

Ontologies:

Dynamic Networks of Formally Represented Meaning

Dieter Fensel

Vrije Universiteit Amsterdam, De Boelelaan 1081a, 1081 HV Amsterdam, The Netherlands,
phone/fax: +31-(0)84-872 27 22, dieter@cs.vu.nl, <http://www.cs.vu.nl/~dieter>.

Abstract. The computer was invented as a device for computation. Meanwhile the “computer” becomes a portal to cyberspace. It has become an entry point to a world-wide network of information exchange and business transactions. Therefore, technology that supports access to unstructured, heterogeneous and distributed information and knowledge sources will become as essential as programming languages were in the 60’s and 70’s. In this essay, we examine some of the essential requirements for such a technology.

1 Ontologies: Formal and Real, based on Consensus

The World-Wide Web (WWW) has drastically changed the availability of electronically available information. Currently there are around one billion documents in the WWW which are used by more than 300 million users internationally. In addition, this number is growing fast. However, this success and exponential growth makes it increasingly difficult to find, to access, to present, and to maintain the information of use to a wide variety of users. Currently, pages on the web must use representation means rooted in format languages such as HTML or SGML and make use of protocols that allow browser to present information to human readers. The information content, however, is mainly presented by natural language. Thus, there is a wide gap between the information available for tools that try to address the problems above and the information kept in human readable form. The current state of Web technology generates serious obstacles to its further growth. The technology's simplicity already caused bottlenecks that hinder searching, extracting, maintaining, and generating information (cf. [Fensel et al., 2000]). Computers are only used as devices that post and render information, but they do not have access to the actual content. Thus, they can only offer limited support in accessing and processing this information.¹ So, the main burden not only of accessing and processing information but also of extracting and interpreting it is on the human user.

Tim Berners-Lee envisioned a *Semantic Web* (cf. [Berners-Lee et al., 2001], [Fensel et al., to appear (b)]) that provides automated information access based on machine-processable semantics of data and heuristics that use these meta data. The explicit representation of the semantics of data, accompanied with domain theories (that is, ontologies), will enable a Web that provides a qualitatively new level of service. It will weave together an incredibly large network of human knowledge and will complement it with machine processability. Various automated services will help the user achieve goals by accessing and providing information in a machine-understandable form. This process might ultimately create an extremely knowledgeable systems with various specialized reasoning services-systems that can support us in nearly all aspects of our life and that will become as necessary to us as access to electric power.

Ontologies (cf. [Fensel, 2001]) are key enabling technology for the semantic web. They need to interweave human understanding of symbols with their machine processability. Therefore, it seems highly justified to take a closer look on the nature of Ontologies and on whether and how they can actually provide such a service. Ontologies were developed in Artificial Intelligence to facilitate knowledge sharing and reuse. Since the beginning of the nineties Ontologies have become a popular research topic investigated by several Artificial Intelligence research communities, including Knowledge Engineering, natural-language processing and knowledge representation. More recently, the notion of Ontology is also becoming widespread in fields such as intelligent information integration, cooperative information systems, information retrieval, electronic commerce, and knowledge management. The reason ontologies are becoming so popular is in large part due to what they promise: *a shared and common understanding of a domain that can be communicated between people and application systems.*

Because Ontologies aim at consensual domain knowledge their development requires a cooperative process. Ontologies are introduced to facilitate knowledge sharing and reuse between various agent, no matter whether they are of human or artificial nature. They should provide this service by providing a consensual and formal conceptualizations of a certain area. Spoken in a nutshell, *Ontologies are formal and consensual specifications of conceptualizations providing a shared and common understanding of a domain that can be communicated across people and application systems.* Therefore, Ontologies glue together *two essential aspects* that help to bring the web to its full potential:

- Ontologies define a *formal* semantics for information allowing information processing by a computer.
- Ontologies define a *real-world semantics* allowing to link machine processable content with meaning for humans based on *consensual* terminologies.

These two orthogonal aspects will be discussed during the following, however, our main focus is devoted to the second aspect.

¹. It is like using a telephone mainly for decorating a living room.

2 Ontologies define formal semantics

Ontologies provide formal semantics enabling machine-processable semantics of information. This aspect is already well-understood and several language proposals have been made (see [Fensel, 2001] and [Fensel et al., 2001]). Formal semantics is achieved by a layered language architecture. At the lowest level, XML² provides a serialized *syntax* for tree structures. RDF³ defines a basic *data model* on top of XML consisting of (object, property, value)-triples. RDF schema (RDFS)⁴ defines basic ontology primitives in RDF: classes with is-a and instance-of relationships, and properties with is-a relationships and domain and range restrictions. OIL⁵ extends RDFS to provide a full-fledged web-based ontology language. One of the central design ideas of OIL is its onion model (see Figure 1). There will never be one language meeting all man purposes. OILs onion model reflects this need. Languages of different complexity are provided allowing applications to select the degree of complexity they require. One of its dialects called **DAML+OIL**⁶ reflects a broad European and (US) American consensus on modeling primitives for the semantic web and is departure point for standardization by the W3C⁷.

