it is important to realize that not all models are exact. For example, suppose we add two more entities to the domain of the first model; that is, we create a new domain $D' = D \cup \{d_5, d_6\}$ say. Suppose that we then define a new interpretation function F' on D' as follows:

$$\begin{aligned} F'(\text{MIA}) &= d_2 \\ F'(\text{HONEY-BUNNY}) &= d_1 \\ F'(\text{VINCENT}) &= d_4 \\ F'(\text{PUMPKIN}) &= d_3 \\ F'(\text{CUSTOMER}) &= \{d_1, d_2, d_4, d_5\} \\ F'(\text{ROBBER}) &= \{d_3, d_6\} \\ F'(\text{LOVE}) &= \{(d_3, d_4), (d_6, d_5)\} \end{aligned}$$

This new model is like our very first model, except that it has two extra entities. Neither of these new entities has a name, but one of them is a customer, one of them is a robber, and the unnamed robber loves the unnamed customer. This *is* a perfectly good first-order model; there is no requirement that every entity in the model must have a name (we only bother to name entities of special interest). Similarly, there is no requirement that each entity in a model must be named by at most one constant. For example, if we wanted to we could give some entity both the name Vincent and Honey Bunny (perhaps one of these names is a nickname).

First-Order Languages

Given some vocabulary, we build *the first-order language over that vocabulary* out of the following ingredients:

1. All the symbols in the vocabulary. We call these symbols the *non-logical* symbols of the language.
2. A countably infinite collection of variables x, y, z, w, \dots , and so on.
3. The Boolean connectives \neg (negation), \rightarrow (implication), \vee (disjunction), and \wedge (conjunction).
4. The quantifiers \forall (the universal quantifier) and \exists (the existential quantifier).
5. The round brackets $)$ and $($ and the comma. (These are essentially punctuation marks; they are used to group symbols.)

Items 2–5 are common to all first-order languages: the only thing that distinguishes one first-order language from another is the choice of non-logical symbols (that is, the choice of vocabulary).

So, suppose we've chosen some vocabulary. How do we mix these ingredients together? That is, what is the *syntax* of first-order languages? First of all, we define a first-order *term* τ to be any constant or any variable. Roughly speaking, terms are the noun phrases of first-order languages: constants can be thought of as first-order analogs of proper names, and variables as first-order analogs of pronouns.

We can then combine our 'noun phrases' with our 'predicates' (that is, the various relation symbols in the vocabulary) to form *atomic formulas*:

If R is a relation symbol of arity n , and τ_1, \dots, τ_n are terms, then $R(\tau_1, \dots, \tau_n)$ is an atomic (or basic) formula.

Intuitively, an atomic formula is the first-order counterpart of a natural language sentence consisting of a single clause (that is, a simple sentence). The intended meaning of $R(\tau_1, \dots, \tau_n)$ is that the entities named by the terms τ_1, \dots, τ_n stand in the relation named by the symbol R . For example

LOVE(PUMPKIN,HONEY-BUNNY)

means that the entity named PUMPKIN stands in the relation denoted by LOVE to the entity named HONEY-BUNNY—or more simply, that Pumpkin loves Honey Bunny.

Now that we know how to build atomic formulas, we can define more complex descriptions. The following inductive definition tells us exactly which *well formed formulas* (or *wffs*, or simply *formulas*) we can form.

1. All atomic formulas are wffs.

2. If ϕ and ψ are wffs then so are $\neg\phi$, $(\phi \rightarrow \psi)$, $(\phi \vee \psi)$, and $(\phi \wedge \psi)$.
3. If ϕ is a wff, and x is a variable, then both $\exists x\phi$ and $\forall x\phi$ are wffs. (We call ϕ the *matrix* of such wffs.)
4. Nothing else is a wff.

Roughly speaking, formulas built using \neg , \rightarrow , \vee and \wedge correspond to natural language expressions of the form **it is not the case that ...**, **if ... then ...**, **... or ...**, and **... and ...**, respectively. (We discuss the strengths and weaknesses of these correspondences at the end of the chapter.) First-order formulas of the form $\exists x\phi$ and $\forall x\phi$ correspond to natural language expressions of the form **some... or all ...**.

In what follows, we occasionally talk of *subformulas*. The subformulas of a formula ϕ are ϕ itself and all the formulas used to build ϕ . For example, the subformulas of

$$\neg\forall y \text{ PERSON}(y)$$

are $\text{PERSON}(y)$, $\forall y \text{ PERSON}(y)$, and $\neg\forall y \text{ PERSON}(y)$. We leave it to the interested reader to give an inductive definition of subformulthood, and turn to a more important topic: the distinction between *free* and *bound* variables.

Consider the following formula:

$$\neg (\text{CUSTOMER}(x) \vee \forall x(\text{ROBBER}(x) \wedge \forall y \text{ PERSON}(y)))$$

The first occurrence of x is *free*. The second and third occurrences of x are *bound*; they are bound by the first occurrence of the quantifier \forall . The first and second occurrences of the variable y are also bound; they are bound by the second occurrence of the quantifier \forall . Here's the full definition:

1. Any occurrence of any variable is free in any atomic formula.
2. No occurrence of any variable is bound in any atomic formula.
3. If an occurrence of any variable is free in ϕ or in ψ , then that same occurrence is free in $\neg\phi$, $(\phi \rightarrow \psi)$, $(\phi \vee \psi)$, and $(\phi \wedge \psi)$.
4. If an occurrence of any variable is bound in ϕ or in ψ , then that same occurrence is bound in $\neg\phi$, $(\phi \rightarrow \psi)$, $(\phi \vee \psi)$, $(\phi \wedge \psi)$. Moreover, that same occurrence is bound in $\forall y\phi$ and $\exists y\phi$, for any choice of variable y .
5. In any formula of the form $\forall y\phi$ or $\exists y\phi$ (here y can be any variable at all) the occurrence of y that immediately follows the initial quantifier symbol is bound.

6. If an occurrence of a variable x is free in ϕ , then that same occurrence is free in $\forall y\phi$ and $\exists y\phi$, for any variable y distinct from x . On the other hand, all occurrences of x that are free in ϕ , are bound in $\forall x\phi$ and in $\exists x\phi$.

If a formula contains no occurrences of free variables then we call it a *sentence*.

Although they are both called variables, free and bound variables are really very different. (In fact, some formulations of first-order logic use two distinct kinds of symbol for what we have lumped together under the heading ‘variable’.) Here’s an analogy. Try thinking of a free variable as something like the pronoun *she* in

She even has a stud in her tongue.

Uttered in isolation, this would be somewhat puzzling, as we don’t know who *she* refers to. But of course, such an utterance would be made in an appropriate *context*. This context might be either non-linguistic (for example, the speaker might be pointing to a heavily tattooed biker, in which case we would say that *she* was being used *deictically* or *demonstratively*) or linguistic (perhaps the speaker’s previous sentence was **Honey Bunny** is heavily into body piercing, in which case the name **Honey Bunny** supplies a suitable anchor for an *anaphoric* interpretation of *she*).

What’s the point of the analogy? Just as the pronoun *she* required something else (namely, contextual information) to supply a suitable referent, so will formulas containing free variables. Simply supplying a model won’t be enough; we need additional information on how to link the free variables to the entities in the model.

Sentences, on the other hand, are relatively self-contained. For example, consider the sentence $\forall x \text{ROBBER}(x)$. This is a claim that *every* individual is a robber. Roughly speaking, the bound variable x in $\text{ROBBER}(x)$ acts as a sort of placeholder. In fact, the use of x is completely arbitrary; the sentence $\forall y \text{ROBBER}(y)$ means exactly the same thing. Both sentences are simply a way of stating that no matter what entity we take the second occurrence of x (or y) as standing for, that entity will be a robber. In any model of appropriate vocabulary this sentence (and in fact *any* sentence over that vocabulary) will either be true or false.

Our discussion of the interpretation of first-order languages in first-order models will make these distinctions precise (indeed, most of the real work involved in interpreting first-order logic centers on the correct handling of free and bound variables). But before turning to semantic issues, one final remark. In what follows, we won’t always keep to the official first-order syntax defined above. In particular, we’ll generally try and use as few brackets as possible, as this tends to improve readability. For example, we would rarely write

(CUSTOMER(VINCENT) \wedge ROBBER(PUMPKIN))

which is the official syntax. Instead, we would (almost invariably) drop the outermost brackets and write

$$\text{CUSTOMER}(\text{VINCENT}) \wedge \text{ROBBER}(\text{PUMPKIN})$$

To help further reduce the bracket count, we assume the following precedence conventions for the Boolean connectives: \neg binds more tightly than \vee and \wedge , both of which in turn bind more tightly than \rightarrow . What this means, for example, is that the formula

$$\forall x (\neg \text{CUSTOMER}(x) \wedge \text{ROBBER}(x) \rightarrow \text{ROBBER}(x))$$

is shorthand for the following:

$$\forall x ((\neg \text{CUSTOMER}(x) \wedge \text{ROBBER}(x)) \rightarrow \text{ROBBER}(x))$$

In addition, we sometimes use the square brackets $]$ and $[$ as well as the official round brackets, as this can make the intended grouping of symbols easier to grasp visually.

The Satisfaction Definition

Given a model of appropriate vocabulary, a sentence such as $\forall x \text{ROBBER}(x)$ is either true or false in that model. To put it more formally, there is a relation called *truth* which holds, or does not hold, between sentences and models of the same vocabulary. Now it is often obvious how to check whether a given sentence is true in a given model (for example, to check the truth of $\forall x \text{ROBBER}(x)$ we simply need to check that every individual in the model is a robber). What is not so clear is how to give a precise definition of this relation for arbitrary sentences.

Note that we cannot give a *direct* inductive definition of truth, for the matrix of a quantified sentence typically won't be a sentence. For example, $\forall x \text{ROBBER}(x)$ is a sentence, but its matrix $\text{ROBBER}(x)$ is not. Thus an inductive truth definition defined solely in terms of sentences couldn't explain why $\forall x \text{ROBBER}(x)$ was true in a model, for there are no sentential subformulas for such a definition to bite on.

Instead we proceed indirectly. We define a three place relation—called *satisfaction*—which holds between a formula, a model, and an *assignment of values to variables*. Given a model $\mathbf{M} = (D, F)$, an assignment of values to variables in \mathbf{M} (or more simply, an *assignment* in \mathbf{M}) is a function g from the set of variables to D . Assignments are a technical device which tell us what the free variables stand for. By making use of assignment functions, we can inductively interpret *arbitrary* formulas in a natural way, and this will make it possible to define the concept of truth for *sentences*.

But before going further, one point is worth stressing: the reader should *not* view assignment functions simply as a technical fix designed to get round the problem of defining truth. Moreover, the reader should *not* think of satisfaction as being a poor relation of truth. If anything, satisfaction, not truth, is the fundamental notion, at least as far as natural language is concerned. Why is this?

The key to the answer is the word *context*. As we said earlier, free variables can be thought of as analogs of pronouns, whose values need to be supplied by context. An assignment of values to variables can be thought of as a (highly idealized) mathematical model of context; it rolls up all the contextual information into one easy to handle unit, specifying a denotation for every free variable. Thus if we want to use first-order logic to model natural language semantics, it is sensible to think in terms of three components: first-order formulas (descriptions), first-order models (situations) and variable assignments (contexts). The idea of assignment-functions-as-contexts is important in contemporary formal semantics; it has a long history and has been explored in a number of interesting directions. We shall periodically return to it in the course of the book, particularly when we discuss Discourse Representation Theory in Part II.

But let's return to the satisfaction definition. Suppose we've fixed our vocabulary. (That is, from now on, when we talk of a model \mathbf{M} , we mean a model of this vocabulary, and whenever we talk of formulas, we mean the formulas built from the symbols in that vocabulary.) We now give two further technical definitions which will enable us to state the satisfaction definition concisely.

First, let $\mathbf{M} = (D, F)$ be a model, let g be an assignment of values to variables in \mathbf{M} , and let τ be a term. Then by the *interpretation* of τ with respect to \mathbf{M} and g is meant $F(\tau)$ if τ is a constant, and $g(\tau)$ if τ is a variable. We denote the interpretation of τ by $I_F^g(\tau)$.

The second idea we need is that of a *variant* of an assignment of values to variables. So, let g be an assignment of values to variables in some model, and let x be a variable. If g' is an assignment of values to variables in the same model, and for all variables y such that $y \neq x$, $g'(y) = g(y)$ then we say that g' is an *x-variant* of g . Variant assignments are the technical tool that allows us to try out new values for a given variable (say x) while keeping the values assigned to all other variables the same.

We are now ready for the satisfaction definition. Let ϕ be a formula, let $\mathbf{M} = (D, F)$ be a model, and let g be an assignment of values to variables in \mathbf{M} . Then the relation $\mathbf{M}, g \models \phi$ (ϕ is satisfied in \mathbf{M} with respect to the assignment of values to variables g) is defined inductively as follows:

$\mathbf{M}, g \models R(\tau_1, \dots, \tau_n)$	<i>iff</i>	$(I_F^g(\tau_1), \dots, I_F^g(\tau_n)) \in F(R)$
$\mathbf{M}, g \models \neg\phi$	<i>iff</i>	not $\mathbf{M}, g \models \phi$
$\mathbf{M}, g \models \phi \wedge \psi$	<i>iff</i>	$\mathbf{M}, g \models \phi$ and $\mathbf{M}, g \models \psi$
$\mathbf{M}, g \models \phi \vee \psi$	<i>iff</i>	$\mathbf{M}, g \models \phi$ or $\mathbf{M}, g \models \psi$
$\mathbf{M}, g \models \phi \rightarrow \psi$	<i>iff</i>	not $\mathbf{M}, g \models \phi$ or $\mathbf{M}, g \models \psi$
$\mathbf{M}, g \models \exists x\phi$	<i>iff</i>	$\mathbf{M}, g' \models \phi$, for some x-variant g' of g
$\mathbf{M}, g \models \forall x\phi$	<i>iff</i>	$\mathbf{M}, g' \models \phi$, for all x-variants g' of g

(Here ‘iff’ is shorthand for ‘if and only if’.) Note the crucial—and indeed, intuitive—role played by the x-variants in the clauses for the quantifiers. For example, what the clause for the existential quantifier boils down to is this: $\exists x\phi$ is satisfied in a given model, with respect to an assignment g , if and only if there is some x-variant g' of g that satisfies ϕ in the model. That is, we have to try to find *some* value for x that satisfies ϕ in the model, while keeping the assignments to all other variables the same.

We can now define what it means for a *sentence* to be true in a model:

A sentence ϕ is true in a model \mathbf{M} if and only if for *any* assignment g of values to variables in \mathbf{M} , we have that $\mathbf{M}, g \models \phi$. If ϕ is true in \mathbf{M} we write $\mathbf{M} \models \phi$

This is an elegant definition of truth that beautifully mirrors the special, self-contained nature of sentences. It hinges on the following observation: *it simply doesn't matter which variable assignment we use to compute the satisfaction of sentences*. Sentences contain no free variables, so the only free variables we will encounter when evaluating one are those produced when evaluating its quantified subformulas (if it has any). But the satisfaction definition tells us what to do with such free variables: simply try out variants of the current assignment and see whether they satisfy the matrix or not. In short, start with whatever assignment you like; the result will be the same. It is reasonably straightforward to make this informal argument precise, and the reader is asked to do so in Exercise 1.1.1.

Still, for all the elegance of the truth definition, satisfaction is the fundamental concept. Not only is satisfaction the technical engine powering the definition of truth, but from the perspective of natural language semantics it is conceptually prior. By making *explicit* the role of variable assignments, it holds up an (admittedly imperfect) mirror to the process of evaluating descriptions in situations while making use of contextual information.

Exercise 1.1.1 We claimed that when evaluating *sentences*, it doesn't matter which variable assignment we start with. Formally, we are claiming that given any *sentence* ϕ and any model \mathbf{M} (of the same vocabulary), and any variable assignments g and g' in \mathbf{M} , then $\mathbf{M}, g \models \phi$ iff $\mathbf{M}, g' \models \phi$. We want the reader to do two things. First, show that the claim is *false* if ϕ is not a sentence but a formula containing free variables. Second, show that the claim is *true* if ϕ is a *sentence*.

Exercise 1.1.2 This exercise shows that free variables and constants are very similar. In particular, if a free variable x and a constant c denote the same individual, we can replace the free variable by the constant without affecting satisfiability. Formally, let $\mathbf{M} = (D, F)$ be a model, let g be an assignment in \mathbf{M} , and suppose that $F(c) = g(x)$. Let ϕ be any formula, and let $\phi[c/x]$ denote the formula obtained by replacing all free occurrences of x in ϕ by c . Then $\mathbf{M}, g \models \phi$ iff $\mathbf{M}, g' \models \phi[c/x]$, where g' is any x -variant of g . It follows that when working with *exact* models, every formula is equivalent to a sentence. Explain why.

Validities and Valid Arguments

In this book we are going to make heavy use of *logical inference*, and by making use of the semantic concepts just introduced we can explain what we mean by this. We do so in two steps. First we say what *valid formulas* (or more simply, *validities*) are, and then we define *valid arguments* (or *valid inferences*).

A valid formula is a formula that is satisfied in *all* models (of the appropriate vocabulary) given *any* variable assignment. That is, if ϕ is a valid formula, it is impossible to find a situation and a context in which ϕ is not satisfied. We indicate that a formula ϕ is valid by writing $\models \phi$.

For example,

$$\models \text{ROBBER}(x) \vee \neg \text{ROBBER}(x).$$

In any model, given any variable assignment, one (and indeed, only one) of the two disjuncts must be true, and hence the whole formula will be satisfied too.

Note that for *sentences* the definition of validity can be rephrased as follows: a valid sentence is a sentence that is true in all models (of the appropriate vocabulary). That is, it is impossible to falsify a valid sentence. For example,

$$\models \forall x(\text{ROBBER}(x) \rightarrow \text{CUSTOMER}(x)) \wedge \text{ROBBER}(\text{MIA}) \rightarrow \text{CUSTOMER}(\text{MIA}).$$

Exercise 1.1.3 This exercise shows that the validity of arbitrary formulas is equivalent to the validity of certain sentences. As a first step, the reader should show that if ϕ is a formula containing x as a free variable, then ϕ is valid iff $\forall x\phi$ is valid. It follows that the validity of formulas is reducible to the validity of sentences. Explain why. [Hint: we've just found a way of reducing number of free variables by one while maintaining validity. Iterate this process.]

Now, validities are clearly in some sense *logical*; they are descriptions which carry a cast-iron guarantee of satisfiability. But logic has traditionally appealed to the more dynamic notion of *valid arguments*, a movement, or inference, from premises to conclusions.

Suppose ϕ_1, \dots, ϕ_n , and ψ are a finite collection of first-order formulas. Then we say that the argument with *premises* ϕ_1, \dots, ϕ_n and *conclusion* ψ is a *valid argument* if and only if whenever all the premises are satisfied in some model, using some variable assignment, then the conclusion is satisfied in the same model using the same variable assignment. The notation

$$\phi_1, \dots, \phi_n \models \psi$$

means that the argument with premises ϕ_1, \dots, ϕ_n and conclusion ψ is valid. There is a wide range of terminology for talking about valid arguments. For example, we sometimes say that ψ is a *valid inference* from the premises ϕ_1, \dots, ϕ_n , or that ψ is a *logical consequence* of ϕ_1, \dots, ϕ_n .

Note that if the premises and conclusion are all *sentences* the definition of valid argument can be rephrased as follows: an argument is valid if whenever the premises are true in some model, the conclusion is too. Or more simply: the truth of the premises guarantees the truth of the conclusion.

Here's an example. The argument with premises $\forall x(\text{ROBBER}(x) \rightarrow \text{CUSTOMER}(x))$ and $\text{ROBBER}(\text{MIA})$ and the conclusion $\text{CUSTOMER}(\text{MIA})$ is valid. That is,

$$\forall x(\text{ROBBER}(x) \rightarrow \text{CUSTOMER}(x)), \text{ROBBER}(\text{MIA}) \models \text{CUSTOMER}(\text{MIA}).$$

As the reader may suspect, there is a connection between the validity of this argument and the fact that

$$\models \forall x(\text{ROBBER}(x) \rightarrow \text{CUSTOMER}(x)) \wedge \text{ROBBER}(\text{MIA}) \rightarrow \text{CUSTOMER}(\text{MIA}).$$

This example suggests that with the help of the Boolean connectives \wedge and \rightarrow we can convert valid arguments into validities. The reader is asked to explore this in Exercise 1.1.5 below.

Validity and valid arguments are the fundamental logical concepts, and they underly the treatment of inference in this book. Both concepts are semantically defined (that is, they are defined in terms of models and variable assignments) but in Chapters 4 and 5 we will find a syntactic way of thinking about them that lends itself to computational implementation (that is, we will develop various *proof systems*). Moreover, as we shall learn in Part II, these concepts provide a useful basis for discussing inference in Discourse Representation Theory.

This concludes our review of the syntax and semantics of first-order logic. We finish with a global remark. As the reader has doubtless observed, most of the ideas we have discussed are fairly straightforward—with the notable exception of variable handling. This required a certain amount of care, and the introduction of concepts such as assignments and variant assignments. Unsurprisingly, if we want to work with first-order logic in Prolog, the key issue that faces us is how to cope with variables. Indeed, the theme of how best to handle variables recurs in various guises throughout the book.

Exercise 1.1.4 We say that two sentences ϕ and ψ are *logically equivalent* if and only if $\phi \models \psi$ and $\psi \models \phi$. Show that $\forall x\phi$ and $\neg\exists x\neg\phi$ are logically equivalent.

Exercise 1.1.5 The *Deduction Theorem* for first-order logic says that $\phi_1, \dots, \phi_n \models \psi$ if and only if $\models (\phi_1 \wedge \dots \wedge \phi_n) \rightarrow \psi$. (That is, there is an intimate link between validities and valid arguments.) Prove the Deduction Theorem.

Exercise 1.1.6 Show that the validity of arguments whose premises or conclusions contain free variables, is reducible to the validity of arguments whose premises and conclusions are sentences. [Hint: think about Exercise 1.1.3.]

Function Symbols and Equality

We now discuss two important extensions of first-order logic; namely first-order logic with function symbols, and first-order logic with equality. Function symbols play an important role when we discuss first-order inference in Chapter 5. Equality plays an important role in our discussion of Discourse Representation Theory in Part II.

Suppose we want to talk about Butch, Butch's father, Butch's grandfather, Butch's great grandfather, and so on. Now, if we know the names of all these people this is easy to do—but what if we don't? A natural solution is to add a 1-place function symbol FATHER to the language. Then if BUTCH is the constant that names Butch, FATHER(BUTCH) is a term that picks out Butch's father, FATHER(FATHER(BUTCH)) picks out Butch's grandfather, and so on. That is, function symbols are a syntactic device that let us form recursively structured terms, thus letting us express many concepts in a natural way.

Let's make this precise. First, we shall suppose that it is the task of the vocabulary to tell us which function symbols we have at our disposal, and what the arity of each of these symbols is. Second, given this information, we say (as before) that a model \mathbf{M} is a pair (D, F) where F interprets the constants and relation symbols as described earlier, and, in addition, F assigns to each function symbol f an appropriate semantic entity. What's an appropriate interpretation for an n -place function symbol? Simply a function that takes n -tuples of elements of D as input, and returns an element of D as output. Third, we need to say what terms we can form using these new symbols. Here's the definition we require:

1. All constants and variables are terms.
2. If f is a function symbol of arity n , and τ_1, \dots, τ_n are terms, then $f(\tau_1, \dots, \tau_n)$ is also a term.
3. Nothing else is a term.

A term is said to be *closed* if and only if it contains no variables.

Only one task remains: interpreting these new terms. In fact we need simply extend our earlier definition of I_F^g in the obvious way. Given a model \mathbf{M} and an assignment g in \mathbf{M} , we define (as before) $I_F^g(\tau)$ to be $F(\tau)$ if τ is a constant, and $g(\tau)$ if τ is a variable. On the other hand, if τ is a term of the form $f(\tau_1, \dots, \tau_n)$, then we define $I_F^g(\tau)$ to be $F(f)(I_F^g(\tau_1), \dots, I_F^g(\tau_n))$. (That is, we apply the n -place function $F(f)$ —the function interpreting f —to the interpretation of the n argument terms.)

Function symbols are certainly a natural extension to first-order languages, as the fatherhood example should suggest. However, they are also important for more technical purposes. In particular, they play a key role in Chapter 5, where they will help us formulate an inference system for first-order logic that is suitable for computational implementation.

The first-order languages we have so far defined have an obvious expressive shortcoming: we have no way to assert that two terms denote the same entity. We will need to be able to express such identities in Chapter ??, when we link Discourse Representation Theory with first-order logic. So what are we to do?

Actually, the solution is straightforward. Given any language of first-order logic (with or without function symbols) we can turn it into a first-order language *with equality* by adding the special two place relation symbol $=$. We use this relation symbol in the natural infix way: that is, if τ_1 and τ_2 are terms then we write $\tau_1 = \tau_2$ rather than the rather ugly $=(\tau_1, \tau_2)$. Beyond this notational convention, there's nothing more to say about the syntax of $=$; it's just a two place relation symbol. But what about its semantics?

Here matters are more interesting. Although, syntactically, $=$ is just a 2-place relation symbol, it is a very special one. In fact (unlike LOVES, or HATES, or any other two place relation symbol) we are *not* free to interpret it how we please. In fact, given any model \mathbf{M} , any assignment g in \mathbf{M} , and any terms τ_1 and τ_2 , we shall insist that

$$\mathbf{M}, g \models \tau_1 = \tau_2 \text{ iff } I_F^g(\tau_1) = I_F^g(\tau_2).$$

That is, the atomic formula $\tau_1 = \tau_2$ is satisfied if and only if τ_1 and τ_2 have exactly the same interpretation. That is, $=$ really means equality. In fact, $=$ is usually regarded as a *logical* symbol on a par with \neg or \forall , for like these symbols it has a fixed interpretation, and a semantically fundamental one at that.

1.2 A Simple Model Checker

We will now implement a simple first-order model checker in Prolog. The checker will take the Prolog representation of a model and the Prolog representation of a formula and test whether or not the formula is satisfied in the model. (Actually, to keep things simple, we are only going to work with exact models, and we are not going to deal with function symbols or equality.) We will learn two important lessons from this exercise:

1. By making use of Prolog variables to represent first-order variables, it is easy to implement the satisfaction definition. This is because Prolog's built-in unification and backtracking mechanisms can take care—more or less automatically—of the process of assigning values to variables.
2. On the other hand, this simple strategy has its dangers: as we shall see, the obvious implementation is incorrect. We demonstrate some of its shortcomings and show how to correct them in the next section.

How should we implement such a model checker? We have four principal tasks: deciding how to represent vocabularies, deciding how to represent models, deciding how to represent formulas, and specifying how (representations of) formulas are to be evaluated in (representations of) models. The representations introduced here will be used throughout the book.

Representing Vocabularies

Vocabularies will be represented as Prolog databases. Let's consider an example. Suppose we are working with the following vocabulary:

```
{ (LOVE,2), (HATE,2),
  (CUSTOMER,1), (ROBBER,1),
  (MIA,0), (VINCENT,0),
  (HONEY-BUNNY,0), (PUMPKIN,0) }
```

This would be represented by the following database:

```
relation(love,2).          constant(mia).
relation(hate,2).          constant(vincent).
relation(customer,1).      constant(honey_bunny).
relation(robber,1).        constant(pumpkin).
```

This should be self explanatory: the database simply lists the relation symbols in the vocabulary (together with their arity), and the constant symbols. Both relation symbols and constant symbols are represented by Prolog atoms in the obvious way.

Representing Models

Suppose we have fixed our vocabulary—for example, suppose we've decided to work with the vocabulary just given. How should we represent models of this vocabulary in Prolog?

Actually, this is an impossibly difficult question. We don't have the remotest chance of dealing with *all* models of this vocabulary, for there are models of this vocabulary based on *any* non-empty domain D whatsoever. In particular, there are lots of *infinite* models. Now, it is possible to give useful finite representations of some infinite models—but most are too big and too unruly to be worked with computationally. Hence we shall confine our attention to *finite* models of this vocabulary. In addition, matters are simpler if we only have to deal with *exact* models. (Recall that an exact model is a model over a vocabulary containing constants such that every element in the model is named by exactly one constant.) So, let's rephrase our question: how should we represent the finite exact models over the above vocabulary?

Here is an example of how we shall do it:

```
[customer(vincent), customer(mia),
  robber(honey_bunny), robber(pumpkin),
  loves(pumpkin, honey_bunny)]
```

Fairly obviously, this list represents a model in which both Vincent and Mia are customers, both Pumpkin and Honey Bunny are robbers, and Pumpkin loves Honey Bunny. But it is important to be aware that other properties are *implicitly* recorded. For example, no facts about the *hate* relationships are given—and *hate* is one of the relation symbols in our vocabulary. This is Prolog's way of recording the fact that the binary relation *hate* is empty; in this little world, nobody hates anybody. Moreover, although Pumpkin loves Honey Bunny, he doesn't love himself nor anybody else, for no such facts are recorded. In fact, no other loving relationships at all hold. Thus, by explicitly listing positive information, and *implicitly* listing negative information, we have completely described a unique exact model over this vocabulary.

Here's a second example, again over the same vocabulary.

```
[customer(vincent),
  robber(honey_bunny), robber(pumpkin),
  loves(pumpkin, honey_bunny)]
```

Now for a trick question: how many individuals are there in the domain of the model that this Prolog term represents? The answer is *four*. If you thought the answer was three—because only Vincent, Pumpkin and Honey Bunny are explicitly mentioned—you haven't properly grasped the role played by the constants listed in the vocabulary. In *any* model of this vocabulary there is an individual that corresponds to Mia. As it happens, in this particular model no positive facts are listed about Mia—but lots of negative facts about Mia are implicitly given. For example, Mia is neither a customer nor a robber (perhaps Mia is a waitress), and Mia neither loves nor hates anyone or anything, and is not loved or hated by anyone or anything.

In short, be careful. The list of constants in the vocabulary plays an important role for us: it tells us exactly how many individuals there are in any exact model of that vocabulary, namely one and only one for each constant. Indeed, we can think of the constants listed in the vocabulary as simply *being* the domain of quantification of exact models.

Exercise 1.2.1 Give the set theoretic description of the models that the two Prolog terms given above represent.

Exercise 1.2.2 Suppose we are working with the following vocabulary:

$$\{ (\text{WORKS-FOR},2), \\ (\text{BOXER},1), (\text{PROBLEM-SOLVER},1), \\ (\text{THE-WOLF},0), (\text{MARSELLUS},0), (\text{BUTCH},0) \}$$

First, represent this vocabulary as a Prolog database. Then, represent each of the following two (exact) models (D, F) over this vocabulary as Prolog terms.

1. $D = \{d_1, d_2, d_3\}$,
 $F(\text{THE-WOLF}) = d_1$,
 $F(\text{MARSELLUS}) = d_2$,
 $F(\text{BUTCH}) = d_3$,
 $F(\text{BOXER}) = \{d_3\}$,
 $F(\text{PROBLEM-SOLVER}) = \{d_1\}$.
2. $D = \{\text{entity-1}, \text{entity-2}, \text{entity-3}\}$
 $F(\text{THE-WOLF}) = \text{entity-3}$,
 $F(\text{MARSELLUS}) = \text{entity-1}$,
 $F(\text{BUTCH}) = \text{entity-2}$,
 $F(\text{BOXER}) = \{\text{entity-2}, \text{entity-3}\}$,
 $F(\text{PROBLEM-SOLVER}) = \{\text{entity-2}\}$,
 $F(\text{WORKS-FOR}) = \{(\text{entity-3}, \text{entity-1}), (\text{entity-2}, \text{entity-1})\}$.

Exercise 1.2.3 Write a Prolog program which when given a vocabulary, and a list, determines whether or not the list represents an (exact) model over that vocabulary. Can your program be used to generate all the exact models over that vocabulary? If not, write a program that can.

Representing Formulas

Let us now decide how to represent first-order formulas (for languages without function symbols or equality) in Prolog. The first (and most fundamental) decision is how to represent first-order variables. We make the following choice:

First-order variables will be represented by Prolog variables.

As has already been mentioned, this decision will enable us to implement a model checker very simply—but it will also give rise to a number of problems.

Next, we must decide how to represent the non-logical symbols. We do so in the obvious way: a first-order constant c will be represented by the Prolog atom c , and a first-order relation symbol R will be represented by the Prolog atom r .

Given this convention, it is obvious how atomic formulas should be represented. For example, $\text{LOVE}(\text{VINCENT}, \text{MIA})$ would be represented by the Prolog term `love(vincent, mia)`, and $\text{HATE}(\text{BUTCH}, x)$ would be represented by `hate(butch, X)`. Note that first-order atomic sentences (for example $\text{BOXER}(\text{BUTCH})$) are represented by exactly the same Prolog term (namely `boxer(butch)`) that is used to represent the fact that Butch is a boxer in our Prolog representation of models. This will enable us to give a simple implementation of the satisfaction clause for atomic formulas.

Next the Booleans. The symbols

`&` `v` `>` `~`

will be used to represent the connectives \wedge , \vee , \rightarrow , and \neg respectively. The following Prolog code ensure that these connectives have their usual precedences:

```
:- op(900, yfx, >).      % implication
:- op(850, yfx, v).     % disjunction
:- op(800, yfx, &).     % conjunction
:- op(750, fy, ~).     % negation
```

Finally, we must decide how to represent the quantifiers. Suppose *formula* is a first-order formula, and `Formula` is its representation as a Prolog term. Then $\forall x \text{ formula}$ will be represented as

`forall(X, Formula)`

and $\exists x \text{ formula}$ will be represented as

`exists(X, Formula)`

Semantic Evaluation

We now turn to the fourth and final task: evaluating (representations of) formulas in (representations of) models. The predicate which carries out the task is called `satisfy/2`,

and the clauses of `satisfy/2` mirror the first-order satisfaction definition in a fairly natural way.

The clause which evaluates (representations of) atomic formulas is:

```
satisfy(Formula,Model):-
    member(Formula,Model).
```

The `member/2` predicate is one of the predicates in the library `comsemPredicates.pl`. It succeeds if its first argument, any Prolog term, is a member of its second argument, which has to be a list. It is defined as follows:

```
member(X,[X|_]).
member(X,[_|L]):-
    member(X,L).
```

Hence an atomic formula is true if and only if that formula is one of the facts recorded in the list `Model`.

Now for Boolean combinations of formulas.

```
satisfy(Formula1 & Formula2,Model):-
    satisfy(Formula1,Model),
    satisfy(Formula2,Model).
```

```
satisfy(Formula1 v Formula2,Model):-
    satisfy(Formula1,Model);
    satisfy(Formula2,Model).
```

```
satisfy(Formula1 > Formula2,Model):-
    satisfy(Formula2,Model);
    \+ satisfy(Formula1,Model).
```

```
satisfy(~ Formula,Model) :-
    \+ satisfy(Formula,Model).
```

These mirror the first-order satisfaction definition in an obvious way. However note that in the clauses for `>` and `~` we have made use of Prolog negation (that is, negation as failure). This is a simple and natural thing to do, but as we shall see in the next section, it can interact with our use of Prolog variables to represent first-order variables in unexpected—and unintended—ways.

Finally come the quantifier clauses. The one for the existential quantifier couldn't be simpler:

```
satisfy(exists(X,Formula),Model):-
    constant(X),
    satisfy(Formula,Model).
```

That is, the program unifies `X` to some entity in the model (recall that because we are working with exact models, the entities we are talking about are essentially the constants) and then tries evaluating `Formula` in `Model`. If `Formula` is *not* satisfied in `Model` the program will backtrack and unify `X` to another entity (if this is still possible) and then evaluate the formula again. Either it eventually succeeds (which means that there exists an entity in the model that does satisfy `Formula`), or it doesn't (which means there isn't). Clearly, this captures the semantics of the existential quantifier. Moreover it does so in a straightforward way: Prolog's unification and backtracking mechanisms are being used to simulate the process of assigning values to first-order variables.

At first glance, the following clause for the universal quantifier may look a little strange—but recall that $\forall x\phi$ is logically equivalent to $\neg\exists x\neg\phi$ (the reader was asked to show this in Exercise 1.1.4). The following clause simply represents this equivalence in Prolog (using negation as failure), and makes use of Prolog's unification and backtracking mechanisms in much the same way as the clause for `exists` does.

```
satisfy(forall(X,Formula),Model):-
    satisfy(~ exists(X,~ Formula),Model).
```

So, it's time to start checking models. Here are some examples of models over our original vocabulary:

```
example(1,[customer(mia),customer(vincent),
           robber(pumpkin),robber(honey_bunny),
           love(pumpkin,honey_bunny)]).
```

```
example(2,[customer(mia),
           robber(pumpkin),robber(honey_bunny),
           love(pumpkin,honey_bunny)]).
```

```
example(3,[customer(mia),customer(vincent),
           robber(pumpkin),robber(honey_bunny),
           love(pumpkin,honey_bunny),love(mia,vincent)]).
```

And here is a driver predicate which evaluates a formula on the desired example model:

```
evaluate(Formula,Example):-
    example(Example,Model),
    satisfy(Formula,Model).
```

Exercise 1.2.4 Systematically test the model checker on these models. Are the results always what you expected? If not, why not?

1.3 Some Refinements

As the reader who has tested the model checker will have learned, the current version of the model checker is far from satisfactory. In fact, it does quite a number of silly things. Let's try pinning down the various problems.

First Problem

What happens if we ask Prolog to evaluate a *variable* in one of our example models:

```
?- evaluate(X,1).
```

This is a perfectly reasonable query; in fact there are two motives we might have for making it. The first is simply to ask Prolog to generate a formula that can be evaluated in model 1. The second—recall that we decided to use Prolog variables for representing first-order variables—is to check whether our predicate successfully recognizes that a variable is not a formula, and hence cannot be evaluated. Both are good ideas, but neither works as planned, for the program immediately goes into an infinite loop.

Exercise 1.3.1 Why does the program get in an infinite loop when posing the above query? If this is not clear, perform a trace.

To solve this problem, we will simply reject queries that try to evaluate variables. This is easily done by adding a filtering clause to the evaluation predicate:

```
evaluate(Formula,Example):-
    \+ var(Formula),
    example(Example,Model),
    satisfy(Formula,Model).
```

This ensures that `evaluate` only succeeds if `Formula` is not a variable. However, this problem is deeper and more structural, and this proposed solution is not satisfactory. To find out why you might try the following exercise.

Exercise 1.3.2 [Easy] Try the above evaluation predicate on formulas that contain subformulas that are represented as Prolog variables, for example the formula `exists(X,robber(X) > Y)`. What happens and why?

Exercise 1.3.3 [Hard] Try to modify the program in such a way that it generates formulas that can be evaluated in a specified model when given the query `evaluate(X,1)` (for example model 1).

Second Problem

The second problem has to do with the way negation as failure interacts with our use of Prolog variables to represent first-order variables.

Suppose we evaluate

```
customer(X)
```

in example 1. That is, we're asking whether it is possible to assign a value to the free variable that satisfies the formula; thus Prolog's task is to find an entity in this model that is a customer. In example 1, of course, it finds two: Mia and Vincent.

Now suppose we evaluate

```
~ customer(X)
```

in example 1. That is, we're again asking whether it is possible to assign a value to the free variable that satisfies the formula, thus Prolog's task is to find an entity in this model that is *not* a customer. In example 1, of course, there are two non-customers, namely Honey Bunny and Pumpkin.

However, Prolog wrongly answers `no`. This is because we defined the interpretation of negated clauses using negation as failure. As `customer(X)` is satisfiable, Prolog (incorrectly) decides that `~ customer(X)` is not satisfiable.

This is unacceptable and has to be fixed. But before doing so, let's look at another problem.

Third Problem

We run into problems if we ask our model checker to verify formulas using completely new non-logical vocabulary. For example, suppose we try evaluating the atomic formula

```
tasty(royale_with_cheese)
```

in any of example models. Then our model checker will say `no`. This response is *not*, strictly speaking, correct. Formally the satisfaction relation is not defined between formulas and

models of different vocabularies. The correct response of our model checker would be to throw out this formula and say something like “Hey, I don’t know anything about these symbols!”.

This example may not strike the reader as too problematic—but worse is in store. If we evaluate

```
~ tasty(royale_with_cheese)
```

in any of example models they will return the answer **yes!** This happens because we used negation as failure to define the evaluation clause for negated formulas. Clearly this needs fixing.

However the problem is not particularly deep; the solution is simply to be explicit about what a well formed formula over a given vocabulary is. That is, we should make use of the information in the vocabulary about which relation symbols we are working with (and, of course, what their arities are), which constants we have, and explicitly state which combinations of these symbols are well formed. (Note that at present we’re not making *any* use of the information stored in our vocabulary. It’s hardly surprising that we’re having these difficulties.) Our driver should check that any sequence of symbols it has to evaluate really is a well formed combination of the chosen relation and constant symbols—and it should refuse to evaluate anything that isn’t.

Now, to do this we’re going to have to recursively work through the structure of the input formula—so this gives us a good opportunity to solve the second problem as well. In fact, we shall define a predicate `sentence/2` whose primary task is to transform a formula with free variables into a sentence by substituting constants of the vocabulary for occurrences of free variables. While doing this it will check that the Prolog terms it encounters represent well formed formulas built out of the given vocabulary. Thus `sentence/2` will solve the second and third problems simultaneously.

We define `sentence/2` as follows. Its first argument will be a list; we use it to keep track of all the bound variables we have encountered in the course of analysis. The second argument is a formula.

The first clause deals with evaluating uninstantiated Prolog variables. This immediately causes failure, and therefore deals with the first problem. (We use the `!, fail` combination to exclude all other clauses for possible solutions.)

```
sentence(_,Var):-
    var(Var), !, fail.
```

The required clauses for non-atomic formulas are fairly obvious:

```

sentence(Bound,forall(X,Formula)):-
    var(X),
    sentence([X|Bound],Formula).

sentence(Bound,exists(X,Formula)):-
    var(X),
    sentence([X|Bound],Formula).

sentence(Bound,Formula1 > Formula2):-
    sentence(Bound,Formula1),
    sentence(Bound,Formula2).

sentence(Bound,Formula1 & Formula2):-
    sentence(Bound,Formula1),
    sentence(Bound,Formula2).

sentence(Bound,Formula1 v Formula2):-
    sentence(Bound,Formula1),
    sentence(Bound,Formula2).

sentence(Bound,~ Formula):-
    sentence(Bound,Formula).

```

Note the way the clauses for `forall` and `exists` push the variable currently being bound onto the bound variable list (after first testing that it really *is* a legitimate Prolog variable).

But that's the easy part. The real work is done at the atomic level, where we substitute constants for free variables. Here's the code:

```

sentence(Bound,Formula):-
    compose(Formula,Symbol,Arguments),
    length(Arguments,Arity),
    relation(Symbol,Arity),
    goodArguments(Bound,Arguments).

goodArguments(_Bound, []).
goodArguments(Bound,[Arg|Others]):-
    member(Var,Bound),
    Var==Arg,!,
    goodArguments(Bound,Others).
goodArguments(Bound,[Arg|Others]):-
    constant(Arg),
    goodArguments(Bound,Others).

```


Here `compose/3` is a predicate in the library file `comsemPredicates.pl`. It is defined as follows:

```
compose(Term,Symbol,ArgList):-
    Term =.. [Symbol|ArgList].
```

In short, `compose/3` uses the built in Prolog `=..` functor to flatten a term into a list.

Thus, at the base of the recursion, `sentence/2` uses `compose/3` to take the atomic formula apart into a list. It then checks that the first item on the list really is one of our relation symbols, and that it has the correct number of arguments. If everything is satisfactory, it hands the list over to `goodArguments/2`. This predicate recursively checks through the tail of the list (that is, the terms used to build the atomic formula). If one of these items is a Prolog variable, it checks to see whether it is listed in `Bound`. If it is, that's fine; we are dealing with a bound variable and don't have to worry. On the other hand, if it's *not* on this list, we have found a free variable which needs replacing by a constant. The call `constant(Arg)` unifies it to one of the constants in the vocabulary. Note that on backtracking this call will successively unify the free variable to every constant in the vocabulary. In effect this code is taking advantage of the similarity between free variables and constants discussed in Exercise 1.1.2; as we observed, when working with exact models, every formula is equivalent to a sentence.

And that's it. It only remains to write a new driver which puts these predicates to work. We'll implement the driver using two clauses. The first clause deals with formulas that are sentences, the second clause with formulas that aren't sentence (in that case it alerts the user with a message saying that the formula cannot be evaluated).

```
evaluate(Formula,Example):-
    sentence([],Formula),
    example(Example,Model),
    satisfy(Formula,Model).

evaluate(Formula,_Example):-
    \+ sentence([],Formula),
    nl,write('Not a wff over the given vocabulary. '),
    nl,write('Cannot be evaluated. '),nl.
```

It should be clear that this revised version of our evaluation predicate fixes the third problem.

This is a much cleaner model checker. Simply by making use of the information stored in vocabularies it tidies up a lot of silly problems in one fell swoop; indeed, it even tidies up a problem not mentioned in the text. It is important that you understand what is going on here, so we suggest you work through the following problems right away.

Exercise 1.3.4 Try out the new evaluation predicate on the query `~ customer(X)`. Why has problem 2 vanished?

Exercise 1.3.5 Suppose we continue to work with the same vocabulary and we add the model \square to our list of examples. That is, we add the model that contains no positive information about any of our four individuals. Does the first version of `evaluate/2` handle queries about this model correctly? Does the revised version? Explain why.

Another Problem

Our new evaluation predicate takes care of the sillier shortcomings of our first version, and we won't bother making any further improvements. Nonetheless, the reader should realize that it is still not perfect. There is another problem, due solely to our decision to use Prolog variables to represent first-order variables.

Although more complex than the original version, the new `evaluate/2` is a fairly simple predicate. This is because all the real work of variable handling is being taken care of—more or less automatically—by Prolog's in-built unification and backtracking mechanisms. But in a sense, we've cheated. Neither the original `evaluate` nor its successor really understand how to treat free and bound variables at all.

An example should make this clear. In first-order logic the formulas

$$\exists x (\text{CUSTOMER}(x)) \wedge \exists y (\text{ROBBER}(y))$$

and

$$\exists x (\text{CUSTOMER}(x)) \wedge \exists x (\text{ROBBER}(x))$$

are equivalent. Now, their Prolog representations are

$$\text{exists}(X,\text{customer}(X)) \ \& \ \text{exists}(Y,\text{robber}(Y))$$

and

$$\text{exists}(X,\text{customer}(X)) \ \& \ \text{exists}(X,\text{robber}(X))$$

respectively. However

$$\text{?- evaluate}(\text{exists}(X,\text{customer}(X)) \ \& \ \text{exists}(Y,\text{robber}(Y)),1).$$

succeeds (it returns all four of the successful ways of instantiating X and Y) while

```
?- evaluate(exists(X,customer(X)) & exists(X,robber(X)),1).
```

returns `no`.

The difficulty is easy to spot. Prolog doesn't know that the X in `customer(X)` is intended to play a different role than the X in `robber(X)`; it will always instantiate them to the same thing.

But fortune seems on our side, as there is a simple Prolog trick that helps us here. The problem lies in instantiating bound variables by one of the constants of our vocabulary using `constant/1` in the predicate `satisfy/2`. Here is where this happens (and note that this is the *only* spot where variables get instantiated):

```
satisfy(exists(X,Formula),Model):-
    constant(X),
    satisfy(Formula,Model).
```

Once X gets instantiated, it remains instantiated. And that's what causes the problem. A simple Prolog-remedy is to use double negation, which ensures that variables inside it are not instantiated (in fact, they are only instantiated during the proof of the negation). Here is a revised clause for handling the existential quantifier:

```
satisfy(exists(X,Formula),Model):-
    \+ \+ (constant(X), satisfy(Formula,Model)).
```

However, this trick doesn't cover all problems of binding. Here is an example of a formula that doesn't get evaluated in model 1 although it's a perfectly fine formula in first-order logic (and satisfied in our example model 1!).

$$\exists x (\text{CUSTOMER}(x) \wedge \exists x \text{ROBBER}(x))$$

The moral is clear. We should either write a better evaluation predicate that guards against this shortcoming (the reader is asked to do this as an exercise) or we should ensure that different occurrences of quantifiers bind distinct variables. More generally, it is clear that 'clever' Prolog solutions to variable manipulation problems require careful scrutiny; we will see another example in the following chapter when we implement β -reduction.

Exercise 1.3.6 [easy] Write a predicate which when given a formula, checks for reused variables. Redefine the predicate `evaluate` that filters out such sentences.

Exercise 1.3.7 [hard] Redefine the predicate `satisfy/2` such that it uses explicit assignment functions.

1.4 First-Order Logic and Natural Language

Expressivity Problems

(to be provided)

Representational Problems

(to be provided)

Software Summary of Chapter 1

`modelChecker.pl` The file that contains the code for the simple model checker for first-order logic. This version neither takes care of free variables, nor does it consult the vocabulary to exclude ill-formed formulas. (page 209)

`modelChecker2.pl` This file contains the code for the revised model checker for first-order logic. This version takes care of free variables and excludes ill-formed formulas. (page 211)

`exampleModels.pl` Contains the specification of the vocabulary and a few example models. (page 214)

Notes

There are many good introductions to first-order logic, and the best advice we can give the reader is to spend some time browsing in a library to see what's on offer. That said, there's three references we particularly recommend. For an unhurried introduction that motivates the subject linguistically, try the first volume of Gamut 1991a. For a wide-ranging discursive overview, we recommend Hodges 1983. This survey article covers a lot of ground—from truth tables to Lindström's celebrated characterization of first-order logic. Parts of the article are quite technical, but, by and large, it's the sort of article that many readers will enjoy browsing through. On the other hand, the reader who wants a more focussed approach may prefer Enderton 1972. This is a good, solid, and extremely readable introduction to first-order logic.

Chapter 2

Lambda Calculus

Now that we know something of first-order logic and how to work with it in Prolog, it is time to turn to the first major theme of the book, namely:

How can we automate the process of associating semantic representations with natural language expressions?

In this chapter we explore the issue concretely. We proceed by trial and error. We first write a simple Prolog program that performs the task in a limited way. We note where it goes wrong, and why, and develop a more sophisticated alternative. These experiments lead us, swiftly and directly, to formulate a version of the *lambda calculus*. The lambda calculus is a tool for controlling the process of making substitutions. With its help, we will be able to describe, neatly and concisely, how semantic representations should be built. The lambda calculus is one of the main tools used in this book, and by the end of the chapter the reader should have a reasonable grasp of why it is useful to computational semantics and how to work with it in Prolog.

2.1 Compositionality

Given a sentence of English, is there a systematic way of constructing its semantic representation? This question is far too general, so let's ask a more specific one: is there a systematic way of translating simple sentences such as 'Vincent likes Mia' and 'A woman snorts' into first-order logic?

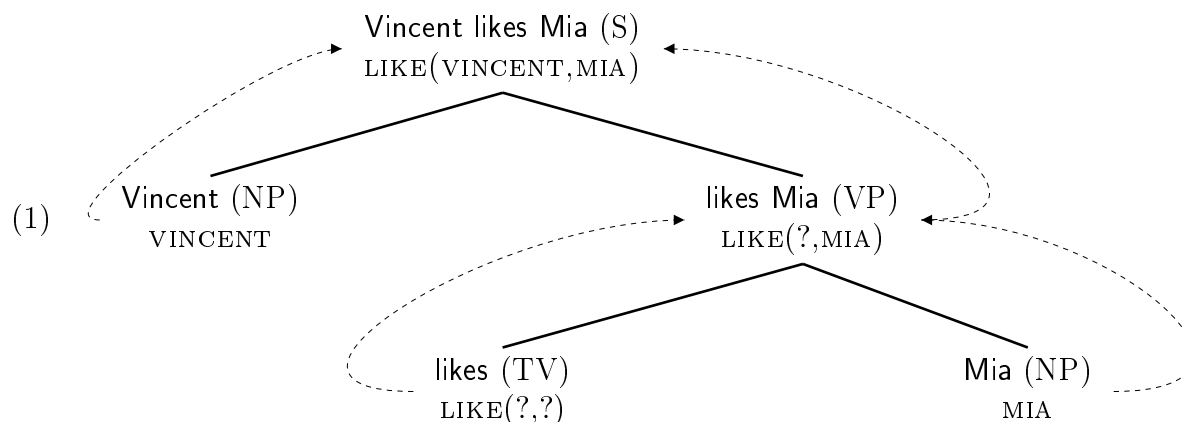
The key to answering this is to be more precise about what we mean by 'systematic'. Consider 'Vincent likes Mia'. Its semantic content is at least partially captured by the first-order formula `LIKE(VINCENT,MIA)`. Now, the most basic observation we can make about systematicity is the following: the proper name 'Vincent' contributes the constant

VINCENT to the representation, the transitive verb ‘likes’ contributes the relation symbol LIKE, and ‘Mia’ contributes MIA.

This first observation is important, and by no means as trivial as it may seem. If we generalize it to the claim that the words making up a sentence contribute *all* the bits and pieces needed to build the sentence’s semantic representation, we have formulated a principle that is a valuable guide to the complexities of natural language semantics. The principle is certainly plausible. Moreover, it has the advantage of forcing us to face up to a number of non-trivial issues sooner rather than later (for example, what exactly does the determiner ‘every’ contribute to the representation of ‘Every woman loves a boxer?’).

Nonetheless, though important, this principle doesn’t tell us everything we need to know about systematicity. For example, from the symbols LIKE, MIA and VINCENT we can also form LIKE(MIA,VINCENT). Why don’t we get this (incorrect) representation when we translate ‘Vincent likes Mia’? What exactly is it about the sentence ‘Vincent likes Mia’ that forces us to translate it as LIKE(VINCENT,MIA)? Note that the answer “But ‘Vincent likes Mia’ means LIKE(VINCENT,MIA), stupid!”, which in many circumstances would be appropriate, isn’t particularly helpful here. Computers *are* stupid. We can’t appeal to their semantic insight, because they don’t have any. If we are to have any hope of automating semantic construction, we must find another kind of answer.

The missing ingredient is a notion of *syntactic structure*. ‘Vincent likes Mia’ isn’t just a string of words: it has a hierarchical structure. In particular, ‘Vincent likes Mia’ is an S (sentence) that is composed of the subject NP (noun phrase) ‘Vincent’ and the VP (verb phrase) ‘likes Mia’. This VP is in turn composed of the TV (transitive verb) ‘likes’ and the direct object NP ‘Mia’. Given this hierarchy, it is easy to tell a coherent story—and indeed, to draw a convincing picture—about why we should get the representation LIKE(VINCENT,MIA), and not anything else:



Why is MIA in the second argument slot of LIKE? Because, when we combine a TV with an NP to form a VP, we have to put the semantic representation associated with the NP (in this case, MIA) in the *second* argument slot of the VP’s semantic representation (in

this case, LIKE(?,?)). Why does VINCENT have to go in the *first* argument slot? Because this is the slot reserved for the semantic representations of NPs that we combine with VPs to form an S. More generally, given that we have some reasonable syntactic story about what the pieces of the sentences are, and which pieces combine with which other pieces, we can try to use this information to explain how the various semantic contributions have to be combined. In short, one reasonable explication of ‘systematicity’ is that it amounts to using the additional information provided by syntactic structure to spell out exactly how the semantic contributions are to be glued together.

Our discussion has led us to one of the key concepts of contemporary semantic theory: *compositionality*. Suppose we have some sort of theory of syntactic structure. It doesn’t matter too much what sort of theory it is, just so long as it is hierarchical in a way that ultimately leads to the lexical items. (That is, our notion of syntactic structure should allow us to classify the sentence into subparts, sub-subpart, and sub-sub-subparts, . . . , and so on—ultimately into the individual words making up the sentence.) Such structures make it possible to tell an elegant story about where semantic representations come from. Ultimately, semantic information flows from the lexicon, thus each lexical item is associated with a representation. How is this information combined? By making use of the hierarchy provided by the syntactic analysis. Suppose the syntax tells us that some kind of sentential subpart (a VP, say) is decomposable into two sub-subparts (a TV and an NP, say). Then our task is to describe how the semantic representation of the VP subpart is to be built out of the representation of its two sub-subparts. If we succeed in doing this for all the grammatical constructions covered by the syntax, we will have given a *compositional semantics* for the language under discussion (or at least, for that fragment of the language covered by our syntactic analysis).

Compositionality is a simple and natural concept that underlies most work in natural language semantics and the semantics of programming languages. Nonetheless, in spite of its simplicity, it raises a number of interesting issues. For example, is every ‘systematic’ semantics a compositional one? In Part II we discuss the standard top-down algorithm for Discourse Representation Theory. This algorithm is a systematic approach to semantic construction (under any reasonable interpretation of the word ‘systematic’) but many semanticists have argued that it is not truly compositional.

Syntax via Definite Clause Grammars

So, is there a systematic way of translating simple sentences such as ‘Vincent likes Mia’ and ‘A woman snorts’ into first-order logic? We don’t yet have a method, but at least we now have a plausible strategy for finding one. We need to:

Task 1 Specify a reasonable syntax for the fragment of natural language of interest.

Task 2 Specify semantic representations for the lexical items.

Task 3 Specify the translation compositionally. That is, we should specify the translation of all expressions in terms of the translation of their parts, where ‘parts’ refers to the substructure given to us by the syntax.

Moreover, all three tasks need to be carried out in a way that leads naturally to computational implementation.

As this is a book on computational semantics, tasks 2 and 3 are where our real interests lie, and most of our energy will be devoted to them. But since compositional semantics presupposes syntax, we need a way of handling task 1. What should we do?

We have opted for a particularly simple solution: in this book the syntactic analysis of a sentence will be a tree whose non-leaf nodes represent *complex syntactic categories* (such as S, NP and VP) and whose leaves represent *lexical items* (these are associated with *basic syntactic categories* such as noun, transitive verb, determiner, proper name and intransitive verb). The tree the reader has just seen is a typical example. This approach has an obvious drawback (namely, the reader won’t learn anything interesting about syntax) but it also has an important advantage: we will be able to make use of Definite Clause Grammars (DCGs), the in-built Prolog mechanism for grammar specification.

Here is a DCG for the fragment of English we shall use in our initial semantic construction experiments. (This DCG, decorated with semantic construction code, can be found in `experiment1.pl` and `experiment2.pl`. The reader unfamiliar with DCGs is advised to consult Appendix D right away, as otherwise this chapter will shortly stop making much sense.)

```

s --> np, vp.           noun --> [woman].
np --> pn.              noun --> [foot,message].
np --> det, noun.       vp --> iv.
pn --> [vincent].       vp --> tv, np.
pn --> [mia].           iv --> [walks].
det --> [a].            tv --> [loves].
det --> [every].        tv --> [likes].

```

This grammar tells us how to build certain kinds of sentences (`s`) out of noun phrases (`np`), verb phrases (`vp`), proper names (`pn`), determiners (`det`), nouns (`noun`), intransitive verbs (`iv`), and transitive verbs (`tv`), and gives us a tiny lexicon to play with. For example, the grammar accepts the simple sentence

Vincent walks

because ‘Vincent’ is declared as a proper name, and proper names are noun phrases according to this grammar; ‘walks’ is an intransitive verb, and hence a verb phrase; and sentences can consist of a noun phrase followed by a verb phrase.

But the real joy of DCGs is that they provide us with a with a lot more than a natural notation for specifying grammars. Because they are part and parcel of Prolog, we can actually compute with them. For example, by posing the query

```
s([mia,likes,a,foot,message],[])
```

we can test whether ‘Mia likes a foot massage’ is accepted by the grammar, and the query

```
s(X,[])
```

generates all grammatical sentences.

Exercise 2.1.1 How many sentences are accepted by this grammar? How many noun phrases? How many verb phrases? Check your answer by generating the relevant items.

With a little effort, we can do a lot more. In particular, by making use of extra arguments (again, see Appendix D if you’re uncertain what this means) we can associate semantic representations with lexical items very straightforwardly. The normal Prolog unification mechanism then gives us the basic tool needed to combine semantic representations, and to pass them up towards sentence level. In short, working with DCGs both frees us from having to implement parsers, and makes available a powerful mechanism for combining representations, so we’ll be able to devote our attention to semantic construction.

The semantic construction methods discussed in this book are compatible with a wide range of syntactical theories. In essence, this book is about exploiting the recursive structure of trees to build representations compositionally; where the trees actually come from is relatively unimportant. We have chosen to fill in the syntactical ‘black box’ using DCGs—but a wide range of more sophisticated options is available and we urge our readers to experiment.

2.2 Two Experiments

How can we systematically associate first-order formulas with the sentences produced by our little grammar? Let’s just plunge in and try, and see how far our knowledge of DCGs and Prolog will take us.

Experiment 1

First the lexical items. We need to associate ‘Vincent’ with the constant `VINCENT`, ‘Mia’ with the constant `MIA`, ‘walks’ with the unary relation symbol `WALK`, and ‘loves’ with

the binary relation symbol `LOVE`. The following piece of DCG code makes these associations. Note that the arity of `walk` and `love` are explicitly included as part of the Prolog representation.

```
pn(vincent)--> [vincent].

pn(mia)--> [mia].

iv(snort(_))--> [snorts].

tv(love(_, _))--> [loves].
```

How do we build semantic representations for sentences? Let's first consider how to build representations for quantifier-free sentences such as 'Mia loves Vincent'. The main problem is to steer the constants into the correct slots of the relation symbol. (Remember, we want 'Vincent loves Mia' to be represented by `LOVE(VINCENT, MIA)`, not `LOVE(MIA, VINCENT)`.) Here's a first (rather naive) attempt. Let's directly encode the idea that when we combine a TV with an NP to form a VP, we have to put the semantic representation associated with the NP in the second argument slot of the VP semantic representation, and that we use the first argument slot for the semantic representations of NPs that we combine with VPs to form Ss.

Recall that Prolog has a built in predicate `arg/3` such that `arg(N,P,I)` is true if `I` is the `N`th argument of `P`. This is a useful tool for manipulating pieces of syntax, and with its help we can cope with simple quantifier free sentences rather easily. Here's the needed extension of the DCG:

```
s(Sem)--> np(SemNP), vp(Sem),
  {
    arg(1, Sem, SemNP)
  }.

np(Sem)--> pn(Sem).

vp(Sem)--> tv(Sem), np(SemNP),
  {
    arg(2, Sem, SemNP)
  }.

vp(Sem)--> iv(Sem).
```

These clauses work by adding an extra argument to the DCG (here, the position filled by the variables `Sem` and `SemNP`) to percolate up the required semantic information using

Prolog unification. Note that while the second and the fourth clauses perform only this percolation task, the first and third clauses, which deal with branching rules, have more work to do: they use `arg/3` to steer the arguments into the correct slots. This is done by associating extra pieces of code with the DCG rules, namely `arg(1,Sem,SemNP)` and `arg(2,Sem,SemNP)`. (These are normal Prolog goals, and are added to the DCG rules in curly brackets to make them distinguishable from the grammar symbols.) This program captures, in a brutally direct way, the idea that the semantic contribution of the object NP goes into the second argument slot of TVs, while the semantic contribution of subject NPs belongs in the first argument slot.

It works. For example, if we pose the query:

```
?- s(Sem,[mia,snorts],[ ]).
```

we obtain the (correct) response:

```
Sem = snort(mia)
```

But this is far too easy—let’s try to extend our fragment with the determiners ‘a’ and ‘every’. First, we need to extend the lexical entries for these words, and the entries for the common nouns they combine with:

```
det(exists(_,&_))--> [a].
```

```
det(forall(_,>_))--> [every].
```

```
noun(woman(_))--> [woman].
```

```
noun(footmassage(_))--> [foot,message].
```

NPs formed by combining a determiner with a noun are called *quantified noun phrases*.

Next, we need to say how the semantic contributions of determiners and noun phrases should be combined. We can do this by using `arg/3` four times to get the instantiation of the different argument positions correct:

```
np(Sem)--> det(Sem), noun(SemNoun),
  {
    arg(1,SemNoun,X),
    arg(1,Sem,X),
    arg(2,Sem,Matrix),
    arg(1,Matrix,SemNoun)
  }.
```

The key idea is that the representation associated with the NP will be the representation associated with the determiner (note that the `Sem` variable is shared between `np` and `det`), but with this representation fleshed out with additional information from the noun. The Prolog variable `X` is a name for the existentially quantified variable the determiner introduces into the semantic representation; the code `arg(1,SemNoun,X)` and `arg(1,Sem,X)` unifies the argument place of the noun with this variable. The code `arg(2,Sem,Matrix)` simply says that the second argument of `Sem` will be the matrix of the NP semantic representation, and `arg(1,Matrix,SemNoun)` then adds more detail: it says that the first slot of the matrix will be filled in by the semantic representation of the noun. So if we pose the query

```
?- np(Sem, [a, woman], []).
```

we obtain the response

```
Sem = exists(X, woman(X)&Y)
```

Note that this representation is an *incomplete* first-order formula. We don't yet have a full first-order formula (the Prolog variable `Y` has yet to be instantiated) but we do know that we are existentially quantifying over the set of women.

Given that such incomplete first-order formulas are the semantic representations associated with quantified NPs, it is fairly clear what must happen when we combine a quantified NP with a VP to form an S: the VP must provide the missing piece of information. (That is, it must provide an instantiation for `Y`.) The following clause does this:

```
s(Sem)--> np(Sem), vp(SemVP),
  {
    arg(1, SemVP, X),
    arg(1, Sem, X),
    arg(2, Sem, Matrix),
    arg(2, Matrix, SemVP)
  }.
```

Unfortunately, while the underlying idea is essentially correct, things have just started going badly wrong. Until now, we've simply been extending the rules of our original DCG with semantic information—and we've already dealt with `s --> np, vp`. If we add this second version of `s --> np, vp` (and it seems we need to) we are duplicating syntactic information. This is uneconomical and inelegant. Worse, this second sentential rule interacts in an unintended way with the rule

```
np(Sem)--> pn(Sem).
```

As the reader should check, as well as assigning the correct semantic representation to ‘A woman snorts’, our DCG also assigns the splendidly useless string of symbols

```
snort(exist(X,woman(X)&Y))
```

But this isn’t the end of our troubles. We already have a rule for forming VPs out of TVs and NPs, but we will need a second rule to cope with quantified NPs in object position, namely:

```
vp(Sem)--> tv(SemTV), np(Sem),
{
  arg(2,SemTV,X),
  arg(1,Sem,X),
  arg(2,Sem,Matrix),
  arg(2,Matrix,SemTV)
}.
```

If we add this rule, we can assign correct representations to all the sentences in our fragment. However we will also produce a lot of nonsense (for example, ‘A woman loves a foot massage’ is assigned four representations, three of which are just jumbles of symbols) and we are being systematically forced into syntactically unmotivated duplication of rules. This doesn’t look promising. Let’s try something else.

Exercise 2.2.1 The code for experiment 1 is in `experiment1.pl`. Generate the semantic representations of the sentences and noun phrases yielded by the grammar.

Experiment 2

Although our first experiment was ultimately unsuccessful, it did teach us something useful: to build representations, we need to work with *incomplete* first-order formulas, and we need a way of manipulating the missing information. Consider the representations associated with determiners. In experiment 1 we associated ‘a’ with `exists(_,&_)`. That is, this determiner contributes the skeleton of a first-order formula whose first slot needs to be instantiated with a variable, whose second slot needs to be filled with the semantic representation of a noun, and whose third slot needs to be filled by the semantic representation of a VP. However in experiment 1 we didn’t manipulate this missing information directly. Instead we took a shortcut: we thought in terms of argument *position* so that we could make use of `arg/3`. Let’s avoid plausible looking short cuts. The idea of missing information is evidently important, so let’s take care to always associate it with an explicit Prolog variable. Perhaps this direct approach will make semantic construction easier.

Let's first apply this idea to the determiners. We shall need three extra arguments: one for the bound variable, one for the contribution made by the noun, and one for the contribution made by the VP. Incidentally, these last two contribution have a standard name: the contribution made by the noun is called the *restriction* and the contribution made by the VP is called the *nuclear scope*. We reflect this terminology in our choice of variable names:

$$\text{det}(X, \text{Restr}, \text{Scope}, \text{exists}(X, \text{Restr} \ \& \ \text{Scope})) \rightarrow [a].$$

$$\text{det}(X, \text{Restr}, \text{Scope}, \text{forall}(X, \text{Restr} \ > \ \text{Scope})) \rightarrow [\text{every}].$$

But the same idea applies to common nouns, intransitive verbs, and transitive verbs too. For example, instead of associating 'woman' with `woman(_)`, we should state that the translation of 'woman' is `WALK(y)` for some particular choice of variable `y`—and we should *explicitly* keep track of which variable we choose. (That is, although the `y` appears free in `WALK(y)`, we actually want to have some sort of hold on it.) Similarly, we want to associate a transitive verb like 'loves' with `LOVE(y,z)` for some particular choice of variables `y` and `z`, and again, we should keep track of the choices we made. So the following lexical entries are called for:

$$\text{noun}(X, \text{woman}(X)) \rightarrow [\text{woman}].$$

$$\text{iv}(Y, \text{snort}(Y)) \rightarrow [\text{snorts}].$$

$$\text{tv}(Y, Z, \text{love}(Y, Z)) \rightarrow [\text{loves}].$$

Given these changes, we need to redefine the rules for producing sentences and verb phrases.

$$\text{s}(\text{Sem}) \rightarrow \text{np}(X, \text{SemVP}, \text{Sem}), \ \text{vp}(X, \text{SemVP}).$$

$$\text{vp}(X, \text{Sem}) \rightarrow \text{tv}(X, Y, \text{SemTV}), \ \text{np}(Y, \text{SemTV}, \text{Sem}).$$

$$\text{vp}(X, \text{Sem}) \rightarrow \text{iv}(X, \text{Sem}).$$

The semantic construction rule associated with the sentential rule, for example, tells us that `Sem`, the semantic representation of the sentence, is essentially going to be that of the noun phrase (that's where the value of the `Sem` variable will trickle up from) but that, in addition, the bound variable `X` used in `Sem` must be the same as the variable used in the verb phrase semantic representation `SemVP`. Moreover, it tells us that `SemVP` will be used to fill in the information missing from the semantic representation of the noun phrase.

So far so good, but now we need a little trickery. Experiment 1 failed because there was no obvious way of making use of the semantic representations supplied by quantified noun

phrases and proper names in a single sentential rule. Here we only have a single sentential rule, so evidently the methods of experiment 2 avoid this problem. Here's how it's done:

$$\text{np}(X, \text{Scope}, \text{Sem}) \rightarrow \text{det}(X, \text{Restr}, \text{Scope}, \text{Sem}), \text{noun}(X, \text{Restr}).$$

$$\text{np}(\text{SemPN}, \text{Sem}, \text{Sem}) \rightarrow \text{pn}(\text{SemPN}).$$

Note that we have given all noun phrases—regardless of whether they are quantifying phrases or proper names—the same arity in the grammar rules. (That is, there really is only *one* np predicate in this grammar, not two predicates that happen to make use of the same atom as their functor name.)

It should be clear how the first rule works. The skeleton of a quantified noun phrase is provided by the determiner. The restriction of this determiner is filled by the noun. The resultant noun phrase representation is thus a skeleton with two marked slots: **X** marks the bound variable, while **Scope** marks the missing scope information. This **Scope** variable will be instantiated by the verb phrase representation when the sentential rule is applied.

The second rule is trickier. The vital work is performed by the doubled **Sem** variable. Roughly speaking, when we combine the noun phrase representations formed by this rule with a verb phrase, the effect of this doubling is to 'move the verb phrase representation rightwards' so that it occupies the slot that will ultimately be used in forming the sentential semantics. Intuitively, whereas the representation for sentences that have a quantified noun phrase as subject are essentially the subject's representation filled out by the verb phrase representation, the reverse is the case when we have a proper name as subject. When that happens, the verb phrase is boss. The sentence representation is essentially the verb phrase representation, and the role of the proper name is simply to obediently instantiate the marked verb phrase slot. The rightwards shuffle performed by the doubled variables is the Prolog mechanism which captures this role-reversal.

Our second experiment has fared far better than our first. It is clearly a good idea to explicitly mark missing information; this gives us the control required to fill it in and maneuver it into place. Nonetheless, experiment 2 uses the idea clumsily. Much of the work is done by the rules. These state how semantic information is to be combined, and (as our NP rule for proper names shows) this may require rule-specific Prolog tricks such as variable doubling. Moreover, it is hard to think about the resulting grammar in a modular way. For example, when we explained why the NP rules took the form they do, we did so by explaining what was eventually going to happen when the S rule was used. Now, perhaps we weren't *forced* to do this—nonetheless, we find it difficult to give an intuitive explanation of our rules on an individual basis.

This suggests we are missing something. Maybe a more disciplined approach to missing information would reduce—or even eliminate—the need for rule-specific combination methods? Indeed, this is exactly happens if we make use of the *lambda calculus*.

Exercise 2.2.2 Using either pen and paper or a tracer, compare the sequence of variable instantiations this program performs when building representations for ‘Vincent snorts’ and ‘A woman snorts’, and ‘Vincent loves Mia’ and ‘Vincent loves a woman’.

2.3 The Lambda Calculus

For present purposes we shall view lambda calculus as a notational extension of first order logic that allows us to bind variables using a new variable binding operator λ . Occurrences of variables bound by λ should be thought of as placeholders for missing information: they *explicitly* mark where we should substitute the various bits and pieces obtained in the course of semantic construction. An operation called β -conversion performs the required substitutions. We suggest that the reader think of the lambda calculus as a special programming language dedicated to a single task: gluing together the items needed to build semantic representations. This glue language turns out to remarkably flexible: it is compatible with the underspecified representations introduced in Chapter 3, and with the Discourse Representation Structures studied in Part II.

The lambda operator marks missing information by binding variables. Here is a simple lambda expression:

$$\lambda x.MAN(x)$$

Here the prefix $\lambda x.$ binds the occurrence of x in $MAN(x)$. We sometimes say that the prefix $\lambda x.$ *abstracts over* the variable x . We call expressions with such prefixes *lambda abstractions* (or, more simply, *abstractions*). In our example, the binding of the free x variable in $MAN(x)$ explicitly indicates that MAN has an argument slot where we may perform substitutions. More generally, the purpose of abstracting over variables is to mark the slots where we want substitutions to be made.

Concatenation indicates that we wish to perform substitution. We’re using a special symbol “@” to indicate concatenation. Consider the following expression:

$$\lambda x.MAN(x)@VINCENT.$$

This compound lambda expression consists of the abstraction $\lambda x.MAN(x)$ written immediately to the left of the expression $VINCENT$, glued together by @ (we’re using an infix notation for the @-operator). Such concatenations are called *functional applications*; the left-hand expression is called the *functor*, and the right-hand expression the *argument*. Such a concatenation is an instruction to throw away the $\lambda x.$ prefix of the functor, and to replace every occurrence of x that was bound by this prefix by the argument. We call this substitution process *β -conversion* (other common names include β -reduction and lambda-conversion). Performing the β -conversion demanded in the previous example yields:

MAN(VINCENT).

Abstraction, functional application, and β -conversion underly much of our subsequent work. The business of specifying semantic representations for lexical items is essentially going to boil down to devising lambda abstractions that specify the missing information, while functional application coupled with β -conversion will be the engine used to combine semantic representations compositionally.

The above example was rather simple, and in one way rather misleading. As our previous experiments have made clear, to deal adequately with noun phrases and determiners (and indeed, many other things) we need to mark more complex kinds of information than that represented by individual variables. We shall use λ for this task as well. For example, our semantic representation of the noun phrase ‘a woman’ will be:

$$\lambda Q.\exists x(\text{WOMAN}(x)\wedge Q@x).$$

Here we are using the variable Q to indicate that some information is missing (namely, the nuclear scope, to use the linguistic terminology mentioned earlier) and to show where this information has to be plugged in when it is found (it will be the conjoined to $\text{WOMAN}(x)$). We’ll use new variable symbols (typically, capital letters such as P , Q , X , and Y) for variables which act as place holder for complex information, and we’ll call these new variables *complex variables*.

We are almost ready to examine some linguistic examples, but let’s first clarify one point. The lambda expressions $\lambda x.\text{MAN}(x)$, $\lambda y.\text{MAN}(y)$, and $\lambda z.\text{MAN}(z)$ are equivalent, as are the expressions $\lambda Q.\exists x(\text{WOMAN}(x)\wedge Q@x)$ and $\lambda Y.\exists x(\text{WOMAN}(x)\wedge Y@x)$. All these expressions are functors which when applied to an argument, replace the bound variable by the argument. No matter which argument \mathcal{A} we choose, the result of applying any of the first three expressions to \mathcal{A} and then β -converting is $\text{MAN}(\mathcal{A})$, and the result of applying either of the last two expressions to \mathcal{A} is $\exists x(\text{WOMAN}(x)\wedge \mathcal{A}@x)$. In short, relabeling bound variables yields lambda expressions which carry out exactly the same glueing task.

The process of relabeling bound variables is called α -conversion. If a lambda expression \mathcal{E} can be obtained from a lambda expression \mathcal{E}' by α -conversion then we say that \mathcal{E} and \mathcal{E}' are α -equivalent. (Thus $\lambda x.\text{MAN}(x)$, $\lambda y.\text{MAN}(y)$, and $\lambda z.\text{MAN}(z)$ are all α -equivalent, as are $\lambda Q.\exists x(\text{WOMAN}(x)\wedge Q@x)$ and $\lambda Y.\exists x(\text{WOMAN}(x)\wedge Y@x)$.) In what follows we often treat α -equivalent expressions as if they were identical. For example, we will sometimes say that the lexical entry for some word is a lambda expression \mathcal{E} , but when we actually work out some semantic construction, we might use an α -equivalent expression \mathcal{E}' instead of \mathcal{E} itself. As λ -bound variables are merely placeholders for substitution slots, this is clearly sensible. But the reader needs to understand that it’s not merely *permissible* to do this, it can be *vital* to do so if β -conversion is to work as intended.

Suppose that the expression \mathcal{F} in $\lambda V.\mathcal{F}$ is a complex expression containing many λ operators. Now, it could happen that when we apply a functor $\lambda V.\mathcal{F}$ to an argument \mathcal{A} , some

occurrence of a variable that is free in \mathcal{A} becomes bound by a lambda operator when we substitute it into \mathcal{F} . *We don't want this to happen.* Such accidental bindings (as they are usually called) defeat the purpose of working with the lambda calculus. The whole point of developing the lambda calculus was to gain control over the process of performing substitutions. We don't want to lose control by foolishly allowing unintended interactions.

Such interactions need never be a problem. We don't need to use $\lambda V.\mathcal{F}$ as the functor; any α -equivalent formula will do. By suitably relabeling the bound variables in $\lambda V.\mathcal{F}$ we can always obtain an α -equivalent functor that doesn't bind any of the variables that occur free in \mathcal{A} , and accidental binding is prevented. Thus, strictly speaking, it is not merely functional application coupled with β -conversion that drives the process of semantic construction in this book, but functional application and β -conversion coupled with (often tacit) use of α -conversion.

That's all we need to know about the lambda calculus for now—though we will mention that lambda calculus can be introduced from a different, more mathematically oriented, perspective. Now, the mathematical perspective is useful (we discuss it briefly in the Notes at the end of the chapter, and present it in depth in Appendix B) but for all its importance, it is *not* the only legitimate perspective on lambda calculus. The computational perspective we have adopted—lambda calculus as glue language—is equally important, and underpins a great deal of work in computational semantics. So let's save the theoretical work for later, and try putting our new tool to work. Here's a good place to start: does lambda calculus solve the problem we started with? That is, does it get rid of the difficulties we encountered in experiments 1 and 2?

Let's see what's involved in building the semantic representation for 'every boxer walks' using lambdas. The first step is to assign lambda expressions to the different basic syntactic categories. We assign the determiner 'every', the common noun 'boxer', and the intransitive verb 'walks' the following lambda expressions:

'every': $\lambda P.\lambda Q.\forall x(P@x \rightarrow Q@x)$

'boxer': $\lambda y.BOXER(y)$

'walks': $\lambda x.WALK(x)$.

Before going further, pause a moment. These expressions should remind the reader of something, namely the representations used in experiment 2. For example, in experiment 2 we gave the determiner 'every' the representation

`det(X, Restr, Scope, forall(X, Restr > Scope))`


If we use the Prolog variable `P` instead of `Restr`, and `Q` instead of `Scope` this becomes

$$\text{det}(X,P,Q,\text{forall}(X,P \rightarrow Q))$$

which is clearly analogous to $\lambda P.\lambda Q.\forall x(P@x \rightarrow Q@x)$.

But there are also important differences. The experiment 2 representation is a Prolog-specific encoding of missing information. In contrast, lambda expressions are programming language independent (we could work with them in C++ or Lisp, for example). Moreover, experiment 2 “solved” the problem of combining missing information on a rule-by-rule basis. As will soon be clear, functional application and β -conversion provide a completely general solution to this problem.

Let’s return to ‘every boxer walks’. According to our grammar, a determiner and a common noun can combine to form a noun phrase. Our semantic analysis couldn’t be simpler: we will associate the NP node with the functional application that has the determiner representation as functor and the noun representation as argument.



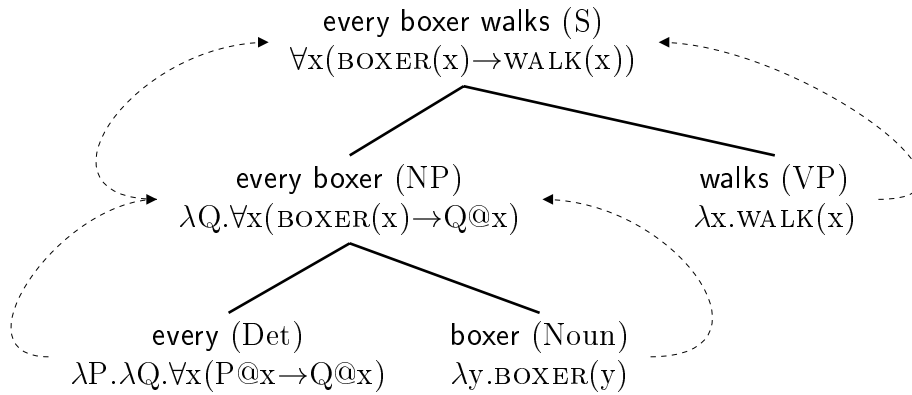
Now, applications are instructions to carry out β -conversion, so let’s do what is required. (Note that as there are no free-occurrences of variables in the argument expression, there is no risk of accidental variable capture, so we don’t need to α -convert the functor.) Performing the demanded substitution yields:

‘every boxer’: $\lambda Q.\forall x((\lambda y.BOXER(y))@x \rightarrow Q@x)$

But this expression contains a subexpression of the form $(\lambda y.BOXER(y))@x$. This is another instruction to perform β -conversion, and when we do so we obtain:

‘every boxer’: $\lambda Q.\forall x(BOXER(x) \rightarrow Q@x)$

We can’t perform any more β -conversions, so let’s carry on with the analysis of the sentence. The following tree shows the final representation we obtain:



Why is the S node associated with $\forall x(\text{BOXER}(x) \rightarrow \text{WALK}(x))$? It's certainly what we want, but where does it come from?

In fact we obtain it by a procedure analogous to that just performed at the NP node. First, we associate the S node with the application that has the NP representation just obtained as functor, and the VP representation as argument:

'every boxer walks': $(\lambda Q. \forall x(\text{BOXER}(x) \rightarrow Q@x))@(\lambda x. \text{WALK}(x))$.

Performing β -conversion yields:

'every boxer walks': $\forall x(\text{BOXER}(x) \rightarrow (\lambda x. \text{WALK}(x))@x)$.

We can then perform β -conversion on the subexpression $(\lambda x. \text{WALK}(x))@x$, and when we do so we obtain the desired representation:

'every boxer walks': $\forall x(\text{BOXER}(x) \rightarrow \text{WALK}(x))$.

It is worth reflecting on this example, for it shows that in two important respects semantic construction is getting simpler. First, the process of combining two representations is now uniform: we simply say which of the representations is the functor and which the argument, whereupon combination is carried out by applying functor to argument and β -converting. Second, more of the load of semantic analysis is now carried by the lexicon: it is here that we use the lambda calculus to make the missing information stipulations.

Are there clouds on the horizon? For example, while the semantic representation of a quantifying noun phrase such as 'a woman' can be used as a functor, surely the semantic representation of an NP like 'Vincent' will have to be used as an argument? We avoided this problem in experiment 2 by the variable doubling trick used in the NP rule for proper names—but that was a Prolog specific approach, incompatible with the use of lambda

calculus. Maybe—horrible thought!—we’re going to be forced to duplicate syntactic rules again, just as we did in experiment 1.

In fact, there’s no problem at all. The lambda calculus offers a delightfully simple functorial representation for proper names, as the following examples show:

‘Mia’: $\lambda P.P@MIA$

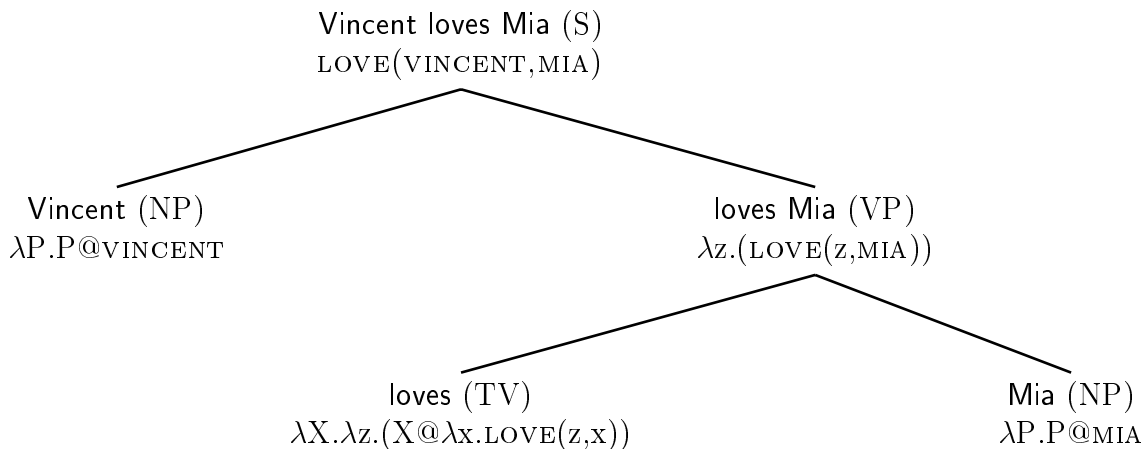
‘Vincent’: $\lambda P.P@VINCENT$

These representations are abstractions, thus they can be used as functors. However note what such functors do. They are essentially instructions to substitute their argument in P, which amounts to applying their own arguments to themselves! Because the lambda calculus offers us the means to specify such role-reversing functors, the specter of syntactic rule duplication vanishes.

As an example of these new representations in action, let us build a representation for ‘Vincent loves Mia’. We shall assume that TV semantic representations take their object NP’s semantic representation as argument, so we assign ‘loves’ the following lambda expression:

$\lambda X.\lambda z.(X@ \lambda x.LOVE(z,x))$.

And as in the previous example, the subject NP semantic representation takes the VP semantic representations as argument, so we can build the following tree:



Exercise 2.3.1 Work through the functional applications and β -conversions required to build the VP and S representations. Make sure you understand the role-reversing idea used in the TV semantic representation.

Let's sum up what we have achieved. Our decision to move beyond the approach of experiment 2 to the more disciplined approach of the lambda calculus was sensible. For a start, we don't need to spend any more time thinking about how to combine two semantic representations—functional application and β -conversion give us a general mechanism for doing so. Moreover, much of the real work is now being done at the lexical level; indeed, even the bothersome problem of finding a decent way of handling NP representations uniformly now has a simple lexical solution.

In fact, for the remainder of this book, the following version of the three tasks listed earlier will underly our approach to semantic construction:

Task 1 Specify a DCG for the fragment of natural language of interest.

Task 2 Specify semantic representations for the lexical items with the help of the lambda calculus.

Task 3 Specify the translation \mathcal{R}' of a syntactic item \mathcal{R} whose parts are \mathcal{F} and \mathcal{A} with the help of functional application and β -conversion. That is, specify which of the subparts is to be thought of as functor (here it's \mathcal{F}), which as argument (here it's \mathcal{A}) and then define \mathcal{R}' to be $\mathcal{F}'@A'$, where \mathcal{F}' is the translation of \mathcal{F} and A' is the translation of \mathcal{A} .

We must now show that the second and third tasks lend themselves naturally to computational implementation.

2.4 Implementing Lambda Calculus

Our decision to perform semantic construction with the aid of an abstract glue language (namely, lambda calculus) has pleasant consequences for grammar writing, so we would like to make the key combinatorial mechanisms (functional application and β -conversion), available as black boxes to the grammar writer. From a grammar engineering perspective, this is a sensible thing to do: when writing fragments we should be free to concentrate on linguistic issues.

In this section we build the required black box. In fact we shall build two versions. First we discuss a strikingly simple (but, alas, not fully correct) version which uses unification to simulate β -conversion. Then, simply by redefining one crucial predicate, we convert this to a correct implementation.

A Unification-Based Implementation

By making ‘clever’ use of Prolog variables, it is almost trivial to implement a black box for performing functional application and β -conversion. And once such a black box is available, we can decorate our little DCG with extremely natural semantic construction code and start building representations. But as the scare quotes indicate, our first implementation of β -conversion is not quite as clever as it seems, and we’ll show where things go wrong and why.

First, we have to decide how to represent lambda expressions in Prolog. Something as simple as the following will do:

```
lambda(L,F)
```

Here `L` is intended to be either a Prolog variable or the Prolog representation of a lambda expression, while `F` is either a first-order formula or the Prolog representation of a lambda expression.

Second, we have to decide how to represent concatenation. Let’s simply transplant our `@`-notation to Prolog by defining `@` as an infix operator:

```
:- op(950,yfx,@).
```

That is, we shall introduce a new Prolog operator `@` to explicitly mark where functional application is to take place: the notation `f @ a` will mean apply function `f` to argument `a`. We will build up our representations using these explicit markings, and delay carrying out β -conversion until all the required information is to hand.

Let’s see how to use this notation in DCGs. We’ll deal with the rules first. Actually, there’s practically nothing that needs to be said here. If we work with rules in the manner suggested by (our new version of) task 3, all we need is the following:

```
s(NP@VP)--> np(NP), vp(VP).
```

```
np(PN)--> pn(PN).
```

```
np(Det@Noun)--> det(Det), noun(Noun).
```

```
vp(IV)--> iv(IV).
```

```
vp(TV@NP)--> tv(TV), np(NP).
```

Note that the unary branching rules just percolate up their semantic representation (here coded as Prolog variables `NP`, `VP` and so on), while the binary branching rules use `@` to build a semantic representations out of their component representations in the manner suggested by task 3. Compared with the code in experiments 1 and 2, this is completely transparent: we simply apply function to argument to get the desired result.

The real work is done at the lexical level. The lexical entries practically write themselves:

```
noun(lambda(X,footmassage(X))--> [foot,massage].
```

```
noun(lambda(X,woman(X))--> [woman].
```

```
iv(lambda(X,walk(X))--> [walks].
```

And here's the code stating that $\lambda P.P@VINCENT$ is the translation of 'Vincent', and $\lambda P.P@MIA$ the translation of 'Mia':

```
pn(lambda(P,P@vincent))--> [vincent].
```

```
pn(lambda(P,P@mia))--> [mia].
```

Recall that the lambda expressions for 'every' and 'a' are $\lambda P.\lambda Q.\forall x.(P@x \rightarrow Q@x)$ and $\lambda P.\lambda Q.\exists x.(P@x \wedge Q@x)$. We express these in Prolog as follows.

```
det(lambda(P,lambda(Q,forall(X,(P@X)>(Q@X))))--> [every].
```

```
det(lambda(P,lambda(Q,exists(X,(P@X)&(Q@X))))--> [a].
```

Now, this makes semantic construction during parsing extremely easy: we simply use `@` to record the required function/argument structure. Here is an example query:

```
?- s(Sem,[mia,snorts],[]).
```

```
Sem = lambda(P,P@mia)@lambda(X,snort(X))
```

But of course, we need to do more work *after* parsing, for we certainly want to reduce these complicated lambda expressions into readable first-order formulas by carrying out β -conversion. The following code does this:


```

betaConvert(Var,Result):-
    var(Var), !, Result=Var.

betaConvert(Functor @ Arg,Result):-
    compound(Functor),
    betaConvert(Functor,ConvertedFunctor),
    apply(ConvertedFunctor,Arg,BetaConverted), !,
    betaConvert(BetaConverted,Result).

betaConvert(Formula,Result):-
    compose(Formula,Functor,Formulas),
    betaConvertList(Formulas,ResultFormulas),
    compose(Result,Functor,ResultFormulas).

```

The first clause of `betaConvert/2` simply records the fact that variables cannot be further reduced (the cut (!) prevents that in this case the other clauses are entered).

The second clause does the most important things: it checks whether the functor is a complex term and, if this is the case, reduces it to a lambda expression (of course it may already be a suitable lambda expression, but it could perfectly well be an application that first had to be reduced). If that succeeds, it applies the converted functor to `Arg` using `apply/3`. The result is reduced as well, as there can be an instruction to convert embedded inside it.

The third and final clause breaks down formulas and predicates and reduces their arguments or subformulas. This is done with the help of:

```

betaConvertList([],[]).
betaConvertList([Formula|Others],[Result|ResultOthers]):-
    betaConvert(Formula,Result),
    betaConvertList(Others,ResultOthers).

```

Now, everything we have done so far is perfectly correct—but we're now going to let ourselves go (temporarily) astray. Only one task remains: we have to define the `apply/3` predicate, which actually carries out the β -conversion. How can we do this? With clever use of Prolog unification we can do so with a single clause:

```

apply(Functor,Argument,Result):-
    Functor=lambda(Argument,Result).

```

The expression `lambda(Argument,Result)` is the functor, the variable `Argument` its argument, and `Result` is the result of the β -conversion. To see why it works, consider the following example. Suppose we wanted to apply $\lambda x.WALK(x)$ to `VINCENT`. We would do this in Prolog by means of the following query:

```
?- apply(lambda(X,walk(X)),vincent,Result).

X = vincent
Result = walk(vincent)

yes
```

Think about it. The definition of `apply` requires its first argument of the form `lambda(Argument,Result)`, and this is the case. As a result, `Argument` gets unified with `vincent`, and `Result` with `walk(Argument)`, hence `walk(vincent)` as final result. In fact, the definition of `apply/3` can even be slightly simplified to:

```
apply(lambda(Argument,Result),Argument,Result).
```

To finish off, let's define a driver predicate that calls the parser to analyze a sentence, reduces the resulting lambda expression into a first-order formula, and directs the result to the standard output:

```
parse:-
    readLine(Sentence),
    s(LambdaExpression,Sentence,[]),
    betaConvert(LambdaExpression,Formula),
    printRepresentation(Formula).
```

Exercise 2.4.1 [intermediate] Give the Prolog code for lexical entries of ditransitive verbs such as 'offer' in 'Vincent offers Mia a drink'.

Exercise 2.4.2 [intermediate] Find a suitable lambda expression for the lexical entry of the determiner 'no', and then give the corresponding Prolog code.

Exercise 2.4.3 [intermediate] Extend the DCG with a simple treatment of adjectives. Your grammar should be able to assign formulas to expressions like 'A pretty woman likes a tough guy'. Hint: first add the rule for adjectives that combine with nouns to the DCG, then think of how semantic representations for adjectives combine with representations for nouns, and finally implement the lexical entries for adjectives.

Exercise 2.4.4 [intermediate] Extend the DCG so that it covers negated sentences such as 'Vincent does not walk' or 'It is not the case that Vincent walks'.

Exercise 2.4.5 [hard] This exercise is about Pereira & Shieber's trick to provide a simple semantics for transitive verbs, i.e., `lambda(X,lambda(Y,love(X,Y)))` for the verb 'love'. The second part of the trick is to change the grammar rule that deals with transitive verbs:

$$\text{vp}(\text{lambda}(X, \text{NP@TV})) \rightarrow \text{tv}(\text{lambda}(X, \text{TV})), \text{np}(\text{NP}).$$

Explain why this works and discuss this particular treatment of transitive verbs.

Exercise 2.4.6 [easy] There is even a way to eliminate the @ instructions and make exclusive use of `lambda/2!` For example, the grammar rule that makes sentences out of noun phrases and verb phrases can be simplified to:

$$s(S) \rightarrow \text{np}(\text{lambda}(\text{VP}, S)), \text{vp}(\text{VP}).$$

Eliminate all @ instructions from the DCG rules given above.

A Substitution-Based Implementation

Our unification-based implementation of β -conversion (and in particular, the definition `apply/3`) is strikingly simple. Unfortunately it is not fully correct, and it fails for reasons directly relevant to semantic construction. We will go through an example to make the problem clear, and then reimplement `apply/3` properly.

So what's wrong? (It probably not obvious that there's *anything* wrong.) Well, our previous implementation is unable to hand *coordination* correctly. Consider the sentence 'Vincent and Mia dance'. We already know what the lambda expressions are for 'Vincent', 'Mia', and 'dance'. Unsurprisingly, we represent 'and' by the following lambda expression:

$$(2) \quad \lambda X. \lambda Y. \lambda P. (X @ P \wedge Y @ P).$$

This is a sensible representation. If we apply it to the representations for 'Mia' and 'Vincent' we get the following result:

$$(3) \quad \lambda P. ((\lambda Q. Q @ \text{VINCENT}) @ P \wedge (\lambda R. R @ \text{MIA}) @ P).$$

Applying β -conversion twice yields:

$$(4) \quad \lambda P. (P @ \text{VINCENT} \wedge P @ \text{MIA}).$$

That is, one and the same lambda expression `P` is to be applied to both `VINCENT` and `MIA`. And this is exactly what we need. In particular, applying this expression to `DANCE` yields the formula `DANCE(VINCENT) \wedge DANCE(MIA)`, which is the representation we would like to have.

So far so good. But now consider what happens when turn this into Prolog. First, the lexical entry for the coordinator 'and' will be as follows:

```
coord(lambda(X,lambda(Y,lambda(P,(X@P) & (Y@P))))--> [and].
```

Next, assume that we want to analyze ‘Vincent and Mia dance’ and that in the process of combining the lexical entries *X* unifies with the semantic representation of ‘Vincent’, *Y* with the semantic representation of ‘Mia’. Further, assume that *P* gets instantiated with `lambda(X,dance(X))`. Then the two `apply/3` statements, which both use a copy of *P*, try to force the occurrence of *X* to unify with both *mia* and *vincent*. As distinct atoms do not unify, this fails.

Exercise 2.4.7 Try to cover verb phrase coordination cases like ‘Vincent walks and talks’ in the unification-based approach. Does the unification problem that appears in noun phrase coordination arise here as well? Give an explanation for your observation.

The lesson is clear. Using Prolog unification to simulate β -conversion is a nice idea—but ultimately inadequate. We need to redefine `apply/3` so that it genuinely substitutes the argument expression for all free occurrences in the scope of the lambda expression, and keeps track of the bound and free variables in a formula.

In `comsemLib.pl` there is a version of the Sterling and Shapiro `substitute/4` predicate. This predicate takes a term, a variable, and a formula as its first three arguments, and returns in its fourth argument the result of substituting the term for each free occurrence of the variable in the formula. This is an important predicate (we shall use it again when we implement a first-order theorem prover) so let’s look at the way it’s defined:

```
substitute(Term,Var,Exp,Result):-
    Exp==Var, !, Result=Term.

substitute(_Term,_Var,Exp,Result):-
    \+ compound(Exp), !, Result=Exp.

substitute(Term,Var,Formula,Result):-
    compose(Formula,Funcor,[Exp,F]),
    member(Funcor,[lambda,forall,exists]), !,
    (
        Exp==Var, !,
        Result=Formula
    );
    substitute(Term,Var,F,R),
    compose(Result,Funcor,[Exp,R])
).

substitute(Term,Var,Formula,Result):-
```

```

compose(Formula, Functor, ArgList),
substituteList(Term, Var, ArgList, ResultList),
compose(Result, Functor, ResultList).

```

The `substituteList/4` predicate used in here recursively applies `substitute/4` down the length of the input list to produce the output list:

```

substituteList(_Term, _Var, [], []).

substituteList(Term, Var, [Exp|Others], [Result|ResultOthers]):-
    substitute(Term, Var, Exp, Result),
    substituteList(Term, Var, Others, ResultOthers).

```

Here is an example that the functionality of this predicate.

```

?- substitute(A,B,love(C,B),Result).

Result = love(C,A)

yes

```

This is exactly what we require. We can we redefine `apply/3` as follows:

```

apply(lambda(X,Formula),Argument,Result):-
    substitute(Argument,X,Formula,Result).

```

And that solves the problem.

Exercise 2.4.8 Check the new grammar rules using β -conversion, especially on cases of noun and verb phrase coordination.

Exercise 2.4.9 Note that we still use Prolog variables to represent first-order variables, although this is not important for the application predicates as we could use basic atoms instead. Why is it still advantageous? Hint: lexical entries.

Exercise 2.4.10 In Exercise 2.4.6 we saw how to dispense with `@` in our unification-based implementation. Can we eliminate `@` in the substitution-based implementation?

2.5 Grammar Engineering

The explicit notation for functional application and the implementation of β -conversion are the basic tools we shall work with in this book, so it is time to define a bigger grammar and start exploring computational semantics; but let's try to observe some basic principles of grammar engineering as we do so. That is, we should strive for a grammar that is *modular* (each component should have a clear role to play and a clean interface with the other components), *extendible* (it should be straightforward to enrich the grammar should the need arise) and *reusable* (we should be able to reuse a significant portion of the grammar, even when we change the underlying representation language).

Grammar engineering principles have strongly influenced the design of our grammars—and *not*, we would like to stress, purely for pedagogic reasons. We will be experimenting with a wide variety of semantic constructions techniques (for example, in the following chapter we will consider four different techniques for coping with scope ambiguities). Moreover, in Part II we will switch from first-order logic to Discourse Representation Structures as our underlying representation language. As we have learned (often the hard way) incorporating these changes, and keeping track of what is going on, requires a disciplined approach towards grammar design.

We have adopted a fairly simple three-level grammar architecture consisting of a collection of semantically annotated *rules*, a *lexicon*, and a set of *semantic macros*. The rules are DCG rules annotated with an extra slot which applies semantic functions to arguments using the @ operator. The rules given below and their annotations will *not* change in the course of the book; in particular, we will reuse them when we work with Discourse Representation Theory in Part II. The lexicon lists information about words belonging to most syntactic categories in an easily extractable form; again, this component will stay fixed throughout the book. Finally, we have the crucial semantic macros. This is a level at which we state what we have previously called 'lexical entries'. It is here that we will do most of our semantic work, and our modifications will largely be confined to this level.

The Rules

Here are the core DCG rules that we would *like* to use. These rules license a number of semantically important constructions, such as proper names, determiners, pronouns, relative clauses, the copula construction, and coordination. In addition, the first two rules let us form discourses by stringing together sentences; we'll need these rules in Part II.

d--> s, d.	noun--> noun, coord, noun.
d--> s.	vp--> vp, coord, vp.
s--> np, vp.	vp--> vbar(fin).
np--> np, coord, np.	vp--> mod, vbar(inf).

np--> det, nbar.	vbar(I)--> vbar(I), coord, vbar(I).
np--> pn.	vbar(I)--> tv(I), np.
np--> pro.	vbar(I)--> iv(I).
nbar--> nbar, coord, nbar.	vbar(fin)--> cop, np.
nbar--> noun.	vbar(fin)--> cop, neg, np.
nbar--> noun, pp.	pp--> prep, np.
nbar--> noun, rc.	rc--> relpro, vp.

(Readers unfamiliar with syntax may be wondering what NBAR and VBAR are. Essentially they are a level of phrase structure intermediate between the full phrasal level (NP and VP respectively) and the lexical level (N and V). There is a variety of evidence for the existence of this intermediate level and we refer the reader to any mainstream syntax textbook for further discussion. We introduce this intermediate level because it will be the source of some nice examples in Part II.)

However these are not quite the rules we shall actually use, for the following reason. The coordination rules are left-recursive, hence the standard Prolog DCG interpreter will loop when given this grammar. As we *do* want to give coordination examples, and as we *don't* want to implement a parser that deals with left-recursive rules, we're simply going to adopt the following simple fix which will make a limited form of coordination available to us. We'll add a auxiliary set of categories named **np2**, **np1**, **v2**, **v1**, etc. These auxiliary categories allow us to specify left-recursive rules to a certain depth of recursion. For example, the rules which have something to say about NPs will be replaced by the following:

```
s--> np2, vp2.
np2--> np1.
np2--> np1, coord, np1.
np1--> det, n2.
np1--> pn.
np1--> pro.
```

Similarly, the rules controlling nouns will become:

```
n2--> n1.
n2--> n1, coord, n1.
n1--> noun.
n1--> noun, pp.
n1--> noun, rc.
```

This is a brutal way of dealing with the problem, but it enables us to generate the examples we want without having to worry about left-recursion.

One other shortcoming should be mentioned. We implemented a limited amount of inflectional morphology—all our examples are relentlessly third-person present-tense. This is a shame (tense and its interaction with temporal reference is a particularly rich source of semantic examples), nonetheless, we shall not be short of interesting things to do.

But for all its shortcomings, this small set of rules assigns tree structures to an interesting range of English sentences or small discourses, including:

Mia knows every owner of a hash bar.
 Vincent or Mia dances.
 Every boxer that kills a criminal loves a woman.
 Vincent does not love a boxer or criminal that snorts.
 She does not love a boxer or criminal that snorts and dances.
 A boxer snorts. He collapses. Mia is a woman. Vincent is not a boxer.

Let's turn to the semantic annotations. Here the news is extremely pleasant. For a start, the required semantic annotations are utterly straightforward; they are simply the obvious “apply the function to the argument statements” expressed with the help of @. Here, for example, is how we semantically annotate `s --> np2, vp2.:`

`s(NP@VP)--> np2(NP), vp2(VP).`

That is, we simply apply the NP semantics to the VP semantics. None of the rules are much more complex than this. The rules for coordination are the most complex, but even these are straightforward as the following example shows:

`n2((C@N1)@N2)--> n1(N1), coord(C), n1(N2).`

The unary rules, of course, are even simpler for they merely pass the representation up to the mother node. For example:

`np1(NP)--> pn(NP).`

Finally, we've got the lexical rules. These are rules that apply to terminal symbols, the actual strings in the input of the parser, and need to call the lexicon to check if a string belongs to the syntactic category searched for.

`noun(Noun)-->`
`{`
`lexicon(noun,Sym,Phrase,_),`
`nounSem(Sym,Noun)`
`},`
`Phrase.`

In this example for nouns, the semantic macro `nounSem` is used to construct the actual semantic representation for a noun. Each lexical category is associated with such a macro, and this enables us to abstract away from specific types of structures. So we have set up the grammar rules in such a way that we are independent from the semantic theory we want to work with, and this is an attractive property as we are going to change the underlying semantic representations in the following chapters.

All these rules will work for us unchanged throughout the book. The reader can find a complete list in the file `englishGrammar.pl`. Furthermore, the file `mainLambda.pl`, that contains the unification- and substitution-based approaches to the lambda calculus of this chapter, already uses these rules in the way described here.

The Lexicon

Our lexicon lists information about the words belonging to most syntactic categories in a form useful for the interface component that we shall shortly define. The general format of a lexical entry is

```
lexicon(Cat,Sem,Phrase,Misc)
```

where `Cat` is the syntactic category, `Sem` the semantic information introduced by the phrase (normally a relation symbol or a constant, though sometimes we leave this field empty), `Phrase` the string of words that span the phrase, and `Misc` miscellaneous information depending on the type of entry. In particular, `Misc` lists gender information for nouns, pronouns, and proper names (we are going to classify nouns as `male`, `female` or `human`, or `nonhuman`, for this information will be important for pronoun resolution in Part II), and inflectional information for verbs.

Typical entries for intransitive verbs are:

```
lexicon(iv,collapse,[collapses],fin).
lexicon(iv,collapse,[collapse],inf).
lexicon(iv,dance,[dances],fin).
lexicon(iv,dance,[dance],inf).
```

Nouns, on the other hand, are listed in the following format:

```
lexicon(noun,boxer,[boxer],human).
lexicon(noun,bkburger,[big,kahuna,burger],nonhuman).
```

Typical entries for pronouns and determiners are as follows:

```

lexicon(det,_, [every], uni).
lexicon(det,_, [a], indef).

lexicon(pro,_, [he], male).
lexicon(pro,_, [she], female).

```

Note that these entries contain no semantic information. This is because the semantic contribution of pronouns and determiners is not simply a constant or predicate symbol, but rather a relatively complex expression that is dependent on the underlying representation language (quantifiers are dealt with very differently in the DRT based work of Part II, for example). Hence we shall specify the semantics of these categories in the semantic macros.

A small number of important words—in particular, copula and the verb phrase modifier construct ‘does not’, are *not* listed in the lexicon at all. This is because they are not associated with either a relation symbol or a constant and they are not marked for gender, thus their lexicon entry would simply consist of a **Phrase** entry. The semantic macros are the sole source of information about such words.

This way of setting up a lexicon offers natural expansion options. For example, if decided to develop a grammar that dealt with inflectional morphology (which we won’t do so in this book), it is just a matter of extending the general format of entries with one or more fields, and if we did this, it would be natural to list all words in the lexicon.

Refer to `englishLexicon.pl` for a complete listing of the lexicon.

Exercise 2.5.1 Find out how copula verbs are handled in the lexicon and grammar, and how the semantic representation for sentences like ‘Mia is a boxer’ and ‘Mia is not Vincent’ are generated.

The Semantic Macros

We now come to the most important part of the grammar, the lexicon/rules interface. Essentially, this component is where we state what we called semantic macros in the text. As the introduction of `apply/3` and `conjoin/3` has reduced the process of combining semantic representations to an elegant triviality, and as the only semantic information the lexicon supplies is the relevant constant and relation symbols, the interface is where the real semantic work will be done. Let’s consider some examples of semantic macros right away.

```

nounSem(Sym, lambda(X, Formula)) :-
    compose(Formula, Sym, [X]).

pnSem(Sym, lambda(P, P@Sym)).

```

```
tvSem(Sym,lambda(K,lambda(Y,K@lambda(X,Formula)))):-
  compose(Formula,Sym,[Y,X]).
```

The first macro, `nounSem/2`, builds a semantic representation for any noun given the predicate symbol `Sym`, turning this into a formula lambda abstracted with respect to a single variable (for example `lambda(X,Formula)`). The representation is built using `compose/3` to incorporate it into a the required lambda expression. The semantic macro for proper names (`pnSem/2`) is even simpler, and the one for transitive verbs (`tvSem/2`) is similar to that of the macro for nouns, apart from handling two variables rather than just one.

As we've already mentioned, the interface also contains self-contained entries for the determiners. Here they are:

```
detSem(uni,lambda(P1,lambda(P2,forall(X,(P1@X) > (P2@X))))).
```

```
detSem(indef,lambda(P1,lambda(P2,exists(X,(P1@X) & (P2@X))))).
```

These, of course, are just the old-style 'lexical entries' we are used to.

From now on, we will always use the lexicon and (with the exceptions of a few pedagogic explorations of other formats) the rules listed above too. To put it another way: *from now on the primary locus of change will be the semantic macros*. For example, it is here that we will develop treatments of quantifier scope, pronoun resolution, presupposition accommodation, and VP-ellipsis. For a complete listing of the macros we have just been discussing, see `semMacrosLambda.pl` (page 222), but we shall see many more types of semantic macros as we work our way through the book.

Software Summary of Chapter 2

`experiment1.pl` The code of our first experiment in semantic construction for a small fragment of English. (page 216)

`experiment2.pl` The second experiment in semantic construction. (page 218)

`mainLambda.pl` The main file for the implementation of the lambda calculus. Consults all necessary files and contains the driver for the parser. You can either choose to make use of the (simple but incorrect) unification-based implementation of application, or use the substitution based approach. (page 220)

`betaConversion.pl` The predicates that implement β -conversion. (page 221)

`semMacrosLambda.pl` The semantic macros for the lambda calculus. (page 222)

`englishLexicon.pl` Our standard English lexical entries. Contains entries for nouns, proper names, intransitive and transitive verbs, prepositions, and pronouns. (page 202)

`englishGrammar.pl` Our standard grammatical rules for a fragment of English. Rules cover basic sentences, noun phrases, relative clauses and modification of prepositional phrases, verb phrases, and a limited form of coordination. (page 207)

Notes

Compositionality is a simple and natural idea—and one capable of arousing an enormous amount of passion and controversy. Traditionally attributed to Gottlob Frege (the formulation “the meaning of the whole is a function of the meaning of its parts” is often called Frege’s principle) it received a precise mathematical formulation in the late 1960s. For detailed and accessible overview of the compositionality concept, the reader should consult Janssen 1997.

The idea of using the simply typed lambda calculus to specify the meanings of lexical entries, and using functional application as the basic mechanism for combining representations, is due to Richard Montague. Indeed, what we have presented in this chapter is simply a computational perspective on a very small part of Montague semantics.

Montague's original papers are well worth reading (they are collected in Thomason 1974) nonetheless they are dense and most readers will be better off approaching Montague semantics via either Dowty, Wall and Peters (1981) or volume 2 of Gamut (1991b). Both contain good, careful, textbook level expositions of Montague's key ideas. A more demanding (and correspondingly more rewarding) exposition is due to Janssen (see Janssen 1986a and Janssen 1986b). These are well worth reading, but rather hard to get hold of. If you see them, grab them. For a good overview of Montague's work, and an account of the major directions in which his work has been developed, Partee 1997 is a must. This account is rich in historic detail, and is an excellent starting point for deeper forays into the semantic literature.

The mathematical and logical literature on the lambda calculus is huge. Here are some of the more obvious points of entry. The bible of untyped lambda calculus is Barendregt 1984. Barendregt 1991 is a good starting point for material on typed systems. However these are very technical accounts, perhaps best suited for occasional reference. For a more approachable account of both typed and untyped systems, try Hindley and Seldin 1986. Turner 1997 is fairly technical, but usefully broad, and notes applications in linguistics and computer science. Finally, for logical work on the systems Montague himself used, and a number of interesting variants, Gallin 1975 is indispensable.

The references given so far approach lambda calculus and its applications in natural language semantics from a logical perspective. In this book we have tried to emphasize that there is a very natural computational perspective too. Not only is the lambda calculus a useful tool for gluing representations together, but the basic idea emerges, with seeming inevitability, when one sits down and actually tries to do semantic construction in Prolog—or at least, that is what we have tried to suggest by approaching lambda calculus via experiments 1 and 2. We hasten to add that this 'seeming inevitability' is clear only with the benefit of hindsight. The links between the ideas of logic programming and Montague semantics seem to have first been explicitly drawn in Pereira and Shieber 1987.

Lambda calculus is probably the mostly widely used tool for semantic construction, and it is the tool of choice in this book—nonetheless, it is *not* the only tool available. Indeed, it has a most interesting rival: the use of *feature structures* and *feature structure unification*. The reader who wants to go further in computational semantics really should be acquainted with this approach. Gazdar and Mellish 1989 contains a good text book level introduction to its use in building syntactic and semantic representations, and indeed, on its uses in pragmatics too. Another readable paper on the topic is Moore 1989. This paper compares the approach with the use of lambda calculus. More recently, linear logic has been proposed as a suitable glue language for semantic construction; see Dalrymple et al. 1997.

Refer to classical DCG-grammar with semantic construction (Pereira?).

Chapter 3

Underspecified Representations

This chapter develops methods for dealing with an important semantic phenomenon: *scope ambiguities*. Sentences with scope ambiguities are often semantically ambiguous (that is, they have at least two non-equivalent first-order representations) but fail to exhibit any syntactic ambiguity. As our approach to semantic construction is firmly based on the idea of using syntactic structure to guide semantic construction, we face an obvious problem here: if there is no syntactic ambiguity, we will only be able to build one of the possible representations. As scope ambiguities are very common, we need to develop ways of coping with them right away.

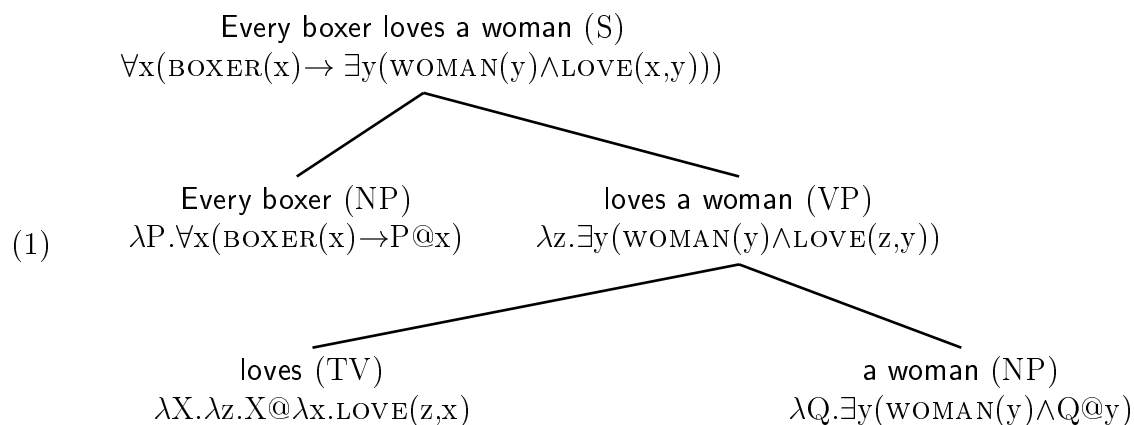
We are going to investigate and implement four different approaches to scope ambiguities: Montague’s original method, two storage based methods, and finally a more recent approach based on the idea of *underspecified representations*. Each approach is interesting in its own right, and the underspecification-based approach we shall eventually describe is a powerful tool that we shall use throughout the book. But as well as developing practical solutions to a pressing problem, this chapter also tells an important story. Computational semanticists are adopting an increasingly abstract perspective on what representations are and how they should be built. Once we have studied the evolutionary line leading from Montague’s method to contemporary underspecification-based methods, we will be in a better position to appreciate why.

3.1 Scope Ambiguities

Scope ambiguity is a common phenomenon and can arise from a wide variety of sources. In this chapter we will mostly be concerned with *quantifier scope ambiguities*. These are ambiguities that arise in sentences containing more than one quantifying noun phrase; for example, ‘Every boxer loves a woman’.

Now, the methods of the previous chapter allow us to assign a representation to this

sentences as follows:



The first-order formula we have constructed states that for each boxer there is a woman that he loves; there might be different women for different boxers. However, ‘Every boxer loves a woman’ has a second meaning (or to use the linguistic terminology, a second *reading*) that is captured by the following formula:

$$(2) \quad \exists y(\text{WOMAN}(y) \wedge \forall x(\text{BOXER}(x) \rightarrow \text{LOVE}(x,y)))$$

This says that there is one woman who is loved by all boxers.

It is clear that these readings are somehow systematically related, and that this relation has something to do with the relative scopes of the quantifiers: both first-order representations have the same components, but somehow the parts contributed by the two quantifying noun phrases have been shuffled around. In the first representation the existential quantifier contributed by ‘a woman’ has ended up inside the scope of the universal quantifier contributed by ‘every boxer’ and in the second representation the nesting is reversed. In fact it is usual to say that in the first reading ‘every boxer’ *has scope over* (or *outscores*) ‘a woman’, while in reading (2) it is the other way around.

Unfortunately, these scoping possibilities are not reflected syntactically: the only plausible parse tree for this sentence is the one just shown. Thus while it makes good semantic sense to say that in reading (2) ‘a woman’ *outscores* ‘every boxer’, we can’t point to any syntactic structure that would explain why this scoping possibility exists. As each word in the sentence is associated with a fixed lambda expression, and as semantic construction is simply functional application guided by the parse tree, this means there is no way for us to produce this second reading. This difficulty clearly strikes at the very heart of our semantic construction methodology. Moreover, scope ambiguities are extremely common, so we urgently need a solution.

In this chapter we examine in detail four increasingly sophisticated (and increasingly abstract) approaches to the problem. The first of these, Montague’s original method, introduces some important ideas, but as it relies on the use of additional rules it isn’t compatible with the approach to grammar engineering adopted in this book. However by introducing a more abstract form of representation—the *store*—we will be able to capture Montague’s key ideas. Stores were introduced by Robin Cooper and are arguably the earliest example of *underspecified representations*. Nonetheless, stores are relatively concrete. By simultaneously moving to more abstract underspecified representations, and replacing the essentially generative perspective underlying storage methods with a constraint based perspective, we arrive at *hole semantics*, the underspecified representation we shall use in this book. This will enable us to handle quantifier scope ambiguities, and the scope ambiguities created by constructs such as negation, in a uniform way.

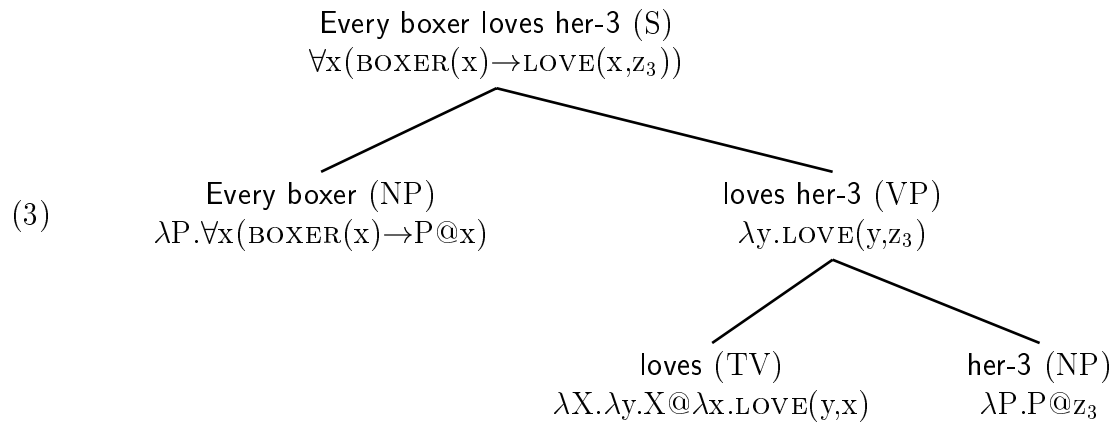
Exercise 3.1.1 Some readers may be wondering whether scope ambiguities really are a genuine problem at all. Consider ‘Every boxer loves a woman’. In a sense, representation (1) is sufficient to cover both readings of our example, since it is entailed by the other reading (2). Arguably, that makes the stronger representation superfluous: perhaps it is pragmatically inferred from the weaker one with the help of contextual knowledge. So do we really need techniques for coping with quantifier ambiguity?

We do, as we want the reader to demonstrate here. First, find examples of quantifier scope ambiguities that give rise to logically *independent* readings (that is, neither implies the other). Secondly, find an example where direct construction doesn’t lead to the weakest reading (that is, the reading entailed by the others). Feel free to use determiners such as ‘one’, ‘many’ and ‘most’ to construct your examples.

3.2 Montague’s Approach

Montague semantics makes use of the direct method of constructing semantic representations for quantified NPs studied in the previous chapter. However, motivated in part by quantifier scope ambiguities, Montague also introduced a rule of quantification (often called *quantifier raising*) that allowed a more indirect approach. The syntactic idea involved is very simple. Instead of directly combining syntactic entities with the quantifying noun phrase we are interested in, we are permitted to choose an ‘indexed pronoun’ and to combine the syntactic entity with the indexed pronoun instead. Intuitively, such indexed pronouns are ‘placeholders’ for the quantifying noun phrase. When this placeholder has moved high enough in the tree to give us the scoping we are interested in, we are permitted to replace it by the quantifying NP of interest.


As an example, let’s consider how to analyze ‘Every boxer loves a woman’. Here’s the first part of the tree we need:



Instead of combining ‘loves’ with the quantifying term ‘a woman’ we have combined it with the placeholder pronoun ‘her-3’. This pronoun bears an ‘index’, namely the numeral ‘3’. The placeholder pronoun is associated with a ‘semantic placeholder’, namely $\lambda P.P @ z_3$. As we shall see, it is the semantic placeholder that does most of the real work for us. Note that the pronoun’s index appears as subscript on the free variable in the semantic placeholder. From a semantic perspective, choosing an indexed pronoun really amounts to opting to work with the semantic placeholder (instead of the semantics of the quantifying NP) and stipulating which free variable the semantic placeholder should contain.

Now, the key point the reader should note about this tree is how *ordinary* it is. True, it contains a weird looking pronoun ‘her-3’—but, that aside, it’s just the sort of structure we’re used to. For a start, the various elements are syntactically combined in the expected way. Moreover semantic construction is also being performed completely standardly: the placeholder pronoun has the semantics $\lambda P.P @ z_3$ —which is the standard pronoun semantics—and (as the reader should check) the representations for ‘loves her-3’ and ‘every boxer loves her-3’ are constructed using functional application just as we discussed in the previous chapter. In short, we have ‘raised’ both the syntactic and semantic placeholders high into the tree, and we have done so in a completely orthodox way.

Now for the next step. We want to ensure that ‘a woman’ outscopes ‘every boxer’. By using the placeholder pronoun ‘her-3’, we have delayed introducing ‘a woman’ into the tree. But ‘every boxer’ is now firmly in place, so if we replaced ‘her-3’ by ‘a woman’ we would have the desired scoping relation. Predictably, there is a rule that lets us do this: given a quantifying NP, and a sentence containing a placeholder pronoun, we are allowed to construct a new sentence by substituting the quantifying NP for the placeholder. In short, we are allowed to extend the previous tree as follows:




But what's happening semantically? We know what formula we want to be assigned to the top S node (namely, the formula 2) but how can we ensure it gets there? Let's think the matter through.

We want 'a woman' to take wide scope over 'every boxer' semantically. Hence we should use the semantic representation associated with 'a woman' as the function. (To see this, simply look at the form of its semantic representation. When we apply it to an argument and lambda convert, we will be left with an existentially quantified expression, which is what we want.) But what should its argument be? There is only one reasonable choice. It must be the representation associated with 'every boxer loves her-3' lambda abstracted with respect to z_3 :

(5) $\lambda z_3.\forall x(\text{BOXER}(x)\rightarrow\text{LOVE}(x,z_3))$

Why is this? Well, right at the bottom of the tree we made use of the semantic placeholder $\lambda P.P@z_3$. When we raised this placeholder up the tree using functional application, we were essentially 'recording' what the semantic representation of 'a woman' would have encountered if we had used it directly. (Remember, we did nothing unusual, either syntactically or semantically, during the raising process.) The formula $\forall x(\text{BOXER}(x)\rightarrow\text{LOVE}(x,z_3))$ is the record of these encounters. When we are ready to 'play back' this recorded information, we lambda abstract with respect to z_3 (thus indicating that this variable is the crucial one, the one originally chosen) and feed the resulting expression as an argument to the semantic representation of 'a woman'. Lambda conversion will glue this record into its rightful place, and, as the following tree shows, everything will work out just right:



That's it. Summing up, Montague's approach makes use of syntactic and semantic placeholders so that we can place quantifying NPs in parse trees at exactly the level required to obtain the desired scope relations. A neat piece of 'lambda programming' (we call it

Montague's trick) ensures that the semantic information recorded by the placeholder is re-introduced into the semantic representation correctly.

Enough theory; can we write a Prolog program which integrates the previous chapter's work on lambda-based semantic construction with quantifier raising? Up to a point, we can. Unfortunately, most of the work is going to have to take place in our DCG rules, and the result is not particularly attractive from a grammar engineering perspective. Still, let's go through it systematically; it's interesting, and we shall be able to re-use some ideas and notation in our later work on storage.

We first add a new grammar rule for noun phrases that introduces the semantic representation of a pronoun (that is, we add a 'placeholder rule'). The pronoun will be associated with an index, and this index, together with the ordinary semantic representation of the noun phrase, will be passed to the mother node of the syntactic tree. (We do this by making use of additional arguments to the DCG rules.)

The following DCG rule implements these ideas. `Det@Noun` is the semantics of the quantifying noun phrase, constructed in the usual way. The term `lambda(P,P@I)`, where `I` is the index, mimics the usual pronoun semantics (it's our semantic placeholder). The first argument of `np1` holds the list of raised quantifiers, together with their associated indexes. We call these elements *indexed binding operators*, and represent them as two-place terms of the form `bo(Sem, Index)`.

```
np1([bo(Det@Noun, I)], lambda(P, P@I)) --> det(Det), n2(Noun).
```

Note that we have used a Prolog list to store the indexed binding operators. This is because it makes it easy to introduce more indexes (if we have to) and straightforward to check whether any indexes were introduced at all.

The other clauses for noun phrases are essentially the ones we had in our original implementation, but we have to add an additional argument (namely, an empty list) to make them compatible with our new rule for quantifier raising.

```
np1([], Det@Noun) --> det(Det), n2(Noun).
```

```
np1([], NP) --> pn(NP).
```

```
np1([], NP) --> pro(NP).
```

Needless to say, the other rules in our grammar fragment need to be rewritten to take our extension into account. First we consider verb phrases. Intransitive verbs don't combine with NPs, so we simply give them the empty list as additional argument. Transitive verbs must pass on the list of raised quantifiers (coded as `Q`) that they get from the noun phrase:

$$v1([], I, V) \rightarrow iv(I, V).$$

$$v1(Q, I, TV@NP) \rightarrow tv(I, TV), np2(Q, NP).$$

We need to be able to replace the placeholder pronouns by the real quantifiers. This happens at the sentential level, thus we need a rule saying that a sentence can be constructed out of sentence with a raised quantifier, by abstracting over the index, and applying the semantics of the raised quantifier to the resulting expression.

$$s([], NP@lambda(I, S)) \rightarrow s([bo(NP, I)], S).$$

The standard rule for forming sentences remains unchanged, save we need (yet again) to add an extra argument and pass through the raised quantifiers.

$$s(Q, NP@VP) \rightarrow np2([], NP), vp2(Q, VP).$$

It is time to test our program (`mainMontague.pl`), using the `parse/0` predicate:

```
?- parse.
```

```
> Every boxer loves a woman.
```

```
Readings:
```

```
1 exists(A, woman(A) & forall(B, boxer(B) > love(B, A)))
```

```
2 forall(A, boxer(A) > exists(B, woman(B) & love(A, B)))
```

```
yes
```

```
?-
```

It gives us exactly the two readings that we want.

What can we say about Montague's approach? It seems fair to say that while the idea of using semantic placeholders regulated by Montague's trick is natural, the quantifier raising mechanism used to exploit these ideas isn't very appealing. For a start, we are placing a heavy burden on the grammar rules. In principal, grammar rules are there to tell us about syntactic structure—but now we're being forced to write additional rules that say something about scope ambiguity. Moreover, as will be apparent to readers who worked through our Prolog code, the obvious implementation simply isn't very attractive. All that fiddling about with extra arguments—and just to cope with sentences containing a transitive verb and two quantifying noun phrases in subject and object position! Extending the program to deal with a wide range of scope ambiguity inducing constructions would not be pleasant. So let us abandon quantifier raising, and turn instead to *storage methods*.

Exercise 3.2.1 Extend the quantifier raising analysis to ditransitive verbs, and check how many readings it gives for sentences like ‘A boxer gives every woman a foot massage’. Are these the readings you would expect?

Exercise 3.2.2 [hard] The file `englishGrammarMontague.pl` contains the revised DCG rules for Montague quantifier raising analysis. The rules for constructing complex noun phrases (such as ‘a man that told every joke’) don’t deal with raising yet. Repair this.

3.3 Storage Methods

Storage methods are an elegant way of coping with quantifier scope ambiguities: they neatly decouple scope considerations from syntactic issues, making it unnecessary to add new grammar rules. Moreover, both historically and pedagogically, they are the natural approach to explore next, for they draw on the key ideas of Montague’s approach (in essence, they exploit semantic placeholders and Montague’s trick in a computationally more natural way) and at the same time they anticipate key themes of modern underspecification-based methods.

Cooper Storage

Cooper storage is a technique developed by Robin Cooper for handling quantifier scope ambiguities. In contrast to the work of the previous section, semantic representations are built directly, without adding to the basic grammar rules. The key idea is to associate each node of a parse tree with a *store*, which contains a ‘core’ semantic representation together with the quantifiers associated with nodes lower in the tree. After the sentence is parsed, the store is used to generate scoped representations. The order in which the stored quantifiers are retrieved from the store and combined with the core representation determines the different scope assignments.

To put it another way, instead of simply associating nodes in parse trees with a single lambda expression (as we have done until now), we are going to associate them with a core semantic representation, together with the information required to turn this core into the kinds of representation we are familiar with. Viewed from this perspective, stores are simply a more abstract form of semantic representation—representations which encode, compactly and without commitment, the various scope possibilities; in short, they are a simple form of underspecified representation.

Let’s make these ideas precise. Formally, a store is an n -place sequence. We represent stores using the angle brackets \langle and \rangle . The first item of the sequence is the core semantic representation; it’s simply a lambda expression. (Incidentally, if we wanted to, we could insist that we’ve been using stores all along: we need merely say that when we previously

talked of assigning a lambda expression ϕ to a node, we *really* meant that we assigned the 1-place store $\langle \phi \rangle$ to a node.) Subsequent elements (if any) are pairs (β, i) , where β is the semantic representation of an NP (that is, another lambda expression) and i is an index. An index is simply a label which picks out a free variable in the core semantic representation. As we shall see, this index has the same purpose as the indexes we used for Montague-style raised quantifiers. We call pairs (β, i) *indexed binding operators*.

How do we use stores for semantic construction? Unsurprisingly, the story starts with the quantified noun phrases (that is, noun phrases formed with the help of a determiner). Instead of simply passing on the store assigned to them, quantified noun phrases are free to ‘re-package’ the information it contains before doing so. (Other sorts of NPs, such as proper names, aren’t allowed to do this.) More precisely, quantified noun phrases are free to make use of the following rule:

Storage (Cooper)

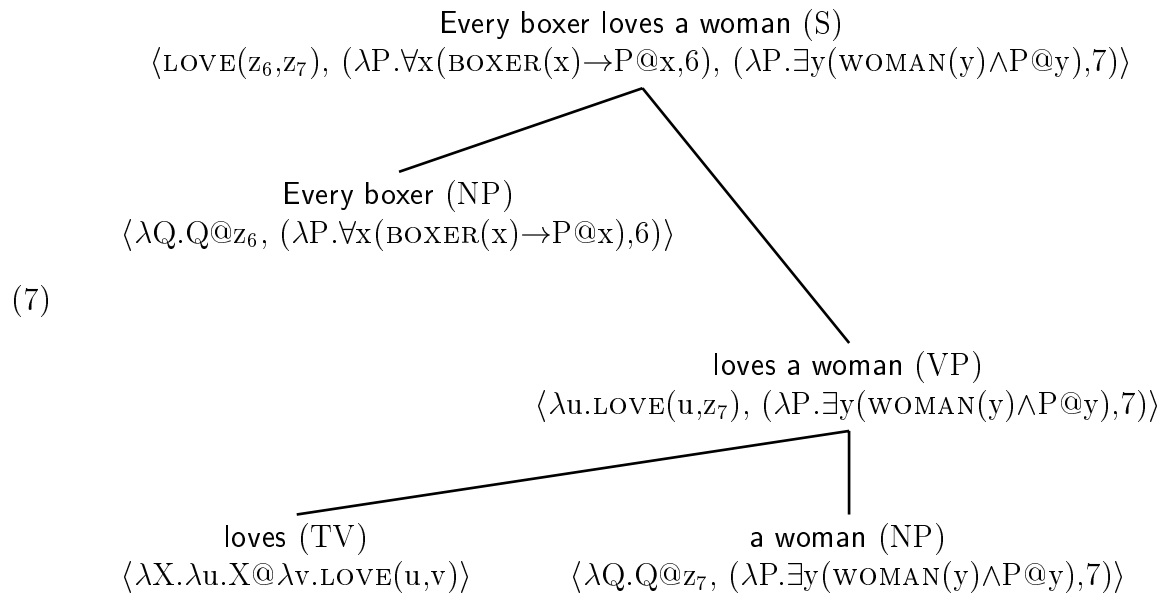
If the store $\langle \phi, (\beta, j), \dots, (\beta', k) \rangle$ is a semantic representation for a quantified NP, then the store $\langle \lambda P.P@z_i, (\phi, i), (\beta, j), \dots, (\beta', k) \rangle$, where i is some unique index, is also a representation for that NP.

The crucial thing to note is that the index associated with ϕ is identical with the subscript on the free variable in $\lambda P.P@z_i$. (After all, if we decide to store ϕ away for later use, it’s sensible to keep track of what its argument is.)

In short, from now on when we encounter a quantified NP we will be faced by a choice. We can either pass on ϕ (together with any previously stored information) straight on up the tree, or we can decide to use $\lambda P.P@z_i$ as the core representation, store the quantifier ϕ on ice for later use (first taking care to record which variable it is associated with) and pass on this new package. The reader should be experiencing a certain feeling of *deja vu*. We’re essentially using the pronoun representation $\lambda P.P@z_i$ as a semantic placeholder, just as we did in the previous section. Indeed, as will presently become clear, our shiny new storage technology re-uses the key ideas of quantifier storage in a fairly direct way.

Incidentally, the storage rule is *not* recursive. It offers a simple two way choice: either pass on the ordinary representation (that is, the store $\langle \phi, (\beta, j), \dots, (\beta', k) \rangle$) or use the storage rule to form $\langle \lambda P.P@z_i, (\phi, i), (\beta, j), \dots, (\beta', k) \rangle$ and pass this new store on up instead. We’re *not* offered—and we don’t want or need—the option of reapplying the storage rule to this new store to form $\langle \lambda P.P@z_m, (\lambda P.P@z_i, m), (\phi, i), (\beta, j), \dots, (\beta', k) \rangle$. Intuitively, we’re offered a straight choice between keeping the lambda expression associated with the quantified NP in the active part of the memory (that is, in the first slot of the store) or placing it, suitably indexed, in the freezer for later consumption.

It’s time for an example. Let’s analyze ‘Every boxer loves a woman’ using Cooper storage. Here’s (part of) the relevant tree.



Note that the nodes for ‘a woman’ and ‘Every boxer’ are both associated with 2-place stores. Why is this? Consider the node for ‘a woman’. We know from our previous work that the lambda expression associated with ‘a woman’ is $\lambda P. \exists y (\text{WOMAN}(y) \wedge P@y)$. In our new representations-are-stores world view, this means that the 1-place store

$$\langle \lambda P. \exists y (\text{WOMAN}(y) \wedge P@y) \rangle$$

is a legitimate interpretation for the NP ‘a woman’. But remember—this is not our only option. We are free to use the storage rule, and this is what we did when building the above tree. We’ve picked a brand new free variable (namely z_7), used the placeholder $\lambda Q. Q@z_7$ as the first item in the store, and ‘iced’ $\lambda P. \exists y (\text{WOMAN}(y) \wedge P@y)$, first recording the fact that z_7 is the variable relevant to this expression. Essentially the same story could be told for the NP ‘every boxer’, save that there we chose the new free variable z_6 .

Once this has been grasped, the rest is easy. In particular, if a functor node \mathcal{F} is associated with a store $\langle \mathcal{F}', (\beta, j), \dots, (\beta', k) \rangle$ and its argument node \mathcal{A} is associated with the store $\langle \mathcal{A}', (\beta'', l), \dots, (\beta''', m) \rangle$, then store associated with the node \mathcal{R} whose parts are \mathcal{F} and \mathcal{A} is

$$\langle \mathcal{F}'@ \mathcal{A}', (\beta, j), \dots, (\beta', k), (\beta'', l), \dots, (\beta''', m) \rangle.$$

That is, the first slot of the store really is the active part: it’s where the core representation is built. If you examine the above tree, you’ll see that the stores associated with ‘loves a woman’ and ‘every boxer loves a woman’ were formed using this ‘functional application in the first slot’ method. Note that in both cases we’ve simplified $\mathcal{F}'@ \mathcal{A}'$ using lambda conversion.

But we’re not yet finished. We now have a sentence, and this sentence is associated with an abstract unscoped representation (that is, a store), but of course, at the end of the day

we really want to get our hands on some ordinary scoped first-order representations. How do we do this?

This is the task of *retrieval*, a rule which is applied to the stores associated with sentences. Retrieval removes one of the indexed binding operators from the freezer, and combines it with the core representation to form a new core representation. (If the freezer is empty, then the store associated with the S node must already be a 1-place sequence, and thus we already have the expression we are looking for.) It continues to do this until it has used up all the indexed binding operators. The last core representation obtained in this way will be the desired scoped semantic representation (we hope!).

What does the combination process involve? Suppose retrieval has removed a binding operator indexed by i . It lambda abstracts the core semantic representation with respect to the variable bearing the subscript i , and then functionally applies the newly retrieved binding operator to the newly lambda-abstracted-over core representation. The result is the new core representation, and it replaces the old core representation as the first item in the store. More precisely:

Retrieval (Cooper)

Let σ_1 and σ_2 be (possibly empty) sequences of binding operators. If the store $\langle \phi, \sigma_1, (\beta, i), \sigma_2 \rangle$ is associated with an expression of category S, then the store $\langle \beta @ \lambda z_i. \phi, \sigma_1, \sigma_2 \rangle$ is also associated with this expression.

Hey—we're simply performing Montague's trick with the aid of stores!

Let's return to our example and apply the retrieval rule to the store associated with the S node. Now, this store contains two indexed binding operators. The retrieval rule allows us to remove either of them and to combine it with the core representation. Suppose we choose to first retrieve the quantifier for 'every boxer' (that is, the indexed binding operator in the second slot of the store). Then the retrieval rule tells us that the following store must be associated with the S node:

$$(8) \quad \langle \lambda P. \forall x (\text{BOXER}(x) \rightarrow P @ x) @ \lambda z_6. \text{LOVE}(z_6, z_7), (\lambda P. \exists y (\text{WOMAN}(y) \wedge P @ y), 7) \rangle$$

Using lambda conversion, this simplifies to:

$$(9) \quad \langle \forall x (\text{BOXER}(x) \rightarrow \text{LOVE}(x, z_7)), (\lambda P. \exists y (\text{WOMAN}(y) \wedge P @ y), 7) \rangle$$

No more lambda conversions are possible, but there's still a quantifier left in store. Retrieving it produces:

$$(10) \quad \langle \lambda P. \exists y (\text{WOMAN}(y) \wedge P @ y) @ \lambda z_7. \forall x (\text{BOXER}(x) \rightarrow \text{LOVE}(x, z_7)) \rangle$$

The result is the reading where ‘a woman’ outscopes ‘every boxer’, as becomes clear if we perform two more lambda conversions to obtain:

$$(11) \langle \exists y(\text{WOMAN}(y) \wedge \forall x(\text{BOXER}(x) \rightarrow \text{LOVE}(x,y))) \rangle$$

How do we get the other reading? We simply retrieve the quantifiers in the other order. We suggest that the reader attempts the following exercise immediately.

Exercise 3.3.1 Show the steps involved in applying retrieval that yield the reading where ‘every boxer’ has scope over ‘a woman’.

Let’s now implement Cooper storage in Prolog.

The first steps are pretty obvious: we’ll represent stores as lists, and indexed binding operators as terms of the form `bo(Quant, Index)`, just as we did in the previous section. Thus the following Prolog list represents a store:

```
[walk(X),bo(lambda(lambda(Y,F),forall(Y,boxer(Y) > F)),X)]
```

But now we need to think a little. The semantic representations we want to work with are stores, not just plain lambda expressions. Moreover, we want to reuse our lexicon and rules, and the idea of combining representations with the `@` operator is hard-wired into our rules. But we have said nothing about combining stores using `@`: do we face a problem here? No: it’s fairly obvious that we simply need to ‘lift’ our use of `@` so that we can meaningfully talk of combining stores with stores—and indeed, stores with quantifiers—using `@`. As stores are just lists of lambda expressions and binding operators, this shouldn’t be too difficult to do. But one matter deserves to be stressed: *the results of these ‘lifted’ uses of @ should always be stores*. This will guarantee that the sentence ends up being associated with a store, and hence that store-retrieval can proceed as explained above. Let’s define a predicate `buildStore/2` that carries out these combinations for us:

```
buildStore(quant(Quant) @ Store, NewStore):-
    buildStore(Store, [Arg|S]),
    npStorage(Quant, [Arg|S], NewStore).

buildStore(Store1 @ Store2, [F@A|S]):-
    buildStore(Store1, [F|S1]),
    buildStore(Store2, [A|S2]),
    append(S1, S2, S).

buildStore([Sem|Store], [Sem|Store]).
```

This predicate has three clauses. Now, the second clause simply combines two stores using @ in the obvious way (it makes use of the familiar fact that the first slot of the store really is where the core representation is built) and the result is clearly a new store. The third clause ends the recursion. So far so good—but what about the first clause, which lets us combine a quantifier and a store using @?

To understand what is going on here, let's examine the predicate `npStorage/3` that the first clause uses:

```
npStorage(Quant, [Arg|Store], [lambda(P, P@X), bo(Quant@Arg, X) | Store]).
```

That is, this predicate simply stores the quantifier in the expected way. But then it is clear that the first clause of `buildStore/2` simply says that combining a quantifier and a store using @ simply amounts to storing that quantifier—thus once again, our extended notion of functional application yields a new store.

Two other remarks should to be made. First, storage should be a possibility *only* available to quantified noun phrases, and this explains the presence of the marker `quant/1`. This marker (or feature) signals that the representation we have in our hands is of a genuine quantifier. Second, note that this version of `buildStore/2` makes storage *compulsory* for quantifiers, not *optional*. This turns out not to be a good idea, and we shall eventually modify our definition of `buildStore/2` accordingly—but more on that topic later.

With `buildStore/2` at our disposal, it now makes sense to talk of combining stores using @, and moreover we know that the result will always be a new store. So we are ready to define the semantic macros. Here's the one for a universal determiner:

```
detSem(uni, quant(lambda(P, lambda(Q, forall(X, (P@X)>(Q@X)))))).
```

Apart from the `quant`-marker, there is nothing really new in this entry—indeed it looks exactly like our previous semantic macros. As for the other categories, the only change we make is that the representations we work with are stores, as the following macros show:

```
tvSem(Sym, [lambda(K, lambda(Y, K@lambda(X, Formula))))):-  
  compose(Formula, Sym, [Y, X]).
```

```
pnSem(Sym, [lambda(P, P@Sym)].
```

Thus we have all that we need to create stores: all that remains is to implement retrieval.

The Prolog predicate that copes with retrieval is `sRetrieval/2`. Its first argument is the store, its second argument the derived scoped representation. If the store contains just one element, then this is the scoped formula; the first clause deals with this case.

Otherwise, it takes an element from the store using the predicate `removeFromStore/3`, lambda abstracts `Sem` (the first item in the list) with respect to the retrieved variable, and applies the retrieved quantifier to it, yielding a new semantic representation that will be β -converted later. Note that the `sRetrieval` predicate is recursive, thus it will eventually reduce the store to a list containing just one item, namely a scoped representation.

```
sRetrieval([S],S).
sRetrieval([Sem|Store],S):-
    removeFromStore(bo(Q,X),Store,NewStore),
    sRetrieval([Q@lambda(X,Sem)|NewStore],S).
```

But why does this generate *all* the possible scoped representations? In fact, we obtain all the possible readings thanks to `removeFromStore/3`. This predicate is defined in such a way that, if there is more than one element in the store, it succeeds when it removes *any* element at all from it. Therefore, `sRetrieval` produces (via the Prolog backtracking mechanism) all the scoped representation possible.

```
removeFromStore(X,[X|T],T).
removeFromStore(X,[Y|T],[Y|R]):-
    removeFromStore(X,T,R).
```

The driver is transparent: we use `buildStore` to build a store, `sRetrieval` to generate the various representations, and then the familiar β -conversion predicate to simplify:

```
parse:-
    readLine(Sentence),
    s(Sem,Sentence,[]),
    setof(Result,Store^Retrieved^(buildStore(Sem,Store),
                                sRetrieval(Store,Retrieved),
                                betaConvert(Retrieved,Result)),
          SemSet),
    printReadings(SemSet).
```

Again, it is time to test our program:

```
?- parse.

> Every boxer loves a woman.

Readings:
1 exists(A,woman(A)&forall(B,boxer(B)>love(B,A)))
2 forall(A,boxer(A)>exists(B,woman(B)&love(A,B)))
```

Unsurprisingly, these are exactly the results we want.

Summing up, Cooper storage is a generalization of Montague’s rule of quantification that makes use of abstract representations called stores. From the perspective of computational semantics, it has a distinct advantage over Montague’s approach: we don’t need to make use of additional grammar rules. Indeed, when you get down to it, all we really needed to do to implement Cooper storage was to define what it meant to combine our new representations using @ (this is what `buildStore/2` did for us) and then provide two additional predicates, `npStorage` and `sRetrieval`, which gave us the ability to manipulate our new representations in ways that plain old functional application couldn’t. Cooper storage is indeed a natural and useful technique.

Exercise 3.3.2 How many scoped representations are retrieved for ‘every piercing that is done with a needle is okay’? (Take ‘is done with’ as a two place relation, and view ‘is okay’ as a one place predicate.) Are they all correct?

Exercise 3.3.3 What does the store (before retrieval) for ‘every piercing is done with a needle’ look like? And after retrieval?

Exercise 3.3.4 In the original discussion of Cooper storage the storage of quantifiers was an *optional* alternative to functional application. In our implementation, all quantified NPs were forced to use storage. Why is there no difference between the two strategies for the given fragment of grammar. Can you think of natural language examples where it actually does differ?

Keller Storage

Cooper storage allows us a great deal of freedom in retrieving information from the store. We are allowed to retrieve quantifiers in any order we like, and the only safety net provided is the use of co-indexed variables and Montague’s trick.

Is this really safe? We haven’t spotted any problems so far—but then we’ve only discussed one kind of scope ambiguity, namely those in sentences containing a transitive verb with quantifying NPs in subject and object position. However there are lots of other syntactic constructions that give rise to quantifier scope ambiguities, for instance relative clauses (12) and prepositional phrases in complex noun phrases (13):

(12) Every piercing that is done with a gun goes against the entire idea behind it.

(13) Mia knows every owner of a hash bar.

Both examples give rise to scope ambiguities. For example, in (13) there is a reading where Mia knows all owners of (possibly different) hash bars, and a reading where Mia knows

all owners that own one and the same hash bar. Moreover, both examples contain nested NPs. In the first example ‘a gun’ is a sub-NP of ‘every piercing that is done with a gun’, while in the second, ‘a hash bar’ is a sub-NP of ‘every owner of a hash bar’.

We’ve never had to deal with nested NPs before. Is Cooper storage delicate enough to cope with them, and does it allow us to generate all possible readings? Let’s examine example (13) more closely and find out. This is the store:

$$\langle \text{KNOW}(\text{MIA}, z_2), (\lambda P. \forall y (\text{OWNER}(y) \wedge \text{OF}(y, z_1) \rightarrow P @ y), 2), (\lambda Q. \exists x (\text{HASHBAR}(x) \wedge Q @ x), 1) \rangle$$

Exercise 3.3.5 Verify this.

There are two ways to perform retrieval: by pulling the universal quantifier off the store before the existential, or vice versa. Let’s explore the first possibility. Pulling the universal quantifier off the store yields (after lambda conversion):

$$(14) \langle \forall y (\text{OWNER}(y) \wedge \text{OF}(y, z_1) \rightarrow \text{KNOW}(\text{MIA}, y)), (\lambda Q. \exists x (\text{HASHBAR}(x) \wedge Q @ x), 1) \rangle$$

Retrieving the existential quantifier then yields (again, after lambda conversion):

$$(15) \langle \exists x (\text{HASHBAR}(x) \wedge \forall y (\text{OWNER}(y) \wedge \text{OF}(y, x) \rightarrow \text{KNOW}(\text{MIA}, y))) \rangle$$

This states that there is a hash bar of which Mia knows every owner. This is one of the readings we would like to have. So let’s explore the other option. If we pull the existential quantifier from the S store first we obtain:

$$(16) \langle \exists x (\text{HASHBAR}(x) \wedge \text{KNOW}(\text{MIA}, z_2)), (\lambda P. \forall y (\text{OWNER}(y) \wedge \text{OF}(y, z_1) \rightarrow P @ y), 2) \rangle$$

Pulling the remaining quantifier off the store then yields:

$$(17) \langle \forall y (\text{OWNER}(y) \wedge \text{OF}(y, z_1) \rightarrow \exists x (\text{HASHBAR}(x) \wedge \text{KNOW}(\text{MIA}, y))) \rangle$$

But this is not at all we wanted! Cooper storage has given us not a sentence, but a formula containing the free variable z_1 . What is going wrong?

Essentially, the Cooper storage mechanism is ignoring the hierarchical structure of the NPs. The sub-NP ‘a hash bar’ contributes the free variable z_1 . However, this free variable does not stay in the core representation: when the NP ‘every owner of a hash bar’ is processed, the variable z_1 is moved out of the core representation and put on ice. Hence lambda abstracting the core representation with respect to z_1 *isn’t* guaranteed to take into account

the contribution that z_1 makes—for z_1 makes its contribution indirectly, via the stored universal quantifier. Everything is fine if we retrieve this quantifier first (since this has the effect of ‘restoring’ z_1 to the core representation) but if we use the other retrieval option it all goes horribly askew. Cooper storage doesn’t impose enough discipline on storage and retrieval, thus when it has to deal with nested NPs, it over-generates.

What are we to do? An easy solution would be to build a ‘free variable check’ into the retrieval process. That is, we might insist that we can only retrieve an indexed binding operator if the variable matching the index occurs in the core representation.

But this isn’t very principled; it deals with the symptoms, not the cause. The heart of the problem is that Cooper storage rides roughshod over the hierarchical structure of NPs; we should try to deal with this head on. (Incidentally, arguably there’s also an empirical problem: the free variable solution isn’t adequate if one extends the grammar somewhat. We won’t discuss this here but refer the reader to the Notes.)

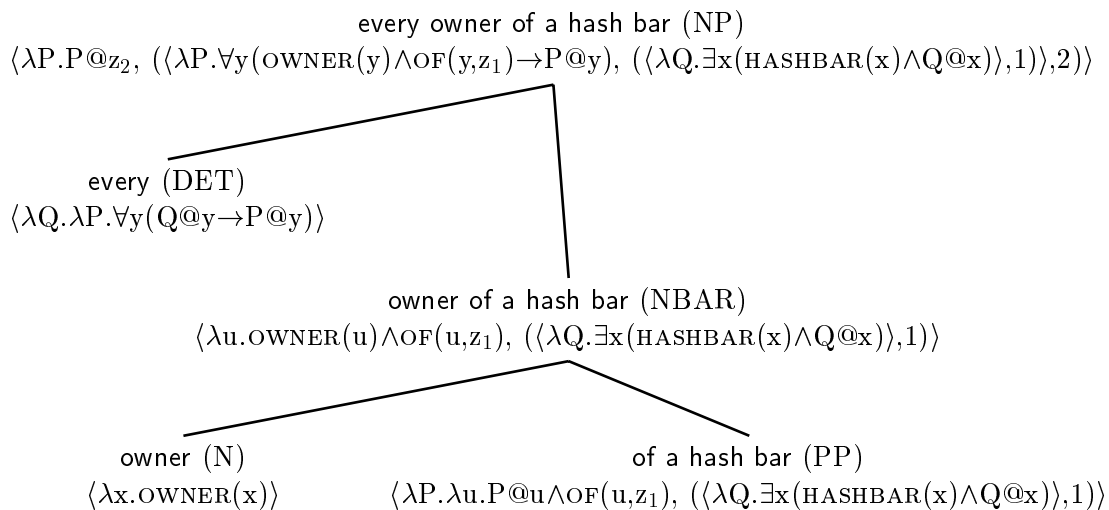
Here’s an elegant solution due to Bill Keller: allow nested stores. That is, allow stores to contain other stores. Intuitively, the nesting structure of the stores should automatically track the nesting structure of NPs in the appropriate way. As an added bonus, nesting is easier to implement than a free variable check.

Here’s the new storage rule:

Storage (Keller)

If the (nested) store $\langle \phi, \sigma \rangle$ is an interpretation for an NP, then the (nested) store $\langle \lambda P.P@z_i, (\langle \phi, \sigma \rangle, i) \rangle$, for some unique index i , is also an interpretation for this NP.

To see how this new storage rule works, consider how we assemble the representation associated with the complex noun phrase ‘every owner of a hash bar’.



As for the retrieval rule, it will now look like this.

Retrieval (Keller)

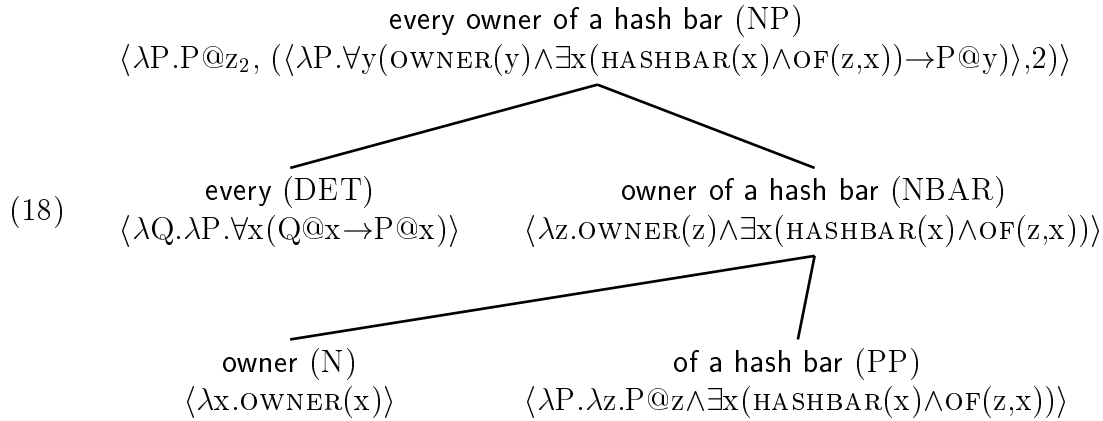
Let σ , σ_1 and σ_2 be (possibly empty) sequences of binding operators. If the (nested) store $\langle \phi, \sigma_1, (\langle \beta, \sigma \rangle, i), \sigma_2 \rangle$ is an interpretation for an expression of category S, then $\langle \beta @ \lambda z_i. \phi, \sigma_1, \sigma, \sigma_2 \rangle$ is too.

The new retrieval rule ensures that any operators stored while processing β become accessible for retrieval only after β itself has been retrieved. Nesting automatically overcomes the problem of generating readings with free variables. To see how it works in practice, let's return to our original example. The nested store associated with 'Mia knows every owner of a hash bar' is

$$\langle \text{KNOW}(\text{MIA}, z_2), (\langle \lambda P. \forall y (\text{OWNER}(y) \wedge \text{OF}(y, z_1) \rightarrow P @ y), (\langle \lambda Q. \exists x (\text{HASHBAR}(x) \wedge Q @ x \rangle, 1) \rangle, 2) \rangle$$

There is only one way to perform retrieval: first pulling of the universal quantifier, followed by the existential quantifier, resulting in (15). Since this is the only possibility, the unwanted reading (17) is not generated.

But wait a minute: how do we get the reading where Mia knows all owners of possible different hash bars? In fact, this couldn't be easier. All we have to do is avoid storing the sub-NP 'a hash bar'. If we do this, we can produce the following tree:



This leads to the following analysis for 'Mia knows every owner of a hash bar':

$$(19) \langle \text{KNOW}(\text{MIA}, z_2), (\langle \lambda P. \forall y (\text{OWNER}(y) \wedge \exists x (\text{HASHBAR}(x) \wedge \text{OF}(y, x)) \rightarrow P @ y \rangle, 2) \rangle$$

There is only one operator the store. Retrieving it yields the reading we want.

(20) $\langle \forall y(\text{OWNER}(y) \wedge \exists x(\text{HASHBAR}(x) \wedge \text{OF}(y,x)) \rightarrow \text{KNOW}(\text{MIA},y)) \rangle$

So our earlier Prolog implementation of Cooper storage—which made storage compulsory—really was linguistically naive. While making storage compulsory didn't have any malign effects in the previous section, this was only because we were working with a very restricted grammar fragment. As the previous example shows, as soon as we move to a richer fragments in which nested NPs are possible, storage really does need to be optional to ensure that all possible readings are generated.

A few rather trivial changes are all that's needed to turn our Cooper storage program into a Keller storage program. Here's the new version of `npStorage/3` we need. It's identical to the previous version, save for the third argument, which now treats stores as lists of lists:

```
npStorage(Quant, [Arg|Store], [lambda(P,P@X), bo([Quant@Arg|Store], X)]).
```

The only other change is in the first clause of `removeFromStore/3`, to enable it to deal with lists of lists:

```
removeFromStore(bo(0, I), [bo([0|T1], I)|T2], T):-
    append(T1, T2, T).
```

```
removeFromStore(X, [Y|T], [Y|R]):-
    removeFromStore(X, T, R).
```

And that's it. The over-generation problem is solved. We'll never see unwanted free variables again.

The last thing to do is thinking of a way to incorporate the optionality of the storage rule into the Prolog program. Naturally, we will implement this by making use of the following disjunctive rule.

```
buildStore(quant(Quant) @ Store, NewStore):-
    buildStore(Store, [Arg|S]),
    (
        NewStore = [Quant@Arg|S]
    ;
        npStorage(Quant, [Arg|S], NewStore)
    ).
```

Let's test our program (we'll use the `parse/0` predicate of the Cooper implementation) and see what happens:

?- parse.

> Mia knows every owner of a hash bar.

Readings:

```
1 exists(A,hashbar(A)&forall(B,of(B,A)&owner(B)>know(mia,B)))
2 forall(A,exists(B,hashbar(B)&of(A,B)&owner(A)>know(mia,A))
```

Let's sum up what we have learned. First, the original version of Cooper storage doesn't handle storage and retrieval in a sufficiently disciplined way, and this causes it to generate spurious readings (in fact, nonsensical readings) when faced with nested NPs. The problem can be elegantly cured by making use of nested stores, and implementing this idea requires only trivial changes to our earlier Prolog code. Second, storage really must be optional. If we are to generate all the required readings in examples involving nested NPs, we must be free not to store the NPs encountered low down in the hierarchy. If these changes are made, we have a flexible tool for handling quantifier scope ambiguities.

Exercise 3.3.6 Give the PL translations for the five readings of 'a man likes every woman with a five-dollar shake'. You might have noticed that the number of correct readings for this example is less than the combinatorial possibilities of the quantifiers that are involved. Naively, one would expect to have $3! = 6$ readings for this example. Why is one reading excluded? Try similar examples on the program.

Nonetheless, storage methods have their limitations. For a start, they are not as expressive as we might wish: although Keller storage predicts the five readings for

(21) One criminal knows every owner of a hash bar

it doesn't allow us to insist that 'every owner' must out-scope 'a hash bar', while at the same time leaving the scope relation between subject and object noun phrase unspecified. To put it another way, storage is essentially a technique which enables us to represent all possible meanings compactly; it doesn't allow us to express additional *constraints* on possible readings, and this is precisely what most modern underspecified representations let us do.

Moreover, storage is a technique specifically designed to handle *quantifier* scope ambiguities. Unfortunately, many other constructs (for example, negation) also give rise to scope ambiguities, and storage has nothing to say about these. Consider the sentence 'Vincent doesn't clean every car'. This has a reading where for each car it is not the case that Vincent cleans it; and a second, possibly preferred, reading where there are some cars (but not all) that Vincent doesn't clean. We would like an uniform approach to scope ambiguity, not a separate mechanism for each construct—and this as another motive for turning to the more abstract view offered by current work on underspecification.

3.4 Underspecification

There has been a great deal of recent interest in the use of *underspecified representations* to cope with scope ambiguities; so much so that it often seems as if semantics has entered an age of underspecification. With the benefit of hindsight, however, we can see that the idea of underspecified representations isn't really new. In our work on storage, for example, we associated stores, not simply lambda expressions, with parse tree nodes; and as we have already remarked, a store is in essence an underspecified representation. What *is* new is both the sophistication of the new generation of underspecified representations (as we shall see, they offer us a great deal of flexibility and expressive power), and, more importantly, the way such representations are now regarded by semanticists.

In the past, storage style (and other 'funny') representations seem to have been regarded with some unease. They were certainly a useful tool, but they appeared to live in a conceptual no-man's-land—not really semantic representations, but not syntax either—that was hard to classify. The key insight that underlies current approaches to underspecification is that it is both *fruitful* and *principled* to view representations more abstractly. That is, it is becoming increasingly clear that the level of representations is richly structured, that doing semantic construction elegantly demands deeper insight into this structure, and that—far from being a sign of some fall from semantic grace—semanticists should learn to play with representations in more refined ways.

In this book, we shall explore an approach to underspecification called *hole semantics*. We have chosen hole semantics because it is the approach we are most familiar with (the method is due to Johan Bos), it illustrates many of the key ideas of current approaches to underspecification in an elegant and simple way, and we will be able to use it when we switch to working with Discourse Representation Structures. Other approaches to underspecification are noted in the Notes at the end of the chapter.

Hole Semantics

Viewed from a distance, hole semantics shares a very obvious similarity with storage methods: at the end of the parsing stage, sentences *won't* be associated with a semantic representation. Rather, they will be associated with abstract representations from which the desired representation can be read off. Viewed closer up, however, it is clear that hole semantics adopts a far more radical perspective on representations than storage does.

Hole semantics is essentially a constraint-based approach to semantic representation. That is, at the end of the parsing process the method delivers a set of *constraints*: any first-order representation which fulfills these constraints—which govern how the various bits and pieces produced during the parsing process can be plugged together—is a permissible semantic representation for the sentence. This contrasts sharply to the essentially generative approach offered by storage (*Here's the store! Enumerate the readings like this!*) and

the source of much of the method's power.

Exactly what kinds constraints can we state? Answering this question is a two-step process.

1. We first need to *unplug* our underlying representation formalism (here, first-order logic) by permitting formulas to be built which contain *holes* (essentially variables over missing structure).
2. Then, with the help of such unplugged formulas, we can define a simple constraint language which governs the way semantic representations are *plugged together*.

(That is, we first break the underlying representation language down into convenient chunks, and then we define a simple constraint language governing how these chunks can be assembled into bigger units. The resulting representations we call underspecified semantic representations, USRs for short.)

Let's take care of the first-step right away. We'll now define what we mean by a PLU formula (this stands for *Predicate Logic Unplugged*). In Part II of the book we'll unplug the language of Discourse Representation Structures in a similar way.

Suppose we have chosen some vocabulary. (We'll assume that the vocabulary is function free; this is purely for simplicity and has no bearing on the work that follows.) We must then chose a countably infinite set H of holes. This is simply any convenient set of new symbols (that is, symbols distinct from those used to build the first-order formulas) which we will usually write as $h_0, h_1, h_2, h_3, \dots$, and so on. The hole h_0 will play a special role in the work that follows—it's used to state constraints, but it *won't* be used in PLU formulas. So let's give a name to the set containing just the ordinary holes (that is h_1, h_2, h_3, \dots , and so on) for these are the ones that are relevant at this stage. We'll use H^+ to denote this set and call it the set of *internal holes*. Note that $H = \{h_0\} \cup H^+$.

Now, what are the PLU-*terms* over this signature and H^+ ? Simply the ordinary first-order terms over this signature: the holes don't play any role at this level. That is, the terms we have at our disposal are simply the constants and variables.

Next, what are the *atomic PLU-formulas*? Again, simply the familiar ones; the holes don't play any role here either:

If R is a relation symbol of arity n , and t_1, \dots, t_n are terms, then $R(t_1, \dots, t_n)$ is a PLU-formula.

Now we are ready for the definition of PLU-*formulas*:

1. All atomic PLU formulas are PLU formulas.
2. If h is an internal hole, then h is a PLU formula.


3. If ϕ and ψ are PLU formulas, then $\neg\phi$, $(\phi \rightarrow \psi)$, $(\psi \vee \phi)$, $(\psi \wedge \phi)$ are PLU formulas.
4. If x is a variable, and ϕ is a PLU formula, then $\forall x\phi$ and $\exists x\phi$ are PLU formulas.
5. Nothing else is a PLU formula.

This completes the first stage in the definition of our constraint language. Let's turn to the second step. Our USRs (underspecified semantic representations) will be built using PLU-formulas, a special new symbol \leq , and a set of new symbols called *labels* which we will write as l_1, l_2, l_3, \dots , and so on. What kind of USRs can we form? Here's an example. As we have already mentioned, sentences are going to be associated with a USR. Here's the USR associated with 'Every boxer loves a woman':

$$(22) \quad < \left\{ \begin{array}{l} l_1 \\ l_2 \\ l_3 \\ h_0 \\ h_1 \\ h_2 \end{array} \right\}, \left\{ \begin{array}{l} l_1 : \forall x(\text{BOXER}(x) \rightarrow h_1) \\ l_2 : \exists y(\text{WOMAN}(y) \wedge h_2) \\ l_3 : \text{LOVE}(x,y) \end{array} \right\}, \left\{ \begin{array}{l} l_1 \leq h_0 \\ l_2 \leq h_0 \\ l_3 \leq h_1 \\ l_3 \leq h_2 \end{array} \right\} >$$

Everything the reader needs to know about constraints can be gleaned from this example. USRs have three components. The left hand side is a set of labels and holes used in the USR. The middle component consists of a set of PLU-formulas preceded by a label. The right hand side consists of a set of constraints of the form $l \leq h$, where l is a label and h is a hole. Note that the special hole h_0 occurs in the right-hand component of the USR.

But what does this USR mean? Essentially it tells us how we are permitted to plug together the labeled formulas to build a first-order formula. The special hole h_0 is a variable that stands for the first-order formula that we will eventually build. This formula will be built by plugging together three labeled PLU-formulas given in the middle component of the USR. Plugging together simply means substituting one of these 3 PLU formulas in another. (Incidentally, when we carry out such a substitution we only substitute the PLU formula, not its label as well. The label is simply an identifier that enables us to state the right hand side constraints succinctly). Any plugging will do—as long as it satisfies the constraints given on the right hand side. These constraints are essentially constraints on *subformulas*, and are best thought of visually. Each arrow in the following diagram corresponds to a \leq expression in an obvious way:



The constraints are set up in such a way that the basic PLU formula $\text{LOVE}(x,y)$ is forced to be out-scoped by the consequent of the universal quantifier's scope, and the second conjunct of the existential quantifier's scope.

Now it's time to say a little more about *resolving* these ambiguous representations. The holes underspecify scope, and in order to give us non-ambiguous interpretations, each hole should be *plugged* with a formula in such a way that all the constraints are satisfied. In other words, we should be sure that each hole gets associated with some label. Obviously, no label can be plugged into two different holes at the same time. Therefore, a so-called *plugging* is a one-to-one correspondence from holes to labels (i.e., a bijective function, with the set of holes as domain and the set of labels as codomain). A plugging for a proper USR is admissible if the instantiations of the holes with labels result in a representation in which there is no contradiction spelled out by the constraints.

There are two pluggings, as is easy to verify:

	h_0	h_1	h_2
P_1	l_1	l_2	l_3
P_2	l_2	l_3	l_1

Plugging P_1 interprets (22) as giving the universal quantifier wide scope, out-scoping the existential quantifier. The corresponding formula in predicate logic is (23), which is true in a model where all boxers love a woman, but not necessarily the same woman.

$$(23) \quad \forall x (\text{BOXER}(x) \rightarrow \exists y (\text{WOMAN}(y) \wedge \text{LOVE}(x,y)))$$

Plugging P_2 interprets (22) as the existential quantifier out-scoping the universal quantifier. In a model where there is some woman that is loved by all boxers, this interpretation denotes truth. A corresponding formula in predicate logic is (24).

$$(24) \quad \exists y (\text{WOMAN}(y) \wedge \forall x (\text{BOXER}(x) \rightarrow \text{LOVE}(x,y)))$$

The Plugging Algorithm

Before turning to the plugging algorithm we have to consider the design of Prolog representations for our PLU-formulas. Labels and holes are represented as Prolog variables (as with first-order variables, this gives us the practical advantage that any introduction of a label or hole appears as a new occurrence). With this we are able to represent PLU formulas. Further, we have `leq(L,H)` as the Prolog notation for the constraint $L \leq H$. And finally, labeled formulas are formed by putting together a label with a formula using the Prolog inbuilt `:` operator.

A USR itself is represented as `usr(D,L,C)`, where `D` is the set of labels and holes (the *Domain*), `L` a list of labeled formulas, and `C` a list of constraints. Here is a sample representation:

```
usr([L1,L2,L3,H0,H1,H2],
    [L1:exists(X,H1&H2),L2:boxer(X),L3:collapse(Y)],
    [leq(L1,H0),leq(L2,H1),leq(L3,H2),leq(L3,H0)])
```

This Prolog term is our representation in PLU for ‘A boxer collapses’. We won’t provide predicates to check for properness of USRs—instead we supply a plugging algorithm straightaway and assume that the USRs we work with are connected, have a unique top, and are acyclic.

The heart of the plugging algorithm is centered around `plugHole/4`. Its first argument is the hole to be plugged (normally we start with the top hole of the USR, and work recursively through the representation until we’ve plugged all holes). The second argument are the labeled formulas (coded as a difference list—used to keep track of which labeled formulas are left for plugging and which are not), and the third argument is the list of constraints imposed by the USR. Its last argument `Scope` is a list of formulas that, under the current reading, outscope `Formula`.

```
plugHole(Formula,LFs1-LFs3,Constraints,Scoped):-
    select(Label:Formula,LFs1,LFs2),
    checkConstraints(Label,Constraints,[Formula|Scoped]),
    variablesInTerm(Formula,[],Arguments),
    checkArguments(Arguments,LFs2-LFs3,Constraints,[Formula|Scoped]).
```

With `select/3` (see `comsemPredicates.pl`) we pick a labeled formula to be plugged and remove it from the list. This predicate is backtrackable—so if it turns out that we picked one that violates the constraints we can pick another instead later. We add `Formula` to the list of scoped formulas and check if no constraints are violated.

Then we continue plugging, by identifying the arguments of the newly plugged formula with `variablesInTerm/2`. This predicate, defined in the library file `comsemPredicates.pl`,

takes a term and returns all its variables. The rest of the work is done with the help of `checkArguments/4`. It traverses through the list of arguments, and according to the type of argument (recall that these can be holes, labels, or ordinary first-order variables), continues the plugging process.

The first clause of `checkArguments/4` takes care of holes. We identify holes by imposing a member-check on the constraints. Next this hole is plugged using `plugHole/4`, and the remaining arguments are checked.

```
checkArguments([Arg|Rest], LFs1-LFs3, Constraints, Scoped):-
    member(leq(_, Hole), Constraints),
    Hole==Arg, !,
    plugHole(Hole, LFs1-LFs2, Constraints, Scoped),
    checkArguments(Rest, LFs2-LFs3, Constraints, Scoped).
```

Clause 2 cares about labels. In that case the corresponding labeled condition is removed from the list of labeled formulas and its arguments are checked as there might be holes among them.

```
checkArguments([Arg|Rest], LFs1-LFs4, Constraints, Scoped):-
    select(Label:Formula, LFs1, LFs2),
    Label==Arg, !,
    Formula=Arg,
    variablesInTerm(Formula, []-Arguments),
    checkArguments(Arguments, LFs2-LFs3, Constraints, Scoped),
    checkArguments(Rest, LFs3-LFs4, Constraints, Scoped).
```

The third clause carries on the recursion if the argument is neither a hole nor a label, hence a first-order variable, which we don't need to bother about (note that the two clauses above use the cut preventing to enter this third clause).

```
checkArguments([_|Rest], LFs1-LFs2, Constraints, Scoped):-
    checkArguments(Rest, LFs1-LFs2, Constraints, Scoped).
```

Finally, there is a clause that ends the recursion:

```
checkArguments([], LFs-LFs, _, _).
```

A typical Prolog call to start the algorithm is `?- plugHole(Top, LFs-[], Cons, [])`, where `Top` is the hole that out-scopes all other formulas, and `LFs-[]` ensures that all labeled-formulas are 'used' for building the fully specified formula, and moreover, prevents us from

getting entangled in cyclic structures. If this goal succeeds, `Top` unifies with the resolved formula.

The constraints are checked each time after a hole is plugged. This is done with the help of `checkConstraints/3`. Its first argument is the label of the formula that is plugged, its second argument is the list of constraints, and its third argument the list of scoped formulas. In fact, it is just matter of working through the \leq -constraints and for each `leq(L,H)`, where the plugged label is identical with `L`, checking whether `H` is part of the scoped formulas.

```
checkConstraints(_, [], _).
checkConstraints(Label, [leq(L,H) | Constraints], Scoped):-
    Label==L, !,
    member(Formula, Scoped), Formula==H,
    checkConstraints(Label, Constraints, Scoped).
checkConstraints(Label, [_ | Constraints], Scoped):-
    checkConstraints(Label, Constraints, Scoped).
```

Note that we use the cut after identifying a relevant constraint. So if a constraint is violated, we force Prolog to backtrack and try a different plugging. And this is all there is to it.

Before putting the plugging algorithm into action, we will show how to build underspecified representations using our standard lexicon and grammar rules.

Building Underspecified Representations

Building USRs is done with the tools we're already familiar way: using lambdas and β -conversion. For USRs, we are going to introduce lambdas for holes and labels as well. Further, we need a *merge*-operation for USRs, that combines two USRs to one big USR.

The way we construct USRs is basically an extension to our normal systematic way of constructing representations. To keep track of the holes and labels, we will decorate *each* USR-representation with two additional lambdas: the 'top hole', and the 'main label'. The top hole is the hole that closes the current scope domain. The main label is the label within the current scope domain that is out-scoped by all elements within its scope domain. Abstracting away from these two concepts, it is straightforward to combine USRs (incidentally, macro definitions might get complicated—drawing a picture often eases matters).

Let's introduce the semantic macro for a determiner:

```
detSem(uni, lambda(P, lambda(Q, lambda(H, lambda(L,
    merge(merge(usr([F, R, S],
```

```
[F:forall(X,R>S)],
  [leq(F,H),leq(L,S)],P@X@H@R),Q@X@H@L)))).
```

The two additional lambdas for H and L close the scope domain of ‘every’. The quantifier itself gets label F with restriction `Restr` and nuclear scope `Scope`. Constraint `leq(F,H)` ensures that the quantifier is out-scoped by the top of the scope domain, and `leq(L,Scope)` says that its verbal argument is in its scope. The applications `P@X@H@R` and `Q@X@H@L` are crucial—these associate the restriction label of the quantifier with the label of the noun representation, and the top and hole label of the verb phrase argument.

Proper names or pronouns, on the other hand do not introduce holes, and therefore need not to bother about leq-constraints. A new top hole H and main label L are introduced and applied to its argument P.

```
pnSem(Sym,lambda(P,lambda(H,lambda(L,P@Sym@H@L))))).
```

Verbs introduce the ‘top’ hole and a leq-constraint between itself and the top. This is of importance for sentences where no quantifiers (or other scope-introducing elements) appear, as it establishes a clean link from the top element to the main label, that of the verb.

```
ivSem(Sym,lambda(X,lambda(H,lambda(L,usr([], [L:F], [leq(L,H)]))))):-
  compose(F,Sym,[X]).
```

Underspecification allows us to be very flexible. We can use one and the same technique for both quantifiers and negation. Here is the semantic macro for verb negation:

```
modSem(neg,lambda(P,lambda(X,lambda(H,lambda(L,
  merge(usr([N,S],[N:(~S)], [leq(N,H),leq(L,S)]),P@X@H@L)))))
```

There is one thing left we haven’t paid any attention to yet. It is combining smaller USRs to bigger ones, by carrying out the stipulated `merge/2` operations in the USR. Actually, this is not a big deal, and done as follows, using `append/3` from the library file.

```
mergeUSR(usr(D,L,C),usr(D,L,C)).

mergeUSR(merge(U1,U2),usr(D3,L3,C3)):-
  mergeUSR(U1,usr(D1,L1,C1)),
  mergeUSR(U2,usr(D2,L2,C2)),
  append(D1,D2,D3),
  append(L1,L2,L3),
  append(C1,C2,C3).
```

As usual we implement a driver that grasps all the important predicates together. This one combines the plugging algorithm and the construction of USRs, displays the USR on the current output device as well as all the readings it has.

```

parse:-
  readLine(Sentence),
  d(Sem,Sentence,[]),
  betaConvert(merge(usr([Top,Main],[],[]),Sem@Top@Main),Reduced),
  mergeUSR(Reducd,usr(D,L,C)),
  printRepresentation(usr(D,L,C)),
  findall(Top,plugHole(Top,L-[],C,[]),Readings),
  printReadings(Readings).

```

Here is an example session:

```

?- parse.

> Every woman does not snort.

usr([A,B,C,D,E,F,G],
     [C:forall(H,D>E),D:woman(H),F:(~G),B:snort(H)],
     [leq(C,A),leq(B,E),leq(F,A),leq(B,G),leq(B,A)])

Readings:
1 forall(A,woman(A)> ~snort(A))
2 ~forall(A,woman(A)>snort(A))

```

Evaluating Underspecification

So far we have formalized semantic underspecified representations and are able to deal with (quantifier) scope ambiguities. But how does this proposal relate to earlier accounts of scope ambiguities? Are there any benefits in using the new streamlined underspecified representation rather than stick to the old fashioned storage methods? The basic answer is that underspecification languages are more powerful: they exhibit far more descriptive potential than storage methods. To make this statement more clear, let's consider a concrete example sentence including three quantifiers:

(25) One criminal knows every owner of a hash bar.

With the tools developed in this chapter it is perfectly possible to represent the relative scopes of the three quantifiers fully underspecified. Moreover, it is possible to partially

specify relative scopes. For some (linguistic) reason, one might wish to express that ‘one criminal’ out-scopes ‘every owner’ and ‘a hash bar’, but that the relative scopes of the latter two quantifying constructs does not matter. We need not add new machinery to our formalism: we can use the \leq -constraint to state such requirements.

How do other approaches to quantifier scope fare with such requirements? Storage techniques (Cooper 1983; Keller 1988), to begin with, are not able to express such constraints. A store is a sequence of the core semantic representation and a list of the binding operators (the quantifiers) which scope is not fixed yet. In the process of constructing a store, quantified noun phrases can either enter the store or combine with the core semantic representation (and form a new core semantic representation, by means of in-situ quantification). At the end of the parse, the different ways in which the quantifiers are retrieved from the store and combined with the core semantic representation result in different readings (the last quantifier which is pulled off the store gets widest scope). Hence, to give a quantifier obligatory wide scope, the only way to pursue is to put it on the store, while performing in-situ application for the remaining quantifiers. But this presumes that the scope order of the remaining quantifiers is fixed!

Software Summary of Chapter 3

`mainMontague.pl` A simple implementation of Montague’s approach to quantification. (page 223)

`englishGrammarMontague.pl` The revised grammar rules for Montague’s quantifier raising. (page 224)

`mainCooperStorage.pl` Implementation of Cooper Storage. Defines storage of noun phrases and retrieval of the store. (page 226)

`mainKellerStorage.pl` Implementation of Nested Storage (Keller). Defines storage of noun phrases and retrieval of the store. (page 228)

`semMacrosStorage.pl` Contains the semantic macros for Cooper and Keller storage. (page 230)

`pluggingAlgorithm.pl` The plugging algorithm for hole semantics. (page 231)

`mainPLU.pl` Implementation of predicate logic unplugged. (page 233)

`semMacrosPLU.pl` Contains the semantic macros for predicate logic unplugged. (page 234)

`mergeUSR.pl` Merging underspecified semantic representations. (page 236)

Notes

More about quantifier scoping in general, and Montague’s approach to quantification in particular, can be found in Dowty, Wall, and Peters 1981 or volume 2 of Gamut 1991b. Cooper storage was first used in Cooper 1975; a more refined version is presented in Cooper 1983. Nested Cooper storage is due to Keller (1988). This very readable paper introduces Cooper storage, explains where the problems lie, gives an example (involving an interaction between scope and anaphoric pronouns) which suggests that a simple check for free variables during retrieval isn’t an adequate solution, and then introduces nested Cooper storage.

There are other approaches to quantifier scope ambiguity the reader should be aware of. Hobbs and Shieber 1987 provide an alternative to Cooper storage that overcomes some of the problems mentioned in this chapter. An improved version of their algorithm has been developed by Lewin (1990).

As we mentioned in the text, stores (and nested stores) are essentially a more abstract form of semantic representation—representations which encode multiple possibilities, without committing us to any particular choice. Viewed in this light, stores are ancestors of the current crop of *underspecification formalisms*, representation languages specifically designed to cope with ambiguity by avoiding overcommitment.

One of the first underspecified semantic representations were underspecified discourse representation structures (UDRSs) proposed by Reyle (1993). In fact, hole semantics can be seen as a generalization of this idea, not stuck to DRSs but compatible to all kinds of logical structures. Hole semantics is also applied in the machine translation system *Verbmobil* (Bos et al. 1996; Bos et al. 1998).

Probably the best way for the reader to find out more about the area of underspecification is to consult the recent collection edited by Van Deemter and Peters (Van Deemter and Peters 1996).

Chapter 4

Propositional Inference

In this chapter we turn to the second major theme of the book, namely:

How can we use the logical representations of natural language expressions to automate the process of drawing inferences?

Recall from Chapter 1 that a valid formula is a formula that cannot be falsified in any model, and that a valid argument (or valid inference) is an argument such that whenever all the premisses are true in some model, the conclusion is true in that model also. These concepts are the fundamental ones underlying the idea of logical inference, and so we need a way of getting to grips with them computationally. But there is a problem. Both concepts are semantic: that is, they are defined in terms of models. Moreover, they use models in a very strong way: their definitions refer to the class of *all* models. Now, the class of all models is a very large and abstract entity. Certainly we cannot put all models in the computer (there are infinitely many of them, most of which are infinite) and check whether an argument is valid by checking its premisses and conclusions in them all. So, if we want programs that perform logical inference, we will have to look for another approach.

Fortunately, alternatives are available. The branch of logic called *proof theory* studies *syntactically* defined inference methods. That is, proof theorists are interested in inference methods which only make use of the syntactic structure of sentences; models play no role. Of course, proof theoretic methods must always be *justifiable* in semantic terms. That is, if someone asks why some proposed syntactic proof method really is a way of performing genuine logical inference—not just some bizarre way of manipulating logical symbols—it must be possible to show that the method is faithful to the fundamental semantic concepts of validity. Nonetheless, *using* proof-theoretic methods requires no appeal to semantic concepts; only syntactic manipulation of formulas is required. Since we are interested in automating inference, it is clearly sensible to try and make use of such techniques, for syntax is pleasantly concrete—it's something that a computer can easily get to grips with.

Proof theorists have studied many different proof methods. Among the better known ones are axiomatic systems, natural deduction systems, resolution-based systems, sequent calculi, and the method we shall discuss here: *tableaux systems*. The various proof methods were developed for different purposes, and have different strengths and weaknesses. For example, axiomatic systems are excellent if one wants to study provability at an abstract level, but awful to actually use. Natural deduction systems, on the other hand, are nice and easy for people to use (as their name suggests, they try to capture something of what is involved when a human being goes about proving something), but, precisely because they demand human insight, are not the best choice for automated theorem proving. Resolution, on the other hand, has proved to be the key method for automated theorem proving (Prolog is based on it), but it is a ‘machine oriented’ rather than a ‘human oriented’ approach. We have chosen to work with tableaux systems because they offer the following blend of advantages:

1. Tableaux are a very intuitive proof method. The tableaux method is certainly syntactic, nonetheless it based on utterly transparent semantic intuitions—indeed, tableaux are often called *semantic tableaux*. They are certainly a ‘human oriented’ approach to inference.
2. Although intuitively natural, the process of finding a tableaux proof does not *depend* on human insight. Certainly human insight can streamline the process of performing tableaux inferences; nonetheless, the basic mechanism is suitable for machine implementation. Indeed, tableaux proof systems are relatively simple to implement, and, resolution aside, are one of the more popular automated theorem proving methods.
3. The tableaux method can be adapted to many different logics, such as higher-order and non-classical logics. The basic ideas are relatively robust, and the reader who masters them in the case of first-order logic has actually mastered ideas of far wider applicability.
4. Tableaux systems are more than just theorem provers: they can also be regarded as *model building tools*. (This is another reason why they are sometimes called semantic tableaux.) For some applications this additional ability is useful.

In this chapter we are going to develop the tableaux method for propositional languages. Recall from Chapter 1 that propositional languages are essentially a user-friendly notation for the quantifier-free fragment of first order languages. It will turn out that this is a very important fragment to examine, for in the next chapter we shall see how to efficiently reduce the tableaux method for a full first-order language to the tableaux method for its propositional fragment.

4.1 Propositional Tableaux

The key intuition underlying the tableaux proof technique revolves around the following *semantic* question:

Suppose we are given a formula, and one of the truth values TRUE or FALSE. Is it possible to find a model in which the given formula has the given truth value?

The tableaux method is essentially a *syntactic* way of answering this question. More precisely, a tableaux proof is a *systematic* check, making use of only syntactic concepts, which lets us know whether or not it is possible to build a model in which some given formula is true or false.

Suppose that we had such a systematic check at our disposal. (We soon will have.) Why would it be useful? First, note that it would give us a way to test formulas for validity: ‘valid’ means ‘true in all models’, so a formula is valid if and only if it is not possible to falsify it in any model. Hence, a formula ϕ would be valid iff the systematic method told us that there was no way to build a model that falsified ϕ .

Second, note that such a method would also give us a way of performing logical inference. For suppose we are given ϕ_1, \dots, ϕ_n are premisses, and we wish to judge whether the argument to the conclusion ψ is valid or not. That is, we wish to know whether or not

$$\phi_1, \dots, \phi_n \models \psi.$$

Now, by the Deduction Theorem (discussed in Chapter 1), this argument is valid precisely when

$$\models \phi_1 \wedge \dots \wedge \phi_n \rightarrow \psi,$$

that is, when $\phi_1 \wedge \dots \wedge \phi_n \rightarrow \psi$ is a valid formula. In short, the problem of deciding whether arguments are valid is reducible to the problem of deciding whether formulas are valid. So, given a method of determining whether a formula is valid, we can use it for judging the validity of arguments as well.

Thus such a systematic check for the existence of satisfying or falsifying models would certainly be worth having, and this is exactly what tableaux systems give us. But what is a tableaux system, and how do they systematically check for the existence of suitable models? In the remainder of this section we informally introduce the key ideas involved. We do this by presenting three examples of tableaux proofs, taking care to point out the underlying semantic intuitions. In the following section we will make our discussion more precise.

Consider the formula $p \vee \neg p$. This is a validity—we certainly can’t falsify it—but what would a systematic search for a falsification look like? Now, we’ve already seen one answer

to this question—simply fill out its truth table! But this answer isn't very appealing. For a start, it's not an answer we can generalise to first-order languages. Moreover, while a truth table is easy to fill out for $p \vee \neg p$, the truth table for a formula containing 8 different propositional variables contains 256 lines, while the table for a formula containing 20 variables contains 2^{20} lines, which is too large for comfort. So we won't bother with truth tables—instead, we'll develop a number of *tableaux expansion rules*. These rules tell us how to make complex formulas true (or false) by breaking them down into their component formulas and giving the components the appropriate truth value. Let's see what sort of rules are needed, and how to use them, by constructing a tableaux proof of $p \vee \neg p$.

We want to try and falsify $p \vee \neg p$, so let's introduce a piece of notation to express this goal. Writing $F\phi$, where ϕ is any formula, will mean that we want to make ϕ false. Similarly, writing $T\phi$ will mean that we want to make ϕ true. (T and F are called *signs*, and a formula preceded by a sign is called a *signed formula*.) Thus, as we are going to try to falsify $p \vee \neg p$, our initial goal is:

$$F(p \vee \neg p).$$

This rather trivial looking object is our first example of a tableau. Actually, when writing out tableaux by hand, it is handy to include a little extra book-keeping information, such as line numbers. So, in practice we'd tend to write the above tableau as:

$$1 \quad F(p \vee \neg p)$$

How do we proceed? Essentially we use the tableaux expansion rules to smash the signed formula into smaller and smaller pieces until we reach the atomic level. So, what expansion rule should we apply here? Obviously we need a rule that tells us how to *falsify a disjunction*. The required rule is clear: to make a disjunction false, falsify both disjuncts. So, applying this rule—let's call it F_{\vee} —we expand our one line tableau to a three line tableau as follows:

$$\begin{array}{lll} 1 & F(p \vee \neg p) & \checkmark \\ 2 & Fp & 1, F_{\vee} \\ 3 & F\neg p & 1, F_{\vee} \end{array}$$

Here, lines 2 and 3 are the extra information we have deduced by applying the F_{\vee} rule. The third column contains more book-keeping information: the \checkmark symbol in line 1 records the fact that we've applied the appropriate rule to line 1, while the annotations '1, F_{\vee} ' in lines 2 and 3 record the fact that these lines were obtained from line 1 using rule F_{\vee} .

What next? In fact, there's only one more thing we can do. We've already applied a rule to line 1, so we've finished with that. Furthermore, line 2 tells us something about *atomic* information (namely that we need to make p false). This is simply a blunt fact, not something that can be further analysed. Thus only line 3 offers us the possibility of further

progress; it tells us that we need to falsify the negation of p . So, we need an expansion rule that tells us how to do this. Again, the required rule is clear: to make the negation of a formula false, make the formula itself true. So, applying this rule—call it F_{\neg} —we can expand our three line tableau to a four line tableau as follows:

1	$F(p \vee \neg p)$	\checkmark
2	Fp	$1, F_{\vee}$
3	$F\neg p$	$1, F_{\vee}, \checkmark$
4	Tp	$3, F_{\neg}$.

Note that we have marked line 3 with a \checkmark (thus recording that the applicable rule has been applied) and have indicated that line 4 was obtained from line 3 using F_{\neg} .

There are two important observations that must be made about this tableau. First, it is *rule saturated*. That is, we cannot expand it any more. Either we've already applied the applicable rule (lines 1 and 3), or the line contains instructions about what to do with atomic information (lines 2 and 4). Second, the tableau is *closed*. That is, it contains contradictory instructions. The tableau tells us that if we want to falsify $p \vee \neg p$, we have to make p false (line 2) and p true (line 4). As this is impossible, and as it should be pretty clear that the above tableau really does indicate all possible ways falsifying $p \vee \neg p$, we conclude that this formula is valid. This closed tableau is called a *tableaux proof* of $p \vee \neg p$.

The previous example is a perfectly good tableaux proof, but it's also just about the simplest one the reader is ever likely to see. So let's consider a slightly more demanding task: testing $\neg(q \wedge r) \rightarrow (\neg q \vee \neg r)$ for validity.

Just as in the previous example, our initial goal is to try and falsify the given formula. Thus our initial tableau is:

1	$F\neg(q \wedge r) \rightarrow (\neg q \vee \neg r)$
---	--

We need to falsify an implication, so we need an expansion rule that tells us how to do this. Again, the required rule is clear: to falsify an implication, make the antecedent true and the consequent false. Applying this rule—let's call it F_{\rightarrow} —yields the following tableau:

1	$F\neg(q \wedge r) \rightarrow (\neg q \vee \neg r)$	\checkmark
2	$T\neg(q \wedge r)$	$1, F_{\rightarrow}$
3	$F(\neg q \vee \neg r)$	$1, F_{\rightarrow}$

Now, line 3 demands that we falsify a disjunction. We've already met the required expansion rule, namely F_{\vee} . Applying it to line 3 yields:

1	$F\neg(q \wedge r) \rightarrow (\neg q \vee \neg r)$	\checkmark
2	$T\neg(q \wedge r)$	1, F_{\rightarrow}
3	$F(\neg q \vee \neg r)$	1, F_{\rightarrow} , \checkmark
4	$F\neg q$	3, F_{\vee}
5	$F\neg r$	3, F_{\vee}

Now at this point the alert reader should be saying “Hold on a minute! Why are we free to apply this rule to line 3? Sure, the rule fits—but look at line 2. There’s a formula there that we need to make true. Don’t we need to take care of that first?”

A sensible question, but the answer is: *no, we don’t*. One of the pleasant things about tableaux is that we are free to apply applicable rules in any order we like. Yes, we could have applied the relevant rule to line 2 at this point, had we wanted to—but we’re equally free to apply a rule to line 3, just as we did above. Tableaux rules just tell us what we’re *permitted* to do to expand a tableau; we’re not forced to apply expansion rules in any particular order.

So let’s move on. Note that both lines 4 and 5 ask us to falsify a negated formula. Again, we have already met the relevant rule, namely F_{\neg} . Applying it first to line 4, and then to line 5 (a completely arbitrary ordering choice) yields:

1	$F\neg(q \wedge r) \rightarrow (\neg q \vee \neg r)$	\checkmark
2	$T\neg(q \wedge r)$	1, F_{\rightarrow}
3	$F(\neg q \vee \neg r)$	1, F_{\rightarrow} , \checkmark
4	$F\neg q$	3, F_{\vee} , \checkmark
5	$F\neg r$	3, F_{\vee} , \checkmark
6	Tq	4, F_{\neg}
7	Tr	5, F_{\neg}

Let’s now deal with line 2. (In fact, there is nothing else we can do.) We need to make a negation true. The required rule is clear: to make the negation of a formula true, make the formula itself false. So, applying this rule—call it T_{\neg} —we can expand our tableau to obtain:

1	$F\neg(q \wedge r) \rightarrow (\neg q \vee \neg r)$	\checkmark
2	$T\neg(q \wedge r)$	1, F_{\rightarrow} , \checkmark
3	$F(\neg q \vee \neg r)$	1, F_{\rightarrow} , \checkmark
4	$F\neg q$	3, F_{\vee} , \checkmark
5	$F\neg r$	3, F_{\vee} , \checkmark
6	Tq	4, F_{\neg}
7	Tr	5, F_{\neg}
8	$F(q \wedge r)$	2, T_{\neg}

Now we have to deal with line 8, which asks us to falsify a conjunction. Doing so leads us to the first real complication in the story we have been telling. The point is this: there's not just one way of making a conjunction false, there's two. Making either conjunct false will falsify the whole formula. As tableaux are meant to be *systematic* searches for falsifications, we're going to have to examine both possibilities. Thus the relevant expansion rule—call it F_{\wedge} —is going to yield two alternative ways of expanding the tableau, and we're going to have to keep track of both.

F_{\wedge} is our first example of a *disjunctive* (or *branching*) expansion rule; we'll encounter more such rules in the following section. Because of such rules, tableaux *won't* generally consist of a single straight line down the page (that is, they're no longer going to consist of a single *branch*, to use the official terminology). Rather, they will be tree-like structures, possibly containing many branches. In the present case, what we get after applying F_{\wedge} is:

1	$F\neg(q \wedge r) \rightarrow (\neg q \wedge \neg r)$	\checkmark
2	$T\neg(q \wedge r)$	$1, F_{\rightarrow}, \checkmark$
3	$F(\neg q \wedge \neg r)$	$1, F_{\rightarrow}, \checkmark$
4	$F\neg q$	$3, F_{\vee}, \checkmark$
5	$F\neg r$	$3, F_{\vee}, \checkmark$
6	Tq	$4, F_{\neg}$
7	Tr	$5, F_{\neg}$
8	$F(q \wedge r)$	$2, T_{\neg}, \checkmark$
9	Fq	$8, F_{\wedge}$
10	Fr	$8, F_{\wedge}$

That is, we have recorded two distinct possibilities: we must either falsify q , or falsify r .

But let's return to our tableaux proof. What do we do next? Actually, we've finished: our two-branch tableau is rule-saturated, as the reader can easily check. So, after all that, is $\neg(q \wedge r) \rightarrow (\neg q \wedge \neg r)$ a validity or not? Yes, it is. This is because the rule-saturated tableau we produced is *closed*, which means that *all the branches it contains are closed*. To see this, note that the left-hand branch is closed because it contains the contradictory instructions Fq (at line 9) and Tq (at line 6); whereas the right-hand branch is closed because it contains Fr (at line 9) and Tr (at line 7). This closed tableau is a tableaux proof of $\neg(q \wedge r) \rightarrow (\neg q \wedge \neg r)$.

Let's consider a final example to illustrate what happens if the formula we are working with is *not* a validity. Now, the formula $(p \wedge q) \rightarrow (r \vee s)$ is certainly not valid; what happens when we try and falsify it using the tableaux method?

As usual, the first step simply states our goal, namely :

$$1 \quad F(p \wedge q) \rightarrow (r \vee s)$$

As we need to falsify an implication, the relevant expansion rule is F_{\rightarrow} :

1	$F(p \wedge q) \rightarrow (r \vee s)$	\checkmark
2	$T(p \wedge q)$	$1, F_{\rightarrow}$
3	$F(r \vee q)$	$1, F_{\rightarrow}$

Now line 2 requires us to make a conjunction true. The expansion rule T_{\wedge} required is clear: we must make both conjuncts true. Applying T_{\wedge} yields:

1	$F(p \wedge q) \rightarrow (r \vee s)$	\checkmark
2	$T(p \wedge q)$	$1, F_{\rightarrow}, \checkmark$
3	$F(r \vee q)$	$1, F_{\rightarrow}$
4	Tp	$2, T_{\wedge}$
5	Tq	$2, T_{\wedge}$

Now let's deal with line 3. The relevant rule is F_{\vee} . Applying it yields:

1	$F(p \wedge q) \rightarrow (r \vee s)$	\checkmark
2	$T(p \wedge q)$	$1, F_{\rightarrow}, \checkmark$
3	$F(r \vee q)$	$1, F_{\rightarrow}, \checkmark$
4	Tp	$2, T_{\wedge}$
5	Tq	$2, T_{\wedge}$
6	Fr	$3, F_{\vee}$
7	Fs	$3, F_{\vee}$

But now the expansion process halts. This tableau is rule-saturated: we've applied the relevant rules to line 1–3, while lines 4–7 simply stipulate what has to be done with atomic information. But note that there is crucial difference between this tableau and the previous rule-saturated tableaux we have seen: the single branch in this tableau is not closed, it's *open*. That is, it does *not* contain contradictory instructions. In fact, it gives us very sensible instructions indeed: lines 4–7 tell us to make p true, q true, r false, and s false. As the reader can easily check, doing this falsifies $(p \wedge q) \rightarrow (r \vee s)$, thus this formula is *not* a validity.

More generally, a rule-saturated tableau is called *open* iff it contains at least one open branch. If we obtain an open tableau when we try to falsify some formula, this means that the formula is *not* a validity. Moreover, just as in the previous example, every open branch on the open tableau actually contains an explicit 'falsification recipe' for the formula: we falsify it by assigning truth values to the atomic symbols in the the way the open branch stipulates. To put it another way, open branches of rule-saturated tableaux tell us to build a model that falsifies the formula we started with.

The reader should now have a fairly clear grasp of the main ideas and intuitions underlying tableaux proofs, so let's now discuss the method more systematically. We begin by listing

the 8 main expansion rules, and classifying them into three types: *unary rules*, *conjunctive rules*, and *disjunctive rules*. This classification isn't vital, but it will help us keep the implementation neat.

For each Boolean connective B , we have two tableaux rules, T_B and F_B . T_B tells us how to make a formula with B as its main connective true, while F_B tells us how to make it false. The rules for negation, T_{\neg} and F_{\neg} , are the simplest:

$$\frac{T\neg\phi}{F\phi} \qquad \frac{F\neg\phi}{T\phi}$$

These rules should be read from top to bottom. Given the examples discussed in the previous section, their meaning should be reasonably clear. In each rule, the signed formula above the horizontal line is the *input* to the rule, and the signed formula below it is the *output*. For example, T_{\neg} takes as input signed formulas of the form $T\neg\phi$, and returns as output signed formulas of the form $F\phi$. Similarly, F_{\neg} takes as input signed formulas of the form $F\neg\phi$ and returns as output signed formulas of the form $T\phi$. In what follows, we shall call these two rules our *unary rules*. This is simply shorthand for the fact that both rules return a single formula as output.

The rules for the binary connectives \wedge , \vee and \rightarrow are rather more interesting. Here are T_{\wedge} and F_{\wedge} , F_{\vee} and T_{\vee} , and F_{\rightarrow} and T_{\rightarrow} , respectively:

$$\frac{T(\phi \wedge \psi)}{T\phi \quad T\psi} \qquad \frac{F(\phi \wedge \psi)}{F\phi \mid F\psi}$$

$$\frac{F(\phi \vee \psi)}{F\phi \quad F\psi} \qquad \frac{T(\phi \vee \psi)}{T\phi \mid T\psi}$$

$$\frac{F(\phi \rightarrow \psi)}{T\phi \quad F\psi} \qquad \frac{T(\phi \rightarrow \psi)}{F\phi \mid T\psi}$$

Again, all six rule should be read from top to bottom, the top being the input to the rule, the bottom the output. We shall call the three rules in the left-hand column (that is, T_{\wedge} , F_{\vee} , and F_{\rightarrow}) *conjunctive rules*. The three rules on the right-hand side (that is, F_{\wedge} , T_{\vee} , and T_{\rightarrow}) are called *disjunctive rules*. Note that each of \wedge , \vee and \rightarrow gives rise to a pair of rules, one of which is conjunctive, the other disjunctive. Both conjunctive and disjunctive rules return two formulas as output—however, as we have already seen, they differ in the effect they have on tableaux. In particular, use of a disjunctive rule splits the branches of the tableaux containing the input formula into distinct branches, each of which records one of the two alternative output formula.

In the discussion that follows, we'll assume that we only have these 8 rules at our disposal. However, the reader should be aware that it is very straightforward to give expansion rules for other connectives, though such rules won't always fit into our three way classification. For example, it can be useful to have the expansion rules for \leftrightarrow at our disposal. Here are the required rules:

$$\frac{T(\phi \leftrightarrow \psi)}{T\phi \mid F\phi} \quad \frac{F(\phi \leftrightarrow \psi)}{T\phi \mid F\phi}$$

$$\frac{}{T\psi \mid F\psi} \quad \frac{}{F\psi \mid T\psi}$$

Note that for this connective we *don't* obtain a neat pair of rules, one of which is conjunctive, the other disjunctive, as we did for \wedge , \vee and \rightarrow . Moreover, though both rules are 'disjunctive' in the sense that both force us to split tableau branches, these rules yield 4 output formulas, 2 for each possible branch. The exercises below ask the reader to develop expansion rules for some other connectives.

Now that we know what our expansion rules are, let's make our previous discussion of tableaux and tableau proofs a little more rigorous.

A (propositional) *tableau* is simply a tree, each of whose nodes is a signed (propositional) formula. A *branch* of a tableau is simply a branch of such a tree—that is, a collection of nodes (that is, signed formula) that contains exactly one leaf node together with all the nodes which dominate it.


An *initial tableau*—that is, the kind of tableau with which we start the tableau expansion process—is a tableau that has only a single branch. The reader may be slightly surprised by this definition. Isn't it too general? Why don't we simply define an initial tableau to be a tableau consisting of a single signed formula of the form $F\phi$? Actually we *could* use this more specific definition—but, as we shall later see, the slightly more general definition given is rather useful.

We carry out the tableau expansion process as follows. Given a tableau, try to find a node in it that (a) isn't a signed atomic formula, and (b) to which we haven't already applied an expansion rule. Let's call such nodes *unexpanded nodes*. If there are no unexpanded nodes, we can't do anything: we have a *rule-saturated* tableau and are finished. So suppose there is at least one unexpanded node. This node has the form $S\phi$, where S is either T or F , and ϕ is a propositional formula. Moreover, as ϕ is not atomic, it has a main connective B , where B is one of the connectives \neg , \wedge , \vee , or \rightarrow . We can then 'read off' the required rule: we have to apply rule S_B .


What happens when we apply a rule to a node? That depends on whether the rule is unary, disjunctive, or conjunctive. If the rule is unary, we extend every branch containing the input node by adding on the output formula of the rule as the new leaf node. If the rule is conjunctive, we extend every branch containing the input node by adding on, again at the leaf node, both output formula of the rule. Finally, if the rule is disjunctive, we extend every branch containing the input node to two distinct branches, one containing

each of two choices of output formulas. Again, both the required additions are carried out at the leaf node of the original branches.

The basic idea of tableau expansion should be clear from the examples we discussed earlier, but there is one point worth emphasizing. A signed formula may belong to several branches. (For example, the root node belongs to every branch of a tableau.) When we perform an expansion, we have to extend *every* branch on which the input formula lies in the appropriate way. For example, consider the following tree:



If we expand the indicated node $T(\phi \vee \psi)$, we obtain:



The tableau expansion process starts when we are given an initial tableau. We apply rules to the initial tableau, and then to the tableaux obtained by earlier rule applications, and so on, until it is not possible to apply any more rules. As we discussed earlier, we can apply the rules in any order we like. The tableau expansion process stops when it is not possible to apply any more rules, that is, when we obtain a rule-saturated tableau. Here an important question looms: is the expansion process always guaranteed to stop? Yes, it is. We will point out why later in the chapter—but it would be a good idea for the reader to try figuring this out now unaided.

We are now almost ready to say what a *tableau proof* is. First, a branch of a tableau is *closed* if it contains both $T\phi$ and $F\phi$, where ϕ is some formula. A branch that is not closed is called *open*. A tableau is closed if *every* branch it contains is closed, and open if it contains *at least one* open branch. Now for the key definition:

A formula ϕ is tableaux-provable (or, more simply: provable) if and only if it is possible to expand the initial tableau consisting of a single node $F\phi$ to a closed tableau. We use the notation $\vdash \phi$ to indicate that ϕ is provable.

And that's the (propositional) tableaux method. To conclude our discussion, two remarks. First, note that the tableaux method really is a purely *syntactic* method. This should be clear. If we want to test whether ϕ is provable we have to start with the initial tableau $F\phi$, and then try and expand it to a closed tableau. The entire expansion process is governed by syntactic ideas. For example, when we expand a node, by simply looking at its sign and main connective we know which rule to apply. Moreover, closure is a purely syntactic concept: it simply amounts to looking for items of the form $T\phi$ and $F\phi$ on some branch. Of course, as our presentation has tried to emphasize, the tableaux method is driven by clear *semantic* ideas. Nonetheless, *using* the method doesn't require any semantic insight at all. Even if we didn't know what the signs T and F were meant to stand for—or indeed, what \neg , \wedge , \vee , and \rightarrow were meant to represent—we would still have a well defined way of manipulating logical formulas. It would probably be an exaggeration to claim that we could train a monkey to carry out tableaux proofs, but as we shall see below, we can certainly implement them in Prolog rather easily.

The second point concerns our definition of initial tableaux. Why did we opt for the general definition, rather simply stating that initial tableau always have the form $F\phi$? Simply because, as we shall now see, by using different initial tableau we can use tableaux for different purposes.

First, we can use tableaux to *directly* test whether arguments are valid. For example, suppose we wanted to test

$$\phi_1, \dots, \phi_n \models \psi,$$

for validity. Now, as we pointed out earlier, by the Deduction Theorem this is equivalent to testing whether $\phi_1 \wedge \dots \wedge \phi_n \rightarrow \psi$ is a valid formula, thus we could simply try using the tableaux method to falsify this formula. But there is a more direct way. An argument is valid if and only if whenever all the premisses are true, the conclusion is true also. Hence, to test whether the above argument is valid, we could form the initial tableau:

$$\begin{array}{c} T\phi_1 \\ \cdot \\ \cdot \\ \cdot \\ T\phi_n \\ F\psi \end{array}$$

That is, we simply use the tableau method to see whether it is possible to make all the premisses true and the conclusion false. If we *can't* do this—that is, if we obtain a closed tableau—the argument is valid.

But we can do more with tableaux than simply testing whether formulas and arguments are valid: we can use them as a tool for telling us how to build models. To give a very trivial example, suppose we wanted to know exactly how to make the formula $\neg p \wedge (r \vee q)$

true. The tableau method will tell us how to do this. As our initial tableau we take:

$$T\neg p \wedge (r \vee q).$$

That is, we are simply asking: “How do we make this formula true?” To get an answer, all we have to do is carry out the expansion process in the usual way:

$$\begin{array}{rcc}
 1 & T\neg p \wedge (r \vee q) & \checkmark \\
 2 & T\neg p & 1, T\wedge, \checkmark \\
 3 & T(r \vee q) & 1, T\wedge, \checkmark \\
 4 & Fp & 2, T\neg, \checkmark \\
 \hline
 5 & Tr \quad 3, T\vee & \\
 & & 6 & Tq \quad 3, T\vee
 \end{array}$$

This tableaux is rule-saturated and contains two open branches. The left hand branch tells us that one way to make the input formula true is to make p false and r true, while the right hand branch tells us that there is also another way to achieve this, namely by making p false and q true.

Exercise 4.1.1 Show that the following formula's are all provable:

1. $\neg\neg p \rightarrow p$
2. $((p \rightarrow q) \rightarrow p) \rightarrow p$
3. $(\neg p \rightarrow \neg q) \rightarrow (q \rightarrow p)$
4. $p \leftrightarrow (p \wedge (q \vee p))$
5. $(p \vee (q \wedge r)) \leftrightarrow ((p \vee q) \wedge (p \vee r))$

Exercise 4.1.2 The connectives *nand* and *nor* are defined by the following truth tables:

p	q	$p \text{ nand } q$	$p \text{ nor } q$
TRUE	TRUE	FALSE	FALSE
TRUE	FALSE	TRUE	FALSE
FALSE	TRUE	TRUE	FALSE
FALSE	FALSE	TRUE	TRUE

Give the tableaux expansion rules for both connectives.

Exercise 4.1.3 There are two ways of using tableaux to show that an argument is valid: the indirect method that appeals to the Deduction Theorem, and the direct method just discussed. Explain why both methods give rise to the same result.

4.2 Implementing Propositional Tableaux

We shall represent tableaux in Prolog as lists containing lists of signed formulas. Signed formulas are represented as normal Prolog terms with functors \mathbf{f} and \mathbf{t} of arity 1. For example, the following is the Prolog representation of a tableau:

$$[[\mathbf{f}(p \rightarrow q)]].$$

This represents the tableau with just one branch, namely a branch containing the single signed formula $F(p \rightarrow q)$.

Here's a second example:

$$[[\mathbf{t}(p \ \& \ q), \ \mathbf{f}(p \ \& \ r)]].$$

Again, this represents a tableau with just one branch. This time, however, the branch contains *two* signed formulas, namely $T(p \wedge q)$ and $F(p \wedge q)$.

The Prolog implementation of propositional tableaux consists of a small collection of predicates which manipulate such lists in accordance with the tableaux expansion rules. For example, the list just given contains $\mathbf{t}(p \ \& \ q)$. Now, to make a conjunction true, we use expansion rule T_{\wedge} , which in this case tells us to add Tp and Tq to the given branch. It is easy to mimic the required expansion in Prolog: all we have to do is define a predicate which when given the previous list as input, returns the following one as output:

$$[[\mathbf{t}(p), \ \mathbf{t}(q), \ \mathbf{f}(p \ \& \ r)]].$$

Note two things. First, as expected, the output list contains $\mathbf{t}(p)$ and $\mathbf{t}(q)$. Second, it does *not* contain $\mathbf{t}(p \ \& \ q)$. Why is this? Because once the relevant expansion rule has been applied to some formula, that formula is no longer relevant and can be removed. That is, removal of signed formulas from lists is the Prolog analog of the \surd marking we used in our handwritten proofs.

Note that our new list contains $\mathbf{f}(p \ \& \ r)$, thus we can process it even further. Now, the rule for falsifying a conjunction is F_{\wedge} , a *disjunctive* rule. How are disjunctive expansions to be handled in Prolog? Again, the basic idea is very simple. To handle F_{\wedge} , for example, we need simply define a predicate which when given the previous list as input, return the following one as output:

$$[[\mathbf{t}(p), \ \mathbf{t}(q), \ \mathbf{f}(p)], \ [\mathbf{t}(p), \ \mathbf{t}(q), \ \mathbf{f}(r)]]$$

Note that this list contains *two* lists. It is thus the Prolog representation of a tableau containing *two* branches, namely the tableau with the signed formulas Tp , Tq and Fp on one (closed) branch, and the formulas Tp , Tq and Fr on the other.

Note that this approach to disjunctive rules differs from our handwritten approach. Our Prolog code make two copies of the branch, one for each of the two disjunctive possibilities. This is somewhat heavy handed—but it makes the code easy to understand. On the other hand, when writing out proofs by hand we want to cut down the writing needed as much as possible—hence our use of tree-like structures in which the information common to two branches is shared. Of course, by taking trees as the fundamental data structure, it is possible to implement this kind of structure sharing in Prolog, and doing so is an extremely valuable exercise.

Summing up, our Prolog implementation is based round the representation of tableaux as lists containing lists of signed formulas. Each list of signed formula corresponds to a branch. We carry out expansions by manipulating these lists of signed formulas in the appropriate way. In particular, conjunctive (and unary) expansions will be carried out by taking the list representation of a branch, removing the relevant conjunctive (or unary) signed formula, and adding its component(s) to the list. Disjunctive expansions are carried out in much the same way, save that the process returns *two* lists, one containing each of the two possible components. Incidentally, the program given below does not handle the \leftrightarrow connective; a good way of testing your understanding of the code is modify it so that it does.

Let's go systematically through the program. The outermost predicate is called `saturate`. This takes a single argument, a tableau, which it recursively attempts to rule-saturate with the help of the `expand` predicate. Note that base clause of `saturate` calls the `closed` predicate, which tests whether the input tableau is closed. By making this test here, we ensure that once a closed tableau is found, no further attempt is made to apply expansion rules.

```
saturate(Tableau):-
    closed(Tableau).

saturate(OldTableau):-
    expand(OldTableau,NewTableau),
    saturate(NewTableau).
```

The `closed` predicate is recursively defined in the expected way. To ground the recursion, the empty list is defined to be closed. A tableau containing branches can then be tested for closure by recursively peeling off its branches and checking for occurrences of both $t(X)$ and $f(X)$ for some formula X .

```
closed([]).

closed([Branch|Rest]):-
    member(t(X),Branch),
```

```

member(f(X), Branch),
closed(Rest).

```

The `expand/2` predicate is a high level ‘organisational’ predicate: it simply works its way recursively through all the branches in the input tableau, looking for branches to which one of the predicates `unaryExpansion`, `conjunctiveExpansion`, or `disjunctiveExpansion` apply. The fourth, recursive clause of `expand/2` takes the next branch of the tableau if no more expansions can be carried out to the current branch.

```

expand([Branch|Tableau], [NewBranch|Tableau]):-
    unaryExpansion(Branch, NewBranch).

expand([Branch|Tableau], [NewBranch|Tableau]):-
    conjunctiveExpansion(Branch, NewBranch).

expand([Branch|Tableau], [NewBranch1, NewBranch2|Tableau]):-
    disjunctiveExpansion(Branch, NewBranch1, NewBranch2).

expand([Branch|Rest], [Branch|Newrest]):-
    expand(Rest, Newrest).

```

The expansion predicates do the real work. Both the predicates `unaryExpansion` and `conjunctiveExpansion` have two arguments, namely an ‘input branch’, and an ‘output branch’, while `disjunctiveExpansion` takes three arguments, an input branch and *two* output branches. All three predicates look for the occurrence of the appropriate type of signed formula (namely `unary`, `conjunctive`, and `disjunctive` respectively) as the first item in the input branch, use the library predicate `removeFirst/3` to remove that occurrence, and then build the required new branch (or branches, in the case of disjunctive formulas) out of the component (or components) of the signed formula.

```

unaryExpansion(Branch, [Component|Temp]):-
    unary(SignedFormula, Component),
    removeFirst(SignedFormula, Branch, Temp).

conjunctiveExpansion(Branch, [Comp1, Comp2|Temp]):-
    conjunctive(SignedFormula, Comp1, Comp2),
    removeFirst(SignedFormula, Branch, Temp).

disjunctiveExpansion(Branch, [Comp1|Temp], [Comp2|Temp]):-
    disjunctive(SignedFormula, Comp1, Comp2),
    removeFirst(SignedFormula, Branch, Temp).

```

Thus, it only remains to define what **conjunctive**, **disjunctive**, and **unary** signed formula are:

```
conjunctive(t(X & Y),t(X),t(Y)).
conjunctive(f(X v Y),f(X),f(Y)).
conjunctive(f(X > Y),t(X),f(Y)).
```

```
disjunctive(f(X & Y),f(X),f(Y)).
disjunctive(t(X v Y),t(X),t(Y)).
disjunctive(t(X > Y),f(X),t(Y)).
```

```
unary(t(~X),f(X)).
unary(f(~X),t(X)).
```

And that's that. Of course, using **saturate** directly is a little clumsy, and it is preferable to call it in a slightly more user-friendly way. For example, we can define a predicate **valid/1**:

```
valid(F):-
    saturate([[f(F]])).
```

The reader should experiment with this program, and attempt at least the first two of the following exercises.

Exercise 4.2.1 Modify the code so that it handles the *nand* and *nor* connectives defined above.

Exercise 4.2.2 Modify the code so that it handles the \leftrightarrow connective.

Exercise 4.2.3 [hard] Reimplement propositional tableaux in a way that is more faithful to our hand-written tree based approach. That is, find a nice way of representing trees so that we don't need to duplicate branches when using disjunctive rules.

Exercise 4.2.4 Experiment with different strategies in the tableaux program, by changing the order of the expansion rules.

Exercise 4.2.5 [project] Add a pretty print predicate to our implementation of propositional tableaux, that shows the branches and the proof steps in a readable way. Test your printer on: $(a \rightarrow \neg(b \wedge c)) \rightarrow ((a \rightarrow \neg b) \wedge (a \rightarrow \neg c))$.

4.3 Equality Constraints

4.4 Theoretical Remarks

To conclude our discussion, we shall discuss four concepts that every reader should have at least a nodding acquaintance with: *soundness*, *completeness*, *termination* and *efficiency*. We won't pursue these topics in depth, nor will we prove any technical results. Rather, our aim is simply to give a clear account of what the concepts are, and how they relate to tableaux systems; for further details, the reader is advised to follow up the references cited in the Notes. Along the way, we shall clarify some points that may be bothering readers.

First of all, the tableaux system presented in this chapter is a *sound and complete* proof system for propositional inference. What does this mean?

As we remarked at the start of the chapter, although a syntactic proof method only *uses* syntactic information, it has to be *justifiable* in semantic terms. That is, when someone proposes some new proof method, we need guarantees that the proposal is semantically sensible.

The most fundamental property we demand of a proof system is that it be *sound*. Soundness is essentially a 'no garbage' demand: if the proof system proves some formula, then that formula must be valid. Sound proof systems are thus semantically justifiable in a very strong sense: they will *never* prove formulas which are falsifiable.

Our tableaux system is sound. That is, for any propositional formula ϕ ,

$$\text{if } \vdash \phi \text{ then } \models \phi.$$

So tableaux proofs will never lead us astray—but this shouldn't really come as a surprise. After all, we developed our tableaux expansion rules by thinking in overtly semantic terms. For example, we asked "How do we go about making a conjunction true?" and gave the (obviously sensible) answer "By making both conjuncts true". Indeed, *all* our expansion rules reflect the semantic definitions of the connectives in an obvious way, so it is hardly surprising that they are 'semantically safe'. And in fact, one proves the soundness of the tableaux method simply by thinking a little more systematically about the consequences of the 'semantic safety' of the expansion rules. We won't spell the details out here, but they are not particularly difficult and we suggest that mathematically inclined readers try to pin down the required argument precisely.

What about completeness? This is a much stronger (and more interesting) demand: a complete proof system is one which is capable of proving *all* valid formula. That is, if ϕ is any valid formula whatsoever, then it must be possible to give a proof of ϕ .

Our tableaux system is complete. That is, for any propositional formula ϕ ,

$$\text{If } \models \phi \text{ then } \vdash \phi.$$

Intuitively, no validity lies beyond the reach of the tableaux proof method: if a formula ϕ is valid, then it is possible to expand the initial tableau $F\phi$ to a closed tableau.¹ Note that because the tableaux method is both sound and complete we have that:

$$\vdash \phi \quad \text{iff} \quad \models \phi.$$

That is, we have a perfect match between the syntactic concept of tableaux-provability and the semantic concept of validity.

Proving that the tableaux method is complete is a non-trivial exercise which lies beyond the scope of this book; readers interested in the details should consult one of the references cited in the Notes. However, it *is* worth knowing that the key step of the completeness proof is to prove the following:

If it is possible to expand an initial tableau $F\phi$ to an open rule-saturated tableau, then one can falsify ϕ by giving the atomic symbols the truth values stipulated on any of the open branches of such a tableau.

This is worth knowing, because it allows us to dispel an issue that may be worrying some readers: how do we know that a formula *isn't* provable?

The problem is this. We defined $\vdash \phi$ to mean that it is possible to extend the initial tableaux $F\phi$ to a closed tableaux. Now, by changing the order we apply expansion rules, it may be possible to expand the initial tableau $F\phi$ to many different rule-saturated tableaux. This is worrying: perhaps some of these tableaux are open and others are closed! So, if we expand $F\phi$ and obtain an open rule-saturated tableau, this doesn't seem to tell us much. Perhaps some other expansion would have led to a closed tableaux. Perhaps—what a horrible thought!—we have to check all possible tableaux expansions before we can safely conclude that ϕ isn't provable.

In fact, we don't need to worry. For suppose some expansion of $F\phi$ leads to an open rule-saturated tableau. Then, by the property just noted, it is possible to falsify ϕ . It follows that no other expansion of $F\phi$ can possibly lead to a closed rule-saturated tableaux. To see this, note that if there was such a closed rule-saturated tableau, this would mean that ϕ was provable. By soundness, this would mean that ϕ was valid—but this would contradict the fact that we can falsify ϕ . In short, as soon as we succeed in expanding $F\phi$ to an open rule-saturated tableau, it is safe to conclude that ϕ is *not* provable.

Soundness and completeness are not the only desirable properties of proof systems. If one is interested in implementation, one wants to know something about a system's computational properties. For example, is the method guaranteed to *terminate* (or *halt* or *stop*) on all possible input formula? And how *efficient* is the method anyway?

¹Incidentally, it is easy to give examples of tableaux systems that are sound but *not* complete—simply throw away expansion rules! For example, if we discard the rule F_{\neg} , we still have a *sound* tableaux system, but we don't have enough power left to prove all validities, as we can no longer have all the rules we need for coping with negated formulas.

The expansion process used by our propositional tableaux is guaranteed to halt on all possible input. To see why, note that every tableaux expansion rule takes a signed formula and returns a *finite* number of signed formulas (in fact, no rule returns more than four signed formula). Moreover, crucially, each of the formula returned contains fewer connectives than the input formula. Thus, as we never have to apply a rule to the same formula twice, after every expansion the collection of unexpanded formulas remains finite, and moreover, any new unexpanded formulas we obtain as a result of rule applications are *simpler*. Thus—no matter which sequence of expansions we choose to make—we will achieve rule-saturation after a finite number of steps. Termination is thus guaranteed. Making this argument really tight requires a little care, but the basic idea should be clear.

Are tableaux an efficient method of doing propositional inference? The honest answer is: *nobody knows*. In fact nobody knows if there is *any* proof method which can prove all valid formulas efficiently. Let's discuss this a little further.

For many formulas, tableaux fare better than other methods. To give an obvious example, suppose we tested the validity of the following formula using truth tables:

$$(q \rightarrow (r \vee (p \wedge \neg t \rightarrow \neg\neg s))) \vee \neg(q \rightarrow (r \vee (p \wedge \neg t \rightarrow \neg\neg s))).$$

The required truth table has 32 rows, which takes quite some time to fill out. The tableaux method, on the other hand, gets this example right in a single step. More generally, the tableaux method is obviously a fairly sensible, and by using a little common-sense when forming tableau (for example, applying as many conjunctive expansion rules as possible before applying disjunctive expansion rules) it is possible to give fairly concise proofs of typical validities.

Unfortunately, not all validities are 'typical'. It is possible to show that there are an infinite number of formulas which create terrible problems for the tableau method: any closed tableau formed by starting with one of these formulas contains a huge number of nodes. Finding such validities isn't easy. Devising them requires careful thought about which kind of formulas require heavy use of disjunctive rules. However, such validities are out there, thus although the tableau method looks better than the truth table method, the bottom line is that it is subject to the same kind of combinatorial explosions.

Now—as far as anybody knows—this isn't because tableaux is a poor proof method. In fact, nearly all the well known proof methods (such as resolution, natural deduction, and so on) have been shown to be subject to similar combinatorial explosions. There simply is no known proof method for propositional languages that avoids combinatorial explosion. In fact, it is widely conjectured that no such proof method exists. For further discussion of this point, consult the references cited in the Notes.

Let us sum up what we have learned in this section. As a proof method, the tableaux system introduced in this chapter is as well-behaved as one could reasonably hope for. It is semantically sensible, and moreover, strong enough to prove all valid formulas. Moreover, it is guaranteed to terminate on all possible input. On 'typical' formulas it seems reasonably

efficient—though the reader should be aware that there are (infinitely many) formulas which, if given as input, will give rise to tableaux that are too big to be practicably computed.

Software Summary of Chapter 4

`propTabl.pl` The file that contains all the predicates for the implementation of our theorem prover. (page 237)

Notes

Proof theory is a rich and fascinating subject. The reader can gain a good overview of various proof systems, and how they relate to each other, by consulting Sundholm 1983. Both signed and unsigned tableaux systems are discussed.

For tableaux systems, the classic source is Smullyan 1995; our discussion of signed tableaux is based on Smullyan's treatment. Good treatments of signed tableaux may be found in Fitting 1996 and Bell and Machover 1977. The reader interested in finding out how to prove the tableaux method complete will find all that is needed in these three books. Our implementation of propositional tableaux was influenced by the implementation of unsigned tableaux given in Fitting (1996). The reader will find it instructive to compare the two implementations. (Apart from anything else, it's a good idea to find out about what unsigned tableaux are and how they work.)

The analysis of the complexity of proof methods is an active field of research. Perhaps the best starting point is Urquhart 1995. This contains a good discussion of why the tableau method sometimes leads to combinatorial explosion. Another interesting (and very readable) paper is D'Agostino 1992, which shows that in certain cases the tableaux method performs *worse* than the truth table method.

Finally, although it takes us far from the concerns of the present book, it is worth noting that proof theory arguably has far deeper connections with natural language semantics than our discussion might suggest. There is an important tradition which claims that meaning shouldn't be explained in terms of truth conditions, but in terms of assertability conditions. (Roughly speaking, on this view the meaning of an utterance is the reason we have for holding it, not the situations in which it is true.) Semanticists in this tradition tend to regard proofs as the primary semantic objects. At first sight this may appear to be a rather strange view, but it has a lot to recommend it. A good introduction to this line of thought is Sundholm 1986. Moreover, as Ranta (1994) demonstrates, the approach is of relevance to computational semantics.

Chapter 5

First-Order Inference

In this chapter we extend our tableaux system to first-order languages. As promised, we do so by reducing first-order formulas to propositional ones. The resulting system is natural and easy for human beings to use. Unfortunately, as a candidate for automation it's awful: it makes use of an infinite space of possible substitutions that requires human insight to navigate.

So we rethink our strategy. Roughly speaking, we decide to delay the problem of finding good substitutions. Our plan is to use free variables to build up a set of constraints on substitutions, and then to hand over the problem of solving the constraints to a unification component. We discuss unification in detail, consider how to combine unification with the tableaux proof method, and implement a first-order theorem prover. We conclude with a discussion of the problems facing first-order theorem provers (and in particular, the problems raised by equality) and discuss alternative approaches.

5.1 First-Order Tableaux

We now extend the signed tableaux proof method to first-order languages. Conceptually, the extension is very simple. Our new tableaux rules will allow us get rid of quantifiers by substituting suitable terms for bound variables. In effect, they let us reduce first-order formulas to propositional ones.

What kinds of rules enable us to do this? Let's consider two examples. Suppose a tableau contains the signed formula $T\forall x\text{KILLER}(x)$, and suppose we are working with a first-order language containing the constants JULES and BUTCH and a 1-place function symbol FATHER. Then it should be legitimate to extend the tableau by adding on any of the following signed formulas: $T\text{KILLER}(\text{JULES})$, $T\text{KILLER}(\text{BUTCH})$, and

$$T\text{KILLER}(\text{FATHER}(\text{FATHER}(\text{FATHER}(\text{JULES}))))).$$

(After all, if everyone is a killer, every term picks out a killer.) More generally, given any universal formula prefixed by the sign T , we should be free to throw away the universal quantifier, substitute any term for the newly freed variable in the matrix, prefix the result by T , and extend the tableau with the result.

Analogous tableaux extensions should be legitimate when we are told that an existential sentence is false. For example, suppose a tableau contains $F\exists x\text{SHOOT}(\text{JULES},x)$. Then it should be legitimate to deduce that $F\text{SHOOT}(\text{JULES},\text{FATHER}(\text{JULES}))$, and that $F\text{SHOOT}(\text{JULES},\text{BUTCH})$, and so on. (After all, if it's false that Jules shoots someone, then it's false that Jules shoots any person we care to name.)

Such examples lead us to formulate the following two tableaux rules, T_{\forall} and F_{\exists} :

$$\frac{T\forall x\phi}{T\phi(\tau)} \qquad \frac{F\exists x\phi}{F\phi(\tau)}$$

Here $\phi(\tau)$ denotes the result of replacing the variable bound by the quantifier by some closed term τ . (Recall that a closed term is a term that does not contain any variables.) We call these two rules *universal rules*.

While both these rules are sound, and clearly trade on the semantic intuitions underlying the tableaux method, they differ from the propositional tableaux rules in important respects. First, each rule licenses as many tableaux extensions as there are closed terms in the language. In particular, if there are infinitely many closed terms in the language, each rule licenses infinitely many ways of extending a tableau. Second, in general we will have to apply a universal rule more than once to a particular occurrence of a signed formula. Here's a simple example. We shall show that

$$\forall x\text{DIE}(x) \rightarrow \text{DIE}(\text{MIA}) \wedge \text{DIE}(\text{ZED})$$

is a theorem.

1	$F(\forall x\text{DIE}(x) \rightarrow \text{DIE}(\text{MIA}) \wedge \text{DIE}(\text{ZED}))$	
2	$T\forall x\text{DIE}(x)$	$1, F_{\rightarrow}$
3	$F(\text{DIE}(\text{MIA}) \wedge \text{DIE}(\text{ZED}))$	$1, F_{\rightarrow}$
4	$T\text{DIE}(\text{MIA})$	$2, T_{\forall}$
5	$T\text{DIE}(\text{ZED})$	$2, T_{\forall}$
6	$F\text{DIE}(\text{MIA})$	$3, F_{\wedge}$
7	$F\text{DIE}(\text{ZED})$	$3, F_{\wedge}$

Note the way we had to apply T_{\forall} twice to line 2: once to get $T\text{DIE}(\text{MIA})$, and once to get $T\text{DIE}(\text{ZED})$. Thus we have lost one of the pleasant properties of propositional tableaux: it is no longer true that we are through with an occurrence of a signed formula once we've applied a rule to it.

Let's turn to the next issue. What sort of tableaux rules are needed to deal with signed formulas of the form $T\exists x\phi$ or $F\forall x\phi$? This is a more subtle matter. Suppose a tableau contains the signed formula $T\exists x\text{KILLER}(x)$. It is *not* legitimate on the basis of this information to deduce that $TKILLER(\text{JULES})$, or that $TKILLER(\text{BUTCH})$, or indeed to deduce that $TKILLER(\text{CLOSED-TERM})$ for *any* closed term of the language we are working with. *Someone* is a killer—but we don't know who. So how are we to eliminate the quantifier?

Actually, the solution is straightforward: we invent a brand new name for the entity whose existence is asserted, and eliminate the quantifier by substituting this new name. Such brand new names are called *parameters*, and by using parameters we can deal with true existential statements, and false universal ones too. For example, if a tableau contains $T\exists x\text{KILLER}(x)$, we give this killer (whoever he or she is) a new name (that is, we choose a new parameter, say c) and extend the tableau by asserting $TKILLER(c)$. Similarly, given the signed formula $F\forall x\text{RELIGIOUS}(x)$, we christen the unbeliever (whoever he or she is) with a new name (say c_8) and extend the tableau by asserting $FRELIGIOUS(c_8)$.

Let's make these ideas precise. Suppose we are working in a first order language over some vocabulary V . Let PAR be a (countably infinite) set of new constant symbols, that is, constant symbols that *don't* belong to V . We'll call these new constant symbols parameters, and reserve the symbols c, c_1, c_2, \dots , and so on, for them. From now on, when we want to do tableaux proofs, we *won't* work in our original language (that is, the language built over the vocabulary V). Rather we'll work in the first-order language whose vocabulary consists of all the original vocabulary V , and in addition, all the new constant symbols in PAR .

Given these ideas, it's easy to define the rules F_{\forall} and T_{\exists} :

$$\frac{F\forall x\phi}{F\phi(c)} \qquad \frac{T\exists x\phi}{T\phi(c)}$$

Here $\phi(c)$ denotes the result of substituting a parameter c that we haven't used so far in the tableau proof, for the newly freed variable in the matrix. We call these two rules *existential rules*.

Two points about these rules must be grasped. First, when we use the existential rules, it is absolutely vital that we substitute parameters that haven't been used so far in the tableau construction. To see why, suppose a tableau contains both $T\exists x\text{KILLER}(x)$ and $T\exists x\text{RELIGIOUS}(x)$. Suppose we first apply T_{\exists} to $T\exists x\text{KILLER}(x)$ using the parameter c_5 (which, let us suppose, hasn't been previously used) to name the killer. That is, we extend the tableau with $TKILLER(c_5)$. Now, if at some later stage we apply T_{\exists} to $T\exists x\text{RELIGIOUS}(x)$, it would be outrageous to re-use the parameter c_5 . If we did this (that is, if we 'deduced' that $TRELIGIOUS(c_5)$) we would in effect be claiming that there is a single individual (namely the one named by c_5) who is both a killer and religious. This simply doesn't follow from the given information. All we know is that there is at least one

killer, and at least one religious person. They may well be different people, hence we need to assign each of them a fresh new name. In short, once we've used an existential rule to replace a quantifier by a parameter, that parameter becomes 'old', and cannot later be re-used.

There is a second important point that the reader should note: the definition of the existential rules has consequences for the universal rules. When carrying out first-order tableaux proofs we no longer work in the original first-order language, but in the original language enriched with an infinite collection PAR of new constants. Now, constant symbols—including parameters—are closed terms, so we should be free to substitute parameters when we use a universal rule. (In fact, it's *crucial* that we be able to do this, as the following example will make clear.) But this means that we will *always* have infinitely many choices when it comes to using a universal rule, even if the original language had no constants or function symbols at all.

Let's have a look at another example. We will show that

$$\exists x \forall y \text{SHOOT}(x,y) \rightarrow \forall y \exists x \text{SHOOT}(x,y)$$

is a theorem. (This is a rather pretty example. It's simple, but puts all four quantifier rules to work.)

1	$F(\exists x \forall y \text{SHOOT}(x,y) \rightarrow \forall y \exists x \text{SHOOT}(x,y))$	
2	$T \exists x \forall y \text{SHOOT}(x,y)$	1, F_{\rightarrow}
3	$F \forall y \exists x \text{SHOOT}(x,y)$	1, F_{\rightarrow}
4	$T \forall y \text{SHOOT}(c_1, y)$	2, T_{\exists}
5	$F \exists x \text{SHOOT}(x, c_2)$	3, F_{\forall}
6	$T \text{SHOOT}(c_1, c_2)$	4, T_{\forall}
7	$F \text{SHOOT}(c_1, c_2)$	5, F_{\exists}

The key point to notice about this proof is the way the existential and universal rules interact. In particular, note the way we used the existential rules to introduce the new parameters (c_1 in line 4 and c_2 in line 5) and then used the universal rules to make further use of these symbols. It should be clear from this example that it is vital that the universal rules have access to the parameters.

How efficient is this tableaux system? This is a question we must answer in two ways.

The first (and most fundamental) thing the reader should know is the following: the first-order tableaux system does *not* give rise to an algorithm for determining which first-order formulas are valid. An algorithm is a recipe which, when given an instance of a problem to solve, halts after a finite number of steps with the correct answer. *There is no algorithm at all for determining the validity of arbitrary first-order formulas.* That is, first-order validity is *undecidable*. It is certainly possible to implement syntactic proof systems (for example, tableaux systems) and we shall do in this chapter, but no implementation of any

system is guaranteed to terminate on all possible input. Incidentally, the tableaux system just described *is* (sound and) complete. That is, if a formula ϕ is valid, then it is possible to construct a (finite) closed tableaux that has $F\phi$ as its root node. Thus the proof system just described is a genuine syntactic analysis of the semantic concept of first-order validity, and indeed a rather natural one. However this analysis does not yield an algorithm for determining which first-order formulas are valid, for no such algorithm exists.

As there is no general computational solution to the problem, it is meaningless to speak of efficiency in an absolute sense, so let's restate the question in a more realistic way. Many first-order theorem provers are useful practical tools. Proof search won't always terminate, but they work well on wide ranges of input, and can be used as components of larger systems quite satisfactorily. So the question we should next pose is: how good is the tableaux system just described as a practical basis for automated first-order theorem proving? This question has a clear answer: *it's terrible*. Although our tableaux system is conceptually simple, it's a computational nightmare.

The heart of the problem lies with the universal rules. They offer us an infinite menu of substitutable terms. If $T\forall x\phi$ belong to a tableau, then we are free to extend it by first adding $\phi(c_1)$ then $\phi(c_2)$ then $\phi(c_3)$, \dots , and so. (We could have done this in the previous tableau, for example, starting at line 4.) Most such extensions will be completely pointless. In the examples given above, it was intuitively clear which substitutions were sensible; we used our sense of what was relevant to guide our choice of substitutions. Unfortunately, computers lack our intuitions. If we want a reasonably practical implementation, we need to find a method choosing substitutions that doesn't depend on human insight.

Here's what we shall do. First, we'll change the universal rules slightly: we'll never substitute closed terms, we'll *always* substitute free variables instead. In a sense, we are not going to make a real choice of substitutions at all—we are going to use free variables as 'dummies' that will enable us to delay making this decision. In this way, we will gradually build up a whole system of 'substitution equations', a set of constraints on variable values that contain a great deal of information about the terms in the tableau we are building. Crucially, there is an algorithm for solving such constraint sets, the famous *unification algorithm*. We will use unification to look for solutions to the constraints that lead to branch closure.

That's our strategy in outline. An awful lot of detail remains to be filled in. For a start, blending unification with tableaux is going to force us to rethink the existential rules, and there are many other details that will require careful attention. Nonetheless, unification is the key to further progress, so let's examine this concept in some detail.

5.2 Unification and Free Variable Tableaux

Unification is the process of carrying out substitutions on two terms so that they become identical. The reader who has made it this far will certainly have a informal idea of what substitutions are, and indeed, since we are working with Prolog, a good working knowledge of what unification involves in practice. However the Prolog version of unification is *not* suitable for theorem proving purposes. Given two terms, Prolog does not make the occurs check, rather it just rushes ahead and tries to unify them. This is fine (and certainly efficient) if the terms are unifiable, but it can lead to non-terminating computations if they are not. So we will need to think carefully about unification and how to implement it. Once we have done this, we will need to think about how to integrate unification into tableaux theorem proving. All in all, we have a lot of work to do, so lets get to start right away.

Unification

Suppose we have chosen the first-order language we are going to work with. Then a *substitution* is a function that maps the set of variables to the terms of this language. We use the notation $x\sigma$ (rather than $\sigma(x)$) to denote the value of x under the substitution σ .

We are most interested in *finite substitutions*. These are substitutions which only assign new terms to a finite number of variables; the rest they leave alone. That is, if σ is a finite substitution, then for all but a finite number of variables, $x\sigma = x$.

The simplest finite substitution is the one which does not assign new terms to *any* of the variables. This substitution is called the *identity substitution*, and we denote it by $\{\}$. We also have a special notation for other finite substitutions, namely:

$$\{x_1/\tau_1, \dots, x_n/\tau_n\}.$$

Here x_1, \dots, x_n are distinct variables, τ_1, \dots, τ_n are terms, and $\tau_i \neq x_i$ for any i from 1 to n . The notation x_i/τ_i means that the variable x_i is mapped to τ_i .

As we have defined them, substitutions only act on variables. In this section we will be exploring the effect of of substitutions on arbitrary terms. Here's a recursive definition of this concept.

Substitutions on Terms. Let σ be a substitution and τ a term. Then:

1. If τ is a variable x , then $\tau\sigma = x\sigma$;
2. If τ is a constant, then $\tau\sigma = \tau$;
3. If τ has the form $f(\tau_1, \dots, \tau_n)$, then $[f(\tau_1, \dots, \tau_n)]\sigma = f(\tau_1\sigma, \dots, \tau_n\sigma)$. (Here f is an n -place function symbol.)

Actually, we need to extend the substitution concept even further, for in the following section we will want to apply substitutions to formulas, and indeed, to entire tableaux. So let's define these concepts right away. First a piece of notation. If σ is a substitution, then by σ_x we mean the substitution that is exactly like σ except that $x\sigma_x = x$. We can now recursively define $\phi\sigma$, the result of applying the substitution σ to the formula ϕ .

Substitutions on formulas.

1. If $R(\tau_1, \dots, \tau_n)$ is an atomic formula, then $[R(\tau_1, \dots, \tau_n)]\sigma$ is $R(\tau_1\sigma, \dots, \tau_n\sigma)$. (Here R is an n -place relation symbol.);
2. $[\neg\phi]\sigma$ is $\neg[\phi\sigma]$;
3. $[\phi \wedge \psi]\sigma$, $[\phi \vee \psi]\sigma$, and $[\phi \rightarrow \psi]\sigma$, are $[\phi\sigma] \wedge [\psi\sigma]$, $[\phi\sigma] \vee [\psi\sigma]$, and $[\phi\sigma] \rightarrow [\psi\sigma]$ respectively;
4. $[\forall x\phi]\sigma$ is $\forall x[\phi\sigma_x]$, and $[\exists x\phi]\sigma$ is $\exists x[\phi\sigma_x]$.

Substitutions on signed tableaux. If σ is a substitution and \mathcal{T} is a signed tableau, then $\mathcal{T}\sigma$ is the signed tableau obtained by replacing every signed formula of the form $T\phi$ in \mathcal{T} by $T[\phi\sigma]$, and every signed formula of the form $F\phi$ in \mathcal{T} by $F[\phi\sigma]$.

With these definitions out of the way, let us return to our main task: understanding unification.

Because substitutions are functions we can compose them in the usual way. That is, if σ_1 and σ_2 are substitutions, then we can define a new substitution $\sigma_1\sigma_2$, which we call the *composition* of σ_1 and σ_2 . For every variable x , we define $x(\sigma_1\sigma_2)$ to be $(x\sigma_1)\sigma_2$. That is, $\sigma_1\sigma_2$ first carries out the substitution σ_1 and then carries out the substitution σ_2 .

There is one more concept we need to discuss before we can define unification. Let's approach it via an example. Suppose we want to make the terms $f(c,y,w)$ and $f(x,y,g(z))$ identical (where c is a constant, w , x , y , and z are variables, f is a 3-place function symbol, and g is a 1-place function symbol). Let σ_1 be the substitution $\{x/c, w/g(z)\}$. Applying σ_1 to these terms has the desired result, for $f(c,y,w)\sigma_1 = f(x,y,g(z))\sigma_1 = f(c,y,g(z))$.

But there are other ways of making these terms identical. For example, let σ_2 be the finite substitution $\{x/c, w/g(z), y/h(u,x)\}$. Applying σ_2 to either term yields $f(c,h(u,x),g(z))$. Nonetheless, clearly σ_1 is a more general solution to the problem: σ_2 does too much work. For suppose we apply σ_1 to some term. Then, if we want to (or need to) we are always free to later give y the value $h(u,x)$. To do so we need simply apply the substitution $\{y/h(u,x)\}$, and this two step process gives us the same effect we would have achieved by directly applying σ_2 . But of course, we might not want to take this further step. (Perhaps mapping y to $h(u,x)$ is incompatible with other substitutions we need to make.) If we use σ_2 to solve the problem, we are overcommitting ourselves.

Such considerations motivate the following definition. A substitution σ_1 is said to be *more general* than a substitution σ_2 if and only if there is some substitution θ such that $\sigma_2 = \sigma_1\theta$. That is, σ_1 is more general than σ_2 if we can get the effect of σ_2 by first carrying out σ_1 and then making a further substitution θ . Thus, reverting to our motivating example, $\{x/c, w/g(z)\}$ is more general than $\{x/c, w/g(z), y/h(u, x)\}$ because there is a substitution θ (namely $\{y/h(u, x)\}$) such that $\{x/c, w/g(z), y/h(u, x)\} = \{x/c, w/g(z)\}\theta$.

Incidentally, note that under this (standard) definition of ‘more general than’, each substitution σ is more general than itself. This is because we can get the effect of σ by first carrying out σ and then carrying out the identity substitution $\{\}$. Thus the ‘more general than’ relation is reflexive. It is also transitive. That is, if σ_1 is more general than σ_2 , and σ_2 is more general than σ_3 , then σ_1 is more general than σ_3 . The reader may like to try proving this.


We are ready for the key definition.

Unification. Let τ_1 and τ_2 be terms. A substitution is a *unifier* for τ_1 and τ_2 if and only if $\tau_1\sigma = \tau_2\sigma$. Terms τ_1 and τ_2 are said to be *unifiable* if and only if they have a unifier. A substitution σ is a *most general unifier* (or *mgu*) for two terms if and only if it is unifier for these terms, and is more general than any other unifier.

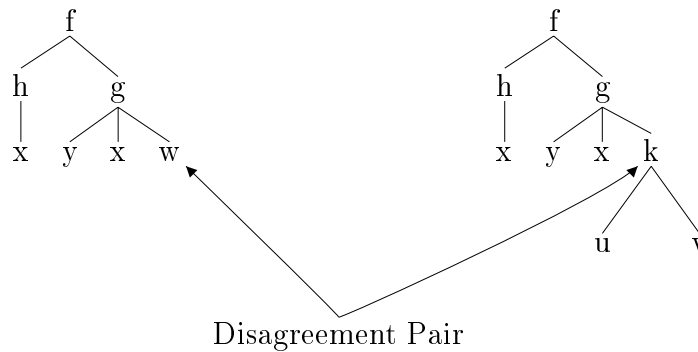
Now we know what unification is—but what is involved in computing unifiers? Ideally, what we want is an algorithm which will take as input two terms and determine whether or not they are unifiable. If the terms are unifiable, it should return their unifier as output. If they are not unifiable, it should halt and tell us so.

Such algorithms exist. Let’s consider one of the more straightforward ones in some detail.

To appreciate what this algorithm does, we really need to think of terms as trees. For example, consider the term $f(h(x), g(y, x, w))$. Its parse tree—that is, the tree showing how it is built up out of sub-terms—looks like this:



When are two terms different? For our purposes, the following answer is the most useful: two terms are different if and only if their parse trees contain at least one *disagreement pair*. What’s a disagreement pair? Here’s an example.



Intuitively, the pair of terms $(w, k(u, v))$ is a disagreement pair for these terms because they are distinct terms that occupy ‘corresponding places’ in the two parse-trees. What is meant by ‘corresponding places’? Simply the nodes that one reaches by following the same sequence of transitions from the root in the two trees. For example, if we follow the transition sequence ⟨second daughter, third daughter⟩ from the root of the first tree we arrive at the node labeled w , and if we follow the same transition sequence in the second tree we arrive at the node labeled k .

Disagreement pairs make terms different, thus unification algorithms should try to eliminate disagreement pairs. When are disagreement pairs eliminable, and how can they be eliminated?

First, suppose that two terms τ_1 and τ_2 are different because there a disagreement pair (d_1, d_2) such that *neither* d_1 nor d_2 is a variable. Then there is nothing we can do. No substitution can help us out. The terms τ_1 and τ_2 are not unifiable.

Now for the tricky question. Suppose that one of these terms (d_1 say) is a variable. Are the two terms unifiable? The answer is ‘yes’, *provided that the variable d_1 does not occur in d_2* . Think about it. Suppose d_1 is a variable, say x , and that x *doesn’t* occur in d_2 . Then we can eliminate this disagreement pair very easily: we simply need to replace x by d_2 (That is, we need simply perform the substitution $\{x/d_2\}$.) To give a concrete example, the disagreement pair $(w, k(u, v))$ shown above is of this form. We eliminate it by replacing w by $k(u, v)$.

On the other hand, if d_1 is a variable (say x) and x *does* occur in d_2 , then unification is impossible. For example, suppose that d_2 is $f(x)$. Then, no matter what value we choose for x , we will *never* render these two terms identical—we’ll always have that extra function symbol f to reckon with.

Let’s make these observations a little more systematic. Suppose we have found a disagreement pair (d_1, d_2) We will call this pair a *simple* disagreement pair if and only if at least one of the terms d_1 or d_2 is a variable that does not occur in the other. (It could happen, of course, that *both* d_1 and d_2 are variables that don’t occur in the other—for example if $d_1 = x$ and $d_2 = y$. That’s fine. As long as there’s at least one we’re happy.) Simple disagreement pairs are the ones we can repair. We do so by carrying out what we shall call

the *relevant repair*. If d_1 is a variable that does not occur in d_2 , then the relevant repair is the substitution $\{d_1/d_2\}$. If d_2 is a variable that does not occur in d_1 , then the relevant repair is the substitution $\{d_2/d_1\}$. If both d_1 and d_2 are variables that don't occur in the other, then there are two substitutions that will repair the problem, namely $\{d_1/d_2\}$ and $\{d_2/d_1\}$. We'll arbitrarily stipulate that in such cases the relevant repair is $\{d_1/d_2\}$.

On the other hand, if the disagreement pair we have found is *not* simple, there's nothing we can do. Either we have that neither d_1 nor d_2 is a variable, or one of them is a variable that occurs in the other term.

We now know which disagreement pairs are eliminable, and how to carry out the elimination. And this means, we are only one small step away from an algorithm for solving the unification problem. To unify two terms, simply try and eliminate *all* the disagreement pairs! This idea immediately suggests the following non-deterministic algorithm:

```

input terms  $\tau_1$  and  $\tau_2$ 
let  $\sigma := \{\}$ 
while  $\tau_1\sigma \neq \tau_2\sigma$ 
    choose a disagreement pair  $(d_1, d_2)$  for  $\tau_1\sigma, \tau_2\sigma$ 
    if  $(d_1, d_2)$  is not simple then
        write  $\tau_1$  and  $\tau_2$  are not unifiable and HALT
    else
        let  $\sigma := \sigma \cup \rho$ , where  $\rho$  is the relevant repair
    endif
endwhile

```

This is a genuine computational solution of the unification problem. No matter which two terms it is given as input, it will halt after finitely many steps. When it halts, it will either have told us that the terms are not unifiable (and if it says this, it's right!) or it will have found out how to build the mgu of the two terms. These claims are not obvious; they require proof. The reader interested in finding out more should consult the references cited in the Notes.

In fact, the algorithm has one additional property: the mgus it produces are *idempotent*. That is, if σ is an mgu produced by this algorithm, then $\sigma\sigma = \sigma$. Why on earth is such an abstract looking property interesting? The answer is: idempotent mgus provide us with an easy way of solving the *simultaneous unification problem*, and actually this problem will be important when working with tableaux.

We will be using unification to try and close tableaux branches. That is, we will look for branches containing pairs of atomic formula $T(R(\tau_1, \dots, \tau_n))$ and $F(R(\tau'_1, \dots, \tau'_n))$. If we can find a substitution σ that makes τ_1 identical to τ'_1 , and \dots , τ_n identical to τ'_n , then by applying σ we obtain a branch containing contradictory formula, that is, a closed branch. Thus we need to solve the problem of finding a single substitution that makes n pairs of terms identical; this is called the simultaneous unification problem.

Now, this problem may look harder than the ordinary unification problem, but actually it's not. It can be solved as follows. To find a substitution that identifies n pairs of terms, first find an *idempotent* mgu σ_1 for τ_1 and τ'_1 . We can do this using the above algorithm. Next, find an *idempotent* mgu σ_2 for $\tau_2\sigma_1$ and $\tau'_2\sigma_1$. Again we can do this using the above algorithm. Next, find an *idempotent* mgu σ_3 for $\tau_3\sigma_1\sigma_2$ and $\tau'_3\sigma_1\sigma_2 \dots$. In fact, all we need to do is keep 'chaining together' the solutions to each individual pair. The substitution $\sigma = \sigma_1\sigma_2 \dots \sigma_{n-1}\sigma_n$ that is obtained in this way is a simultaneous mgu for the n pairs of terms. For this 'chaining together' method to work, the substitutions constructed at each step must be idempotent, and as the above algorithm yields idempotent mgus, it really does deliver everything we shall need for tableaux theorem proving.

The basic concepts should now be clear, so let's turn to a more practical issue: how can we use Prolog to unify terms? Now, we could approach this problem from scratch. However (as the reader is undoubtedly aware) one of the main mechanisms underlying Prolog is a form of unification. Perhaps we can use Prolog's in-built unification mechanism directly, or adapt it straightforwardly, to perform the term unifications we require? This certainly seems a sensible choice of strategy. Let's see what it involves.

The way Prolog performs unification is reasonably close to the above algorithm. There are two main differences. First, the algorithm given above is non-deterministic, whereas Prolog's in-built mechanism for unifying terms is deterministic. Clearly this is not a difference that need concern us. The second difference, however, is important. In the interests of efficiency, Prolog does not bother making the 'occurs check'. That is, given a disagreement pair (d_1, d_2) , one of which is a variable, Prolog does *not* check whether this variables occurs in the other term, but will go straight ahead and attempt to carry out what it thinks the required repair is. This leads it to attempt to unify terms that aren't unifiable, and thus to stack overflows and other undesirable behaviour.

Now, we are interested in using unification as part of a first-order theorem prover. We really need to know whether or not two terms are unifiable, and we certainly *don't* want Prolog to mess things up with its 'Hey, just go for it!' behaviour. Can Prolog be tamed? That is, can we make use of its in-built unification mechanism, but modify it so that makes an occurs check? The answer is 'yes'. The following code, adapted from Chapter 10 of Sterling and Shapiro (Sterling and Shapiro 1986), does precisely this.

The main predicate `unify` uses the in-built unification predicate `=` to carry out the required unifications, but carefully inserts an occurs check (via the `notOccursIn/2` predicate) where required.

```

unify(X,Y):-
    var(X), var(Y), X=Y.

unify(X,Y):-
    var(X), nonvar(Y), notOccursIn(X,Y), X=Y.

unify(X,Y):-
    var(Y), nonvar(X), notOccursIn(Y,X), X=Y.

unify(X,Y):-
    nonvar(X), nonvar(Y), atomic(X), atomic(Y), X=Y.

unify(X,Y):-
    nonvar(X), nonvar(Y), compound(X), compound(Y), termUnify(X,Y).

```

So how is `notOccursIn/2` to be defined? Well, if the term is a variable, then the key idea is to make use of the in-built Prolog metalogical predicate `\==` (Recall that query `X \== Y` succeeds if and only if `X` and `Y` are not identical.) This is coded in the first clause. The second clause handles the case where the term is atomic (we don't have to do anything else in such a case), and the third clause takes care of the case where the term is complex.

```

notOccursIn(X,Term):-
    var(Term), X \== Term.

notOccursIn(_,Term):-
    nonvar(Term), atomic(Term).

notOccursIn(X,Term):-
    nonvar(Term), compound(Term),
    functor(Term,_,Arity), notOccursInComplexTerm(Arity,X,Term).

```

The `notOccursInComplexTerm` predicate recursively checks whether the variable `X` does not occur in either of the arguments of the complex term `Term`, making use of the `notOccursIn` predicate as you would have expected.

```

notOccursInComplexTerm(N,X,Y):-
    N > 0, arg(N,Y,Arg), notOccursIn(X,Arg),
    M is N - 1, notOccursInComplexTerm(M,X,Y).

notOccursInComplexTerm(0,_,_).

```


All that remains is to define the predicates which were used in the definition of `unify/2`. The predicate `termUnify/2` first checks if the terms have the same functor symbols and arity, and then uses `unifyArgs/3` to unify their arguments.

```
termUnify(X,Y):-
    functor(X,Func,A), functor(Y,Func,A),
    unifyArgs(A,X,Y).

unifyArgs(N,X,Y):-
    N > 0, M is N - 1,
    arg(N,X,ArgX), arg(N,Y,ArgY),
    unify(ArgX,ArgY), unifyArgs(M,X,Y).

unifyArgs(0,_,_).
```

Free variable Tableaux

Let's see how we can use unification to define a more computationally realistic signed tableaux proof system.

As we discussed earlier, the problem with our first attempt at a first-order tableaux system lay with the universal rules. Because they offered us too many ways of making substitutions (in fact, infinitely many) human insight was required to find proofs. By substituting free variables, and gradually building up a system of constraints which we will solve using unification, we hope to bypass the need for human insight. Let's work through this idea in detail, and see where it leads. As a first step, here are our new universal rules:

$$\frac{T\forall x\phi}{T\phi(\bar{v})} \qquad \frac{F\exists x\phi}{F\phi(\bar{v})}$$

Here $\phi(\bar{v})$ denotes the result of replacing all instances of the variable that the quantifier bound by a new variable \bar{v} that does not occur bound anywhere in the tableau. (This restriction is simply to prevent any 'accidental bindings' taking place. In fact, every time we apply these rules, we're going to substitute a previously unused variable.)

So far so good—but a moment's thought will convince the reader that our new strategy could lead to serious problems with the existential rules. Recall that the basic idea behind the existential rules was to invent a brand new name for the entity asserted to exist (we called these new names 'parameters') and to eliminate the quantifier by substituting these parameters. Unification threatens to undercut this strategy completely: the substitutions it makes may undo all our careful choices of new names! In short, we seem to be in a dilemma. The use of free variables in the universal rules is essentially a 'delaying' device;

we don't want to make a real choices of what to substitute, we want unification to sort it all out for us. But if we delay in this way, how can we guarantee that unification will respect the 'new names' concept that is crucial to the existential rules?

A very clever—and very simple—idea allows us a way round this problem. We are going to substitute *structured* terms when we use the existential rules; the term structure itself will ensure that unification cannot spoil anything. To be more precise, we are to use what are known as *Skolem terms*. Without further ado, here are the new existential rules:

$$\frac{F\forall x\phi}{F\phi(s(x_1, \dots, x_n))} \qquad \frac{T\exists x\phi}{T\phi(s(x_1, \dots, x_n))}$$

Here s is a new *Skolem function symbol*, and x_1, \dots, x_n are all the free variables in ϕ . (If there are no free variables, s is a new *Skolem constant*.) What does this mean, and why does it work?

The basic idea is a straightforward generalisation of what we did earlier with parameters. Once again, instead of working with the first-order language built over the original vocabulary V , we are going to choose a set of new symbols SKO consisting of a countably infinite set of Skolem constants (these are essentially the same as our earlier parameters) and for every natural number n a countably infinite set of Skolem function symbols of arity n . When carrying out tableaux proofs, we won't work in the original first-order language, but first-order language whose vocabulary consists of all the original vocabulary V plus all the new symbols in SKO .

The symbols in SKO enable us to manufacture new names. Crucially, however, because we now have Skolem *function* symbols at our disposal, we can do something that we couldn't do with parameters: we can 'build in newness' in a way that will survive the unification operation! Look at the term the existential rules demands we substitute: $s(x_1, \dots, x_n)$, where x_1, \dots, x_n are all the free variables in ϕ . Now recall our discussion of the occurs check. Quite simply, $s(x_1, \dots, x_n)$ *cannot* unify with any of x_1, \dots, x_n . Our new 'Skolem structured' term really will be new.

"Oh really?" the skeptical reader will demand. "The basic idea is clear, but the restriction simply isn't strong enough! When we used parameters, we had to pick a parameter that was brand new *to the entire tableau*. But look: the above rule only ensures that the new structured name won't unify with any of the free variables in ϕ itself! We should really insist that the new instantiated name be $s(x_1, \dots, x_n)$, where x_1, \dots, x_n are all the free variables *in the entire tableau!*"

This is an intelligent observation. Our skeptical reader is certainly thinking about the rule in the right way. The suggestion is sensible: instantiations certainly *could* be carried out that way. Nonetheless, we don't *have* to (and for this we should be grateful, since it means we can substitute smaller Skolem terms). The simpler rule stated above is sound. This is not obvious, and for a proof (and further discussion) of this matter the reader should

consult the sources cited in the Notes.¹

We now have the new quantifier rules we need. Only one task remains: bringing unification into play.

We want unification to be the ‘intelligence’ guiding the proof search. We’re going to use our quantifier rules essentially ‘blindly’: we’re simply going to build up a set of constraints and hope that unification can do something useful with them. In particular, we hope that unification will be able to close branches for us. How could it do this for us? Let’s consider an example. We shall redo our proof of

$$\exists x \forall y \text{SHOOT}(x, y) \rightarrow \forall y \exists x \text{SHOOT}(x, y)$$

using our new quantifier rules and unification. Here are the first 7 steps of the construction.

1	$F(\exists x \forall y \text{SHOOT}(x, y) \rightarrow \forall y \exists x \text{SHOOT}(x, y))$	
2	$T \exists x \forall y \text{SHOOT}(x, y)$	1, F_{\rightarrow}
3	$F \forall y \exists x \text{SHOOT}(x, y)$	1, F_{\rightarrow}
4	$T \forall y \text{SHOOT}(s_1, y)$	2, T_{\exists}
5	$F \exists x \text{SHOOT}(x, s_2)$	3, F_{\forall}
6	$T \text{SHOOT}(s_1, v_1)$	4, T_{\forall}
7	$F \text{SHOOT}(v_2, s_2)$	5, F_{\exists}

The first 5 lines are essentially identical with the previous version, save that we have used the Skolem constants s_1 and s_2 (rather than the parameters c_1 and c_2) in lines 4 and 5. (Note that we don’t need to use Skolem functions as neither formula contains free variables.) The real difference occurs in lines 6 and 7. In both lines we have ‘blindly’ instantiated in new variables, namely v_1 and v_2 . So we don’t (yet) have closure.

But closure is easy to get. Consider the substitution $\{v_1/s_2, v_2/s_1\}$. If we apply this to the tableau (recall that we defined the concept of applying a substitution to a tableau in the previous section) we get the following tableau:

1	$F(\exists x \forall y \text{SHOOT}(x, y) \rightarrow \forall y \exists x \text{SHOOT}(x, y))$	
2	$T \exists x \forall y \text{SHOOT}(x, y)$	1, F_{\rightarrow}
3	$F \forall y \exists x \text{SHOOT}(x, y)$	1, F_{\rightarrow}
4	$T \forall y \text{SHOOT}(s_1, y)$	2, T_{\exists}
5	$F \exists x \text{SHOOT}(x, s_2)$	3, F_{\forall}
6	$T \text{SHOOT}(s_1, s_2)$	4, T_{\forall}
7	$F \text{SHOOT}(s_1, s_2)$	5, F_{\exists}

This tableau *is* closed.

¹Readers with some background in logic should try proving it themselves. Once you’ve been alerted to the fact that the simpler rule is sound, it’s not that difficult to work out why. (Like so many things in logic and computer science, it’s spotting the simplification in the first place that’s tricky.)

This example motivates the addition of the following *Atomic MGU Closure Rule*:

Suppose that \mathcal{T} is a tableau formed from some initial tableau \mathcal{I} , and that some branch of \mathcal{T} contains a pair of atomic formulas of the form $T(R(\tau_1, \dots, \tau_n))$ and $F(R(\tau'_1, \dots, \tau'_n))$. Then $\mathcal{T}\sigma$ is also a tableau formed from the initial tableau \mathcal{I} , where σ is a simultaneous mgu of τ_1 and τ'_1, \dots , and τ_n and τ'_n .

This rule is rather different from the other rules we've seen: it's not an extension rule, rather it's a transformation rule. It tells us that if we have tableau formed from an initial set \mathcal{I} , and we transform it by applying a substitution having certain properties, we obtain another tableau for the same initial set.

The basic idea guiding the choice of transformation should be clear: we want to apply substitutions that could lead to branch closure. One point, however, may be puzzling. Obviously we are interested in pairs of formulas of the form $T\phi$ and $F\phi$, but why have we restricted our attention to *atomic* formulas? The answer is: *simplicity*. In fact we could have formulated a more general rule, but if we did so we would have to state it carefully (we would need to avoid 'accidental capture' of variables) and checking for accidental capture would be computationally expensive. By restricting our attention to atomic formulas we avoid these difficulties.

And that's our free variable tableaux proof system. Frankly, it's not nearly as nice as our previous system if one wants to prove things by hand—playing with Skolem functions swiftly get unwieldy, and thinking in terms of unification is cumbersome. But it wasn't designed with the needs of humans in mind, it was designed for automation. And, as we shall now see, for this purpose it is really rather good.

5.3 Implementing Free Variable Tableaux

We shall now present an implementation of the free variable signed tableaux proof system. The implementation is an extension of the propositional tableaux implementation—but it can't be described as a straightforward extension. While the basic ideas underlying signed free variable tableaux make it possible to devise practical implementations, there is more detail to take care of than in the propositional case. Moreover, the fact remains that first-order logic is undecidable, so the implementation has to defuse the (very real) threat of non-terminating tableaux constructions, and this requires a little care.

Our implementation draws on the Fitting (Fitting 1996) implementation of free variable unsigned propositional tableaux. Roughly speaking, we take our signed propositional tableaux implementation as the starting point, and extend it to a first-order system by adopting many of Fitting's design choices. For further information, see the Notes.

Before examining the main body of the code, let's take a quick look at the main supporting routines we shall use. Obviously we need predicates to handle substitution. The workhorse is the `substitute` predicate (supplied in `comsemPredicates.pl`), which we examined in Section 2.4. Recall that the `substitute/4` predicate takes a term, a variable, and a formula as its first three arguments, and returns in its fourth argument the result of substituting the term for each free occurrence of the variable in the formula. With it predicate at our disposal, it is straightforward to define what we mean by instances of quantified formulas:

```
instance(t(forall(X,F)),Term,t(NewF)):-
    substitute(Term,X,F,NewF).

instance(f(exists(X,F)),Term,f(NewF)):-
    substitute(Term,X,F,NewF).

instance(t(exists(X,F)),Term,t(NewF)):-
    substitute(Term,X,F,NewF).

instance(f(forall(X,F)),Term,f(NewF)):-
    substitute(Term,X,F,NewF).
```

To handle the existential rule correctly, we need to be able generate new Skolem function symbols on demand. The following predicate does this. (The predicates it calls may be found in `comsemPredicates.pl`.)

```
skolemFunction(VarList,SkolemTerm) :-
    newFunctionCounter(N),
    compose(SkolemTerm,fun,[N|VarList]).
```

We shall also need to know what free variables a formula contains. We associate this information explicitly with each formula. The following predicate does this.

```
notatedFormula(n(Free,Formula),Free,Formula).
```

We are now ready to discuss the main code. As in the propositional case, the outermost predicate is called `saturate`. This recursively attempts to rule-saturate the input tableau with the aid of the `expand` predicate. Moreover, as in the propositional case, the base clause of `saturate` tests for closure via a predicate called `closed`.

But there is an important difference. The first-order version `saturate` has a second argument, a number called `Qdepth` (which can be read as 'quantification depth'). This number is the maximum number of times that we are allowed to apply universal rules in the course of constructing a tableau. This is the mechanism which wards off the threat of non-terminating tableaux constructions.

```

saturate(Tableau,_) :-
    closed(Tableau).

saturate(OldTableau,Qdepth):-
    expand(OldTableau,Qdepth,NewTableau,NewQdepth),!,
    saturate(NewTableau,NewQdepth).

```

The `closed/1` predicate tests for branch closure using unification. It attempts to find a pair of signed formulas of the form $t(X)$ and $f(X)$, and then uses the `unify` predicate discussed in the previous section to test for branch closure.

```

closed([]).

closed([Branch|Rest]):-
    member(NotatedOne,Branch),
    notatedFormula(NotatedOne,_,t(X)),
    member(NotatedTwo,Branch),
    notatedFormula(NotatedTwo,_,f(Y)),
    unify(X,Y),
    closed(Rest).

```

Like its propositional cousin, the `expand/4` predicate is a high level organisational predicate which works its way recursively through the branches of the input tableau, and tries to apply the various kinds of expansion. Its two extra arguments keep track of the input and output quantification depths. Quantification depth is unaffected by all expansions save universal expansions. Each universal expansion ‘uses up’ one of our predetermined quota of expansions and thus decrease the quantification depth by 1.

In addition, `expand/4` performs a more interesting task. Note the use of `append` (in the clause handling universal expansions) to glue the new branch back onto the *end* of the tableau. Why do this? Essentially, it’s an attempt to ensure that our predetermined quota of universal expansions are ‘spread around fairly’. It would be rather silly to use up our entire quota on a single branch. By using `append` to ‘rotate’ the branches on the tableau, we ensure that every branch receives its fair share of universal expansions.

```

expand([Branch|Tableau],QD,[NewBranch|Tableau],QD):-
    unaryExpansion(Branch,NewBranch).

expand([Branch|Tableau],QD,[NewBranch|Tableau],QD):-
    conjunctiveExpansion(Branch,NewBranch).

expand([Branch|Tableau],QD,[NewBranch|Tableau],QD):-

```

```

    existentialExpansion(Branch, NewBranch) .

expand([Branch|Tableau], OldQD, NewTableau, NewQD) :-
    universalExpansion(Branch, OldQD, NewBranch, NewQD),
    append(Tableau, [NewBranch], NewTableau) .

expand([Branch|Tableau], QD, [NewBranch1, NewBranch2|Tableau], QD) :-
    disjunctiveExpansion(Branch, NewBranch1, NewBranch2) .

expand([Branch|Rest], OldQD, [Branch|Newrest], NewQD) :-
    expand(Rest, OldQD, Newrest, NewQD) .

```

Now for the predicates that actually carry out the expansions. The expansion predicates for unary, conjunctive, and disjunctive formulas are essentially the same as in the propositional case. The only difference is that we have to translate a notated formula into a signed formula to calculate its components, and then translate the components back again into notated components.

```

unaryExpansion(Branch, [NotatedComponent|Temp]) :-
    unary(SignedFormula, Component),
    notatedFormula(NotatedFormula, Free, SignedFormula),
    removeFirst(NotatedFormula, Branch, Temp),
    notatedFormula(NotatedComponent, Free, Component) .

conjunctiveExpansion(Branch, [NotatedComp1, NotatedComp2|Temp]) :-
    conjunctive(SignedFormula, Comp1, Comp2),
    notatedFormula(NotatedFormula, Free, SignedFormula),
    removeFirst(NotatedFormula, Branch, Temp),
    notatedFormula(NotatedComp1, Free, Comp1),
    notatedFormula(NotatedComp2, Free, Comp2) .

disjunctiveExpansion(Branch, [NotComp1|Temp], [NotComp2|Temp]) :-
    disjunctive(SignedFormula, Comp1, Comp2),
    notatedFormula(NotatedFormula, Free, SignedFormula),
    removeFirst(NotatedFormula, Branch, Temp),
    notatedFormula(NotComp1, Free, Comp1),
    notatedFormula(NotComp2, Free, Comp2) .

```

The most interesting are `existentialExpansion/2` and `universalExpansion/4`. Note the use of `skolemFunction/2` to instantiate a new Skolem term in `existentialExpansion`. In `universalExpansion/4`, note that quantification depth is decremented, a new variable `V` is ‘blindly’ instantiated, and then—on the very last line—note the way `append` is used

to replace the universal formula we have been working with back on *end* the branch. Why do this? Well, we need to replace the formula because (as we have already discussed) we will often have to re-use universal formulas. But then it makes very good sense to replace the formula at the *end* of the branch. If we leave it where it is, we run the risk of using up our entire quota of universal rule applications on this one formula. It is far more sensible to ensure fairness by ‘rotating’ the universal formulas on the branch, much as we rotated the branches of the tableau earlier.

```

existentialExpansion(Branch, [NotatedInstance|Temp]):-
    notatedFormula(NotatedFormula,Free,SignedFormula),
    existential(SignedFormula),
    removeFirst(NotatedFormula,Branch,Temp),
    skolemFunction(Free,Term),
    instance(SignedFormula,Term,Instance),
    notatedFormula(NotatedInstance,Free,Instance).

universalExpansion(Branch,OldQD,NewBranch,NewQD):-
    OldQD > 0, NewQD is OldQD - 1,
    member(NotatedFormula,Branch),
    notatedFormula(NotatedFormula,Free,SignedFormula),
    universal(SignedFormula),
    removeFirst(NotatedFormula,Branch,Temp),
    instance(SignedFormula,V,Instance),
    notatedFormula(NotatedInstance,[V|Free],Instance),
    append([NotatedInstance|Temp],[NotatedFormula],NewBranch).

```

There is one tricky point in the definition of the universal expansion predicate: the use of `member/2` in the very first lines of it. At first sight this line of code seems superfluous (note that we don’t have it in the other expansion predicates). However, if we leave out the call to `member/2`, `universal/1` will pick the first of the two universal formulas format, adds it component to the beginning of the branch, and puts the universal formula back at the end of the branch (via `removeFirst` and `append`). In a case where there are two different universal formulas on a branch, one of them will never be touched, and the Q-depth will be wasted on the other. Adding `member` at the beginning of the clause overcomes this problem.

Now it only remains to spell out what kinds of signed formula we are working with.

```

conjunctive(t(X & Y),t(X),t(Y)).
conjunctive(f(X v Y),f(X),f(Y)).
conjunctive(f(X > Y),t(X),f(Y)).

```



```

disjunctive(f(X & Y),f(X),f(Y)).
disjunctive(t(X v Y),t(X),t(Y)).
disjunctive(t(X > Y),f(X),t(Y)).

unary(t(~X),f(X)).
unary(f(~X),t(X)).

universal(t(forall(_,_))).
universal(f(exists(_,_))).

existential(t(exists(_,_))).
existential(f(forall(_,_))).

```

As with the propositional implementation, it's nice to have a driver that translates the formula into a notated, signed formula, and calls the `saturate` predicate. We call this predicate `valid/2`, as in our implementation for propositional tableaux. Its second additional argument is the value of Q-depth.

```

valid(X,Qdepth):-
    notatedFormula(NotatedFormula,[],f(X)),
    saturate([[NotatedFormula]],Qdepth).

```

We conclude with a warning. Remember that this program can only construct those tableaux which require fewer applications of the universal rules than the user-imposed `Qdepth` limit. Thus if checking a formula with `valid/2` fails, this most emphatically does *not* mean “Not a first-order tableau theorem”! The formula in question may well be a tableaux theorem—perhaps we just set `Qdepth` to too small a value for a proof to be found. (Incidentally, while trying this program out on the kinds of formula found in introductory logic books, we usually had `Qdepth` set to 25.) On the other hand, if the program tells us that ϕ is a first-order tableaux theorem, this is not open to question. If a closed tableau has been formed from the initial tableau $F\phi$, then ϕ is provable, and that's that.

Exercise 5.3.1 Add a pretty print predicate to our implementation of first-order tableaux, that shows the branches and the proof steps in a readable way.

5.4 Off-the-shelf Inference Tools

To be provided....

Software Summary of Chapter 5

`freeVarTabl.pl` The file that contains the code for free variable tableaux, using a number of ideas from Melvin Fitting's implementation for first-order logic. (page 240)

`callTheoremProver.pl` The Prolog-interface to the theorem prover Otter. (page 245)

`callModelBuilder.pl` The Prolog-interface to the model builder Mace. (page 246)

`fol2otter.pl` Translates a formula in Otter syntax to standard output. (page 247)

Notes

The human-oriented tableaux system presented at the start of the chapter is a typical first-order signed tableaux system; essentially identical systems may be found in Smullyan 1995 and Sundholm 1983, and Bell and Machover 1977, Fitting 1996, Smullyan 1995, Sundholm 1983 discuss it's unsigned counterpart. Incidentally, we really meant what we said in the text. If you want to do tableaux proofs by hand, use *this* system; its free-variable cousin wasn't designed for people but for computers.

Unification and the unification algorithm were introduced in Robinson 1965 for the purposes of automated theorem proving. The algorithm discussed in the text is essentially Robinson's original. The literature unification is now huge, and a wide variety of unification algorithms have been investigated. We suggest that the reader who wants to know more try Fitting 1996 or Apt 1995. Fitting's discussion is very clear, and should be accessible to most readers of this book. Apt's discussion is more advanced, but still very approachable. He gives a nice account of idempotent substitutions, and analyses three unification algorithms, including the non-deterministic Robinson algorithm presented here. This is also a good source for further pointers to the literature.

The free variable proof system presented in the text is a signed version of the unsigned free variable tableaux system presented in Fitting 1996, and the reader interested in learning more about the soundness and completeness of the system is advised to start there (perhaps using Smullyan as a backup). Any skeptical readers still worried by the soundness of the existential rule should consult Hähnle and Schmitt (1994). A useful general source is Gallier (1986). This is a detailed introduction to the theory of first-order theorem proving that covers tableaux, unification and much else besides.

To extend our propositional implementation to first-order logic, we have by and large

followed the Fitting (1996) strategy for unsigned tableaux. In particular, our use of `Qdepth`, and the subsequent decision to use `append` to ‘rotate’ branches of tableaux, and universal signed formulas on branches, follows Fitting. (Incidentally, this technique has a name. We are essentially treating both tableaux and branches as *priority queues*. The idea of doing this in tableaux proofs dates back to Smullyan (1995)). The implementation of `skolemFunction` is taken from Fitting, and the `notatedFormula` predicate is Fitting’s `notation` and `fmla` predicates rolled into one.

Chapter 6

Putting It All Together

We now have some answers to the two questions with which these notes began.

1. *We can build first-order representations in a compositional way for simple natural language expressions. (Moreover, we are able to do in a way that takes scope ambiguities into account.)*
2. *We have automated the process of performing reasoning with (equality free) first-order representations.*

Along the way, we've developed a number of useful tools. Now it's time to have a little fun and bring the various bits and pieces together. As we shall see, by 'plugging together' our lambda calculus, quantifier storage, model checking and theorem proving programs we are able to define simple question handling and argumentation predicates. All that's involved is some minor extensions to our DCGs, and a few simple driver definitions.

6.1 Natural Language Questions

From Chapter 2 we know how to compositionally build semantic representations for simple expressions of English, and from Chapter 1 we know how to evaluate first-order formulas in well-named models. Now, well-named models are just (rather primitive) databases. Let us extend our grammar coverage so that we can pose simple natural language database queries to them.

First we introduce some new syntactic categories: **q** for questions, and **wh** for wh-phrases. Second, we define a DCG rule that makes a question out of a wh-phrase and a verb phrase:

$$q(\text{WH@VP}) \text{--> } wh(\text{WH}), vp2(\text{VP}).$$

(Note that this rule is the essentially the same as the rule that makes sentences out of noun phrases and verb phrases.) The third step is to add the rules for the wh-phrases:

```
wh(lambda(P,lambda(X,P@X))--> [who] .
```

```
wh(lambda(P,lambda(X,(N@X) & (P@X))))--> [which], n2(N).
```

With these entries we can handle questions like ‘Who is a customer?’ or ‘Which robber loves Honey Bunny?’. The semantic representations that are derived for questions are lambda expressions, for example:

```
?- q(Sem,[who,is,a,customer],[ ]).
```

```
Sem = lambda(X,customer(X))
```

The final step is to write a driver that translates a question into a lambda expression, feeds it to an example model, and returns the answer. (The example models we use here are those of Chapter 1.) The driver can be coded as follows:

```
query(Example):-
    readLine(Question),
    q(Q,Question,[ ]),
    betaConvert(Q,lambda(Answer,Sem)),
    evaluate(Sem,Example),
    write(Answer), nl, fail.
```

With `readLine/1` Prolog waits for the user to enter a question. The user’s input is given to the DCG, which returns the representation for the question entered. After conversion, we require the semantic representation being of the form `lambda(Answer,Sem)`, that is, a question.

Now, the way we designed the treatment of questions, `Sem` will contain a free occurrence of the variable `Answer`. Feeding the formula `Sem` to the evaluation predicate, after making it a sentence, generates all possible (and correct) assignments for `Answer`, as we have learned from Chapter 1.

Here is an example query, evaluated to example model 1:

```
?- query(1).
```

```
> Who loves Honey Bunny?
pumpkin
```

In model 1, only Pumpkin loves Honey Bunny. Our driver tries to find other instantiations for answers, but it—correctly—fails. So, all it displays is the answer `pumpkin`.

We could extend the driver, and let it ask again for a question, by adding a second clause for `query/1`. After all, the first clause will *always* fail (due to the usage of `fail` in it).

```
query(Example):-  
    query(Example).
```

Now we have changed our driver in a “never ending conversation”! Here is a snapshot of an example session:

```
?- query(1).  
  
> Who loves Honey Bunny?  
pumpkin  
  
> Who does not love Honey Bunny?  
mia  
vincent  
honey_bunny  
  
> Which robber loves Honey Bunny?  
pumpkin  
  
> Which robber does not love Honey Bunny?  
honey_bunny  
  
>
```

Recall the vocabulary and the content of example model 1, and note the way negated questions are answered. Handling negated questions in this correct way is due to our revised model checker of Chapter 1.

Exercise 6.1.1 [intermediate] Extend the grammar fragment and the driver predicate in such a way that yes-no questions are covered as well.

Exercise 6.1.2 [hard] Reimplement wh-questions by using the gap-threading technique to pass on the variable of the WH-phrase down the syntax tree, as required in ‘Who did Vincent kill?’ (For gap-threading, consult Pereira & Shieber.)

6.2 Natural Language Argumentation

We might want to let the computer decide whether a natural language argument is valid. (A natural language argument is a sequence of sentences, the last of which is called the *conclusion*, the rest, *premises*.) Here, for example, is a simple two premise natural language argument:

- (1)
$$\frac{\text{Vincent knows every boxer.} \\ \text{Butch is a boxer.}}{\text{Vincent knows Butch.}}$$

This argument is *valid*. If the premises are true, then the conclusion is true too.

Now, we already have the basic tool needed to test arguments for validity, namely our first-order tableaux theorem prover. What we need to add is a driver which hands the theorem prover a list of formulas instead of just a single formula. The theorem prover will then attempt to make the premises true and the conclusion false. If it can't do this—that is, if it terminates with a closed tableau—then the argument is valid.

Let's implement a predicate which does this. We'll call it `validArgument/2`, and it will feed the theorem prover with a list of formulas, where the first formula of the list is the conclusion, and the rest of the list are premises. We will also need a predicate `premises/3` that turns a list of formulas in a list of signed, notated formulas.

```
validArgument([Conclusion|Premises],Qdepth):-
    notatedFormula(NotatedConclusion,[],f(Conclusion)),
    premises(Premises,[],NotatedPremises),
    saturate([[NotatedConclusion|NotatedPremises]],Qdepth).

premises([],N,N).
premises([P|Premises],SoFar,NotatedPremises):-
    notatedFormula(NotatedP,[],t(P)),
    premises(Premises,[NotatedP|SoFar],NotatedPremises).
```

The only thing missing is a driver that takes a natural language argument, gives the sentences in the argument to a DCG that builds first order representations, and then feeds the representations to `validArgument`. But this is easily done by a predicate called `argument/0`:

```
argument:-
    enterPremises(Premises),
    enterConclusion(Conclusion),
```



```
(
    validArgument([Conclusion|Premises],10),
    nl,write('This is a valid argument!'),nl,!
;
    nl,write('Not a valid argument...'),nl
).
```

When starting this predicate, the user is asked to enter some premises, and a conclusion, which are then passed through as argument of `validArgument/2` (we allow 10 applications of the universal rules). If a proof is found, this predicate prompts “This is a valid argument!”, otherwise it prints “Not a valid argument...”.

Entering the premises is done with the help of the following (recursive) predicate:

```
enterPremises(Premises):-
    nl, write('Enter premises (or Return to continue):'),
    readLine(Input),
    (
        Input=[],!,
        Premises=[]
    ;
        d(SemPremise,Input,[]),
        betaConvert(SemPremise,ConvertedPremise),
        enterPremises(Others),
        Premises=[ConvertedPremise|Others]
    ).
```

This predicate—recursively—asks the user to enter premises and translates them into first-order logic, and returns a list of formulas. The recursion finishes if the user enters an empty sentence.

A similar predicate does this job for entering the conclusion:

```
enterConclusion(Conclusion):-
    nl, write('Enter conclusion:'),
    readLine(Input),
    d(SemConclusion,Input,[]),
    betaConvert(SemConclusion,Conclusion).
```

Time for a test session to see if this all works.

?- argument.

Enter premises (or Return to continue):

> Butch is a boxer.

Enter premises (or Return to continue):

> Vincent knows every boxer.

Enter premises (or Return to continue):

>

Enter conclusion:

> Vincent knows a boxer.

This is a valid argument!

Equally well, we can try and see if the whole thing works for arguments that are *not* valid. The following, for instance, is not a valid argument.

(2)
$$\frac{\begin{array}{l} \text{If Mia snorts, then Vincent smokes.} \\ \text{Vincent smokes.} \end{array}}{\text{Mia snorts.}}$$

So let's try:

?- argument.

Enter premises (or Return to continue):

> If Mia snorts, then Vincent smokes.

Enter premises (or Return to continue):

> Vincent smokes.

Enter premises (or Return to continue):

>

Enter conclusion:

> Mia snorts.

Not a valid argument...

Now Prolog gives the appropriate response after entering the third sentence, as expected.

Argumentation with Scope Ambiguities

Let's have a look at the following argument:

- (3)
$$\frac{\begin{array}{l} \text{A woman loves every man.} \\ \text{Every boxer is a man.} \end{array}}{\text{A woman loves every boxer.}}$$

Is this a valid argument? The answer is: *sometime yes, sometimes no*. The premises are ambiguous, so it depends on which reading we give them. It *is* a valid argument if the first premise is understood as: there is a woman such that she loves every man. (That is, the reading where 'a woman' out-scopes 'every man'.)

But, this is not the reading that `mainLambda.pl` gives us. And trying this example on our driver for natural language argumentation will generate an "invalid" as response. It would be nice to be able to test whether a natural language argument is valid under some reading of the premises. Obviously our program from Chapter 2 won't help us here, since it can't cope with scope ambiguities. However, if we consult `mainCooperStorage.pl` instead, we get both readings for the first premise:

```
?- s(Sem,[every,man,is,loved,by,a,woman],[ ]).

Sem = exists(Y,woman(Y)&forall(X,man(X)>love(Y,X))) ;

Sem = forall(X,man(X)>exists(Y,woman(Y)&love(Y,X))) ;

no
```

The first one is the reading we are interested in, the second isn't. So now we are able, in principle, to test arguments with scope ambiguities, and show, using our tools, that this argument has a valid reading.

Exercise 6.2.1 Revise the `argument/1` driver in such a way that it handles scope ambiguities. Incorporate one of the programs from Chapter 3 (Cooper Storage, Keller Storage, or Hole Semantics).


6.3 Building Ontologies

Although we have discussed in detail how to build meaning representations for natural language expressions out of the representations of the words they consist of, so far we haven't

paid any attention to the meaning of the words themselves. For example, we simply assumed that the meaning of ‘boxer’ was adequately represented by the symbol BOXER. Even worse, we left ‘five dollar shake’ unanalyzed and represented it by FIVE_DOLLAR_SHAKE. In short, we have said nothing about *lexical semantics*.

Lexical semantics is a fascinating (and difficult) subject, and one we won’t be able to explore in detail in this book. Nonetheless, it is worth knowing that it is possible to say something useful about word meaning by thinking about the relationships that hold *between* words. To illustrate this, we will present a simple treatment of *nouns*. We will describe the fine structure of entities (that is, the things nouns denote) by designing an *ontology of concepts*. The work involved is essentially *classification*. That is, we are going to divide the concepts that make up the world into different classes. Some of these classes will be *disjoint*. (For example, we could decide that male and female entities are disjoint.) On the other hand, some classes will *inherit* properties of other classes. (For example, we could classify the noun ‘man’ as male. As male is disjoint from female, we could then infer that ‘man’ inherits the property of being disjoint from female.)

These two relations—disjointness and inheritance—are closely interconnected. Here is a tree structure of four nouns (‘organism’, ‘plant’, ‘animal’, and ‘human’) that shows why and how.

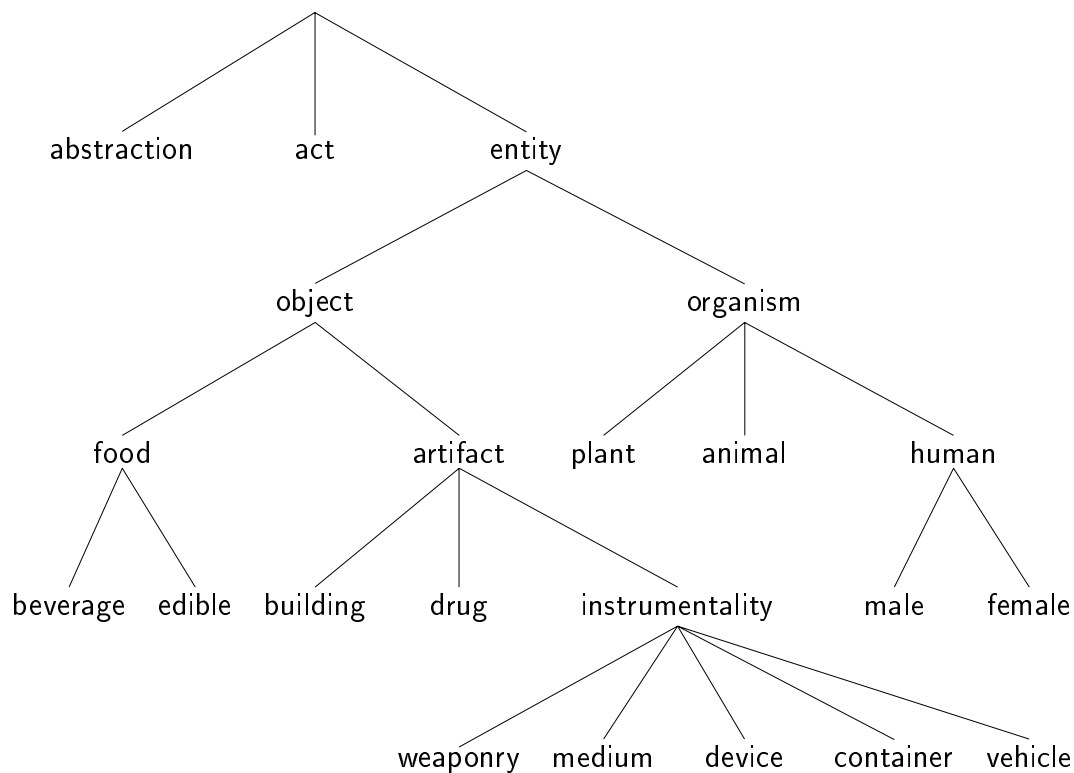


These are the kind of tree-structures that we will use to classify nouns, so it is important to know how to interpret them. First, all daughters of a mother node are disjoint. So, the tree above states that ‘plant’, ‘animal’, and ‘human’ are pairwise disjoint concepts. That is, according to this partial picture of the world, nothing can be a plant and an animal, nothing can be an animal and a human, and nothing can be a plant and a human. This classification (which is based on that used in the lexical database WordNet) seems to make sense. Of course one could quibble with the terminology—a biologist would insist that humans *are* animals—but for many purposes it is clearly a sensible classification.

The second kind of information this picture gives us concerns inheritance. All the daughter nodes inherit information from the mother nodes. According to the tree above, every ‘plant’ is an ‘organism’, every ‘animal’ is an ‘organism’, and every ‘human’ is an ‘organism’. Or to use the standard linguistic terminology for these relationships, ‘organism’ is a *hypernym* of ‘animal’, and ‘animal’ is a *hyponym* of ‘organism’. These relationships are *not* transitive. That is, while ‘object’ is a *hypernym* of ‘food’, and ‘food’ is a *hypernym* of ‘beverage’, the linguist would not say that ‘object’ was a *hypernym* of ‘beverage’. In graphical terms, to move to the hypernym we take one step up the tree to the mother node, and to move to a hyponym we move one step down to a daughter.

That's the basic idea. Let's apply it on a bigger scale, by classifying all the nouns in our lexicon. That is, we shall draw a tree ontology that covers all of the nouns in our lexicon. Now, obviously there are many ways to carry out this task, and it is also highly likely that if two different people classified the same set of nouns, they'd each come up with a different ontology, for everyone has a slightly different conception of the world. Is this a serious obstacle? For our purposes, no. We simply want to illustrate how important it is to provide *some* (consistent) picture of the world, and how to put this picture to work. But of course, if you want to adapt the tools provided in this book to some particular domain you will almost certainly want to extend the ontology provided below in various ways.

The classification of nouns that we will adopt is based on that provided by the lexical database WordNet. Here is the top structure of our ontology:




At the top of this structure we have ‘abstraction’ (a general concept formed by extracting common features from specific examples), ‘act’ (things that people do) and ‘entity’ (any concretely existing object). (You might want to argue that any ‘abstraction’ is a concretely existing entity. Fine: if you feel this way, change the tree so that ‘abstraction’ becomes a daughter of ‘entity’.) Then ‘entity’ is divided into ‘object’ (any non-living entity) and ‘organism’ (a living entity). The remaining classifications should be fairly self-explanatory.

This is the top structure of our ontology. And now, taking this tree as our point of departure, let's try to find our way through the world of nouns. That is, let's go through


our lexicon and relate each of the nouns to one of these nodes. Sometimes this is easy. For example, ‘episode’ and ‘joke’ are abstractions, ‘foot massage’ and ‘piercing’ are acts, a ‘five dollar shake’ is a beverage, and so on. Moreover, the framework outlined so far does enable us to deal with homonyms (words that have multiple unconnected meanings), as the reader is now asked to show.

Exercise 6.3.1 [easy] Think of a way to put homonyms (words that have multiple unconnected meanings) in the ontology of nouns. For instance, according to WordNet, ‘boxer’ could be someone who fights with his fists for sport, a worker who packs things into containers, or a breed of stocky medium-sized short-haired dog with a brindled coat and square-jawed muzzle developed in Germany.

But what about nouns like ‘boss’, ‘boxer’, ‘owner’, or ‘criminal’? These should definitely be classified as human, but they can’t be classified as simply as male or female: they can be either. However if we extend the tree for ‘human’ using the ideas introduced so far, the only possibility is:



This is clearly *wrong*. It states that boxers are disjoint from bosses, and that owners are disjoint from male and female — in short, precisely what we need to avoid. What are we to do? A simple solution (and the one we shall adopt) is to add yet another kind of relation to the ontology tree: one that includes inheritance, but excludes disjointness. We will use dotted lines for this kind of relation. Applying this idea gives us the following tree:



Using this extension, it is reasonably straightforward to classify all the nouns in our little lexicon. And, as we shall see, this information will be useful in at least two ways: First, we will use it as background knowledge for general first-order inference tasks (we shall do so by translating the inheritance and disjointness relations into first-order formulas). Second, it will be the basic source for determining sortal conflicts, a simple form of inference we can encode directly in Prolog.

But before turning to inference, let’s turn to a more pressing question: how shall we represent the ontology? Actually, there’s an even better question: where shall we *store*

the ontology? And, since we want to state the relations between *words*, what could be a better place to choose than the the lexicon (after all, this was specifically designed to contain information about words)? So let's agree to code the ontology tree in our lexicon. Here's how we shall do this:

```
lexicon(noun,car,[car],[vehicle]).
lexicon(noun,chainsaw,[chainsaw],[device]).
lexicon(noun,criminal,[criminal],human).
lexicon(noun,customer,[customer],human).
lexicon(noun,drug,[drug],[artifact]).
lexicon(noun,entity,[entity],[top]).
lexicon(noun,episode,[episode],abstraction).
lexicon(noun,female,[female],[human]).
```


For each noun in the lexicon, we fill the last slot with its hypernym noun. With the exception of `top`, all the hypernyms are part of the lexicon, too. Words that encode their hypernyms in square brackets (such as `car` and `chainsaw` above) are disjoint with respect to other words with the same hypernym (that is, square brackets encode the solid lines in the tree). On the other hand, words that don't do this (such as `criminal` and `customer`) are not required to be disjoint from other words with the same hypernym (that is, hypernyms listed without square brackets indicate that the dotted line relationship introduced in the previous tree diagram holds).

An important remark. As we've already remarked, the hypernym relation is not transitive. These relations are *not* transitive. That is, while 'object' is a *hyponym* of 'food', and 'food' is a *hyponym* of 'beverage', 'object' is not a *hyponym* of 'beverage'. Nonetheless, it is certainly true that anything that is a beverage is an object, and when performing inference we will be interested in drawing such conclusions, and this means that we will need to chain our way through the hypernym relation. But this leads to a practical point: if we have been careless, and specified the hypernym relation wrongly, we may get silly answers when we carry out such chaining. In particular we need to take care that there are no words w, v_1, v_2, \dots, v_n in our lexicon such that w is a hypernym of v_1 , v_1 is a hypernym of v_2 , \dots , and v_n is a hypernym of w . It would be useful to have a Prolog utility that checks that our ontological specifications don't contain cycles, and the next exercise asks the reader to provide one.

Exercise 6.3.2 Write a checker that browses through the lexicon, checking that it is never possible to cycle back to a word by moving along the hypernym relation for nouns.

Once we've stored our ontology in the lexicon, we want to use the information it contains to perform inference. Actually, we're going to do so in two distinct ways: by compiling the ontology into first-order formulas (so that we can use general first-order inference methods) and by using more restricted (but more efficient) Prolog-based inference.

Let's first see how to compile ontologies into first-order formulas. If P is a mother of Q, then we can infer that for all things that are Q, they also are P. And if Q and R have the same mother P, and both relationships are shown with a *solid* line, then we infer that no thing that is Q, is R. So, from the following tree



we can derive the formulas: $\forall x(Q(x) \rightarrow P(x))$, $\forall x(R(x) \rightarrow P(x))$, $\forall x(S(x) \rightarrow P(x))$, $\forall x(T(x) \rightarrow P(x))$, and also $\forall x(Q(x) \rightarrow \neg R(x))$.

Exercise 6.3.3 Do we need, in addition to the formulas above, also the formula $\forall x(R(x) \rightarrow \neg Q(x))$? Why or why not?

In Prolog, this translation is implemented as follows:

```

generateOntology(Formulas):-
    generateIsa(I0),
    generateDisjoint(I0-I1, I2),
    isa2fol(I1, []-F),
    isa2fol(I2, F-Formulas).
  
```

With the help of the Prolog built-in `setof` predicate, we collect all the inheritance relations in terms of the form `isa(Hypo,Hyper)`.

```

generateIsa(I):-
    setof(isa(Hypo,Hyper), Words^lexicon(noun, Hypo, Words, Hyper), I).
  
```

On the basis of this list, we calculate a list of `disjoint(A,B)` terms, and update the list of `isa/2` relations, by removing the square brackets from the hypernym (we don't need these anymore).

```

generateDisjoint([], [], []).

generateDisjoint([isa(A, [Hyper])|L1]-[isa(A, Hyper)|L2], I3):-!,
    forall(disjoint(A,B), member(isa(B, [Hyper]), L1), I1),
    generateDisjoint(L1-L2, I2),
    append(I1, I2, I3).

generateDisjoint([isa(A, Hyper)|L1]-[isa(A, Hyper)|L2], I):-
    generateDisjoint(L1-L2, I).
  
```


These lists contain the inheritance information (`isa/2`) and the disjointness relations (by `disjoint/2`). Now the translation into first-order representations is straightforward:

```
isa2fol([],A-A):- !.

isa2fol([isa(S1,[S2])|L],A1-[forall(X,F1 > F2)|A2]):- !,
    compose(F1,S1,[X]),
    compose(F2,S2,[X]),
    isa2fol(L,A1-A2).

isa2fol([isa(S1,S2)|L],A1-[forall(X,F1 > F2)|A2]):-
    compose(F1,S1,[X]),
    compose(F2,S2,[X]),
    isa2fol(L,A1-A2).

isa2fol([disjoint(S1,S2)|L],A1-[forall(X,F1 > ~ F2)|A2]):-
    compose(F1,S1,[X]),
    compose(F2,S2,[X]),
    isa2fol(L,A1-A2).
```

Running this predicate gives us a list of terms that look like:

```
?- generateOntology(0).

0 = [forall(A,abstraction(A)>top(A)),
     forall(B,act(B)>top(B)),
     forall(C,animal(C)>organism(C)),
     forall(D,artifact(D)>object(D)),
     forall(E,beverage(E)>food(E)),
     forall(F,bkburger(F)>edible(F)),
     forall(G,boss(G)>human(G)),.....]
```

This compilation opens the door to using general first-order inference techniques.

But sometimes we will want to use the ontology to perform inference directly in Prolog; while less general, this can be far more efficient. For instance, we would like to test whether something can be both a ‘radio’ and a ‘gun’, or both an ‘artifact’ and ‘suitcase’. We do this by the predicate `consistent/2`.

```
consistent(X,Y):-
    generateIsa(I),
    generateDisjoint(I-Isa,Disjoint),
    \+ inconsistent(X,Y,Isa,Disjoint).
```

We shall say that two items from our ontology are consistent if we can't prove that they are inconsistent. So when are two items inconsistent? First of all, if they are disjoint:

```
inconsistent(X,Y,_,Disjoint):-
    member(disjoint(X,Y),Disjoint).

inconsistent(X,Y,_,Disjoint):-
    member(disjoint(Y,X),Disjoint).
```

Second, they are inconsistent if by chaining up the tree from the first item, we can find a a hypernym that is inconsistent with the other item.

```
inconsistent(X,Y,Isa,Disjoint):-
    member(isa(X,Z),Isa),
    inconsistent(Z,Y,Isa,Disjoint).

inconsistent(X,Y,Isa,Disjoint):-
    member(isa(Y,Z),Isa),
    inconsistent(X,Z,Isa,Disjoint).
```

Now we can test this predicate with queries such as:

```
?- consistent(female,boss).

yes

?- consistent(male,radio).

no
```

Exercise 6.3.4 Our Prolog implementation of the `consistent/2` predicate is rather inefficient. Every time this predicate is put to work, it traverses through the entire lexicon to build up the set of `isa/2` and `disjoint/2` terms. But as the lexicon is quite likely to be static during parsing or semantic analysis, the same work is done for every call on `consistent/2`. Change this, by declaring `isa/2` and `disjoint/2` as dynamic predicates, that are asserted to Prolog's database at initialization of the module.

Exercise 6.3.5 [project] Use the ontology for nouns above and add sortal constraints to the semantics of verbs, such that 'Mia drinks a restaurant' is rejected, but 'Mia drinks a five dollar shake' is accepted.

6.4 Logical Redundancies

There are many important issues to do with quantifier scoping which we have ignored. One such issue has to do with logical redundancies: our algorithm will generate two logically equivalent readings for ‘Every boxer loves every woman’, or for ‘A man saw a woman’.

Exercise 6.4.1 [project] Write an algorithm that, after generating all the scoped representations with storage or hole semantics, checks whether these are logically equivalent. Use either our first-order tableaux program, or an “off-the-shelf” theorem prover to carry out the reasoning tasks.

Further, as we mentioned in Chapter 3, simply generating all possible permutations of quantifier scope is naive—ideally, we want world knowledge to filter out some of the possibilities. But if can think of a portion of world knowledge coded as first-order formulas, that are relevant to this problem, then we can use automated reasoning to perform the required filtration. For example, given a sentence with the noun phrase ‘a fist of every boxer’, our ‘generate all possibility’ methods will happily generate readings in which a fist is shared by several boxers, which is biologically implausible. Another example is the sentence ‘Every car has a radio’, in which the reading where all the cars in the domain share the same radio is quite unlikely.

Exercise 6.4.2 [project] Design a first-order theory (a set of formulas that put further constraints on the objects in our domain), that eliminate unwanted readings that are produced by scoping algorithms. Use either our first-order tableaux program, or an “off-the-shelf” theorem prover to carry out the reasoning tasks.

Software Summary of Chapter 6

`nlQuestions.pl` Simple example that shows how to combine model checking and semantic construction to get a system for natural language queries. (page 249)

`nlArgumentation.pl` Combines the theorem prover with a natural language interface to handle natural language argumentation. (page 251)

`semOntology` Predicates for working with the semantic ontology (page 253)

Notes

Wordnets

WordNet is a lexical database for English nouns, verbs, adjectives, and adverbs (Miller 1995). What WordNet makes a winner is that it has a huge coverage, and that all its sources are electronically available, in different kinds of frameworks (including Prolog). Probably the best introduction to WordNet is the recently published book (Fellbaum 1998). If you can't get hold of this book, read for an introduction the "Five papers on WordNet" (available on the Web), which explains the basic philosophy behind WordNet, and the structure and organization of its components.

Eliminating logically equivalent readings

Removing non-redundant scopings is not only interesting, it also reduces processing time. Nevertheless, relatively little work is done on eliminating logically equivalent readings. Vestre describes an algorithm that generates non-redundant quantifier scopings (Vestre 1991), without appealing to any kind of automated reasoning. A more general strategy is approached by Gabsdil and Striegnitz (1999), who use off-the-shelf inference engines to eliminate equivalent readings generated by a standard scoping algorithm, such as those provided in Chapter 3.

Afterword: Where To From Here?

To be supplied...

Bibliography

- Apt, K. (1995). *From Logic Programming to Prolog*. International series in computer science. Prentice Hall.
- Barendregt, H. (1984). *The Lambda Calculus: Its Syntax and Semantics*. North Holland.
- Barendregt, H. (1991). Lambda calculi with types. In D. G. S. Abramsky and T. Maibaum (Eds.), *Handbook of Logic in Computer Science*, pp. 118–279. Oxford University Press.
- Bell, J. and M. Machover (1977). *A Course in Mathematical Logic*. North Holland.
- Bos, J., B. Buschbeck-Wolf, M. Dorna, and C. Rupp (1998). Managing information at linguistic interfaces. In *The 17th International Conference on Computational Linguistics*, Montreal, pp. 160–166.
- Bos, J., B. Gambäck, C. Lieske, Y. Mori, M. Pinkal, and K. Worm (1996). Compositional Semantics in Verbmobil. In *The 16th International Conference on Computational Linguistics*, Copenhagen, Denmark, pp. 131–136.
- Bratko, I. (1990). *Prolog; Programming for Artificial Intelligence* (Second Edition ed.). Addison-Wesley.
- Clocksin, W. and C. Mellish (1987). *Programming in Prolog* (Third, Revised and Extended Edition ed.). Springer Verlag.
- Cooper, R. (1975). *Montague's Semantic Theory and Transformational Syntax*. Ph. D. thesis, University of Massachusetts.
- Cooper, R. (1983). *Quantification and Syntactic Theory*. Dordrecht: Reidel.
- Cooper, R., I. Lewin, and A. W. Black (1993). Prolog and Natural Language Semantics. Notes for AI3/4 Computational Semantics, University of Edinburgh.
- D'Agostino, M. (1992). Are tableaux an improvement on truth tables? *Journal of Logic, Language and Information* 1, 235–252.
- Dalrymple, M., J. Lamping, F. Pereira, and V. Saraswat (1997). Quantifiers, anaphora, and intentionality. *Journal of Logic, Language, and Information* 6, 219–273.
- Deransart, P., A. Ed-Dbali, and L. Cervoni (1996). *Prolog: The Standard. Reference Manual*. Springer.

- Dowty, D., R. Wall, and S. Peters (1981). *Introduction to Montague Semantics*. Dordrecht: Reidel.
- Enderton, H. (1972). *A Mathematical Introduction to Logic*. New York and London: Academic Press.
- Fellbaum, C. (Ed.) (1998). *WordNet. An Electronic Lexical Database*. The MIT Press.
- Fitting, M. (1996). *First-Order Logic and Automated Theorem Proving* (second ed.). Springer-Verlag.
- Franke, A. and M. Kohlhase (1999). System description: Mathweb, an agent-based communication layer for distributed automated theorem proving. In *16th International Conference on Automated Deduction CADE-16*.
- Gabsdil, M. and K. Striegnitz (1999). Classifying Scope Ambiguities. In C. Monz and M. de Rijke (Eds.), *ICOS-1, Inference in Computational Semantics*, Institute for Logic, Language and Computation (ILLC), Amsterdam, pp. 125–131.
- Gallier, J. (1986). *Logic for Computer Science*. New York: Harper Row.
- Gallin, D. (1975). *Intentional and Higher-Order Modal Logic*. North Holland.
- Gamut, L. (1991a). *Logic, Language, and Meaning. Volume I. Introduction to Logic*. Chicago and London: The University of Chicago Press.
- Gamut, L. (1991b). *Logic, Language, and Meaning. Volume II. Intensional Logic and Logical Grammar*. Chicago and London: The University of Chicago Press.
- Gazdar, G. and C. Mellish (1989). *Natural Language Processing in Prolog*. Addison-Wesley Publishing Company.
- Hähnle, R. and P. Schmitt (1994). The liberalized δ -rule in free variable semantic tableaux. *Journal of Automated Reasoning* 13, 211–221.
- Hindley, J. and J. Seldin (1986). *Introduction to Combinators and the Lambda Calculus*. London Mathematical Society Student Texts vol. 1. Cambridge University Press.
- Hobbs, J. R. and S. M. Shieber (1987). An algorithm for generating quantifier scopings. *Computational Linguistics* 13(1-2), 47–55.
- Hodges, W. (1983). Elementary predicate logic. In D. Gabbay and F. Guenther (Eds.), *Handbook of Philosophical Logic. Volume I.*, pp. 1–131. Reidel.
- Janssen, T. (1986a). *Foundations and Applications of Montague Grammar, Part 1: Philosophy, Framework, Computer Science*. CWI Tracts no. 19. Amsterdam: Centre for Mathematics and Computer Science.
- Janssen, T. (1986b). *Foundations and Applications of Montague Grammar, Part 2: Applications to Natural Language*. CWI Tracts no. 28. Amsterdam: Centre for Mathematics and Computer Science.
- Janssen, T. (1997). Compositionality. In J. van Benthem and A. ter Meulen (Eds.), *Handbook of Logic and Language*, pp. 417–473. Elsevier.

- Keller, W. R. (1988). Nested coooper storage: The proper treatment of quantification in ordinary noun phrases. In U. Reyle and C. Rohrer (Eds.), *Natural Language Parsing and Linguistic Theories*. Dordrecht: Reidel.
- Lewin, I. (1990). A Quantifier Scoping Algorithm without A Free Variable Constraint. In *Coling-90. Papers presented to the 13th International Conference on Computational Linguistics*, University of Helsinki, pp. 190–193.
- Miller, G. (1995). WordNet: A lexical database for English. *Communications of the ACM* 38(11), 39–41.
- Moore, R. (1989). Unification-Based Semantic Interpretation. In *Proceedings of the 27th Annual Meeting of the Association for Computational Linguistics*, Vancouver.
- Partee, B. (1997). Montague grammar. In J. van Benthem and A. ter Meulen (Eds.), *Handbook of Logic and Language*, pp. 5–91. Elsevier.
- Pereira, F. and S. Shieber (1987). *Prolog and Natural Language Analysis*. CSLI Lecture Notes 10. Stanford: Chicago University Press.
- Ranta, A. (1994). *Type-Theoretical Grammar*. Oxford: Clarendon Press.
- Reyle, U. (1993). Dealing with Ambiguities by Underspecification: Construction, Representation and Deduction. *Journal of Semantics* 10, 123–179.
- Robinson, J. (1965). A machine oriented logic based on the resolution principle. *Journal of the ACM* 12, 23–61.
- Smullyan, R. (1995). *First-Order Logic*. Dover Publications. This is a reprinted and corrected version of the 1968 Springer-Verlag edition.
- Sterling, L. and E. Shapiro (1986). *The Art of Prolog; Advanced Programming Techniques*. The MIT Press.
- Sundholm, G. (1983). Systems of deduction. In D. Gabbay and F. Guentner (Eds.), *Handbook of Philosophical Logic. Volume I.*, pp. 133–188. Reidel.
- Sundholm, G. (1986). Proof theory and meaning. In D. Gabbay and F. Guentner (Eds.), *Handbook of Philosophical Logic. Volume III.*, pp. 471–506. Reidel.
- Thomason, R. (1974). *Formal Philosophy. Selected Papers of Richard Montague*. New Haven: Yale University Press.
- Turner, R. (1997). Types. In J. van Benthem and A. ter Meulen (Eds.), *Handbook of Logic and Language*, pp. 535–586. Elsevier.
- Urquhart, A. (1995). The complexity of propositional proofs. *Bulletin of Symbolic Logic* 1, 425–467.
- Van Deemter, K. and S. Peters (Eds.) (1996). *Semantic Ambiguity and Underspecification*. CSLI Lecture Notes 55. Stanford: Chicago University Press.
- Vestre, E. J. (1991). An Algorithm for Generating Non-Redundant Quantifier Scopings. In *Fifth Conference of the European Chapter of the ACL*, Berlin, Germany, pp. 251–256.

Appendix A

Propositional Languages

The quantifier-free fragment of any first-order language (as the terminology suggests) simply consists of all formulas of the language that contain no occurrences of the symbols \exists or \forall . For example, $\text{ROBBER}(\text{PUMPKIN})$, $\text{CUSTOMER}(\text{MIA})$, $\text{CUSTOMER}(y)$, and

$$\text{CUSTOMER}(y) \rightarrow \text{LOVE}(\text{VINCENT},y)$$

are all quantifier-free formulas. On the other hand,

$$\forall y[\text{CUSTOMER}(y) \rightarrow \text{LOVE}(\text{VINCENT},y)]$$

isn't, as it contains an occurrence of the quantifier \forall .

The key thing to note about quantifier-free formulas is the following. Suppose we are given a model \mathbf{M} (of appropriate vocabulary) and an assignment g in \mathbf{M} . Now, normally we need to work with two semantic notions: satisfaction for arbitrary formulas and truth for sentences. However, when working with quantifier-free formulas, there are no bound variables to complicate matters, so this distinction is unnecessary. In fact, when working with a quantifier-free fragment, we may as well view each variable x as a constant interpreted by $g(x)$. If we do this, then every quantifier-free formula is either true or false in \mathbf{M} with respect to g . Moreover, it is obvious how to calculate the semantic value of complex sentences: conjunctions will be true if and only if both conjuncts are true, disjunctions will be true if and only if at least one disjunct is true, a negated formula will be true if and only if the formula itself is not true, and so on. (In short, we basically need to make the truth table calculations, which the reader is probably familiar with.)

This means we can simplify our notation somewhat. Because we don't have quantifiers, the internal structure of atomic formulas is irrelevant, for we're never going to bind any free variables they may contain. All that is important is whether the atomic symbols are true

or false, and how they are joined together using Boolean operators. For example, while it may be mnemonically helpful to choose propositional symbols such as `CUSTOMER(x)` or `LOVE(VINCENT,MIA)`, we lose nothing if we replace them by simpler symbols such as p and q . Following this line of thought leads us to define *propositional languages*. To specify a propositional language, we first say which symbols we are going to start with. (A fairly standard choice is p, q, r, s, t , and so on, often decorated with superscripts and subscripts: for example, p'' , r''' , or q_2 .) The chosen symbols are called proposition symbols, or atomic sentences. Complex sentences are built up using the standard Boolean connectives (for example, \neg, \wedge, \vee and \rightarrow) in the obvious way, and the truth values of complex sentences are calculated using the familiar truth table rules.

In short, propositional languages are essentially a simple notation for representing the quantifier-free formulas of first-order languages. When we devise inference mechanisms for first order logic in Chapters 4 and 5, it turns out to be sensible to first investigate inference methods for the quantifier free fragment (we do this in Chapter 4), and only then turn to the problem for the full first order language (the task of Chapter 5). Thus in Chapter 4, we make use of the simpler propositional notation just defined.

Appendix B

Type Theory

But What Does It All Mean?

We have given an informal, computationally oriented, introduction to the lambda calculus and its applications in semantic construction. We adopted a rather procedural perspective, encouraging the reader to think of the lambda calculus as a programming language—indeed, as the sort of language that emerges when one tries to generalize straightforward logic programming approaches to semantic construction (such as that of experiment 2).

However our account hasn't discussed one interesting issue: what do lambda expressions actually mean? Hopefully the reader now has a pretty firm grasp of what one can *do* with lambda expressions—but is one forced to think of lambda expressions purely procedurally? As we are associating lambda expressions with expressions of natural language, it would be nice if we could give them some kind of model theoretic interpretation.

Actually, there's something even more basic we haven't done: we haven't been precise about what counts as a λ -expression! Moreover—as the industrious reader may already have observed—if one takes an 'anything goes' attitude, it is possible to form some pretty wild (and indeed, wildly pretty) expressions. For example, consider the following expression:

$$\lambda x.(x@x)@x$$

Is this a legitimate lambda expression? Suppose we functionally apply this expression to itself (after all, nothing we have said rules out self-application). That is:

$$(\lambda x.(x@x)@x)@(\lambda x.(x@x)@x)$$

If we now apply β -conversion we get:

$$((\lambda x.(x@x)@x)@(\lambda x.(x@x)@x))@(\lambda x.(x@x)@x)$$

But note this is just the functional application followed by an additional occurrence of $\lambda x.(x@x)@x$! Obviously we can go on applying β -conversion as often as we like, producing ever longer expressions as we do so. Now, this is very interesting—but is it the sort of thing we had in mind when we decided to use the lambda calculus?

In this appendix, we shall briefly discuss some of these issues. The main points we wish to make are that there are different versions of the lambda calculus, useful for different purposes, and that both major variants of the lambda calculus can be given model theoretic interpretations in terms of functions and arguments. The reader interested in learning more is encouraged to follow up the references given in the Notes of Chapter 2.

The lambda calculus comes in two main varieties: the *untyped* lambda calculus, and the *typed* lambda calculus. Both can be regarded as programming languages, but they are very different.

The untyped lambda calculus adopts an ‘anything goes’ attitude to functional application. For example, $\lambda x.(x@x)@x$ is a perfectly reasonable expression in untyped lambda calculus, and it is fine to apply it to itself as we did above. The untyped lambda calculus is relevant to the study of natural language (explaining various natural language phenomena seems to demand some form of self-application). Moreover, it is relevant to functional programming (the core of the programming language LISP is essentially the untyped lambda calculus). Finally, it does have a model-theoretic semantics, indeed a very beautiful one. As one might suspect, this semantics interprets abstractions as certain kinds of functions. The primary difficulty is to find suitable collections of functions in which the idea of self application can be captured. (In standard set theory, functions can’t apply to themselves, so constructing function spaces with the structure necessary to model self application is a non-trivial exercise.) There are a variety of solutions to this problem, some of which are very elegant indeed.

There are many kinds of typed lambda calculi. The one we shall discuss is called the *simply typed lambda calculus*.

Natural language semanticists generally make use of some version of the simply typed lambda calculus. The key feature of the simply typed lambda calculus is that it adopts a very restrictive approach to functional application. “If it doesn’t fit, don’t force it”, and typed systems have exacting notions about what fits. Let’s discuss the idea of simple typing in a little more detail.

To build the kinds of representations we have been making use of in simply typed lambda calculus, we would proceed as follows. First we would specify the set of types. There would be infinitely many of these, namely (1) the type e of individuals, (2) the type t of truth values, and (3) for any types τ_1 and τ_2 , the function type $\langle \tau_1, \tau_2 \rangle$. Second, we would specify our logical language. This would contain all the familiar first order symbols, but in addition it would contain an infinite collection of variables of each type (the ordinary first

order variables, which range over individuals, would now be thought of as the set of type e variables) together with the λ operator.

In the typed lambda calculus, every expression receives a *unique* type. The key clauses that ensures this are the definitions of abstraction and functional application. First, if \mathcal{E} is a lambda expression of type τ_2 , and v a variable of type τ_1 then $\lambda v.\mathcal{E}$ is a lambda expression of type $\langle\tau_1, \tau_2\rangle$. In short, it matters which type of variable we abstract over. Abstracting with respect to different types of variables results in abstractions with different types. Second, if \mathcal{E} is a lambda expression of type $\langle\tau_1, \tau_2\rangle$, and \mathcal{E}' is a lambda expression of type τ_1 then we are permitted to functionally apply \mathcal{E} to \mathcal{E}' . If we do this, the application has type τ_2 . In short, we are only allowed to perform functional application when the types \mathcal{E} and \mathcal{E}' fit together correctly, *and only then*. The intuition is that the types tell us what the domain and range of each expression is, and if these don't match, application is not permitted. Note, moreover, that the type of the application is determined by the types of \mathcal{E} and \mathcal{E}' . In effect, we have imposed a very strict type discipline on our formalism. The typed lambda calculus is a programming language, but when we use it we have to abide by very strict guidelines.

This has a number of consequences. For a start, self-application is impossible. (To use \mathcal{E} as a functor, it must have a function type, say $\langle\tau_1, \tau_2\rangle$. But then its arguments must have type τ_1 . So \mathcal{E} can't be one of its own arguments.) Moreover, it is very straightforward to give a semantics to such systems. Given any model M , the denotation D_e of type e expressions are the elements of the model, the permitted denotations D_t of the type t expressions are TRUE and FALSE, and the permitted denotations $D_{\langle\tau_1, \tau_2\rangle}$ of type $\langle\tau_1, \tau_2\rangle$ expressions are functions whose domain is D_{τ_1} and whose range is D_{τ_2} . In short, expressions of the simply typed lambda calculus are interpreted using a straightforward, inductively defined, function hierarchy.

Which particular functions are actually used? Consider $\lambda x.MAN(x)$, where x is an ordinary first order variable. Now, $MAN(x)$ is a formula, something that can be TRUE or FALSE, so this has type t . As was already mentioned, first-order variables are viewed as type e variables, hence it follows that the abstraction $\lambda x.MAN(x)$ has type $\langle e, t\rangle$. That is, it must be interpreted by a function from the set of individuals to the set of truth values. But which one? In fact, it would be interpreted by the function which, when given an individual from the domain of quantification as argument, returns TRUE if that individual is a man, and FALSE otherwise. To put it another way, it is interpreted using the function which exactly characterizes the subset of the model consisting of all men. But this subset is precisely the subset that the standard first-order semantics uses to interpret MAN . In short, the 'functional interpretation' of lambda expressions is set up so that, via the mechanism of such *characteristic functions*, it meshes perfectly with the ordinary first order interpretation.

By building over this base in a fairly straightforward way, the interpretation can be extended to cover all lambda expressions. For example, the expression

$$\lambda P.\exists x(WOMAN(x)\wedge P@x)$$

would be interpreted as a function which when given the type of function that P denotes as argument (and P denotes type $\langle e, t \rangle$ functions, like the characteristic function of MAN) returns a type t value (that is, either `TRUE` or `FALSE`). Admittedly, thinking in terms of functions that take other functions as arguments and return functions as values can get pretty tedious, but the basic idea is straightforward, and for applications in natural language semantics, only a very small part of the function hierarchy tends to be used.

In short, throughout this book we have been talking about the lambda calculus as a mechanism for marking ‘missing information’, and we have exploited the mechanisms of functional application and β -conversion as a way of moving such missing information to where we want it. But in fact there is nothing at all mysterious about this ‘missing information’ metaphor. It is possible to give precise mathematical models of missing information in terms of functions and arguments. An abstraction is interpreted as a function, and the ‘missing information’ is simply the argument we will later supply it with. Indeed, a variety of models are possible, depending on whether one wants to work with typed or untyped versions of the lambda calculus.

Exercise B.0.1 Does our implementation of β -conversion allow self-application?

Exercise B.0.2 What happens when you functionally apply the formula $\lambda x.x@x$ to itself?

Appendix C

Theorem Provers and Model Builders

Automated reasoning has seen an enormous increase of performance of (especially first-order) inference engines and model builders. This appendix is a guide to a number of useful reasoning engines for computational semantics.

Theorem Provers

- **BLIKSEM**
Resolution based theorem prover for first order logic with equality (Hans de Nivelle). Almost as fast as the speed of light.
URL: <http://turing.wins.uva.nl/~mdr/ACLG/Provers/Bliksem/bliksem.html>
- **SPASS**
First-order theorem prover (Christoph Weidenbach et al.)
Winner at the CASC-15 in the divisions FOF and SAT.
URL: <http://spass.mpi-sb.mpg.de/>.
- **FDPLL:**
Or, if you prefer, the “First-Order Davis-Putnam-Loveland-Logeman” theorem prover (Peter Baumgartner). Good at satisfiable problems.
URL: <http://www.uni-koblenz.de/~peter/FDPLL/>
- **OTTER**
The ‘classical’ theorem prover by W. McCune.
URL: <http://www-unix.mcs.anl.gov/AR/otter/>

Model Builders

- **MACE**
Short for “Model and Counter-Examples”. Mace is a Model builder for first-order logic with equality (McCune). Got the first price at CASC-16 in division SAT (working in tandem with OTTER).
URL: <http://www-unix.mcs.anl.gov/AR/mace/>
- **SATCHMO**
Satchmo is a model generator for first-order theories, implemented in Prolog.
URL: <http://www.pms.informatik.uni-muenchen.de/software/>
- **KIMBA**
Kimba is an Higher-Order Model Generator, a deduction system for abductive or circumscriptive reasoning within linguistic applications (Karsten Konrad).
URL: <http://www.ags.uni-sb.de/~konrad/kimba.html>

Systems

- **MATHWEB**
This system for automated theorem proving connects a wide-range of mathematical services by a common, mathematical software bus (Franke and Kohlhase 1999).
URL: <http://www.ags.uni-sb.de/~omega/www/mathweb.html>

Appendix D

Prolog in a Nutshell

What is Prolog?

Prolog is one of the most important programming languages in Computational Linguistics and Artificial Intelligence. Two key features distinguish Prolog from other programming languages: its *declarative* nature, and its extensive use of *unification*. Ideally, in Prolog we simply state the nature of the problem, and let the Prolog unification-driven inference engine search for a solution.

The Basics

There are actually only three basic constructs in Prolog: *facts*, *rules*, and *queries*. A set of facts and rules — that is, a *knowledge base* — is all a Prolog program consists of. Facts are used to state things that are unconditionally true of the domain of interest. For example, we can state that Mia and Jody are women by putting the facts

```
woman(mia).  
woman(jody).
```

in our knowledge base. Rules relate facts by logical implications. We can add to our knowledge base the conditional information that Mia plays air-guitar *if* she listens to music as follows:

```
playsAirGuitar(mia):-  
    listensToMusic(mia).
```

The `:-` should be read as “if”, or “is implied by”. The part on the left hand side of the `:-` is called the *head* of the rule, the part on the right hand side the *body*. (Incidentally, note that we can view a fact as a rule with an empty body. That is, we can think of facts as ‘conditionals that don’t have any antecedent conditions’.) The facts and rules contained in a knowledge base are called *clauses*. The collection of all clauses in a knowledge base that have the same head is called a *predicate*.

Posing queries makes the Prolog inference engine try to deduce a positive answer from the information contained in the knowledge base. There are two basic circumstances under which Prolog can return a positive answer. The first, and simplest, is when the question posed is one of the facts listed in the knowledge base. The second is when Prolog can deduce the positive answer by using the deduction rule called *modus ponens*. That is, if `head :- body` and `body` are both in the knowledge base, then Prolog can deduce that `head` is true.

Let’s consider an example. We can ask Prolog whether Mia is a woman by posing the query

```
?- woman(mia).
```

and Prolog will answer “yes”, since this is a fact in the knowledge base. However, if we ask whether Mia plays air-guitar by posing the query

```
?- playsAirGuitar(mia).
```

its answer is “no”. First, this fact is not recorded in the knowledge base. Second, it cannot infer that Mia plays air-guitar as there is nothing in the knowledge base stating that Mia is listening to music (thus Prolog assumes that this is false) and hence we cannot make use of the only rule we have in our knowledge base. On the other hand, if the knowledge base had contained the additional fact

```
listensToMusic(mia).
```

then Prolog would have responded “yes”, since it could then have deduced by *modus ponens* that Mia plays air-guitar.

Jumping ahead slightly, there are two things that make Prolog so powerful. The first is that it is capable of ‘chaining together’ uses of *modus ponens*. The second is that Prolog has a powerful mechanism (described later) called *unification*, which lets it handle variables. This combination of chained *modus ponens* and sophisticated variable handling enables it to draw far more interesting inferences than our rather trivial examples might suggest.

The comma `,` expresses logical conjunction in Prolog. We can change the rule above by adding another condition to its body:

```
playsAirGuitar(mia):-  
    listensToMusic(mia),  
    happy(mia).
```

Now the rule reads: Mia plays air-guitar if she listens to music, *and* if she is happy. We can also express disjunction in Prolog. Let's change the rule to:

```
playsAirGuitar(mia):-  
    listensToMusic(mia);  
    happy(mia).
```

The ; should be read as *or*, so here we have a rule stating that Mia plays air-guitar if she listens to music, *or* if she is happy.

However there is another (much more commonly used) way of expressing disjunctive conditions: we simply list a number of clauses that have the same head. For example, the following two rules mean exactly the same as the previous rule:

```
playsAirGuitar(mia):-  
    listensToMusic(mia).  
playsAirGuitar(mia):-  
    happy(mia).
```

In fact, disjunctions are almost always expressed using such multiple rules; extensive use of semicolon makes Prolog code pretty hard to read.

Syntax

What exactly are the syntactic entities such as `woman(jody)` and `playsAirGuitar(mia)` that we use to build facts, rules, and queries? They are called *terms*, and there are three kinds of terms in Prolog: atoms, variables, and complex terms. An atom is a sequence of characters starting with a lowercase character. A variable is also a sequence of characters, but it must start with an uppercase character or an underscore. So `mia` and `airGuitar` are atoms, while `X`, `Mayonnaise`, and `_mayonnaise` are variables.

Complex terms are build out of a *functor* and a sequence of *arguments*. The arguments are put in ordinary brackets, separated by commas, and placed after the functor. The functor *must* be an atom. That is, variables *cannot* be used as functors. On the other hand, arguments can be any kind of term. For example

```
hide(X,father(butch))
```

is a complex term. Its functor is `hide`, and it has two arguments: the variable `X`, and the complex term `father(butch)`.

The number of arguments that a complex term has is called its *arity*. For instance, `woman(mia)` is a complex term with arity 1, while `loves(vincent,mia)` is a complex term with arity 2.

Arity is important to Prolog. Prolog would be quite happy for us to use two predicates with the same functor but with a different number of arguments (for example, `love(vincent,mia)`, and `love(vincent,marcellus,mia)`) in the same program. However, if we insisted on doing this, Prolog would treat the two place `love` and the three place `love` as completely different predicates.

When we need to talk about predicates and how we intend to use them (for example, in documentation) it is usual to use a suffix `/` followed by a number to indicate the predicate's arity. For example, if we were talking about the `playsAirGuitar` predicate we would write it as `playsAirGuitar/1` to indicate that it takes one argument. We make use of this convention in the book.

Unification

Variables allow us to make general statements in Prolog. For example, to declare that every woman likes a foot massage, we add the following rules to the database:

```
likeFootMessage(X):-  
    woman(X).
```

It can be read as: `X` likes a foot massage, if `X` is a woman. The nice thing about variables is that they have no fixed values. They can be instantiated with a value by the process of unification. If we pose the query

```
?- likeFootMessage(mia).
```

the variable `X` is *unified* with the atom `mia`, and Prolog tries to prove that `woman(mia)` can be inferred from the knowledge base. Two terms are unifiable if they are the same atoms, or if one of them is a variable, or if they are both complex terms with the same functor name and arity, and all corresponding arguments unify. Unification makes terms identical. It is the heart of the engine that drives Prolog.

Prolog has a built-in operator for unification: the `=`. By querying the goal

```
?- X = butch.
```

the variable `X` gets unified with the atom `butch` and the goal succeeds.

Recursion

Predicates can be defined recursively. Roughly speaking, a predicate is recursively defined if one or more rules in its definition refers to itself. Here's an example:

```
moreExpensive(X,Y):-
    costsaLittleMore(X,Y).

moreExpensive(X,Y):-
    costsaLittleMore(X,Z),
    moreExpensive(Z,Y).
```

The definition of the `moreExpensive/2` predicate is fairly typical of the way recursive predicates are defined in Prolog. Clearly, `moreExpensive/2` is (at least partially) defined in terms of itself, as the more `moreExpensive` functor occurs on both the left and right hand sides of the second rule. Note, however, that there is an 'escape' from this circularity. This is provided by the `costsaLittleMore` predicate, which occurs in both the first and second rules. (Significantly, the right hand side of the first rule makes no mention of `moreExpensive`.)

Let's see how Prolog makes use of such definitions. Suppose we had the following facts in our knowledge base:

```
costsaLittleMore(royaleWithCheese,bigKahunaBurger).

costsaLittleMore(fiveDollarShake,royaleWithCheese).
```

If we pose the query

```
?- moreExpensive(fiveDollarShake,bigKahunaBurger).
```

then Prolog goes to work as follows. First, it tries to make use of the first rule listed concerning `moreExpensive`. This tells it that X is more expensive than Y if X costs a little more than Y , but as the knowledge base doesn't contain the information that a five dollar shake costs a little more than a big kahuna burger, this is no help. So, Prolog tries to make use of the second rule. By unifying X with `fiveDollarShake` and Y with `bigKahunaBurger` it obtains the following goal:

```
?- costsaLittleMore(fiveDollarShake,Z),
    moreExpensive(Z,bigKahunaBurger).
```

Prolog deduces `moreExpensive(fiveDollarShake,bigKahunaBurger)` if it can find a value for Z such that, firstly,

```
?- costsaLittleMore(fiveDollarShake,Z).
```

is deducible, and secondly,

```
?- moreExpensive(Z,bigKahunaBurger).
```

is deducible too. But there is such a value for Z: `royaleWithCheese`. It is immediate that

```
?- costsaLittleMore(fiveDollarShake,royaleWithCheese).
```

will succeed, for this fact is listed in the knowledge base. Deducing

```
?- moreExpensive(royaleWithCheese,bigKahunaBurger).
```

is almost as simple, for the first clause of `moreExpensive/2` reduces this goal to deducing

```
?- costsaLittleMore(royaleWithCheese,bigKahunBurger).
```

and this is a fact listed in the knowledge base.

There is one very important thing you should bear in mind when working with recursion: a recursive predicate should always have at least two clauses: a non-recursive one (the clause that stops the recursion at some point, otherwise Prolog would never be able to finish a proof!), and one that contains the recursion. In our example, the first clause of `more_expensive/2` is the non-recursive clause, and the second clause contains the recursion. (Note that the order of these two clauses in the knowledge base is not important!)

Lists

Lists are recursive data structures in Prolog. The recursive definition of a list runs as follows. First, the empty list is a list. Second, a complex term is a list if it consists of two items, the first of which is a term (often referred to as *first*), and the second of which is a list (often referred to as *rest* or *the rest list*).

Square brackets indicate lists. The empty list is written as `[]`. The list operator `|` separates the first item of a list from the rest list. For example, here is a list with three items:

```
[butch|[pumpkin|[marsellus|[]]]]
```


When working with lists, Prolog always makes use of such recursive first/rest representations, and for some purposes it is important to know this. Mercifully, however, Prolog also offers a more readable form of list representation. The same list can be declared as:

```
[butch,pumpkin,marsellus]
```

Prolog will quite happily accept lists in this more palatable notation as input, and moreover, it does its best to use this notation for its output.

Note that the items in a list need not only be atoms: they can be any Prolog term, including lists. For example

```
[vincent,[honey_bunny,pumpkin],[marsellus,mia]]
```

is a perfectly good list.

Since lists are recursive data structures, most of the predicates that work on lists, are defined using recursive predicates. The simplest example is the `member/2` predicate, which is given in the program library.

Operators

Many newcomers to Prolog find the word *operator* rather misleading, for Prolog's operators don't actually operate on anything, or indeed do anything at all. They are simply a notational device that Prolog offers to represent complex terms in a more readable fashion.

For example, suppose that we want to use `not` as the functor expressing sentence negation, and `and` as the functor expressing sentence conjunction. Then the term representing `Butch boxes and Vincent doesn't dance` would be:

```
and(butch_boxes,not(vincent_dances)).
```

It would be nicer if we could use the more familiar notation in which the conjunction symbol stands between the two sentences it conjoins. That is, we would prefer the following representation:

```
butch_boxes and not vincent_dances
```

The following *operator definitions* let us do precisely this.

```
:- op(30,yfx,and).
:- op(20,fy,not).
```

The `and` is declared as an infix operator by `yfx` with precedence value of 30. The `y` represents an argument (in this case the left argument of `&`) whose precedence is lower or equal of the operator. The `x` represents an argument (the right argument of `and`) whose precedence value must be strictly lower than that of the operator (30, in this case). The `not` is defined as a prefix operator with precedence 20, and an argument that should have a precedence value lower or equal to 20.

More generally, Prolog allows us to define our own infix, prefix and postfix operators. Operator names must be atoms. When we declare a new operator, we also have to state its precedence and those of its arguments (in order for Prolog to disambiguate expressions with more than one operator).

A very important final remark about operators: Prolog is featured with a set of special predefined operators, that *have* a meaning. Among these is the `=` for unification as we saw before and `\+` for negation as we will see later. Other predefined operators we use in the programs of this book are the infix operators `==` and `\==`. The goal `X == Y` succeeds if `X` and `Y` are identical terms (That is, they should have the same structure, even variables should have the same name — unification plays no role here!). On the other hand, `\==` checks whether its arguments are *not* identical.

Arithmetic

Prolog contains some built-in operators for handling integer arithmetic. These include `*`, `/`, `+`, `-` (for multiplication, division, addition, and subtraction, respectively) and `>`, `<` for comparing numbers.

These symbols, however, are just ordinary Prolog operators. That is, they are just a user friendly notation for writing arithmetic expressions: they *don't* carry out the actual arithmetic. For example, posing the query

```
?- X = 1 + 1.
```

unifies the variable `X` with the complex term `1 + 1`, not with `2`, which, for people unused to Prolog's little ways, tends to come as a bit of a surprise.

If we want to carry out the actual arithmetic involved, we have to explicitly force evaluation by making use of the very special inbuilt 'operator' `is/2`. This calls an inbuilt mechanism which carries out the arithmetic evaluation of its second argument, and then unifies the result with its first argument.

```
?- X is 1 + 1.  
X = 2  
yes
```

Negation

There is no explicit negation in Prolog. Something is regarded as false if it cannot be proved. This is the so-called closed world assumption of Prolog.

Prolog does have an inbuilt mechanism for “negation as failure”. That is, we can ask it whether something cannot be proved by using the built-in prefix operator `\+`. The goal `?- \+ man(jules)` succeeds if and only if the goal `?- man(jules)` fails.

Variables in negated goals are not instantiated. Therefore, the following two goals are not fully equivalent.

```
?- likeFootMessage(Who)  
Who = mia  
yes  
  
?- \+ \+ likeFootMessage(Who)  
yes
```

Backtracking and the Cut

When Prolog tries to prove a goal, the structure underlying it attempts can be regarded as a tree: a tree with branches that lead (or do not lead) to possible proofs. The Prolog inference engine essentially searches for a branch that makes the main goal true. Of course it may, and usually does, find itself in a branch that does not lead to a proof (that is, a branch that *fails*). Then Prolog automatically *backtracks* to the previous node in the tree and tries another (the next) branch. (If there is no previous node, this means that the goal is not provable, and Prolog tells us “no”.)

Backtracking can be forced by the user by entering a semicolon after Prolog gives a solution. This allows us to try generating more than one solution to a query. For example (reverting to our original knowledge base) we can demand that Prolog finds all the women as follows:

```
?- woman(Who).  
Who = mia;  
Who = jody;  
no
```

There are two ways to influence Prolog's search strategy: using the *cut* to suppress backtracking, or changing the order of the clauses in the database.

The *cut* (written `!`, it is a built-in Prolog predicate with arity 0) removes certain branches from the proof tree. If a *cut* is put in a clause, and Prolog encounters it during a proof, it removes all the clauses of the same predicate that are listed further down in the knowledge base, and moreover, removes all alternative solutions to conjuncts to the left of the *cut* in its clause.

The order of clauses play a role in Prolog's search strategy, since Prolog works on the database in a sequential way, and subgoals are proved from left to right in the search tree. (Sometimes the order even drastically determines the outcome of a proof.)

Built-in Predicates

There is a set of built-in predicates available in Prolog. We briefly discuss the ones that we use in our programs. It should be noted though, that it depends on the Prolog version that you use, which predicates are built-in.

First we have the ones for controlling output. The predicate `write/1` displays a term to the current output device, and `nl/0` creates a new line.

Then there is whole family of predicates that are used for term manipulation. With these you can break complex terms into pieces or build new ones. The predicate `functor(T,F,A)` succeeds if `F` is the functor of term `T` with arity `A`. And `arg(N,T,A)` is true if `A` is the `N`th argument of `T`. In some case we prefer to make use of the so-called "univ" predicate, weirdly written as `=..`, that transforms a complex term into a list consisting of the functor followed by its arguments. For instance:

```
?- love(pumpkin,honey_bunny) =.. List.
List = [love,pumpkin,honey_bunny]
yes
```

For checking the types of terms: `nonvar(X)` is true if `X` is instantiated, `var(X)` is true if `X` is not instantiated. The predicate `simple/1` succeeds if its argument is either an atom or a variable, and `compound/1` if its argument is a complex term.

With `assert/1` and `retract/1` it is possible to change the knowledge base while executing a goal. The former asserts a clause to the database, the latter removes it. Many versions of Prolog require a `dynamic` declaration of those predicates that are modified by other predicates, as we do with the counter for the skolem index in the library file.

Finally, there is some predefined stuff that gives controll over backtracking. The very rude `fail/0` predicate causes Prolog to backtrack (this is used to generate more solutions). With

`bagof/3` it is possible to generate all solutions of a specific goal, and put the instantiations of a certain variable (or complex term with variables) in a list. E.g., the goal

```
?- bagof(Who, woman(Who), Answer).
Answer = [mia, jody].
yes
```

tries to satisfies the goal in the second argument of `bagof`, and for each solution it puts the value of its first argument (`Who`) in its third argument, the list `Answer`. The predicate `setof/3` functions similarly, but it removes duplicate answers.

Difference Lists

A difference list `List1-List2` is a way of using two lists (namely, `List1` and `List2`) to represent one single list (namely, the *difference* of the two lists).

For example, the three element list `[pumpkin, butch, jimmy]` can be represented as the following difference list:

```
[pumpkin, butch, jimmy, lance]-[lance]
```

or as the following one:

```
[pumpkin, butch, jimmy]-[]
```

Indeed, there are infinitely many difference list representations of `[pumpkin, butch, jimmy]`. In all of them, the first list in the difference list representation consists of the list we are interested in (namely `[pumpkin, butch, jimmy]`) followed by some suffix, while the second list is that suffix. So a difference list is simply a pair of list structures, the second of which is a suffix from the other. We follow the usual convention of using the built-in operator `-` to group the two list together. (However, note that we don't have to do this: any other operator would do.)

Let's go a step further. Suppose we take the suffix to be a variable. Then any other difference-list encoding of `[pumpkin, butch, jimmy]` is an instance of the following *most general difference list*.

```
[pumpkin, butch, jimmy|X]-X
```

In practice, we will use the term *difference list* to mean the most general difference list. The empty list is represented by the difference list $X-X$.

Why on earth should anyone want to represent lists as difference lists? There is a simple answer: efficiency. In particular, when we use the difference list representation, Prolog can perform concatenation of lists much faster. For example, the usual `append/3` for normal lists is a recursively defined predicate that can be inefficient for large lists. (This is because Prolog must recursively work its way down the entire collection of first/rest pairs.) On the other hand, when we make use of difference lists `append/3` can be defined as follows.

```
append(X-Y, Y-Z, X-Z).
```

Consider how this works. Suppose we want to append `[pumpkin, butch]` to `[jody, lance]`. Then we make the following query:

```
?- append([pumpkin, butch|E]-E, [jody, lance|F]-F, A-B).
```

This causes `A` to unify with `[pumpkin, butch|E]`, `E` with `[jody, lance|F]`, and `B` with `F`. As a result, `A-B` unifies with `[pumpkin, butch, jody, lance|B]-B`, which is a difference list of the four items `pumpkin`, `butch`, `jody`, and `lance`.

The reader who works through this example will see that what makes difference list representations so efficient is that the suffix variable gives us direct access to the end of the list. In the conventional first/rest representation we have to work our way recursively down towards the end of the list. Difference list representations avoid this overhead. We have to pay a price for this gain in efficiency (difference list representation is less transparent) but in many circumstances this is a price worth paying.

Definite Clause Grammars

Definite clause grammars (DCGs) are the in built Prolog mechanism for defining grammars. Actually, they are really a syntactically sugared way of working with certain difference lists. With DCGs you can kill two birds with one stone: if you define the grammar rules you'll get the parser for free!

Here is a DCG for a very small fragment of english grammar. We have defined syntactic categories `s`, `np`, `vp`, `det`, `noun`, and `tv`, standing for sentence, noun phrase, verb phrase, determiner, common noun, and transitive verb. These are also called the *non-terminal* symbols. Every rule in the DCG has a non-terminal symbol on its left hand side.

```
s --> np, vp.
```

```

np --> det, noun.
np --> [mia].
det --> [a].
noun --> [man].
noun --> [five,dollar,shake].
vp --> tv, np.
vp --> [drinks].
tv --> [loves].

```

On the right hand side of the rules, you'll either find one or more non-terminal symbols, or a terminal symbol. The terminal symbols are the lexical entries, the actual words of the language of our interest.

Now, what do these rules mean? They are very intuitive indeed. The rule `s --> np, vp` says that a syntactic category called `s`, consists of an `np`, followed by a `vp`. Similarly, according to this DCG, the category `noun` is either the string `man`, or the sequence of strings `five dollar shake` (represented in lists). And so on.

This DCG covers sentences like `Mia loves a five dollar shake`, `A man drinks`, but not: `A woman loves` or `Mia drinks a one dollar shake`.

To parse sentences, we can pose a query like:

```
?- s([mia,loves,a,five,dollar,shake], []).
```

This goal is satisfied if the sequence of words in the first argument belongs to `s`. There is only one rule for `s` in our DCG. It says that an `s` can be replaced by an `np` followed by a `vp`. That is, we have to take some items of the input list that form a noun phrase, in such a way that the rest of the items on the list form a verb phrase. Since there is the rule

```
np --> [mia].
```

in the knowledge base we now have to proof whether

```
?- vp([loves,a,five,dollar,shake], []).
```

which is provable indeed (we leave this is an exercise to the reader). And this is basically how parsing with DCGs takes place on the surface level. However, Prolog doesn't use our DCG rules directly for the purpose of parsing. The DCG is translated internally into the following clauses:

```

s(A,B) :- np(A,C), vp(C,B).
np(A,B) :- det(A,C), noun(C,B).
np([mia|B],B).
det([a|B],B).
noun([man|B],B).
noun([five,dollar,shake|B],B).
vp(A,B) :- tv(A,C), np(C,B).
vp([drinks|B],B).
tv([loves|B],B).

```

The alert reader will notice that these clauses look surprisingly familiar. In fact, they are the direct encodings of difference lists (in the facts, for example `noun([man|B],B)`) and the appending of two difference lists (in the rules, for instance

```

s(A,B) :- np(A,C), vp(C,B).

```

can be read as: `s` is a difference list `A-B` if we can prove that is the result of appending difference list `A-C` to `C-B`).

Since, as we just noticed, DCG rules are normal Prolog clauses, it is perfectly allowed to add arguments to the rules. Some useful stuff we can add to our grammar is information on agreement. Suppose we want to include noun phrases like *all boxers* in our grammar by adding entries for *all* and *boxers*:

```

det --> [all].
noun --> [boxers].

```

Be careful though! Adding these clauses make it possible to parse *Mia loves all boxers*, but also non-grammatical *All man drinks* or *A boxers loves all woman*. Clearly, our grammar lacks information about agreement. However, this information can be added very easily to the rules. Let's do it for the determiners and nouns first (and why not add some extra entries at the same time):

```

det(plural) --> [all].
det(singular) --> [a].
noun(singular) --> [boxer].
noun(plural) --> [boxers].
noun(singular) --> [man].
noun(plural) --> [men].
noun(singular) --> [five,dollar,shake].
noun(plural) --> [five,dollar,shakes].

```


Since `det` and `noun` have an extra argument now, all the rules that have these symbols at their right hand side (this is the `np --> det, noun.` rule in our current grammar) should get an extra argument as well. As can be seen from the entries above, the value of this argument is either `singular` or `plural`. Since we want the determiner to have the same agreement as the noun when they are combined to a noun phrase, we could write the following code:

```
np --> det(singular), noun(singular).
np --> det(plural), noun(plural).
```

But there is much more elegant way of encoding this. We can bring in a variable that, during parsing, gets instantiated with the agreement value of the determiner and noun, and collapse the above rules into one:

```
np --> det(Agr), noun(Agr).
```

Nevertheless, from the examples given earlier, we also want to add agreement features to noun phrases and verb phrases. We won't do explain this step by step — the principle should be clear now — but list the entire rewritten grammar including agreement, and extended with some new entries, below.

```
s --> np(Agr), vp(Agr).
np(Agr) --> det(Agr), noun(Agr).
np(singular) --> [mia].
det(plural) --> [all].
det(singular) --> [a].
noun(singular) --> [boxer].
noun(plural) --> [boxers].
noun(singular) --> [man].
noun(plural) --> [men].
noun(singular) --> [five,dollar,shake].
noun(plural) --> [five,dollar,shakes].
vp(Agr) --> tv(Agr), np(_).
vp(singular) --> [drinks].
vp(plural) --> [drink].
tv(singular) --> [loves].
tv(plural) --> [love].
```

Now look at our grammar. From an aesthetic point of view, there certainly is some space for improvement! Everything is mixed up: the rules and the lexical entries together form

one chaotic DCG. Consider what these grammar would look like if we extend its coverage to a more serious fragment of english!

One way to bring in some organisation in the grammar, is to make a physical distinction between the lexicon and the grammar rules. The lexicon might be designed as follows:

```
lexicon(np,singular,mia).
lexicon(det,plural,all).
lexicon(det,singular,a).
lexicon(noun,singular,boxer).
lexicon(noun,plural,boxers).
lexicon(noun,singular,man).
lexicon(noun,plural,men).
lexicon(vp,singular,drinks).
lexicon(vp,plural,drink).
lexicon(tv,singular,loves).
lexicon(tv,plural,love).
```

That is, normal Prolog facts of the form `lexicon/3`, with the first argument stating the syntactic category, the second argument the agreement value, and the third the word itself. The only thing left to do is to make a connection between this lexicon and the grammar rules. DCGs have a neat way to do this, normal Prolog goals can be included in the rules included in curly brackets. This is the result:

```
s --> np(Agr), vp(Agr).
np(Agr) --> [X], {lexicon(np,Agr,X)}.
np(Agr) --> det(Agr), noun(Agr).
det(Agr) --> [X], {lexicon(det,Agr,X)}.
vp(Agr) --> tv(Agr), np(_).
vp(Agr) --> [X], {lexicon(vp,Agr,X)}.
tv(Agr) --> [X], {lexicon(tv,Agr,X)}.
```

Notes

In this appendix we summarized the basic concepts of Prolog. We hope it will be a handy reference, however it is *not* intended as a substitute for good introduction to Prolog. The reader who wants to learn about computational semantics, but who knows no Prolog, is strongly advised to put this book aside for a while and study one of the many excellent Prolog texts currently available. We particularly recommend the following ones. For a succinct, no-frills overview, try Clocksin and Mellish (Clocksin and Mellish 1987). For a leisurely, in-depth introduction to programming in Prolog, try Bratko (Bratko 1990). For

a more theoretically oriented introduction, try Sterling and Shapiro (Sterling and Shapiro 1986). Finally, for an introduction specially geared towards computational linguistics, try Pereira and Shieber (Pereira and Shieber 1987).

Appendix E

Listing of Programs

This appendix includes the full program listings that are developed in this book. Most predicates are decorated with a short documentation, using the following notational conventions for its required argument instantiations:

`:Arg` Argument `Arg` should be instantiated to a term denoting a goal

`+Arg` Argument `Arg` should be instantiated

`-Arg` Argument `Arg` should not be instantiated

`?Arg` Argument `Arg` may or may not be instantiated

Since Prolog's birth in the beginning of the seventies, a number of Prolog dialects emerged, and not all agree on a syntax or the in-built predicates. The programs in this book follow the conventions of "Standard Prolog", the ISO international standard on Prolog (Deransart, Ed-Dbali, and Cervoni 1996), as close as possible. The following predicates are assumed to be built-in in your version of Prolog (such as for example Quintus or Sicstus Prolog):

All solutions

`bagof/3`

`findall/3`

Arithmetic comparison

`>` (arithmetic greater than)

< (arithmetic less than)

Arithmetic evaluation

is/2 (evaluate expression)

Atomic term processing

atom_chars/2 (conversion of atoms to character codes and vice versa)

Character input

get0/1

Clause creation and destruction

asserta/1 (clause creation)

retract/1 (clause destruction)

File consultation

[File|Files] (consult list of files)

List operation

length/2] (determine length of a list)

Logic and control

,/2 (conjunction)

;/2 (disjunction)

!/0 (cut)

fail/0

true/0

\+ (not provable)

Operator definition

op/3 (extending operator table)

Term comparison

==/2 (term identical)

Term creation and decomposition

arg/3

functor/3

=../2 (the “univ”)

Term unification

= (unify)

Type testing

atom/1

atomic/1

compound/1

nonvar/1

var/1

Term output

write/1

nl/0

The following is a practical overview of the programs in this appendix. We start with the library files consulted by most of the other programs, and then give a chapter by chapter break-down description of the files.

All the files at one glance

File Name	Chapter	Page
comsemPredicates.pl	Chapter 1-12	p. 194
comsemOperators.pl	Chapter 1-12	p. 193
readLine.pl	Chapter 2-3, 6, 8-12	p. 200
englishLexicon.pl	Chapter 2-3, 6, 8-12	p. 202
englishGrammar.pl	Chapter 2-3, 6, 8-12	p. 207
modelChecker.pl	Chapter 1, 6	p. 209
modelChecker2.pl	Chapter 1, 6	p. 211
exampleModels.pl	Chapter 1, 6, 7	p. 214
experiment1.pl	Chapter 2	p. 216
experiment2.pl	Chapter 2	p. 218
mainLambda.pl	Chapter 2	p. 220
semMacrosLambda.pl	Chapter 2	p. 222
betaConversion.pl	Chapter 2-3, 8	p. 221
mainMontague.pl	Chapter 3	p. 223
englishGrammarMontague.pl	Chapter 3	p. 224
mainCooperStorage.pl	Chapter 3	p. 226
mainKellerStorage.pl	Chapter 3	p. 228
semMacrosStorage.pl	Chapter 3	p. 230
pluggingAlgorithm.pl	Chapter 3, 8	p. 231
mainPLU.pl	Chapter 3	p. 233
semMacrosPLU.pl	Chapter 3	p. 234
mergeUSR.pl	Chapter 3, 8, 12	p. 236
propTabl.pl	Chapter 4	p. 237
freeVarTabl.pl	Chapter 5	p. 240
callTheoremProver.pl	Chapter 5	p. 245
callModelBuilder.pl	Chapter 5	p. 246
fol2otter.pl	Chapter 5	p. 247
nlQuestions.pl	Chapter 6	p. 249
nlArgumentation.pl	Chapter 6	p. 251
SemOntology.pl	Chapter 6	p. 253


```

/*****
    name: comsemOperators.pl
    version: May 25, 1999
    description: Operator definitions
    authors: Patrick Blackburn & Johan Bos
*****/

:- module(comsemOperators, []).

/*=====
   Operator Definitions
   =====*/

:- op(950,yfx,@).      % application
:- op(900,yfx,'<>').  % bin impl
:- op(900,yfx,>).      % implication
:- op(850,yfx,v).      % disjunction
:- op(800,yfx,&).      % conjunction
:- op(750, fy,~).     % negation

```

```

/*****

    name: comsemPredicates.pl
    version: November 8, 1997
    description: Set of Prolog predicates
    authors: Patrick Blackburn & Johan Bos

*****/

:- module(comsemPredicates,
    [member/2,
     select/3,
     append/3,
     simpleTerms/1,
     compose/3,
     unify/2,
     removeFirst/3,
     substitute/4,
     variablesInTerm/2,
     newFunctionCounter/1,
     printReadings/1,
     printRepresentation/1]).

/*=====

    List Manipulation
    -----

member(?Element,?List)
    Element occurs in List.

select(?X,+OldList,?NewList).
    X is removed from OldList, result in NewList.

append(?List1,?List2,?List3)
    List3 is the concatenation of List1 and List2.

allExactMember(?Elements,?List)
    Elements all occur in List (no unification, exact match).

removeAll(?Item,?List,?Newlist)
    Newlist is the result of removing all occurrences of Item from List.

removeFirst(?Item,?List,?Newlist)
    Newlist is the result of removing the first occurrence of Item from
    List. Fails when Item is not member of List.

=====*/

member(X,[X|_]).

```

```

member(X,[_|Tail]):-
    member(X, Tail).

select(X,[X|L],L).
select(X,[Y|L1],[Y|L2]):-
    select(X,L1,L2).

append([],List,List).
append([X|Tail1],List,[X|Tail2]):-
    append(Tail1,List,Tail2).

allExactMember([],_).
allExactMember([X|R],L):-
    memberOfList(Y,L),
    X==Y,
    allExactMember(R,L).

removeAll(_,[],[]).
removeAll(X,[X|Tail],Newtail):-
    removeAll(X,Tail,Newtail).
removeAll(X,[Head|Tail],[Head|Newtail]):-
    X \== Head,
    removeAll(X,Tail,Newtail).

removeFirst(X,[X|Tail],Tail) :- !.
removeFirst(X,[Head|Tail],[Head|NewTail]):-
    removeFirst(X,Tail,NewTail).

/*=====

    Term Manipulation
    -----

simpleTerms(?List)
    List is a list of elements that are currently uninstantiated or
    instantiated to an atom or number. Uses built-in Quintus/Sicstus
    predicate simple/1.

compose(?Term,+Symbol,+ArgList)
compose(+Term,?Symbol,?ArgList)
    Composes a complex Term with functor Symbol and arguments ArgList.
    Uses the Prolog built-in =.. predicate.

variablesInTerm(+Term,?InList-?OutList)
    Adds all occurrences of variables in Term (arbitrarily deeply
    nested to the difference list InList-OutList.

=====*/

simpleTerms([]).

```

```
simpleTerms([X|Rest]):-
    simple(X), simpleTerms(Rest).
```

```
compose(Term,Symbol,ArgList):-
    Term =.. [Symbol|ArgList].
```

```
variablesInTerm(Term,Var1-Var2):-
    compose(Term,_,Args),
    countVar(Args,Var1-Var2).
```

```
countVar([],Var-Var).
countVar([X|Rest],Var1-Var2):-
    var(X),!,
    countVar(Rest,[X|Var1]-Var2).
countVar([X|Rest],Var1-Var3):-
    variablesInTerm(X,Var1-Var2),
    countVar(Rest,Var2-Var3).
```

```
/*=====
```

```
Unification Predicates
```

```
-----
```

```
unify(Term1,Term2)
    Unify Term1 with Term2 including occurs check. Adapted from
    "The Art of Prolog" by Sterling & Shapiro, MIT Press 1986, page 152.
```

```
notOccursIn(X,Term)
    Succeeds if variable X does not occur in Term.
```

```
notOccursInComplexTerm(N,X,Term)
    Succeeds if variable X does not occur in complex Term with arity N
```

```
termUnify(Term1,Term2)
    Unify the complex terms Term1 and Term2.
```

```
=====*/
```

```
unify(X,Y):-
    var(X), var(Y), X=Y.
unify(X,Y):-
    var(X), nonvar(Y), notOccursIn(X,Y), X=Y.
unify(X,Y):-
    var(Y), nonvar(X), notOccursIn(Y,X), X=Y.
unify(X,Y):-
    nonvar(X), nonvar(Y), atomic(X), atomic(Y), X=Y.
unify(X,Y):-
    nonvar(X), nonvar(Y), compound(X), compound(Y), termUnify(X,Y).
```

```

notOccursIn(X,Term):-
    var(Term), X \== Term.
notOccursIn(_,Term):-
    nonvar(Term), atomic(Term).
notOccursIn(X,Term):-
    nonvar(Term), compound(Term),
    functor(Term,_,Arity), notOccursInComplexTerm(Arity,X,Term).

notOccursInComplexTerm(N,X,Y):-
    N > 0, arg(N,Y,Arg), notOccursIn(X,Arg),
    M is N - 1, notOccursInComplexTerm(M,X,Y).
notOccursInComplexTerm(0,_,_).

termUnify(X,Y):-
    functor(X,Func,Ari), functor(Y,Func,Ari),
    unifyArgs(Ari,X,Y).

unifyArgs(N,X,Y):-
    N > 0, M is N - 1,
    arg(N,X,ArgX), arg(N,Y,ArgY),
    unify(ArgX,ArgY), unifyArgs(M,X,Y).
unifyArgs(0,_,_).

/*=====

    Substitution Predicates
    -----

substitute(?Term,?Variable,+Exp,-Result)
    Result is the result of substituting occurrences of Term for each
    free occurrence of Variable in Exp.

=====*/

substitute(Term,Var,Exp,Result):-
    Exp==Var, !, Result=Term.
substitute(_Term,_Var,Exp,Result):-
    \+ compound(Exp), !, Result=Exp.
substitute(Term,Var,Formula,Result):-
    compose(Formula,Func,[Exp,F]),
    member(Func,[lambda forall exists]), !,
    (
        Exp==Var, !,
        Result=Formula
    );
    substitute(Term,Var,F,R),
    compose(Result,Func,[Exp,R])
).
substitute(Term,Var,Formula,Result):-
    compose(Formula,Func,ArgList),

```

```

    substituteList(Term,Var,ArgList,ResultList),
    compose(Result,Funcor,ResultList).

substituteList(_Term,_Var,[],[]).
substituteList(Term,Var,[Exp|Others],[Result|ResultOthers]):-
    substitute(Term,Var,Exp,Result),
    substituteList(Term,Var,Others,ResultOthers).

/*=====

    Skolem Function Counter
    -----

functionCounter(?N)
    N is the current Skolem function index. Declared as dynamic,
    and set to value 1.

newFunctionCounter(?N)
    Unifies N with the current Skolem function index, and increases
    value of the counter.

=====*/

:- dynamic(functionCounter/1).

functionCounter(1).

newFunctionCounter(N):-
    functionCounter(N), M is N+1,
    retract(functionCounter(N)),
    asserta(functionCounter(M)).

/*=====

    Pretty Print Predicates
    -----

=====*/

printRepresentation(Rep):-
    nl, \+ \+ (numbervars(Rep,0,_), write(Rep)), nl.

printReadings(Readings):-
    nl, write('Readings: '), nl, printReading(Readings,0).

printReading([],N):-
    nl, (N=0, write('no readings'); true), nl.
printReading([Reading|OtherReadings],M):-
    N is M + 1, write(N), tab(1),
    \+ \+ (numbervars(Reading,0,_), write(Reading)), nl,

```

```
printReading(OtherReadings,N).
```

```

/*****

    name: readLine.pl
    version: March 31, 1998
    description: Converting input line to list of atoms, suitable for
                 DCG input.
    authors: Patrick Blackburn & Johan Bos

*****/

:- module(readLine,[readLine/1]).

/*=====

    Read Predicates
    -----

readLine(-WordList)
    Outputs a prompt, reads a sequence of characters from the standard
    input and converts this to WordList, a list of strings. Punctuation
    is stripped.

readWords(-WordList)
    Reads in a sequence of characters, until a return is registered,
    and converts this to WordList a list of strings.

readWord(+Char,-Chars,?State)
    Read a word coded as Chars (a list of ascii values), starting
    with with ascii value Char, and determine the State of input
    ('ended' = end of line, 'notended' = not end of line).
    Blanks and full stops split words, a return ends a line.

checkWords(+OldWordList,-NewWordList)
    Check if all words are unquoted atoms, if not convert them
    into atoms.

convertWord(+OldWord,-NewWord)
    OldWord and NewWord are words represented as lists of ascii values.
    Converts upper into lower case characters, and eliminates
    non-alphabetic characters.

=====*/

readLine(WordList):-
    nl, write('> '),
    readWords(WordList),
    checkWords(WordList,WordList).

readWords([Word|Rest]):-
    get0(Char),

```



```
    readWord(Char, Chars, State),
    atom_chars(Word, Chars),
    readRest(Rest, State).

readRest([], ended).
readRest(Rest, notended):-
    readWords(Rest).

readWord(32, [], notended):-!.      %%% blank
readWord(46, [], notended):-!.      %%% full stop
readWord(10, [], ended):-!.         %%% return
readWord(Code, [Code|Rest], State):-
    get0(Char),
    readWord(Char, Rest, State).

checkWords([], []):-!.
checkWords([''|Rest1], Rest2):-
    checkWords(Rest1, Rest2).
checkWords([Atom|Rest1], [Atom2|Rest2]):-
    atom_chars(Atom, Word1),
    convertWord(Word1, Word2),
    atom_chars(Atom2, Word2),
    checkWords(Rest1, Rest2).

convertWord([], []):-!.
convertWord([Capital|Rest1], [Small|Rest2]):-
    Capital > 64, Capital < 91, !,
    Small is Capital + 32,
    convertWord(Rest1, Rest2).
convertWord([Weird|Rest1], Rest2):-
    (Weird < 97; Weird > 122), !,
    convertWord(Rest1, Rest2).
convertWord([Char|Rest1], [Char|Rest2]):-
    convertWord(Rest1, Rest2).
```

```

/*****

```

```

    name: englishLexicon.pl
    version: November 12, 1997; March 9, 1999.
    description: Lexical entries for a small coverage of English
    authors: Patrick Blackburn & Johan Bos

```

This file contains the lexical entries for a small fragment of English. Entries have the form `lexicon(Cat,Sym,Phrase,Misc)`, where `Cat` is the syntactic category, `Sym` the predicate symbol introduced by the phrase, `Phrase` a list of the words that form the phrase, and `Misc` miscellaneous information depending on the the type of entry.

```

*****/

```

```

/*=====
   Determiners: lexicon(det,_,Words,Type)
=====*/

```

```

lexicon(det,_, [every], uni).
lexicon(det,_, [a], indef).
lexicon(det,_, [the], def).
lexicon(det,_, [one], card(1)).
lexicon(det,_, [another], alt).
lexicon(det,_, [his], poss(male)).
lexicon(det,_, [her], poss(female)).
lexicon(det,_, [its], poss(nonhuman)).

```

```

/*=====
   Nouns: lexicon(noun,Symbol,Words, {[], [Hypernym], Hypernym})
=====*/

```

```

lexicon(noun,abstraction, [abstraction], [top]).
lexicon(noun,act, [act], [top]).
lexicon(noun,animal, [animal], [organism]).
lexicon(noun,artifact, [artifact], [object]).
lexicon(noun,beverage, [beverage], [food]).
lexicon(noun,building, [building], [artifact]).
lexicon(noun,container, [container], [instrumentality]).
lexicon(noun,cup, [cup], [container]).
lexicon(noun,device, [device], [instrumentality]).
lexicon(noun,edible, [edible,food], [food]).
lexicon(noun,bkburger, [big,kahuna,burger], [edible]).
lexicon(noun,boxer, [boxer], human).
lexicon(noun,boss, [boss], human).
lexicon(noun,car, [car], [vehicle]).
lexicon(noun,chainsaw, [chainsaw], [device]).
lexicon(noun,criminal, [criminal], human).
lexicon(noun,customer, [customer], human).
lexicon(noun,drug, [drug], [artifact]).

```

```

lexicon(noun,entity,[entity],[top]).
lexicon(noun,episode,[episode],abstraction).
lexicon(noun,female,[female],[human]).
lexicon(noun,fdshake,[five,dollar,shake],[beverage]).
lexicon(noun,food,[food],[object]).
lexicon(noun,footmassage,[foot,massage],[act]).
lexicon(noun,gimp,[gimp],human).
lexicon(noun,glass,[glass],[container]).
lexicon(noun,gun,[gun],[weaponry]).
lexicon(noun,hammer,[hammer],[device]).
lexicon(noun,hashbar,[hash,bar],[building]).
lexicon(noun,human,[human],[organism]).
lexicon(noun,husband,[husband],male).
lexicon(noun,instrumentality,[instrumentality],artifact).
lexicon(noun,joke,[joke],abstraction).
lexicon(noun,man,[man],male).
lexicon(noun,male,[male],[human]).
lexicon(noun,medium,[medium],[instrumentality]).
lexicon(noun,needle,[needle],[device]).
lexicon(noun,object,[object],[entity]).
lexicon(noun,organism,[organism],[entity]).
lexicon(noun,owner,[owner],human).
lexicon(noun,piercing,[piercing],[act]).
lexicon(noun,plant,[plant],[organism]).
lexicon(noun,qpwc,[quarter,pounder,with,cheese],[edible]).
lexicon(noun,radio,[radio],[medium]).
lexicon(noun,restaurant,[restaurant],[building]).
lexicon(noun,robber,[robber],human).
lexicon(noun,suitcase,[suitcase],[container]).
lexicon(noun,shotgun,[shotgun],[weaponry]).
lexicon(noun,sword,[sword],[weaponry]).
lexicon(noun,vehicle,[vehicle],[instrumentality]).
lexicon(noun,weaponry,[weaponry],[instrumentality]).
lexicon(noun,woman,[woman],female).

```

```

/*=====
   Proper Names: lexicon(pn,Symbol,Words,{male,female})
=====*/

```

```

lexicon(pn,butch,[butch],male).
lexicon(pn,honey_bunny,[honey,bunny],male).
lexicon(pn,jimmy,[jimmy],male).
lexicon(pn,jody,[jody],female).
lexicon(pn,jules,[jules],male).
lexicon(pn,lance,[lance],male).
lexicon(pn,marsellus,[marsellus],male).
lexicon(pn,marsellus,[marsellus,wallace],male).
lexicon(pn,marvin,[marvin],male).
lexicon(pn,mia,[mia],female).
lexicon(pn,mia,[mia,wallace],female).

```

```
lexicon(pn,pumpkin,[pumpkin],male).
lexicon(pn,thewolf,[the,wolf],male).
lexicon(pn,vincent,[vincent],male).
lexicon(pn,vincent,[vincent,vega],male).
```

```
/*=====
   Intransitive Verbs: lexicon(iv,Symbol,Words,{fin,inf})
=====*/
```

```
lexicon(iv,collapse,[collapses],fin).
lexicon(iv,collapse,[collapse],inf).
lexicon(iv,dance,[dances],fin).
lexicon(iv,dance,[dance],inf).
lexicon(iv,die,[dies],fin).
lexicon(iv,die,[die],inf).
lexicon(iv,growl,[growls],fin).
lexicon(iv,growl,[growl],inf).
lexicon(iv,okay,[is,okay],fin).
lexicon(iv,outoftown,[is,out,of,town],fin).
lexicon(iv,married,[is,married],fin).
lexicon(iv,playairguitar,[plays,air,guitar],fin).
lexicon(iv,playairguitar,[play,air,guitar],inf).
lexicon(iv,smoke,[smokes],fin).
lexicon(iv,smoke,[smoke],inf).
lexicon(iv,snort,[snorts],fin).
lexicon(iv,snort,[snort],inf).
lexicon(iv,shriek,[shrieks],fin).
lexicon(iv,shriek,[shriek],inf).
lexicon(iv,walk,[walks],fin).
lexicon(iv,walk,[walk],inf).
```

```
/*=====
   Transitive Verbs: lexicon(tv,Symbol,Words,{fin,inf})
=====*/
```

```
lexicon(tv,clean,[cleans],fin).
lexicon(tv,clean,[clean],inf).
lexicon(tv,drink,[drinks],fin).
lexicon(tv,drink,[drink],inf).
lexicon(tv,date,[dates],fin).
lexicon(tv,date,[date],inf).
lexicon(tv,discard,[discards],fin).
lexicon(tv,discard,[discard],inf).
lexicon(tv,eat,[eats],fin).
lexicon(tv,eat,[eat],inf).
lexicon(tv,enjoy,[enjoys],fin).
lexicon(tv,enjoy,[enjoy],inf).
lexicon(tv,hate,[hates],fin).
lexicon(tv,hate,[hate],inf).
lexicon(tv,have,[has],fin).
```

```

lexicon(tv,have,[have],inf).
lexicon(tv,donewith,[is,done,with],fin).
lexicon(tv,kill,[kills],fin).
lexicon(tv,kill,[kill],inf).
lexicon(tv,know,[knows],fin).
lexicon(tv,know,[know],inf).
lexicon(tv,like,[likes],fin).
lexicon(tv,like,[like],inf).
lexicon(tv,love,[loves],fin).
lexicon(tv,love,[love],inf).
lexicon(tv,pickup,[picks,up],fin).
lexicon(tv,pickup,[pick,up],inf).
lexicon(tv,shoot,[shot],fin).
lexicon(tv,shoot,[shoot],inf).
lexicon(tv,tell,[told],fin).
lexicon(tv,tell,[tell],inf).
lexicon(tv,worksfor,[works,for],fin).
lexicon(tv,worksfor,[work,for],inf).

/*=====
   Copula
=====*/

lexicon(cop,'=',[is],fin).

/*=====
   Prepositions: lexicon(pre,Symbol,Words,_)
=====*/

lexicon(pre,in,[in],_).
lexicon(pre,of,[of],_).
lexicon(pre,with,[with],_).

/*=====
   Pronouns: lexicon(pro,Sym,Words,{refl,nonrefl})
=====*/

lexicon(pro,male,[he],nonrefl).
lexicon(pro,female,[she],nonrefl).
lexicon(pro,nonhuman,[it],nonrefl).
lexicon(pro,male,[him],nonrefl).
lexicon(pro,female,[her],nonrefl).
lexicon(pro,male,[himself],refl).
lexicon(pro,female,[herself],refl).
lexicon(pro,nonhuman,[itself],refl).

/*=====
   Relative Pronouns: lexicon(relpro,_,Words,_)
=====*/

```

```
lexicon(relpro,_,[who],_).
lexicon(relpro,_,[that],_).

/*=====
   Coordinations: lexicon(coord,_,Words,{conj,disj})
=====*/

lexicon(coord,_,[and],conj).
lexicon(coord,_,[or],disj).

/*=====
   Discontinious Coordinations: lexicon(dcoord,W1,W2,{conj,cond,disj})
=====*/

lexicon(dcoord,[if],[then],cond).
lexicon(dcoord,[if],[],cond).
lexicon(dcoord,[either],[or],disj).
lexicon(dcoord,[],[or],disj).
lexicon(dcoord,[],[and],conj).
lexicon(dcoord,[],[],conj).

/*=====
   Modifiers: lexicon(mod,_,Words,Type)
=====*/

lexicon(mod,_,[does,not],neg).
lexicon(mod,_,[did,not],neg).
```

```

/*****
    name: englishGrammar.pl
    version: November 12, 1997
    description: Grammar rules for a small coverage of English
    authors: Patrick Blackburn & Johan Bos
*****/

/*=====
    Grammar Rules
=====*/

d(S) --> s(S).
d((C@S1)@S2)--> dcoord(D,C), s(S1), D, d(S2).

s(NP@VP)--> np2(NP), vp2(VP).

np2(NP)--> np1(NP).
np2((C@NP1)@NP2)--> np1(NP1), coord(C), np1(NP2).

np1(Det@Noun)--> det(Det), n2(Noun).
np1(NP)--> pn(NP).
np1(NP)--> pro(NP).

n2(N)--> n1(N).
n2((C@N1)@N2)--> n1(N1), coord(C), n1(N2).

n1(N)--> noun(N).
n1(PP@N)--> noun(N), pp(PP).
n1(RC@N)--> noun(N), rc(RC).

vp2(VP)--> vp1(VP).
vp2((C@VP1)@VP2)--> vp1(VP1), coord(C), vp1(VP2).

vp1(Mod@VP)--> mod(Mod), v2(fin,VP).
vp1(VP)--> v2(fin,VP).

v2(fin,Cop@NP)--> cop(Cop), np2(NP).
v2(fin,Neg@(Cop@NP))--> cop(Cop), neg(Neg), np2(NP).

v2(I,V)--> v1(I,V).
v2(I,(C@V1)@V2)--> v1(I,V1), coord(C), v1(I,V2).

v1(I,V)--> iv(I,V).
v1(I,TV@NP)--> tv(I,TV), np2(NP).

pp(Prep@NP)--> prep(Prep), np2(NP).

rc(RP@VP)--> relpro(RP), vp2(VP).

```

iv(I,IV)--> {lexicon(iv,Sym,Word,I),ivSem(Sym,IV)}, Word.
tv(I,TV)--> {lexicon(tv,Sym,Word,I),tvSem(Sym,TV)}, Word.
cop(Cop)--> {lexicon(cop,Sym,Word,_),tvSem(Sym,Cop)}, Word.
det(Det)--> {lexicon(det,_,Word,Type),detSem(Type,Det)}, Word.
pn(PN)--> {lexicon(pn,Sym,Word,G),pnSem(Sym,G,PN)}, Word.
pro(Pro)--> {lexicon(pro,Gender,Word,Type),proSem(Gender,Type,Pro)}, Word.
noun(N)--> {lexicon(noun,Sym,Word,_),nounSem(Sym,N)}, Word.
relpro(RP)--> {lexicon(relpro,_,Word,_),relproSem(RP)}, Word.
prep(Prep)--> {lexicon(prepare,Sym,Word,_),prepSem(Sym,Prep)}, Word.
mod(Mod)--> {lexicon(mod,_,Word,Type),modSem(Type,Mod)}, Word.
neg(Neg)--> [not], {modSem(neg,Neg)}.
coord(C)--> {lexicon(coord,_,Word,Type), coordSem(Type,C)}, Word.
dcoord(D,C)--> {lexicon(dcoord,Word,D,Type), dcoordSem(Type,C)}, Word.


```

/*****

    name: modelChecker.pl (Chapter 1)
    version: June 19, 1997; March 9, 1999.
    description: A model checker for first-order logic.
    authors: Patrick Blackburn & Johan Bos

*****/

:- module(modelChecker, [evaluate/2]).

:- use_module(comsemOperators),
   use_module(comsemPredicates, [member/2]),
   use_module(exampleModels, [constant/1, example/2]).

/*=====
   Evaluate a formula in a model
   =====*/

evaluate(Formula, Example):-
    example(Example, Model),
    satisfy(Formula, Model).

/*=====
   Semantic Evaluation
   =====*/

satisfy(exists(X, Formula), Model):-
    constant(X),
    satisfy(Formula, Model).

satisfy(forall(X, Formula), Model):-
    satisfy(~ exists(X, ~ Formula), Model).

satisfy(Formula1 & Formula2, Model):-
    satisfy(Formula1, Model),
    satisfy(Formula2, Model).

satisfy(Formula1 v Formula2, Model):-
    satisfy(Formula1, Model);
    satisfy(Formula2, Model).

satisfy(Formula1 > Formula2, Model):-
    satisfy(Formula2, Model);
    \+ satisfy(Formula1, Model).

satisfy(~ Formula, Model):-
    \+ satisfy(Formula, Model).

satisfy(Formula, Model):-

```

`member(Formula,Model).`

```

/*****

    name: modelChecker2.pl (Chapter 1)
    version: November 3, 1998
    description: Extention of modelChecker.pl
    authors: Patrick Blackburn & Johan Bos

*****/

:- module(modelChecker, [evaluate/2]).

:- use_module(comsemOperators),
   use_module(comsemPredicates, [member/2, compose/3]),
   use_module(exampleModels, [constant/1, relation/2, example/2]).

/*=====
   Evaluate a formula in a model
   =====*/

evaluate(Formula, Example):-
    sentence([], Formula),
    example(Example, Model),
    satisfy(Formula, Model).

evaluate(Formula, _Example):-
    \+ sentence([], Formula),
    nl, write('Not a wff over the given vocabulary.'),
    nl, write('Cannot be evaluated.'), nl.

/*=====
   Check if a formula is a sentence conform the vocabulary
   =====*/

sentence(_, Var):-
    var(Var), !, fail.

sentence(Bound, forall(X, Formula)):-
    var(X),
    sentence([X|Bound], Formula).

sentence(Bound, exists(X, Formula)):-
    var(X),
    sentence([X|Bound], Formula).

sentence(Bound, Formula1 > Formula2):-
    sentence(Bound, Formula1),
    sentence(Bound, Formula2).

sentence(Bound, Formula1 & Formula2):-
    sentence(Bound, Formula1),

```

```

    sentence(Bound,Formula2).

sentence(Bound,Formula1 v Formula2):-
    sentence(Bound,Formula1),
    sentence(Bound,Formula2).

sentence(Bound,~ Formula):-
    sentence(Bound,Formula).

sentence(Bound,Formula):-
    compose(Formula,Symbol,Arguments),
    length(Arguments,Arity),
    relation(Symbol,Arity),
    goodArguments(Bound,Arguments).

/*=====
    Check arguments (must be either constants or bound variables)
=====*/

goodArguments(_Bound, []).
goodArguments(Bound,[Arg|Others]):-
    member(Var,Bound),
    Var==Arg,!,
    goodArguments(Bound,Others).
goodArguments(Bound,[Arg|Others]):-
    constant(Arg),
    goodArguments(Bound,Others).

/*=====
    Semantic Evaluation
=====*/

satisfy(exists(X,Formula),Model):-
    \+ \+ (constant(X), satisfy(Formula,Model)).

satisfy(forall(X,Formula),Model):-
    satisfy(~ exists(X,~ Formula),Model).

satisfy(Formula1 & Formula2,Model):-
    satisfy(Formula1,Model),
    satisfy(Formula2,Model).

satisfy(Formula1 v Formula2,Model):-
    satisfy(Formula1,Model);
    satisfy(Formula2,Model).

satisfy(Formula1 > Formula2,Model):-
    satisfy(Formula2,Model);
    \+ satisfy(Formula1,Model).

```

```
satisfy(~ Formula,Model):-  
  \+ satisfy(Formula,Model).
```

```
satisfy(Formula,Model):-  
  member(Formula,Model).
```

```

/*****

    name: exampleModels.pl (Chapter 1)
    version: March 9, 1999
    description: Some example models defined over a vocabulary
    authors: Patrick Blackburn & Johan Bos

*****/

:- module(exampleModels,[constant/1,relation/2,example/2]).

/*=====
   Vocabulary
=====*/

relation(love,2).
relation(hate,2).
relation(tell,2).
relation(episode,2).
relation(in,2).
relation(customer,1).
relation(robber,1).
relation(joke,1).
relation(woman,1).
relation(man,1).

constant(mia).
constant(vincent).
constant(honey_bunny).
constant(pumpkin).
constant(jules).
constant(jody).
constant(j1).
constant(j2).
constant(e15).
constant(e13).

/*=====
   Example Models
=====*/

example(1,[customer(mia),customer(vincent),
           robber(pumpkin),robber(honey_bunny),
           love(pumpkin,honey_bunny)]).

example(2,[customer(mia),
           robber(pumpkin),robber(honey_bunny),
           love(pumpkin,honey_bunny)]).

```

```
example(3,[customer(mia),customer(vincent),
           robber(pumpkin),robber(honey_bunny),
           love(pumpkin,honey_bunny),love(mia,vincent)]).
```

```
example(4,[woman(mia),
           man(vincent),man(jules),
           joke(j1),joke(j2),
           in(j1,e15),
           episode(e13),episode(e15),
           tell(mia,j1)]).
```

```
example(5,[woman(mia),woman(jody),
           man(vincent),man(jules),
           joke(j1),joke(j2),
           in(j1,e15),in(j1,e13),
           episode(e13),episode(e15),
           tell(mia,j1),tell(jody,j2)]).
```

```

/*****

    name: experiment1.pl (Chapter 2)
    version: July 1, 1997
    description: This is the code of the first experiment
    authors: Patrick Blackburn & Johan Bos

*****/

:- use_module(comsemOperators).

/*=====
   Syntax-Semantics Rules
   =====*/

s(Sem)--> np(Sem), vp(SemVP),
{
    arg(1,SemVP,X),
    arg(1,Sem,X),
    arg(2,Sem,Matrix),
    arg(2,Matrix,SemVP)
}.

s(Sem)--> np(SemNP), vp(Sem),
{
    arg(1,Sem,SemNP)
}.

np(Sem)--> pn(Sem).

np(Sem)--> det(Sem), noun(SemNoun),
{
    arg(1,SemNoun,X),
    arg(1,Sem,X),
    arg(2,Sem,Matrix),
    arg(1,Matrix,SemNoun)
}.

vp(Sem)--> iv(Sem).

vp(Sem)--> tv(SemTV), np(Sem),
{
    arg(2,SemTV,X),
    arg(1,Sem,X),
    arg(2,Sem,Matrix),
    arg(2,Matrix,SemTV)
}.

vp(Sem)--> tv(Sem), np(SemNP),
{

```



```

    arg(2, Sem, SemNP)
  }.

/*=====
  Proper Names
=====*/

pn(vincent)--> [vincent].

pn(mia)--> [mia].

/*=====
  Transitive Verbs
=====*/

tv(love(_, _))--> [loves].

/*=====
  Intransitive Verbs
=====*/

iv(snort(_))--> [snorts].

/*=====
  Determiners
=====*/

det(exists(_, _ & _))--> [a].

det(forall(_, _ > _))--> [every].

/*=====
  Nouns
=====*/

noun(woman(_))--> [woman].

noun(footmassage(_))--> [foot, massage].

```

```

/*****

    name: experiment2.pl (Chapter 2)
    version: July 2, 1997
    description: This is the code of the second experiment
    authors: Patrick Blackburn & Johan Bos

*****/

:- use_module(comsemOperators).

/*=====
   Syntax-semantics rules
   =====*/

s(Sem)--> np(X,SemVP,Sem), vp(X,SemVP).

np(X,Scope,Sem)--> det(X,Restr,Scope,Sem), noun(X,Restr).

np(SemPN,Sem,Sem)--> pn(SemPN).

vp(X,Sem)--> iv(X,Sem).

vp(X,Sem)--> tv(X,Y,SemTV), np(Y,SemTV,Sem).

/*=====
   Proper Names
   =====*/

pn(vincent)--> [vincent].

pn(mia)--> [mia].

/*=====
   Transitive Verbs
   =====*/

tv(Y,Z,love(Y,Z))--> [loves].

/*=====
   Intransitive Verbs
   =====*/

iv(Y,snort(Y))--> [snorts].

/*=====
   Determiners
   =====*/

det(X,Restr,Scope,exists(X,Restr & Scope))--> [a].

```

det(X,Restr,Scope,forall(X,Restr > Scope))--> [every].

```
/*=====
  Nouns
=====*/
```

noun(X,woman(X))--> [woman].

noun(X,footmassage(X))--> [foot,massage].

```
/******  
    name: mainLambda.pl (Chapter 2)  
    version: May 15, 1997  
    description: Semantic Construction with Beta Conversion  
    authors: Patrick Blackburn & Johan Bos  
*****/  
  
:- module(mainLambda,[parse/0]).  
  
:- use_module(readLine,[readLine/1]),  
   use_module(comsemOperators),  
   use_module(comsemPredicates,[printRepresentation/1,compose/3]),  
   use_module(betaConversion,[betaConvert/2]).  
  
:- [englishGrammar], [englishLexicon], [semMacrosLambda].  
  
/*======  
   Driver Predicate  
=====*/  
  
parse:-  
    readLine(Sentence),  
    s(Formula,Sentence,[]),  
    betaConvert(Formula,Converted),  
    printRepresentation(Converted).
```

```

/*****

    name: betaConversion.pl (Chapter 2)
    version: March 11, 1998
    description: Implementation of Beta-Conversion
    authors: Patrick Blackburn & Johan Bos

*****/

:- module(betaConversion,[betaConvert/2]).

:- use_module(comsemOperators),
   use_module(comsemPredicates,[compose/3,substitute/4]).

/*=====
   Beta-Conversion
=====*/

betaConvert(Var,Result):-
    var(Var), !, Result=Var.
betaConvert(Functor @ Arg,Result):-
    compound(Functor),
    betaConvert(Functor,ConvertedFunctor),
    apply(ConvertedFunctor,Arg,BetaConverted), !,
    betaConvert(BetaConverted,Result).
betaConvert(Formula,Result):-
    compose(Formula,Functor,Formulas),
    betaConvertList(Formulas,ResultFormulas),
    compose(Result,Functor,ResultFormulas).

betaConvertList([],[]).
betaConvertList([Formula|Others],[Result|ResultOthers]):-
    betaConvert(Formula,Result),
    betaConvertList(Others,ResultOthers).

/*=====
   Application (Unification-Based)
=====*/

%apply(lambda(Argument,Result),Argument,Result).

/*=====
   Application (Substitution-Based)
=====*/

apply(lambda(X,Formula),Argument,Result):-
    substitute(Argument,X,Formula,Result).

```

```

/*****
    name: semMacrosLambda.pl (Chapter 2)
    version: March 10, 1999
    description: Semantic Macros for the Lambda Calculus
    authors: Patrick Blackburn & Johan Bos
*****/

/*=====
    Semantic Macros
=====*/

detSem(uni,lambda(P,lambda(Q,forall(X,(P@X)>(Q@X))))).
detSem(indef,lambda(P,lambda(Q,exists(X,(P@X)&(Q@X))))).
nounSem(Sym,lambda(X,Formula)):-
    compose(Formula,Sym,[X]).
pnSem(Sym,_Gender,lambda(P,P@Sym)).
proSem(_Gender,_Type,lambda(P,P@_)).
ivSem(Sym,lambda(X,Formula)):-
    compose(Formula,Sym,[X]).
tvSem(Sym,lambda(K,lambda(Y,K @ lambda(X,Formula)))):-
    compose(Formula,Sym,[Y,X]).
relproSem(lambda(P,lambda(Q,lambda(X,(P@X)&(Q@X))))).
prepSem(Sym,lambda(K,lambda(P,lambda(Y,(K@lambda(X,F)) & (P@Y)))):-
    compose(F,Sym,[Y,X]).
modSem(neg,lambda(P,lambda(X,~(P@X)))).
coordSem(conj,lambda(X,lambda(Y,lambda(P,(X@P) & (Y@P))))).
coordSem(disj,lambda(X,lambda(Y,lambda(P,(X@P) v (Y@P))))).
dcoordSem(cond,lambda(X,lambda(Y,X > Y))).
dcoordSem(conj,lambda(X,lambda(Y,X & Y))).
dcoordSem(disj,lambda(X,lambda(Y,X v Y))).

```

```
/******  
    name: mainMontague.pl (Chapter 3)  
    version: May 25, 1999  
    description: Montague's Rule of Quantification  
    authors: Patrick Blackburn & Johan Bos  
*****/  
  
:- module(mainMontague, [parse/0]).  
  
:- use_module(comsemOperators),  
   use_module(readLine, [readLine/1]),  
   use_module(comsemPredicates, [printReadings/1, compose/3]),  
   use_module(betaConversion, [betaConvert/2]).  
  
:- [englishGrammarMontague], [englishLexicon], [semMacrosLambda].  
  
/*=====Driver Predicate=====*/  
  
parse:-  
    readLine(Sentence),  
    findall(Sem2, (s([], Sem1, Sentence, []), betaConvert(Sem1, Sem2)), Readings),  
    printReadings(Readings).
```

```

/*****

    name: englishGrammarMontague.pl (Chapter 3)
    version: May 25, 1999
    description: Grammar rules for Montague quantifier raising
    authors: Patrick Blackburn & Johan Bos

*****/

/*=====
   Grammar Rules
   =====*/

d(S) --> s([],S).
d((C@S1)@S2)--> dcoord(D,C), s([],S1), D, d(S2).

s([],NP@lambda(I,S))--> s([bo(NP,I)],S).

s(Q,NP@VP)--> np2([],NP), vp2(Q,VP).

np2(Q,NP)--> np1(Q,NP).

np1([],Det@Noun)--> det(Det), n2(Noun).
np1([bo(Det@Noun,I)],lambda(P,P@I))--> det(Det), n2(Noun).
np1([],NP)--> pn(NP).
np1([],NP)--> pro(NP).

n2(N)--> n1(N).
n2((C@N1)@N2)--> n1(N1), coord(C), n1(N2).

n1(N)--> noun(N).
n1(PP@N)--> noun(N), pp(PP).
n1(RC@N)--> noun(N), rc(RC).

vp2(Q,VP)--> vp1(Q,VP).

vp1(Q,Mod@VP)--> mod(Mod), v2(Q,inf,VP).
vp1(Q,VP)--> v2(Q,fin,VP).

v2(Q,fin,Cop@NP)--> cop(Cop), np2(Q,NP).
v2(Q,fin,Neg@(Cop@NP))--> cop(Cop), neg(Neg), np2(Q,NP).

v2(Q,I,V)--> v1(Q,I,V).

v1([],I,V)--> iv(I,V).
v1(Q,I,TV@NP)--> tv(I,TV), np2(Q,NP).

pp(Prep@NP)--> prep(Prep), np2([],NP).

rc(RP@VP)--> relpro(RP), vp2([],VP).

```


iv(I,IV)--> {lexicon(iv,Sym,Word,I),ivSem(Sym,IV)}, Word.
tv(I,TV)--> {lexicon(tv,Sym,Word,I),tvSem(Sym,TV)}, Word.
cop(TV)--> {lexicon(cop,Sym,Word,_),tvSem(Sym,TV)}, Word.
det(Det)--> {lexicon(det,_,Word,Type),detSem(Type,Det)}, Word.
pn(PN)--> {lexicon(pn,Sym,Word,Gender),pnSem(Sym,Gender,PN)}, Word.
pro(Pro)--> {lexicon(pro,Gender,Word,Type),proSem(Gender,Type,Pro)}, Word.
noun(N)--> {lexicon(noun,Sym,Word,_),nounSem(Sym,N)}, Word.
relpro(RP)--> {lexicon(relpro,_,Word,_),relproSem(RP)}, Word.
prep(Prep)--> {lexicon(prepare,Sym,Word,_),prepSem(Sym,Prep)}, Word.
mod(Mod)--> {lexicon(mod,_,Word,Type),modSem(Type,Mod)}, Word.
neg(Neg)--> [not], {modSem(neg,Neg)}.
coord(C)--> {lexicon(coord,_,Word,Type), coordSem(Type,C)}, Word.
dcoord(D,C)--> {lexicon(dcoord,Word,D,Type), dcoordSem(Type,C)}, Word.

```

/*****

    name: mainCooperStorage.pl (Chapter 3)
    version: November 14, 1997
    description: Cooper Storage Implementation
    authors: Patrick Blackburn & Johan Bos

*****/

:- module(mainCooperStorage, [parse/0]).

:- use_module(readLine, [readLine/1]),
   use_module(comsemOperators),
   use_module(comsemPredicates, [printReadings/1, compose/3, append/3]),
   use_module(betaConversion, [betaConvert/2]).

:- [englishGrammar], [englishLexicon], [semMacrosStorage].

/*=====
   Driver Predicate
   =====*/

parse:-
    readLine(Sentence),
    s(Sem, Sentence, []),
    setof(Result, Store^Retrieved^(buildStore(Sem, Store),
                                   sRetrieval(Store, Retrieved),
                                   betaConvert(Retrieved, Result)),
          SemSet),
    printReadings(SemSet).

/*=====
   Quantifier Storage
   =====*/

npStorage(Quant, [Arg|Store], [lambda(P, P@X), bo(Quant@Arg, X)|Store]).

/*=====
   Quantifier Retrieval
   =====*/

sRetrieval([S], S).

sRetrieval([Sem|Store], S):-
    removeFromStore(bo(Q, X), Store, NewStore),
    sRetrieval([Q@lambda(X, Sem)|NewStore], S).

removeFromStore(X, [X|T], T).

removeFromStore(X, [Y|T], [Y|R]):-

```

```
removeFromStore(X,T,R).

/*=====
  Store Constructing
=====*/

buildStore(quant(Quant) @ Store,NewStore):-
  buildStore(Store,[Arg|S]),
  npStorage(Quant,[Arg|S],NewStore).

buildStore([Sem|Store],[Sem|Store]).

buildStore(Store1 @ Store2,[F@A|S]):-
  buildStore(Store1,[F|S1]),
  buildStore(Store2,[A|S2]),
  append(S1,S2,S).
```

```

/*****

    name: mainKellerStorage.pl (Chapter 3)
    version: November 14, 1997
    description: Keller Storage Implementation
    authors: Patrick Blackburn & Johan Bos

*****/

:- module(mainKellerStorage, [parse/0]).

:- use_module(readLine, [readLine/1]),
   use_module(comsemOperators),
   use_module(comsemPredicates, [printReadings/1, compose/3, append/3]),
   use_module(betaConversion, [betaConvert/2]).

:- [englishGrammar], [englishLexicon], [semMacrosStorage].

/*=====
   Driver Predicate
=====*/

parse:-
    readLine(Sentence),
    s(Sem, Sentence, []),
    setof(Result, Store^Retrieved^(buildStore(Sem, Store),
                                   sRetrieval(Store, Retrieved),
                                   betaConvert(Retrieved, Result)),
          SemSet),
    printReadings(SemSet).

/*=====
   Quantifier Storage
=====*/

npStorage(Quant, [Arg|Store], [lambda(P, P@X), bo([Quant@Arg|Store], X)]).

/*=====
   Quantifier Retrieval
=====*/

sRetrieval([S], S).

sRetrieval([Sem|Store], S):-
    removeFromStore(bo(Q, X), Store, NewStore),
    sRetrieval([Q@lambda(X, Sem)|NewStore], S).

removeFromStore(bo(0, I), [bo([0|T1], I)|T2], T):-
    append(T1, T2, T).

```

```
removeFromStore(X, [Y|T], [Y|R]):-
    removeFromStore(X,T,R).

/*=====
   Store Constructing
=====*/

buildStore(quant(Quant) @ Store,NewStore):-
    buildStore(Store, [Arg|S]),
    (
        NewStore = [Quant@Arg|S]
    ;
        npStorage(Quant, [Arg|S],NewStore)
    ).

buildStore([Sem|Store],[Sem|Store]).

buildStore(Store1 @ Store2,[F@A|S]):-
    buildStore(Store1, [F|S1]),
    buildStore(Store2, [A|S2]),
    append(S1,S2,S).
```

```

/*****
    name: semMacrosStorage.pl (Chapter 3)
    version: May 25, 1999
    description: Semantic Macros for Storage
    authors: Patrick Blackburn & Johan Bos
*****/

/*=====
    Semantic Macros
=====*/

detSem(uni, quant(lambda(P, lambda(Q, forall(X, (P@X)>(Q@X)))))).

detSem(indef, quant(lambda(P, lambda(Q, exists(X, (P@X)&(Q@X)))))).

nounSem(Sym, [lambda(X, Formula)]):-
    compose(Formula, Sym, [X]).

pnSem(Sym, _Gender, [lambda(P, P@Sym)]).

proSem(_Gender, _Type, [lambda(P, P@_)]).

ivSem(Sym, [lambda(X, Formula)]):-
    compose(Formula, Sym, [X]).

tvSem(Sym, [lambda(K, lambda(Y, K@lambda(X, Formula))))):-
    compose(Formula, Sym, [Y, X]).

relproSem([lambda(P, lambda(Q, lambda(X, (P@X)&(Q@X))))]).

prepSem(Sym, [lambda(K, lambda(P, lambda(Y, (K@lambda(X, F))&(P@Y))))):-
    compose(F, Sym, [Y, X]).

modSem(neg, [lambda(P, lambda(X, ~(P@X))]).

coordSem(conj, [lambda(X, lambda(Y, lambda(P, (X@P) & (Y@P))))]).
coordSem(disj, [lambda(X, lambda(Y, lambda(P, (X@P) v (Y@P))))]).

dcoordSem(cond, lambda(X, lambda(Y, X > Y))).
dcoordSem(conj, lambda(X, lambda(Y, X & Y))).
dcoordSem(disj, lambda(X, lambda(Y, X v Y))).

```

```

/*****

    name: pluggingAlgorithm.pl (Chapter 3)
    version: June 4, 1998
    description: Plugging Algorithm
    author: Patrick Blackburn & Johan Bos

*****/

:- module(pluggingAlgorithm, [plugHole/4]).

:- use_module(comsemPredicates, [member/2, select/3, variablesInTerm/2]).

/*=====
   Plugging Predicates
=====*/

plugHole(Formula, LFs1-LFs3, Constraints, Scoped):-
    select(Label:Formula, LFs1, LFs2),
    checkConstraints(Label, Constraints, [Formula|Scoped]),
    variablesInTerm(Formula, []-Arguments),
    checkArguments(Arguments, LFs2-LFs3, Constraints, [Formula|Scoped]).

checkArguments([], LFs-LFs, _, _).

checkArguments([Arg|Rest], LFs1-LFs3, Constraints, Scoped):-
    member(leq(_, Hole), Constraints),
    Hole==Arg, !,
    plugHole(Hole, LFs1-LFs2, Constraints, Scoped),
    checkArguments(Rest, LFs2-LFs3, Constraints, Scoped).

checkArguments([Arg|Rest], LFs1-LFs4, Constraints, Scoped):-
    select(Label:Formula, LFs1, LFs2),
    Label==Arg, !,
    Formula=Arg,
    variablesInTerm(Formula, []-Arguments),
    checkArguments(Arguments, LFs2-LFs3, Constraints, Scoped),
    checkArguments(Rest, LFs3-LFs4, Constraints, Scoped).

checkArguments([_|Rest], LFs1-LFs2, Constraints, Scoped):-
    checkArguments(Rest, LFs1-LFs2, Constraints, Scoped).

/*=====
   Constraint Checking
=====*/

checkConstraints(_, [], _).

checkConstraints(Label, [leq(L,H)|Constraints], Scoped):-
    Label==L, !,

```

```
member(Formula,Scoped), Formula==H,  
checkConstraints(Label,Constraints,Scoped).  
  
checkConstraints(Label,[_|Constraints],Scoped):-  
checkConstraints(Label,Constraints,Scoped).
```



```

/*****

    name: mainPLU.pl (Chapter 3)
    version: Jan 31, 1998
    description: Predicate Logic Unplugged
    author: Patrick Blackburn & Johan Bos

*****/

:- module(mainPLU,[parse/0]).

:- use_module(readLine,[readLine/1]),
   use_module(comsemOperators),
   use_module(betaConversion,[betaConvert/2]),
   use_module(pluggingAlgorithm,[plugHole/4]),
   use_module(mergeUSR,[mergeUSR/2]),
   use_module(comsemPredicates,[append/3,printRepresentation/1,
                                printReadings/1,compose/3]).

:- [englishGrammar], [englishLexicon], [semMacrosPLU].

/*=====
   Driver Predicate
   =====*/

parse:-
  readLine(Sentence),
  d(Sem,Sentence,[]),
  betaConvert(merge(usr([Top,Main],[],[ ]),Sem@Top@Main),Reduced),
  mergeUSR(Reducd,usr(D,L,C)),
  printRepresentation(usr(D,L,C)),
  findall(Top,plugHole(Top,L-[ ],C,[ ]),Readings),
  printReadings(Readings).

```

```

/*****

    name: semMacrosPLU.pl (Chapter 3)
    version: May 26, 1999
    description: Semantic Macros for Predicate Logic Unplugged
    author: Patrick Blackburn & Johan Bos

*****/

/*=====
   Semantic Macros
=====*/

detSem(uni, lambda(P, lambda(Q, lambda(H, lambda(L,
    merge(merge(usr([F, R, S],
                [F:forall(X, R>S)],
                [leq(F, H), leq(L, S)])), P@X@H@R), Q@X@H@L))))).

detSem(indef, lambda(P, lambda(Q, lambda(H, lambda(L,
    merge(merge(usr([E, R, S],
                [E:exists(X, R&S)],
                [leq(E, H), leq(L, S)])), P@X@H@R), Q@X@H@L))))).

nounSem(Sym, lambda(X, lambda(_, lambda(L, usr([], [L:Formula], []))))):-
    compose(Formula, Sym, [X]).

pnSem(Sym, _Gender, lambda(P, lambda(H, lambda(L, P@Sym@H@L)))).

proSem(_Gender, _Type, lambda(P, lambda(H, lambda(L, P@_@H@L)))).

ivSem(Sym, lambda(X, lambda(H, lambda(L, usr([], [L:F], [leq(L, H)]))))):-
    compose(F, Sym, [X]).

tvSem(Sym, lambda(K, lambda(Y, K@lambda(X, lambda(H, lambda(L,
    usr([], [L:Formula], [leq(L, H)]))))))):-
    compose(Formula, Sym, [Y, X]).

relproSem(lambda(P, lambda(Q, lambda(X, lambda(H, lambda(L,
    merge(usr([L2, L3, H1],
                [L:(L3&H1)],
                [leq(L2, H1)])),
    merge(P@X@H@L2, Q@X@H@L3)))))).

prepSem(Sym, lambda(K, lambda(P, lambda(Y, lambda(H, lambda(L3,
    merge(K@lambda(X, lambda(H, lambda(L1,
    usr([L2, H1],
        [L3:(L2&H1), L1:Formula],
        [leq(L1, H1)]))))@H@L1, P@Y@H@L2))))))):-
    compose(Formula, Sym, [Y, X]).

```

```
modSem(neg, lambda(P, lambda(X, lambda(H, lambda(L,
  merge(usr([N, S], [N: (~S)], [leq(N, H), leq(L, S)]), P@X@H@L)))))).
```

```
coordSem(conj, lambda(X, lambda(Y, lambda(P, lambda(H, lambda(L,
  merge(usr([L1, L2], [L: (L1&L2)], [leq(L, H)]),
  merge(X@P@H@L1, Y@P@H@L2)))))).
```

```
coordSem(disj, lambda(X, lambda(Y, lambda(P, lambda(H, lambda(L,
  merge(usr([L1, L2], [L: (L1 v L2)], [leq(L, H)]),
  merge(X@P@H@L1, Y@P@H@L2)))))).
```

```
dcoordSem(cond, lambda(C1, lambda(C2, lambda(H, lambda(L,
  merge(usr([H1, H2], [L: (H1 > H2)], [leq(L, H)]),
  merge(C1@H1@_, C2@H2@_)))))).
```

```
dcoordSem(conj, lambda(C1, lambda(C2, lambda(H, lambda(L,
  merge(usr([H1, H2], [L: (H1 & H2)], [leq(L, H)]),
  merge(C1@H1@_, C2@H2@_)))))).
```

```
dcoordSem(disj, lambda(C1, lambda(C2, lambda(H, lambda(L,
  merge(usr([H1, H2], [L: (H1 v H2)], [leq(L, H)]),
  merge(C1@H1@_, C2@H2@_)))))).
```

```
/******  
    name: mergeUSR.pl (Chapter 3)  
    version: June 18, 1999  
    description: Definition of the merge for USRs  
    authors: Patrick Blackburn & Johan Bos  
*****/  
  
:- module(mergeUSR, [mergeUSR/2]).  
  
:- use_module(comsemPredicates, [append/3]).  
  
/*===== Merge for Underspecified Semantic Representations =====*/  
  
mergeUSR(usr(D,L,C),usr(D,L,C)).  
  
mergeUSR(merge(U1,U2),usr(D3,L3,C3)):-  
    mergeUSR(U1,usr(D1,L1,C1)),  
    mergeUSR(U2,usr(D2,L2,C2)),  
    append(D1,D2,D3),  
    append(L1,L2,L3),  
    append(C1,C2,C3).
```

```

/*****
    name: propTabl.pl (Chapter 4)
    version: Dec 3, 1998
    description: Propositional Tableaux Program
    authors: Patrick Blackburn & Johan Bos
*****/

:- module(propTabl,[valid/1,saturate/1]).

:- use_module(comsemOperators),
   use_module(comsemPredicates,[member/2,removeFirst/3]).

/*=====

    General Tableau Predicates
    -----

saturate(+Tableau)
    Expand Tableau until it is closed.

closed(+Tableau)
    Every branch of Tableau contains a contradiction.

valid(+Formula)
    Try to create a closed tableau expansion for f(Formula).

=====*/

saturate(Tableau):-
    closed(Tableau).

saturate(OldTableau):-
    expand(OldTableau,NewTableau),
    saturate(NewTableau).

closed([]).

closed([Branch|Rest]):-
    member(t(X),Branch),
    member(f(X),Branch),
    closed(Rest).

valid(F):-
    saturate([[f(F]])).

/*=====

    Tableau Expansion Predicates

```

```

-----

expand(+Oldtableau,-Newtableau)
  Newtableaux is the result of applying a tableaux expansion
  rule to Oldtableaux.

unaryExpansion(+Branch,-NewBranch)
  Take Branch as input, and return NewBranches if a tableau rule
  allows unary expansion.

conjunctiveExpansion(+Branch,-NewBranch)
  Take Branch as input, and return the NewBranch if a tableau rule
  allows conjunctive expansion.

disjunctiveExpansion(+Branch,-NewBranch1,-NewBranch2)
  Take Branch as input, and return the NewBranch1 and NewBranch2
  if a tableau rule allows disjunctive expansion.

=====*/

expand([Branch|Tableau],[NewBranch|Tableau]):-
  unaryExpansion(Branch,NewBranch).

expand([Branch|Tableau],[NewBranch|Tableau]):-
  conjunctiveExpansion(Branch,NewBranch).

expand([Branch|Tableau],[NewBranch1,NewBranch2|Tableau]):-
  disjunctiveExpansion(Branch,NewBranch1,NewBranch2).

expand([Branch|Rest],[Branch|Newrest]):-
  expand(Rest,Newrest).

unaryExpansion(Branch,[Component|Temp]):-
  unary(SignedFormula,Component),
  removeFirst(SignedFormula,Branch,Temp).

conjunctiveExpansion(Branch,[Comp1,Comp2|Temp]):-
  conjunctive(SignedFormula,Comp1,Comp2),
  removeFirst(SignedFormula,Branch,Temp).

disjunctiveExpansion(Branch,[Comp1|Temp],[Comp2|Temp]):-
  disjunctive(SignedFormula,Comp1,Comp2),
  removeFirst(SignedFormula,Branch,Temp).

/*=====

  Formula Identification
  -----

conjunctive(?F,?Comp1,?Comp2)

```

F is a conjunctive signed formula with components Comp1 and Comp2.

disjunctive(?F,?Comp1,?Comp2)

F is a disjunctive signed formula with components Comp1 and Comp2.

unary(?F,?Comp)

F is a signed formula with component Comp.

=====*/

conjunctive(t(X & Y),t(X),t(Y)).

conjunctive(f(X v Y),f(X),f(Y)).

conjunctive(f(X > Y),t(X),f(Y)).

disjunctive(f(X & Y),f(X),f(Y)).

disjunctive(t(X v Y),t(X),t(Y)).

disjunctive(t(X > Y),f(X),t(Y)).

unary(t(~X),f(X)).

unary(f(~X),t(X)).

```

/*****

    name: freeVarTabl.pl (Chapter 5)
    version: Dec 10, 1998
    description: Free Variable Semantic Tableaux
                 Uses a number of ideas from Melvin Fitting's
                 implementation of an unsigned tableaux theorem prover
                 for first-order logic, in "First-Order Logic and
                 Automated Theorem Proving", Second Edition (1996),
                 Graduate Texts in Computer Science, Springer.
                 For more details, see the Notes to Chapter 5.
    authors: Patrick Blackburn & Johan Bos

*****

:- module(freeVarTabl,[valid/2,saturate/2,notatedFormula/3]).

:- use_module(comsemOperators),
   use_module(comsemPredicates,[member/2,removeFirst/3,unify/2,append/3,
                                compose/3,newFunctionCounter/1,
                                substitute/4]).

/*=====

    General Tableaux Predicates
    -----

saturate(+Tableau,+Qdepth)
    Expand Tableau until it is closed, allowing Qdepth
    applications of the universal rule.

closed(+Tableau)
    Every branch of Tableau can be made to contain a contradiction,
    after a suitable free variable substitution.

skolemFunction(+VarList,-SkoTerm)
    VarList is a list of free variables, and SkoTerm is a previously
    unused Skolem function symbol fun(N) applied to those free variables.

valid(+X,+Qdepth)
    Try to create a tableau expansion for f(X) that is closed allowing
    Qdepth applications of the universal rule.

notatedFormula(Notated,Free,SignedFormula)
    Notated is a notated formula, Free its associated free variable
    list, and SignedFormula the signed formula part.

=====*/

saturate(Tableau,_Q):-

```



```

closed(Tableau).

saturate(OldTableau,Qdepth):-
  expand(OldTableau,Qdepth,NewTableau,NewQdepth),!,
  saturate(NewTableau,NewQdepth).

closed([]).
closed([Branch|Rest]):-
  member(NotatedOne,Branch),
  notatedFormula(NotatedOne,_,t(X)),
  member(NotatedTwo,Branch),
  notatedFormula(NotatedTwo,_,f(Y)),
  unify(X,Y),
  closed(Rest).

skolemFunction(VarList,SkolemTerm) :-
  newFunctionCounter(N),
  compose(SkolemTerm,fun,[N|VarList]).

valid(X,Qdepth):-
  notatedFormula(NotatedFormula,[],f(X)),
  saturate([[NotatedFormula]],Qdepth).

notatedFormula(n(Free,Formula),Free,Formula).

/*=====

  Tableaux Expansion Predicates
  -----

expand(+Oldtableau,+OldQdepth,-Newtableau,-NewQdepth)
  Newtableaux with Q-depth NewQdepth is the result of applying
  a tableaux expansion rule to Oldtableaux with a Q-depth of OldQdepth.

unaryExpansion(+Branch,-NewBranch)
  Take Branch as input, and return NewBranches if a tableau rule
  allows unary expansion.

conjunctiveExpansion(+Branch,-NewBranch)
  Take Branch as input, and return the NewBranch if a tableau rule
  allows conjunctive expansion.

disjunctiveExpansion(+Branch,-NewBranch1,-NewBranch2)
  Take Branch as input, and return the NewNbranch1 and NewBranch2
  if a tableau rule allows disjunctive expansion.

existentialExpansion(+Branch,-NewBranch)
  Take Branch as input, and return the NewBranch if a tableau rule
  allows existential expansion.

```

```
universalExpansion(+Branch,+OldQDepth,-NewBranch,-NewQDepth)
  Take Branch and OldQD as input, and return the NewBranch and
  NewQDepth if a tableau rule allow universal expansion.
```

```
=====*/
```

```
expand([Branch|Tableau],QD,[NewBranch|Tableau],QD):-
  unaryExpansion(Branch,NewBranch).
```

```
expand([Branch|Tableau],QD,[NewBranch|Tableau],QD):-
  conjunctiveExpansion(Branch,NewBranch).
```

```
expand([Branch|Tableau],QD,[NewBranch|Tableau],QD):-
  existentialExpansion(Branch,NewBranch).
```

```
expand([Branch|Tableau],QD,[NewBranch1,NewBranch2|Tableau],QD):-
  disjunctiveExpansion(Branch,NewBranch1,NewBranch2).
```

```
expand([Branch|Tableau],OldQD,NewTableau,NewQD):-
  universalExpansion(Branch,OldQD,NewBranch,NewQD),
  append(Tableau,[NewBranch],NewTableau).
```

```
expand([Branch|Rest],OldQD,[Branch|Newrest],NewQD):-
  expand(Rest,OldQD,Newrest,NewQD).
```

```
unaryExpansion(Branch,[NotatedComponent|Temp]) :-
  unary(SignedFormula,Component),
  notatedFormula(NotatedFormula,Free,SignedFormula),
  removeFirst(NotatedFormula,Branch,Temp),
  notatedFormula(NotatedComponent,Free,Component).
```

```
conjunctiveExpansion(Branch,[NotatedComp1,NotatedComp2|Temp]) :-
  conjunctive(SignedFormula,Comp1,Comp2),
  notatedFormula(NotatedFormula,Free,SignedFormula),
  removeFirst(NotatedFormula,Branch,Temp),
  notatedFormula(NotatedComp1,Free,Comp1),
  notatedFormula(NotatedComp2,Free,Comp2).
```

```
disjunctiveExpansion(Branch,[NotComp1|Temp],[NotComp2|Temp]) :-
  disjunctive(SignedFormula,Comp1,Comp2),
  notatedFormula(NotatedFormula,Free,SignedFormula),
  removeFirst(NotatedFormula,Branch,Temp),
  notatedFormula(NotComp1,Free,Comp1),
  notatedFormula(NotComp2,Free,Comp2).
```

```
existentialExpansion(Branch,[NotatedInstance|Temp]) :-
  notatedFormula(NotatedFormula,Free,SignedFormula),
  existential(SignedFormula),
  removeFirst(NotatedFormula,Branch,Temp),
  skolemFunction(Free,Term),
```

```

instance(SignedFormula,Term,Instance),
notatedFormula(NotatedInstance,Free,Instance).

universalExpansion(Branch,OldQD,NewBranch,NewQD):-
OldQD > 0, NewQD is OldQD - 1,
member(NotatedFormula,Branch),
notatedFormula(NotatedFormula,Free,SignedFormula),
universal(SignedFormula),
removeFirst(NotatedFormula,Branch,Temp),
instance(SignedFormula,V,Instance),
notatedFormula(NotatedInstance,[V|Free],Instance),
append([NotatedInstance|Temp],[NotatedFormula],NewBranch).

/*=====

Formula Identification
-----

conjunctive(?F,?Comp1,?Comp2)
  F is a conjunctive signed formula with components Comp1 and Comp2

disjunctive(?F,?Comp1,?Comp2)
  F is a disjunctive signed formula with components Comp1 and Comp2

unary(?F,?Comp)
  F is a signed formula with component Comp

universal(?F)
  F is a universal formula.

existential(?F)
  F is an existential formula.

instance(F,Term,Ins)
  F is a signed quantified formula, and Ins is the result of
  removing the quantifier and replacing all free occurrences of
  the quantified variable by occurrences of Term.

=====*/

conjunctive(t(X & Y),t(X),t(Y)).
conjunctive(f(X v Y),f(X),f(Y)).
conjunctive(f(X > Y),t(X),f(Y)).

disjunctive(f(X & Y),f(X),f(Y)).
disjunctive(t(X v Y),t(X),t(Y)).
disjunctive(t(X > Y),f(X),t(Y)).

unary(t(~X),f(X)).
unary(f(~X),t(X)).

```

```
universal(t(forall(_,_))).
universal(f(exists(_,_))).

existential(t(exists(_,_))).
existential(f(forall(_,_))).

instance(t(forall(X,F)),Term,t(NewF):-
  substitute(Term,X,F,NewF).
instance(f(exists(X,F)),Term,f(NewF):-
  substitute(Term,X,F,NewF).
instance(t(exists(X,F)),Term,t(NewF):-
  substitute(Term,X,F,NewF).
instance(f(forall(X,F)),Term,f(NewF):-
  substitute(Term,X,F,NewF).
```

```
/******  
    name: callTheoremProver.pl (Chapter 5)  
    version: June 18, 1998  
    description: Prolog Interface to Otter (Sicstus required)  
    authors: Patrick Blackburn & Johan Bos  
*****/  
  
:- module(callTheoremProver,[callTheoremProver/3]).  
  
:- use_module(library(system)),  
   use_module(fol2otter,[fol2otter/2]).  
  
/*=====  
   Calls to Theorem Prover (Otter)  
=====*/  
  
callTheoremProver(Axioms,Formula,Proof):-  
    fol2otter(Axioms,Formula),  
    shell('./otter < temp.in > temp.out 2> /dev/null',X),  
    (X=26368,Proof=yes,!;X=26624,Proof=no).
```

```

/*****
    name: callModelBuilder.pl (Chapter 5)
    version: September 3, 1999
    description: Prolog Interface to Mace (Sicstus required)
    authors: Patrick Blackburn & Johan Bos
*****/

:- module(callModelBuilder,[callModelBuider/4]).

:- use_module(library(system)),
   use_module(comsemPredicates,[append/3]),
   use_module(fol2otter,[fol2otter/2]).

/*=====
    Calls to Model Generator (Mace)

    Changed in MACE (generate.c):
    #define MAX_SYMBOLS 100 number of functors (was 50)
=====*/

callModelBuilder(Axioms,Formula,DomainSize,Model):-
    fol2otter(Axioms,Formula),
    name('./mace -n',C1),
    name(DomainSize,C2),
    append(C1,C2,C3),
    name(' -p -t2 -m1 < temp.in > temp.out 2> /dev/null',C4),
    append(C3,C4,C5),
    name(Shell,C5),
    shell(Shell,X),
    % this is not correct. It always returns 0!
    (X=0,Model=1,!;write(X),Model=0).

```

```

/*****

    name: fol2otter.pl (Chapter 11)
    version: June 18, 1998
    description: Translates a formula in otter syntax to standard output
    authors: Patrick Blackburn & Johan Bos

*****/

:- module(fol2otter, [fol2otter/2]).

:- use_module(comsemOperators).

/*=====
   Translates formula to otter syntax in file 'temp.in'
=====*/

fol2otter(Axioms,Formula):-
    tell('temp.in'),
    format('set(auto).~n~n',[]),
    format('clear(print_proofs).~n~n',[]),
    format('set(prolog_style_variables).~n~n',[]),
    format('formula_list(usable).~n~n~n',[]),
    printOtterList(Axioms),
    printOtterFormula(Formula),
    format('~nend_of_list.~n',[]),
    told.

/*=====
   Print a list of Otter formulas
=====*/

printOtterList([]).
printOtterList([X|L]):-
    printOtterFormula(X),
    printOtterList(L).

/*=====
   Print an Otter formula
=====*/

printOtterFormula(F):-
    \+ \+ (numbervars(F,0,_), printOtter(F,5)),
    format('~n',[]).

printOtter(exists(X,Formula),Tab):-
    write('(exists '),write(X),write(' '),!,
    printOtter(Formula,Tab),write(')').

printOtter(forall(X,Formula),Tab):-

```

```
write('(all '),write(X),write(' '),!,
printOtter(Formula,Tab),write(')')).

printOtter(Phi & Psi,Tab):-
write('('),!,
printOtter(Phi,Tab),
write(' & '), nl, tab(Tab),
NewTab is Tab + 5,
printOtter(Psi,NewTab), write(')').

printOtter(Phi v Psi,Tab):-
write('('),!,
printOtter(Phi,Tab), write(' | '),
printOtter(Psi,Tab), write(')').

printOtter(Phi <> Psi,Tab):-
write('('),!,
printOtter(Phi,Tab), write(' <-> '),
printOtter(Psi,Tab), write(')').

printOtter(Phi > Psi,Tab):-
write('('),!,
printOtter(Phi,Tab), write(' -> '),
printOtter(Psi,Tab), write(')').

printOtter(~ Phi,Tab):-
write('~('),!,
printOtter(Phi,Tab), write(')').

printOtter(Phi,_):-
write(Phi).
```



```

/*****

    name: nlQuestions.pl (Chapter 6)
    version: Dec 16, 1998
    description: Combining the model checker and the lambda calculus
                 to answer simple natural language questions.
    authors: Patrick Blackburn & Johan Bos

*****/

:- module(nlQuestions,[query/1,q/3]).

:- use_module(modelChecker2,[evaluate/2]),
   use_module(readLine,[readLine/1]),
   use_module(comsemOperators),
   use_module(comsemPredicates,[compose/3]),
   use_module(betaConversion,[betaConvert/2]).

:- [englishGrammar], [englishLexicon], [semMacrosLambda].

/*=====

    Driver Predicate
    -----

query(?Example)
    Tries to translate the user's input Question, into a lambda
    expression, and evaluate the Sem-part of it in the model named Example.
    Prints all the answers for which the formula can be evaluated, and
    then asks again for a new question.

=====*/

query(Example):-
    readLine(Question),
    q(Q,Question,[]),
    betaConvert(Q,lambda(Answer,Sem)),
    evaluate(Sem,Example),
    write(Answer), nl, fail.

query(Example):-
    query(Example).

/*=====

    Grammar rules for Wh-Questions
    =====*/

q(WH@VP)--> wh(WH), vp2(VP).

```

$\text{wh}(\lambda(P, \lambda(X, P@X))) \rightarrow [\text{who}]$.

$\text{wh}(\lambda(P, \lambda(X, (N@X) \& (P@X)))) \rightarrow [\text{which}], n2(N)$.

```

/*****

    name: nlArgumentation.pl (Chapter 6)
    version: December 16, 1998
    description: Natural Language Argumentation
    authors: Patrick Blackburn & Johan Bos

*****/

:- module(nlArgumentation, [argument/0]).

:- use_module(freeVarTabl, [notatedFormula/3, saturate/2]),
   use_module(readLine, [readLine/1]),
   use_module(comsemOperators),
   use_module(comsemPredicates, [compose/3]),
   use_module(betaConversion, [betaConvert/2]).

:- [englishGrammar], [englishLexicon], [semMacrosLambda].

/*=====
   Interface to the free variable tableaux
=====*/

validArgument([Conclusion|Premises], Qdepth):-
    notatedFormula(NotatedConclusion, [], f(Conclusion)),
    premises(Premises, [], NotatedPremises),
    saturate([[NotatedConclusion|NotatedPremises]], Qdepth).

/*=====
   Auxiliary predicate to translate premises
=====*/

premise([], N, N).
premise([P|Premises], SoFar, NotatedPremises):-
    notatedFormula(NotatedP, [], t(P)),
    premise(Premises, [NotatedP|SoFar], NotatedPremises).

/*=====
   The driver predicate
=====*/

argument:-
    enterPremises(Premises),
    enterConclusion(Conclusion),
    (
        validArgument([Conclusion|Premises], 10),
        nl, write('This is a valid argument!'), nl, !
    ;
        nl, write('Not a valid argument...'), nl

```

```
) .

/*=====
   Predicate for entering the premises
=====*/

enterPremises(Premises):-
  nl, write('Enter premises (or Return to continue):'),
  readLine(Input),
  (
    Input= [], !,
    Premises= []
  ;
    d(SemPremise, Input, []),
    betaConvert(SemPremise, ConvertedPremise),
    enterPremises(Others),
    Premises=[ConvertedPremise|Others]
  ).

/*=====
   Predicate for entering the conclusion
=====*/

enterConclusion(Conclusion):-
  nl, write('Enter conclusion:'),
  readLine(Input),
  d(SemConclusion, Input, []),
  betaConvert(SemConclusion, Conclusion).
```

```

/*****

    name: semOntology.pl (Chapter 6)
    version: July 10, 1999
    description: Predicates for working with the semantic ontology
    authors: Patrick Blackburn & Johan Bos

*****/

:- module(semOntology,[generateOntology/1,consistent/2]).

:- use_module(comsemPredicates,[member/2,append/3,compose/3]),
   use_module(comsemOperators).

:- [englishLexicon].

/*=====
   Generating Ontology in First-Order Formulas
   =====*/

generateOntology(Formulas):-
    generateIsa(I0),
    generateDisjoint(I0-I1,I2),
    isa2fol(I1,[]-F),
    isa2fol(I2,F-Formulas).

/*=====
   Generating isa/2 relations
   =====*/

generateIsa(I):-
    setof(isa(Hypo,Hyper),Words^lexicon(noun,Hypo,Words,Hyper),I).

/*=====
   Generating disjoint/2 relations (on the basis of isa/2)
   =====*/

generateDisjoint([]-[],[]).

generateDisjoint([isa(A,[Hyper])|L1]-[isa(A,Hyper)|L2],I3):-!,
    findall(disjoint(A,B),member(isa(B,[Hyper]),L1),I1),
    generateDisjoint(L1-L2,I2),
    append(I1,I2,I3).

generateDisjoint([isa(A,Hyper)|L1]-[isa(A,Hyper)|L2],I):-
    generateDisjoint(L1-L2,I).

/*=====
   Translating ISA-relations to first-order formulas
   =====*/

```

```

=====*/
isa2fol([],A-A):- !.

isa2fol([isa(S1,[S2])|L],A1-[forall(X,F1 > F2)|A2]):- !,
    compose(F1,S1,[X]),
    compose(F2,S2,[X]),
    isa2fol(L,A1-A2).

isa2fol([isa(S1,S2)|L],A1-[forall(X,F1 > F2)|A2]):-
    compose(F1,S1,[X]),
    compose(F2,S2,[X]),
    isa2fol(L,A1-A2).

isa2fol([disjoint(S1,S2)|L],A1-[forall(X,F1 > ~ F2)|A2]):-
    compose(F1,S1,[X]),
    compose(F2,S2,[X]),
    isa2fol(L,A1-A2).

/*=====
    Consistency Check
=====*/

consistent(X,Y):-
    generateIsa(I),
    generateDisjoint(I-Isa,Disjoint),
    \+ inconsistent(X,Y,Isa,[disjoint(human,nonhuman)|Disjoint]).

inconsistent(X,Y,_,Disjoint):-
    member(disjoint(X,Y),Disjoint).

inconsistent(X,Y,_,Disjoint):-
    member(disjoint(Y,X),Disjoint).

inconsistent(X,Y,Isa,Disjoint):-
    member(isa(X,Z),Isa),
    inconsistent(Z,Y,Isa,Disjoint).

inconsistent(X,Y,Isa,Disjoint):-
    member(isa(Y,Z),Isa),
    inconsistent(X,Z,Isa,Disjoint).

```

Representation and Inference for Natural Language

A First Course in Computational Semantics

Volume II

Working with Discourse Representation Structures

Patrick Blackburn & Johan Bos

September 3, 1999

Contents

Preface	iii
1 Discourse Representation Theory	1
1.1 An Overview of DRT	2
1.2 Interpreting DRSs	17
1.3 DRT and First-Order Logic	22
1.4 DRT in Prolog	27
2 Building Discourse Representations	37
2.1 DRS-Threading	37
2.2 Building DRSs with Lambdas	47
2.3 Underspecified DRSs	53
2.4 Merging into Darkness?	57
3 Pronoun Resolution	61
3.1 The Nature of Pronouns	61
3.2 Implementing Pronoun Resolution in DRT	63
3.3 Adding Reflexive Pronouns	69
3.4 The Focusing Algorithm	72
4 Presupposition Resolution	85
4.1 Introducing Presuppositions	85
4.2 Dealing with Presupposition in DRT	88

4.3	Presupposition Resolution in Prolog	96
4.4	Optimizing the Algorithm	103
5	Acceptability Constraints	109
5.1	Maxims of Conversation	109
5.2	Choosing a Theorem Prover	110
5.3	Chatting with Curt	113
5.4	Adding Local Constraints	116
6	Putting It All Together	121
6.1	Combining Scope and Presupposition	121
6.2	Adding Focus	121
	Bibliography	123
A	Propositional Languages	127
B	Type Theory	129
C	Theorem Provers and Model Builders	133
D	Prolog in a Nutshell	135
E	Listing of Programs	153

Preface

In Volume II we change our underlying logical representation formalism: instead of using first-order representations we will use Discourse Representation Structures (DRSs). DRSs are the representations used in *Discourse Representation Theory* (DRT), one of the most influential (and certainly one of the most interesting) current approaches to the semantics of natural language. In spite of this change, we will be able to reuse the grammars and tools developed in Volume I without difficulty. We then investigate some interesting semantic phenomena, namely pronoun resolution and presupposition projection, and integrate—at least partially—our work on inference with our work on semantic construction. Chapter by chapter, Volume II looks like this:

Chapter 1. Discourse Representation Theory. This chapter introduces DRT and prepares the way for our later computational work. We discuss the notion of context change potential, define DRS languages, discuss accessibility and the standard DRS construction algorithm, give two (equivalent) semantics for DRS languages, and show how to translate DRSs into first-order logic. We then implement a simple model checker, and a DRS to first-order logic compiler.

Chapter 2. Building Discourse Representations. We first examine three distinct techniques for constructing the DRSs. First, we discuss threading, an elegant and efficient approach that is essentially the standard DRS construction algorithm viewed declaratively. Second, we show how straightforward it is to transplant the lambda calculus technology developed in Volume I to the setting of DRT; we call the resulting blend λ -DRT. Third, we integrate DRS construction with underspecification. Finally, we discuss the merging of DRSs.

Chapter 3. Pronoun Resolution. Here we turn to a key topic for DRT: pronoun resolution. We develop a Prolog program which determines which antecedent NPs are accessible, and performs pronoun resolution on the basis of this information. We will then extend this program to deal with reflexive pronouns. Finally, we use a focusing algorithm to get—in case ambiguities arise—a preferred reading.

Chapter 4. Presupposition Resolution. What exactly are presuppositions, and how can we solve the projection problem for presupposition? In this chapter we discuss and implement an elegant DRT-based approach to these problems due to Rob van der Sandt. This approach treats presuppositions as anaphors which may be accommodated, offering a novel perspective on presupposition projection. In this chapter we show how the pronoun resolution techniques introduced in Chapter 3 can be extended to handle presupposition accommodation.

Chapter 5. Acceptability Constraints. The algorithms we have supplied for dealing with pronouns and presupposition are genuinely over-generating. But on the basis of a simple maxims of conversations, such as consistency and informativity, some of them can be ruled out. In this chapter we show how to do this.

Chapter 6. Putting It All Together. Here we combine the techniques we have acquired in this book. We'll try to put together a system that integrates our techniques for dealing with scope, anaphora, and presupposition. We show that *hole semantics* developed in Volume I adapts straightforwardly to DRT, thus giving us a mechanism for coping with scope ambiguities in DRT.

Each chapter concludes with two sections: *Software Summary* lists the programs developed in the chapter, and *Notes* lists references the reader may find helpful, and discusses more advanced topics.

This book developed out of material for a course on Computational Semantics we regularly offer at the Department of Computational Linguistics, University of the Saarland. We also taught a preliminary version (essentially the material that now makes up Volume I) as an introductory course at ESSLLI'97, the *European Summer School in Logic, Language and Information* held at Aix-en-Provence, France in August 1997. When designing these courses, we found no single source which contained all the material we wanted to present. At that time, the only notes solely devoted to computational semantics we knew of were Cooper et al. 1993. These notes, which we recommend to our readers, were developed at the University of Edinburgh, and are probably the first systematic introduction to modern computational semantics. Like the present book they are Prolog based, and cover some of the same ground using interestingly different tools and techniques. However we wanted to teach the subject in a way that emphasized such ideas as inference, underspecification, and architectural issues. This led us to a first version of the book, which was heavily influenced by Pereira and Shieber 1987 for semantic construction, Fitting 1996 and Smullyan 1995 for tableaux systems, and Kamp and Reyle 1993 for DRT. Since then, the project has taken on a life of its own, and grown in a variety of (often unexpected) directions. Both the code and the text has been extensively rewritten and we are now (we hope!) in the final stretch of producing the kind of introduction to computational semantics that we wanted all along.

Acknowledgments

We would like to thank Manfred Pinkal and all our colleagues at the Department of Computational Linguistics, University of Saarland, Saarbrücken, Germany. Thanks to Manfred, the department has become an extremely stimulating place for working on computational semantics; it was the ideal place to write this book.

Johan Bos would like to thank his colleagues in the *Verbmobil* Project (grant 01 IV 701 R4 and 01 IV 701 N3) at IMS-Stuttgart, TU-Berlin, CSLI Stanford, the DFKI, and the Department of Computational Linguistics at Saarbrücken, and those in the Trindi Project (Task Oriented Instructional Dialogue; LE4-8314) at the University of Gothenburg, SRI Cambridge, University of Edinburgh, and Xerox Grenoble.

Conversations with Paul Dekker, Andreas Franke, Martin Kay, Hans Kamp, Michael Kohlhase, Karsten Konrad, Christof Monz, Reinhard Muskens, Maarten de Rijke, and Henk Zeevat were helpful in clarifying our goals. We're grateful to David Beaver, Judith Baur, Björn Gambäck, Ewan Klein, Emiel Kraemer, Rob van der Sandt, and Frank Schilder for their comments on an early draft of this book. We're also very grateful to our students both at Saarbrücken and at ESSLLI, who has provided us with invaluable feedback. In particular we would like to thank Aljoscha Burchardt, Gerd Fliedner, Malte Gabsdil, Kristina Striegnitz, Stefan Thater, and Stephan Walter. Patrick Blackburn would like to thank Aravind Joshi and the staff of Institute for Research in Cognitive Science at the University of Pennsylvania for their hospitality while the Montreal version was being prepared.

Finally, we are deeply grateful to Robin Cooper, who taught a course based on the previous draft of this book, and provided us with extremely detailed feedback on what worked and what didn't. His comments have greatly improved the book.

*Patrick Blackburn and Johan Bos,
Computerlinguistik, Universität des Saarlandes,
September 1999.*

Chapter 1

Discourse Representation Theory

In Part I we used first-order formulas as our underlying representations; in Part II we switch to *Discourse Representation Structures (DRSs)*, the representations used in *Discourse Representation Theory (DRT)*. DRSs will enable us to encode the information contained in multi-sentence discourses, and to do so in a way that yields a natural treatment of such phenomena as anaphoric pronouns and presuppositions. As we shall see, DRT is compatible with a wide range of implementation strategies, and we will be able to reuse much of our work from Part I with little or no fuss.

This chapter introduces the basic ideas of DRT. The first section gives a swift overview of DRT that encourages the reader to think of DRSs (or *boxes*) as pictures. We show that it is useful to view discourse comprehension as a process of constructing a picture of the changing context, define *DRS languages*, informally sketch the *embedding semantics*, introduce the concept of *accessibility*, and discuss the *standard construction algorithm*. The remaining sections then explore semantics of DRS languages in more detail, both theoretically and computationally. First, we present two alternative interpretations of DRS languages: the embedding semantics, which emphasizes the intuition that DRSs are pictures, and *dynamic semantics*, which emphasizes the view that DRSs are context-transforming programs. We then examine the relationship between first-order languages and DRS languages. We conclude the chapter with two simple Prolog implementations: a model checker for DRSs, and a DRS to first-order logic compiler. By the end of the chapter it will be clear that DRT offers an entire architecture for thinking about semantics, an architecture that blends formal, empirical, and computational ideas in a flexible and suggestive way.

1.1 An Overview of DRT

In this section we introduce a number of key concepts of DRT. Much of our discussion revolves around the representation language DRT employs, a language based on box-like structures called DRSs. We will be making heavy use of DRSs in subsequent chapters, for many different purposes, so it is important that the reader has a good grasp of what they are and why they are so useful. The present section is largely based around the following intuition:

DRSs are Pictures.

That is, here we shall present DRT from a *representational* perspective, encouraging the reader to view DRSs as something like mental models constructed during the process of discourse comprehension. This isn't the only way of thinking about DRSs; there is also a natural *dynamic* perspective which insists that

DRSs are Programs.

This perspective will be introduced in the following section.

DRT and Context Change Potential

Imagine some agent—perhaps a little robot—about to receive an important message in natural language. Unless the message is very simple, a sequence of several natural language sentences will be required, and we immediately enter into a fascinating new area: *discourse*.

In this book we adopt a very simple model of discourse: discourses will simply be sequences of natural language sentences. (Of course, we certainly *don't* assume the converse: not every sequence of natural language sentences counts as a genuine discourse.) And the question that interests us is: how can we represent the meaning of a discourse? This is difficult to answer, but one negative observation can be made immediately: whatever the meaning of a discourse is, it *isn't* simply the conjunction of the first-order representations of its individual sentences. To see this, let's assume that our little robot is programmed to build first-order representations using the kind of tools developed in Part I. Now consider the following discourse:

'Mia is a woman. She loves Vincent'.

Our robot correctly parses the first sentence and obtains the representation WOMAN(MIA). Then along comes the second sentence, and with it one of the most obvious problems of discourse representation: how to cope with pronouns. Now, as we mentioned in Chapter ??,

in certain respects free variables are like pronouns, thus a first attempt at representation might be $\text{LOVE}(x, \text{VINCENT})$, and we certainly could build this using the kinds of techniques developed in Part I. However the meaning of the (first two items of) this discourse is *not* correctly captured by the following conjunction:

$$\text{WOMAN}(\text{MIA}) \wedge \text{LOVE}(x, \text{VINCENT}).$$

This representation misses the obvious fact that the pronoun ‘she’ ‘refers back’ to Mia. To use the linguistic terminology, ‘she’ is an *anaphoric* pronoun that needs to be *resolved*; we need to link it to a suitable discourse *antecedent* which will tell us what this pronoun is meant to refer to.

What are we to do? Well, perhaps we could simply bite the bullet and say that this free-variable-containing conjunction really *is* the underlying representation, and that pronoun resolution is essentially a post-processing (or transformational) step that converts this crude, unresolved, formula into the final representation. That is, perhaps we should simply form the above conjunction and then find a strategy for converting it into

$$\text{WOMAN}(\text{MIA}) \wedge \text{LOVE}(\text{MIA}, \text{VINCENT}).$$

This strategy is just about plausible for simple discourses like the previous one, but it soon collapses. Consider instead the following discourse:

‘A woman snorts. She collapses’.

Our naive first-order approach would build the following representation:

$$\exists z(\text{WOMAN}(z) \wedge \text{SNORT}(z)) \wedge \text{COLLAPSE}(x).$$

Again, this representation contains a free variable x , so our post-processing will need to do something about this. The obvious strategy is to try converting it into

$$\exists z(\text{WOMAN}(z) \wedge \text{SNORT}(z) \wedge \text{COLLAPSE}(z)).$$

This at least gets the truth conditions right—but there are two problems. For a start, our post-processing step is beginning to look rather complex: sometimes it has to substitute a constant (such as MIA) for a free variable, while sometimes it has to play around with quantifier scopes. It isn’t particularly plausible that transforming first-order representations in this way is compatible with the demands of good grammar engineering. As we saw in Chapter ?? (where we developed the lambda calculus after realizing that naive computational approaches to noun phrase representations led to trouble) and Chapter ??

(where we examined a number of increasingly sophisticated approaches to quantifier scope) the key to good grammar engineering is to start with the right representation. The two (extremely simply) discourses we have examined make it clear that merging the information contained in successive utterances is a complex process, one that seems to involve both quantifier scope and free variable manipulation—and that’s simply to allow us to cope with anaphoric pronouns! Once (say) presuppositions have been added to the brew, it becomes obvious that we need to adopt a more disciplined approach to representing and merging the information delivered in the course of a discourse. First order languages simply weren’t designed with these kinds of issues in mind.

But there’s a second problem: neither the alleged underlying representation

$$\exists z(\text{WOMAN}(z) \wedge \text{SNORT}(z)) \wedge \text{COLLAPSE}(x)$$

nor the final resolved representation

$$\exists z(\text{WOMAN}(z) \wedge \text{SNORT}(z) \wedge \text{COLLAPSE}(z))$$

corresponds well with our intuitions about the way the second discourse actually *works*. Intuitively, the sentence ‘A woman snorts’ sets the agenda. It paints a simple picture: we are talking about a woman, a woman who snorts, and presumably the subsequent discourse will tell us more about her (perhaps it will give her a name). The sentence ‘She collapses’ then goes on to exploit this: the ‘she’ in a very real sense *refers* to the woman we’re talking about.

These are strong intuitions, and first-order representations simply don’t get to grips with them. The first-order representation of ‘A woman snorts’ is $\exists z(\text{WOMAN}(z) \wedge \text{SNORT}(z))$. This is truth-conditionally adequate—that is, it captures the fact that the sentence asserts that a woman exists and that she snorts—but it gives us no grip on the fact that a woman has—somehow or other—been made contextually relevant. That is, although the first-order representation handles the truth-conditional dimension of meaning well, something is missing.

This missing aspect of meaning is often called *Context Change Potential (CCP)*. When we utter a ‘A woman snorts’ we don’t simply make a claim about the world, we also *change the context in which subsequent utterances will be interpreted*. If we are to make headway in computing discourse representations, we need to take account of the context change potential of sentences right from the start. That is, instead of starting with naive representations and hoping to post-process the contextual effects back in, we need to build more sophisticated representations which mirror context change potential in a natural way. Of course, we want these more sophisticated representations to get the truth conditions right as well—but, intuitively, it should not be too difficult to ensure this. It certainly seems far more sensible to try making truth conditional omelettes out of sophisticated eggs rather than first making the omelette and then trying get the eggs back!


Summing up, we would like a semantic architecture that enables us to build representations that naturally mirror CCP, but which also allows us easy access to the truth-conditional dimension of meaning. And in fact, this is precisely what Discourse Representation Theory offers. For our first encounter with the world of DRT, let's consider how a DRT-based robot would handle the previous discourse.

As soon as it realized that a new discourse had started, our robot would open a DRS, or box. To keep things simple, let's assume the robot is programmed to open the *empty box*, that is, the box containing no prior information:



In contrast, to the first-order robot, designed to operate on a “sentence-by-sentence first-order representations + post-processing” approach, the DRT robot, right from the start, is programmed to take an inter-sentential perspective on the discourse: it will progressively expand this empty box with information from the entire discourse. From the DRT perspective, discourse processing is all about carefully filling this initial box, observing various constraints while doing so. Moreover, the box filling process is highly intuitive: it builds a sort of picture (or perhaps, a mental model) of the incoming information. This picture records how the initial context (represented by the empty box) is changed in the course of the discourse.

When the ‘A woman snorts’ arrives, the robot parses it, and uses the parse tree to guide the way it expands the empty DRS. (We'll explain *how* the parse tree guides DRS construction when we discuss the standard construction algorithm; for the time being we'll ignore the underlying details.) The result of this process is the following DRS:



The x in the top compartment of this box is a *discourse referent*. The expressions $\text{WOMAN}(x)$ and $\text{SNORT}(x)$ in the bottom compartment are *conditions*. The discourse referent x is placed in the top compartment by the NP ‘a woman’; in fact, the primary function of NPs in DRT is to introduce new discourse referents. The conditions $\text{WOMAN}(x)$ and $\text{SNORT}(x)$ were placed in the bottom compartment by the NP ‘a woman’ and the VP ‘snorts’ respectively. The reader should think of the top half of this DRS as the domain of a little model (the discourse model, if you like) and the discourse referents that sit there as entities we introduce to help us record the way the context changes as the information streams in; they are potentially salient objects, available for later anaphoric reference. The conditions

in the bottom half of the picture record what we learn about these entities and their inter-relationships as the discourse proceeds. The following metaphor is quite useful: think of the discourse referents as *pegs*, and the conditions as coats we hang on these pegs as best we can.

But to return to our example: what happens when our robot then receives the sentence ‘She snorts’? The robot parses it and does three things. First, it adds a new discourse referent (y , say) to the top part of the DRS; it does this because ‘she’ is an NP, and the primary function of NPs is to introduce new discourse referents. Second, it adds the condition $\text{COLLAPSE}(y)$; it is the VP that contributes this condition. Thirdly, and most interestingly, it adds a further condition, $x=y$. That is, after the second sentence has been processed, we have the following DRS:

(1) A woman snorts. She collapses.

x y
$\text{WOMAN}(x)$ $\text{SNORT}(x)$ $\text{COLLAPSE}(y)$ $y=x$

It should be intuitively clear that the condition $y=x$ is doing something very sensible: it identifies the collapsing entity with the snorting woman introduced in the first sentence, thus correctly resolving the pronoun ‘she’. But *why* did the robot do this? Well, ‘she’ is not only an NP, it is a pronoun. Let us assume our robot is programmed to treat all pronouns as anaphoric. In DRT, anaphoric pronouns are handled as followed: the discourse referent introduced by the pronoun must be identified with an *accessible* discourse referent. Accessibility is a key concept in DRT and we’ll be working with it a lot in this book—but for now it’s enough to say that there is only one other discourse referent in our example DRS (namely x), and as a matter of fact x *is* accessible (we’ll see why later), thus in this example the robot is *forced* to add the condition $x=y$. In short, our robot resolves pronouns by exploiting the picture of the context it has built.

The idea of building pictures of the changing context by introducing discourse referents and stating constraints on them is fundamental to DRT. For example, had the first sentence been ‘Mia snorts’, the robot would have built the following DRS:

x
$x=\text{MIA}$ $\text{SNORT}(x)$

That is, the DRT perspective is that this sentence introduces a new discourse referent (namely x) and simultaneously records that this new ‘peg’ has to be identified with the entity named Mia. If you compare this example with the previous one, you will see that

DRT handles proper names and existentially quantified NPs in essentially the same way: both introduce a discourse referent which is then constrained by adding conditions. Moreover, note that there is an even closer parallelism between the treatment of anaphoric pronouns and proper names: both introduce new discourse referents, and both constrain them *equationally* (that is, using the equality symbol).

Summing up, DRT offers a natural architecture for thinking about the way information is accumulated in the course of discourse processing: essentially it allows us to draw pictures of the changing context. DRSs distinguish two types of information: information about which discourse entities we have at our disposal (recorded in the top compartment of DRSs) and information about the properties these entities have and the way they are interrelated (recorded in the bottom). As we have already seen, such pictures give us a natural way of thinking about anaphoric dependencies—but the DRT perspective turns out to be useful for a lot more besides: because it focuses on modeling how the context is transformed as the discourse proceeds, DRT is a good framework for thinking about many devices in natural language (such as presupposition, ellipsis, tense, and temporal reference) which give natural language discourses their rich intersentential structure. Switching to DRSs as the basic building blocks of our semantic representations will enable us to take some important first steps in discourse processing.

DRS Languages

We have just given our first examples of DRSs, and encouraged the reader to think of them as pictures—but clearly the simple DRSs we have seen so far can't cope with such obvious features of natural language as universal quantification and negation. How are we to deal with these? In fact, there are entire *languages* of DRSs: DRSs can be combined using various connectives and thus nested one inside another; these nested pictures offer us all representational power we would expect. So before going any further, let's be precise about the facilities DRS languages offer us.

Discourse Representation Structure languages (or *DRS languages*, or *box languages*) share many of the ingredients of first-order languages. Like first-order languages, they are built over *vocabularies* (see Chapter ??). Like first-order languages, they contain the symbols \neg , \vee , \rightarrow (though they normally don't contain \wedge), and like first-order languages with equality, they contain a special symbol $=$. In addition, DRS languages contain the symbols x , y , z , and so on, though these are called *discourse referents*, not variables. Nonetheless, there are important differences. Most importantly, DRS languages *don't* contain the symbols \forall or \exists . Instead, they approach quantification via the idea of *Discourse Representation Structures* (*DRSs*) or *boxes*.

DRSs are pairs consisting of a finite set of discourse referents and a finite set of conditions. Some of these conditions are primitive—these are the only kind of conditions we have seen so far—but it is also possible to build complex conditions; these are built out of boxes using

the connectives \neg , \vee , and \Rightarrow . As boxes may contain complex conditions, and as complex conditions are themselves defined in terms of boxes, boxes and conditions must be defined by mutual induction; let us give this crucial definition right away.

Suppose we have chosen some vocabulary of interest. We assume that the vocabulary contains no function symbols (this is simply because we won't be needing function symbols in what follows; it's easy to add them if they're wanted), thus a *term* τ is either a constant (if the vocabulary contains constants, which for natural language applications it normally would) or a discourse referent. We then build DRSs and conditions over our chosen vocabulary as follows:

1. If x_1, \dots, x_n are discourse referents (here $n \geq 0$) and $\gamma_1, \dots, \gamma_m$ (here $m \geq 0$) are conditions then

x_1, \dots, x_n
γ_1
\cdot
\cdot
\cdot
γ_m

is a DRS.

2. If R is a relation symbol of arity n , and τ_1, \dots, τ_n are terms, then $R(\tau_1, \dots, \tau_n)$ is a condition.
3. If τ_1 and τ_2 are terms then $\tau_1 = \tau_2$ is a condition.
4. If B is a DRS, then $\neg B$ is a condition.
5. If B_1 and B_2 are DRSs, then $B_1 \vee B_2$ is a condition.
6. If B_1 and B_2 are DRSs, then $B_1 \Rightarrow B_2$ is a condition.
7. Nothing else is a DRS or a condition.

Clause 1 specifies that a DRS is a structure consisting of two components. The top component (which may be empty) consists of a set of discourse referents; we call this component the *universe*. The bottom component (which may also be empty) consists of a set of conditions. The DRS

(that is, the DRS with the empty universe that contains no conditions) is called the *empty DRS*, or the *empty box*.

Conditions licensed by clauses 2 and 3 are called *primitive* conditions; they are obvious analogs of first-order atomic formulas, and we have already seen examples of them. Conditions licensed by clauses 4, 5 and 6 are called *complex* conditions. As the reader probably suspects, \neg handles negative information and \vee handles disjunctive information. The \Rightarrow construct handles both conditional and universally quantified information. Note that any DRS will have an outermost box (we call this the *main DRS*) and may in addition have a number of *sub-DRSs*; these sub-DRSs (if any) are the DRSs used to build complex conditions.

The definition of DRS languages may well have raised more questions than it answers; there's no denying that—at least at first glance—its two-dimensional box-based format is somewhat unusual. However, as will become increasingly apparent, DRS languages are extremely natural. For a start, the intuition that DRSs are pictures is a robust one; it extends in a natural way to DRSs containing complex conditions, and gives us a good way of defining the semantics of DRSs. Furthermore, the way boxes are nested one inside another gives rise to a geometrical notion called *accessibility*, and this will play an important role in later work.

Boxes as Pictures

When is a DRS satisfied? Given that a DRS is supposed to be a picture, it seems natural to say that a DRS is satisfied in a model if and only if it is an accurate image of the information recorded inside the model. For example, the DRS


x y
WOMAN(x)
BOXER(y)
ADMIRE(x,y)

should be satisfied with respect to a model if and only if it is possible to associate the discourse referents x and y with a pair of entities in the model such that the first entity is a woman, the second is a boxer, and the first stands in the admires relation to the second. If we can associate the discourse referents with such a pair of entities, it is as if the little picture the DRS gives us is “embedded” inside the (possibly very large) model, and it is natural to regard such a picture as satisfied.

In fact, the idea of associating discourse referents with model-theoretic entities, thereby embedding DRS in models, is the cornerstone of the embedding semantics we shall develop in detail in the following section. But the reader may be skeptical. For sure, thinking in

terms of pictures is a natural way of viewing boxes that contain only basic conditions—but does it help us with the semantics of complex conditions?

Actually, it does. Roughly speaking, a negated DRS will be satisfied if it is *not* possible to find the picture the DRS gives us embedded inside the model, and a pair of DRSs joined by a disjunction symbol will be satisfied if at least one of these pictures can be found embedded inside the model. But what about conditions of the form $B_1 \Rightarrow B_2$? This is nice: such conditions will be satisfied if every way of embedding the antecedent picture B_1 gives rise to an embedding of the consequent picture B_2 . For example, the DRS



which captures the narrow scope reading of the sentence ‘Every woman loves a boxer’, will be satisfied if no matter what entity we choose to associate with the discourse referent y (for example, let’s suppose we choose an entity w), we will always be able associate x with an entity (let’s call it b) such that w loves b . In short, no matter which entities we use to embed the antecedent picture, we will be able to embed the consequent picture too.

Thus DRS languages can be given a semantics that is faithful to the pictorial intuitions we have been emphasizing, and this is one good reason for being interested in them. But there is another: box syntax gives rise to the important notion of *accessibility*.

Accessibility

Our little DRT robot resolved the anaphoric pronoun ‘she’ by adding an equality condition to the picture it was building. Equality conditions are the mechanism used in DRT to resolve anaphors, and their use is subject to an interesting constraint: if y is the discourse referent introduced by a pronoun, and x is a previously introduced discourse referent, then we are only allowed to add the condition $y=x$ if x is in the universe of a DRS that is *accessible* from the DRS whose universe contains y . This accessibility constraint not only grounds many of DRT’s empirical claims about anaphora, it will also play an important role in our later discussion of presupposition.

What is accessibility? Actually, it’s a simple geometric concept, defined in terms of the way DRSs are nested one inside another. Here’s the definition. DRS B_1 is accessible from DRS B_2 when B_1 *equals* B_2 or when B_1 *subordinates* B_2 . The subordinates relation is defined as follows: B_1 subordinates B_2 if and only if

1. B_1 *immediately subordinates* B_2 ; or
2. There is some DRS B such that B_1 *subordinates* B and B *subordinates* B_2 .

That is, the subordinates relation is the transitive closure of the immediately subordinates relation. So to complete the definition we need only stipulate how immediate subordination is defined. We say: B_1 immediately subordinates B_2 if and only if

1. B_1 contains a condition of the form $\neg B_2$; or
2. B_1 contains a condition of the form $B_2 \vee B$ or $B \vee B_2$, for some DRS B ; or
3. B_1 contains a condition of the form $B_2 \Rightarrow B$, for some DRS B ; or
4. $B_1 \Rightarrow B_2$ is a condition in some DRS B .

Exercise 1.1.1 Draw the configurations of DRSS listed in the definition of immediate subordination. Suppose that B is a DRS and that B_s is a sub-DRS of B (that is, B_s is nested, perhaps very deeply, somewhere inside B). Prove that B accessible to B_s .

Now, the previous definition tells us what it means for one DRS to be accessible from another—but we really want accessibility to be a relation between *occurrences of discourse referents in universes*. However it is now straightforward to define this concept: we say that an occurrence of a discourse referent (say x) in the universe of some box (say B_1) is accessible from an occurrence of a discourse referent (say y) in the universe of some box (say B_2) if and only if B_1 is accessible from B_2 . And given this concept we can now state the accessibility constraint on pronoun resolution more precisely:

Suppose a pronoun has introduced a new discourse referent (say y) into the universe of a DRS B_2 . Then we are only free to add the condition $y=x$ to the condition set of B_2 if some occurrence of x in some universe is accessible from the occurrence of y in the universe of B_2 .

This can be simplified. When using DRSS to analyze natural language, we usually end up working with DRSS in which each discourse referent occurs in exactly one universe. When this is the case we don't need to bother talking about occurrences of discourse referents in universes—in effect we can identify a discourse referent with unique universe it occurs in, and simply talk of one discourse referent being accessible from another. This enables us to restate the accessibility constraint as follows:

Suppose a pronoun has introduced a new discourse referent (say y) into the universe of some DRS B . Then we are only free to add the condition $y=x$ to the condition set of B if x is accessible from y .

We'll usually talk about accessibility in this simpler way in what follows.

Let's see some examples of the accessibility constraint at work. First, it enables DRT to correctly predict that an anaphoric link is permitted between 'she' and 'a woman' in the first little discourse that follows, but not in the second:

(2) A woman snorts. She collapses.	x y
	WOMAN(x)
	SNORT(x)
	COLLAPSE(y)
	y=x

Here we are free to add the constraint $y=x$, for the (unique) occurrences of x and y belong to the same universe, and thus x is accessible to y . On the other hand, consider the following two sentence sequence and its (only partially completed) DRS representation:

(3) Every woman snorts. ? She collapses.	y	
	x	
	WOMAN(x)	SNORT(x)
	COLLAPSE(y)	
	y=?	

In this example there is no accessible discourse referent for y —the only candidate is x , and x is not accessible from y . Hence we cannot complete the representation, and thus (correctly) predict that the pronoun 'she' does not have an anaphoric interpretation. Here are some other examples for which DRT correctly predicts felicity or infelicity on the basis of accessibility arguments:

- (4) Mia ordered a five dollar shake. Vincent tasted it.
- (5) Mia didn't order a five dollar shake. * Vincent tasted it.
- (6) Butch stole a chopper. It belonged to Zed.
- (7) Butch stole a chopper or a motor cycle. ? It belonged to Zed.
- (8) Butch stole a chopper or a motor cycle. * The chopper belonged to Zed.

Exercise 1.1.2 Use accessibility arguments to explain why anaphoric reference is or is not possible in these examples.

To conclude this discussion, let's consider how DRT represents the infamous "donkey sentences"; for example, 'If a man eats a big Kahuna burger, he enjoys it'. (Donkey sentences are so-called because the standard examples are 'If a farmer owns a donkey, he beats it' and 'Every farmer who owns a donkey beats it'.)

Why are such sentences considered difficult? Because it is not clear how the required semantic representations can be built compositionally using the traditional tools of Montague grammar (roughly speaking, the λ -based approach to semantic construction introduced in Chapter ??). Now, it's no problem finding a first-order formula which captures the truth conditions of this sentence; the following formula clearly gets it right:

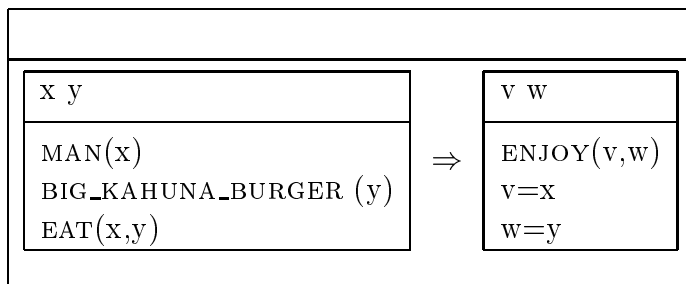
$$\forall x \forall y [\text{MAN}(x) \wedge \text{BIG_KAHUNA_BURGER}(y) \wedge \text{EAT}(x,y) \rightarrow \text{ENJOY}(x,y)].$$

But how can we build this representation? For a start, how have the *existential* NPs 'a man' and 'a big Kahuna burger' given rise to the *universal* quantifiers $\forall x$ and $\forall y$? And how did these two quantifiers come to take wide scope over the \Rightarrow contributed by the 'If...' that glues the antecedent and consequent sentences together? If you attempt to build first-order representations using the methods of Chapter ??, it's quite likely that you'll wind up generating:

$$\exists x [\text{MAN}(x) \wedge \exists y [\text{BIG_KAHUNA_BURGER}(y) \wedge \text{EAT}(x,y)] \rightarrow \text{ENJOY}(x,y)].$$

Exercise 1.1.3 Under the standard semantics of first-order logic, this 'representation' does not correctly capture the meaning of the sentence. Explain why.

In DRT this sentence would be represented as follows:



This looks more promising. For a start, note that outermost level of structure is a box containing a condition of the form $K_1 \Rightarrow K_2$, which clearly mirrors the fact that this sentence is a conditional. Moreover, note that the effect of each of the existential NPs

in the antecedent (that is, ‘a man’ and ‘a big Kahuna burger’) is to contribute a discourse referent (namely x and y) just as they are supposed to. Further, the pronouns ‘he’ and ‘it’ in the consequent do this too (they introduce v and w), and the linkage between the pronouns and their antecedents is neatly captured by the two equalities. Moreover, this representation really does capture the meaning of the donkey sentence: it should be clear that this DRS will be satisfied if and only if whenever we find a little picture in which a man is eating a big Kahuna burger, we can extend this to a picture in which that same man is enjoying that same Kahuna burger. Furthermore, we will soon have a DRS to first-order logic compiler at our disposal, and this compiler will translate the above DRS to:

$$\forall x \forall y [\text{MAN}(x) \wedge \text{BIG_K_BURGER}(y) \wedge \text{EAT}(x, y) \rightarrow \exists v \exists w (\text{ENJOY}(v, w) \wedge v = x \wedge w = y)],$$


which is easily seen to be equivalent to the earlier first-order representation.

So the only question that remains to be answered is: can this DRS be built systematically? The answer is “yes”—indeed, from the DRT perspective there really isn’t anything particularly special about donkey sentences at all; they are handled using exactly the same construction methods that handle other sentence types. So the really important question is not “How do we handle donkey sentences in DRT?” but “How do we set about building DRSs at all?”. The following chapter is entirely devoted to this question, but it’s high time we had a first glimpse of the issues involved.


The Standard Construction Algorithm

In this section we’ll show, step-by-step, how the DRS corresponding to the discourse ‘A woman snorts. She collapses’ can be constructed using a simple top-down strategy. This strategy, which we call the *standard construction algorithm*, is the account of DRS construction the reader is most likely to meet in other presentations of DRT.


First we start feeding information from the initial sentence into a box:




We proceed as follows: we are going to work our way around the tree in a top-down, left-to-right fashion, adding discourse referents as we go. So first we work down from the S node to the NP node. We then peek inside the NP : what sort of NP is it? In fact, it’s an indefinite noun phrase; indefinites introduce a new discourse referent into the upper part of the box, so we introduce the discourse referent x :




What conditions attach to this discourse referent? To find out, we move around the NP subtree to the N node. We find the word ‘woman’, thus we constrain x by putting the condition $WOMAN(x)$ in the lower part of the box (Step 3).




We’ve now explored the entire NP subtree; there’s nothing left to do here. So it’s time to move back up to S and then down into the VP subtree. Here we encounter the verb phrase that consists of the intransitive verb ‘snorts’. The standard algorithm has kept track of the fact that we’ve basically moved from the subject position of the tree to the predicate position, so it instructs us to view as a further constraint on x . Accordingly we add the condition $SNORT(x)$ to the DRS:




This is the DRS for the first sentence of discourse (2). Let’s continue straight on with the second sentence. As with the first sentence, we will explore its parse tree in a top-down, left-to-right fashion. Crucially, the standard construction algorithm insists that we add all the information we obtain from the second sentence straight into the DRS we have already built using the first sentence (after all, the main DRS is meant to be a picture of the entire discourse). This will make it very easy to perform the required anaphora resolution.



Again we work our way down from S, to NP, and then peek to see what kind of NP we are dealing with. It's a pronoun, so this means we must introduce a new discourse referent (we have chosen y) and a condition of the form $y=?$, where the question mark will be identified with an accessible, earlier introduced, discourse referent. (Actually, the standard algorithm insists on a further constraint: the discourse referent we use must not only be accessible, it must also be *suitable*—but we'll defer our discussion of suitability till Chapter 9.) There is only one accessible discourse referent available, namely the discourse referent x , which was introduced by 'a woman', so we are forced to build the following DRS:



This finishes the NP subtree. We then work our way round to the VP subtree and handle it in essentially the same way as did in the first sentence, so this is the final DRS we obtain:



This DRS captures the truth conditions of the discourse appropriately. Two entities are introduced, x and y , where x possesses the properties of being a woman and snorting, and y the property of collapsing, and these two entities are identified, just as the anaphoric link between the pronoun 'she' and 'a woman' would lead us to expect.

The standard construction algorithm is simple and easy to understand. But a worrying thought may have occurred to some readers: is this the *only* way we have of constructing DRSs? If so, then DRT isn't very computationally promising. The heart of computational semantics is developing flexible semantic construction strategies; we simply can't afford to get locked into a top-down (or indeed, any other) strait-jacket.

In fact, there's no problem: DRT turns out to be compatible with a wide range of implementation strategies. Nor is the standard construction algorithm a dead end: on the contrary, appropriately viewed it provides the basis for an efficient approach to DRS construction known as threading. These issues will be explored in detail in the following chapter.

1.2 Interpreting DRSs

It's time to be precise about the interpretation of DRSs. There are two popular ways of doing this: one is to use embedding semantics, which we have already briefly discussed, the other is to use dynamic semantics. These interpretations are equivalent—at least for the ‘core’ language of DRT we are discussing here—but both are worth getting to know, for each emphasizes different intuitions about what DRSs are.

Embedding Semantics

Suppose we are working with some DRS language, and that $\mathbf{M} = (D, F)$ is a model for that language (that is, the language and the model have the same vocabulary). An *embedding* is a partial function from the set of discourse referents to D , the domain of \mathbf{M} . Note that as they are *partial* functions, embeddings need not associate *all* discourse referents with elements of the model. Indeed there is a special embedding, the *empty embedding*, which does not associate *any* discourse referents with elements of the model.

Suppose i and j are embeddings in some model \mathbf{M} . We say that j *extends* i if and only if whenever $i(x)$ is defined for some discourse referent x , then $j(x)$ is defined too, and $j(x) = i(x)$. In set-theoretic terms this simply means that viewed as sets of ordered pairs, $i \subseteq j$. It is also useful to think of this concept in information-oriented terms: intuitively, j extends i means that j contains at least as much information as i . Note that every embedding extends itself, and that every embedding extends the empty embedding (think of the empty embedding as carrying no information).

We can now say what it means for an i embedding to satisfy a DRS B in a model \mathbf{M} :

$$\mathbf{M}, i \models \begin{array}{|c|} \hline x_1, \dots, x_n \\ \hline K_1 \\ \cdot \\ \cdot \\ \cdot \\ K_m \\ \hline \end{array} \quad \text{iff} \quad \begin{array}{l} \text{there is an embedding } j \text{ such that } j \text{ extends } i, \text{ and} \\ j(x_1), \dots, j(x_n) \text{ are all defined, and} \\ j \text{ satisfies conditions } K_1, \dots, K_m \text{ in } \mathbf{M}. \end{array}$$

This definition captures the intuition of “finding the picture inside the model” in a fairly obvious way: it demands that the embedding i matches discourse referents with elements of the model in such a way that all the conditions in the DRS are satisfied. (Note, incidentally, that this clause explains why we don't need an explicit conjunction symbol \wedge in DRT; any conditions listed in the same DRS are, in effect, conjoined.) There is only one complication: why do we shift from the original embedding i to an extended embedding j ? Because some of the conditions listed in the DRS we are interpreting (let's call it B) may be complex,

and the DRSs these complex conditions are formed from may contain discourse referents that do not belong to the universe of the original DRS; we need to be able to successfully extend embedding i to cover all the discourse referents introduced in sub-DRSs.

Now for the next step: we need to say what it means for a model \mathbf{M} and an embedding i to satisfy a condition K . For basic conditions, this is straightforward. First, for any term τ , by the interpretation of τ with respect to \mathbf{M} and i we mean $F(\tau)$ if τ is a constant, and $i(\tau)$ if τ is a discourse referent; we denote the interpretation of τ by $I_F^i(\tau)$. Note that as embeddings are partial, $i(\tau)$ may be undefined; if this is the case we say that $I_F^i(\tau)$ is undefined too. We can now state what it means for basic condition to be satisfied:

$$\mathbf{M}, i \models R(\tau_1, \dots, \tau_n) \quad \text{iff} \quad I_F^i(\tau_1), \dots, I_F^i(\tau_n) \text{ are all defined, and} \\ (I_F^i(\tau_1), \dots, I_F^i(\tau_n)) \in F(R)$$

$$\mathbf{M}, i \models \tau_1 = \tau_2 \quad \text{iff} \quad I_F^i(\tau_1) \text{ and } I_F^i(\tau_2) \text{ are defined, and} \\ I_F^i(\tau_1) = I_F^i(\tau_2)$$

Note that neither of these clauses makes use of the idea of extending embeddings: to interpret basic conditions we simply compare condition and model in the obvious way.

Now for the complex conditions. Let us say that an embedding j extends an embedding i on a box B if the arguments j is defined on are precisely the arguments i is defined on together with the discourse referents in the universe of B . Using this terminology we define:

$$\mathbf{M}, i \models \neg B \quad \text{iff} \quad \text{for all } j \text{ such that } j \text{ extends } i \text{ on } B, \\ \text{it is not the case that } \mathbf{M}, j \models B$$

$$\mathbf{M}, i \models B_1 \vee B_2 \quad \text{iff} \quad \text{there is an } j \text{ such that } j \text{ extends } i \text{ on } B_1, \text{ and} \\ \mathbf{M}, j \models B_1; \text{ or} \\ \text{there is an } k \text{ such that } k \text{ extends } i \text{ on } B_2, \text{ and} \\ \mathbf{M}, k \models B_2$$

$$\mathbf{M}, i \models B_1 \Rightarrow B_2 \quad \text{iff} \quad \text{for all } j \text{ such that } j \text{ extends } i \text{ on } B_1 \text{ and} \\ \mathbf{M}, j \models B_1, \\ \text{there is a } k \text{ such that } k \text{ extends } j \text{ on } B_2 \text{ and} \\ \mathbf{M}, k \models B_2$$

That is, a negative condition is satisfied precisely when we have painted ourselves into a corner: no matter how we extend the embedding i on the box in question—note that this may include doing nothing at all to i , since every embedding extends itself—we simply can't embed the DRS behind the negation symbol. Similarly, a disjunctive condition is satisfied precisely if it is possible to extend i so that at least one of the disjoined DRSs is

embedded. Finally, a conditional condition is satisfied if every successful embedding of the antecedent leads to a successful embedding of the consequent, just as we discussed earlier.

We are almost there. We now know what it means for a DRS to be satisfied in a model with respect to a given assignment function—but what does it mean simply to say that a DRS is satisfied in a model? The following:

A DRS B is embedding satisfied in a model \mathbf{M} if and only if there is an embedding i such that $\mathbf{M}, i \models B$.

And that’s the embedding semantics. As we promised, it makes precise the intuition that DRSs are pictures in a very direct way. Nonetheless some readers may have detected a second idea emerging in the above discussion. Many of the key definitions—including the one for boxes—made use of the idea of *extending* an embedding. There is something very procedural about this idea: indeed extending an embedding i to an embedding j sounds very like “adding information by performing a computation”. In fact, the dynamic semantics we shall now define has its roots in this observation.

Exercise 1.2.1 Both the \Rightarrow symbol and the \vee symbol are *eliminable* from languages of DRSs. That is, we can throw away *both* the \Rightarrow and \vee from the language of DRSs without losing expressive power for we can express both types of condition using boxes and negative conditions. Prove this with the help of the embedding semantics. (We return to the eliminability of \Rightarrow and \vee , from a somewhat different perspective, when we discuss the relationship of DRS languages and first-order languages; see in particular Exercise 1.3.5.)

Dynamic Semantics

The embedding semantics explains what it means for a picture to be verified. But although this tells us something about the notion of context change potential—after all, these pictures are pictures of changing context—it is based on an essentially static intuition. Dynamic semantics foregrounds the idea of change: it insists that DRSs are programs that modify context, and that only conditions should be viewed statically.

What kind of mathematical object could we use to model contexts? The traditional answer given in dynamic semantics is: *total embeddings*. Suppose we have fixed some (function symbol free) signature, and suppose that $\mathbf{M} = (D, F)$ is a model of this signature. A total embedding in \mathbf{M} is a (total) function from the set of all discourse referents to D .

A moment’s thought shows this is a reasonable notion of context. Given that our task is to explain the kind of context change potential that brings objects into salience and enables anaphoric links to be established, contexts are essentially objects that link discourse referents with model theoretic entities—and this is precisely what total embeddings are. In fact we’ve met this idea before. In Chapter 1 we noted that an assignment of values to

variables could be viewed as a context; but of course, ‘discourse referent’ is simply DRT terminology for ‘variable’, hence a total embedding is nothing but an assignment function.

There is one point which may be bothering some readers: do we really need to work with *total* embeddings? The answer is, “No, we don’t”; but doing so simplifies the technicalities, and it is pleasant to have an interpretation for DRS languages that makes use of *exactly* the same semantic machinery—models and assignment functions—as first order languages.

So how does the the dynamic put this machinery to work? First, it draws a distinction between the semantics of conditions and the semantics of DRSs. Conditions are regarded as *tests* and treated statically: $\mathbf{M}, g \models K$ will mean that K (a test) is satisfied in a model \mathbf{M} (a situation) under a total embedding g (a context). Conceptually and notationally this is like the treatment of satisfaction in first-order logic; tests don’t change contexts—their semantics is static. The dynamics comes in the treatment of DRSs: we are going to think of DRSs as non-deterministic programs which act on contexts to produce new contexts. That is, we will define what it means for a DRS B (a *program*) to be satisfied in a model \mathbf{M} (a situation) with respect to a *pair* of total embeddings g and h , (the *input context* and the *output context*); this is written $\mathbf{M}, g, h \models B$. Thus dynamic semantics will blend four ideas: programs, tests, situations, and contexts. Incidentally, *blends* is the key word here. Much of the interest of dynamic semantics lies in the harmonious way it combines a dynamic analysis of context change with traditional static ideas.

So let’s turn to the key definition. The following piece of notation will be helpful. If g and h are total embeddings on the same model we shall write $g[x_1, \dots, x_n]h$ to indicate that g differs from h , if at all, only in the values it assigns to the discourse referents x_1, \dots, x_n . (Note that this is just an n -place version of the idea of variant assignments used in the definition of first-order satisfaction in Chapter ??.) As before, we’ll let $I_F^g(\tau_n)$, the denotation of τ be $F(\tau)$ if τ is a constant, and $g(\tau)$ if τ is a discourse referent. We can now define the static and dynamic notions we need by mutual indication:

$\mathbf{M}, g \models R(\tau_1, \dots, \tau_n)$	<i>iff</i>	$(I_F^g(\tau_1), \dots, I_F^g(\tau_n)) \in F(R)$
$\mathbf{M}, g \models \tau_1 = \tau_n$	<i>iff</i>	$I_F^g(\tau_1) = I_F^g(\tau_n)$
$\mathbf{M}, g \models \neg B$	<i>iff</i>	for all h , it is not the case that $\mathbf{M}, g, h \models B$
$\mathbf{M}, g \models B_1 \vee B_2$	<i>iff</i>	there is an h such that $\mathbf{M}, g, h \models B_1$ or $\mathbf{M}, g, h \models B_2$
$\mathbf{M}, g \models B_1 \Rightarrow B_2$	<i>iff</i>	for all f such that $\mathbf{M}, g, f \models B_1$ there is an h such that $\mathbf{M}, f, h \models B_2$

$\mathbf{M}, g, h \models$	x_1, \dots, x_n K_1 \cdot \cdot \cdot K_m	<i>iff</i>	$g[x_1, \dots, x_n]h$ and $\mathbf{M}, h \models K_1 \ \& \ \dots \ \& \ \mathbf{M}, h \models K_m$
----------------------------	--	------------	--

The first two clauses simply say that basic conditions are satisfied in a situation \mathbf{M} given contextual information g precisely when the interpretations of the discourse reference guarantees that we get the answer “Yes” to these tests.

The third clause says that a condition $\neg B$ is satisfied in situation \mathbf{M} in context g precisely when the *program* B cannot be successfully executed when g is used as the input context. (Note the flip-flop, characteristic of DRS languages, from conditions to programs.) That is, when we run program B on input g , there is no state h that is output by this computation. In the fourth clause we again see the same flip-flop: a disjunctive condition is satisfied precisely when at least one of its component programs can be executed on the input context g , producing a context h as output. (That is, at least one of the component programs is capable of transforming the input context in a satisfactory way.)

The clause for implicational conditions says that an implication is satisfied in situation \mathbf{M} in context g precisely when every possible output f of this computation can be successfully acted on by the consequent program B_2 to produce output context. As a well known slogan puts it: *every successful execution of the antecedent leads to a successful execution of the consequent*. To appreciate the force of this definition the reader must recall that as we are thinking of DRSs as *non-deterministic* programs, a DRS may be able to modify the input context in various ways. This definition demands that no matter which context the DRS gives back, it must always be possible to execute the consequent program on it.

The last clause is the key one: it tells us what these programs actually do, and why they are non-deterministic. Given an input context g , a DRS must try to find interpretations

for the discourse referents in its universe so that all its conditions are satisfied. That is, boxes are programs that update the links between discourse referents and the world. Note that such programs are intrinsically non-deterministic: after all, there may be many ways of updating these links that result in all conditions being satisfied.

One task remains: defining what it means for a DRS to be satisfied in a model.

A DRS B is *dynamically satisfied* in a model \mathbf{M} with respect to assignment f if there is an assignment f' such that $\mathbf{M}, f, f' \models B$.

And that's dynamic semantics. It should be quite clear that embedding semantics and dynamic semantics are closely linked—indeed they are essentially two different perspectives on the same set of ideas. The best way to get to grips with the way these two interpretations fit together technically is to attempt the following exercise right away.

Exercise 1.2.2 Show that any DRS B can be dynamically satisfied in a model \mathbf{M} if and only if it can be embedding satisfied in model \mathbf{M} .

1.3 DRT and First-Order Logic

In spite of their differences, DRT and first-order logic are obviously related. Given any vocabulary without function symbols, we can use it to build either a DRS language or a first-order language. Moreover, both languages will be interpreted on exactly the same models. Thus it is natural to try to relate DRS and first-order languages via satisfaction preserving translations. Such translations essentially tell us how to compile one language down into the other in a sensible way.

As we shall now show, any DRS language and its corresponding first-order language are intertranslatable: either can be compiled down into the other, and the translations involved are straightforward and efficient.

From DRT to First-Order Logic

We first show how to translate DRSs into formulas of first-order logic with equality. Suppose we are working with a vocabulary that contains no function symbols. We shall define a translation function fo which takes any DRS built over this vocabulary and maps it to a formula of the first-order language (with equality) built over the same vocabulary. This function works by recursively mapping the DRSs and conditions that make up the input DRS to first-order formulas. It is easy to understand and (as we shall see in the following section) easy to implement in Prolog.

First we need to know how to translate boxes. Here is the general translation schema:

$$\left(\begin{array}{|l} \mathbf{x}_1, \dots, \mathbf{x}_n \\ \hline \gamma_1 \\ \cdot \\ \cdot \\ \gamma_m \end{array} \right)^{fo} = \exists \mathbf{x}_1 \dots \exists \mathbf{x}_n ((\gamma_1)^{fo} \wedge \dots \wedge (\gamma_m)^{fo})$$

That is, this clause maps the discourse referents in the universe of the box to existentially quantified variables, and then recursively translates the conditions. The basic idea should be clear, but to avoid misunderstanding it is worth spelling out what this schema means in the following special cases. First, if there is only one condition γ_1 in the condition set, then the translation is $\exists \mathbf{x}_1 \dots \exists \mathbf{x}_n (\gamma_1)^{fo}$. Second, if there are no conditions in the condition set, then the translation is $\exists \mathbf{x}_1 \dots \exists \mathbf{x}_n \top$ (recall that \top stands for *true*). Third, if the DRS we are translating has an empty universe, then the translation is $(\gamma_1)^{fo} \wedge \dots \wedge (\gamma_m)^{fo}$. It follows from these conventions that the translation of the empty DRS is \top .

Now for the conditions. The treatment of the basic conditions couldn't be simpler; they simply map to themselves-viewed-as-first-order-atomic-formulas:

$$\begin{aligned} (R(\mathbf{x}_1, \dots, \mathbf{x}_n))^{fo} &= R(\mathbf{x}_1, \dots, \mathbf{x}_n) \\ (\tau_1 = \tau_2)^{fo} &= \tau_1 = \tau_2 \end{aligned}$$

Moreover, complex conditions formed using \neg and \vee are also straightforwardly handled; we simply push the translation function in over the connective, leaving the connective unchanged:

$$\begin{aligned} (\neg B)^{fo} &= \neg(B)^{fo} \\ (B_1 \vee B_2)^{fo} &= (B_1)^{fo} \vee (B_2)^{fo} \end{aligned}$$

Now for complex conditions formed using \Rightarrow . Here's the general schema:

$$\left(\begin{array}{|l} \mathbf{x}_1, \dots, \mathbf{x}_n \\ \hline \gamma_1 \\ \cdot \\ \cdot \\ \gamma_m \end{array} \right) \Rightarrow B)^{fo} = \forall \mathbf{x}_1 \dots \forall \mathbf{x}_n (((\gamma_1)^{fo} \wedge \dots \wedge (\gamma_m)^{fo}) \rightarrow (B)^{fo})$$

This clause clearly captures the idea that every possible embedding of the antecedent DRS should lead to a embedding of the conclusion. However, as with the translation of boxes,

it is a good idea to be explicit about what this schema means in the following special cases. First, if there is only one condition γ_1 in the condition set of the antecedent, then the translation is $\forall x_1 \cdots \forall x_n ((\gamma_1)^{fo} \rightarrow (B)^{fo})$. Second, if there are no conditions in the condition set of the antecedent, then the translation is $\forall x_1 \cdots \forall x_n (\top \rightarrow (B)^{fo})$, which is logically equivalent to $\forall x_1 \cdots \forall x_n (B)^{fo}$. Third, when the antecedent DRS has an empty universe, the translation is $(\gamma_1)^{fo} \wedge \cdots \wedge (\gamma_m)^{fo} \rightarrow B^{fo}$. It follows from these conventions that if the antecedent DRS is the empty DRS, we obtain $\top \rightarrow (B)^{fo}$, which is logically equivalent to $(B)^{fo}$, as the translation.

This translation is not only semantically sensible, it is also *efficient*: the number of symbols in the first-order formula it returns as output is of the same order of magnitude as the number of symbols in the DRS. This means that we have an easy way of turning DRT inference problems into first-order inference problems, a possibility we shall exploit heavily in later chapters.

Exercise 1.3.1 Use *fo* to translate the DRS representing if a man eats a big Kahuna burger, he enjoys it into first-order logic with equality.

Exercise 1.3.2 We claimed that *fo* was a satisfaction preserving translation. Show that this claim can be made precise in two ways. Let B be a DRS. First, show that B is *dynamically* satisfied in a model \mathbf{M} if and only if $\mathbf{M} \models (B)^{fo}$. Second, show that B is *embedding* satisfied in \mathbf{M} if and only if $\mathbf{M} \models (B)^{fo}$.

From First-Order Logic to DRT

We now go in the reverse direction: we shall show how given any vocabulary at all (even one containing function symbols) we can map first-order formulas over this vocabulary to a DRS built over the same vocabulary (minus the function symbols) in a satisfaction preserving way. We shall present a *lean* version of the translation; that is, one that does not operate directly on all first-order connectives. The exercises below ask the reader to fatten it up, and also show why lean can be useful.

The translation requires a preprocessing stage, which performs three functions. First, it rewrites the formula to an equivalent formula that contains no function symbols. (We saw how to do this in Exercise ??.) Having done this, it rewrites this function-free formula to an equivalent formula that contains only the connectives \neg , \vee and \exists . (Again this is possible; see Exercise ??.) Finally, the resulting formula is rewritten to an equivalent formula in which distinct occurrences of \exists bind distinct variables, an easy (but important) step. We are only going to translate the formulas that are output by this three-step preprocessing stage.

Given a preprocessed first-order formula ϕ , how do we translate it? First we form the following structure:

$(\phi)^{dr}$

That is, we take an empty box, and put the expression $(\phi)^{dr}$ in the compartment reserved for conditions. The superscripted dr is our translation function, and we shall simply keep recursively applying this function; as we do so we will fill the box with conditions and discourse referents until we finish building the DRS we require.

So, what does dr let us do? Obviously we need to specify a translation step for each kind of first-order formula we could encounter, and as ϕ has been preprocessed this means atomic formulas, conjunctions, negations, and existential quantifications. Let's go through these possibilities in turn, starting with the atomic formulas.

Suppose we see a box of the following form somewhere in the DRS we are building (it may be the main box itself, or a sub-box):

<i>Discourse Referents</i>
<i>Conditions</i> $(\tau_1 = \tau_2)^{dr}$ <i>Conditions</i>

(The notation is meant to emphasize that the important thing is that we see a box containing $(\tau_1 = \tau_2)^{dr}$: it *doesn't* matter if this box contains other conditions, either before or after the $(\tau_1 = \tau_2)^{dr}$, and it *doesn't* matter which—if any—discourse referents this box has in its universe.) If we find such a box, it should be clear how to handle it: we simply drop the brackets. Doing so yields a legitimate DRS condition and gives us the following box:

<i>Discourse Referents</i>
<i>Conditions</i> $\tau_1 = \tau_2$ <i>Conditions</i>

In a similar fashion, if we see a box containing the expression $(R(x_1, \dots, x_n))^{dr}$, then we simplify this to $R(x_1, \dots, x_n)$ and carry on.

So what happens if we see a box containing a conjunction that needs to be translated? That is, what do we do if we find something of the following form:

<i>Discourse Referents</i>
<i>Conditions</i> $(\phi \wedge \psi)^{dr}$ <i>Conditions</i>

Fairly obviously, we simply break this expression in two, forming the following box:

<i>Discourse Referents</i>
<i>Conditions</i> $(\phi)^{dr}$ $(\psi)^{dr}$ <i>Conditions</i>

So let's consider the clause for negation, which is rather more interesting. Suppose we find a box of the following form:

<i>Discourse Referents</i>
<i>Conditions</i> $(\neg\phi)^{dr}$ <i>Conditions</i>

This triggers sub-DRS construction—we form a negative condition by opening up a negated DRS:

<i>Discourse Referents</i>		
<i>Conditions</i> <table border="1" style="margin-left: 20px;"> <tr> <td style="text-align: center;"> \neg <table border="1" style="margin-left: 20px;"> <tr> <td style="text-align: center;">$(\phi)^{dr}$</td> </tr> </table> </td> </tr> </table>	\neg <table border="1" style="margin-left: 20px;"> <tr> <td style="text-align: center;">$(\phi)^{dr}$</td> </tr> </table>	$(\phi)^{dr}$
\neg <table border="1" style="margin-left: 20px;"> <tr> <td style="text-align: center;">$(\phi)^{dr}$</td> </tr> </table>	$(\phi)^{dr}$	
$(\phi)^{dr}$		
<i>Conditions</i>		

Finally, suppose we encounter an existentially quantified expression that needs translating:

<i>Discourse Referents</i>
<i>Conditions</i> $(\exists x\phi)^{dr}$ <i>Conditions</i>

We simply add the variable x to the universe of the DRS and form:

x , <i>Discourse Referents</i>
<i>Conditions</i> $(\phi)^{dr}$ <i>Conditions</i>

Note that this step exploits our preprocessing: we know that each quantifier bind a *unique* variable, so, given the inductive nature of our construction, x is guaranteed to be distinct from any of the discourse referents already in this box's universe.

And that's the translation process—we simply start at the top and carry out this simplification process until we bottom out in the atomic formulas; it easy to show that this process terminates yielding a legitimate DRS. Note that this translation is *non-deterministic*: we are allowed to simplify in any order we like, wherever we see an appropriate configuration. It should also be clear that this translation is just begging to be implemented in Prolog, and after we have discussed how to represent DRS in Prolog we shall get the reader to do this (see Exercise 1.4.4).

Between them, *fo* and *dr* unambiguously show that—*regarded simply as tools for talking about models*—DRS languages and first-order languages are expressively equivalent. This strongly underlines the point made at the start of the chapter: the move to DRT does *not* reflect dissatisfaction with first-order logic as a tool for representing truth conditions. First-order logic is excellent at this task (it's not called classical logic for nothing) and DRS languages add *nothing* on this level. Rather, the gains from DRS languages reflect a completely different set of ideas. Although they offer the same model-theoretic expressivity as first order languages, they package this expressivity differently. As both embedding and dynamic semantics in their different ways show, this repackaging allows us to model the idea of context change potential in a straightforward and natural way. DRS languages and first-order languages don't differ in what they say about models; they differ in what they teach us about context.

Exercise 1.3.3 Preprocess $\forall x(\text{WOMAN}(x) \rightarrow \text{SNORT}(x))$ and use *dr* to translate it into a DRS.

Exercise 1.3.4 As the previous exercise shows, a lean translation has its unpleasant aspects: it requires us to eliminate symbols before we can use it. Fatten *dr* up by adding clauses that operate directly on \vee , \rightarrow , and \forall formulas, and then retranslate $\forall x(\text{WOMAN}(x) \rightarrow \text{SNORT}(x))$.

Exercise 1.3.5 But lean translations can be useful. For example, using the fact that *fo* and *dr* are both satisfaction preserving, it is easy to show that the symbols \vee and \Rightarrow can be eliminated from DRS languages without losing expressivity. Indeed we can even pull out explicit paraphrases of \vee and \Rightarrow conditions. Why is this? [Hint: suppose a DRS B contains occurrences of \Rightarrow or \vee . What do you get if you first apply *fo* and then apply *dr* to the result?]

1.4 DRT in Prolog

The architecture that is offered by DRT is gradually coming into focus—but so far our discussion has been purely theoretical. What does DRT look like from the perspective of *computational* semantics? The remainder of this book is essentially an extended answer to

this question; in the present section we lay the computational foundations for the chapters that follow.


We shall do two things. First, we define a simple model checker for DRS languages. This will enable us to fix our basic Prolog conventions for DRT, and gives us an interesting new perspective on the semantics of DRSs. Following this, we implement the translation of DRT into first-order logic given in Section 1.3. This simple program opens the door to the world of first-order inference techniques, and underpins much of our later work.

A Simple Model Checker

We want to define a model checker for DRSs. As we know from Chapter 1, where we defined a model checker for first-order logic, there are four things on our to-do list. We have to: (1) decide how to represent vocabularies in Prolog, (2) decide how to represent models, (3) decide how to represent DRSs, and (4) specify the model checker for DRSs. The good news is that we can reuse the vocabularies and model representations that we introduced in Chapter 1, so let's forget (1) and (2) and turn straight to (3): fixing a Prolog notation for boxes.

We will represent DRSs in Prolog as terms of the form `drs(D,C)`, where `D` is a list of terms representing the discourse referents, and `C` is a list of other terms representing the DRS conditions. To represent complex conditions we use the same operator definitions as for first-order logic (but note that we are not going to use the operator for conjunction, since we don't need it in DRT).

Let's consider an example. The DRS for 'Every gimp shrieks' is:



And in Prolog we represent this DRS as:

```
drs([], [drs([X], [gimp(X)]) > drs([], [shriek(X)])]).
```

Our choice of representation is actually rather arbitrary. For a start, we didn't have to use `drs` as the functor name—pretty much anything else would have done. Moreover, we could have devised a glossy operator notation for DRSs had we wanted. Basically, we chose this aspect of our representation because it is simple and fairly readable, but nothing much hangs on it. Note, however, that discourse referents are represented as Prolog

variables. That is, we have made the same choice here that we made when deciding how to represent first-order variables. As we have already discussed, this choice has a number of advantages (in particular, we won't have to worry about α -conversion, for Prolog handles that automatically) but it can be dangerous if used carelessly.

Exercise 1.4.1 What are the Prolog representations of the DRSs for (1) 'Vincent does not die', (2) 'Vincent dances and Mia dances', (3) 'Vincent cleans the back seat or Jules cleans the back seat', and (4) 'If Butch wins, Marsellus loses'?

We now turn to task (4): defining a simple model checker for DRSs. We do so in the spirit of Chapter 1, where we first defined a (rather naive) `satisfy/2` predicate that succeeded if its first argument, a formula of first-order logic, could be satisfied in the model specified by the second argument.

Now, DRSs and conditions are mutually recursive constructs. Unsurprisingly, this is reflected in the implementation of our model checker: we will specify a predicate `satisfyDrs` that succeeds if a DRS can be satisfied in a model, and a predicate `satisfyCondition` that succeeds if a DRS-condition can be satisfied in a model. These predicates will be mutually recursive. Here is the code for `satisfyDrs`:

```
satisfyDrs(drs(Dom,C),Model):-
    assignReferents(Dom),
    satisfyConditions(C,Model).

assignReferents([]).
assignReferents([Referent|Others]):-
    constant(Referent),
    assignReferents(Others).
```

Using the predicate `assignReferents/1`, `satisfyDrs/2` assigns elements of the model (which are specified as constants in our vocabulary) to the discourse referents of the DRS. The next step involves checking the conditions of this DRS, and this is done with the help of `satisfyConditions/2`.

So how do we define `satisfyConditions/2`? Obviously we will need a recursive clause that works us through the condition list, and this is straightforward to define:

```
satisfyConditions([],_).
satisfyConditions([Condition|Others],Model):-
    satisfyCondition(Condition,Model),
    satisfyConditions(Others,Model).
```

So it remains to specify what `satisfyConditions/2` does to individual conditions. Let's first consider negative and disjunctive conditions. In fact, nothing much needs to be said here; we can deal with them simply by using Prolog negation (negation as failure) and Prolog disjunction respectively:

```
satisfyCondition(~ Drs,Model):- \+ satisfyDrs(Drs,Model).
```


```
satisfyCondition(Drs1 v Drs2,Model):-
  (
    satisfyDrs(Drs1,Model)
  ;
    satisfyDrs(Drs2,Model)
  ).
```

The clause for \Rightarrow is more interesting; here's how we'll do it:


```
satisfyCondition(Drs1 > Drs2,Model):-
  (
    satisfyDrs(Drs1,Model),
    \+ satisfyDrs(Drs2,Model),
    !,
    fail
  ;
    true
  ).
```

What's going on here? Well, from a computational perspective, the following: this clause tries to satisfy the antecedent DRS (that is, `Drs1`), and, at the same time, tries to prove that the consequent DRS (that is, `Drs2`) *cannot* be evaluated. If this is the case, the implication as a whole cannot be satisfied, and we use the cut-fail combination to inform Prolog of this. Now, using backtracking, Prolog will try out every possible assignment of constants of the domain to the discourse referents of the antecedent DRS; if it does not succeed in showing that the consequent DRS cannot be satisfied, then we land in the second disjunct of the clause, where the built-in Prolog predicate `true` tells us that the implication as a whole *can* be satisfied.

But there's a more abstract perspective on this code. As we asked the reader to show in Exercises 1.2.1 and 1.3.5, \Rightarrow conditions can be eliminated without loss of expressive power. In fact, a configuration of the form



is equivalent to a configuration of the form



A moment's thought reveals that the code above exploits this equivalence.

We come at last to the basic conditions. Note that identity conditions are simply checked by trying to unify them—this only succeeds if X and Y are identified with the same constants out of the model's domain.

```
satisfyCondition(X=Y,_Model):-
    X=Y.
```

```
satisfyCondition(BasicCond,Model):-
    memberOfList(BasicCond,Model).
```

It only remains to wrap this all up in a driver. We'll do this in the way we did in Chapter ??; we'll define a driver that, given a DRS, picks an example model and tries to satisfy it:

```
evaluate(Drs,Example):-
    example(Example,Model),
    satisfyDrs(Drs,Model).
```

And now for the \$64,000 question: *what have we actually implemented?* Clearly we've implemented some kind of DRS model-checker—but is it an implementation of embedding semantics or of dynamic semantics?

It's impossible to decide—it's both. Read declaratively, the code seems a fairly direct implementation of embedding semantics. Of course, we have to make allowances for the fact that we're working with lists and not sets—nonetheless, the basic correspondence is clear. On the other hand, if you think of what the Prolog interpreter is actually doing with this code, one is drawn to dynamic semantics; that is, read procedurally, this code seems more like an implementation of dynamic semantics. In short, simply expressing the basic semantic ideas of DRT in Prolog blurs the distinction between the dynamic and embedding semantics, which emphasizes how easy it is to view DRS languages either statically or dynamically.

In our view, this is just the way it should be. DRT supports a rich and productive interplay of intuitions. Simple though it is, our little model checker underlines the fundamental unity of some of its key supporting ideas.

Exercise 1.4.2 Explain why, in the clause for checking identity of two discourse referents, there is no need to use information from the model, as supplied by the second argument of `satisfyCondition/2`.

Exercise 1.4.3 As we know from Chapter 1, using Prolog unification in combination with negation as failure can have nasty consequences. Test how the DRS model checker fares with respect to the problems we noted for the naive first-order logic model checker developed in Chapter 1. Define a refined version of the DRS-evaluation predicate that deals with these problems, reusing as much Prolog code as possible.

Compiling DRSs into First-Order Logic

In this section we define a simple Prolog predicate called `drs2fo1` which takes a DRS and produces an equivalent first-order formula; this predicate is a straightforward implementation of the translation function fo discussed in Section 1.3.

Let's first consider how to translate boxes. Recall that fo does this as follows:

$$\left(\begin{array}{|c|} \hline x_1, \dots, x_n \\ \hline \gamma_1 \\ \cdot \\ \cdot \\ \cdot \\ \gamma_m \\ \hline \end{array} \right)^{fo} = \exists x_1 \dots \exists x_n ((\gamma_1)^{fo} \wedge \dots \wedge (\gamma_m)^{fo})$$

Given our list-based Prolog representation of DRSs, the following is a natural implementation:

```

drs2fol(drs([], [Cond]), Formula) :- cond2fol(Cond, Formula).

drs2fol(drs([], [Cond1, Cond2|Conds]), Formula1 & Formula2) :-
    cond2fol(Cond1, Formula1),
    drs2fol(drs([], [Cond2|Conds]), Formula2).

drs2fol(drs([X|Referents], Conds), exists(X, Formula)) :-
    drs2fol(drs(Referents, Conds), Formula).

```

That is, we work through the list of discourse referents, recursively pumping out the existential quantifiers we need; the third clause does this. In addition, we work our way through the condition list, recursively conjoining their translations as we do so; the second clause does this. Of course, as DRSs and conditions are mutually inductive concepts, the second clauses need to call on a predicate `cond2fol` which tells us how to translate conditions to first-order formulas; we will define this predicate shortly. The first clause grounds this process with a call to `cond2fol`.

So our next task is to specify what `cond2fol` does. The clauses for disjunctive and negative conditions are transparent implementations of what *fo* does in these cases:

```

cond2fol(~ Drs, ~ Formula) :-
    drs2fol(Drs, Formula).

cond2fol(Drs1 v Drs2, Formula1 v Formula2) :-
    drs2fol(Drs1, Formula1),
    drs2fol(Drs2, Formula2).

```

But now we must deal with \Rightarrow conditions. Recall that *fo* handles these as follows:

$$\left(\begin{array}{|c|} \hline X_1, \dots, X_n \\ \hline \gamma_1 \\ \cdot \\ \cdot \\ \cdot \\ \gamma_m \\ \hline \end{array} \right) \Rightarrow B^{fo} = \forall X_1 \dots \forall X_n ((\gamma_1)^{fo} \wedge \dots \wedge (\gamma_m)^{fo}) \rightarrow (B)^{fo}$$

We implement this by working recursively through the discourse referent list, pumping out the universally quantified formulas we need, and then placing the implication in place:

```

cond2fol(drs([], Conds) > Drs2, Formula1 > Formula2) :-
    drs2fol(drs([], Conds), Formula1),
    drs2fol(Drs2, Formula2).

```

```
cond2fol(drs([X|Referents],Conds) > Drs2, forall(X,Formula)):-
  cond2fol(drs(Referents,Conds) > Drs2,Formula).
```

Only one task remains: we need to say how basic conditions are translated. This is the clause on which the entire mutual induction rests, and it is defined as follows:

```
cond2fol(BasicCondition,AtomicFormula):-
  BasicCondition =.. [_Predicate|[FirstArg|OtherArgs]],
  simpleTerms([FirstArg|OtherArgs]),
  AtomicFormula=BasicCondition.
```

That is, when we reach the bottom level, we check that what we have found is really a basic condition; if it is, we have the atomic first-order formula we need.

And that's it. Very simple, and as we shall see, very useful. We now have a bridge between the world of DRT and the world of first-order inference methods, and we shall make heavy use of it in Chapter 6.

Exercise 1.4.4 Implement the translation *dr* discussed in Section 1.3 that maps first-order expressions to DRSs. Don't just implement the lean version: add clauses to handle \vee , \Rightarrow and \forall directly (see Exercise 1.3.4).

Exercise 1.4.5 Is `drs2fol/2` reversible? It clearly isn't. Check the predicate on all input modes, and report cases where it fails to give the right result.

Software Summary of Chapter 1

`modelDRT.pl` Contains the predicate that evaluate a DRS with respect to a first-order model. (page 173)

`drs2fol.pl` The translation from DRSs to first-order formulas. (page 175)

Notes

DRT was originally developed by Kamp (1984) and Heim (1982) (who called it “File Change Semantics”). DRT is blessed with an excellent textbook, namely Kamp and Reyle

1993. The book patiently describes the underlying ideas and formal details of the theory, and applies it to a wide range of natural language phenomena. The standard construction algorithm is developed in detail and DRS are interpreted via the embedding semantics; the reader who wants to find out more about DRT is advised to start here. Another text-book level account well worth looking at is given in the second volume of Gamut (Gamut 1991). This account is more critical of DRT, and develops an alternative called Dynamic Predicate Logic (DPL); this interprets the standard syntax of first-order using a variant of the dynamic semantics discussed in the text. However the box-based syntax of DRT is now widely accepted as standard notation in linguistics and computational linguistics, and dynamic semantics seems to be increasingly viewed as an alternative interpretation of DRSs. The translation *fo* of DRSs into first-order logic with equality given in the text was taken from Kamp and Reyle (Kamp and Reyle 1993). It's probably the best known translation, but there are others; Muskens (Muskens 1996), for example, gives an interesting alternative. For a more advanced discussion of DRT, try Van Eijck and Kamp 1997.

Chapter 2

Building Discourse Representations

We now know what DRT is, and have developed some simple Prolog programs for working with DRS languages. But how are we to construct the DRSs that correspond to sentences, or indeed, entire discourses? This chapter explores the issue in detail, and the basic message that emerges is straightforward:

DRT is a highly flexible theory, compatible with a wide range of semantic construction strategies.

We begin by exploring three alternative approaches to building discourse representations. First, we introduce a technique called *DRS-threading*. This technique, which can be viewed as a declarative approach to the standard algorithm, is elegant, efficient, and highly instructive. However, our initial implementation of threading is open to the same sort of difficulties we encountered in Experiment 2 in Chapter ???. So we abandon threading and turn to a second technique, the lambda-based methods developed in Part I. As we shall see, the techniques and tools developed in our earlier work adapt straightforwardly to DRT; we will be able to reuse the rules in `englishGrammar.pl`, and the required interface is very simple.

We then change gears. Basic semantic construction is all very well, but how are we to handle quantifier scope ambiguities? As we shall show, the hole semantics approach to underspecification (Chapter ??) adapts easily DRS languages, and implementing it requires only routine adaptations of our earlier code.

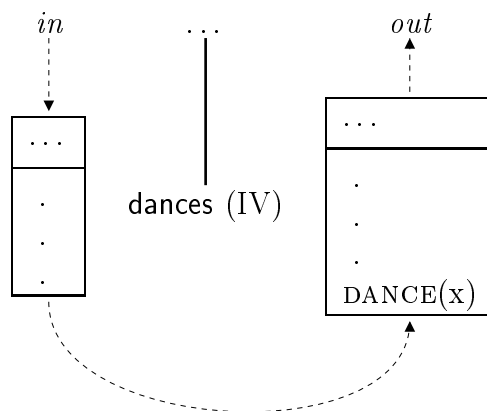
2.1 DRS-Threading

Threading is an interesting technique for building DRSs. For a start, as it is based on the use of *difference structures*, threading is extremely efficient. Difference structure programming uses the difference between incomplete data structures to represent partial results of

a computation; the classic Prolog example is the use of difference lists to combine lists efficiently. As we have chosen to represent DRSs in Prolog as a pair of lists, and as we want to combine the small DRSs associated with lexical entries into a big DRS that represents the entire discourse, it seems sensible to experiment with difference structures. But there is an even more compelling reason for thinking about DRS construction in terms of threading: it is enlightening and instructive to do so. Threading embodies a clear and remarkably simple idea—an idea that will enable to peel away the inessential top-down paraphernalia of the standard DRS construction algorithm and reveal the elegant declarative core that lies at its heart. Moreover it gives us (yet another) computational perspective on the idea of context change potential.

When we introduced the standard construction algorithm in Section 1.1, we did so with the help of a series of diagrams showing a parse tree, a DRS, and arrows linking them. We explained how we moved around the tree in a top-down left-to-right direction, gathering the semantic contributions of the various constituents and packing them into a DRS. Let's change this picture slightly. Think of the parse tree as a rigid structure, perhaps made of wire. Think of the DRS as a little plastic bubble which we can freely slide round this rigid tree (rather like the beads on an abacus, or a children's game). At the start of a 'calculation' or 'game' we place the DRS representing the prior discourse at the top of this tree and then slide it smoothly from node to node, taking care to slide it over every single node in the tree. When the DRS slides over a node (or to put it another way, when a node is *threaded through* the DRS) the node places its semantic contribution inside the DRS.

Here's an example. What happens when we move the DRS over a node labeled by an intransitive verb, say 'dances'? Now, the contribution of an intransitive verb to the overall DRS is just a basic condition; it doesn't introduce a discourse referent. So when the DRS threads its way around a node labeled by 'dances', the following occurs:



Note that the domain of the outgoing DRS is equal to the domain of the ingoing DRS, and that the condition set of the outgoing DRS consists of the conditions of the ingoing DRS *plus* the basic condition introduced by the verb. It is clear that the node has simply made the expected semantic contribution as it was threaded through the DRS.

As far as the basic idea of threading is concerned, there really is not a lot more to say than is suggested by this diagram. All the work in this section is directly based on the intuition of a DRS sliding its way round a parse tree, collecting the semantic contributions as it goes, and we shall simply work out the consequences of this idea for the various syntactic categories. But then, why is this new picture important? After all, sliding a DRS around a tree sounds every bit as procedural as the original presentation of the standard construction algorithm!

However we really have taken an important step forward, for threading focuses attention on what is truly important: at every node there is an *ingoing* DRS, an *outgoing* DRS, and the difference between the ingoing and outgoing DRS is exactly the information that is contributed by the syntactic category. To put it another way: threading is really about thinking in terms of *pairs* of DRSs, for these give us ‘before-and-after’ pictures of the various stages of semantic construction.

And now for the payoff. If we think in terms of ingoing and outgoing DRSs, the procedural aspects of the standard construction algorithm simply melt away. We are left with its declarative core, and this can be expressed using simple Prolog constraints. To see how, let’s return to our ‘dances’ example.


We shall associate each category of the syntactic analysis with a complex Prolog term `DrsIn-DrsOut`, representing the ingoing and outgoing DRSs. So, to capture the effect of the previous diagram, the lexical entry for the intransitive verb ‘dances’ must be:

```
iv(X, DrsIn-DrsOut)-->
{
  DrsIn=drs(Dom, Conds),
  DrsOut=drs(Dom, [dance(X) | Conds])
},
[dances].
```

The most important point to observe is two unification equations: these constraints constitute the declarative core of the previous diagram. However the reader should also note that we are keeping track of the discourse referent `X` used as the argument of `dance`. This is clearly sensible, but note that we have done so by adding it as an extra argument to the grammar rule and letting the DCG take care of it. Recall that this is essentially the strategy we adopted in Experiment 2 of Chapter ??, and given the outcome of that experiment, our first hint that trouble is brewing; but let’s not worry about it for now. Finally, note that the code can be simplified by replacing `DrsIn` and `DrsOut` by the terms they have to unify with:

```
iv(X, drs(Dom, Conds)-drs(Dom, [dance(X) | Conds]))-->
[dances].
```

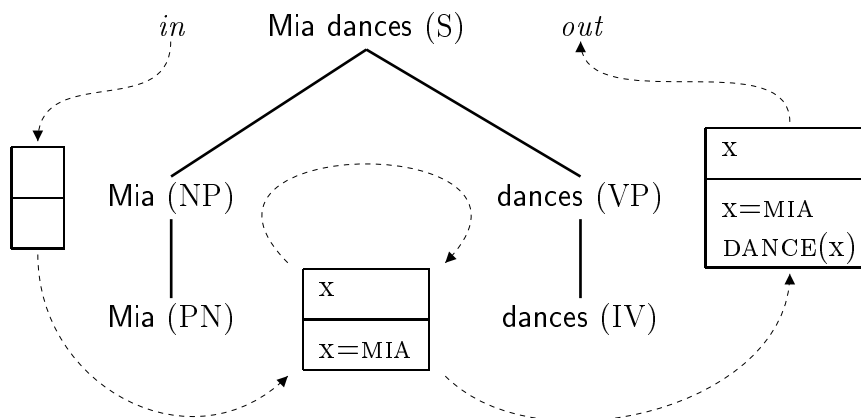
Lets now give a threading analysis of the sentence ‘Mia dances’. First, we need to add a lexical entry for the proper name ‘Mia’. Proper names introduce both a discourse referent and a condition. The following diagram shows what happens when a node labeled by ‘Mia’ is threaded through a DRS:



The following Prolog code turns this picture into a simple constraint (once again, note that we’ve also added an extra DCG argument, Experiment 2 style, to keep track of the discourse referent):

```
pn(X, drs(Dom, Conds) - drs([X|Dom], [X=mia|Conds])) -->
  [mia].
```

That’s the situation in the lexicon. So let’s now see what happens as the DRS representing the previous discourse slides its way around the parse tree for ‘Mia dances’ (in the following diagram we assume that this initial DRS is empty):



There are two unary branching rules in this example ($NP \rightarrow PN$ and $VP \rightarrow IV$), and they are easy to handle in Prolog—they simply pass on the pair of DRSs, since no new information is added. That is, the ingoing DRS of the daughter category is the ingoing DRS of the mother, and the outgoing DRS of the mother category is the outgoing DRS of its daughter:

```
np(X, Drs In-DrsOut)-->
  pn(X, Drs In-DrsOut).

vp(X, Drs In-DrsOut)-->
  iv(X, Drs In-DrsOut).
```

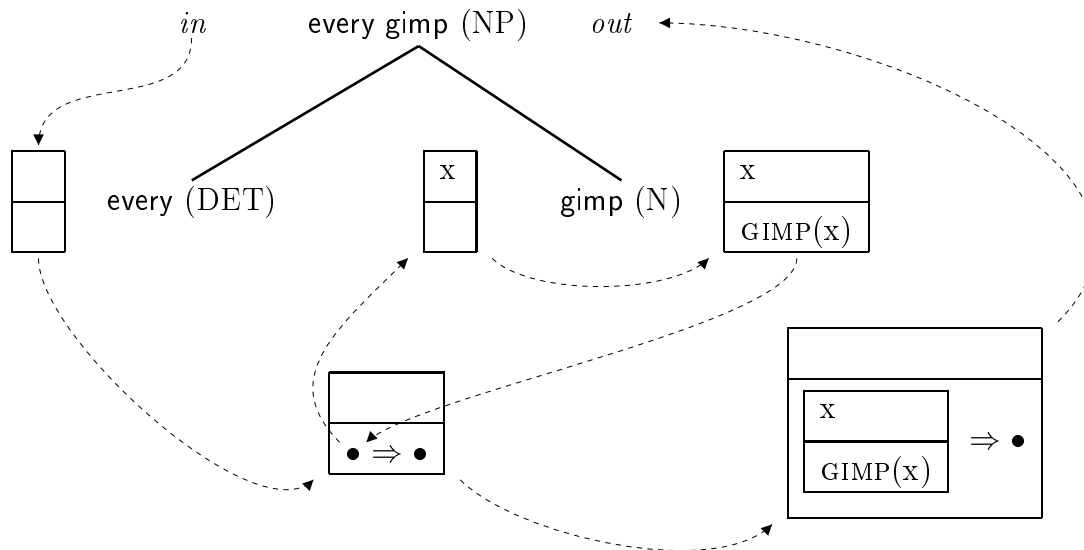
For the binary branching rule $S \rightarrow NP VP$, the following serves our needs:

```
s(Drs In-DrsOut)-->
  np(X, Drs In-Drs),
  vp(X, Drs-DrsOut).
```

Here two points need to be made here. First, note that the outgoing DRS of the NP node (that is, *Drs*) unifies with the ingoing DRS of the VP node, the ingoing DRS of the sentence is the ingoing node of the NP, and the outgoing DRS of the sentence is the outgoing DRS of the VP. These constraints ensure that the information is packed into the DRS correctly—in effect, they perform the same function as the processing insisted on in the standard construction algorithm, but abstract away from its (unnecessarily restrictive) procedural stipulations. Second, note that the extra arguments percolated upwards by the noun phrase and verb phrase are unified; this ensures the correct bindings of the discourse referent with its conditions, very much in the spirit of Experiment 2 of Chapter ??.

And that completes our threading analysis of ‘Mia dances’. But this is a very simple sentence; how does DRS-threading extend to sentences containing determiners such as ‘every’. Let’s give a threading analysis of ‘Every gimp runs’.

The semantic information associated with the the determiner ‘every’ is a complex structure: an implicational condition consisting of two DRSs. This is shown in the analysis for the noun phrase ‘every gimp’ (we assume that the ingoing DRS is empty):



The outgoing DRS for ‘every’ is the ingoing one to which implicational condition (that is, an arrow linking two sub-DRSs) has been added. We don’t yet know what the contents of these sub-DRSs are, so we represent each of them with a black hole •. Actually, we *do* know a little more: this implication is a universal quantification, hence the antecedent black hole must contain a discourse referent x ; this is indicated by the third threading arrow. How is the condition set of the antecedent to be filled? By the contribution of the noun (‘gimp’ in this example), as fourth threading arrow indicates. In short, once we’ve threaded every node in the noun phrase through the initial DRS, we will have successfully filled in the antecedent black hole and now are ready to pass on the resulting incomplete DRS (incomplete because the consequent is still just a black hole) to the rest of the sentence for further threading. But before going any further, let’s express this noun phrase analysis in Prolog. First, the lexical entry for ‘every’:

```
det(X,DrsIn-DrsOut,RestrIn-RestrOut,ScopeIn-ScopeOut)-->
  [every],
  {
    DrsIn = drs(Dom,Conds),
    DrsOut = drs(Dom,[RestrOut > ScopeOut|Conds]),
    RestrIn = drs([X],[ ]),
    ScopeIn = drs([],[ ])
  }.
```

Note that we have added two DRS-threading pairs: **RestrIn-RestrOut**, which builds the DRS for the restriction (the first •), and **ScopeIn-ScopeOut** which builds the DRS for the nuclear scope (the second •). The outgoing DRS equals the ingoing DRS plus a new implicational condition. A new box **RestrIn** is opened in which information about the

restriction can be placed, and a new discourse referent is placed in it. Similarly, a new box **ScopeIn** is provided for the nuclear scope, but no information can be placed in this box by ‘every’.

Now for the DCG rule that combines a determiner and a noun to form an NP:

```
np(X, DrsIn-DrsOut, ScopeIn-ScopeOut)-->
  det(X, DrsIn-DrsOut, RestrIn-RestrOut, ScopeIn-ScopeOut),
  noun(X, RestrIn-RestrOut).
```

Quite simply, the noun fills in the restriction information, just as we saw in the previous diagram. Note that that this rule can be simplified without affecting the way it works:

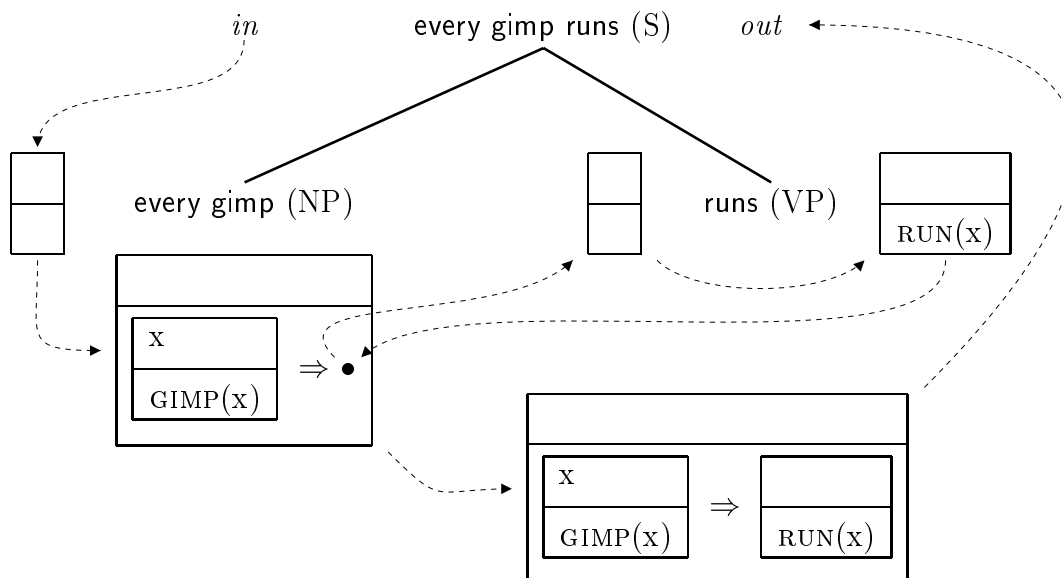
```
np(X, Drs, Scope)-->
  det(X, Drs, Restr, Scope),
  noun(X, Restr).
```

Here **Drs** stands for the DRS-pair corresponding to the the overall DRS, **Restr** for the DRS-pair that fills the first hole (the restriction of the quantifier), and **Scope** for the DRS-pair that fills the second hole, the nuclear scope of the quantifier. Now this is nice—but observe that trouble is on its way. In particular, note that whereas in this rule **np** has *three* arguments, in the earlier proper name rule used to analyze ‘Mia dances’ it only had *two* arguments. In short, we are starting to run into the same sort of unpleasantness we encountered in Experiment 2: the need to adapt all our rules to a common format. We will do this shortly (for the small collection of rules we’re working with here, it’s not difficult) but clearly our DCG-based handling of missing information is raising problems that will require further attention.

Exercise 2.1.1 [Easy] Simplify the DCG rules that deal with determiners and nouns by introducing the discourse referent in the rule for nouns, instead of the rule for determiners.

Exercise 2.1.2 [Intermediate] Draw a picture that illustrates how to give a threading analysis the noun phrase ‘a gimp’. Using this picture as a guide, give a DCG rule for the determiner ‘a’.

Lets return to our analysis of ‘Every gimp runs’. We now have an analysis for ‘Every gimp’, but this contains a black hole marking the missing nuclear scope information. How do we fill it? The verb phrase takes care of this, as the following analysis shows:



That is, we start threading the VP with an empty DRS and substitute the result for the black hole in the consequent of the implicational condition introduced by ‘every’. (Incidentally, note that the rule for ‘every’ already states that threading for the scope commences with the empty DRS, for it contains the constraint $\text{ScopeIn} = \text{drs}([], [])$.)

Now the basic idea is clear, but once again Experiment 2 strikes, and to get it working properly we have to revise the DCG rule for sentences (given in our earlier analysis of ‘Mia dances’) as follows:

```
s(DrsIn-DrsOut)-->
  np(X,DrsIn-DrsOut,ScopeIn-ScopeOut),
  vp(X,ScopeIn-ScopeOut).
```

Putting this in a more readable format we get:

```
s(Drs)-->
  np(X,Drs,Scope),
  vp(X,Scope).
```

This completes our threading analysis of ‘Every gimp runs’, and all important the ingredients of DRS-threading have now been discussed. To finish off, we’ll extend our DCG with some new grammar rules (including the required reformulation of the rule for proper names), see how to make use of the lexicon defined in Chapter ??, and define a simple driver. First, the rules:


```

d(DrsIn-DrsOut)-->
  s(DrsIn-Drs),
  d(Drs-DrsOut).

d(Drs-Drs)-->
  [].

vp(X, DrsIn-DrsOut)-->
  tv(X,Y,Scope),
  np(Y,DrsIn-DrsOut,Scope).

np(X, DrsIn-DrsOut, Drs-DrsOut)-->
  pn(X,DrsIn-Drs).

```

The first two rules allows us to make small discourses by sequencing sentences (we have introduced a new category D for ‘discourse’). We also give a DCG rule for combining a transitive verb and a noun phrase to form a verb phrase. Finally, as promised, we have given a revised version of the DCG rule for NP \rightarrow PN so that we have one uniform treatment of noun phrases within the grammar. This new rule may be a little puzzling, and the reader is asked to study it in the following exercise.

Exercise 2.1.3 [Intermediate] Using either pen and paper or a tracer, compare the sequence of variable instantiations this program performs when building representations for ‘Mia dances’ and ‘A gimp dances’. Compare the new rule given for proper names with the corresponding rule of Experiment 2 in Chapter ?? . Make sure you understand how the use of variable doubling in the earlier chapter relates to the code given above.

It is straightforward to make use of the the English lexicon defined in Chapter ?? with the threading DCG, For example, the entry for a transitive verb is simply:

```

tv(X, Y, drs(Dom, Conds)-drs(Dom, [Cond|Conds]))-->
  {
    lexicon(tv, Sym, Phrase),
    compose(Cond, Sym, [Y, X])
  },
  Phrase.

```

It remains to design a driver predicate that feeds the DCG with an ingoing DRS and returns the outgoing DRS. We shall simply give it the empty DRS, but we could just as easily give it the DRS produced by previously analyzed discourse.

```

parse:-
  readLine(Discourse),
  d(drs([],[])-Drs,Discourse,[]),
  printRepresentation(Drs).

```

For instance, suppose we want to build the DRS for the discourse that contains of the sentence ‘Vincent snorts’ followed by ‘A woman collapses’.

```

?- parse.

> Vincent snorts. A woman collapses.

drs([A,B],[collapse(A),woman(A),snort(B),B=vincent])

```

Note that the lists are built from right to left; that is, the most recent information comes first in the list.

And that’s DRS threading. What can we say about it at a more general level?

Threading clearly has a lot going for it. For a start, the way it arises as declarative abstraction from the standard construction algorithm is instructive. Moreover, it is suggestive in other ways: it mirrors the idea that the meaning of an utterance lies in its context change potential in way that is reminiscent of dynamic semantics. Recall that dynamic semantics interprets DRSs with respect to *pairs* of embeddings, an input embedding and an output embedding—but whereas dynamic semantics exploit the use of input/output pairs to define an alternative semantics, threading uses the same idea to place constraints of semantic construction. The reader will find it instructive to re-examine our presentation of dynamic semantics in the light of the work of the present section.

But now for the bad side. Our approach to DRS-threading faces two problems. Actually, the first of these, which concerns our treatment of proper names, is more interesting than problematic. As we defined DRS-threading, proper names add a discourse referent and condition to the ingoing DRS. But the ingoing DRS need *not* be the main DRS: it could easily be a sub-DRS (for example, the consequent DRS of an implicational DRS) and this gives rise to an an unwanted result. The discourse referent of a proper name and its condition should float to the top of the DRS, but in our presentation of threading they may stay down at a lower level. This, however, is not a serious problem.

Exercise 2.1.4 [Intermediate] Change the program in such a way that proper names are floated to the top DRS. Hint: use a second DRS pair.

But there is a problem with our work that is potentially far more damaging. Quite simply, all the problems associated with the Experiment 2—the problems which motivated our

introduction of the lambda calculus—are back in force and as annoying as ever. Although we solved the problems for the present (very simple) grammar, the solution is clearly Prolog specific and rather ad-hoc. Thus an obvious question faces us. Is it possible to work with variables in DRT in more disciplined, less Prolog oriented way?

2.2 Building DRSs with Lambdas

In Chapter 2 we took first-order logic as our basic representation language, and allowed ourselves to mark missing information with the help of the λ -operator. This enabled us to give a compositional semantics to sentences containing quantified NPs, and opened the door to more sophisticated ideas, such as the use of nested Cooper storage in Chapter 3. Let's explore a parallel strategy here: we will build up sentence representations with the help of the lambda calculus, but will use λ to mark information missing from DRSs rather than first-order formulas. We often call this combination as λ -DRT, and we shall see, it enables us to reuse the tools and techniques developed in Part I.

In order to make this approach work, we are going to have to enrich the box languages with a new construct: the *merge* operation \oplus . This combines two DRSs by taking the union of the two universes and the two lists of conditions. For example,

$$\begin{array}{|c|} \hline x \\ \hline \text{BOXER}(x) \\ \text{LOSE}(x) \\ \hline \end{array} \oplus \begin{array}{|c|} \hline y \\ \hline \text{DIE}(y) \\ y=x \\ \hline \end{array} = \begin{array}{|c|} \hline x \ y \\ \hline \text{BOXER}(x) \\ \text{LOSE}(x) \\ \text{DIE}(y) \\ y=x \\ \hline \end{array}$$

There are two reasons for adding the merge. The first is simply that it makes *explicit* something we have been doing *implicitly* all along. In particular, note that whenever we use the standard construction algorithm (or indeed, the threading implementation of the previous section) we take care to add the semantic contribution of the latest sentence to the DRS associated with the discourse so far—after all, that's one of the key ideas of DRT. The merge gives us some natural terminology and notation for talking about this idea. Indeed, it lets us break it down this process into slightly smaller steps, for we can now say that the semantic representation for a discourse is obtained by building a (completely new) DRS for the incoming sentence and then merging this new DRS with the DRS associated with the discourse so far. The second reason is deeper: when used in combination with λ , the merge is precisely the operation on DRS we need to state lexical entries clearly and concisely. We shall shortly see why.

What sort of λ -DRSs (as we usually call them) should be assigned to lexical items? Let's consider some examples. Nouns are associated with the following kind of representation:

$$\text{'boxer': } \lambda y. \begin{array}{|c|} \hline \\ \hline \text{BOXER}(y) \\ \hline \end{array}$$

$$\text{'woman': } \lambda y. \begin{array}{|c|} \hline \\ \hline \text{WOMAN}(y) \\ \hline \end{array}$$

That is, just as with our earlier work with first-order representations, we are using lambdas as a tool to explicitly mark missing information—a tool that has nothing to do with the peculiarities of Prolog, and DCGs, or indeed an other language. In the examples just given, the bound occurrences of y indicate that nouns are associated with DRS—but DRSs that need to be supplied with a term, in the indicated position, to become a full DRS.

The lexical entries for the determiners ‘a’ and ‘every’ are more interesting. In particular, they require us to use both lambdas and the merge:

$$\text{'a': } \lambda P. \lambda Q. \begin{array}{|c|} \hline z \\ \hline \\ \hline \end{array} \oplus P@z \oplus Q@z$$

$$\text{'every': } \lambda P. \lambda Q. \begin{array}{|c|} \hline \\ \hline \begin{array}{|c|} \hline z \\ \hline \\ \hline \end{array} \oplus P@z \Rightarrow Q@z \\ \hline \end{array}$$

These are clearly analogous to the lexical entries given for these determiner given in Chapter ???. For example, we there assigned the following representation to ‘every’:

$$\text{'every': } \lambda P. \lambda Q. \forall z (P@z \rightarrow Q@z)$$

In both representations, the abstracted variable P marks the missing restriction information, while Q marks the missing nuclear scope. Moreover, the quantificational force of the λ -DRS representation is obtained by using a combination of $\begin{array}{|c|} \hline z \\ \hline \\ \hline \end{array}$ and \Rightarrow , in a fashion that is clearly analogous to the use of $\forall z$ and \rightarrow in the other. Note that \oplus is useful because it lets us state which DRS need to be glued together in the course of semantic construction.

Finally, let’s consider the lexical entries for intransitive and transitive verbs:

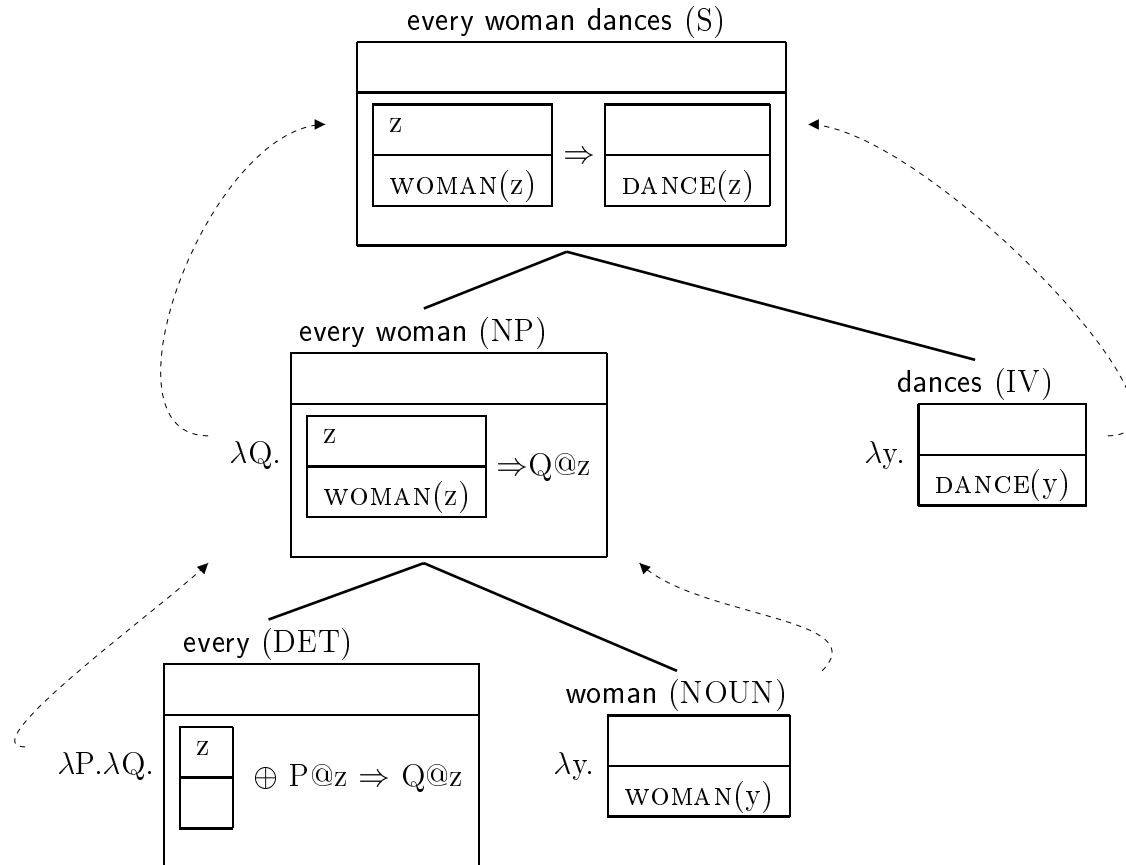
$$\text{'dances': } \lambda y. \begin{array}{|c|} \hline \\ \hline \text{DANCE}(y) \\ \hline \end{array}$$

‘loves’: $\lambda X.\lambda y.X@ \lambda x.$

LOVE(y,x)

Again, the reader should compare these entries with their analogs in Chapter ??.

It’s time to put these entries to work. As a sample, here’s how we build a DRS for ‘Every woman dances’:



Note that the fundamental mechanism that drives this construction is the one we are familiar with from Chapter ??: the lambda calculus plugs things together in the expected way. Note, however, that merge plays an important supporting role. For example, the λ -DRS associated with ‘every woman’ in the above tree is formed from the λ -DRSs associated with ‘every’ and ‘woman’ by a lambda conversion (using the bound variable P), followed by another lambda conversion (using the bound variable z), followed by a merge.

Exercise 2.2.1 [Easy] Write out the steps involved in obtaining the λ -DRS associated with ‘every woman’ in the above tree.

Exercise 2.2.2 [Easy] Build the DRS representing ‘A woman snorts’ using the lexical representations given above.

Exercise 2.2.3 [Intermediate] State the needed lexical representations, and then build the DRS representation of the donkey sentence ‘If a boxer loses, he dies’.

Exercise 2.2.4 [Easy] Suppose that we tried to avoid using the merge by giving ‘a’ the following lexical entry instead of the one given above:

$$\text{‘a’}: \lambda P.\lambda Q. \begin{array}{|c|} \hline x \\ \hline P@x \\ \hline Q@x \\ \hline \end{array}$$

This won’t work. Why?

In short, the main ideas underlying semantic λ -DRT are those we are familiar with from Part I. Moreover, as we shall now see, implementing λ -DRT in Prolog is rather easy—indeed the only part of the implementation that requires thought is how to handle the merge. Just about everything else is a matter of reusing the tools developed in Part I.

Let’s start with the really good news: *the grammar requires no work whatsoever*. A moments thought shows why. Our grammar is a DCG which uses our predicates for functional application to manipulate the representations handed up by the lexical entries. It doesn’t much matter to this DCG whether these representations are first-order formulas mixed with lambdas, or DRSs mixed with lambdas. In short, as we claimed in Chapter ??, lambda-based construction methods make good sense from a grammar engineering perspective.

So what do we actually have to do? Just two things: define the macros for the lexical entries, and wrap everything up in a new driver. Let’s start with thinking about the semantic macros. It follows from our earlier discussion that they will look like this:

```
nounSem(Sym,lambda(X,drs([], [Cond]))) :-
    compose(Cond, Sym, [X]).

ivSem(Sym,lambda(X,drs([], [Cond]))) :-
    compose(Cond, Sym, [X]).

modSem(neg,lambda(P,lambda(X,drs([], [~(P@X)])))).
```

But now we have more work to do. Up till now we have not had to use the merge. But what are we to do with merges which cannot be conjured away like this, as is the case for the merges in the lexical entry for the determiner ‘a’? One option would be to make use of `append/3` in the semantic macro for determiners, but this doesn’t work. Suppose that we use the Prolog variables `P` and `Q` for the λ -DRSs for the ‘N’ and the ‘VP’ part, and introduce `X` as the new discourse referent ‘a’ generates. Then we need to merge the DRSs

that result from, say $P@X$ and $Q@X$. But from Prolog's point of view, these are not DRSs (for Prolog DRSs are of the form $\text{drs}/2$).

So we need to think of something else. We shall explore a more straightforward alternative. We shall decide on a Prolog notation for the merge of two DRSs, build Prolog representations which contain such explicit merge markings, and then define a predicate which carries out the (explicitly marked) merges required. That is, we are going to adopt the same strategy we used to implement the substitution based approach to β conversion.

Suppose $B1$ and $B2$ are DRSs. We can represent the merge of these DRSs as $\text{merge}(B1, B2)$. Then the revised macro for indefinite determiners looks like:

```
detSem(indef, lambda(P, lambda(Q, merge(merge(drs([X], []), P@X), Q@X))))).
```

What happens when we query a DCG containing such entries? Here's the output that we get for a simple sentence with two indefinite noun phrases:

```
?- s(S, [a, boxer, eats, a, big, kahuna, burger], []), betaConvert(S, B).
```

```
B = merge(drs([X], [boxer(X)]),
          merge(drs([Y], [big_kahuna_burger(Y)]), drs([], [eat(X, Y)])))
```

That is, we obtain the merge of a DRS with a DRS that is itself the result of another merge. Now this is correct, but not very readable, so let's write a predicate $\text{mergeDrs}/2$ that will actually carry out the stipulated merges. Note that as DRSs can be used to build complex DRSs, merges can also appear in complex conditions of DRSs, thus we need the following recursive predicate:

```
mergeDrs(drs(D, C1), drs(D, C2)) :-
    mergeDrs(C1, C2).
```

```
mergeDrs(merge(B1, B2), drs(D3, C3)) :-
    mergeDrs(B1, drs(D1, C1)),
    mergeDrs(B2, drs(D2, C2)),
    append(D1, D2, D3),
    append(C1, C2, C3).
```

```
mergeDrs([B1 > B2 | C1], [B3 > B4 | C2]) :- !,
    mergeDrs(B1, B3), mergeDrs(B2, B4), mergeDrs(C1, C2).
```

```
mergeDrs([B1 v B2 | C1], [B3 v B4 | C2]) :- !,
    mergeDrs(B1, B3), mergeDrs(B2, B4), mergeDrs(C1, C2).
```

```

mergeDrs([~ B1|C1],[~ B2|C2]):-!,
    mergeDrs(B1,B2), mergeDrs(C1,C2).

mergeDrs([C|C1],[C|C2]):-
    mergeDrs(C1,C2).

mergeDrs([],[]).

```

With the help of this predicate, we can pursue the following strategy: we first parse a sentence or discourse. This gives us a DRS with explicit merge statements. We then use `mergeDRS` to arrive at simpler DRSs containing no merge statements.

We can now define a driver that parses the discourse and returns the associated DRS for it. By incorporating a call to `mergeDrs`, we ensure that our DRSs are presented to us nicely simplified:

```

parse:-
    readLine(Discourse),
    d(MergeDrs,Discourse,[]),
    betaConvert(MergeDrs,ReducedDrs),
    mergeDrs(ReducedDrs,Drs),
    printRepresentation(Drs).

```

For example:

```

?- parse.

> Every boxer growls

drs([], [drs([A], [boxer(A)])>drs([], [growl(A)])])

```

And that's all there is to it. The exercises at the end of the section extend this work in various directions.

What can we say of a more general nature concerning λ -DRT? Clearly the approach is natural. For a start, it imitates the Montagovian approach to semantic construction, and extends it to discourse in a very straightforward way; the merge plays a key role here, giving us an operation which can be explicitly invoked to define the required DRSs. Moreover, the fact that we can reuse so much of our earlier work is a powerful testimony both to the robustness of Montague's original methods, and to the naturalness of the ideas underlying DRT.

Exercise 2.2.5 [Intermediate] Extend program `lambdaDRT.pl` with a treatment of the determiner ‘no’. (Try to handle sentences like ‘No boxer retires’.)

Exercise 2.2.6 [Intermediate] Extend the program to handle adjectives. (Try to handle sentences like ‘A rich boxer retires’.)

Exercise 2.2.7 [Intermediate] Extend the program with ditransitive verbs (for example, ‘give’).

Exercise 2.2.8 [Project] Rework the programs given in this section so that they are based on the use of difference lists.


2.3 Underspecified DRSs

Until now we have been concerned with how to use syntactic structure to guide DRS construction. But as we know from Chapter 3, this is only the start of the story, for we also need a way to cope with scope ambiguities. What are we to do here?


In fact both the storage and underspecification methods developed in Chapter 3 extend straightforwardly to λ -DRT. As far as storage methods are concerned, this is probably clear: from a computational perspective storage methods, essentially required us to work with a rather simple list-based extension of our normal semantic representation. Making these ideas work for DRS-based representations involves no really new ideas, and we ask the reader to explore this option in Exercise 2.3.2. But what about underspecified representations based on hole semantics? In fact, as we shall now show, most of the work has already been done: the underlying ideas transfer straightforwardly, we can reuse the plugging algorithm introduced in Chapter 3, and we can reuse the lexicon and rules for our fragment of English introduced in Chapter 2.

Combining hole semantics with DRS involves two actions: we first need to *unplug* our DRS-language, by permitting DRSs to be built which contain *holes*. Then, we make use of the simple constraint language defined in Chapter 3 that governs the way how the DRS-chunks can be *plugged together*. Let’s work through a simple example, ‘every boxer loves a woman’, to show what’s involved.

Unplugging the DRS representation for this sentence yields the following chunks, labeled l_1 , l_2 , and l_3 :




Adding the constraints $l_3 \leq h_1$, $l_3 \leq h_2$ ensures that the main verb is out-scoped by the DRSs introduced by the NPs, and $l_1 \leq h_0$, $l_2 \leq h_0$ enclose all the chunks within the scope domain h_0 . Note that we arrive at the essentially the same underspecified representations as in Chapter ??, but instead of dealing with first-order formula's, we now deal with DRSs.




And resolving these underspecified DRS-representations is not different from resolving underspecified first-order representations either. Again, the holes underspecify scope, and in order to give us non-ambiguous interpretations, each hole should be *plugged* with a DRS in such a way that all the constraints are satisfied. In other words, a plugging for a USR is admissible if the instantiations of the holes with labels result in a DRS in which satisfy the constraints. To return to our example, here we have two pluggings:

	h_0	h_1	h_2
P_1	l_1	l_2	l_3
P_2	l_2	l_3	l_1

Plugging P_1 interprets the USR as giving the universally quantifying NP wide scope, out-scoping the indefinite NP. The corresponding DRS is:



Plugging P_2 yields the other reading, where the indefinite NP has wide scope. Here is the corresponding DRS:



That's the idea—so how do we implement it? As in Chapter ??, a USR is represented in Prolog as `usr(D,L,C)`, where `D` is a list of holes and labels of the USR, `L` a list of the labeled DRSs, and `C` a list of constraints. Here is an example:

```
usr([A,C,E,G,I,J,D,H],
    [A:drs([], [merge(drs([B], []), C)>D]),
      C:drs([], [boxer(B)]),
      E:merge(merge(drs([F], []), G), H),
      G:drs([], [woman(F)]),
      I:drs([], [love(B,F)])],
    [leq(A,J), leq(I,D), leq(E,J), leq(I,H), leq(I,J)])
```

As remarked earlier, we can reuse the plugging algorithm introduced in Chapter 3. Moreover, we can reuse the lexicon and rules for our fragment of English introduced in Chapter 2. In fact, all that's left to do is defining the lexical macros and the basic semantic operations.

Let's go through a few examples of macro definitions. They should be essentially familiar to the reader. Here, for example, are the definitions for nouns, proper names, and verbs.

```
nounSem(Sym, lambda(X, lambda(_, lambda(L, usr([], [L:drs([], [Cond])], []))))):-
  compose(Cond, Sym, [X]).

pnSem(Sym, _, lambda(P, lambda(H, lambda(M,
  merge(usr([], [M:merge(drs([X], [X=Sym]), L)], [leq(M,H)]), P@X@H@L))))).

ivSem(Sym, lambda(X, lambda(_, lambda(L, usr([], [L:drs([], [Cond])], []))))):-
  compose(Cond, Sym, [X]).

tvSem(Sym, lambda(K, lambda(Y, K@lambda(X, lambda(_, lambda(L,
  usr([], [L:drs([], [Cond])], []))))))):-
  compose(Cond, Sym, [Y,X]).
```

Nouns and proper names introduce one labeled DRS, and impose no constraints. Verbs, on the other hand, introduce a labeled DRS together with a \leq -constraint that sets down the scope domain by `leq(L,H)`: both `H` and `L` are lambda arguments, eventually used by other scope introducing elements of the sentence.

Determiners function add two constraints to the USR, one that states that the main verb (labeled L) should be out-scoped by its nuclear scope, and one that ensures that its box stays within the actual scope domain, imposed by H:

```
detSem(uni, lambda(P, lambda(Q, lambda(H, lambda(L,
    merge(merge(usr([F,R,S],
                [F:drs([], [merge(drs([X], []), R)>S])),
                [leq(F,H), leq(L,S)]), P@X@H@R), Q@X@H@L)))))).
```

And here is the macro for verb negation. From the kind of constraints it introduces, it is very similar to the definitions of determiners.

```
modSem(neg, lambda(P, lambda(X, lambda(H, lambda(L,
    merge(usr([N,S],
            [N:drs([], [~S])],
            [leq(N,H), leq(L,S)]),
    P@X@H@L)))))).
```

By putting together a driver predicate, we combine our underspecified DRs with the plugging algorithm. Here it is:

```
parse:-
    readLine(Sentence),
    d(Sem, Sentence, []),
    betaConvert(merge(usr([Top,Main], [], []), Sem@Top@Main), Reduced),
    mergeUSR(Reduced, usr(D,L,C)),
    printRepresentation(usr(D,L,C)),
    findall(Drs, (plugHole(Top, L-[ ], C, []), mergeDrs(Top, Drs)), DRs),
    printReadings(DRs).
```

When calling this predicate, Prolog asks the user to type in a sentence, and then gives this sentence to the grammar, which returns a USR. This USR then is β -converted, the stipulated merges are carried out, and all possible pluggings are computed. Here is an example session:

```
?- parse.
> Every boxer or criminal knows a woman.
usr([A:drs([], [merge(drs([B], []), C)>D]),
    C:drs([], [E v F]),
    E:drs([], [boxer(B)]),
    F:drs([], [criminal(B)]),
    G:merge(merge(drs([H], []), I), J),
```

```

I:drs([], [woman(H)]),
K:drs([], [know(B,H)]),
[leq(A,L), leq(K,D), leq(C,L), leq(G,L), leq(K,J), leq(K,L)]

Readings:
1 drs([], [drs([A], [drs([], [boxer(A))]v drs([], [criminal(A)])]>drs([B], [woman(B), know(A,B)]))])
2 drs([A], [woman(A), drs([B], [drs([], [boxer(B))]v drs([], [criminal(B)])]>drs([], [know(B,A)]))])

```

Exercise 2.3.1 [hard] Running the program on coordinated NPs sometimes brings problems. Find out which cases do this, explain why, and propose a solution.

Exercise 2.3.2 [easy] Implement both Cooper storage and Keller storage for λ -DRT. Reuse as much code from Chapter 3 as possible.

Exercise 2.3.3 Discuss the possibilities to integrate an account for scope ambiguities (Storage or Hole Semantics) in the threading approach to DRS-construction.

2.4 Merging into Darkness?

This section will be devoted solely to the problems involved by merging DRSs (and indeed, underspecified DRSs). So far we took for granted that we can carry out any merge-instruction without problems. But we didn't pay attention to special occasions. What would we do, for instance, if we're ought to merge two DRSs that both contain an occurrence of the same discourse referent? Like for example:

$$(1) \quad \begin{array}{|c|} \hline x \\ \hline \text{WOMAN}(x) \\ \text{WALK}(x) \\ \hline \end{array} \oplus \begin{array}{|c|} \hline x \\ \hline \text{WOMAN}(x) \\ \text{TALK}(x) \\ \hline \end{array}$$

If we take our initial definition of Section 8.2, which says that we have to take the unions of the discourse referents and the conditions, we'll get the following DRS:

$$(2) \quad \begin{array}{|c|} \hline x \\ \hline \text{WOMAN}(x) \\ \text{WALK}(x) \\ \text{TALK}(x) \\ \hline \end{array}$$

The two discourse referents, because we gave them the same name, collapsed into one! Before thinking whether this would be a correct way of dealing with such cases, we should answer the following question: Is it possible that we, in the course of DRS construction,

arrive at such expressions? If we can answer this question negatively, we don't need to bother about it (because we're basically interested in the computational aspects of DRS-construction, and less in its theoretical considerations).

Interestingly enough, such cases appear. A classic example is coordination: 'A woman walks and a woman talks' gives exactly the initial representation (1) above. But is the reduced DRS in (2), that describes that there is a woman that walks and talks, the final representation we would like to have? Clearly not! So there is something dramatically wrong with our merge-operation, that needs to be fixed.

Exercise 2.4.1 Think of other examples in natural language (non-coordination) where this problem appears.

More to be added...

Software Summary of Chapter 2

`threadingDRT.pl` Implementation of the Johnson & Klein's threading technique for constructing discourse representation structures. (page 177)

`mainLambdaDRT.pl` Building DRSs with lambdas. (page 180)

`semMacrosLambdaDRT.pl` The semantic macros for lambda-DRT. (page 181)

`mergeDRT.pl` The predicates for the DRS-merge. (page 182)

`mainDRTU.pl` Main program for DRT Unplugged. (page 183)

`semMacrosDRTU.pl` Semantic macros for DRTU. (page 184)

Notes

The idea of DRS threading was first introduced by Johnson & Klein (1986). This elegant paper remains one of the best discussions the semantic construction in DRT and we strongly recommend it to our readers. Threading is combined with lambdas in Johnson and Kay (Johnson and Kay 1990); this paper addresses a number of important grammar engineering issues for computational semantics—required reading.

There are a number of “Box + λ ” systems. The earliest we know of is due to Frey, who proposes an implementation of DRS construction based on the f-structures of Lexical Functional Grammar (Frey 1985; Wada and Asher 1986). This bottom-up approach specifies

partial DRSs for lexical entries in much the same spirit as we do with lambdas, and includes a modest treatment of quantifier scope ambiguities. More recent approaches include λ -DRT (Bos et al. 1994; Kohlhase et al. 1996) and the system of Asher (1993); moreover, Dynamic Montague Grammar (Groenendijk and Stokhof 1990) has much in common with these systems. Most Box + λ systems make use of some kind of merge operator; an explicit merge operation seems to have been first introduced into DRT by Zeevat (1989). The theoretical aspects of Box + λ systems have been increasingly explored in recent years; Muskens (Muskens 1996) is probably the best starting point for readers interested in this topic. This (very readable) paper shows how to construct Box + λ system in standard type theory.

Combining boxes with lambdas has proved a good approach to defining and implementing compositional semantics for large scale grammars. For example, λ -DRT is used as a semantic representation formalism for Verbmobil, a large spoken-language machine-translation system (see Bos et al. 1994 for an account of this application).

Chapter 3

Pronoun Resolution

Pronoun resolution is an extremely lively and fascinating area and it has drawn strong interests from artificial intelligence, computational, and theoretical linguistics. In this chapter we will investigate and combine different approaches to pronoun resolution, by extending our Prolog programs for working with DRSs from Chapter 8. We will look at both linguistic and AI-oriented procedures for pronoun resolution.

3.1 The Nature of Pronouns

A lot of factors come together in processing pronouns. In Chapter 7 we already demonstrated how the geometrical structure of DRSs constrains the resolution of pronouns. We will certainly make use of this observation in our implementation of pronoun resolution, but it is overwhelmingly wrong to think that only the structure of discourse, as we canned them in DRSs, will determine the correct antecedent for a pronoun. Discourse structure undoubtedly plays an important part in pronoun resolution, but goes far beyond the simple structures we assigned to discourses so far.


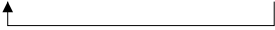
In fact, there is a lot more to pronoun resolution than the structure of discourse. Obviously, syntactic and semantic agreement play a role, depending on the natural language that we want to analyze (even closely related language as Dutch, German and English show significant differences in behavior). But also the way human beings communicate affects resolution: normally people talk about one thing at a time. World knowledge is important too, which requires, eventually, to integrate some kind of inference component in our algorithm for pronoun resolution. (We will do this in Chapter 11.)

Pronouns are, by their nature, context-sensitive expressions, and the way they appear syntactically (in some language they are simply left out—this phenomena is known as pro-drop or zero anaphora) and operate semantically varies in a number of ways. We won't treat all of this kinds, this goes far behind our purposes—instead we will concentrate

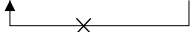
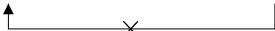
on, admittedly, the most simplest cases of pronouns, the third person singular personal pronouns ‘he/him’, ‘she/her’, and ‘it’, and their reflexive forms ‘himself’, ‘herself’, and ‘itself’.

We will only consider their *anaphoric* use (pronouns that relate back to other objects in the discourse so far processed), as opposed to their *deictic* use (relating to objects in a situation, for instance by pointing at something), or *cataphoric* use (when pronouns refer to items not yet mentioned in the text, as in ‘After he won the match, Butch left town.’). Neither will we discuss nor implement pleonastic use of pronouns as in ‘It’s about nine o’clock in the morning’.

As we have already mentioned in the introductory chapter to DRT, a discourse referent can play the role of antecedent for an anaphor only if it is accessible. Recall that DRT predicts that in the following examples anaphoric links are allowed.

- (1) A woman snorts. She collapses.

- (2) Mia ordered a five dollar shake. Vincent tasted it.


However, by changing the examples slightly, trying to establish anaphoric links results in weird or even unacceptable discourses.

- (3) Every woman snorts. She collapses.

- (4) Mia didn’t order a five dollar shake. Vincent tasted it.


Universally quantified noun phrase, such as ‘every woman’, and indefinite noun phrases in the scope of negation, introduce their discourse referent in a subordinated DRS, and are hence not accessible as antecedents for pronouns in following sentences.

But the structure of DRSs is, although a very important one, by no reason the *only* constraint on pronoun resolution. There is a number of other constraints. An obvious one is grammatical agreement: in English, pronouns come with a gender and number feature and can only refer to antecedents that carry the same features values. Consider


- (5) Mia ordered a five dollar shake. It tasted delicious.

where the pronoun it can only refer to the five dollar shake, although there is in principal another discourse referent available: the one introduced by the proper name Mia. But since this is a name for a female entity, one would normally uses the correct pronoun she

to refer her. This sound obvious, but languages behave quite differently with respect to gender agreement.

Pronouns notoriously introduce ambiguity. Often there is more than one interpretation possible.


(6) Butch threw a TV at the window. It broke.



Our program will eventually generate loads of readings (too many to deal with in an efficient natural language processing system). Some of these are more likely than others, so one way of dealing with it is adding preferences to the readings.


World knowledge might help us in some cases to select an antecedent. So, in (7) the pronoun *it* is likely to refer to a *vase* introduced in the first sentence (and not to the linguistically equally well *the wall*).

(7) Butch threw a vase at the wall. It broke.



Selectional constraints play a role here. The pronoun in *it broke* can only refer to breakable things. Here is another example that illustrates this.

(8) Butch walks into his modest kitchen. He opens the refrigerator.
He takes out a milk and drinks it.



This might sound intuitively correct, matters are however, far more complicated. Consider for example ‘drinking a cup of coffee’. Here, the cup is the argument of drinking, but selectional constraints would erroneously rule out this, as strictly spoken, only *beverages* can be drinken.

It should be clear that incorporating world knowledge is far more difficult than integrating pure linguistic constraints in the resolution component. We will therefore concentrate on purely linguistic constraints first.

Exercise 3.1.1 If you’re not a native speaker of English, find out how pronouns are used in your mother tongue, and compare it with the use of pronouns in English.

3.2 Implementing Pronoun Resolution in DRT

This section discusses a basic implementation of pronoun resolution in DRT. We will focus on the linguistic sources to constrain resolution. As noted above, we take into account the geometrical structure of DRSs. Further, we will use *gender* as sortal information to rule

out some anaphoric links. Normally, in English, ‘he’ refers to male, ‘she’ to female, and ‘it’ to non-human entities (there are clear exceptions to this rule, but we won’t deal with these in our implementation). Recall the grammar rule for handling pronouns we have:

```

pro(Pro)-->
  {
    lexicon(pro,Gender,Phrase,Type),
    proSem(Gender,Type,Pro)
  },
  Phrase.

```

This rule declares `Phrase` a pronoun (i.e., being a member of the syntactic category ‘Pro’) if it is listed in the lexicon as such, and associates the gender information (male, female, or non-human) and type information (reflexive or non-reflexive) with it. Further, the semantic macro `proSem/3` generates the semantics according to this gender and type information. (We’ll define this predicate shortly below.)

The lexicon entries are divided into non-reflexive and reflexive pronouns:

```

lexicon(pro,male,[he],nonrefl).
lexicon(pro,female,[she],nonrefl).
lexicon(pro,nonhuman,[it],nonrefl).
lexicon(pro,male,[him],nonrefl).
lexicon(pro,female,[her],nonrefl).

lexicon(pro,male,[himself],refl).
lexicon(pro,female,[herself],refl).
lexicon(pro,nonhuman,[itself],refl).

```

This is how pronouns are integrated in our fragment of English grammar. But we face a practical problem when resolving pronouns. As they are introduced in the lexicon, we don’t have access to the context and cannot determine their antecedents. How are we to cope with this?

Our strategy will be the following: we will ‘mark’ the discourse referents that are introduced by pronouns, and postpone unifying them with candidate antecedents (for example, when we have parsed the entire discourse). So let’s extend our Prolog DRS language with the notation

```

alfa(X,Type,Gender,Drs)

```

where `X` is the (still unknown) value of an antecedent discourse referent, `Type` either `refl` (reflexive) or `nonrefl` (non-reflexive), and `Gender`, either `male(X)`, `female(X)`, or `nonhuman(X)`. The `Drs` in this notation, is the DRS in which the pronoun is used.

We'll call such a term an α -DRS. For example, **she snorts** is represented by the following α -DRS:

```
alfa(Y,nonrefl,female(Y),drs([],[snorts(Y)]))
```

or in our familiar boxed notation:

α_y	<div style="border: 1px solid black; width: 100%; height: 100%; display: flex; align-items: center; justify-content: center;"> <div style="border-bottom: 1px solid black; width: 80%;"></div> </div>
	<div style="border: 1px solid black; width: 100%; height: 100%; display: flex; align-items: center; justify-content: center;"> SNORT(y) </div>

Now this should remind the reader of something—namely, the use of the `lambda` predicate! In fact, we are really just re-applying the strategy that grounded our earlier work on first order representations: we are carefully marking the ‘missing information’ which later need to be filled in. Here, of course, the missing information is of a rather special type (namely anaphoric information) so we mark it with a special new `alfa` (for ‘anaphoric’) binder rather than with `lambda`. The semantic macro that generates the α -DRS is:


```
proSem(Gender,Type,lambda(P,alfa(X,Type,Cond,P@X))):-
  compose(Cond,Gender,[X]).
```

Before we start thinking about how to implement anaphora resolution, let's try out the parser on the discourse ‘A woman snorts. She collapses.’ by posing the following Prolog query:

```
?- d(B,[a,woman,snorts,she,collapses],[]).

B = merge(
  merge(
    merge(
      drs([A],[]),
      drs([],[woman(A),female(A)])
    ),
    drs([],[snort(A)])
  ),
  alfa(B,nonrefl,female(B),drs([],[collapse(B)]))
)
```

In box-notation this Prolog-expression is translated as:



There are some DRSs in the output that still have to be merged, and there is also an α -DRS introduced by the pronoun *she*. To arrive at a pure DRS (a DRS that contains no merges or α -expressions) we need to carry out procedures that resolve these two kinds of expressions. Note that it is sensible to first resolve the pronouns before performing the merge, as otherwise we would have to revise the `mergeDrs/2` predicate (it is not specified for α -expressions). However, it turns out that it is most efficient to combine the two resolution procedures into one predicate, because for both of them we have to traverse the entire structure of the DRS; combining pronoun resolution with merge reduction kills two birds with one stone.

So let's define a predicate `resolveDrs/1` that will do this for us. Recall that a discourse referent of a DRS B_1 is accessible from DRS B_2 when B_1 subordinates B_2 , or when B_1 equals B_2 . (For the definition of subordination, see Chapter 7.) The predicate `resolveDrs/1` encodes this concept of accessibility in a fairly transparent way.

The predicate `resolveDrs/1` works with difference lists of stacks of DRSs. The input list is a stack of DRSs, of which the first item is the DRS currently under examination, and the other all subordinating it. (That means, if we work on a DRS which is the consequent of a conditional, the main DRS and the DRS of the antecedent are on this list.) The output list is a copy of the input list, but with the merges and instructions to resolve pronouns carried out. So, eventually (if `resolveDrs/1` succeeds), the output list will contain ordinary DRSs. Here are the three and only clauses that define it:

```
resolveDrs([merge(B1,B2)|A1]-[drs(D,C)|A3]):-
    resolveDrs([B1|A1]-A2),
    resolveDrs([B2|A2]-[drs(D2,C2),drs(D1,C1)|A3]),
    append(D1,D2,D),
    append(C1,C2,C).
```

```
resolveDrs([alfa(Referent,Type,Gender,B1)|A1]-A2):-
    potentialAntecedent(A1,Referent,Gender),
    resolveDrs([B1|A1]-A2).
```

```
resolveDrs([drs(D1,C1)|A1]-A2):-
    resolveConds(C1,[drs(D1,[])|A1]-A2).
```

The first clause handles the merge by resolving its argument DRSs (after all, they can contain merges or α -DRSs), and then concatenating the lists of discourse referents and conditions. Note that B_1 is put on the stack before resolving B_2 , so the discourse referents of B_1 will become automatically accessible for B_2 .

The second clause carries out pronoun resolution. It takes a DRS from the accessible DRSs hold in list `A1`, and tries to find a unification of the α -bound variable and an accessible discourse referent with the help of the following predicates:

```
potentialAntecedent(A,X,Gender):-
  member(drs(Dom,Conds),A),
  member(X,Dom),
  compose(Gender,Symbol1,_),
  \+ (
    member(Cond,Conds),
    compose(Cond,Symbol2,[Y]),
    Y==X,
    \+ consistent(Symbol1,Symbol2)
  ).
```

The predicate `potentialAntecedent/3` succeeds only if there is an accessible discourse referent (here this is `X`, a discourse referent of one of the DRSs on the stack `A`), such that the gender information is not inconsistent with the conditions imposed on discourse referent `X`. By backtracking on the `member/2` goals, it could find more potential candidates. Note that we make use of `consistent/2`, a predicate that we defined in Chapter 6.

The third and last clause of `resolveDrs/1` deals with ordinary DRSs, and uses another predicate `resolveConds/2` that works through the DRS-conditions:

```
resolveConds([~B1|Conds],A1-A3):-
  resolveDrs([B1|A1]-[B2,drs(D,C)|A2]),
  resolveConds(Conds,[drs(D,[~B2|C])|A2]-A3).

resolveConds([B1 > B2|Conds],A1-A4):-
  resolveDrs([B1|A1]-A2),
  resolveDrs([B2|A2]-[B4,B3,drs(D,C)|A3]),
  resolveConds(Conds,[drs(D,[B3 > B4|C])|A3]-A4).

resolveConds([B1 v B2|Conds],A1-A4):-
  resolveDrs([B1|A1]-[B3|A2]),
  resolveDrs([B2|A2]-[B4,drs(D,C)|A3]),
  resolveConds(Conds,[drs(D,[B3 v B4|C])|A3]-A4).

resolveConds([Cond|Conds],[drs(D,C)|A1]-A2):-
  compose(Cond,_Symbol,Arguments),
  simpleTerms(Arguments),
  resolveConds(Conds,[drs(D,[Cond|C])|A1]-A2).
```

```
resolveConds([],A-A).
```

The clauses of `resolveConds/2` work recursively through the DRS-conditions, by placing sub-DRSs on the stack of DRSs and calling back to `resolveDrs/1`.

So, what left to do is defining a driver predicate that ties together the parser, lambda-DRT, and our pronoun resolution module:

```
parse:-
  readLine(Discourse),
  d(MergeDrs,Discourse,[]),
  betaConvert(MergeDrs,Drs),
  resolveDrs([Drs]-[ResolvedDrs]),
  printRepresentation(ResolvedDrs).
```

After consulting the complete program (`mainPronounsDRT.pl`) we test it as follows:

```
?- parse.

> A woman snorts. She collapses.

drs([A],[collapse(A),snort(A),female(A),woman(A)])
```

In short, we now have a simple implementation of pronoun resolution in DRT. We will shortly see that our implementation has a number of problems. Before coming to that, the reader is asked to do the following exercises.

Exercise 3.2.1 [easy] Imagine this program processing ‘If a man eats a big kahuna burger, he enjoys it.’. How many readings do you think it will generate? Why? And how many readings will the program give for ‘Vincent knows Butch. He likes him.’?

Exercise 3.2.2 [hard] The current implementation of `resolveDrs/1` makes use of `append/3`. Explain why this is not efficient and make the program more efficient by using difference lists to concatenate lists.

A Problem

Pronouns are resolved on the basis of its gender information. A discourse referent is considered a proper candidate antecedent for a pronoun if its sortal information (as presented in the ontology of nouns, Chapter 6) is not inconsistent with the gender information specified

for the pronoun. To perform this test, we use the `consistent/2` predicate developed in Chapter 6.

This works well for male and female pronouns. But it doesn't work for nonhuman pronouns for a simple reason: the sort `NONHUMAN` is not part of our ontology. So, in the current setting, there are no sortal constraint on the pronoun 'it'. Can't we just add `NONHUMAN` to our ontology of nouns? This is difficult? What would it be? We know what it is not (human), but it can be a plant, an animal, an artifact, an abstraction, and so on.

Actually, the only thing we want to say about the sort `NONHUMAN` is that is disjoint from `HUMAN`. We can't state this in the lexicon directly (at least not in the way we designed the lexicon). The trick to incorporate disjointness information is to directly code in the predicate `consistent/2`:

```
consistent(X,Y):-
    generateIsa(I),
    generateDisjoint(I-Isa,Disjoint),
    \+ inconsistent(X,Y,Isa,[disjoint(human,nonhuman)|Disjoint]).
```

This works. Now the neuter pronouns only match with nonhuman objects. There is another—related—problem. We'll solve this in Chapter 10, but the reader is asked to get aware of this problem by trying the following exercise.

Exercise 3.2.3 [easy] Try the program (`mainPronounsDRT.pl`) on the discourse 'A boxer grows. She collapses. He dies.', and explain what is wrong with the analysis our program gives for it. Revise the semantic macro `proSem/3` to do away with this problem.


3.3 Adding Reflexive Pronouns

Another unavoidable requirement in pronoun resolution is to cope with the difference in anaphoric behavior between reflexive and non-reflexive pronouns. Reflexive pronouns appearing in object NP position of a verb phrase can only be anaphorically linked to the subject NP of this verb phrase:

- (9) Vincent goes to the toilet, and Jules enjoys himself.
-

Non-reflexive pronouns show a complementary behavior to reflexive pronouns: they *cannot* refer to their subject! The following examples make this fact clear:

(10) Vincent enters the restaurant, and Jules watches him.



In sum, pronouns obey rules of binding, and we will incorporate these constraints in our resolution algorithm. We will make use of the fact that a non-reflexive pronoun and its antecedent may not occur in the same simplex sentence, and that a reflexive pronoun and its antecedent should do so. We appeal to this linguistic rule of thumb purely on a “semantic representational” basis: non-reflexive pronouns with discourse referent x cannot be identified with a discourse referent y if both appear in a condition $P(x,y)$ introduced by the transitive verb. In contrast, reflexive pronouns *require* such a condition for a suitable antecedent.

We will refine our Prolog program with a treatment of reflexive pronouns, by appealing to an extra test in our resolution algorithm. This test is implemented by the predicate `properBinding/3`). Here is the revised clause for the predicate `resolveDrs` that integrates this new predicate:

```
resolveDrs([alfa(Referent,Type,Gender,B1)|A1]-A2):-
    potentialAntecedent(A1,Referent,Gender),
    properBinding(Type,Referent,B1),
    resolveDrs([B1|A1]-A2).
```

What does `properBinding/3` do? If the pronoun is of reflexive type, it succeeds if the binding constraints for reflexives are not violated. (This is carried out by `reflexiveBinding/2`).

```
properBinding(Type,X,Drs):-
    Type=refl,
    reflexiveBinding(X,Drs).
```

If, on the other hand, a pronoun is non-reflexive, we try to prove if the constraints for reflexive pronouns are not violated (by using the same predicate `reflexiveBinding/2`), and if this is the case, use the `!, fail` Prolog combination to let `properBinding` fail. Otherwise, it will succeed.

```
properBinding(Type,X,Drs):-
    \+ Type=refl,
    (
        reflexiveBinding(X,Drs),
        !, fail
    );
    true
).
```

In the case of a reflexive binding constraints, we would like to find a basic condition in which the two discourse referents corresponding to the subject and object (the reflexive pronoun) of the transitive verb appear. If we find such a condition we stop, otherwise we continue by recursing through the list of other conditions of the DRS.

```

reflexiveBinding(X,[Basic|Conds]):- !,
    (
        compose(Basic,_Sym,[Subj,Obj]),
        Subj==Obj,
        X==Obj, !
    );
    reflexiveBinding(X,Conds)
).

```

If we don't find such a condition, then the test fails (so we we're not allowed to use a reflexive pronoun here). There are two ways in which it could fail: when finishing checking all the DRS-conditions, or when the argument DRS of the reflexive pronoun is another α -DRS, a merge, or a negation. (In the second case the reflexive pronoun is surely not in object position, so we fail here too). Here is the code:

```

reflexiveBinding(_,[]):- fail.
reflexiveBinding(_,alfa(_,_,_,_)):- fail.
reflexiveBinding(_,merge(_,_)):- fail.
reflexiveBinding(_,_):- !, fail.
reflexiveBinding(X,drs(_,_Conds)):-
    reflexiveBinding(X,Conds).

```

This completes our way of dealing with reflexive pronouns. In principle, it works for most of the examples you will encounter in texts. But it also has its shortcomings. We only discussed the interaction of reflexive pronouns in transitive verbs, but we didn't deal with the ditransitive verb. Extending our analysis of reflexive pronouns to ditransitive verbs is not too hard, and is left as an exercise to the reader.

Exercise 3.3.1 [easy] Extend the program to handle reflexive pronouns appearing as arguments of ditransitive verbs, as in 'Vincent walks to the bar and pours himself a drink' and 'Vincent has given himself a little pep talk'.

Nor did we consider cases where reflexive pronouns appear with nouns. Reflexive pronouns are fine to use in combination with nouns that express *information*, such as 'story', 'episode', 'movie', 'picture', 'image', and so on. Here is an example that shows this behavior:

(11) Mia hates every episode about herself.

Our analysis of reflexive pronouns incorrectly rejects this sentence. Notice that for these cases, a use of a non-reflexive pronoun seems to be permitted too.

Exercise 3.3.2 [hard] Examine the example above and similar ones, explain why our program deals with them incorrectly, and sketch a solution that overcomes this problem.

3.4 The Focusing Algorithm

So far we have implemented a pronoun resolution algorithm that does its work under the supervision of syntactic and semantic constraints. This won't be enough in serious discourse analysis, though. Each new sentence in a discourse introduced new entities in the discourse that can be picked up by pronouns used later in the text. For this simple reason, the constraints that our resolution algorithm appeals to cannot avoid the problem of producing an explosion of potential readings. This leads us to one obvious choice to deal with this: adding preferences to some of the solutions we produce.

How are we going to do this? Various proposals are available (see also the Notes at the end of this chapter for a brief overview). We choose to extend our pronoun resolution program with the focusing algorithm (due to Candy Sidner). The task of the focusing algorithm is to predict an antecedent for a pronoun. The syntactic or semantic constraints we already have confirm or reject this choice.

What is focusing and focus? Given a discourse or a dialogue, the speaker or speakers talk about something, one thing at a time. The element or object of discourse on which the attention of the speakers is centered is called the *focus*. The focus is one of the very important threads that makes a series of sentences a coherent piece of text, i.e., a discourse. *Focusing* is the process of keeping track (and updating) the focus of the discourse.

What has focusing to do with pronoun resolution? Well, speakers who talk in several sentences about one focus, do not re-introduce the focus in its full syntactic form in each sentence. Instead they use pronouns (or other anaphoric expressions) under the assumption that the focus is part of the shared knowledge of the participants in the conversation.

We will distinguish two kinds of foci in a discourse: the discourse focus, and the actor focus. Here is an example that shows the use of discourse focus:

- (12) The Captain pulls a gold wrist watch out of his pocket. This watch was first purchased by Butch's great-granddaddy. It was bought during the First World War in a little general store in Knoxville, Tennessee.

The three sentences of this discourse give information about the focused element: the gold watch. It is introduced with an indefinite noun phrase in the first sentence, and referred to by 'this watch' and 'it' in the succeeding sentences. Note that there also noun phrases

introduced in the text that were not referred to in succeeding sentences (the Captain, his pocket, Butch's great-granddaddy), and further notice that each new sentence makes a reference to the discourse focus.

In addition to the *discourse focus*, focusing must take into account the actors of the discourse. We will call this the *actor focus*. Pronouns in agent position tend to refer to the actor focus, whereas pronouns in other position mostly co-refer with the discourse focus. Sometimes the actor focus can become the discourse focus when no other candidate is available, but under normal circumstances it plays a subordinated role in focusing.

The actor focus is needed besides the discourse focus because the focus of the discourse is often distinguished from the actor, and both can be referred to with pronouns. Here is an example that illustrates this:

(13) Jules grabs the burger. He takes a bite of it.

Here the discourse focus is assigned to the burger, the actor focus is Jules. The occurrence of 'he' refers to Jules, the pronoun 'it' to the burger.

The final thing to mention on focusing is that focus can shift from one object to another in the course of a text or dialogue. The following example shows how the discourse focus changes from the hammer to the chainsaw (the actor focus remains assigned to Butch):

(14) Butch picks up a big destructive-looking hammer, then discards it: not destructive enough. He picks up a chainsaw, thinks about it for a moment, then puts it back.

The last two occurrences of 'it' both refer to the chainsaw, not to the (yet out of focus) hammer.

Both the discourse and actor focus are ordinary discourse referents, but we have to distinguish them from the other referents to signify their special status. We'll do this by adding a new structure to the DRSs we have worked with so far. Ordinary DRSs consists of a set of discourse referents D and a set of conditions C , and are represented in Prolog as $\text{drs}(D, C)$. From now on we will also consider Focus-DRSs, that consists of the familiar set of discourse referents, the actor focus AF (a stack of discourse referents), the discourse focus DF (a stack of discourse referents), and the familiar set of DRS-conditions. In Prolog we will code Focus-DRSs as $\text{drs}(D, AF, DF, C)$, where both the actor and discourse focus are lists of discourse referents.

We will transfer this newly structured DRSs also to our familiar box-notation. So the Focus-DRS for 'Butch picks up a hammer. He discards it.' is

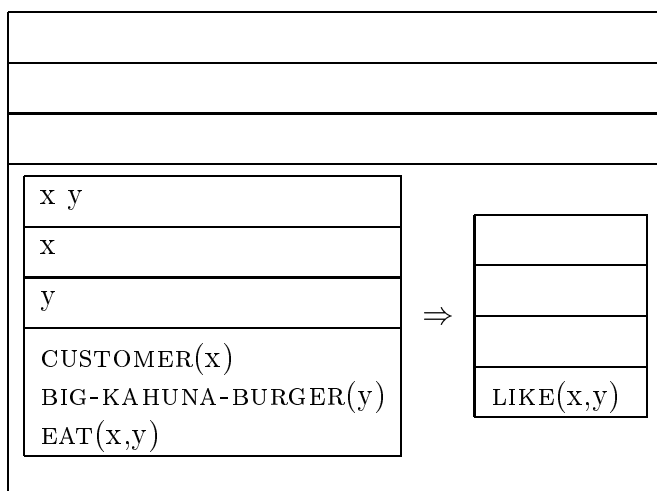
x y
x
y
x=BUTCH HAMMER(y) PICK-UP(x,y) DISCARD(x,y)

where x is the actor focus (Butch), and y the discourse focus (the hammer picked up by Butch). Of course, in a progressing discourse, there can be more of these labels, as the focus of attention can shift. Continuing this discourse with ‘He picks up a chainsaw’ triggers a shift of the discourse focus.

x y z
x
z y
x=BUTCH HAMMER(y) PICK-UP(x,y) DISCARD(x,y) PICK-UP(x,z) CHAINSAW(z)

The discourse focus now has two elements on its stack. But note that z is the first element of the stack, representing the *current* discourse focus.

Note that Focus-DRSs can play a role on the subordinated level as well. For example, the conditional ‘If a customer eats a big Kahuna burger, he likes it.’ leads to the construction of the following Focus-DRS:



This is the basic idea of focusing—we will get a more detailed understanding while integrating the focusing algorithm directly into our DRT-implementation for pronoun resolution. This algorithm consists of three procedural steps:

1. choose foci based on the first sentence of the discourse
2. use foci to resolve pronouns
3. update the foci on the basis of the current DRS

Step 1 guesses the foci on the basis of syntactic and semantic structure of the first sentence. Step 2 involves an extension to the resolution algorithm we have so far—we will give preference to potential antecedents that are in focus. Step 3 retains or resets the focus, using the information how pronouns got resolved. The focus will be moved to another discourse referent if the discourse referent previously in focus is not referred to in the current sentence.

To choose the foci, we will make an initial prediction based on simple heuristics. These, we call it the expected foci, will either be confirmed or rejected by step 3 of the algorithm. We might be wrong in choosing an expected discourse focus (there might be more than one choice), but still the prediction helps us to rank the readings on preference.

Step 1: Choosing Foci

Let's first have a closer look at step 1. How do we identify foci? For the actor focus, this turns out to be rather simple: it is the agent of the sentence if there is one (when the agent of the next sentence is a pronoun, the actor focus is usually the antecedent).

This idea is implemented straightforwardly as follows (we will discuss the predicate `role/3` below):

```


expectedActorFocus(drs(Dom, Conds), Focus) :-
    member(Focus, Dom),
    role(agent, Focus, Conds).

```

Determining the discourse focus is somehow more problematic. It boils down to characterizing what is talked about in a discourse. Language seems to have different phenomena for marking focus, deviating strongly from language to language. In English, there-insertion, clefts, and the prosodic structure of the sentence indicate what the sentence is about. Further, usage of grammatical roles give a clue.

The modest fragment of English that we cover does not include clefts or there-insertion constructions (actually, clefts are rarely used in the beginning of a discourse). So we will assign the focus solely on the basis of semantic categories of the verb. Often, the discourse focus is the thematic role of the verb, and the actor focus the agent. Here is an example:

(15) A hand lays an envelope full of money on the table in front of Butch. Butch picks it up.



The first sentence introduces four discourse referents that are candidates for discourse and actor focus. The hand is proposed as actor focus, but not retained, as there is no reference to it in the follow-up sentence. The discourse focus is the envelope with money, because that's what Butch picks up. (Note that world knowledge would not exclude that the pronoun 'it' refers to the table. But this is not the preferred antecedent, as predicted by the focusing theory.) Copular sentences form an exception to this rule, they take their subject as discourse focus.

Translated into Prolog this gives us:

```

expectedDiscourseFocus(drs(Dom, Conds), Focus) :-
    member(Focus, Dom),
    role(theme, Focus, Conds).

```

The thematic roles are determined by traversing through the DRS looking for a condition that matches the properties for a discourse referent for being in agent or theme position. Here are the clauses of `role/3` that deal with Focus-DRSs and DRSs:

```

role(Role, X, drs(_, _, _, Conds)) :-
    role(Role, X, Conds).

role(Role, X, drs(_, Conds)) :-
    role(Role, X, Conds).

role(Role, X, merge(B1, B2)) :-

```



```

role(Role,X,B1);
role(Role,X,B2).

role(Role,X,alfa(?,?,?,B)):-
  role(Role,X,B).

```

And here are the conditions that deal with the DRS-conditions:

```

role(Role,X,[~ B|_]):-
  role(Role,X,B).

role(Role,X,[B1 > B2|_]):-
  role(Role,X,B1);
  role(Role,X,B2).

role(Role,X,[B1 v B2|_]):-
  role(Role,X,B1);
  role(Role,X,B2).

role(Role,X,[_|Conds]):-
  role(Role,X,Conds).

role(agent,X,[C|_]):-
  agent(X,C).

role(theme,X,[C|_]):-
  theme(X,C).

```

We use a simple rule of thumb for determining the agent role of a discourse referent: if it is the subject of a transitive verb. Therefore, in our Prolog code that implements this rule, we appeal to our lexicon:

```

agent(Focus,Cond):-
  compose(Cond,Sym,[Subject,_]),
  Focus==Subject,
  lexicon(tv,Sym,_,inf).

```

Similarly, we implement thematic positions, as either being the subject of an intransitive verb, the object of a transitive verb, or the subject of a copular sentence. This is coded as follows:

```

theme(Focus,Cond):-
  compose(Cond,Sym,[Subject]),
  Focus==Subject,
  lexicon(iv,Sym,_,inf).

theme(Focus,Cond):-
  compose(Cond,Sym,[_ ,Object]),
  Focus==Object,
  lexicon(tv,Sym,_,inf).

theme(Focus,Cond):-
  compose(Cond,'=',[Subject,Object]),
  var(Object),
  Focus==Subject.

```

Exercise 3.4.1 Extend `theme/2` by adding clauses for ditransitive verbs.

Step 2: Pronoun Resolution using Focus

The focus information is used to resolve pronouns as formulated by the following two rules:

If the pronoun appears in a sentence thematic relation other than agent, choose the discourse focus as antecedent.

If the pronoun appears in agent position, choose the actor focus as antecedent.

These rules apply only if other constraints do not rule out these choices. Here is the implementation of the first rule (pronoun appearing in agent-position):

```

resolveDrs([alfa(X,Type,Gender,B1)|A1]-A2,R1-[Type:X|R2]):-
  role(agent,X,B1), !,
  (
    currentActorFocus(A1,AF),
    potentialAntecedent(A1,X,Gender),
    X == AF,
    properBinding(Type,X,B1)
  );
  currentDiscourseFocus(A1,DF),
  potentialAntecedent(A1,X,Gender),
  X == DF,

```

```

    properBinding(Type,X,B1)
;
    potentialAntecedent(A1,X,Gender),
    properBinding(Type,X,B1)
),
resolveDrs([B1|A1]-A2,R1-R2).

```

The order of the three disjuncts mirror the preferences in a clear way: first try to actor focus, then the discourse focus, and then any other discourse referent.

The implementation of the second rule (the pronoun is not in agent position) is set up in a similar way, but with the preference order changed:

```

resolveDrs([alfa(X,Type,Gender,B1)|A1]-A2,R1-[Type:X|R2]):-
(
    currentDiscourseFocus(A1,DF),
    potentialAntecedent(A1,X,Gender),
    X == DF,
    properBinding(Type,X,B1)
;
    currentActorFocus(A1,AF),
    potentialAntecedent(A1,X,Gender),
    X == AF,
    properBinding(Type,X,B1)
;
    potentialAntecedent(A1,X,Gender),
    properBinding(Type,X,B1)
),
resolveDrs([B1|A1]-A2,R1-R2).

```

The predicates that define the current actor and discourse focus are specified as follows:

```

currentDiscourseFocus(DRSs,Focus):-
    member(drs(_,[Focus|_],_),DRSs), !.

currentActorFocus(DRSs,Focus):-
    member(drs(_,[Focus|_],_),DRSs), !.

```

What these predicates do is find a DRS on the stack that have a non-empty focus. The cut prevents backtracking after succeeding in finding a focus.

Step 3: Updating the Focus

After each analysis of a sentence, the focus information needs to be updated. This applies also to sentence-internally, for example in conditional sentences. For updating the focus we will use the predicates `expectedActorFocus` and `expectedDiscourseFocus` for an initial sentence of a discourse; otherwise we use the information of how the pronouns in the current sentence are resolved. Recall from Step 2 that we represent the resolution information as a list with terms `Alfa-Type:Antecedent`.

The body of the focusing update function can be designed as follows. First update the actor focus, and then the discourse focus.

```
updateFocus (Anaphora,DRSs1,Updated):-
    updateActorFocus (DRSs1,DRSs2),
    updateDiscourseFocus (Anaphora,DRSs2,Updated).
```

The actor focus is updated by taking the expected actor focus (if there is one).

```
updateActorFocus ([drs (D,AF,DF,C) | A], Updated):-
    expectedActorFocus (drs (D,C),Focus), !,
    addActorFocus (Focus, [drs (D,AF,DF,C) | A], Updated).
```

```
updateActorFocus (DRSs,DRSs).
```

```
addActorFocus (Focus,DRSs,Updated):-
    (
        currentActorFocus (DRSs,Current),
        Current==Focus, !,
        Updated=DRSs
    );
    DRSs=[drs (D,AF,DF,C) | A],
    Updated=[drs (D,[Focus|AF],DF,C) | A]
).
```

This is how the discourse focus gets updated. If there is a previous discourse focus and there is an anaphor binding it, then the focus stays like it was. If there is no previous discourse focus, then the current focus is the expected discourse focus. In all other cases, we leave it as it is.

```
updateDiscourseFocus (Anaphora,DRSs,Updated):-
    currentDiscourseFocus (DRSs,Current),
    member(_:A,Anaphora), A==Current, !,
```

```

Updated=DRSs.

updateDiscourseFocus(_Anaphora, [drs(D, AF, DF, C) | A], Updated) :-
    expectedDiscourseFocus(drs(D, C), Focus), !,
    addDiscourseFocus(Focus, [drs(D, AF, DF, C) | A], Updated).

updateDiscourseFocus(_, DRSs, DRSs).

addDiscourseFocus(Focus, DRSs, Updated) :-
    (
        currentDiscourseFocus(DRSs, Current),
        Current==Focus, !,
        Updated=DRSs
    );
    DRSs=[drs(D, AF, DF, C) | A],
    Updated=[drs(D, AF, [Focus | DF], C) | A]
).

```

Here is how the focus updating predicates are integrated into our resolution algorithm (for a merge of DRSs, for Focus-DRSs, and ordinary DRSs):

```

resolveDrs([merge(B1, B2) | A1] - [drs(D, AF, DF, C) | A5], R1-R3) :-
    resolveDrs([B1 | A1] - A2, R1-R2),
    updateFocus(R2, A2, A3),
    resolveDrs([B2 | A3] - A4, R2-R3),
    updateFocus(R3, A4, [drs(D1, AF1, DF1, C1), drs(D2, AF2, DF2, C2) | A5]),
    append(D1, D2, D),
    append(AF1, AF2, AF),
    append(DF1, DF2, DF),
    append(C1, C2, C).

resolveDrs([drs(D, AF, DF, C) | A1] - A3, R1-R2) :-
    resolveConds(C, [drs(D, AF, DF, []) | A1] - A2, R1-R2),
    updateFocus(R2, A2, A3).

resolveDrs([drs(D, C) | A1] - A3, R1-R2) :-
    resolveConds(C, [drs(D, [], [], []) | A1] - A2, R1-R2),
    updateFocus(R2, A2, A3).

```

Resolving DRS-conditions is exactly like that in the basic pronoun resolution algorithm, except for conditionals. Here we also need to introduce a focus update. Because in a ‘if

S1 then S2' sentence, pronouns appearing in the second sentence (S2) will be resolved also depending on the information of sentence (S1).

```
resolveConds([B1 > B2|Conds],A1-A5,R1-R4):-
  resolveDrs([B1|A1]-A2,R1-R2),
  updateFocus(R2,A2,A3),
  resolveDrs([B2|A3]-[B4,B3,drs(D,AF,DF,C)|A4],R2-R3),
  resolveConds(Conds,[drs(D,AF,DF,[B3 > B4|C])|A4]-A5,R3-R4).
```

Finally, we provide a driver predicate that makes it easy to experiment with our focusing algorithm. After each sentence it analyses, it prints the DRS, after which the user can extend the discourse with a new sentence.

```
parse:-
  parse(drs([],[])).

parse(Old):-
  readLine(Discourse),
  d(SDrs,Discourse,[]),
  betaConvert(SDrs,Drs),
  resolveDrs([merge(Old,Drs)]-A,[]-Anaphora),
  updateFocus(Anaphora,A,[New]),
  printRepresentation(New),!,
  parse(New).
```

Exercise 3.4.2 [easy] Try the program with the following example: ‘Butch picks up a hammer. He discards it. He picks up a chainsaw, and he likes it.’ Explain the focus shifts.

Exercise 3.4.3 [hard] Write a driver that generates all readings, ordered on preference. Assign scores to the readings, according to the focus rules.

Software Summary of Chapter 3

`mainPronounsDRT.pl` The basic implementation of pronoun resolution for DRs.
(page 186)

`bindingDRT.pl` Predicates that check binding of pronouns in DRT. (page 188)

`mainFocusDRT.pl` The focus algorithm. (page 190)

Notes

The binding behavior of pronouns (coreference within a sentence) is discussed in the linguistic literature in great detail, starting with Chomsky's Binding Theory (Chomsky 1988). A rich source of information on this topic is provided in Reinhart's work (Reinhart 1983). For more inspiration on (especially reflexive) pronouns the reader is encouraged to check out the first two chapters of Fauconnier's book on Mental Spaces (Fauconnier 1985).

Pronoun resolution procedures that go beyond the sentence level are widely available in the field of Artificial Intelligence. Pioneers in defining heuristics for pronoun resolution are Terry 'the blocks world' Winograd (Winograd 1972) and Eugene Charniak. Charniak's dissertation (Charniak 1972) includes a proposal for pronoun resolution included syntactic (gender, number) as well as semantic properties. He was the first that made aware the need of common-sense knowledge reasoning to determine antecedents of pronouns. See also (Charniak and Wilks 1976; Charniak and McDermott 1985).

Hobbs presents a syntactic and semantic approach to pronoun resolution (Hobbs 1986). The syntactic approach offers what he called the 'naive' algorithm, which traverse the parse tree to determine potential antecedents, thereby using constraints known from transformational grammar to deal with reflexive pronouns and cataphora. Within the semantic approach he proposes different strategies using inference to determine antecedents. This paper also give a modest but useful overview of the problems that appear with pronoun resolution.

Wada and Asher included in their DRS construction implementation a module that takes care of pronoun resolution (Wada and Asher 1986). Our basic implementation borrows a lot from this work as it: (1) marks pronouns in the DRS by a special condition; (2) associates discourse referents with gender and number information stored in a tree structure that parallels the subordination relation of a DRS; (3) resolves this tree by finding the closest antecedent that matches gender and number; and (4) uses additional constraints (e.g., the reflexive/non-reflexive distinction).

The focusing algorithm as presented here is originally due to Candy Sidner (Sidner 1986). Some of the ideas in this paper—that appeared in the late seventies—are strikingly similar to the principles underlying DRT: using an intermediate level of representation, where what Sidner calls *specifiers* play a similar role as the discourse referents in DRSs. Sidner motivates the rules in her focusing algorithm by a number of examples, and also discusses the role of plural pronouns in focusing, and that of the demonstrative pronouns 'this' and 'that'. The implementation that we gave here is a slight variant of Sidner's model. We deal with sentence-internal discourse structure (which Sidner doesn't), but, on the other hand, simplified the set of update rules.

The focus model is a precursor of what is called the *centering* approach to pronoun resolution. Centering elaborates on the basic idea of focusing. This approach distinguishes *forward looking centers* (i.e., in the DRT-terminology, a set of available discourse referents,

ordered on grammatical relations or other criteria) and the *backward looking center* (the highest ranked forward looking center of the previous utterance, to which a pronoun or other anaphoric expressions in the current utterance refers to). Useful pointers are (Grosz, Joshi, and Weinstein 1995) and (Kehler 1997). Beaver (unpublished) proposes a model that integrates centering within a formal update semantics framework.

Focus information, as we plainly added it to the evolving DRS of a text, adds another dimension to the structure of the discourse. A more elaborate model of discourse structure can be found in the classical paper of Grosz and Sidner (Grosz and Sidner 1986), or a formalized ‘boxed’ version, in Asher’s Segmented DRT (Asher 1993).

Pronouns are a motley crew of context-sensitive expressions. We’ve only dealt with third person singular pronouns, and didn’t consider other uses of anaphoric expressions, such as impersonal use of pronouns, plural anaphora, or pronouns that refer to abstract entities such as eventualities or propositions. A recommended DRT-based introduction to plural anaphora is Chapter 4 of (Kamp and Reyle 1993). The undoubtedly best overview of propositional and event-type anaphora, as well as the related phenomenon known as verb phrase ellipsis, is (Asher 1993), also presented in DRT.

Chapter 4

Presupposition Resolution

In this chapter we are going to deal with a stubborn obstacle to the computational analysis of discourse: presuppositions. Roughly speaking, presuppositions are pieces of information that are taken for granted in a context. In this chapter we will learn what the typical problems associated with presuppositions are, and how to deal with them in Discourse Representation Theory.

Presupposition is a huge topic. It has been discussed in detail by both the philosophical and linguistic communities, and we won't be able to cover the wide variety of analyses such discussions have inspired. In fact, we shall concentrate exclusively on the DRT based approach of Rob van der Sandt. There are a number of reasons for this choice. Empirically, it is one of the most successful approaches to the problems posed by presuppositions (for example, the *projection problem* introduced below). Conceptually the theory is also extremely natural: it is simple, allows the important notion of *accommodation* to be modeled very directly, and links presupposition and anaphora in a suggestive way. All in all, Van der Sandt's account both draws on and illuminates the intuitions about discourse processing that DRT embodies. Finally, the approach is interesting from a computational perspective: it demands an extension of our earlier implementation of pronoun resolution, and access to inference methods.

4.1 Introducing Presuppositions

In the previous chapter we worked with elements of natural language that were obviously context dependent, namely pronouns. Pronoun interpretation can only be determined with help of the previous discourse, and we gave a DRT-based account of anaphora resolution. But there is another natural language phenomenon, possibly related to pronoun use, that behaves in a similar way. Consider the following example sentences, uttered out of the blue:

- (1) The couple that won the dance contest was pleased.
- (2) Jody loves her husband.
- (3) Vincent regrets that Mia is married.

There is something interesting about these examples: they force us to take something for granted if we want to accept them as natural contributions to a discourse or conversation. For (1), we need to suppose that there actually is a couple that won the dance contest, for (2) that Jody has a husband, and for (3) that Mia is married.

Moreover, given a context containing contrary information, these sentences do not make sense at all. For example, the following sequences are unacceptable:

- (4) Jody has no husband. ? Jody loves her husband.
- (5) Mia is not married. ? Vincent regrets that Mia is married.

Furthermore, whatever we are dealing with here, it behaves completely differently from ordinary entailment. Not only does ‘Jody loves her husband’ take it for granted that Jody has a husband, but so does its negation ‘Jody does not love her husband’; note that

- (6) Jody has no husband. ? Jody does not love her husband.

is just as bad as (4). Whatever this phenomenon is, it seems to follow its own set of laws.

In fact we are dealing with *presupposition*, and this is a very common natural language phenomenon that any analysis of discourse has to cope with. The sentence ‘Jody loves her husband’ *presupposes* that Jody is married, as does ‘Jody does not love her husband’. Any context in which Jody is not married is inappropriate for accepting either the sentence ‘Jody loves her husband’ or its negation.

Is there a systematic way to compute the presuppositions of a sentence? To answer this question we need to investigate how presuppositions come to life. It turns out that, in most cases, they are *lexically* triggered. (At least this is the case for English, certain prosodic or syntactic structures can also introduce presuppositions, though we won’t deal with such triggers in this book.) In (1), the definite article ‘the’ is the source of the presupposition (the reader might verify that by changing ‘the’ into ‘a’), in (2) it is the possessive ‘her’, and in (3) the verb ‘regret’. Such lexical items are called *presupposition triggers*: once they are used, they introduce a presupposition, thereby putting extra constraints on the context.

At first glance this problem seems easy to deal with. Surely it is just a matter of going through our lexicon, marking all presupposition triggers and the presupposition they induce, and then when analyzing a sentence, checking if all presuppositions triggered by the

complete sentence are fulfilled by the discourse. And in fact, this *is* a good way of thinking of it, and we will handle presupposition triggers along these lines in this chapter. However, matters will not always be so simple. There are three issues we need to pay attention to: the *binding* problem, the *projection problem*, and *accommodation*. All three issues are well known in the (vast) literature on presuppositions.

Let's first turn to the *binding* problem. An example like

A boxer nearly escaped from his apartment.

should clarify what is at stake here. The trigger 'his' induces the presupposition that someone has an apartment. But it does not presuppose that just anyone has an apartment, nor that some boxer or other owns an apartment. No, it is *the boxer we are actually talking about* who has an apartment. That is, the existentially quantified NP ties together two types of information—ordinary factual information, and presuppositional information. As these two types of information obey different laws (recall the way presuppositions survive negation) it is no trivial matter to tie them together, and many otherwise interesting accounts of presupposition have been shipwrecked on precisely this rock. However, as we shall soon see, Van der Sandt's DRT-based account handles it rather elegantly. Because DRT was designed right from the start to deal with *integrating* information, the binding problem simply won't be an issue.

The *projection problem* has to do with complex sentences. We have already seen that

Mia's husband is out of town.

is a sentence presupposing that Mia has a husband. But if we use this sentence to build the more complex

If Mia has a husband, then her husband is out of town.

then the result does *not* presuppose that she has a husband. Neither does

If Mia is married, then her husband is out of town.

although

If Mia dates Vincent, then her husband is out of town.

clearly does. The moral is clear. In complex sentences we need to be careful when dealing with presupposition triggers, as sometimes subparts of complex sentence carry presuppositions which are “neutralized” in the main sentence. A natural solution to this problem would seem to be to take linguistic context into account, and this is also what we are going to do, again within the framework of DRT.

Finally a few words on *accommodation*. This can be thought of as a way of obtaining a robust and realistic treatment of presupposition. Analyzing

Vincent informed his boss.

yields the elementary presupposition that Vincent has a boss. If the discourse built up so far is incompatible with Vincent having a boss (maybe he is a free-lancer), then we could just say: “Hey, what are you telling me, Vincent has no boss”, and not accept this contribution to the discourse. But what if we don’t have a clue whether Vincent has a boss or not? There are two obvious ways to proceed here: the simplistic way is simply to refuse to accept the utterance because its accompanying presuppositions are not supported by the context. This might seem adequate, but it isn’t what we would typically do in such a situation—instead we would try to add the presupposition to the context. The addition process is called *accommodation*, and it is best viewed as a repair strategy. We are trying to build up a picture of a particular situation, but it turns out that (somehow or other) we seem have missed some piece of information. OK—let’s try and incorporate this information the best we can and carry on. This is an inherently robust and realistic way of dealing with presuppositions and it is the option we shall explore in this book.

4.2 Dealing with Presupposition in DRT

We are now going to show how to deal with these three problems in DRT using a method due to Rob van der Sandt. In his view, presuppositions are essentially extremely *rich* pronouns. Like ordinary pronouns, they make use of the notion of accessibility in the way we are familiar with. But presuppositions are rich in a way that pronouns aren’t: they have descriptive content. Thus presuppositions can’t simply introduce discourse referents, so what do they do? Here’s Van der Sandt’s answer: *presuppositions introduce new DRSs*.

This is a strikingly simple and promising answer. DRSs are pictures of the evolving discourse; what better way to incorporate presuppositional information than by adding a new picture? The key word here is *incorporate*. Simply by placing the DRS yielded by a presupposition into a larger DRS we locate it in a web of discourse referents regulated by accessibility constraints. That is, by proceeding in this way we automatically *contextualize* presuppositional information. It seems plausible that this is a good way of coping with the binding, projection, and accommodation problems, and this in fact is the way things turn out. Let’s turn to the details.


To begin with, we have to carry out two tasks:

1. select presupposition triggers in the lexicon; and
2. indicate what they presuppose.

Task 1 appeals to our empirical knowledge of presupposition. For our fragment of English, we classify the definite article, possessive constructions, and proper names as presupposition triggers. For task 2 we extend the core DRS language with an α operator. The α operator takes two normal DRSs as arguments forming a new special kind of DRS; it is simply a marker (or if you prefer, a *semantic feature*) that indicates that one DRS is presupposed information for the other. So, we will write $B_1 \alpha_i B_2$ to say that B_1 is presupposed information for B_2 , or put differently, B_2 presupposed B_1 . We use the index i to express that it is the principal discourse referent that is anaphoric, as there could be more discourse referents in the presupposed DRS (this is the discourse referent for which we will try to find an antecedent). We will call B_1 the α -DRS (this is the “anaphoric” DRS for which we have to resolve), B_2 is a DRS that remains at its place after resolution.


Exercise 4.2.1 Find natural language examples where there are more discourse referents in the α -DRS (the DRS that contains presupposed information).

Once we’ve dealt with the lexicon, what’s the next step? Well, using our familiar DRS construction technology, we build a DRS for the sentence. If the sentence contains presupposition triggers, the DRS we output will generally contain α -DRSs. Here’s the DRS we obtain for ‘The woman collapses’:



The best way to view this DRS is as an ordinary DRS, but an ordinary DRS marked as being unresolved with respect to presuppositions. (Another way to look at this is to say that it is a kind of underspecified DRS; we’ll see why later.) What are we going to do with this DRS?

First we merge this new DRS with the DRS that represents the discourse so far; this merging process takes place while the presuppositions are still unresolved. For example, suppose that the previous utterance was ‘A woman snorts’. Then after merging we obtain:



Only after merging of the previous DRS with the new DRS we attempt to resolve the presuppositions. We recursively travel through this new merged DRS and, for each α -marked DRS we encounter, we try to find a suitable ‘anchor’ to resolve to. That is, *we try to match the context of the α -DRS with that of superordinated DRSs*. Intuitively this is a natural thing to do; after all, presupposed information is supposed to be available in the previous context.

Let’s see how this works. In our example, we only have one presupposition trigger with α -DRS:

x
WOMAN(x)

To resolve α -DRS, there is only one superordinated DRS (with the candidate antecedent discourse referent y) for resolution.

y
WOMAN(y)
SNORT(y)

First we identify the discourse referents x and y by adding the condition $x=y$ to the α -DRS:

y
WOMAN(y)
SNORT(y)

 \oplus

x
WOMAN(x)
x=y

 α_x

COLLAPSE(x)

Then we move the information of the α -DRSs to the superordinated DRS:

y
WOMAN(y)
SNORT(y)

 \oplus

x
WOMAN(x)
x=y

 α_x

COLLAPSE(x)

Finally we replace the operator α_x by a merge-instruction and arrive at:

y x
WOMAN(y)
SNORT(y)
WOMAN(x)
x=y

 \oplus

 \oplus

COLLAPSE(x)

This is a DRS that contains two instructions for merging. We can reduced it to:

y x
WOMAN(y)
SNORT(y)
COLLAPSE(y)
WOMAN(x)
x=y

But note that we had have could carry out the identification of x and y explicitly (by unification) and compute the following, equivalent DRS:

y
WOMAN(y)
SNORT(y)
COLLAPSE(y)

In the other examples we will discuss, and also in our Prolog-implementation of presupposition resolution, we will perform this unification operation instead of adding an equality condition to the DRS, as this will lead to much simpler DRSs.

In short, we have successfully dealt with the presupposition induced by the definite article ‘the’ by identifying the discourse referent x it introduced with the woman-denoting discourse referent in the preceding context.

That’s the basic idea, but things don’t always go this smoothly. Sometimes we can’t find the presupposed information in the preceding context. (Maybe, we missed a bit of a conversation; and anyway, people typically make different assumptions about what the assumed context actually is.) To deal with such cases we make use of *accommodation*: if we can’t link our elementary presuppositions to a suitable element in the context, we don’t give up. Instead we simply add the required background information.

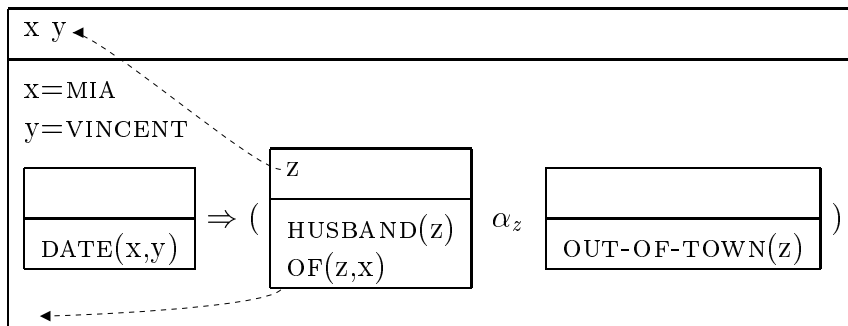
Here’s an example. Consider the sentence ‘If Mia dates Vincent, then her husband is out of town’. Concentrating only on the trigger ‘her husband’, we get:

x y							
x=MIA							
y=VINCENT							
<table border="1"> <tr> <td></td> </tr> <tr> <td>DATE(x,y)</td> </tr> </table> \Rightarrow <table border="1"> <tr> <td>z</td> </tr> <tr> <td>HUSBAND(z)</td> </tr> <tr> <td>OF(z,x)</td> </tr> </table> α_z <table border="1"> <tr> <td></td> </tr> <tr> <td>OUT-OF-TOWN(z)</td> </tr> </table>		DATE(x,y)	z	HUSBAND(z)	OF(z,x)		OUT-OF-TOWN(z)
DATE(x,y)							
z							
HUSBAND(z)							
OF(z,x)							
OUT-OF-TOWN(z)							

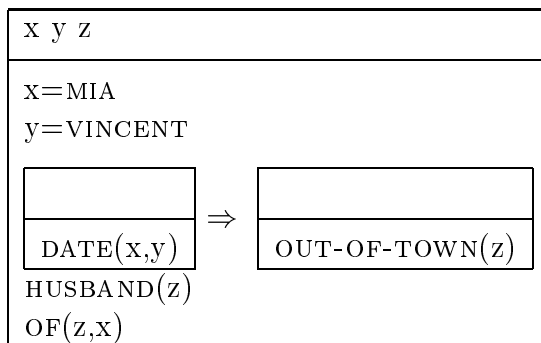
Assuming this is the first DRS we have to process (that is, that the DRS built up so far is still empty), there is no candidate DRS for matching the presupposed information that Mia has a husband, which is coded by the following DRS:

z
HUSBAND(z) OF(z,x)


In such cases our robust way of dealing with discourse comes into action: we *accommodate* the presupposed information. Accommodation is similar to resolution. The big difference is that we *don't* identify discourse referents of the α -DRS with a superordinated DRS. For the rest, accommodation is exactly the same as resolution, and we move the α -DRS to a superordinated DRS:



As we moved to the presupposed information to the outermost DRS, we will call this *global accommodation*. After replacing the α -operator by \oplus and carrying out the merges we get the following, final, DRS:



So far so good: our “try to match the presupposed information, and if that fails, accommodate it at the global DRS” strategy has survived its first test. But we were lucky with our example; consider instead ‘If Mia is married, then her husband is out of town’. Does this presuppose that Mia has a husband? No, it doesn’t—but applying our global accommodation strategy yields:



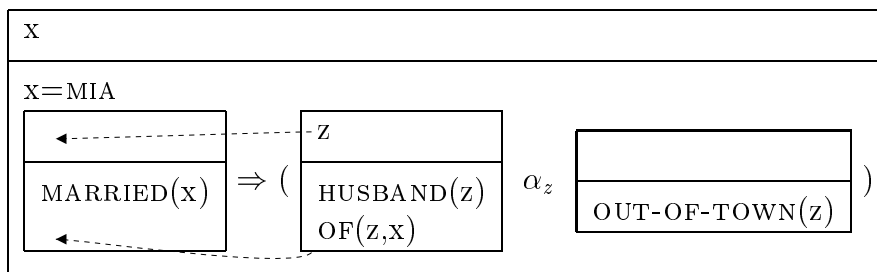
The content of this DRS is paraphrased as

- (7) Mia has a husband. If she is married, then he is out of town.
-

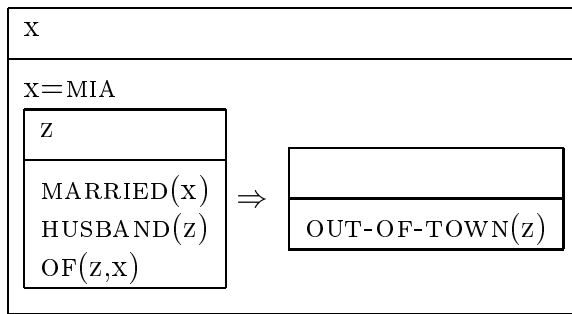
which is a spectacularly wrong result, as it asserts that Mia is married.

But wait on—who said we needed to perform accommodation in the outermost DRS? For sure, all discourse referents in the universe of outermost DRS are accessible—but one of the most pleasant aspects of DRT is that it shows how pictures can be nested one inside the other, and tells us which discourse referents are available at arbitrary levels of nesting. Why not pursue a more flexible attitude towards accommodation and allow accommodation at other DRS levels as well? There is certainly no technical barrier to doing this, and arguably it is the natural option to explore: bluntly insisting that all information has to be added globally doesn't seem to reflect the myriad possibilities offered by natural language.

Let's see how this more flexible strategy can help us with our previous example. In fact it licenses accommodation on a non-global, intermediate, level:

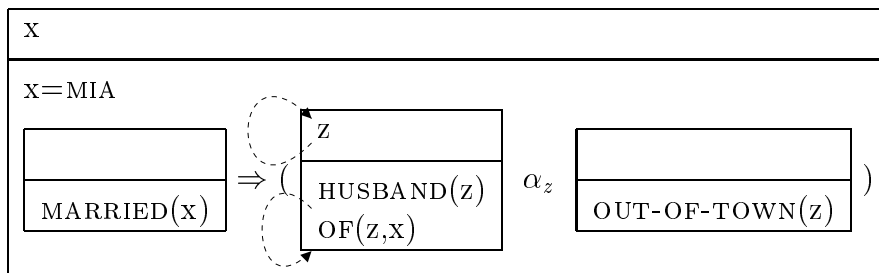


This gives us (again, after replacing the α -operator by a merge) the following reading:

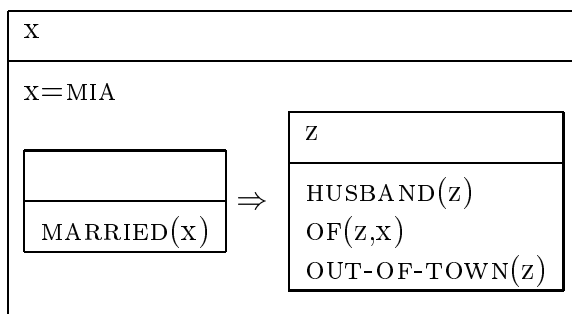


Here the presupposed information is accommodated in the left-hand-side DRS of the conditional. This is the reading we would like to get: it doesn't demand the existence of a husband for Mia.

Indeed, in this example, there is another possibility for accommodation. It is called local accommodation, and sketched as follows:



Local accommodation is actually not performing any move operation. The presupposed information just stays where it originated. After replacing α_z by a merge we get:



That's the basic idea—but there are two further points that need to be mentioned. First, it's clear that we can't perform accommodation arbitrarily; we've already seen what happens when we try to inappropriately accommodate at the global level. Rather, just like ordinary pronoun resolution, accommodation is subject to a variety of acceptability constraints. We will briefly discuss the required constraints shortly; Chapter 11 is largely devoted to implementing them. Second, accommodation is a non-deterministic process.


That is, we are free to accommodate where we like, so long as we don't violate an acceptability constraint. Note that this has linguistic consequences: we are predicting that presupposition is an additional source of ambiguity. Incidentally, this is why we claimed earlier that α -DRSs could be viewed as underspecified representations: essentially α offers a compact encoding of accommodation possibilities.

Summing up, both here and in Chapter 11, we are working with the following non-deterministic algorithm:

1. Generate a DRS for the input sentence with the tools of Chapter 2, with all elementary presuppositions given as an α -DRS.
2. Merge this DRS with the DRS of the discourse so far processed.
3. Traverse the DRS, and on encountering an α -DRS try to
 - (a) link the presupposed information to an accessible antecedent (partial match);
 - (b) if that fails, accommodate the information to a superordinated level of discourse.
4. Remove those DRS from the set of potential readings that violate the acceptability constraints.

What sort of constraints do we appeal to at Step 4? Just like pronoun resolution, accommodation is subject to two kinds of constraints. The first is essentially syntactic, or formal (that is, like the accessibility constraint on pronoun resolution it simply imposes a well-formedness condition on the resulting DRSs—only simple syntactic checking is required to check whether it is fulfilled). This constraint is called the *free variable check*. The following example shows what this involves; it also shows the kinds of ambiguities predicted by the Van der Sandtian approach to presupposition.

Consider the sentence 'Every woman likes her husband'. There are three candidate accommodation strategies: local, intermediate, and global. Here are the DRS these options yield:



In fact, both the local and the intermediate are acceptable: neither violates an acceptability constraint.

Exercise 4.2.2 Spell out the readings given by local and intermediate accommodation.

But now consider the DRS obtained by global accommodation. Clearly it is rather strange; one of the conditions (namely, $\text{OF}(z,x)$) ascribes a property to a discourse referent (namely, x) that does not occur in the universe. So to speak, x is a free discourse referent, and in fact if we used the translation function fo to compile this DRS into a first-order formula, we would obtain not a sentence but a formula in which the variable x occurred free. The free variable constraint simply rules out such DRSs.

Exercise 4.2.3 Although we can appeal to fo to define what we mean by a ‘free discourse referent’, we certainly don’t need to do this. Give a direct definition of this concept.

But what are the other acceptability constraints? Here matters get a lot more interesting (and a lot more difficult) for these constraints require us to perform *inference*, and a great deal of inference at that. Let’s briefly note the two most obvious constraints we should observe.

First, we need to obey the *consistency* constraint; obviously it would be foolish to accommodate information in such a way that the resulting DRS was inconsistent. Note, however, that we can’t test for inconsistency simply by inspecting the DRS; we have to do some real work. Second, we should obey the *informativity* constraint. It would be redundant to accommodate information in such a way that the resulting DRS actually followed from the DRS we started with. That is, we should only carry out accommodation if it results in something genuinely new. Once again, however, it requires inference to determine whether or not this constraint is obeyed. Actually, Van der Sandt requires *local* versions of consistency and informativity as well, and like the constraints just mentioned, these require inference to determine whether or not they are fulfilled. We shall discuss these constraints further in Chapter 11.

Here’s our plan. In the remainder of this chapter we implement the basic architecture described above. That is, we shall learn how to generate candidate DRSs (as we shall see, this part of the process is closely related to the basic pronoun resolution method implemented in the previous chapter) but we *won’t* perform the filtering demanded by the acceptability check. In Chapter 11 we shall add the required checks. We will do so by using `drs2fo1`, our DRS to first-order logic compiler, and then exploiting standard first-order inference techniques.

4.3 Presupposition Resolution in Prolog

Implementing accommodation has two main steps. First we must take care of the lexicon: we need to say exactly what sort of presuppositions the various entries trigger. Secondly we need to think about how the actual accommodation process is to be performed.

Extending Lexical Macros

All presupposition triggers in our fragment of English introduce an α -condition. We shall use a prolog term

```
alfa(X,Type,ADrs,Drs),
```

where *ADrs* is the α -DRSs, and *Drs* its scope. (So, we altered the representations for pronouns of the previous chapter into one where all we did is replacing the gender information by the α -DRS. In fact, we will use this notation for pronouns as well.) Here is the semantic macro for the determiner ‘the’:

```
detSem(def,lambda(P,lambda(Q,alfa(X,def,merge(drs([X],[ ]),P@X),Q@X))))).
```

That is, the scope of the presupposition operator α is restricted to the ‘noun part’ of the definite article. Compare this with the rule for the determiner ‘a’:

```
detSem(indef,lambda(P,lambda(Q,merge(merge(drs([X],[ ]),P@X),Q@X))))).
```

Proper names are also viewed as presupposition triggers. This give us an important practical advantage: we don’t need an additional device that moves the discourse referents to the outermost DRS, thereby allowing accessibility for later pronouns: global accommodation deals with this automatically. So here is the revised macro for proper names:

```
pnSem(Sym,Gender,lambda(P,alfa(X,name,drs([X],[X=Sym,Cond]),P@X))):-
  compose(Cond,Gender,[X]).
```

Since linking is part of our presupposition resolution mechanism, pronouns are dealt with in the same spirit. So pronouns also add an α , with a DRS that only has a discourse referent and no conditions.

```
proSem(Gender,Type,lambda(P,alfa(X,Type,drs([X],[Cond]),P@X))):-
  compose(Cond,Gender,[X]).
```

Possessives like ‘her X’ and ‘his X’ are analyzed as ‘the X of her’ and ‘the X of him’. Syntactically, they are determiners. Here is the semantic macro; note that this is a rather complex entry, for we have nested α ’s.

```
detSem(poss(Gender),lambda(P,lambda(Q,
  alfa(X,nonrefl,drs([X],[Cond]),
    alfa(Y,nonrefl,merge(drs([Y],[of(Y,X)]),P@Y),Q@X))))):-
  compose(Cond,Gender,[X]).
```

With one exception, these are the presupposition triggers we deal with in our grammar fragment of English. The remaining one is left as an exercise for the reader.

Exercise 4.3.1 Fill in the the interface definition for the determiner ‘another’. (It is already in the lexicon). What kind of presupposition does it trigger?

Exercise 4.3.2 Adjectives like ‘old’ and ‘new’ are under one reading presupposition triggers. Spell out the presuppositions for these kind of triggers and specify their lexical entries.

Implementing Linking and Accommodation

The initial output of our parser is a DRS with α -DRSs just like our implementation for pronoun resolution. This will be the input to the predicate `projectDrs/1`, that performs the linking and accommodation operations and finally returns a normal DRS.

Notably, `projectDrs/1` works along the same DRT-principles as `resolveDrs/1` from the previous chapter. (The reader who doesn’t know about this predicate, is advised to learn about from the previous chapter, before proceeding). As its closely related `resolveDrs/1`, `projectDrs/1` works with stacks of DRSs. The three clauses of `projectDrs/1` resemble those of `resolveDrs/1`:

```
projectDrs([merge(B1,B2)|A1]-[drs(D3,C3)|A3]):-
    projectDrs([B1|A1]-A2),
    projectDrs([B2|A2]-[drs(D1,C1),drs(D2,C2)|A3]),
    append(D1,D2,D3),
    append(C1,C2,C3).
```

```
projectDrs([alfa(X,Type,B1,B3)|A1]-A5):-
    projectDrs([B1|A1]-[B2|A2]),
    resolveAlfa(X,B2,A2-A3),
    projectDrs([B3|A3]-A4),
    addBindingInfo(X,Type,A4-A5).
```

```
projectDrs([drs(D1,C1)|A1]-A2):-
    projectConds(C1,[drs(D1,[])|A1]-A2).
```

But there is more to it. We want to modify the stack more drastically, for we want to allow accommodation. This comes to play in the clauses that deal with the DRS-conditions, for which we introduce the predicate `projectConds/2` (indeed, as you might have expected, this predicate mirrors `resolveConds/2` of the previous chapter).

Accommodation is allowed at subordinated levels of discourse, and at the global level. But not inside nested merges. So we have to be careful when to apply accommodation, and

avoid spurious accommodation. To manage accommodation, we introduce a special object on the stack, coded as `pre(Pres)`, that is a place holder for an accommodation site. Here, `Pres` is a list of presuppositions (naturally coded as DRSs) that are accommodated at this site (at the beginning, this is the empty list). As we will see later, `resolveAlfa` will fill this slot.

The clauses of `projectConds/2` make clear where these accommodation sites are introduced. The clause for negation, for instance, puts an empty site `pre([])` on the stack, which is filled with locally accommodated presuppositions `pre(Pres)` after processing the negated box, which are then, finally, accommodated to the negated box itself by use of `accommodate/2`. (We will come to this predicate shortly.)

```
projectConds([~B1|Conds],A1-A3):-
  projectDrs([B1,pre([])|A1]-[B2,pre(Pres),drs(D,C)|A2]),
  accommodate(Pres,B2-B3),
  projectConds(Conds,[drs(D,[~B3|C])|A2]-A3).

projectConds([B1 > B2|Conds],A1-A4):-
  projectDrs([B1,pre([])|A1]-A2),
  projectDrs([B2,pre([])|A2]-[B4,pre(P2),B3,pre(P1),drs(D,C)|A3]),
  accommodate(P1,B3-B5),
  accommodate(P2,B4-B6),
  projectConds(Conds,[drs(D,[B5 > B6|C])|A3]-A4).

projectConds([B1 v B2|Conds],A1-A4):-
  projectDrs([B1,pre([])|A1]-[B3,pre(P1)|A2]),
  projectDrs([B2,pre([])|A2]-[B4,pre(P2),drs(D,C)|A3]),
  accommodate(P1,B3-B5),
  accommodate(P2,B4-B6),
  projectConds(Conds,[drs(D,[B5 v B6|C])|A3]-A4).

projectConds([Basic|Conds],[drs(D,C)|A1]-A2):-
  compose(Basic,_Symbol,Arguments),
  simpleTerms(Arguments),
  projectConds(Conds,[drs(D,[Basic|C])|A1]-A2).

projectConds([],A-A).
```

Gathering all accommodated presuppositions on accommodation sites is very useful indeed. It is a means to compute the presuppositions for each sub-DRS. After the gathering is done, the presuppositions are added to the actual DRS they belong to with the help of `accommodate/2`. There is nothing special to this predicate: it recursively cycles through

the list of accommodated DRSs and adds them to the actual DRS, by straightforward use of `append/3`.

```
accommodate([],B-B).

accommodate([drs(D1,C1)|Presups],drs(D2,C2)-B):-
    append(D1,D2,D3),
    append(C1,C2,C3),
    accommodate(Presups,drs(D3,C3)-B).
```

This is how, basically, the implementation for presupposition projection is arranged. But we haven't explained how the actual resolving and accommodation takes place. It is the predicate `resolveAlfa/4` that carries out this task. (It is part of the clause of `projectDrs` that handles α -DRSs).

The first two clause of `resolveAlfa` deal with linking (partially matching) the α -DRS with one of the DRS on the stack. (Here we use `matchDrs/5`, which we explain in a minute.) By the recursive nature of this predicate, partial matching is allowed with any of the DRSs on the stack by use of backtracking.

```
resolveAlfa(X,AlfaDrs,[drs(D,C)|Others]-[New|Others]):-
    matchDrs(X,AlfaDrs,drs(D,C),New).

resolveAlfa(X,AlfaDrs,[AnteDrs|Others]-[AnteDrs|NewOthers]):-
    resolveAlfa(X,AlfaDrs,Others-NewOthers).
```

But not only linking is possible. Accommodation is allowed on the sites on the stack marked by the term `pre/1`. This is straightforward too, as the argument of `pre/1` is a list of DRSs, and we can use the list operator and simply add the α -DRS to the accommodation site. Again, with the recursive set-up of this predicate, all possible accommodation sites are given a chance under backtracking.

```
resolveAlfa(_,AlfaDrs,[pre(A)|Others]-[pre([AlfaDrs|A])|Others]).

resolveAlfa(X,Alfa,[pre(A1)|Others]-[pre([New|A2])|Others]):-
    select(Drs,A1,A2),
    matchDrs(X,Alfa,Drs,New).
```

Finally, we have to code the partial match of DRSs. The predicate `matchDrs/5` matches two DRSs yielding a new one, by unifying the anaphoric referent `X` with one of the discourse referents of the antecedent DRS (so, it will fail if the domain of this DRS is empty), and then take the union of the discourse referents and DRS-conditions to form the new resolved DRS.


```

matchDrs (X, drs (D1, C1), drs (D2, C2), drs (D3, C3)) :-
    member(X, D2),
    mergeLists (D1, D2, D3),
    mergeLists (C1, C2, C3),
    consistentConditions (X, Conds).

```

This is done with the help of `mergeLists/3`. It works very much like `append/3`, but it prevents that multiple occurrences of items will turn up in the resulting list. This is how it is coded:

```

mergeLists ([], L, L).

mergeLists ([X|R], L1, L2) :-
    member(Y, L1),
    X==Y, !,
    mergeLists (R, L1, L2).

mergeLists ([X|R], L1, [X|L2]) :-
    mergeLists (R, L1, L2).

```

Matching of two DRSs only succeeds if the resulting DRS is internally consistent. We'll use the `consistent` predicate from Chapter 6 for a first check (for complex cases, we have to rely on deep automated reasoning—see Chapter 11). A DRS is consistent with respect to a discourse referent `x`, if it is not the case that we can find two conditions on `x` that are not consistent. Here is the implementation in Prolog:

```

consistentConditions (X, Conds) :-
    \+ (
        member(Cond1, Conds),
        member(Cond2, Conds), \+ Cond1=Cond2,
        compose(Cond1, Symbol1, [Y]), Y==X,
        compose(Cond2, Symbol2, [Z]), Z==X,
        \+ consistent(Symbol1, Symbol2)
    ).

```

And this finishes our basic implementation of Van der Sandt's algorithm. Of course we follow our computational tradition and design a driver predicate that allows easy input and provides us with readable output. The following driver computes all readings that our implementation generates (using the built-in Prolog predicate `findall/3`):

```

parse :-

```

```

readLine(Discourse),
d(Drs1, Discourse, []),
betaConvert(Drs1, Drs2),
findall(Drs4, (
    projectDrs ([Drs2, pre([])]-[Drs3, pre(P)]),
    accommodate(P, Drs3-Drs4),
    checkBinding(Drs4)
),
    Readings),
printReadings(Readings).

```

Let's see how our implementation works for simple discourses. Here is an example of a simple definite. Two readings are generated, one caused by accommodation, and one caused by linking.

```
?- parse.
```

```
> A woman smokes. The woman collapses.
```

```
Readings:
```

```
1 drs([A],[collapse(A),smoke(A),woman(A)])
2 drs([A,B],[woman(A),collapse(A),smoke(B),woman(B)])
```

Experimenting with this basic implementation will probably surprise you. Even for simple examples, the program will produce a high number of readings. Here is an example that has only two anaphoric expressions (the proper name 'Vincent' and the pronoun 'it') but produces four readings—our implementation seems to over-generate readings:

```
?- parse.
```

```
> Vincent eats a big kahuna burger. Every criminal likes it.
```

```
Readings:
```

```
1 drs([A,B],[male(A),A=vincent,drs([C],[criminal(C)])>drs([D],[nonhuman(D),like(C,D)]),eat(A,B),bkburger(B)])
2 drs([A,B],[male(A),A=vincent,drs([C,D],[nonhuman(C),criminal(D)])>drs([], [like(D,C)]),eat(A,B),bkburger(B)])
3 drs([A,B],[male(A),A=vincent,drs([C],[criminal(C)])>drs([], [like(C,B)]),nonhuman(B),eat(A,B),bkburger(B)])
4 drs([A,B,C],[male(A),A=vincent,nonhuman(B),drs([D],[criminal(D)])>drs([], [like(D,B)]),eat(A,C),bkburger(C)])
```

Exercise 4.3.3 Inspect the readings above and explain why they are generated by our implementation of Van der Sandt's algorithm.

Note that the problem is not purely one of producing too many readings, but one of producing readings which are unequally acceptable, although not impossible to rule out at all. We will attack this problem in the next session, by using a series of linguistic heuristics to judge the readings.

Exercise 4.3.4 Change the program in such a way that it doesn't support intermediate accommodation (for example, accommodation in the antecedent DRS of an implicational condition).

Free Variable Trapping

Nested presuppositions can have free variables in their α -DRS, which can potentially escape from their original binder when accommodated on a more global level. One of the notorious examples of this branch of problems is the solutions we generate for ‘Every boxer enjoys his big kahuna burger’. The presupposition trigger ‘his big kahuna burger’ gives rise to local, intermediate, and global accommodation. Under global accommodation, a free variable will appear in the main DRS. In fact, this reading is not available at all.

This is serious problem, and we have to deal with it in a satisfactory way. But note that this problem is not entirely new, we discussed it in the description of Van der Sandt’s algorithm, and proposed the free variable check as a solution.

Exercise 4.3.5 Add a free variable check on the generated solutions and test the modified programs on examples like ‘every boxer likes his wife’.

4.4 Optimizing the Algorithm

The current implementation of Van der Sandt’s algorithm faces the problem that it generates a high number of readings, of which some are more likely to appear than others, but this is not stated in the output.

This comes from being too tolerant with regard to linking and accommodation. Given a presupposition or anaphor, we are equally happy with accommodation or linking. Similarly, we have equal opinion whether it is resolved locally or globally. We just don’t care. But this is not the intuitive way of dealing with it. Linking, for instance, is preferred to accommodation. And global accommodation is preferable to local accommodation. Moreover, these heuristics depend on the kind of trigger we deal with: proper names behave quite differently from pronouns or definite descriptions in this respect.

Let’s have a look how the different anaphoric constructs behave, and what kind of heuristic we could attach to it.

Proper Names

Proper names are—just like any other presupposition trigger—allowed to accommodate anywhere. But this is far too loosely formulated. Local or intermediate accommodation is not something one would expect for proper names (unlike other definite descriptions), and neither is linking on a non-global accommodation (although it is not to be totally excluded). A simple way of dealing with this is to give global accommodation and linking a much higher preference to local operations.

Exercise 4.4.1 Try to think of examples in English where proper names accommodate or link at a non-local level of discourse structure.

Related to this problem, is the fact that, at least for proper names, accommodate should only be an option if linking is impossible. Otherwise, for a discourse ‘Vincent knows Butch. Butch is a boxer.’, there will be a reading where there are two Butches, each having a different discourse referent. Again, this is not impossible, but it certainly is the less preferred option.

Pronouns

Pronouns have less accommodation power than proper names. On subordinate levels accommodation is perhaps totally impossible (well, try to think of an example), and on the global level it is possible (you might have missed part of the text or conversation) but it is certainly less preferred than linking.

Exercise 4.4.2 Discuss in what cases one would perhaps allow global or local accommodation of pronouns.

Definite Descriptions

These easily tend to accommodate, also on non-global levels. But of course, linking, if possible, is to be preferred to accommodation.

Adding Scores

What we are going to do is to attach to each reading a judgement. This judgement is a score between 0 and 1, the higher the score, the more preferred this reading is.

Intuitively, the score says something about the processing costs involved for a reading. The ideal score of 1 means that there are no difficulties whatsoever to process the sentence under this reading. The extreme low score of 0, on the contrary, means that it is impossible to process the sentence under this reading.

So what we do is simple and straightforward: we extend the implementation of presupposition projection with one extra parameter: the score. At the beginning it gets assigned the value 1, and as we proceed our analysis, it will get updated by the different operations we have to perform. Accommodation will change the score in a negative way, but linking, since it is preferred to accommodation, in a less negative way.

We will use the notation **ScoreIn-ScoreOut** in the predicates, where **ScoreIn** is the score before the predicate is proved, and **ScoreOut** is the new value of the score if the proof was

successful. All the clauses for `projectDrs` and `projectConds` get this extra argument. For instance:

```
projectDrs([alfa(X,Type,B1,B3)|A1]-A5,S1-S4):-
  projectDrs([B1|A1]-[B2|A2],S1-S2),
  resolveAlfa(X,Type,B2,A2-A3,S2-S3),
  projectDrs([B3|A3]-A4,S3-S4),
  addBindingInfo(X,Type,A4-A5).
```

The value of the score is changed when we resolve the α -DRS, because that's exactly the point what kind of anaphoric construct we deal with, where it is resolved, and how it is resolved. All these factors play a role for determining the new score.

Here is how the scores are changed. Basically, it is done by multiplying the old score with a factor (ranging from 0 to 1), and assigning the result to the new score. (Note that this ensures that the resulting value will never outrange the scope 0-1.)

```
resolveAlfa(X,Type,AlfaDrs,[drs(D,C)|Others]-[New|Others],S1-S2):-
  global(Others),
  matchDrs(X,AlfaDrs,drs(D,C),New),
  (Type=refl, S2 = S1;
   Type=nonrefl, S2 = S1;
   Type=def, S2 = S1;
   Type=name, S2 = S1).
```

```
resolveAlfa(X,Type,AlfaDrs,[drs(D,C)|Others]-[New|Others],S1-S2):-
  nonglobal(Others),
  matchDrs(X,AlfaDrs,drs(D,C),New),
  (Type=refl, S2 is S1 * 1;
   Type=nonrefl, S2 is S1 * 1;
   Type=def, S2 is S1 * 0.5;
   Type=name, S2 is S1 * 0.2).
```

```
resolveAlfa(_,Type,Alfa,[pre(A)]-[pre([Alfa|A])],S1-S2):-
  global([pre(A)]),
  (Type=nonrefl, S2 is S1 * 0.5;
   Type=def, S2 is S1 * 0.9;
   Type=name, S2 is S1 * 0.9).
```

```
resolveAlfa(_,Type,Alfa,[pre(A)|Others]-[pre([Alfa|A])|Others],S1-S2):-
  nonglobal([pre(A)|Others]),
  (Type=nonrefl, S2 is S1 * 0.1;
   Type=def, S2 is S1 * 0.7;
   Type=name, S2 is S1 * 0.2).
```

```

resolveAlfa(X,Type,Alfa,[pre(A1)|Others]-[pre([New|A2])|Others],S1-S2):-
    global([pre(A1)|Others]),
    deleteFromList(Drs,A1,A2),
    matchDrs(X,Alfa,Drs,New),
    (Type=refl, S2 = S1;
     Type=nonrefl, S2 = S1;
     Type=def, S2 = S1;
     Type=name, S2 = S1).

resolveAlfa(X,Type,Alfa,[pre(A1)|Others]-[pre([New|A2])|Others],S1-S2):-
    nonglobal([pre(A1)|Others]),
    deleteFromList(Drs,A1,A2),
    matchDrs(X,Alfa,Drs,New),
    (Type=refl, S2 is S1 * 1;
     Type=nonrefl, S2 is S1 * 1;
     Type=def, S2 is S1 * 0.5;
     Type=name, S2 is S1 * 0.2).

```

Our revised predicate `resolveAlfa` is now turned into a *tuner* for presupposition projection. If you don't like the value, just take a screwdriver and change them what you think is more accurate (maybe you want to add more anaphoric types).

To distinguish between global and local accommodation, `resolveAlfa` has more clauses than its predecessor. It uses two auxiliary predicates to determine the level of DRS given the current situation of the stack. The global level of DRS is a situation where the stack contains exactly one `pre/1` term (and this is the last item on the stack). This is coded by `global/1`:

```

global([pre(_)]).
global([drs(_,_)|Stack]):- global(Stack).

```

The stack determines a non-global level of discourse if there at least two occurrences of `pre/1` on it. This is coded by `nonglobal/1`:

```

nonglobal([pre(_),drs(_,_)|_]).
nonglobal([drs(_,_)|Stack]):- nonglobal(Stack).

```

The driver has to be adapted too. All we do is add the calculated score to the generated readings.

```

parse:-

```

```

readLine(Discourse),
d(Drs1,Discourse,[]),
betaConvert(Drs1,Drs2),
findall((Score:Drs4),(
    projectDrs([Drs2,pre([])]-[Drs3,pre(P)],1-Score),
    accommodate(P,Drs3-Drs4),
    checkBinding(Drs4)
),
    Readings),
printReadings(Readings).

```

Here is the same example as we run with our first implementation, but now the scores are made visible:

```

?- parse.

> A woman smokes. The woman collapses.

Readings:
1 1:drs([A],[collapse(A),smoke(A),woman(A)])
2 0.9:drs([A,B],[woman(A),collapse(A),smoke(B),woman(B)])

```

Our implementation with scores gives a perfect score of 1 to the reading where ‘the woman’ is linked to the discourse referent introduced by ‘a woman’, and a score of 0.9 to the reading where ‘the woman’ accommodates a discourse referent. How nice!

Now for the other example:

```

?- parse.

> Vincent eats a big kahuna burger. Every criminal enjoys it.

Readings:
1 0.9:drs([A,B],[male(A),A=vincent,drs([C],[criminal(C)])>drs([],enjoy(C,B)),nonhuman(B),eat(A,B),bkburger(B)])
2 0.45:drs([A,B,C],[male(A),A=vincent,nonhuman(B),drs([D],[criminal(D)])>drs([],enjoy(D,B)),eat(A,C),bkburger(C)])

```

Exercise 4.4.3 Compare the readings for ‘Vincent eats a big kahuna burger. Every criminal enjoys it.’ with the ones that we got with our first version of presupposition resolution. Which readings are suppressed and why

Exercise 4.4.4 [easy] Change the program such that it allows local accommodation of proper names.

Software Summary of Chapter 4

`mainPresupDRT.pl` Implementation of Van der Sandt's presupposition resolution algorithm for DRSs. (page 196)

`mainPresupScoreDRT.pl` Implementation of Van der Sandt's presupposition resolution algorithm for DRSs, with score computation for each generated reading. (page 200)

`resolvePresup.pl` Implementation of Van der Sandt's presupposition resolution algorithm for DRSs. (page 197)

`resolvePresupScore.pl` Implementation of Van der Sandt's presupposition resolution algorithm for DRSs, with score computation for each generated reading. (page 201)

`matchDRT.pl` Code for partial matching of DRSs. (page 199)

`semMacrosPresupDRT.pl` Definitions of the semantic macros for the presupposition projection programs. (page 204)

Notes

By far the most important reference for the work of this chapter is Van der Sandt's classic paper: (Van der Sandt 1992). This contains a wealth of examples, and locates the approach, both intellectually and historically, in the complex terrain of presupposition theory. For a recent overview of this terrain from a rather different perspective, see David Beaver's survey article: (Beaver 1997).

Van der Sandt attributes the term accommodation for presupposition to Lewis (Lewis 1979). Accommodation is a hot topic in presupposition theory. Although most tend to agree with global and local forms of accommodation, *intermediate* accommodation is heavily disputed.

Chapter 5

Acceptability Constraints

This chapter is devoted to implementing the acceptability constraints imposed by Van der Sandt on presupposition resolution and accommodation. We begin by explaining how to implement the consistency and informativity constraints. We then describe and implement the local versions of these constraints that Van der Sandt also imposes.

In essence this chapter revolves around a single theme: integrating semantic construction and inference. As we shall show, it is possible to develop an interesting interactive discourse system by hooking together a theorem prover (here we use the widely available Otter system) with the semantic construction tools we have been developing in this book.

5.1 Maxims of Conversation

Humans communicate successfully on the basis of a simple conversational principle: cooperation between speaker and hearer. Normally, speakers do not give less or more information than is appropriate, they don't say something they believe is false or already known. Two very general maxims which we plan to implement in our DRS-toolkit are:

1. each contribution to the discourse should be consistent with it, and
2. each contribution should introduce new information.

The first maxim, 'be consistent!', is an obvious constraint on discourse contributions. One should not contradict oneself. Obeying this rule, the following discourse are unacceptable:

Mia is a man. Mia is a woman.
Jody is married. Jody does not have a husband.

Incidentally, note that the inconsistency of these discourses is not a fact of pure logic; it depends on additional information, namely that men are not women, that women are not men, that Jody is a woman, and that married women have husbands.

The second maxim, ‘be informative!’, is a constraint on the newness of information. Every contribution to the discourse should introduce new information. This too is quite natural. Discourses that violate this principle are:

Jody is a boxer. Jody is a boxer.
Mia is a married. She has a husband.

Again, note that while the first inference is purely logical, the second inference hinges on our knowledge that Mia is a woman, and that married women have husbands.

It is important to realize that these maxims are not always obeyed by humans. (Obviously, liars do not.) They are, however, the default setting in which humans analyze utterances. Violation of consistency appears when somebody *corrects* himself/herself, as for instance in:

Jody is a boxer. No, Jody is *not* a boxer.

The informativity constraint (which, perhaps, has less priority than consistency) can be violated as well, and assertions may well be part of the established conversation. Utterances may be repeated for emphasis:

Jody is *not* a boxer, she is *not* a boxer.

Consistency and informativity are two of the acceptability constraints Van der Sandt imposes on accommodation, and we are going to build these constraints into a small interactive discourse system. The system is set up as follows. The user can build up a discourse by inputting a sentence to the system. Using the DCG-parser and the DRS-construction program, we build a DRS for this sentence, and integrate it with the DRS from the previous contribution. Presuppositions will be handled as in Chapter 10, but now we’re going to use a theorem prover to check which accommodation possibilities obey the maxims.

5.2 Choosing a Theorem Prover

Which theorem prover are we are going to use? For a start, standard theorem provers are not able to process DRSs directly, so we need to translate DRSs to first-order logic. But we know from Chapter 7 that this is no big deal: we have a compiler, `drs2fol` which handles this efficiently.

Can we use our theorem prover from Chapter 5? Mostly, yes—but not always. DRS languages translate to first-order logic *with equality*. Our theorem prover cannot deal with equality, so we have to look for something different.

Exercise 5.2.1 [hard] Actually, for the grammar fragment under consideration, most equality conditions in DRSs can be compiled away in the translation to first-order logic. Write a predicate that does this.

Exercise 5.2.2 [easy] Think of examples that invoke equality conditions that cannot be compiled out.

Fortunately, there are a number of freely available provers that handle equality. One of the most widely available is Otter, and we will use Otter here for our theorem proving work.

We'll hook up Otter to our Prolog-programs by defining an interface predicate called `callTheoremProver/3`. (This is done on the basis of a Unix platform, so you might want to redefine this interface if you're using another operating system.) The interface works as follows: with `fol2otter/2` the axioms and formulas are written to the file `temp.in` in Otter-syntax. (We won't discuss this little translator. It simply translates our first-order notation to the notation Otter expects and is a straightforward piece of code.) Then, using the Sicstus in-built predicate `shell/2`, we start Otter, supplying it with the input file, and let it output `temp.out`. The second argument of `shell/2` returns the state of the process. For some mysterious reasons, this state variable returns the value 26386 if Otter found a proof, and 26624 when it didn't succeed in finding a proof, and that's why we'll implement the interface like this:

```
callTheoremProver(Axioms,Formula,Proof):-
    fol2otter(Axioms,Formula),
    shell('./otter < temp.in > temp.out',X),
    (X=26368,Proof=yes;X=26624,Proof=no).
```

Checking for consistency boils down to translating the DRS to first-order logic—call this formula Φ —and then checking that Φ is consistent. For the theorem proving case this means we need to find out whether $\neg\Phi$ is valid: if it is, then Φ is inconsistent, and if it is not, then Φ is consistent. For the model building case this involves finding out whether Φ is satisfiable: if it is not, then Φ is inconsistent. This is summarized in the following table:

Now, Otter is a refutation based theorem prover, which means that we have to give it the *negation* of what we are trying to prove, thus we have to give it $\neg\neg\Phi$, or equivalently Φ , as input.

How do we test for informativity? Well, if the new DRS follows from the DRS of the previous discourse, then the new DRS does not encode new information. Given B_{new}

Table 5.1: Consistency checking for theorem provers and model builders

$\Phi \setminus \neg\Phi$	valid	?	invalid
satisfiable	–	consistent	consistent
?	inconsistent	?	consistent
not satisfiable	inconsistent	inconsistent	–

encoding the new information, and B_{old} the information contributed so far, we translate a DRS with solely the condition $B_{old} \Rightarrow B_{new}$ to first-order logic, say to a formula Φ . If Φ is a theorem or not satisfiable, we have no new information. Or for a complete picture of this setting, consult Table 5.2:

Table 5.2: Informativity checking for theorem provers and model builders

$\neg\Phi \setminus \Phi$	valid	?	invalid
satisfiable	–	informative	informative
?	not informative	?	informative
not satisfiable	not informative	not informative	–

One other thing needs to be done. We are *not* interested merely in what can be proved without supporting assumptions. Quite the contrary: we are interested in what can be proved given background knowledge (for example, that married woman have husbands, that Mia is a woman, and so on). Using `backgroundKnowledge(Drs)`, an interface to certain facts of the world coded in first-order formulas the following code for our consistency check takes such knowledge into account.

```
consistent(NewDrs):-
  backgroundKnowledge(Chi),
  drs2fol(NewDrs,Phi),
  callTheoremProver(Chi,Phi,Proof),
  (Proof=yes, !, fail; true).
```

There is a similar predicate for informativity. Here, too, we need to take background knowledge into account. So the actual coding for the check on informativity is as follows:

```

informative(drs([],[]),_):-!.
informative(OldDrs,NewDrs):-
    backgroundKnowledge(Chi),
    drs2fol(OldDrs,Phi),
    drs2fol(NewDrs,Psi),
    callTheoremProver(Chi,~(Phi > Psi),Proof),
    (Proof=yes,!,fail;true).

```

What about background knowledge? For a start, we use the information our ontology of noun objects gives us (see Chapter 6). When we discuss the local constraints, we'll extend the background knowledge with facts that are not covered by the ontology information (for instance, that married women have husbands).

```

backgroundKnowledge(Formulas):-
    generateOntology(Formulas).

```

Summing up, we now have predicates which take ordinary DRSs as input and tell us whether the acceptability constraints are violated or not. So now we're ready to combine our DRS-construction tools with our inference-tools, and integrate them into a small interactive system.

5.3 Chatting with Curt

The best way to show how the acceptability constraints work is to implement them in a small interactive program. We'll develop such a program that accepts utterances of the user and comments on them according to the maxims of conversation—we'll call this system Curt. While Curt is short for "Clever use of reasoning tools", it really owes its name to the brief (and indeed, rather rude) replies it gives.

Two (mutually recursive) predicates implement Curt: `curtInput` and `curtOutput`. The user's input is handled by `curtInput/1`:

```

curtInput(Readings):-
    readLine(Input),
    curtOutput(Input,Readings).

```

By `readLine/1` the user's input is read, and passed on (together with the current list of readings, represented as DRSs) to `curtOutput`. The predicate `curtOutput/2` checks the input and accordingly gives an appropriate response. If Curt finds a parse, it checks whether the readings it obtained are consistent (otherwise it objects), and informative

(otherwise it will say so). The readings that remain are passed on to `curtInput/1`, and the user can enter a new sentence, which will be integrated into the DRS encoding the previous utterances.

```

curtOutput(Input, Readings):-
    d(MergeDrs, Input, []),!,
    betaConvert(MergeDrs, ReducedDrs),
    getReadings(ReducedDrs, Readings, PotentialReadings),
    consistentReadings(PotentialReadings, ConsistentReadings),
    (
        ConsistentReadings=[],
        curtSays('No! I do not believe that!'),!,
        curtInput(Readings)
    );
    informativeReadings(ConsistentReadings, InformativeReadings),
    (
        InformativeReadings=[],
        curtSays('Yes, I knew that!'),!,
        curtInput(Readings)
    );
    curtSays('Ok. '),
    curtInput(InformativeReadings)
).

```

The DRSs of the input sentence are computed with the help of the predicate `getReadings`. This predicate uses the module for presupposition projection (Chapter 10) to deal with pronouns and presuppositions.

```

getReadings(ProtoDrs, Readings, PotentialReadings):-
    findall((OldDrs, NewDrs),
        (
            member(OldDrs, Readings),
            projectDrs([merge(OldDrs, ProtoDrs), pre([])]-[Drs, pre(P)]),
            accommodate(P, Drs-NewDrs)
        ),
        PotentialReadings).

```

What if our parser fails to produce a result? The next clause implements Curt's behaviour for cases where it is not able to analyse the user's input.

```

curtOutput(_, Readings):-

```

```

    curtSays('What?'),
    curtInput(Readings).

```

In addition, we have some meta-commands to steer Curt. The input 'bye' ends the conversation, 'new' clears the DRSs and starts a new chat, and 'drs' displays the current readings in DRS format.

```

    curtOutput([bye],_):- !,
        curtSays('Bye bye!').

    curtOutput([new],_):- !,
        curt.

    curtOutput([drs],Readings):- !,
        printReadings(Readings),
        curtInput(Readings).

```

Finally, there is a driver predicate that starts a chat-session with Curt:

```

    curt:-
        curtInput([drs([],[])]).

```

Here is an example session. (The input of the user is preceded by a >, the rest is output of the system.)

```

?- curt.

> Marsellus loves Butch and Mia.

Ok.

> Mia is a boxer.

Ok.

> Butch is a boxer.

Ok.

> Marsellus loves one boxer.

No! I do not believe that!

>

```

Is Otter a useful prover for this application? It's pretty good—indeed it's light years ahead of the simple theorem prover we first experimented with. In particular, when something is a theorem it's very good at finding a proof. However when something is not a theorem, it may run into problems. The reader is asked to do the following exercise to experience this.

Exercise 5.3.1 Consider the sequence 'Marsellus loves a boxer' followed by 'Marsellus loves one boxer'. Is this a consistent and informative discourse? Try the program on the input sequence and find out which formula Otter tries to prove.


One line of experiment has been to simply substitute other theorem provers for Otter. This alone has led to interesting results, but there is another line of development that deserves much more thorough exploration, namely the use of model builders to filter out non-theorems. It is true that, when all is said and done, first-order inference is an undecidable problem—nonetheless, our experiments so far indicate that with a decent theorem prover to pin down the theorems, and a model builder to weed out the non-theorems, a striking amount can be done very speedily.

5.4 Adding Local Constraints

Van der Sandt's algorithm for presupposition projection actually uses the two maxims given above together with the the following extensions: superordinated DRSs should neither imply a subordinated DRS, nor a negated subordinated DRS. This constraint rules out the reading where the presupposition that Mia has a husband is globally accommodated:

If Mia is married then her husband is out of town.

This is because the fact that Mia has a husband follows from our background knowledge that she is a woman and is married. We will extend our discourse system to deal with this. To apply the local constraints to a certain DRS B , we have to calculate ordered pairs of DRSs, where each pair contains super- and subordinated DRSs of B . But what exactly these pairs of super-sub DRSs, and how are they computed? The best way to show the idea is giving an example. Assume we have the following DRS with sub-DRSs A , B , and C , structured as follows:



Then the pairs of super-sub DRSs are:

1. $\left(\frac{\quad}{B \Rightarrow C} \right), A$
2. $\left(\frac{\quad}{\neg A} \right), B$
3. $\left(\frac{\quad}{\neg A} \right) \oplus B, C$

This is coded in Prolog using `superSubDrs/3`. The first argument is the DRS that is traversed, the second argument the superordinated, and the last argument the subordinated DRS. Note that the superordinated DRS is represented as a difference list—this is done while conditions that contain the sub-DRS are removed from the superordinated DRS.

```

superSubDrs(drs(D, [Sub > _|C]), Drs-Super, Sub):-
    mergeDrs(merge(Drs, drs(D, C)), Super).
superSubDrs(drs(D, [B > Sub|C]), Drs-Super, Sub):-
    mergeDrs(merge(merge(drs(D, C), B), Drs), Super).
superSubDrs(drs(D, [B1 > B2|C]), Drs-Super, Sub):-
    superSubDrs(B2, merge(Drs, merge(merge(drs(D, C), B1), B2))-Super, Sub).
superSubDrs(drs(D, [Sub v _|C]), Drs-Super, Sub):-
    mergeDrs(merge(Drs, drs(D, C)), Super).
superSubDrs(drs(D, [_ v Sub|C]), Drs-Super, Sub):-
    mergeDrs(merge(Drs, drs(D, C)), Super).
superSubDrs(drs(D, [B v _|C]), Drs-Super, Sub):-
    superSubDrs(B, merge(Drs, merge(drs(D, C), B))-Super, Sub).
superSubDrs(drs(D, [_ v B|C]), Drs-Super, Sub):-
    superSubDrs(B, merge(Drs, merge(drs(D, C), B))-Super, Sub).
superSubDrs(drs(D, [~ Sub|C]), Drs-Super, Sub):-
    mergeDrs(merge(Drs, drs(D, C)), Super).
superSubDrs(drs(D, [~ B|C]), Drs-Super, Sub):-
    superSubDrs(B, merge(Drs, merge(drs(D, C), B))-Super, Sub).
superSubDrs(drs(D, [Cond|C]), Drs-Super, Sub):-
    superSubDrs(drs([], C), Drs-B, Sub),
    mergeDrs(merge(drs(D, [Cond]), B), Super).

```

We'll use this new predicate `superSubDrs` to traverse through the DRS and compute pairs of super- and subordinated DRSs in our definition of the predicates for the local constraints. Here is the predicate for local informativity:

```

localInformative(Drs):-

```

```

findall((Super,Sub),superSubDrs(Drs,drs([],[])-Super,Sub),List),
allLocalInformative(List).

```

```

allLocalInformative([]).
allLocalInformative([(Super,Sub)|Others]):-
  backgroundKnowledge(Chi),
  drs2fol(drs([], [Super>Sub]),Phi),
  callTheoremProver(Chi,~Phi,Proof),
  (Proof=yes,!,fail;allLocalInformative(Others)).

```

Similarly, we can code local consistency:

```

localConsistent(Drs):-
  findall((Super,Sub),superSubDrs(Drs,drs([],[])-Super,Sub),List),
  allLocalConsistent(List).

allLocalConsistent([]).
allLocalConsistent([(Super,Sub)|Others]):-
  backgroundKnowledge(Chi),
  drs2fol(drs([], [Super>drs([],[~Sub])]),Phi),
  callTheoremProver(Chi,~Phi,Proof),
  (Proof=yes,!,fail;allLocalConsistent(Others)).

```

How do we integrate our appeal to the local constraints in Curt? We'll use these as a filter on the readings, and only apply them to readings that are consistent and informative. The local constraints are less important than their global versions. This is how we integrate it in the predicate `curtOutput/2`:

```

curtOutput(Input,Readings):-
  d(MergeDrs,Input,[],!),
  betaConvert(MergeDrs,ReducedDrs),
  getReadings(ReducedDrs,Readings,PotentialReadings),
  consistentReadings(PotentialReadings,ConsistentReadings),
  (
    ConsistentReadings=[],
    curtSays('No! I do not believe that!'),!,
    curtInput(Readings)
  );
  informativeReadings(ConsistentReadings,InformativeReadings),
  (
    InformativeReadings=[],
    curtSays('Yes, I knew that!'),!,
    curtInput(Readings)
  );

```

```

    curtSays('Ok. '),
    localInformativReadings(InformativeReadings,LocalInformative),
    localConsistentReadings(LocalInformative,LocalConsistent),
    (
      LocalConsistent=[],!,
      SelectedReadings=InformativeReadings
    );
    SelectedReadings=LocalConsistent
  ),
  curtInput(SelectedReadings)
).

```

Exercise 5.4.1 Change this predicate in the following way. Only accept readings that do not violate (global) consistency and (global) informativity, and from these, select those with the *least* number of violations of the local constraints. (This strategy for applying the local constraints is due to David Beaver).

Background knowledge is encoded in first-order formulas. Of course, we want to extend the knowledge derived from the lexicon (Chapter 6) with common-sense facts of the world, such as married persons (if they are female) have husbands, and so on. Here is the new encoding of background knowledge.

```

backgroundKnowledge(Formulas):-
  knowledge(Formulas1),
  generateOntology(Formulas2),
  append(Formulas1,Formulas2,Formulas).

knowledge([
  forall(X,forall(Y,have(X,Y)>of(Y,X))),
  forall(X,forall(Y,of(Y,X)>have(X,Y))),
  forall(X,female(X)&married(X)>exists(Y,husband(Y)&have(X,Y))),
  forall(X,forall(Y,husband(Y)&have(X,Y)>married(X)&female(X)))
]).

```

Exercise 5.4.2 Try Curt on examples ‘If Mia is married then her husband is out of town’ and ‘Either Vincent does not have a car, or he cleans his car’, and see how Curt handles the presupposition triggers ‘her husband’ and ‘his car’.

Software Summary of Chapter 5

`curt.pl` The main predicates that build the framework of our interactive discourse system. (page 206)

`acceptabilityConstraints.pl` The implementation of consistency, informativity, and the local constraints. (page 209)

`callTheoremProver.pl` The Prolog-interface to the theorem prover Otter. (page 213)

`fol2otter.pl` Translates a formula in Otter syntax to standard output. (page 215)

Notes

The automatic deduction system Otter (McCune and Padmanabhan 1996) that we put together with our Prolog programs (we used version 3.0.5) is freely available on the World Wide Web, via:

<http://www-c.mcs.anl.gov/home/mccune/ar/otter/>

The DORIS system is an extended version of Curt (Blackburn, Bos, Kohlhase, and de Nivelle 1999). DORIS uses the MathWeb society of inference engines (currently, it uses the theorem provers Spass and Bliksem) for its reasoning tasks (Franke and Kohlhase 1999). Using the internet to communicate the inference problems and their answers, the theorem provers work on the problems in parallel running of different machines.

Chapter 6

Putting It All Together

6.1 Combining Scope and Presupposition

Exercise 6.1.1 Add the scores.

6.2 Adding Focus

Exercise 6.2.1 Use a theorem prover to sort the readings generated by the presupposition projection program.

Software Summary of Chapter 6

`mainPresupDRTU.pl` DRTU combined with presupposition resolution. (page 217)

`semMacrosPresupDRTU.pl` Semantic macros for DRTU and presupposition resolution. (page 218)

Notes

We won't give the reader any references here—indeed we're going swap roles and ask for some help. We would like to obtain references for programs that work more or less like the ones sketched in this chapter. That is, we're interested in hearing about examples of programs that 'plug together' parsers, theorem provers, and so on. All help greatly appreciated!

Bibliography

- Asher, N. (1993). *Reference to Abstract Objects in Discourse*. Dordrecht: Kluwer.
- Beaver, D. I. (1997). Presupposition. In J. Van Benthem and A. Ter Meulen (Eds.), *Handbook of Logic and Language*, Chapter 17, pp. 939–1008. Elsevier, MIT.
- Blackburn, P., J. Bos, M. Kohlhase, and H. de Nivelte (1999). Inference and Computational Semantics. In H. Bunt and E. Thijsse (Eds.), *Third International Workshop on Computational Semantics (IWCS-3)*, pp. 5–19. Computational Linguistics, Tilburg University.
- Bos, J., E. Mastenbroek, S. McGlashan, S. Millies, and M. Pinkal (1994). A Compositional DRS-based Formalism for NLP Applications. In H. Bunt, R. Muskens, and G. Rentier (Eds.), *International Workshop on Computational Semantics*. University of Tilburg.
- Bratko, I. (1990). *Prolog; Programming for Artificial Intelligence* (Second Edition ed.). Addison-Wesley.
- Charniak, E. (1972). *Towards a Model of Children's Story Comprehension*. Ph. D. thesis, Massachusetts Institute of Technology.
- Charniak, E. and D. McDermott (1985). *Introduction to Artificial Intelligence*. Addison-Wesley.
- Charniak, E. and Y. Wilks (Eds.) (1976). *Computational Semantics. An Introduction to Artificial Intelligence and Natural Language Comprehension*, Volume 4 of *Fundamental Studies in Computer Science*. Amsterdam/New York/Oxford: North-Holland Publishing Company.
- Chomsky, N. (1988). *Lectures on Government and Binding. The Pisa Lectures*. Studies in Generative Grammar. Foris Publications.
- Clocksin, W. and C. Mellish (1987). *Programming in Prolog* (Third, Revised and Extended Edition ed.). Springer Verlag.
- Cooper, R., I. Lewin, and A. W. Black (1993). Prolog and Natural Language Semantics. Notes for AI3/4 Computational Semantics, University of Edinburgh.
- Deransart, P., A. Ed-Dbali, and L. Cervoni (1996). *Prolog: The Standard. Reference Manual*. Springer.

- Fauconnier, G. (1985). *Mental Spaces. Aspects of Meaning Construction in Natural Language*. Bradford Books. The MIT Press.
- Fitting, M. (1996). *First-Order Logic and Automated Theorem Proving* (second ed.). Springer-Verlag.
- Franke, A. and M. Kohlhase (1999). System description: Mathweb, an agent-based communication layer for distributed automated theorem proving. In *16th International Conference on Automated Deduction CADE-16*.
- Frey, W. (1985). Syntax and Semantics of some Noun Phrases. In J. Laubsch (Ed.), *8th German Workshop on Artificial Intelligence (GWAI-84)*, Informatik Fachberichte 103, pp. 49–63.
- Gamut, L. (1991). *Logic, Language, and Meaning. Volume II. Intensional Logic and Logical Grammar*. Chicago and London: The University of Chicago Press.
- Groenendijk, J. and M. Stokhof (1990). Dynamic Montague Grammar. In L. Kálmán and L. Pólos (Eds.), *Papers from the Second Symposium on Logic and Language*, pp. 3–48.
- Grosz, B. J., A. K. Joshi, and S. Weinstein (1995). Centering: A framework for modeling the local coherence of discourse. *Computational Linguistics* 21(2), 203–225.
- Grosz, B. J. and L. Sidner, Candace (1986). Attention, Intentions, and the Structure of Discourse. *Computational Linguistics* 12, 175–204.
- Heim, I. (1982). *The Semantics of Definite and Indefinite Noun Phrases*. Ph. D. thesis, University of Massachusetts.
- Hobbs, J. R. (1986). Resolving Pronoun References. In B. J. Grosz, K. Sparck Jones, and B. L. Webber (Eds.), *Readings in Natural Language Processing*, pp. 339–352. Morgan Kaufmann.
- Johnson, M. and M. Kay (1990). Semantic Abstraction and Anaphora. In *COLING-90. Papers presented to the 13th International Conference on Computational Linguistics*, University of Helsinki, pp. 17–27.
- Johnson, M. and E. Klein (1986). Discourse, anaphora and parsing. In *11th International Conference on Computational Linguistics. Proceedings of Coling '86*, University of Bonn, pp. 669–675.
- Kamp, H. (1984). A Theory of Truth and Semantic Representation. In J. Groenendijk, T. M. Janssen, and M. Stokhof (Eds.), *Truth, Interpretation and Information; Selected Papers from the Third Amsterdam Colloquium*, pp. 1–41. Dordrecht - Holland/Cinnaminson - U.S.A.: Foris.
- Kamp, H. and U. Reyle (1993). *From Discourse to Logic*. Dordrecht: Kluwer.
- Kehler, A. (1997). Current theories of centering for pronoun interpretation. *Computational Linguistics* 23(3), 467–475.

- Kohlhase, M., S. Kuschert, and M. Pinkal (1996). A Type-Theoretic Semantics for λ -DRT. In P. Dekker and M. Stokhof (Eds.), *Proceedings of the Tenth Amsterdam Colloquium*, ILLC/Department of Philosophy, University of Amsterdam, pp. 479–498.
- Lewis, D. (1979). Scorekeeping in a Language Game. In R. Bäuerle, U. Egli, and A. von Stechow (Eds.), *Semantics from a different Point of View*, Volume 6 of *Springer Series in Language and Communication*, pp. 172–187. Springer-Verlag.
- McCune, W. and R. Padmanabhan (1996). *Automated Deduction in Equational Logic and Cubic Curves*. Lecture Notes in Computer Science (AI subseries). Springer-Verlag.
- Muskens, R. (1996). Combining montague semantics and discourse representation. *Linguistics and Philosophy* 19, 143–186.
- Pereira, F. and S. Shieber (1987). *Prolog and Natural Language Analysis*. CSLI Lecture Notes 10. Stanford: Chicago University Press.
- Reinhart, T. (1983). *Anaphora and Semantic Interpretation*. Chicago: The University of Chicago Press.
- Sidner, C. L. (1986). Focusing in the Comprehension of Definite Anaphora. In B. J. Grosz, K. S. Jones, and B. L. Webber (Eds.), *Readings in Natural Language Processing*, pp. 363–394. Morgan Kaufman.
- Smullyan, R. (1995). *First-Order Logic*. Dover Publications. This is a reprinted and corrected version of the 1968 Springer-Verlag edition.
- Sterling, L. and E. Shapiro (1986). *The Art of Prolog; Advanced Programming Techniques*. The MIT Press.
- Van der Sandt, R. A. (1992). Presupposition Projection as Anaphora Resolution. *Journal of Semantics* 9, 333–377.
- Van Eijck, J. and H. Kamp (1997). Representing discourse in context. In J. van Benthem and A. ter Meulen (Eds.), *Handbook of Logic and Language*, pp. 179–237. Elsevier.
- Wada, H. and N. Asher (1986). BUILDERS: An Implementation of DR Theory and LFG. In *11th International Conference on Computational Linguistics. Proceedings of Coling '86*, University of Bonn, pp. 540–545.
- Winograd, T. (1972). *Understanding Natural Language*. Academic Press.
- Zeevat, H. (1989). A compositional approach to discourse representation theory. *Linguistics and Philosophy* 12, 95–131.

Appendix A

Propositional Languages

The quantifier-free fragment of any first-order language (as the terminology suggests) simply consists of all formulas of the language that contain no occurrences of the symbols \exists or \forall . For example, $\text{ROBBER}(\text{PUMPKIN})$, $\text{CUSTOMER}(\text{MIA})$, $\text{CUSTOMER}(y)$, and

$$\text{CUSTOMER}(y) \rightarrow \text{LOVE}(\text{VINCENT},y)$$

are all quantifier-free formulas. On the other hand,

$$\forall y[\text{CUSTOMER}(y) \rightarrow \text{LOVE}(\text{VINCENT},y)]$$

isn't, as it contains an occurrence of the quantifier \forall .

The key thing to note about quantifier-free formulas is the following. Suppose we are given a model \mathbf{M} (of appropriate vocabulary) and an assignment g in \mathbf{M} . Now, normally we need to work with two semantic notions: satisfaction for arbitrary formulas and truth for sentences. However, when working with quantifier-free formulas, there are no bound variables to complicate matters, so this distinction is unnecessary. In fact, when working with a quantifier-free fragment, we may as well view each variable x as a constant interpreted by $g(x)$. If we do this, then every quantifier-free formula is either true or false in \mathbf{M} with respect to g . Moreover, it is obvious how to calculate the semantic value of complex sentences: conjunctions will be true if and only if both conjuncts are true, disjunctions will be true if and only if at least one disjunct is true, a negated formula will be true if and only if the formula itself is not true, and so on. (In short, we basically need to make the truth table calculations, which the reader is probably familiar with.)

This means we can simplify our notation somewhat. Because we don't have quantifiers, the internal structure of atomic formulas is irrelevant, for we're never going to bind any free variables they may contain. All that is important is whether the atomic symbols are true

or false, and how they are joined together using Boolean operators. For example, while it may be mnemonically helpful to choose propositional symbols such as `CUSTOMER(x)` or `LOVE(VINCENT,MIA)`, we lose nothing if we replace them by simpler symbols such as p and q . Following this line of thought leads us to define *propositional languages*. To specify a propositional language, we first say which symbols we are going to start with. (A fairly standard choice is p, q, r, s, t , and so on, often decorated with superscripts and subscripts: for example, p'' , r''' , or q_2 .) The chosen symbols are called proposition symbols, or atomic sentences. Complex sentences are built up using the standard Boolean connectives (for example, \neg, \wedge, \vee and \rightarrow) in the obvious way, and the truth values of complex sentences are calculated using the familiar truth table rules.

In short, propositional languages are essentially a simple notation for representing the quantifier-free formulas of first-order languages. When we devise inference mechanisms for first order logic in Chapters 4 and 5, it turns out to be sensible to first investigate inference methods for the quantifier free fragment (we do this in Chapter 4), and only then turn to the problem for the full first order language (the task of Chapter 5). Thus in Chapter 4, we make use of the simpler propositional notation just defined.

Appendix B

Type Theory

But What Does It All Mean?

We have given an informal, computationally oriented, introduction to the lambda calculus and its applications in semantic construction. We adopted a rather procedural perspective, encouraging the reader to think of the lambda calculus as a programming language—indeed, as the sort of language that emerges when one tries to generalize straightforward logic programming approaches to semantic construction (such as that of experiment 2).

However our account hasn't discussed one interesting issue: what do lambda expressions actually mean? Hopefully the reader now has a pretty firm grasp of what one can *do* with lambda expressions—but is one forced to think of lambda expressions purely procedurally? As we are associating lambda expressions with expressions of natural language, it would be nice if we could give them some kind of model theoretic interpretation.

Actually, there's something even more basic we haven't done: we haven't been precise about what counts as a λ -expression! Moreover—as the industrious reader may already have observed—if one takes an 'anything goes' attitude, it is possible to form some pretty wild (and indeed, wildly pretty) expressions. For example, consider the following expression:

$$\lambda x.(x@x)@x$$

Is this a legitimate lambda expression? Suppose we functionally apply this expression to itself (after all, nothing we have said rules out self-application). That is:

$$(\lambda x.(x@x)@x)@(\lambda x.(x@x)@x)$$

If we now apply β -conversion we get:

$$((\lambda x.(x@x)@x)@(\lambda x.(x@x)@x))@(\lambda x.(x@x)@x)$$

But note this is just the functional application followed by an additional occurrence of $\lambda x.(x@x)@x$! Obviously we can go on applying β -conversion as often as we like, producing ever longer expressions as we do so. Now, this is very interesting—but is it the sort of thing we had in mind when we decided to use the lambda calculus?

In this appendix, we shall briefly discuss some of these issues. The main points we wish to make are that there are different versions of the lambda calculus, useful for different purposes, and that both major variants of the lambda calculus can be given model theoretic interpretations in terms of functions and arguments. The reader interested in learning more is encouraged to follow up the references given in the Notes of Chapter 2.

The lambda calculus comes in two main varieties: the *untyped* lambda calculus, and the *typed* lambda calculus. Both can be regarded as programming languages, but they are very different.

The untyped lambda calculus adopts an ‘anything goes’ attitude to functional application. For example, $\lambda x.(x@x)@x$ is a perfectly reasonable expression in untyped lambda calculus, and it is fine to apply it to itself as we did above. The untyped lambda calculus is relevant to the study of natural language (explaining various natural language phenomena seems to demand some form of self-application). Moreover, it is relevant to functional programming (the core of the programming language LISP is essentially the untyped lambda calculus). Finally, it does have a model-theoretic semantics, indeed a very beautiful one. As one might suspect, this semantics interprets abstractions as certain kinds of functions. The primary difficulty is to find suitable collections of functions in which the idea of self application can be captured. (In standard set theory, functions can’t apply to themselves, so constructing function spaces with the structure necessary to model self application is a non-trivial exercise.) There are a variety of solutions to this problem, some of which are very elegant indeed.

There are many kinds of typed lambda calculi. The one we shall discuss is called the *simply typed lambda calculus*.

Natural language semanticists generally make use of some version of the simply typed lambda calculus. The key feature of the simply typed lambda calculus is that it adopts a very restrictive approach to functional application. “If it doesn’t fit, don’t force it”, and typed systems have exacting notions about what fits. Let’s discuss the idea of simple typing in a little more detail.

To build the kinds of representations we have been making use of in simply typed lambda calculus, we would proceed as follows. First we would specify the set of types. There would be infinitely many of these, namely (1) the type e of individuals, (2) the type t of truth values, and (3) for any types τ_1 and τ_2 , the function type $\langle \tau_1, \tau_2 \rangle$. Second, we would specify our logical language. This would contain all the familiar first order symbols, but in addition it would contain an infinite collection of variables of each type (the ordinary first

order variables, which range over individuals, would now be thought of as the set of type e variables) together with the λ operator.

In the typed lambda calculus, every expression receives a *unique* type. The key clauses that ensures this are the definitions of abstraction and functional application. First, if \mathcal{E} is a lambda expression of type τ_2 , and v a variable of type τ_1 then $\lambda v.\mathcal{E}$ is a lambda expression of type $\langle\tau_1, \tau_2\rangle$. In short, it matters which type of variable we abstract over. Abstracting with respect to different types of variables results in abstractions with different types. Second, if \mathcal{E} is a lambda expression of type $\langle\tau_1, \tau_2\rangle$, and \mathcal{E}' is a lambda expression of type τ_1 then we are permitted to functionally apply \mathcal{E} to \mathcal{E}' . If we do this, the application has type τ_2 . In short, we are only allowed to perform functional application when the types \mathcal{E} and \mathcal{E}' fit together correctly, *and only then*. The intuition is that the types tell us what the domain and range of each expression is, and if these don't match, application is not permitted. Note, moreover, that the type of the application is determined by the types of \mathcal{E} and \mathcal{E}' . In effect, we have imposed a very strict type discipline on our formalism. The typed lambda calculus is a programming language, but when we use it we have to abide by very strict guidelines.

This has a number of consequences. For a start, self-application is impossible. (To use \mathcal{E} as a functor, it must have a function type, say $\langle\tau_1, \tau_2\rangle$. But then its arguments must have type τ_1 . So \mathcal{E} can't be one of its own arguments.) Moreover, it is very straightforward to give a semantics to such systems. Given any model M , the denotation D_e of type e expressions are the elements of the model, the permitted denotations D_t of the type t expressions are TRUE and FALSE, and the permitted denotations $D_{\langle\tau_1, \tau_2\rangle}$ of type $\langle\tau_1, \tau_2\rangle$ expressions are functions whose domain is D_{τ_1} and whose range is D_{τ_2} . In short, expressions of the simply typed lambda calculus are interpreted using a straightforward, inductively defined, function hierarchy.

Which particular functions are actually used? Consider $\lambda x.MAN(x)$, where x is an ordinary first order variable. Now, $MAN(x)$ is a formula, something that can be TRUE or FALSE, so this has type t . As was already mentioned, first-order variables are viewed as type e variables, hence it follows that the abstraction $\lambda x.MAN(x)$ has type $\langle e, t\rangle$. That is, it must be interpreted by a function from the set of individuals to the set of truth values. But which one? In fact, it would be interpreted by the function which, when given an individual from the domain of quantification as argument, returns TRUE if that individual is a man, and FALSE otherwise. To put it another way, it is interpreted using the function which exactly characterizes the subset of the model consisting of all men. But this subset is precisely the subset that the standard first-order semantics uses to interpret MAN . In short, the 'functional interpretation' of lambda expressions is set up so that, via the mechanism of such *characteristic functions*, it meshes perfectly with the ordinary first order interpretation.

By building over this base in a fairly straightforward way, the interpretation can be extended to cover all lambda expressions. For example, the expression

$$\lambda P.\exists x(WOMAN(x)\wedge P@x)$$

would be interpreted as a function which when given the type of function that P denotes as argument (and P denotes type $\langle e, t \rangle$ functions, like the characteristic function of MAN) returns a type t value (that is, either `TRUE` or `FALSE`). Admittedly, thinking in terms of functions that take other functions as arguments and return functions as values can get pretty tedious, but the basic idea is straightforward, and for applications in natural language semantics, only a very small part of the function hierarchy tends to be used.

In short, throughout this book we have been talking about the lambda calculus as a mechanism for marking ‘missing information’, and we have exploited the mechanisms of functional application and β -conversion as a way of moving such missing information to where we want it. But in fact there is nothing at all mysterious about this ‘missing information’ metaphor. It is possible to give precise mathematical models of missing information in terms of functions and arguments. An abstraction is interpreted as a function, and the ‘missing information’ is simply the argument we will later supply it with. Indeed, a variety of models are possible, depending on whether one wants to work with typed or untyped versions of the lambda calculus.

Exercise B.0.1 Does our implementation of β -conversion allow self-application?

Exercise B.0.2 What happens when you functionally apply the formula $\lambda x.x@x$ to itself?

Appendix C

Theorem Provers and Model Builders

Automated reasoning has seen an enormous increase of performance of (especially first-order) inference engines and model builders. This appendix is a guide to a number of useful reasoning engines for computational semantics.

Theorem Provers

- **BLIKSEM**
Resolution based theorem prover for first order logic with equality (Hans de Nivelle). Almost as fast as the speed of light.
URL: <http://turing.wins.uva.nl/~mdr/ACLG/Provers/Bliksem/bliksem.html>
- **SPASS**
First-order theorem prover (Christoph Weidenbach et al.)
Winner at the CASC-15 in the divisions FOF and SAT.
URL: <http://spass.mpi-sb.mpg.de/>.
- **FDPLL**:
Or, if you prefer, the “First-Order Davis-Putnam-Loveland-Logeman” theorem prover (Peter Baumgartner). Good at satisfiable problems.
URL: <http://www.uni-koblenz.de/~peter/FDPLL/>
- **OTTER**
The ‘classical’ theorem prover by W. McCune.
URL: <http://www-unix.mcs.anl.gov/AR/otter/>

Model Builders

- **MACE**
Short for “Model and Counter-Examples”. Mace is a Model builder for first-order logic with equality (McCune). Got the first price at CASC-16 in division SAT (working in tandem with OTTER).
URL: <http://www-unix.mcs.anl.gov/AR/mace/>
- **SATCHMO**
Satchmo is a model generator for first-order theories, implemented in Prolog.
URL: <http://www.pms.informatik.uni-muenchen.de/software/>
- **KIMBA**
Kimba is an Higher-Order Model Generator, a deduction system for abductive or circumscriptive reasoning within linguistic applications (Karsten Konrad).
URL: <http://www.ags.uni-sb.de/~konrad/kimba.html>

Systems

- **MATHWEB**
This system for automated theorem proving connects a wide-range of mathematical services by a common, mathematical software bus (Franke and Kohlhase 1999).
URL: <http://www.ags.uni-sb.de/~omega/www/mathweb.html>

Appendix D

Prolog in a Nutshell

What is Prolog?

Prolog is one of the most important programming languages in Computational Linguistics and Artificial Intelligence. Two key features distinguish Prolog from other programming languages: its *declarative* nature, and its extensive use of *unification*. Ideally, in Prolog we simply state the nature of the problem, and let the Prolog unification-driven inference engine search for a solution.

The Basics

There are actually only three basic constructs in Prolog: *facts*, *rules*, and *queries*. A set of facts and rules — that is, a *knowledge base* — is all a Prolog program consists of. Facts are used to state things that are unconditionally true of the domain of interest. For example, we can state that Mia and Jody are women by putting the facts

```
woman(mia).  
woman(jody).
```

in our knowledge base. Rules relate facts by logical implications. We can add to our knowledge base the conditional information that Mia plays air-guitar *if* she listens to music as follows:

```
playsAirGuitar(mia):-  
    listensToMusic(mia).
```

The `:-` should be read as “if”, or “is implied by”. The part on the left hand side of the `:-` is called the *head* of the rule, the part on the right hand side the *body*. (Incidentally, note that we can view a fact as a rule with an empty body. That is, we can think of facts as ‘conditionals that don’t have any antecedent conditions’.) The facts and rules contained in a knowledge base are called *clauses*. The collection of all clauses in a knowledge base that have the same head is called a *predicate*.

Posing queries makes the Prolog inference engine try to deduce a positive answer from the information contained in the knowledge base. There are two basic circumstances under which Prolog can return a positive answer. The first, and simplest, is when the question posed is one of the facts listed in the knowledge base. The second is when Prolog can deduce the positive answer by using the deduction rule called modus ponens. That is, if `head :- body` and `body` are both in the knowledge base, then Prolog can deduce that `head` is true.

Let’s consider an example. We can ask Prolog whether Mia is a woman by posing the query

```
?- woman(mia).
```

and Prolog will answer “yes”, since this is a fact in the knowledge base. However, if we ask whether Mia plays air-guitar by posing the query

```
?- playsAirGuitar(mia).
```

its answer is “no”. First, this fact is not recorded in the knowledge base. Second, it cannot infer that Mia plays air-guitar as there is nothing in the knowledge base stating that Mia is listening to music (thus Prolog assumes that this is false) and hence we cannot make use of the only rule we have in our knowledge base. On the other hand, if the knowledge base had contained the additional fact

```
listensToMusic(mia).
```

then Prolog would have responded “yes”, since it could then have deduced by modus ponens that Mia plays air-guitar.

Jumping ahead slightly, there are two things that make Prolog so powerful. The first is that it is capable of ‘chaining together’ uses of modus ponens. The second is that Prolog has a powerful mechanism (described later) called *unification*, which lets it handle variables. This combination of chained modus ponens and sophisticated variable handling enables it to draw far more interesting inferences than our rather trivial examples might suggest.

The comma `,` expresses logical conjunction in Prolog. We can change the rule above by adding another condition to its body:

```
playsAirGuitar(mia):-  
    listensToMusic(mia),  
    happy(mia).
```

Now the rule reads: Mia plays air-guitar if she listens to music, *and* if she is happy. We can also express disjunction in Prolog. Let's change the rule to:

```
playsAirGuitar(mia):-  
    listensToMusic(mia);  
    happy(mia).
```

The ; should be read as *or*, so here we have a rule stating that Mia plays air-guitar if she listens to music, *or* if she is happy.

However there is another (much more commonly used) way of expressing disjunctive conditions: we simply list a number of clauses that have the same head. For example, the following two rules mean exactly the same as the previous rule:

```
playsAirGuitar(mia):-  
    listensToMusic(mia).  
playsAirGuitar(mia):-  
    happy(mia).
```

In fact, disjunctions are almost always expressed using such multiple rules; extensive use of semicolon makes Prolog code pretty hard to read.

Syntax

What exactly are the syntactic entities such as `woman(jody)` and `playsAirGuitar(mia)` that we use to build facts, rules, and queries? They are called *terms*, and there are three kinds of terms in Prolog: atoms, variables, and complex terms. An atom is a sequence of characters starting with a lowercase character. A variable is also a sequence of characters, but it must start with an uppercase character or an underscore. So `mia` and `airGuitar` are atoms, while `X`, `Mayonnaise`, and `_mayonnaise` are variables.

Complex terms are build out of a *functor* and a sequence of *arguments*. The arguments are put in ordinary brackets, separated by commas, and placed after the functor. The functor *must* be an atom. That is, variables *cannot* be used as functors. On the other hand, arguments can be any kind of term. For example

```
hide(X,father(butch))
```

is a complex term. Its functor is `hide`, and it has two arguments: the variable `X`, and the complex term `father(butch)`.

The number of arguments that a complex term has is called its *arity*. For instance, `woman(mia)` is a complex term with arity 1, while `loves(vincent,mia)` is a complex term with arity 2.

Arity is important to Prolog. Prolog would be quite happy for us to use two predicates with the same functor but with a different number of arguments (for example, `love(vincent,mia)`, and `love(vincent,marcellus,mia)`) in the same program. However, if we insisted on doing this, Prolog would treat the two place `love` and the three place `love` as completely different predicates.

When we need to talk about predicates and how we intend to use them (for example, in documentation) it is usual to use a suffix `/` followed by a number to indicate the predicate's arity. For example, if we were talking about the `playsAirGuitar` predicate we would write it as `playsAirGuitar/1` to indicate that it takes one argument. We make use of this convention in the book.

Unification

Variables allow us to make general statements in Prolog. For example, to declare that every woman likes a foot massage, we add the following rules to the database:

```
likeFootMessage(X):-  
    woman(X).
```

It can be read as: `X` likes a foot massage, if `X` is a woman. The nice thing about variables is that they have no fixed values. They can be instantiated with a value by the process of unification. If we pose the query

```
?- likeFootMessage(mia).
```

the variable `X` is *unified* with the atom `mia`, and Prolog tries to prove that `woman(mia)` can be inferred from the knowledge base. Two terms are unifiable if they are the same atoms, or if one of them is a variable, or if they are both complex terms with the same functor name and arity, and all corresponding arguments unify. Unification makes terms identical. It is the heart of the engine that drives Prolog.

Prolog has a built-in operator for unification: the `=`. By querying the goal

```
?- X = butch.
```

the variable `X` gets unified with the atom `butch` and the goal succeeds.

Recursion

Predicates can be defined recursively. Roughly speaking, a predicate is recursively defined if one or more rules in its definition refers to itself. Here's an example:

```
moreExpensive(X,Y):-
    costsaLittleMore(X,Y).

moreExpensive(X,Y):-
    costsaLittleMore(X,Z),
    moreExpensive(Z,Y).
```

The definition of the `moreExpensive/2` predicate is fairly typical of the way recursive predicates are defined in Prolog. Clearly, `moreExpensive/2` is (at least partially) defined in terms of itself, as the more `moreExpensive` functor occurs on both the left and right hand sides of the second rule. Note, however, that there is an 'escape' from this circularity. This is provided by the `costsaLittleMore` predicate, which occurs in both the first and second rules. (Significantly, the right hand side of the first rule makes no mention of `moreExpensive`.)

Let's see how Prolog makes use of such definitions. Suppose we had the following facts in our knowledge base:

```
costsaLittleMore(royaleWithCheese,bigKahunaBurger).

costsaLittleMore(fiveDollarShake,royaleWithCheese).
```

If we pose the query

```
?- moreExpensive(fiveDollarShake,bigKahunaBurger).
```

then Prolog goes to work as follows. First, it tries to make use of the first rule listed concerning `moreExpensive`. This tells it that X is more expensive than Y if X costs a little more than Y , but as the knowledge base doesn't contain the information that a five dollar shake costs a little more than a big kahuna burger, this is no help. So, Prolog tries to make use of the second rule. By unifying X with `fiveDollarShake` and Y with `bigKahunaBurger` it obtains the following goal:

```
?- costsaLittleMore(fiveDollarShake,Z),
    moreExpensive(Z,bigKahunaBurger).
```

Prolog deduces `moreExpensive(fiveDollarShake,bigKahunaBurger)` if it can find a value for Z such that, firstly,

```
?- costsaLittleMore(fiveDollarShake,Z).
```

is deducible, and secondly,

```
?- moreExpensive(Z,bigKahunaBurger).
```

is deducible too. But there is such a value for Z: `royaleWithCheese`. It is immediate that

```
?- costsaLittleMore(fiveDollarShake,royaleWithCheese).
```

will succeed, for this fact is listed in the knowledge base. Deducing

```
?- moreExpensive(royaleWithCheese,bigKahunaBurger).
```

is almost as simple, for the first clause of `moreExpensive/2` reduces this goal to deducing

```
?- costsaLittleMore(royaleWithCheese,bigKahunBurger).
```

and this is a fact listed in the knowledge base.

There is one very important thing you should bear in mind when working with recursion: a recursive predicate should always have at least two clauses: a non-recursive one (the clause that stops the recursion at some point, otherwise Prolog would never be able to finish a proof!), and one that contains the recursion. In our example, the first clause of `more_expensive/2` is the non-recursive clause, and the second clause contains the recursion. (Note that the order of these two clauses in the knowledge base is not important!)

Lists

Lists are recursive data structures in Prolog. The recursive definition of a list runs as follows. First, the empty list is a list. Second, a complex term is a list if it consists of two items, the first of which is a term (often referred to as *first*), and the second of which is a list (often referred to as *rest* or *the rest list*).

Square brackets indicate lists. The empty list is written as `[]`. The list operator `|` separates the first item of a list from the rest list. For example, here is a list with three items:

```
[butch|[pumpkin|[marsellus|[]]]]
```


When working with lists, Prolog always makes use of such recursive first/rest representations, and for some purposes it is important to know this. Mercifully, however, Prolog also offers a more readable form of list representation. The same list can be declared as:

```
[butch,pumpkin,marsellus]
```

Prolog will quite happily accept lists in this more palatable notation as input, and moreover, it does its best to use this notation for its output.

Note that the items in a list need not only be atoms: they can be any Prolog term, including lists. For example

```
[vincent,[honey_bunny,pumpkin],[marsellus,mia]]
```

is a perfectly good list.

Since lists are recursive data structures, most of the predicates that work on lists, are defined using recursive predicates. The simplest example is the `member/2` predicate, which is given in the program library.

Operators

Many newcomers to Prolog find the word *operator* rather misleading, for Prolog's operators don't actually operate on anything, or indeed do anything at all. They are simply a notational device that Prolog offers to represent complex terms in a more readable fashion.

For example, suppose that we want to use `not` as the functor expressing sentence negation, and `and` as the functor expressing sentence conjunction. Then the term representing `Butch boxes and Vincent doesn't dance` would be:

```
and(butch_boxes,not(vincent_dances)).
```

It would be nicer if we could use the more familiar notation in which the conjunction symbol stands between the two sentences it conjoins. That is, we would prefer the following representation:

```
butch_boxes and not vincent_dances
```

The following *operator definitions* let us do precisely this.

```
:- op(30,yfx,and).
:- op(20,fy,not).
```

The `and` is declared as an infix operator by `yfx` with precedence value of 30. The `y` represents an argument (in this case the left argument of `&`) whose precedence is lower or equal of the operator. The `x` represents an argument (the right argument of `and`) whose precedence value must be strictly lower than that of the operator (30, in this case). The `not` is defined as a prefix operator with precedence 20, and an argument that should have a precedence value lower or equal to 20.

More generally, Prolog allows us to define our own infix, prefix and postfix operators. Operator names must be atoms. When we declare a new operator, we also have to state its precedence and those of its arguments (in order for Prolog to disambiguate expressions with more than one operator).

A very important final remark about operators: Prolog is featured with a set of special predefined operators, that *have* a meaning. Among these is the `=` for unification as we saw before and `\+` for negation as we will see later. Other predefined operators we use in the programs of this book are the infix operators `==` and `\==`. The goal `X == Y` succeeds if `X` and `Y` are identical terms (That is, they should have the same structure, even variables should have the same name — unification plays no role here!). On the other hand, `\==` checks whether its arguments are *not* identical.

Arithmetic

Prolog contains some built-in operators for handling integer arithmetic. These include `*`, `/`, `+`, `-` (for multiplication, division, addition, and subtraction, respectively) and `>`, `<` for comparing numbers.

These symbols, however, are just ordinary Prolog operators. That is, they are just a user friendly notation for writing arithmetic expressions: they *don't* carry out the actual arithmetic. For example, posing the query

```
?- X = 1 + 1.
```

unifies the variable `X` with the complex term `1 + 1`, not with `2`, which, for people unused to Prolog's little ways, tends to come as a bit of a surprise.

If we want to carry out the actual arithmetic involved, we have to explicitly force evaluation by making use of the very special inbuilt 'operator' `is/2`. This calls an inbuilt mechanism which carries out the arithmetic evaluation of its second argument, and then unifies the result with its first argument.

```
?- X is 1 + 1.  
X = 2  
yes
```

Negation

There is no explicit negation in Prolog. Something is regarded as false if it cannot be proved. This is the so-called closed world assumption of Prolog.

Prolog does have an inbuilt mechanism for “negation as failure”. That is, we can ask it whether something cannot be proved by using the built-in prefix operator `\+`. The goal `?- \+ man(jules)` succeeds if and only if the goal `?- man(jules)` fails.

Variables in negated goals are not instantiated. Therefore, the following two goals are not fully equivalent.

```
?- likeFootMessage(Who)  
Who = mia  
yes  
  
?- \+ \+ likeFootMessage(Who)  
yes
```

Backtracking and the Cut

When Prolog tries to prove a goal, the structure underlying it attempts can be regarded as a tree: a tree with branches that lead (or do not lead) to possible proofs. The Prolog inference engine essentially searches for a branch that makes the main goal true. Of course it may, and usually does, find itself in a branch that does not lead to a proof (that is, a branch that *fails*). Then Prolog automatically *backtracks* to the previous node in the tree and tries another (the next) branch. (If there is no previous node, this means that the goal is not provable, and Prolog tells us “no”.)

Backtracking can be forced by the user by entering a semicolon after Prolog gives a solution. This allows us to try generating more than one solution to a query. For example (reverting to our original knowledge base) we can demand that Prolog finds all the women as follows:

```
?- woman(Who).  
Who = mia;  
Who = jody;  
no
```

There are two ways to influence Prolog's search strategy: using the *cut* to suppress backtracking, or changing the order of the clauses in the database.

The *cut* (written `!`, it is a built-in Prolog predicate with arity 0) removes certain branches from the proof tree. If a *cut* is put in a clause, and Prolog encounters it during a proof, it removes all the clauses of the same predicate that are listed further down in the knowledge base, and moreover, removes all alternative solutions to conjuncts to the left of the *cut* in its clause.

The order of clauses play a role in Prolog's search strategy, since Prolog works on the database in a sequential way, and subgoals are proved from left to right in the search tree. (Sometimes the order even drastically determines the outcome of a proof.)

Built-in Predicates

There is a set of built-in predicates available in Prolog. We briefly discuss the ones that we use in our programs. It should be noted though, that it depends on the Prolog version that you use, which predicates are built-in.

First we have the ones for controlling output. The predicate `write/1` displays a term to the current output device, and `nl/0` creates a new line.

Then there is whole family of predicates that are used for term manipulation. With these you can break complex terms into pieces or build new ones. The predicate `functor(T,F,A)` succeeds if `F` is the functor of term `T` with arity `A`. And `arg(N,T,A)` is true if `A` is the `N`th argument of `T`. In some case we prefer to make use of the so-called "univ" predicate, weirdly written as `=..`, that transforms a complex term into a list consisting of the functor followed by its arguments. For instance:

```
?- love(pumpkin,honey_bunny) =.. List.
List = [love,pumpkin,honey_bunny]
yes
```

For checking the types of terms: `nonvar(X)` is true if `X` is instantiated, `var(X)` is true if `X` is not instantiated. The predicate `simple/1` succeeds if its argument is either an atom or a variable, and `compound/1` if its argument is a complex term.

With `assert/1` and `retract/1` it is possible to change the knowledge base while executing a goal. The former asserts a clause to the database, the latter removes it. Many versions of Prolog require a `dynamic` declaration of those predicates that are modified by other predicates, as we do with the counter for the skolem index in the library file.

Finally, there is some predefined stuff that gives controll over backtracking. The very rude `fail/0` predicate causes Prolog to backtrack (this is used to generate more solutions). With

`bagof/3` it is possible to generate all solutions of a specific goal, and put the instantiations of a certain variable (or complex term with variables) in a list. E.g., the goal

```
?- bagof(Who, woman(Who), Answer).
Answer = [mia, jody].
yes
```

tries to satisfies the goal in the second argument of `bagof`, and for each solution it puts the value of its first argument (`Who`) in its third argument, the list `Answer`. The predicate `setof/3` functions similarly, but it removes duplicate answers.

Difference Lists

A difference list `List1-List2` is a way of using two lists (namely, `List1` and `List2`) to represent one single list (namely, the *difference* of the two lists).

For example, the three element list `[pumpkin, butch, jimmy]` can be represented as the following difference list:

```
[pumpkin, butch, jimmy, lance]-[lance]
```

or as the following one:

```
[pumpkin, butch, jimmy]-[]
```

Indeed, there are infinitely many difference list representations of `[pumpkin, butch, jimmy]`. In all of them, the first list in the difference list representation consists of the list we are interested in (namely `[pumpkin, butch, jimmy]`) followed by some suffix, while the second list is that suffix. So a difference list is simply a pair of list structures, the second of which is a suffix from the other. We follow the usual convention of using the built-in operator `-` to group the two list together. (However, note that we don't have to do this: any other operator would do.)

Let's go a step further. Suppose we take the suffix to be a variable. Then any other difference-list encoding of `[pumpkin, butch, jimmy]` is an instance of the following *most general difference list*.

```
[pumpkin, butch, jimmy|X]-X
```

In practice, we will use the term *difference list* to mean the most general difference list. The empty list is represented by the difference list $X-X$.

Why on earth should anyone want to represent lists as difference lists? There is a simple answer: efficiency. In particular, when we use the difference list representation, Prolog can perform concatenation of lists much faster. For example, the usual `append/3` for normal lists is a recursively defined predicate that can be inefficient for large lists. (This is because Prolog must recursively work its way down the entire collection of first/rest pairs.) On the other hand, when we make use of difference lists `append/3` can be defined as follows.

```
append(X-Y, Y-Z, X-Z).
```

Consider how this works. Suppose we want to append `[pumpkin, butch]` to `[jody, lance]`. Then we make the following query:

```
?- append([pumpkin, butch|E]-E, [jody, lance|F]-F, A-B).
```

This causes `A` to unify with `[pumpkin, butch|E]`, `E` with `[jody, lance|F]`, and `B` with `F`. As a result, `A-B` unifies with `[pumpkin, butch, jody, lance|B]-B`, which is a difference list of the four items `pumpkin`, `butch`, `jody`, and `lance`.

The reader who works through this example will see that what makes difference list representations so efficient is that the suffix variable gives us direct access to the end of the list. In the conventional first/rest representation we have to work our way recursively down towards the end of the list. Difference list representations avoid this overhead. We have to pay a price for this gain in efficiency (difference list representation is less transparent) but in many circumstances this is a price worth paying.

Definite Clause Grammars

Definite clause grammars (DCGs) are the in built Prolog mechanism for defining grammars. Actually, they are really a syntactically sugared way of working with certain difference lists. With DCGs you can kill two birds with one stone: if you define the grammar rules you'll get the parser for free!

Here is a DCG for a very small fragment of english grammar. We have defined syntactic categories `s`, `np`, `vp`, `det`, `noun`, and `tv`, standing for sentence, noun phrase, verb phrase, determiner, common noun, and transitive verb. These are also called the *non-terminal* symbols. Every rule in the DCG has a non-terminal symbol on its left hand side.

```
s --> np, vp.
```

```
np --> det, noun.  
np --> [mia].  
det --> [a].  
noun --> [man].  
noun --> [five,dollar,shake].  
vp --> tv, np.  
vp --> [drinks].  
tv --> [loves].
```

On the right hand side of the rules, you'll either find one or more non-terminal symbols, or a terminal symbol. The terminal symbols are the lexical entries, the actual words of the language of our interest.

Now, what do these rules mean? They are very intuitive indeed. The rule `s --> np, vp` says that a syntactic category called `s`, consists of an `np`, followed by a `vp`. Similarly, according to this DCG, the category `noun` is either the string `man`, or the sequence of strings `five dollar shake` (represented in lists). And so on.

This DCG covers sentences like `Mia loves a five dollar shake`, `A man drinks`, but not: `A woman loves` or `Mia drinks a one dollar shake`.

To parse sentences, we can pose a query like:

```
?- s([mia,loves,a,five,dollar,shake], []).
```

This goal is satisfied if the sequence of words in the first argument belongs to `s`. There is only one rule for `s` in our DCG. It says that an `s` can be replaced by an `np` followed by a `vp`. That is, we have to take some items of the input list that form a noun phrase, in such a way that the rest of the items on the list form a verb phrase. Since there is the rule

```
np --> [mia].
```

in the knowledge base we now have to proof whether

```
?- vp([loves,a,five,dollar,shake], []).
```

which is provable indeed (we leave this is an exercise to the reader). And this is basically how parsing with DCGs takes place on the surface level. However, Prolog doesn't use our DCG rules directly for the purpose of parsing. The DCG is translated internally into the following clauses:

```

s(A,B) :- np(A,C), vp(C,B).
np(A,B) :- det(A,C), noun(C,B).
np([mia|B],B).
det([a|B],B).
noun([man|B],B).
noun([five,dollar,shake|B],B).
vp(A,B) :- tv(A,C), np(C,B).
vp([drinks|B],B).
tv([loves|B],B).

```

The alert reader will notice that these clauses look surprisingly familiar. In fact, they are the direct encodings of difference lists (in the facts, for example `noun([man|B],B)`) and the appending of two difference lists (in the rules, for instance

```

s(A,B) :- np(A,C), vp(C,B).

```

can be read as: `s` is a difference list `A-B` if we can prove that is the result of appending difference list `A-C` to `C-B`).

Since, as we just noticed, DCG rules are normal Prolog clauses, it is perfectly allowed to add arguments to the rules. Some useful stuff we can add to our grammar is information on agreement. Suppose we want to include noun phrases like *all boxers* in our grammar by adding entries for *all* and *boxers*:

```

det --> [all].
noun --> [boxers].

```

Be careful though! Adding these clauses make it possible to parse *Mia loves all boxers*, but also non-grammatical *All man drinks* or *A boxers loves all woman*. Clearly, our grammar lacks information about agreement. However, this information can be added very easily to the rules. Let's do it for the determiners and nouns first (and why not add some extra entries at the same time):

```

det(plural) --> [all].
det(singular) --> [a].
noun(singular) --> [boxer].
noun(plural) --> [boxers].
noun(singular) --> [man].
noun(plural) --> [men].
noun(singular) --> [five,dollar,shake].
noun(plural) --> [five,dollar,shakes].

```


Since `det` and `noun` have an extra argument now, all the rules that have these symbols at their right hand side (this is the `np --> det, noun.` rule in our current grammar) should get an extra argument as well. As can be seen from the entries above, the value of this argument is either `singular` or `plural`. Since we want the determiner to have the same agreement as the noun when they are combined to a noun phrase, we could write the following code:

```
np --> det(singular), noun(singular).
np --> det(plural), noun(plural).
```

But there is much more elegant way of encoding this. We can bring in a variable that, during parsing, gets instantiated with the agreement value of the determiner and noun, and collapse the above rules into one:

```
np --> det(Agr), noun(Agr).
```

Nevertheless, from the examples given earlier, we also want to add agreement features to noun phrases and verb phrases. We won't do explain this step by step — the principle should be clear now — but list the entire rewritten grammar including agreement, and extended with some new entries, below.

```
s --> np(Agr), vp(Agr).
np(Agr) --> det(Agr), noun(Agr).
np(singular) --> [mia].
det(plural) --> [all].
det(singular) --> [a].
noun(singular) --> [boxer].
noun(plural) --> [boxers].
noun(singular) --> [man].
noun(plural) --> [men].
noun(singular) --> [five,dollar,shake].
noun(plural) --> [five,dollar,shakes].
vp(Agr) --> tv(Agr), np(_).
vp(singular) --> [drinks].
vp(plural) --> [drink].
tv(singular) --> [loves].
tv(plural) --> [love].
```

Now look at our grammar. From an aesthetic point of view, there certainly is some space for improvement! Everything is mixed up: the rules and the lexical entries together form

one chaotic DCG. Consider what these grammar would look like if we extend its coverage to a more serious fragment of english!

One way to bring in some organisation in the grammar, is to make a physical distinction between the lexicon and the grammar rules. The lexicon might be designed as follows:

```
lexicon(np,singular,mia).
lexicon(det,plural,all).
lexicon(det,singular,a).
lexicon(noun,singular,boxer).
lexicon(noun,plural,boxers).
lexicon(noun,singular,man).
lexicon(noun,plural,men).
lexicon(vp,singular,drinks).
lexicon(vp,plural,drink).
lexicon(tv,singular,loves).
lexicon(tv,plural,love).
```

That is, normal Prolog facts of the form `lexicon/3`, with the first argument stating the syntactic category, the second argument the agreement value, and the third the word itself. The only thing left to do is to make a connection between this lexicon and the grammar rules. DCGs have a neat way to do this, normal Prolog goals can be included in the rules included in curly brackets. This is the result:

```
s --> np(Agr), vp(Agr).
np(Agr) --> [X], {lexicon(np,Agr,X)}.
np(Agr) --> det(Agr), noun(Agr).
det(Agr) --> [X], {lexicon(det,Agr,X)}.
vp(Agr) --> tv(Agr), np(_).
vp(Agr) --> [X], {lexicon(vp,Agr,X)}.
tv(Agr) --> [X], {lexicon(tv,Agr,X)}.
```

Notes

In this appendix we summarized the basic concepts of Prolog. We hope it will be a handy reference, however it is *not* intended as a substitute for good introduction to Prolog. The reader who wants to learn about computational semantics, but who knows no Prolog, is strongly advised to put this book aside for a while and study one of the many excellent Prolog texts currently available. We particularly recommend the following ones. For a succinct, no-frills overview, try Clocksin and Mellish (Clocksin and Mellish 1987). For a leisurely, in-depth introduction to programming in Prolog, try Bratko (Bratko 1990). For

a more theoretically oriented introduction, try Sterling and Shapiro (Sterling and Shapiro 1986). Finally, for an introduction specially geared towards computational linguistics, try Pereira and Shieber (Pereira and Shieber 1987).

Appendix E

Listing of Programs

This appendix includes the full program listings that are developed in this book. Most predicates are decorated with a short documentation, using the following notational conventions for its required argument instantiations:

`:Arg` Argument `Arg` should be instantiated to a term denoting a goal

`+Arg` Argument `Arg` should be instantiated

`-Arg` Argument `Arg` should not be instantiated

`?Arg` Argument `Arg` may or may not be instantiated

Since Prolog's birth in the beginning of the seventies, a number of Prolog dialects emerged, and not all agree on a syntax or the in-built predicates. The programs in this book follow the conventions of "Standard Prolog", the ISO international standard on Prolog (Deransart, Ed-Dbali, and Cervoni 1996), as close as possible. The following predicates are assumed to be built-in in your version of Prolog (such as for example Quintus or Sicstus Prolog):

All solutions

`bagof/3`

`findall/3`

Arithmetic comparison

`>` (arithmetic greater than)

< (arithmetic less than)

Arithmetic evaluation

is/2 (evaluate expression)

Atomic term processing

atom_chars/2 (conversion of atoms to character codes and vice versa)

Character input

get0/1

Clause creation and destruction

asserta/1 (clause creation)

retract/1 (clause destruction)

File consultation

[File|Files] (consult list of files)

List operation

length/2] (determine length of a list)

Logic and control

,/2 (conjunction)

;/2 (disjunction)

!/0 (cut)

fail/0

true/0

\+ (not provable)

Operator definition

op/3 (extending operator table)

Term comparison

==/2 (term identical)

Term creation and decomposition

arg/3

functor/3

=../2 (the “univ”)

Term unification

= (unify)

Type testing

atom/1

atomic/1

compound/1

nonvar/1

var/1

Term output

write/1

nl/0

The following is a practical overview of the programs in this appendix. We start with the library files consulted by most of the other programs, and then give a chapter by chapter break-down description of the files.

All the files at one glance

File Name	Chapter	Page
<code>modelDRT.pl</code>	Chapter 1	p. 173
<code>drs2fol.pl</code>	Chapter 1	p. 175
<code>threadingDRT.pl</code>	Chapter 2	p. 177
<code>mainLambdaDRT.pl</code>	Chapter 2	p. 180
<code>semMacrosLambdaDRT.pl</code>	Chapter 2	p. 181
<code>mergeDRT.pl</code>	Chapter 2	p. 182
<code>mainDRTU.pl</code>	Chapter 2	p. 183
<code>semMacrosDRTU.pl</code>	Chapter 2	p. 184
<code>mainPronounsDRT.pl</code>	Chapter 3	p. 186
<code>bindingDRT.pl</code>	Chapter 3	p. 188
<code>mainFocusDRT.pl</code>	Chapter 3	p. 190
<code>matchDRT.pl</code>	Chapter 4	p. 199
<code>mainPresupDRT.pl</code>	Chapter 4	p. 196
<code>resolvePresup.pl</code>	Chapter 4	p. 197
<code>semMacrosPresupDRT.pl</code>	Chapter 4	p. 204
<code>mainPresupScoreDRT.pl</code>	Chapter 4	p. 200
<code>resolvePresupScore.pl</code>	Chapter 4	p. 201
<code>curt.pl</code>	Chapter 5	p. 206
<code>acceptabilityConstraints.pl</code>	Chapter 5	p. 209
<code>callTheoremProver.pl</code>	Chapter 5	p. 213
<code>callModelBuilder.pl</code>	Chapter 5	p. 214
<code>fol2otter.pl</code>	Chapter 5	p. 215
<code>mainPresupDRTU.pl</code>	Chapter 6	p. 217
<code>semMacrosPresupDRTU.pl</code>	Chapter 6	p. 218


```
/******  
  
    name: comsemOperators.pl  
    version: May 25, 1999  
    description: Operator definitions  
    authors: Patrick Blackburn & Johan Bos  
  
*****/  
  
:- module(comsemOperators, []).  
  
/*=====Operator Definitions=====*/  
  
:- op(950,yfx,@).      % application  
:- op(900,yfx,'<>').  % bin impl  
:- op(900,yfx,>).     % implication  
:- op(850,yfx,v).     % disjunction  
:- op(800,yfx,&).     % conjunction  
:- op(750, fy,~).     % negation
```

```

/*****

    name: comsemPredicates.pl
    version: November 8, 1997
    description: Set of Prolog predicates
    authors: Patrick Blackburn & Johan Bos

*****/

:- module(comsemPredicates,
    [member/2,
     select/3,
     append/3,
     simpleTerms/1,
     compose/3,
     unify/2,
     removeFirst/3,
     substitute/4,
     variablesInTerm/2,
     newFunctionCounter/1,
     printReadings/1,
     printRepresentation/1]).

/*=====

    List Manipulation
    -----

member(?Element,?List)
    Element occurs in List.

select(?X,+OldList,?NewList).
    X is removed from OldList, result in NewList.

append(?List1,?List2,?List3)
    List3 is the concatenation of List1 and List2.

allExactMember(?Elements,?List)
    Elements all occur in List (no unification, exact match).

removeAll(?Item,?List,?Newlist)
    Newlist is the result of removing all occurrences of Item from List.

removeFirst(?Item,?List,?Newlist)
    Newlist is the result of removing the first occurrence of Item from
    List. Fails when Item is not member of List.

=====*/

member(X,[X|_]).

```

```

member(X,[_|Tail]):-
    member(X, Tail).

select(X,[X|L],L).
select(X,[Y|L1],[Y|L2]):-
    select(X,L1,L2).

append([],List,List).
append([X|Tail1],List,[X|Tail2]):-
    append(Tail1,List,Tail2).

allExactMember([],_).
allExactMember([X|R],L):-
    memberOfList(Y,L),
    X==Y,
    allExactMember(R,L).

removeAll(_,[],[]).
removeAll(X,[X|Tail],Newtail):-
    removeAll(X,Tail,Newtail).
removeAll(X,[Head|Tail],[Head|Newtail]):-
    X \== Head,
    removeAll(X,Tail,Newtail).

removeFirst(X,[X|Tail],Tail) :- !.
removeFirst(X,[Head|Tail],[Head|NewTail]):-
    removeFirst(X,Tail,NewTail).

/*=====

    Term Manipulation
    -----

simpleTerms(?List)
    List is a list of elements that are currently uninstantiated or
    instantiated to an atom or number. Uses built-in Quintus/Sicstus
    predicate simple/1.

compose(?Term,+Symbol,+ArgList)
compose(+Term,?Symbol,?ArgList)
    Composes a complex Term with functor Symbol and arguments ArgList.
    Uses the Prolog built-in =.. predicate.

variablesInTerm(+Term,?InList-?OutList)
    Adds all occurrences of variables in Term (arbitrarily deeply
    nested to the difference list InList-OutList.

=====*/

simpleTerms([]).

```

```
simpleTerms([X|Rest]):-
    simple(X), simpleTerms(Rest).
```

```
compose(Term,Symbol,ArgList):-
    Term =.. [Symbol|ArgList].
```

```
variablesInTerm(Term,Var1-Var2):-
    compose(Term,_,Args),
    countVar(Args,Var1-Var2).
```

```
countVar([],Var-Var).
countVar([X|Rest],Var1-Var2):-
    var(X),!,
    countVar(Rest,[X|Var1]-Var2).
countVar([X|Rest],Var1-Var3):-
    variablesInTerm(X,Var1-Var2),
    countVar(Rest,Var2-Var3).
```

```
/*=====
```

```
    Unification Predicates
    -----
```

```
unify(Term1,Term2)
    Unify Term1 with Term2 including occurs check. Adapted from
    "The Art of Prolog" by Sterling & Shapiro, MIT Press 1986, page 152.
```

```
notOccursIn(X,Term)
    Succeeds if variable X does not occur in Term.
```

```
notOccursInComplexTerm(N,X,Term)
    Succeeds if variable X does not occur in complex Term with arity N
```

```
termUnify(Term1,Term2)
    Unify the complex terms Term1 and Term2.
```

```
=====*/
```

```
unify(X,Y):-
    var(X), var(Y), X=Y.
unify(X,Y):-
    var(X), nonvar(Y), notOccursIn(X,Y), X=Y.
unify(X,Y):-
    var(Y), nonvar(X), notOccursIn(Y,X), X=Y.
unify(X,Y):-
    nonvar(X), nonvar(Y), atomic(X), atomic(Y), X=Y.
unify(X,Y):-
    nonvar(X), nonvar(Y), compound(X), compound(Y), termUnify(X,Y).
```

```

notOccursIn(X,Term):-
    var(Term), X \== Term.
notOccursIn(_,Term):-
    nonvar(Term), atomic(Term).
notOccursIn(X,Term):-
    nonvar(Term), compound(Term),
    functor(Term,_,Arity), notOccursInComplexTerm(Arity,X,Term).

notOccursInComplexTerm(N,X,Y):-
    N > 0, arg(N,Y,Arg), notOccursIn(X,Arg),
    M is N - 1, notOccursInComplexTerm(M,X,Y).
notOccursInComplexTerm(0,_,_).

termUnify(X,Y):-
    functor(X,Func,Ari), functor(Y,Func,Ari),
    unifyArgs(Ari,X,Y).

unifyArgs(N,X,Y):-
    N > 0, M is N - 1,
    arg(N,X,ArgX), arg(N,Y,ArgY),
    unify(ArgX,ArgY), unifyArgs(M,X,Y).
unifyArgs(0,_,_).

/*=====

    Substitution Predicates
    -----

substitute(?Term,?Variable,+Exp,-Result)
    Result is the result of substituting occurrences of Term for each
    free occurrence of Variable in Exp.

=====*/

substitute(Term,Var,Exp,Result):-
    Exp==Var, !, Result=Term.
substitute(_Term,_Var,Exp,Result):-
    \+ compound(Exp), !, Result=Exp.
substitute(Term,Var,Formula,Result):-
    compose(Formula,Func,[Exp,F]),
    member(Func,[lambda forall exists]), !,
    (
        Exp==Var, !,
        Result=Formula
    );
    substitute(Term,Var,F,R),
    compose(Result,Func,[Exp,R])
).
substitute(Term,Var,Formula,Result):-
    compose(Formula,Func,ArgList),

```

```

    substituteList(Term,Var,ArgList,ResultList),
    compose(Result,Funcor,ResultList).

substituteList(_Term,_Var,[],[]).
substituteList(Term,Var,[Exp|Others],[Result|ResultOthers]):-
    substitute(Term,Var,Exp,Result),
    substituteList(Term,Var,Others,ResultOthers).

/*=====

    Skolem Function Counter
    -----

functionCounter(?N)
    N is the current Skolem function index. Declared as dynamic,
    and set to value 1.

newFunctionCounter(?N)
    Unifies N with the current Skolem function index, and increases
    value of the counter.

=====*/

:- dynamic(functionCounter/1).

functionCounter(1).

newFunctionCounter(N):-
    functionCounter(N), M is N+1,
    retract(functionCounter(N)),
    asserta(functionCounter(M)).

/*=====

    Pretty Print Predicates
    -----

=====*/

printRepresentation(Rep):-
    nl, \+ \+ (numbervars(Rep,0,_), write(Rep)), nl.

printReadings(Readings):-
    nl, write('Readings: '), nl, printReading(Readings,0).

printReading([],N):-
    nl, (N=0, write('no readings'); true), nl.
printReading([Reading|OtherReadings],M):-
    N is M + 1, write(N), tab(1),
    \+ \+ (numbervars(Reading,0,_), write(Reading)), nl,

```

```
printReading(OtherReadings,N).
```

```

/*****
    name: readLine.pl
    version: March 31, 1998
    description: Converting input line to list of atoms, suitable for
                DCG input.
    authors: Patrick Blackburn & Johan Bos
*****/

:- module(readLine,[readLine/1]).

/*=====

    Read Predicates
    -----

readLine(-WordList)
    Outputs a prompt, reads a sequence of characters from the standard
    input and converts this to WordList, a list of strings. Punctuation
    is stripped.

readWords(-WordList)
    Reads in a sequence of characters, until a return is registered,
    and converts this to WordList a list of strings.

readWord(+Char,-Chars,?State)
    Read a word coded as Chars (a list of ascii values), starting
    with with ascii value Char, and determine the State of input
    ('ended' = end of line, 'notended' = not end of line).
    Blanks and full stops split words, a return ends a line.

checkWords(+OldWordList,-NewWordList)
    Check if all words are unquoted atoms, if not convert them
    into atoms.

convertWord(+OldWord,-NewWord)
    OldWord and NewWord are words represented as lists of ascii values.
    Converts upper into lower case characters, and eliminates
    non-alphabetic characters.

=====*/

readLine(WordList):-
    nl, write('> '),
    readWords(WordList),
    checkWords(WordList,WordList).

readWords([Word|Rest]):-
    get0(Char),

```



```

    readWord(Char, Chars, State),
    atom_chars(Word, Chars),
    readRest(Rest, State).

readRest([], ended).
readRest(Rest, notended):-
    readWords(Rest).

readWord(32, [], notended):-!.      %%% blank
readWord(46, [], notended):-!.      %%% full stop
readWord(10, [], ended):-!.         %%% return
readWord(Code, [Code|Rest], State):-
    get0(Char),
    readWord(Char, Rest, State).

checkWords([], []):-!.
checkWords([''|Rest1], Rest2):-
    checkWords(Rest1, Rest2).
checkWords([Atom|Rest1], [Atom2|Rest2]):-
    atom_chars(Atom, Word1),
    convertWord(Word1, Word2),
    atom_chars(Atom2, Word2),
    checkWords(Rest1, Rest2).

convertWord([], []):-!.
convertWord([Capital|Rest1], [Small|Rest2]):-
    Capital > 64, Capital < 91, !,
    Small is Capital + 32,
    convertWord(Rest1, Rest2).
convertWord([Weird|Rest1], Rest2):-
    (Weird < 97; Weird > 122), !,
    convertWord(Rest1, Rest2).
convertWord([Char|Rest1], [Char|Rest2]):-
    convertWord(Rest1, Rest2).

```

```

/*****

```

```

    name: englishLexicon.pl
    version: November 12, 1997; March 9, 1999.
    description: Lexical entries for a small coverage of English
    authors: Patrick Blackburn & Johan Bos

```

This file contains the lexical entries for a small fragment of English. Entries have the form `lexicon(Cat,Sym,Phrase,Misc)`, where `Cat` is the syntactic category, `Sym` the predicate symbol introduced by the phrase, `Phrase` a list of the words that form the phrase, and `Misc` miscellaneous information depending on the the type of entry.

```

*****/

```

```

/*=====
   Determiners: lexicon(det,_,Words,Type)
=====*/

```

```

lexicon(det,_, [every], uni).
lexicon(det,_, [a], indef).
lexicon(det,_, [the], def).
lexicon(det,_, [one], card(1)).
lexicon(det,_, [another], alt).
lexicon(det,_, [his], poss(male)).
lexicon(det,_, [her], poss(female)).
lexicon(det,_, [its], poss(nonhuman)).

```

```

/*=====
   Nouns: lexicon(noun,Symbol,Words, {[], [Hypernym], Hypernym})
=====*/

```

```

lexicon(noun,abstraction, [abstraction], [top]).
lexicon(noun,act, [act], [top]).
lexicon(noun,animal, [animal], [organism]).
lexicon(noun,artifact, [artifact], [object]).
lexicon(noun,beverage, [beverage], [food]).
lexicon(noun,building, [building], [artifact]).
lexicon(noun,container, [container], [instrumentality]).
lexicon(noun,cup, [cup], [container]).
lexicon(noun,device, [device], [instrumentality]).
lexicon(noun,edible, [edible,food], [food]).
lexicon(noun,bkburger, [big,kahuna,burger], [edible]).
lexicon(noun,boxer, [boxer], human).
lexicon(noun,boss, [boss], human).
lexicon(noun,car, [car], [vehicle]).
lexicon(noun,chainsaw, [chainsaw], [device]).
lexicon(noun,criminal, [criminal], human).
lexicon(noun,customer, [customer], human).
lexicon(noun,drug, [drug], [artifact]).

```

```

lexicon(noun,entity,[entity],[top]).
lexicon(noun,episode,[episode],abstraction).
lexicon(noun,female,[female],[human]).
lexicon(noun,fdshake,[five,dollar,shake],[beverage]).
lexicon(noun,food,[food],[object]).
lexicon(noun,footmassage,[foot,massage],[act]).
lexicon(noun,gimp,[gimp],human).
lexicon(noun,glass,[glass],[container]).
lexicon(noun,gun,[gun],[weaponry]).
lexicon(noun,hammer,[hammer],[device]).
lexicon(noun,hashbar,[hash,bar],[building]).
lexicon(noun,human,[human],[organism]).
lexicon(noun,husband,[husband],male).
lexicon(noun,instrumentality,[instrumentality],artifact).
lexicon(noun,joke,[joke],abstraction).
lexicon(noun,man,[man],male).
lexicon(noun,male,[male],[human]).
lexicon(noun,medium,[medium],[instrumentality]).
lexicon(noun,needle,[needle],[device]).
lexicon(noun,object,[object],[entity]).
lexicon(noun,organism,[organism],[entity]).
lexicon(noun,owner,[owner],human).
lexicon(noun,piercing,[piercing],[act]).
lexicon(noun,plant,[plant],[organism]).
lexicon(noun,qpwc,[quarter,pounder,with,cheese],[edible]).
lexicon(noun,radio,[radio],[medium]).
lexicon(noun,restaurant,[restaurant],[building]).
lexicon(noun,robber,[robber],human).
lexicon(noun,suitcase,[suitcase],[container]).
lexicon(noun,shotgun,[shotgun],[weaponry]).
lexicon(noun,sword,[sword],[weaponry]).
lexicon(noun,vehicle,[vehicle],[instrumentality]).
lexicon(noun,weaponry,[weaponry],[instrumentality]).
lexicon(noun,woman,[woman],female).

```

```

/*=====
   Proper Names: lexicon(pn,Symbol,Words,{male,female})
=====*/

```

```

lexicon(pn,butch,[butch],male).
lexicon(pn,honey_bunny,[honey,bunny],male).
lexicon(pn,jimmy,[jimmy],male).
lexicon(pn,jody,[jody],female).
lexicon(pn,jules,[jules],male).
lexicon(pn,lance,[lance],male).
lexicon(pn,marsellus,[marsellus],male).
lexicon(pn,marsellus,[marsellus,wallace],male).
lexicon(pn,marvin,[marvin],male).
lexicon(pn,mia,[mia],female).
lexicon(pn,mia,[mia,wallace],female).

```

```
lexicon(pn,pumpkin,[pumpkin],male).
lexicon(pn,thewolf,[the,wolf],male).
lexicon(pn,vincent,[vincent],male).
lexicon(pn,vincent,[vincent,vega],male).
```

```
/*=====
   Intransitive Verbs: lexicon(iv,Symbol,Words,{fin,inf})
=====*/
```

```
lexicon(iv,collapse,[collapses],fin).
lexicon(iv,collapse,[collapse],inf).
lexicon(iv,dance,[dances],fin).
lexicon(iv,dance,[dance],inf).
lexicon(iv,die,[dies],fin).
lexicon(iv,die,[die],inf).
lexicon(iv,growl,[growls],fin).
lexicon(iv,growl,[growl],inf).
lexicon(iv,okay,[is,okay],fin).
lexicon(iv,outoftown,[is,out,of,town],fin).
lexicon(iv,married,[is,married],fin).
lexicon(iv,playairguitar,[plays,air,guitar],fin).
lexicon(iv,playairguitar,[play,air,guitar],inf).
lexicon(iv,smoke,[smokes],fin).
lexicon(iv,smoke,[smoke],inf).
lexicon(iv,snort,[snorts],fin).
lexicon(iv,snort,[snort],inf).
lexicon(iv,shriek,[shrieks],fin).
lexicon(iv,shriek,[shriek],inf).
lexicon(iv,walk,[walks],fin).
lexicon(iv,walk,[walk],inf).
```

```
/*=====
   Transitive Verbs: lexicon(tv,Symbol,Words,{fin,inf})
=====*/
```

```
lexicon(tv,clean,[cleans],fin).
lexicon(tv,clean,[clean],inf).
lexicon(tv,drink,[drinks],fin).
lexicon(tv,drink,[drink],inf).
lexicon(tv,date,[dates],fin).
lexicon(tv,date,[date],inf).
lexicon(tv,discard,[discards],fin).
lexicon(tv,discard,[discard],inf).
lexicon(tv,eat,[eats],fin).
lexicon(tv,eat,[eat],inf).
lexicon(tv,enjoy,[enjoys],fin).
lexicon(tv,enjoy,[enjoy],inf).
lexicon(tv,hate,[hates],fin).
lexicon(tv,hate,[hate],inf).
lexicon(tv,have,[has],fin).
```

```

lexicon(tv,have,[have],inf).
lexicon(tv,donewith,[is,done,with],fin).
lexicon(tv,kill,[kills],fin).
lexicon(tv,kill,[kill],inf).
lexicon(tv,know,[knows],fin).
lexicon(tv,know,[know],inf).
lexicon(tv,like,[likes],fin).
lexicon(tv,like,[like],inf).
lexicon(tv,love,[loves],fin).
lexicon(tv,love,[love],inf).
lexicon(tv,pickup,[picks,up],fin).
lexicon(tv,pickup,[pick,up],inf).
lexicon(tv,shoot,[shot],fin).
lexicon(tv,shoot,[shoot],inf).
lexicon(tv,tell,[told],fin).
lexicon(tv,tell,[tell],inf).
lexicon(tv,worksfor,[works,for],fin).
lexicon(tv,worksfor,[work,for],inf).

/*=====
   Copula
=====*/

lexicon(cop,'=',[is],fin).

/*=====
   Prepositions: lexicon(preop,Symbol,Words,_)
=====*/

lexicon(preop,in,[in],_).
lexicon(preop,of,[of],_).
lexicon(preop,with,[with],_).

/*=====
   Pronouns: lexicon(pro,Sym,Words,{refl,nonrefl})
=====*/

lexicon(pro,male,[he],nonrefl).
lexicon(pro,female,[she],nonrefl).
lexicon(pro,nonhuman,[it],nonrefl).
lexicon(pro,male,[him],nonrefl).
lexicon(pro,female,[her],nonrefl).
lexicon(pro,male,[himself],refl).
lexicon(pro,female,[herself],refl).
lexicon(pro,nonhuman,[itself],refl).

/*=====
   Relative Pronouns: lexicon(relpro,_,Words,_)
=====*/

```

```
lexicon(relpro,_, [who],_).
lexicon(relpro,_, [that],_).

/*=====
   Coordinations: lexicon(coord,_,Words,{conj,disj})
=====*/

lexicon(coord,_, [and], conj).
lexicon(coord,_, [or], disj).

/*=====
   Discontinious Coordinations: lexicon(dcoord,W1,W2,{conj,cond,disj})
=====*/

lexicon(dcoord, [if], [then], cond).
lexicon(dcoord, [if], [], cond).
lexicon(dcoord, [either], [or], disj).
lexicon(dcoord, [], [or], disj).
lexicon(dcoord, [], [and], conj).
lexicon(dcoord, [], [], conj).

/*=====
   Modifiers: lexicon(mod,_,Words,Type)
=====*/

lexicon(mod,_, [does,not], neg).
lexicon(mod,_, [did,not], neg).
```

```

/*****
    name: englishGrammar.pl
    version: November 12, 1997
    description: Grammar rules for a small coverage of English
    authors: Patrick Blackburn & Johan Bos
*****/

/*=====
    Grammar Rules
=====*/

d(S) --> s(S).
d((C@S1)@S2)--> dcoord(D,C), s(S1), D, d(S2).

s(NP@VP)--> np2(NP), vp2(VP).

np2(NP)--> np1(NP).
np2((C@NP1)@NP2)--> np1(NP1), coord(C), np1(NP2).

np1(Det@Noun)--> det(Det), n2(Noun).
np1(NP)--> pn(NP).
np1(NP)--> pro(NP).

n2(N)--> n1(N).
n2((C@N1)@N2)--> n1(N1), coord(C), n1(N2).

n1(N)--> noun(N).
n1(PP@N)--> noun(N), pp(PP).
n1(RC@N)--> noun(N), rc(RC).

vp2(VP)--> vp1(VP).
vp2((C@VP1)@VP2)--> vp1(VP1), coord(C), vp1(VP2).

vp1(Mod@VP)--> mod(Mod), v2(fin,VP).
vp1(VP)--> v2(fin,VP).

v2(fin,Cop@NP)--> cop(Cop), np2(NP).
v2(fin,Neg@(Cop@NP))--> cop(Cop), neg(Neg), np2(NP).

v2(I,V)--> v1(I,V).
v2(I,(C@V1)@V2)--> v1(I,V1), coord(C), v1(I,V2).

v1(I,V)--> iv(I,V).
v1(I,TV@NP)--> tv(I,TV), np2(NP).

pp(Prep@NP)--> prep(Prep), np2(NP).

rc(RP@VP)--> relpro(RP), vp2(VP).

```

iv(I,IV)--> {lexicon(iv,Sym,Word,I),ivSem(Sym,IV)}, Word.
tv(I,TV)--> {lexicon(tv,Sym,Word,I),tvSem(Sym,TV)}, Word.
cop(Cop)--> {lexicon(cop,Sym,Word,_),tvSem(Sym,Cop)}, Word.
det(Det)--> {lexicon(det,_,Word,Type),detSem(Type,Det)}, Word.
pn(PN)--> {lexicon(pn,Sym,Word,G),pnSem(Sym,G,PN)}, Word.
pro(Pro)--> {lexicon(pro,Gender,Word,Type),proSem(Gender,Type,Pro)}, Word.
noun(N)--> {lexicon(noun,Sym,Word,_),nounSem(Sym,N)}, Word.
relpro(RP)--> {lexicon(relpro,_,Word,_),relproSem(RP)}, Word.
prep(Prep)--> {lexicon(prepare,Sym,Word,_),prepSem(Sym,Prep)}, Word.
mod(Mod)--> {lexicon(mod,_,Word,Type),modSem(Type,Mod)}, Word.
neg(Neg)--> [not], {modSem(neg,Neg)}.
coord(C)--> {lexicon(coord,_,Word,Type), coordSem(Type,C)}, Word.
dcoord(D,C)--> {lexicon(dcoord,Word,D,Type), dcoordSem(Type,C)}, Word.


```

/*****

    name: modelDRT.pl (Chapter 7)
    version: July 24, 1997
    description: Model Evaluation for DRSs
    authors: Patrick Blackburn & Johan Bos

*****/

:- module(modelDRT,[evaluate/2]).

:- use_module(comsemOperators),
   use_module(comsemPredicates,[member/2]),
   use_module(exampleModels,[constant/1,example/2]).

/*=====
   Semantic Interpretation
   =====*/

satisfyDrs(drs(Dom,C),Model):-
    assignReferents(Dom),
    satisfyConditions(C,Model).

assignReferents([]).
assignReferents([Referent|Others]):-
    constant(Referent),
    assignReferents(Others).

satisfyConditions([],_).
satisfyConditions([Condition|Others],Model):-
    satisfyCondition(Condition,Model),
    satisfyConditions(Others,Model).

satisfyCondition(~ Drs,Model):-
    \+ satisfyDrs(Drs,Model).

satisfyCondition(Drs1 > Drs2,Model):-
    (
        satisfyDrs(Drs1,Model),
        \+ satisfyDrs(Drs2,Model),
        !,
        fail
    );
    true
).

satisfyCondition(Drs1 v Drs2,Model):-
    (
        satisfyDrs(Drs1,Model)
    );

```

```
    satisfyDrs(Drs2,Model)
  ).

satisfyCondition(X=Y,_Model):-
  X=Y.

satisfyCondition(BasicCond,Model):-
  member(BasicCond,Model).

/*=====
  Evaluation
=====*/

evaluate(Drs,Example):-
  example(Example,Model),
  satisfyDrs(Drs,Model).
```

```

/*****
    name: drs2fol.pl (Chapter 7)
    version: June 8, 1998
    description: From DRSs to First-Order Logic
    authors: Patrick Blackburn & Johan Bos
*****/

:- module(drs2fol, [drs2fol/2]).

:- use_module(comsemOperators),
   use_module(comsemPredicates, [compose/3, simpleTerms/1]).

/*=====

    Translation Predicates
    -----

drs2fol(+Drs, ?F)
    converts DRS to formula F. DRS is a term drs(Dom,Cond), where
    Dom is a list of discourse referents, and Cond is a non empty
    list of DRS-conditions.

=====*/

drs2fol(drs([], [Cond]), Formula):-
    cond2fol(Cond, Formula).

drs2fol(drs([], [Cond1, Cond2|Conds]), Formula1 & Formula2):-
    cond2fol(Cond1, Formula1),
    drs2fol(drs([], [Cond2|Conds]), Formula2).

drs2fol(drs([X|Referents], Conds), exists(X, Formula)):-
    drs2fol(drs(Referents, Conds), Formula).

cond2fol(~ Drs, ~ Formula):-
    drs2fol(Drs, Formula).

cond2fol(Drs1 v Drs2, Formula1 v Formula2):-
    drs2fol(Drs1, Formula1),
    drs2fol(Drs2, Formula2).

cond2fol(drs([], Conds) > Drs2, Formula1 > Formula2):-
    drs2fol(drs([], Conds), Formula1),
    drs2fol(Drs2, Formula2).

cond2fol(drs([X|Referents], Conds) > Drs2, forall(X, Formula)):-
    cond2fol(drs(Referents, Conds) > Drs2, Formula).

```

```
cond2fol(BasicCondition,AtomicFormula):-  
  AtomicFormula=BasicCondition,  
  compose(BasicCondition,_Symbol,Arguments),  
  simpleTerms(Arguments).
```

```

/*****

    name: threadingDRT.pl (Chapter 8)
    version: June 8, 1998
    description: DRS-threading (Johnson & Klein 1986)
    authors: Patrick Blackburn & Johan Bos

*****/

:- module(threadingDRT, [parse/0]).

:- use_module(readLine, [readLine/1]),
   use_module(comsemOperators),
   use_module(comsemPredicates, [printRepresentation/1, compose/3]).

:- [englishLexicon].

/*=====
   Driver Predicate
=====*/

parse:-
  readLine(Discourse),
  d(drs([], [])-Drs, Discourse, []),
  printRepresentation(Drs).

/*=====
   Grammar Rules
=====*/

d(Drs In-Drs Out)-->
  s(Drs In-Drs),
  d(Drs-DrsOut).
d(Drs-Drs)-->
  [].

s(Drs In-Drs Out)-->
  np(X, Drs In-Drs Out, Scope),
  vp(X, Scope).

np(X, Sem, Scope)-->
  det(X, Sem, Restr, Scope),
  noun(X, Restr).

np(X, Drs In-Drs Out, Drs-DrsOut)-->
  pn(X, Drs In-Drs).

vp(X, Drs In-Drs Out)-->
  iv(X, Drs In-Drs Out).

```

```

vp(X, DrsIn-DrsOut)-->
  tv(X, Y, Scope),
  np(Y, DrsIn-DrsOut, Scope).

/*=====
  Determiners
=====*/

det(X, DrsIn-DrsOut, RestrIn-RestrOut, ScopeIn-ScopeOut)-->
{
  lexicon(det, _, Phrase, indef),
  DrsIn = drs(Dom, Conds),
  DrsOut = ScopeOut,
  RestrIn = drs([X|Dom], Conds),
  RestrOut = ScopeIn
},
Phrase.

det(X, DrsIn-DrsOut, RestrIn-RestrOut, ScopeIn-ScopeOut)-->
{
  lexicon(det, _, Phrase, uni),
  DrsIn = drs(Dom, Conds),
  DrsOut = drs(Dom, [RestrOut > ScopeOut|Conds]),
  RestrIn = drs([X], []),
  ScopeIn = drs([], [])
},
Phrase.

/*=====
  Common Nouns
=====*/

noun(X, drs(Dom, Conds)-drs(Dom, [Cond|Conds])) -->
{
  lexicon(noun, Sym, Phrase, _Type),
  compose(Cond, Sym, [X])
},
Phrase.

/*=====
  Proper Names
=====*/

pn(X, drs(Dom, Conds)-drs([X|Dom], [X=Sym|Conds]))-->
{
  lexicon(pn, Sym, Phrase, _Gender)
},
Phrase.

/*=====

```

Intransitive Verbs

```
=====*/  
iv(X,drs(Dom,Conds)-drs(Dom,[Cond|Conds]))-->  
  {  
    lexicon(iv,Sym,Phrase,_),  
    compose(Cond,Sym,[X])  
  },  
  Phrase.
```

```
/*=====/  
  Transitive Verbs  
=====*/
```

```
tv(X,Y,drs(Dom,Conds)-drs(Dom,[Cond|Conds]))-->  
  {  
    lexicon(tv,Sym,Phrase,_),  
    compose(Cond,Sym,[Y,X])  
  },  
  Phrase.
```

```

/*****

    name: mainLambdaDRT.pl (Chapter 8)
    version: March 22, 1999
    description: Semantic Construction with Beta Conversion for DRT
    authors: Patrick Blackburn & Johan Bos

*****/

:- module(mainLambdaDRT,[parse/0]).

:- use_module(readLine,[readLine/1]),
   use_module(comsemOperators),
   use_module(comsemPredicates,[printRepresentation/1,compose/3]),
   use_module(betaConversion,[betaConvert/2]),
   use_module(mergeDRT,[mergeDrs/2]).

:- [englishGrammar], [englishLexicon], [semMacrosLambdaDRT].

/*=====
   Driver Predicate
   =====*/

parse:-
  readLine(Discourse),
  statistics(walltime,_),
  d(MergeDrs,Discourse,[]),
  betaConvert(MergeDrs,ReducedDrs),
  mergeDrs(ReducedDrs,Drs),
  statistics(walltime,[_,Time]),
  nl, write(Time),
  printRepresentation(Drs).

```



```

/*****
    name: semMacroslambdaDRT.pl (Chapter 8)
    version: June 9, 1998
    description: Compositional DRS construction for a fragment of English
    authors: Patrick Blackburn & Johan Bos
*****/

/*=====
    Semantic Macros
=====*/

detSem(uni,lambda(P,lambda(Q,drs([], [merge(drs([X], []),P@X)>(Q@X)])))).
detSem(indef,lambda(P,lambda(Q,merge(merge(drs([X], []),P@X),Q@X))).

nounSem(Sym,lambda(X,drs([], [Cond]))):-
    compose(Cond,Sym,[X]).

pnSem(Sym,Gender,lambda(P,merge(drs([X],[Cond,X=Sym]),P@X))):-
    compose(Cond,Gender,[X]).

proSem(Gender,Type,lambda(P,alfa(X,Type,Cond,P@X))):-
    compose(Cond,Gender,[X]).

ivSem(Sym,lambda(X,drs([], [Cond]))):-
    compose(Cond,Sym,[X]).

tvSem(Sym,lambda(K,lambda(Y,K @ lambda(X,drs([], [Cond]))))):-
    compose(Cond,Sym,[Y,X]).

relproSem(lambda(P1,lambda(P2,lambda(X,merge(P1@X,P2@X)))).

prepSem(Sym,lambda(K,lambda(P,lambda(Y,Drs))):-
    Drs=merge(K@lambda(X,drs([], [Cond])),P@Y),
    compose(Cond,Sym,[Y,X]).

modSem(neg,lambda(P,lambda(X,drs([], [~(P @ X)])))).

coordSem(conj,lambda(X,lambda(Y,lambda(P,merge(X@P,Y@P)))).
coordSem(disj,lambda(X,lambda(Y,lambda(P,drs([], [(X@P) v (Y@P)])))).

dcoordSem(cond,lambda(X,lambda(Y,drs([], [X > Y]))).
dcoordSem(conj,lambda(X,lambda(Y,merge(X,Y))).
dcoordSem(disj,lambda(X,lambda(Y,drs([], [X v Y]))).

```

```

/*****
    name: mergeDRT.pl (Chapter 8)
    version: June 9, 1998
    description: Definition of the merge for DRsSs
    authors: Patrick Blackburn & Johan Bos

*****/

:- module(mergeDRT, [mergeDrs/2]).

:- use_module(comsemOperators),
   use_module(comsemPredicates, [append/3]).

/*=====
   DRS-merge
   =====*/

mergeDrs(drs(D,C1),drs(D,C2)):-
    mergeDrs(C1,C2).
mergeDrs(merge(B1,B2),drs(D3,C3)):-
    mergeDrs(B1,drs(D1,C1)), mergeDrs(B2,drs(D2,C2)),
    append(D1,D2,D3), append(C1,C2,C3).
mergeDrs([B1 > B2|C1],[B3 > B4|C2]):- !,
    mergeDrs(B1,B3), mergeDrs(B2,B4), mergeDrs(C1,C2).
mergeDrs([B1 v B2|C1],[B3 v B4|C2]):- !,
    mergeDrs(B1,B3), mergeDrs(B2,B4), mergeDrs(C1,C2).
mergeDrs([~ B1|C1],[~ B2|C2]):- !,
    mergeDrs(B1,B2), mergeDrs(C1,C2).
mergeDrs([C|C1],[C|C2]):-
    mergeDrs(C1,C2).
mergeDrs([],[]).

```

```

/*****

    name: mainDRTU.pl (Chapter 8)
    version: May 28, 1998
    description: Plugging Stuff for Underspecified DRSs
    author: Patrick Blackburn & Johan Bos

*****/

:- module(mainDRTU,[parse/0]).

:- use_module(readLine,[readLine/1]),
   use_module(comsemOperators),
   use_module(betaConversion,[betaConvert/2]),
   use_module(pluggingAlgorithm,[plugHole/4]),
   use_module(mergeDRT,[mergeDrs/2]),
   use_module(mergeUSR,[mergeUSR/2]),
   use_module(comsemPredicates,[append/3,printRepresentation/1,
                                printReadings/1,compose/3]).

:- [englishGrammar],[englishLexicon],[semMacrosDRTU].

/*=====
   Driver Predicate
   =====*/

parse:-
  readLine(Sentence),
  d(Sem,Sentence,[]),
  betaConvert(merge(usr([Top,Main],[],[]),Sem@Top@Main),Reduced),
  mergeUSR(Reduced,usr(D,L,C)),
  printRepresentation(usr(D,L,C)),
  findall(Drs,(plugHole(Top,L-[],C,[]),mergeDrs(Top,Drs)),DRSs),
  printReadings(DRSs).

```

```

/*****

    name: semMacrosDRTU.pl (Chapter 8)
    version: May 28, 1998
    description: Semantic Macros for DRT Unplugged
    author: Patrick Blackburn & Johan Bos

*****/

/*=====
   Semantic Macros
=====*/

detSem(uni, lambda(P, lambda(Q, lambda(H, lambda(L,
    merge(merge(usr([F, R, S],
                [F:drs([], [merge(drs([X], []), R)>S])),
                [leq(F, H), leq(L, S)]), P@X@H@R), Q@X@H@L)))))).

detSem(indef, lambda(P, lambda(Q, lambda(H, lambda(L,
    merge(merge(usr([E, R, S],
                [E:merge(merge(drs([X], []), R), S)],
                [leq(E, H), leq(L, S)]), P@X@H@R), Q@X@H@L)))))).

nounSem(Sym, lambda(X, lambda(_, lambda(L, usr([], [L:drs([], [Cond])], []))))):-
    compose(Cond, Sym, [X]).

pnSem(Sym, _, lambda(P, lambda(H, lambda(M,
    merge(usr([L], [M:merge(drs([X], [X=Sym]), L)], [leq(M, H)]),
    P@X@H@L))))).

proSem(_, _, lambda(P, lambda(H, lambda(M,
    merge(usr([L], [M:merge(drs([X], []), L)], [leq(M, H)]),
    P@X@H@L))))).

ivSem(Sym, lambda(X, lambda(_, lambda(L,
    usr([], [L:drs([], [Cond])], []))))):-
    compose(Cond, Sym, [X]).

tvSem(Sym, lambda(K, lambda(Y, K@lambda(X, lambda(_, lambda(L,
    usr([], [L:drs([], [Cond])], []))))))):-
    compose(Cond, Sym, [Y, X]).

relproSem(lambda(P, lambda(Q, lambda(X, lambda(H, lambda(L,
    merge(usr([L2, L3, H1],
                [L:merge(L3, H1)],
                [leq(L2, H1)]),
    merge(P@X@H@L2, Q@X@H@L3)))))))).

prepSem(Sym, lambda(K, lambda(P, lambda(Y, lambda(H, lambda(L3,
    merge(K@lambda(X, lambda(H, lambda(L1,

```

```

        usr([L2,H1],
            [L3:merge(L2,H1),L1:drs([],[Cond])],
            [leq(L1,H1)]))@H@L1,P@Y@H@L2)))))):-
compose(Cond,Sym,[Y,X]).

modSem(neg,lambda(P,lambda(X,lambda(H,lambda(L,
    merge(usr([N,S],[N:drs([],[~S])],[leq(N,H),leq(L,S)]),
    P@X@H@L)))))).

coordSem(conj,lambda(X,lambda(Y,lambda(P,lambda(H,lambda(L,
    merge(usr([L1,L2],[L:merge(L1,L2)],[leq(L,H)]),
    merge(X@P@H@L1,Y@P@H@L2)))))))).

coordSem(disj,lambda(X,lambda(Y,lambda(P,lambda(H,lambda(L,
    merge(usr([L1,L2],[L:drs([],[L1 v L2])],[leq(L,H)]),
    merge(X@P@H@L1,Y@P@H@L2)))))))).

dcoordSem(cond,lambda(C1,lambda(C2,lambda(H,lambda(L,
    merge(usr([H1,H2],[L:drs([],[H1 > H2])],[leq(L,H)]),
    merge(C1@H1@_,C2@H2@_)))))).

dcoordSem(conj,lambda(C1,lambda(C2,lambda(H,lambda(L,
    merge(usr([H1,H2],[L:merge(H1,H2)],[leq(L,H)]),
    merge(C1@H1@_,C2@H2@_)))))).

dcoordSem(disj,lambda(C1,lambda(C2,lambda(H,lambda(L,
    merge(usr([H1,H2],[L:drs([],[H1 v H2])],[leq(L,H)]),
    merge(C1@H1@_,C2@H2@_)))))).

```

```

/*****

    name: mainPronounsDRT.pl (Chapter 9)
    version: Feb 3, 1999
    description: Pronoun Resolution
    authors: Patrick Blackburn & Johan Bos

*****/

:- module(mainPronounsDRT, [parse/0, resolveDrs/1]).

:- use_module(readLine, [readLine/1]),
   use_module(comsemOperators),
   use_module(comsemPredicates, [printRepresentation/1, simpleTerms/1,
                                  compose/3, append/3, member/2]),
   use_module(betaConversion, [betaConvert/2]),
   use_module(bindingDRT, [potentialAntecedent/3, properBinding/3]).

:- [englishGrammar], [englishLexicon], [semMacrosLambdaDRT].

/*=====
   Driver Predicate
=====*/

parse:-
  readLine(Discourse),
  d(MergeDrs, Discourse, []),
  betaConvert(MergeDrs, Drs),
  resolveDrs([Drs]-[ResolvedDrs]),
  printRepresentation(ResolvedDrs).

/*=====
   Pronoun Resolution
=====*/

resolveDrs([merge(B1, B2) | A1]-[drs(D, C) | A3]):-
  resolveDrs([B1 | A1]-A2),
  resolveDrs([B2 | A2]-[drs(D1, C1), drs(D2, C2) | A3]),
  append(D1, D2, D),
  append(C1, C2, C).

resolveDrs([alfa(Referent, Type, Gender, B1) | A1]-A2):-
  potentialAntecedent(A1, Referent, Gender),
  properBinding(Type, Referent, B1),
  resolveDrs([B1 | A1]-A2).

resolveDrs([drs(D1, C1) | A1]-A2):-
  resolveConds(C1, [drs(D1, []) | A1]-A2).

resolveConds([~B1 | Conds], A1-A3):-

```

```
    resolveDrs ([B1 | A1] - [B2, drs(D, C) | A2]),
    resolveConds (Conds, [drs(D, [~B2 | C]) | A2] - A3).

resolveConds ([B1 > B2 | Conds], A1-A4) :-
    resolveDrs ([B1 | A1] - A2),
    resolveDrs ([B2 | A2] - [B4, B3, drs(D, C) | A3]),
    resolveConds (Conds, [drs(D, [B3 > B4 | C]) | A3] - A4).

resolveConds ([B1 v B2 | Conds], A1-A4) :-
    resolveDrs ([B1 | A1] - [B3 | A2]),
    resolveDrs ([B2 | A2] - [B4, drs(D, C) | A3]),
    resolveConds (Conds, [drs(D, [B3 v B4 | C]) | A3] - A4).

resolveConds ([Basic | Conds], [drs(D, C) | A1] - A2) :-
    compose (Basic, _Symbol, Arguments),
    simpleTerms (Arguments),
    resolveConds (Conds, [drs(D, [Basic | C]) | A1] - A2).

resolveConds ([], A-A).
```

```

/*****

    name: bindingDRT.pl (Chapter 9)
    version: Feb 3, 1999
    description: Check Binding Constraints
    authors: Patrick Blackburn & Johan Bos

*****/

:- module(bindingDRT,[potentialAntecedent/3,
                    properBinding/3]).

:- use_module(comsemOperators),
   use_module(semOntology,[consistent/2]),
   use_module(comsemPredicates,[compose/3,member/2]).

/*=====
   Potential Antecedent (Ordinary DRSs)
   =====*/

potentialAntecedent(A,X,Gender):-
  member(drs(Dom,Conds),A),
  member(X,Dom),
  compose(Gender,Symbol1,_),
  \+ (
    member(Cond,Conds),
    compose(Cond,Symbol2,[Y]),
    Y==X,
    \+ consistent(Symbol1,Symbol2)
  ).

/*=====
   Potential Antecedent (Focus DRSs)
   =====*/

potentialAntecedent(A,X,Gender):-
  member(drs(Dom,_,_,Conds),A),
  member(X,Dom),
  compose(Gender,Symbol1,_),
  \+ (
    member(Cond,Conds),
    compose(Cond,Symbol2,[Y]),
    Y==X,
    \+ consistent(Symbol1,Symbol2)
  ).

/*=====
   Check Binding violation.
   =====

```



```
=====*/

properBinding(Type,X,Drs):-
    Type=refl,
    reflexiveBinding(X,Drs).

properBinding(Type,X,Drs):-
    \+ Type=refl,
    (
        reflexiveBinding(X,Drs),
        !, fail
    ;
        true
    ).

reflexiveBinding(_,[ ]):- fail.
reflexiveBinding(_,_alfa(_,_,_,_)):- fail.
reflexiveBinding(_,_merge(_,_)):- fail.
reflexiveBinding(_,_~_):- !, fail.
reflexiveBinding(X,drs(_,_Conds)):-
    reflexiveBinding(X,Conds).

reflexiveBinding(X,[Basic|Conds]):- !,
    (
        compose(Basic,_Sym,[Subj,Obj]),
        Subj==Obj,
        X==Obj, !
    ;
        reflexiveBinding(X,Conds)
    ).
```

```

/*****

    name: mainFocusDRT.pl (Chapter 9)
    version: Feb 3, 1999
    description: Pronoun Resolution with the Focusing Algorithm
    authors: Patrick Blackburn & Johan Bos

*****/

:- module(mainFocusDRT, [parse/0, resolveDrs/2]).

:- use_module(readLine, [readLine/1]),
   use_module(comsemOperators),
   use_module(comsemPredicates, [printRepresentation/1, simpleTerms/1,
                                  compose/3, append/3, member/2]),
   use_module(betaConversion, [betaConvert/2]),
   use_module(bindingDRT, [potentialAntecedent/3, properBinding/3]).

:- [englishGrammar], [englishLexicon], [semMacrosLambdaDRT].

/*=====
   Driver Predicate
   =====*/

parse:-
    parse(drs([], [], [], [])).

parse(Old):-
    readLine(Discourse),
    d(SDrs, Discourse, []),
    betaConvert(SDrs, Drs),
    resolveDrs([merge(Old, Drs)]-A, []-Anaphora),
    updateFocus(Anaphora, A, [New]),
    printRepresentation(New), !,
    parse(New).

/*=====
   Pronoun Resolution: DRSS
   =====*/

resolveDrs([merge(B1, B2) | A1]-[drs(D, AF, DF, C) | A5], R1-R3):-
    resolveDrs([B1 | A1]-A2, R1-R2),
    updateFocus(R2, A2, A3),
    resolveDrs([B2 | A3]-A4, R2-R3),
    updateFocus(R3, A4, [drs(D1, AF1, DF1, C1), drs(D2, AF2, DF2, C2) | A5]),
    append(D1, D2, D),
    append(AF1, AF2, AF),
    append(DF1, DF2, DF),
    append(C1, C2, C).

```

```
resolveDrs([drs(D,AF,DF,C)|A1]-A3,R1-R2):-
  resolveConds(C,[drs(D,AF,DF,[])|A1]-A2,R1-R2),
  updateFocus(R2,A2,A3).
```

```
resolveDrs([drs(D,C)|A1]-A3,R1-R2):-
  resolveConds(C,[drs(D,[],[],[])|A1]-A2,R1-R2),
  updateFocus(R2,A2,A3).
```

```
/*=====
  Pronoun Resolution: if pronoun appears as agent, then choose the
  actor focus as preferred antecedent
=====*/
```

```
resolveDrs([alfa(X,Type,Gender,B1)|A1]-A2,R1-[Type:X|R2]):-
  role(agent,X,B1), !,
  (
    currentActorFocus(A1,AF),
    potentialAntecedent(A1,X,Gender),
    X == AF,
    properBinding(Type,X,B1)
  ;
    currentDiscourseFocus(A1,DF),
    potentialAntecedent(A1,X,Gender),
    X == DF,
    properBinding(Type,X,B1)
  ;
    potentialAntecedent(A1,X,Gender),
    properBinding(Type,X,B1)
  ),
  resolveDrs([B1|A1]-A2,R1-R2).
```

```
/*=====
  Pronoun Resolution: if pronoun appears as non-agent, then choose
  the discourse focus as preferred antecedent
=====*/
```

```
resolveDrs([alfa(X,Type,Gender,B1)|A1]-A2,R1-[Type:X|R2]):-
  (
    currentDiscourseFocus(A1,DF),
    potentialAntecedent(A1,X,Gender),
    X == DF,
    properBinding(Type,X,B1)
  ;
    currentActorFocus(A1,AF),
    potentialAntecedent(A1,X,Gender),
    X == AF,
    properBinding(Type,X,B1)
  ;
    potentialAntecedent(A1,X,Gender),
```

```

    properBinding(Type,X,B1)
  ),
  resolveDrs([B1|A1]-A2,R1-R2).

/*=====
   Pronoun Resolution: DRS-conditions
=====*/

resolveConds([~B1|Conds],A1-A3,R1-R3):-
  resolveDrs([B1|A1]-[B2,drs(D,AF,DF,C)|A2],R1-R2),
  resolveConds(Conds,[drs(D,AF,DF,[~B2|C])|A2]-A3,R2-R3).

resolveConds([B1 > B2|Conds],A1-A5,R1-R4):-
  resolveDrs([B1|A1]-A2,R1-R2),
  updateFocus(R2,A2,A3),
  resolveDrs([B2|A3]-[B4,B3,drs(D,AF,DF,C)|A4],R2-R3),
  resolveConds(Conds,[drs(D,AF,DF,[B3 > B4|C])|A4]-A5,R3-R4).

resolveConds([B1 v B2|Conds],A1-A4,R1-R4):-
  resolveDrs([B1|A1]-[B3|A2],R1-R2),
  resolveDrs([B2|A2]-[B4,drs(D,AF,DF,C)|A3],R2-R3),
  resolveConds(Conds,[drs(D,AF,DF,[B3 v B4|C])|A3]-A4,R3-R4).

resolveConds([Basic|Conds],[drs(D,AF,DF,C)|A1]-A2,R1-R2):-
  compose(Basic,_Symbol,Arguments),
  simpleTerms(Arguments),
  resolveConds(Conds,[drs(D,AF,DF,[Basic|C])|A1]-A2,R1-R2).

resolveConds([],A-A,R-R).

/*=====
   Focus Update
=====*/

updateFocus(Anaphora,DRSs1,Updated):-
  updateActorFocus(DRSs1,DRSs2),
  updateDiscourseFocus(Anaphora,DRSs2,Updated).

/*=====
   Update Actor Focus
=====*/

updateActorFocus([drs(D,AF,DF,C)|A],Updated):-
  expectedActorFocus(drs(D,C),Focus), !,
  addActorFocus(Focus,[drs(D,AF,DF,C)|A],Updated).

updateActorFocus(DRSs,DRSs).

```

```
/*=====
Update Discourse Focus
=====*/
```

```
updateDiscourseFocus(Anaphora,DRSs,Updated):-
  currentDiscourseFocus(DRSs,Current),
  member(_:A,Anaphora), A==Current, !,
  Updated=DRSs.
```

```
updateDiscourseFocus(_Anaphora,[drs(D,AF,DF,C)|A],Updated):-
  expectedDiscourseFocus(drs(D,C),Focus), !,
  addDiscourseFocus(Focus,[drs(D,AF,DF,C)|A],Updated).
```

```
updateDiscourseFocus(_,DRSs,DRSs).
```

```
/*=====
Add new focus to DRS
=====*/
```

```
addActorFocus(Focus,DRSs,Updated):-
  (
    currentActorFocus(DRSs,Current),
    Current==Focus, !,
    Updated=DRSs
  ;
    DRSs=[drs(D,AF,DF,C)|A],
    Updated=[drs(D,[Focus|AF],DF,C)|A]
  ).
```

```
addDiscourseFocus(Focus,DRSs,Updated):-
  (
    currentDiscourseFocus(DRSs,Current),
    Current==Focus, !,
    Updated=DRSs
  ;
    DRSs=[drs(D,AF,DF,C)|A],
    Updated=[drs(D,AF,[Focus|DF],C)|A]
  ).
```

```
/*=====
Current Foci
=====*/
```

```
currentDiscourseFocus(DRSs,Focus):-
  member(drs(_,[Focus|_],_),DRSs), !.
```

```
currentActorFocus(DRSs,Focus):-
```

```

    member(drs( _, [Focus | _], _, _ ), DRSs), !.

/*=====
    Expected Focus
=====*/

expectedActorFocus(drs(Dom, Conds), Focus):-
    member(Focus, Dom),
    role(agent, Focus, Conds).

expectedDiscourseFocus(drs(Dom, Conds), Focus):-
    member(Focus, Dom),
    role(theme, Focus, Conds).

/*=====
    Thematic Roles
=====*/

role(Role, X, drs( _, _, _, Conds )):-
    role(Role, X, Conds).

role(Role, X, drs( _, Conds )):-
    role(Role, X, Conds).

role(Role, X, merge(B1, B2)):-
    role(Role, X, B1);
    role(Role, X, B2).

role(Role, X, alfa( _, _, _, B )):-
    role(Role, X, B).

role(Role, X, [ ~ B | _ ]):-
    role(Role, X, B).

role(Role, X, [ B1 > B2 | _ ]):-
    role(Role, X, B1);
    role(Role, X, B2).

role(Role, X, [ B1 v B2 | _ ]):-
    role(Role, X, B1);
    role(Role, X, B2).

role(Role, X, [ _ | Conds ]):-
    role(Role, X, Conds).

role(agent, X, [ C | _ ]):-
    agent(X, C).

role(theme, X, [ C | _ ]):-

```

```
theme(X,C).

/*=====
   Agent Position
=====*/

agent(X,C):-
  compose(C,Sym,[Y,_]),
  X==Y,
  lexicon(tv,Sym,_,inf).

/*=====
   Theme Position
=====*/

theme(X,C):-
  compose(C,Sym,[Y]),
  X==Y,
  lexicon(iv,Sym,_,inf).

theme(X,C):-
  compose(C,Sym,[_ ,Y]),
  X==Y,
  lexicon(tv,Sym,_,inf).

theme(X,C):-
  compose(C,'=' , [Y,Z]),
  var(Z),
  X==Y.
```

```

/*****

    name: mainPresupDRT.pl
    version: May 5, 1998
    description: Presupposition Projection
    authors: Patrick Blackburn & Johan Bos

*****/

:- module(mainPresupDRT, [parse/0, projectDrs/1]).

:- use_module(readLine, [readLine/1]),
   use_module(comsemOperators),
   use_module(comsemPredicates, [printReadings/1, compose/3]),
   use_module(betaConversion, [betaConvert/2]),
   use_module(resolvePresup, [projectDrs/1, accommodate/2]).

:- [englishGrammar], [englishLexicon], [semMacrosPresupDRT].

/*=====
   Driver Predicate
   =====*/

parse:-
  readLine(Discourse),
  d(Drs1, Discourse, []),
  betaConvert(Drs1, Drs2),
  findall(Drs4, (
    projectDrs([Drs2, pre([])]-[Drs3, pre(P)]),
    accommodate(P, Drs3-Drs4)
  ),
  Readings),
  printReadings(Readings).

```



```

/*****

    name: resolvePresup.pl
    version: May 5, 1998
    description: Presupposition Resolution
    authors: Patrick Blackburn & Johan Bos

*****/

:- module(resolvePresup, [projectDrs/1, accommodate/2]).

:- use_module(comsemOperators),
   use_module(comsemPredicates, [simpleTerms/1, compose/3,
                                  append/3, select/3]),
   use_module(bindingDRT, [properBinding/3]),
   use_module(matchDRT, [matchDrs/4]).

/*=====
   Presupposition Projection for DRSs
   =====*/

projectDrs([merge(B1, B2) | A1] - [drs(D3, C3) | A3]) :-
    projectDrs([B1 | A1] - A2),
    projectDrs([B2 | A2] - [drs(D1, C1), drs(D2, C2) | A3]),
    append(D1, D2, D3),
    append(C1, C2, C3).

projectDrs([alfa(X, Type, B1, B3) | A1] - A4) :-
    projectDrs([B1 | A1] - [B2 | A2]),
    resolveAlfa(X, B2, A2 - A3),
    properBinding(Type, X, B3),
    projectDrs([B3 | A3] - A4).

projectDrs([drs(D1, C1) | A1] - A2) :-
    projectConds(C1, [drs(D1, []) | A1] - A2).

/*=====
   Presupposition Projection for DRS-Conditions
   =====*/

projectConds([~B1 | Conds], A1 - A3) :-
    projectDrs([B1, pre([]) | A1] - [B2, pre(Pres), drs(D, C) | A2]),
    accommodate(Pres, B2 - B3),
    projectConds(Conds, [drs(D, [~B3 | C]) | A2] - A3).

projectConds([B1 > B2 | Conds], A1 - A4) :-
    projectDrs([B1, pre([]) | A1] - A2),
    projectDrs([B2, pre([]) | A2] - [B4, pre(P2), B3, pre(P1), drs(D, C) | A3]),
    accommodate(P1, B3 - B5),
    accommodate(P2, B4 - B6),

```

```

    projectConds(Conds, [drs(D, [B5 > B6|C])|A3]-A4).

projectConds([B1 v B2|Conds], A1-A4):-
    projectDrs([B1,pre([])|A1]-[B3,pre(P1)|A2]),
    projectDrs([B2,pre([])|A2]-[B4,pre(P2),drs(D,C)|A3]),
    accommodate(P1,B3-B5),
    accommodate(P2,B4-B6),
    projectConds(Conds, [drs(D, [B5 v B6|C])|A3]-A4).

projectConds([Basic|Conds], [drs(D,C)|A1]-A2):-
    compose(Basic,_Symbol,Arguments),
    simpleTerms(Arguments),
    projectConds(Conds, [drs(D, [Basic|C])|A1]-A2).

projectConds([],A-A).

/*=====
   Resolving Alfa-DRSs
=====*/

resolveAlfa(X,AlfaDrs, [drs(D,C)|Others]-[New|Others]):-
    matchDrs(X,AlfaDrs,drs(D,C),New).

resolveAlfa(_,AlfaDrs, [pre(A)|Others]-[pre([AlfaDrs|A])|Others]).

resolveAlfa(X,Alfa, [pre(A1)|Others]-[pre([New|A2])|Others]):-
    select(Drs,A1,A2),
    matchDrs(X,Alfa,Drs,New).

resolveAlfa(X,AlfaDrs, [AnteDrs|Others]-[AnteDrs|NewOthers]):-
    resolveAlfa(X,AlfaDrs,Others-NewOthers).

/*=====
   Accommodation
=====*/

accommodate([],B-B).

accommodate([drs(D1,C1)|Presups],drs(D2,C2)-B):-
    append(D1,D2,D3),
    append(C1,C2,C3),
    accommodate(Presups,drs(D3,C3)-B).

```

```

/*****

    name: matchDRT.pl
    version: Feb 15, 1999
    description: Matching for DRSs, including consistency check
    authors: Patrick Blackburn & Johan Bos

*****/

:- module(matchDRT, [matchDrs/4]).

:- use_module(comsemPredicates, [member/2, compose/3]),
   use_module(semOntology, [consistent/2]).

/*=====
   Partial Match
   =====*/

matchDrs(X, drs(D1, C1), drs(D2, C2), drs(D3, C3)) :-
    member(X, D2),
    mergeLists(D1, D2, D3),
    mergeLists(C1, C2, C3),
    consistentConditions(X, C3).

consistentConditions(X, Conds) :-
    \+ (
        member(Cond1, Conds),
        member(Cond2, Conds), \+ Cond1=Cond2,
        compose(Cond1, Symbol1, [Y]), Y==X,
        compose(Cond2, Symbol2, [Z]), Z==X,
        \+ consistent(Symbol1, Symbol2)
    ).

/*=====
   Merging of Lists
   =====*/

mergeLists([], L, L).

mergeLists([X|R], L1, L2) :-
    member(Y, L1),
    X==Y, !,
    mergeLists(R, L1, L2).

mergeLists([X|R], L1, [X|L2]) :-
    mergeLists(R, L1, L2).

```

```

/*****

    name: mainPresupScoreDRT.pl
    version: Feb 6, 1999
    description: Presupposition Resolution with Score Calculation
    authors: Patrick Blackburn & Johan Bos

*****/

:- module(mainPresupScoreDRT, [parse/0, projectDrs/1]).

:- use_module(readLine, [readLine/1]),
   use_module(comsemOperators),
   use_module(comsemPredicates, [printReadings/1, compose/3]),
   use_module(betaConversion, [betaConvert/2]),
   use_module(resolvePresupScore, [projectDrs/2, accommodate/2]).

:- [englishGrammar], [englishLexicon], [semMacrosPresupDRT].

/*=====
   Driver Predicate
   =====*/

parse:-
  readLine(Discourse),
  d(Drs1, Discourse, []),
  betaConvert(Drs1, Drs2),
  findall((Score:Drs4), (
    projectDrs([Drs2, pre([])]-[Drs3, pre(P)], 1-Score),
    accommodate(P, Drs3-Drs4)
  )),
  Readings),
  printReadings(Readings).

```

```

/*****

    name: resolvePresupScore.pl
    version: Feb 6, 1999
    description: Presupposition Projection with Score Calculation
    authors: Patrick Blackburn & Johan Bos

*****/

:- module(resolvePresupScore, [projectDrs/2, accommodate/2]).

:- use_module(comsemOperators),
   use_module(comsemPredicates, [simpleTerms/1, compose/3,
                                   append/3, select/3]),
   use_module(bindingDRT, [properBinding/3]),
   use_module(matchDRT, [matchDrs/4]).

/*=====
   Presupposition Projection for DRSs
=====*/

projectDrs([merge(B1, B2) | A1] - [drs(D3, C3) | A3], S1-S3) :-
    projectDrs([B1 | A1] - A2, S1-S2),
    projectDrs([B2 | A2] - [drs(D1, C1), drs(D2, C2) | A3], S2-S3),
    append(D1, D2, D3),
    append(C1, C2, C3).

projectDrs([alfa(X, Type, B1, B3) | A1] - A4, S1-S4) :-
    projectDrs([B1 | A1] - [B2 | A2], S1-S2),
    resolveAlfa(X, Type, B2, A2-A3, S2-S3),
    properBinding(Type, X, B3),
    projectDrs([B3 | A3] - A4, S3-S4).

projectDrs([drs(D1, C1) | A1] - A2, S1-S2) :-
    projectConds(C1, [drs(D1, []) | A1] - A2, S1-S2).

/*=====
   Presupposition Projection for DRS-Conditions
=====*/

projectConds([~B1 | Conds], A1-A3, S1-S3) :-
    projectDrs([B1, pre([]) | A1] - [B2, pre(Pres), drs(D, C) | A2], S1-S2),
    accommodate(Pres, B2-B3),
    projectConds(Conds, [drs(D, [~B3 | C]) | A2] - A3, S2-S3).

projectConds([B1 > B2 | Conds], A1-A4, S1-S4) :-
    projectDrs([B1, pre([]) | A1] - A2, S1-S2),
    projectDrs([B2, pre([]) | A2] - [B4, pre(P2), B3, pre(P1), drs(D, C) | A3], S2-S3),

```

```

    accommodate(P1,B3-B5),
    accommodate(P2,B4-B6),
    projectConds(Conds,[drs(D,[B5 > B6|C])|A3]-A4,S3-S4).

projectConds([B1 v B2|Conds],A1-A4,S1-S4):-
    projectDrs([B1,pre([])|A1]-[B3,pre(P1)|A2],S1-S2),
    projectDrs([B2,pre([])|A2]-[B4,pre(P2),drs(D,C)|A3],S2-S3),
    accommodate(P1,B3-B5),
    accommodate(P2,B4-B6),
    projectConds(Conds,[drs(D,[B5 v B6|C])|A3]-A4,S3-S4).

projectConds([Basic|Conds],[drs(D,C)|A1]-A2,S1-S2):-
    compose(Basic,_Symbol,Arguments),
    simpleTerms(Arguments),
    projectConds(Conds,[drs(D,[Basic|C])|A1]-A2,S1-S2).

projectConds([],A-A,S-S).

/*=====
   Resolving Alfa-DRSs
   =====*/

resolveAlfa(X,Type,AlfaDrs,[drs(D,C)|Others]-[New|Others],S1-S2):-
    global(Others),
    matchDrs(X,AlfaDrs,drs(D,C),New),
    (Type=refl, S2 = S1;
     Type=nonrefl, S2 = S1;
     Type=def, S2 = S1;
     Type=name, S2 = S1).

resolveAlfa(X,Type,AlfaDrs,[drs(D,C)|Others]-[New|Others],S1-S2):-
    nonglobal(Others),
    matchDrs(X,AlfaDrs,drs(D,C),New),
    (Type=refl, S2 is S1 * 1;
     Type=nonrefl, S2 is S1 * 1;
     Type=def, S2 is S1 * 0.5;
     Type=name, S2 is S1 * 0.2).

resolveAlfa(_,Type,Alfa,[pre(A)]-[pre([Alfa|A])],S1-S2):-
    global([pre(A)]),
    (Type=nonrefl, S2 is S1 * 0.5;
     Type=def, S2 is S1 * 0.9;
     Type=name, S2 is S1 * 0.9).

resolveAlfa(_,Type,Alfa,[pre(A)|Others]-[pre([Alfa|A])|Others],S1-S2):-
    nonglobal([pre(A)|Others]),
    (
    %Type=nonrefl, S2 is S1 * 0.1;
     Type=def, S2 is S1 * 0.7;

```

```

    Type=name, S2 is S1 * 0.2).

resolveAlfa(X,Type,Alfa,[pre(A1)|Others]-[pre([New|A2])|Others],S1-S2):-
    global([pre(A1)|Others]),
    select(Drs,A1,A2),
    matchDrs(X,Alfa,Drs,New),
    (Type=refl, S2 = S1;
     Type=nonrefl, S2 = S1;
     Type=def, S2 = S1;
     Type=name, S2 = S1).

resolveAlfa(X,Type,Alfa,[pre(A1)|Others]-[pre([New|A2])|Others],S1-S2):-
    nonglobal([pre(A1)|Others]),
    select(Drs,A1,A2),
    matchDrs(X,Alfa,Drs,New),
    (Type=refl, S2 is S1 * 1;
     Type=nonrefl, S2 is S1 * 1;
     Type=def, S2 is S1 * 0.5;
     Type=name, S2 is S1 * 0.2).

resolveAlfa(X,Type,AlfaDrs,[AnteDrs|Others]-[AnteDrs|NewOthers],S1-S2):-
    resolveAlfa(X,Type,AlfaDrs,Others-NewOthers,S1-S2).

/*=====
   Global or non-global position with respect to stack
   =====*/

global([pre(_)]).
global([drs(_,_)|Stack]):- global(Stack).

nonglobal([pre(_),drs(_,_)|_]).
nonglobal([drs(_,_)|Stack]):- nonglobal(Stack).

/*=====
   Accommodation
   =====*/

accommodate([],B-B).

accommodate([drs(D1,C1)|Presups],drs(D2,C2)-B):-
    append(D1,D2,D3),
    append(C1,C2,C3),
    accommodate(Presups,drs(D3,C3)-B).

```

```

/*****
    name: semMacrosPresupDRT.pl
    version: May 5, 1998
    description: Semantic Macros for Presupposition Projection in DRT
    authors: Patrick Blackburn & Johan Bos
*****/

/*=====
   Semantic Macros
=====*/

detSem(uni,lambda(P,lambda(Q,drs([], [merge(drs([X], []),P@X)>(Q@X)])))).
detSem(indef,lambda(P,lambda(Q,merge(merge(drs([X], []),P@X),Q@X))).
detSem(def,lambda(P,lambda(Q,alfa(X,def,merge(drs([X], []),P@X),Q@X))).
detSem(card(1),lambda(P,lambda(Q,
    merge(drs([X], [merge(drs([Y], []),merge(P@Y,Q@Y)]>drs([], [X=Y])),
    merge(P@X,Q@X)])))).
detSem(poss(Gender),lambda(P,lambda(Q,
    alfa(Y,def,alfa(X,nonrefl,drs([X], [Basic]),
    merge(drs([Y], [of(Y,X)],P@Y),Q@Y))):-
    compose(Basic,Gender,[X])).
nounSem(Sym,lambda(X,drs([], [Cond]))):-
    compose(Cond,Sym,[X]).
pnSem(Sym,Gender,lambda(P,alfa(X,name,drs([X], [X=Sym,Cond]),P@X))):-
    compose(Cond,Gender,[X]).
proSem(Gender,Type,lambda(P,alfa(X,Type,drs([X], [Cond]),P@X))):-
    compose(Cond,Gender,[X]).
ivSem(Sym,lambda(X,drs([], [Cond]))):-
    compose(Cond,Sym,[X]).
tvSem(Sym,lambda(K,lambda(Y,K @ lambda(X,drs([], [Cond]))))):-
    compose(Cond,Sym,[Y,X]).
relproSem(lambda(P,lambda(Q,lambda(X,merge(P@X,Q@X)))).
prepSem(Sym,lambda(K,lambda(P,lambda(Y,Drs))):-
    Drs=merge(K@lambda(X,drs([], [Cond])),P@Y),
    compose(Cond,Sym,[Y,X]).
modSem(neg,lambda(P,lambda(X,drs([], [~(P @ X)])))).

```



```
coordSem(conj, lambda(X, lambda(Y, lambda(P, merge(X@P, Y@P))))).  
coordSem(disj, lambda(X, lambda(Y, lambda(P, drs([], [(X@P) v (Y@P)]))))) .  
  
dcoordSem(cond, lambda(X, lambda(Y, drs([], [X > Y])))).  
dcoordSem(conj, lambda(X, lambda(Y, merge(X, Y)))).  
dcoordSem(disj, lambda(X, lambda(Y, drs([], [X v Y])))).
```

```

/*****

    name: curt.pl (Chapter 11)
    version: June 18, 1998
    description: User-System discourse interaction
    authors: Patrick Blackburn & Johan Bos

*****/

:- module(curt, [curt/0]).

:- use_module(readLine, [readLine/1]),
   use_module(comsemOperators),
   use_module(comsemPredicates, [printReadings/1, member/2, compose/3]),
   use_module(betaConversion, [betaConvert/2]),
   use_module(resolvePresup, [projectDrs/1, accommodate/2]),
   use_module(acceptabilityConstraints, [consistentReadings/2,
                                         informativeReadings/2,
                                         localConsistentReadings/2,
                                         localInformativeReadings/2]).

:- [englishGrammar], [englishLexicon], [semMacrosPresupDRT].

/*=====
   Start Curt
   =====*/

curt:-
    curtInput([drs([], [])]).

/*=====
   System requests user
   =====*/

curtInput(Readings):-
    readLine(Input),
    curtOutput(Input, Readings).

/*=====
   Systems responds to user
   =====*/

curtOutput([], Readings):- !,
    curtSays('Want to tell me something?'),
    curtInput(Readings).

curtOutput([bye], _):- !,
    curtSays('Bye bye!').

```

```

curtOutput([new],_):- !,
    curt.

curtOutput([help],Readings):- !,
    nl, write('bye - no more talking'),
    nl, write('drs - prints current readings in DRS-format'),
    nl, write('new - starts a new discourse'), nl,
    curtInput(Readings).

curtOutput([drs],Readings):- !,
    printReadings(Readings),
    curtInput(Readings).

curtOutput(Input,Readings):-
    d(MergeDrs,Input,[]),!,
    betaConvert(MergeDrs,ReducedDrs),
    getReadings(ReducedDrs,Readings,PotentialReadings),
    consistentReadings(PotentialReadings,ConsistentReadings),
    (
        ConsistentReadings=[],
        curtSays('No! I do not believe that!'),!,
        curtInput(Readings)
    );
    informativeReadings(ConsistentReadings,InformativeReadings),
    (
        InformativeReadings=[],
        curtSays('Yes, I knew that!'),!,
        curtInput(Readings)
    );
    curtSays('Ok. '),
    localInformativeReadings(InformativeReadings,LocalInformative),
    localConsistentReadings(LocalInformative,LocalConsistent),
    (
        LocalConsistent=[], !,
        SelectedReadings=InformativeReadings
    );
    SelectedReadings=LocalConsistent
    ),
    curtInput(SelectedReadings)
    ).

curtOutput(_,Readings):-
    curtSays('What?'),
    curtInput(Readings).

/*=====
    Compute all readings
=====*/

```

```
getReadings(ProtoDrs,Readings,PotentialReadings):-
  findall((OldDrs,NewDrs),
    (
      member(OldDrs,Readings),
      projectDrs([merge(OldDrs,ProtoDrs),pre([])]-[Drs,pre(P)]),
      accommodate(P,Drs-NewDrs)
    ),
    PotentialReadings).

/*=====
  Curt's output
=====*/

curtSays(X):- nl, write(X), nl.
```

```

/*****

    name: acceptabilityConstraints.pl (Chapter 11)
    version: June 18, 1998
    description: Consistency and Informativity
    authors: Patrick Blackburn & Johan Bos

*****/

:- module(acceptabilityConstraints, [consistentReadings/2,
                                   informativeReadings/2,
                                   localConsistentReadings/2,
                                   localInformativeReadings/2]).

:- use_module(drs2fol, [drs2fol/2]),
   use_module(comsemOperators),
   use_module(comsemPredicates, [append/3]),
   use_module(mergeDRT, [mergeDrs/2]),
   use_module(callTheoremProver, [callTheoremProver/3,
                                  callModelBuilder/4]),
   use_module(semOntology, [generateOntology/1]).

/*=====
   Select Readings
=====*/

informativeReadings([], []).
informativeReadings([(DrsSoFar, Drs) | Potential], [Drs | Selected]):-
    informative(DrsSoFar, Drs), !,
    informativeReadings(Potential, Selected).
informativeReadings([_ | Potential], Selected):-
    informativeReadings(Potential, Selected).

consistentReadings([], []).
consistentReadings([(Old, Drs) | Potential], [(Old, Drs) | Selected]):-
    consistent(Drs), !,
    consistentReadings(Potential, Selected).
consistentReadings([_ | Potential], Selected):-
    consistentReadings(Potential, Selected).

localConsistentReadings([], []).
localConsistentReadings([Drs | Potential], [Drs | Selected]):-
    localConsistent(Drs), !,
    localConsistentReadings(Potential, Selected).
localConsistentReadings([_ | Potential], Selected):-
    localConsistentReadings(Potential, Selected).

localInformativeReadings([], []).
localInformativeReadings([Drs | Potential], [Drs | Selected]):-
    localInformative(Drs), !,

```

```

    localInformativeReadings(Potential,Selected).
localInformativeReadings([_|Potential],Selected):-
    localInformativeReadings(Potential,Selected).

/*=====
    Informativity
=====*/

informative(drs([],[]),_):-!.
informative(OldDrs,NewDrs):-
    backgroundKnowledge(Chi),
    drs2fol(OldDrs,Phi),
    drs2fol(NewDrs,Psi),
    callTheoremProver(Chi,~(Phi > Psi),Proof),
    (Proof=yes, !, fail; true).

/*=====
    Consistency
=====*/

consistent(NewDrs):-
    domainSize(NewDrs,Size),
    backgroundKnowledge(Chi),
    drs2fol(NewDrs,Phi),
    (
        callModelBuilder(Chi,Phi,Size,Model),
        Model=1,!
    ;
        callTheoremProver(Chi,Phi,Proof),
        (
            Proof=yes, !, fail
        ;
            true
        )
    ).

/*=====
    Local Informativity
=====*/

localInformative(Drs):-
    findall((Super,Sub),superSubDrs(Drs,drs([],[])-Super,Sub),List),
    allLocalInformative(List).

allLocalInformative([]).
allLocalInformative([(Super,Sub)|Others]):-
    backgroundKnowledge(Chi),
    drs2fol(drs([],[Super>Sub]),Phi),
    callTheoremProver(Chi,~Phi,Proof),
    (Proof=yes, !, fail; allLocalInformative(Others)).

```

```

/*=====
  Local Consistency
=====*/

localConsistent(Drs):-
  findall((Super,Sub),superSubDrs(Drs,drs([],[])-Super,Sub),List),
  allLocalConsistent(List).

allLocalConsistent([]).
allLocalConsistent([(Super,Sub)|Others]):-
  backgroundKnowledge(Chi),
  drs2fol(drs([], [Super>drs([], [~Sub])]),Phi),
  callTheoremProver(Chi,~Phi,Proof),
  (Proof=yes,!,fail;allLocalConsistent(Others)).

/*=====
  Background Knowledge
=====*/

backgroundKnowledge(Formulas):-
  knowledge(Formulas1),
  generateOntology(Formulas2),
  append(Formulas1,Formulas2,Formulas).

knowledge([
  forall(X,forall(Y,have(X,Y)>of(Y,X))),
  forall(X,forall(Y,of(Y,X)>have(X,Y))),
  forall(X,female(X)&married(X)>exists(Y,husband(Y)&have(X,Y))),
  forall(X,forall(Y,husband(Y)&have(X,Y)>married(X)&female(X)))
]).

/*=====
  Domain Size
=====*/

domainSize(drs(Dom,_),Size):-
  (
    Dom=[],
    Size=1,!
  ;
    length(Dom,Size)
  ).

/*=====
  Compute super- and subordinated DRs
=====*/

superSubDrs(drs(D,[Sub > _|C]),Drs-Super,Sub):-
  mergeDrs(merge(Drs,drs(D,C)),Super).

```

```

superSubDrs (drs (D, [B > Sub|C]), Drs-Super, Sub) :-
    mergeDrs (merge (merge (drs (D, C), B), Drs), Super).
superSubDrs (drs (D, [B1 > B2|C]), Drs-Super, Sub) :-
    superSubDrs (B2, merge (Drs, merge (merge (drs (D, C), B1), B2))-Super, Sub).
superSubDrs (drs (D, [Sub v _|C]), Drs-Super, Sub) :-
    mergeDrs (merge (Drs, drs (D, C)), Super).
superSubDrs (drs (D, [_ v Sub|C]), Drs-Super, Sub) :-
    mergeDrs (merge (Drs, drs (D, C)), Super).
superSubDrs (drs (D, [B v _|C]), Drs-Super, Sub) :-
    superSubDrs (B, merge (Drs, merge (drs (D, C), B))-Super, Sub).
superSubDrs (drs (D, [_ v B|C]), Drs-Super, Sub) :-
    superSubDrs (B, merge (Drs, merge (drs (D, C), B))-Super, Sub).
superSubDrs (drs (D, [~ Sub|C]), Drs-Super, Sub) :-
    mergeDrs (merge (Drs, drs (D, C)), Super).
superSubDrs (drs (D, [~ B|C]), Drs-Super, Sub) :-
    superSubDrs (B, merge (Drs, merge (drs (D, C), B))-Super, Sub).
superSubDrs (drs (D, [Cond|C]), Drs-Super, Sub) :-
    superSubDrs (drs ([], C), Drs-B, Sub),
    mergeDrs (merge (drs (D, [Cond]), B), Super).

```



```
/******  
    name: callTheoremProver.pl (Chapter 5)  
    version: June 18, 1998  
    description: Prolog Interface to Otter (Sicstus required)  
    authors: Patrick Blackburn & Johan Bos  
*****/  
  
:- module(callTheoremProver,[callTheoremProver/3]).  
  
:- use_module(library(system)),  
   use_module(fol2otter,[fol2otter/2]).  
  
/*=====  
   Calls to Theorem Prover (Otter)  
=====*/  
  
callTheoremProver(Axioms,Formula,Proof):-  
    fol2otter(Axioms,Formula),  
    shell('./otter < temp.in > temp.out 2> /dev/null',X),  
    (X=26368,Proof=yes,!;X=26624,Proof=no).
```

```

/*****

    name: callModelBuilder.pl (Chapter 5)
    version: September 3, 1999
    description: Prolog Interface to Mace (Sicstus required)
    authors: Patrick Blackburn & Johan Bos

*****/

:- module(callModelBuilder,[callModelBuider/4]).

:- use_module(library(system)),
   use_module(comsemPredicates,[append/3]),
   use_module(fol2otter,[fol2otter/2]).

/*=====
   Calls to Model Generator (Mace)

   Changed in MACE (generate.c):
   #define MAX_SYMBOLS 100 number of functors (was 50)
=====*/

callModelBuilder(Axioms,Formula,DomainSize,Model):-
    fol2otter(Axioms,Formula),
    name('./mace -n',C1),
    name(DomainSize,C2),
    append(C1,C2,C3),
    name(' -p -t2 -m1 < temp.in > temp.out 2> /dev/null',C4),
    append(C3,C4,C5),
    name(Shell,C5),
    shell(Shell,X),
    % this is not correct. It always returns 0!
    (X=0,Model=1,!;write(X),Model=0).

```

```

/*****

    name: fol2otter.pl (Chapter 11)
    version: June 18, 1998
    description: Translates a formula in otter syntax to standard output
    authors: Patrick Blackburn & Johan Bos

*****/

:- module(fol2otter, [fol2otter/2]).

:- use_module(comsemOperators).

/*=====
   Translates formula to otter syntax in file 'temp.in'
=====*/

fol2otter(Axioms,Formula):-
    tell('temp.in'),
    format('set(auto).~n~n',[]),
    format('clear(print_proofs).~n~n',[]),
    format('set(prolog_style_variables).~n~n',[]),
    format('formula_list(usable).~n~n~n',[]),
    printOtterList(Axioms),
    printOtterFormula(Formula),
    format('~nend_of_list.~n',[]),
    told.

/*=====
   Print a list of Otter formulas
=====*/

printOtterList([]).
printOtterList([X|L]):-
    printOtterFormula(X),
    printOtterList(L).

/*=====
   Print an Otter formula
=====*/

printOtterFormula(F):-
    \+ \+ (numbervars(F,0,_), printOtter(F,5)),
    format('~n',[]).

printOtter(exists(X,Formula),Tab):-
    write('(exists '),write(X),write(' '),!,
    printOtter(Formula,Tab),write(')').

printOtter(forall(X,Formula),Tab):-

```

```
write('(all '),write(X),write(' '),!,
printOtter(Formula,Tab),write(')')).

printOtter(Phi & Psi,Tab):-
write('('),!,
printOtter(Phi,Tab),
write(' & '), nl, tab(Tab),
NewTab is Tab + 5,
printOtter(Psi,NewTab), write(')').

printOtter(Phi v Psi,Tab):-
write('('),!,
printOtter(Phi,Tab), write(' | '),
printOtter(Psi,Tab), write(')').

printOtter(Phi <> Psi,Tab):-
write('('),!,
printOtter(Phi,Tab), write(' <-> '),
printOtter(Psi,Tab), write(')').

printOtter(Phi > Psi,Tab):-
write('('),!,
printOtter(Phi,Tab), write(' -> '),
printOtter(Psi,Tab), write(')').

printOtter(~ Phi,Tab):-
write('~('),!,
printOtter(Phi,Tab), write(')').

printOtter(Phi,_):-
write(Phi).
```

```

/*****

    name: mainPresupDRTU.pl (Chapter 12)
    version: May 28, 1998
    description: Combining DRTU with presupposition resolution
    author: Patrick Blackburn & Johan Bos

*****/

:- module(mainPresupDRTU, [parse/0]).

:- use_module(readLine, [readLine/1]),
   use_module(comsemOperators),
   use_module(comsemPredicates, [printRepresentation/1,
                                  printReadings/1, compose/3]),
   use_module(betaConversion, [betaConvert/2]),
   use_module(pluggingAlgorithm, [plugHole/4]),
   use_module(mergeDRT, [mergeDrs/2]),
   use_module(mergeUSR, [mergeUSR/2]),
   use_module(resolvePresup, [projectDrs/1, accommodate/2]).

:- [englishGrammar], [englishLexicon], [semMacrosPresupDRTU].

/*=====
   Driver Predicate
=====*/

parse:-
  readLine(Discourse),
  d(Sem, Discourse, []),
  betaConvert(merge(usr([Top, Main], [], []), Sem@Top@Main), Reduced),
  mergeUSR(Reduced, usr(D, L, C)),
  printRepresentation(usr(D, L, C)),
  findall(Drs,
    (
      plugHole(Top, L-[], C, []),
      projectDrs([Top, pre([])]- [PresupDrs, pre(P)]),
      accommodate(P, PresupDrs-Drs)
    ),
    Readings),
  printReadings(Readings).

```

```

/*****

    name: semMacrosPresupDRTU.pl (Chapter 12)
    version: June 18, 1999
    description: Semantic Macros for DRTU with Presupposition Resolution
    author: Patrick Blackburn & Johan Bos

*****/

/*=====
   Semantic Macros
=====*/

detSem(uni, lambda(P, lambda(Q, lambda(H, lambda(L,
    merge(merge(usr([F, R, S],
                [F:drs([], [merge(drs([X], []), R)>S])),
                [leq(F, H), leq(L, S)]), P@X@H@R), Q@X@H@L))))).

detSem(indef, lambda(P, lambda(Q, lambda(H, lambda(L,
    merge(merge(usr([E, R, S],
                [E:merge(merge(drs([X], []), R), S)],
                [leq(E, H), leq(L, S)]), P@X@H@R), Q@X@H@L))))).

detSem(def, lambda(P, lambda(Q, lambda(H, lambda(L,
    merge(merge(usr([R, S],
                [L:alfa(X, def, merge(drs([X], []), R), S)],
                [leq(L, H)]), P@X@H@R), Q@X@H@S))))).

detSem(poss(Gender), lambda(P, lambda(Q, lambda(H, lambda(L,
    merge(merge(usr([R, S],
                [L:alfa(X, def, alfa(Y, nonrefl, drs([Y], [Cond]),
                merge(drs([X], [of(X, Y)]), R)), S)],
                [leq(L, H)]), P@X@H@R), Q@X@H@S))))):-
    compose(Cond, Gender, [Y]).

nounSem(Sym, lambda(X, lambda(_, lambda(L, usr([], [L:drs([], [Cond])], []))))):-
    compose(Cond, Sym, [X]).

pnSem(Sym, Gender, lambda(P, lambda(H, lambda(L,
    merge(usr([S],
                [L:alfa(X, name, drs([X], [X=Sym, Cond]), S)],
                [leq(L, H)]),
                P@X@H@S))))):-
    compose(Cond, Gender, [X]).

proSem(Gender, Type, lambda(P, lambda(H, lambda(L,
    merge(usr([S],
                [L:alfa(X, Type, drs([X], [Cond]), S)],
                [leq(L, H)]),
                P@X@H@S))))):-

```

```

compose(Cond, Gender, [X]).

ivSem(Sym, lambda(X, lambda(_, lambda(L,
  usr([], [L:drs([], [Cond])], [])])))):-
  compose(Cond, Sym, [X]).

tvSem(Sym, lambda(K, lambda(Y, K@lambda(X, lambda(_, lambda(L,
  usr([], [L:drs([], [Cond])], [])]))))):-
  compose(Cond, Sym, [Y, X]).

relproSem(lambda(P, lambda(Q, lambda(X, lambda(H, lambda(L,
  merge(usr([L2, L3, H1],
    [L:merge(L3, H1)],
    [leq(L2, H1)]),
  merge(P@X@H@L2, Q@X@H@L3)))))))).

prepSem(Sym, lambda(K, lambda(P, lambda(Y, lambda(H, lambda(L3,
  merge(K@lambda(X, lambda(H, lambda(L1,
  usr([L2, H1],
    [L3:merge(L2, H1), L1:drs([], [Cond])],
    [leq(L1, H1)]))@H@L1, P@Y@H@L2))))))):-
  compose(Cond, Sym, [Y, X]).

modSem(neg, lambda(P, lambda(X, lambda(H, lambda(L,
  merge(usr([N, S], [N:drs([], [~S])], [leq(N, H), leq(L, S)]),
  P@X@H@L)))))).

coordSem(conj, lambda(X, lambda(Y, lambda(P, lambda(H, lambda(L,
  merge(usr([L1, L2], [L:merge(L1, L2)], [leq(L, H)]),
  merge(X@P@H@L1, Y@P@H@L2)))))))).

coordSem(disj, lambda(X, lambda(Y, lambda(P, lambda(H, lambda(L,
  merge(usr([L1, L2], [L:drs([], [L1 v L2])], [leq(L, H)]),
  merge(X@P@H@L1, Y@P@H@L2)))))))).

dcoordSem(cond, lambda(C1, lambda(C2, lambda(H, lambda(L,
  merge(usr([H1, H2], [L:drs([], [H1 > H2])], [leq(L, H)]),
  merge(C1@H1@_, C2@H2@_)))))).

dcoordSem(conj, lambda(C1, lambda(C2, lambda(H, lambda(L,
  merge(usr([H1, H2], [L:merge(H1, H2)], [leq(L, H)]),
  merge(C1@H1@_, C2@H2@_)))))).

dcoordSem(disj, lambda(C1, lambda(C2, lambda(H, lambda(L,
  merge(usr([H1, H2], [L:drs([], [H1 v H2])], [leq(L, H)]),
  merge(C1@H1@_, C2@H2@_)))))).

```