3 Ontologies define real-world semantics

This aspect is still far from being studied properly. In this essay, we will focus on it, i.e., on how can Ontologies be used to communicate real-world semantics between human




Fig 1. The onion model to control complexity.

² <http://www.w3.org/XML/>

³ <http://www.w3.org/RDF/>

⁴ <http://www.w3.org/TR/2000/CR-rdf-schema-20000327/>

⁵ <http://www.ontoknowledge.org/oil>

⁶ <http://www.daml.org>

⁷ <http://www.w3c.org>

and artificial agents. For understanding this potential, we have to bring in an important point on how to look at them. This point of view is required to bring Ontology technology to its full potential and also brings into mind that most of the work on Ontologies is partially miss-focussed, i.e., ignores the main problems in building and using them.

Every first years philosophy student may have heart about the evils circle in trying to explain our ability of communication as a way to exchange meaning and to create understanding between human beings. On the one hand, people can only communicate and exchange meaning based on a common understanding of symbols and intensions. Therefore, a joined set of symbols and a consensual interpretation is the pre-request for communication. On the other hand, such a joined set of symbols and a consensual interpretation can only be established as a result of communication.⁸ Therefore, what is a result of successful communication is at the same time a pre-requisite for it. In consequence, its existence is required for explaining its existence. Our first years philosophy student may also have learned how to overcome such a paradoxical situation. There must be an underlying *process* that takes both sides as intermediate and repeatedly taken sub-steps relying on something that mediates between its extremes.⁹ Then successful communication and a joined set of understandings are just two sites of the same coin. The reader may found our arguments “too” philosophical. However, we want to undermine the principal difference between viewing ontologies as “*true*” *models of the real world* or steps in a *process of organizing evolving consensus*. Therefore, a brief argument on the cyclic nature of understanding and communication seems appropriate.

Viewed from an abstract philosophical point of view it looks like a miracle that two humans are able to understand each other. Taken in the extreme we cannot even be sure about our mutual existence. Since Descartes we take the fact that we are aware of our own thinking as proof of our own existence.¹⁰ However, we make notice of the existence of other agents via our perception and it is their existence in our perception and not their actual existence that follows from it.¹¹ Again we have to make the doubtful deduction that their existence in our perception reflects their actual behavior and existence. Even taking this assumption to be grounded we are still far away from explaining on how meaning can be exchanged between such brittle agents. Meaning and intention cannot be exchanged *directly*. Neither can it be expressed directly nor can we access the actual meaning that is perceived and understood by our counter part. We can only express our intension by some action that influences the perception of our counter part. And we can only guess what this is supposed to mean to him by analyzing his behavior as much as it is reflected in our perception.¹² In consequence, establishing meaning and communication (to exchange

⁸. At least as long as this interpretation is not hard-coded via instincts.

⁹. Cf. G. W. F. Hegel: *Wissenschaft der Logik*.

¹⁰. Already this conclusion could be viewed as being doubtful, however, its discussion would leave the scope of this paper.

¹¹. See for example I. Kant: *Critik der reinen Vernunft*.

¹². In principle, it is not even important whether another agent actually thinks. He “understands” our communicative acts properly if it is properly contained in the way he cooperates with us.

meaning) is per definition a *process*. People can only establish joined meaning and communicate it to each other in a process where they co-ordinate some of their actions to achieve common goals. Therefore, from the early beginning¹³ it can only be a social process that creates a joined understanding that is the basis for exchanging meaning with communicative symbols.

Following this argument it is also rather clear that there will be neither such a thing as THE Ontology where everybody subscribes to. Instead, ontologies arise as pre-requisite and result of cooperation in certain areas reflecting task, domain, and sociological boundaries. In the same way as the web weaves billions of people together to support them in their information needs, Ontologies can only be thought as a network of interweaved Ontologies. This network of Ontologies may have overlapping and excluding pieces, and it must be as dynamic in nature as the dynamics of the process it underlies. This view on *Ontologies as dynamic networks of formally represented meaning* is what we want to stress in the essay. Most work on Ontologies view Ontologies as a isolated theory containing possible large number of concepts, relationships, and constraints that further detach formal semantics to them. Here we take a much broader view on Ontologies. Basically, there are two main dimensions in which these mediators of communication differ from current work on Ontologies: *Ontologies must have a network architecture* and *Ontologies must be dynamic*.

3.1 Heterogeneity in Space: Ontology as Networks of Meaning

Island of meaning must be interwoven to form more complex structures enabling exchange of information beyond domain, task, and sociological boundaries. This implies two efforts. Tool support must be provided to define local domain models that express a commitment of a group of agents that share a certain domain and task and that can agree on a joined world view for this purpose. Here much work has already been spent and significant methodological support is available (see [Fensel et al., to appear (a)] for a survey). Second, these local models must be interwoven with other models like the social practice of the agents that use Ontologies to facilitate their communicational needs. Here not much work has been spent. We do no longer talk about a single Ontology but rather about a network of Ontologies. Links must be defined between these Ontologies and this network must allow overlapping Ontologies with conflicting and even contradictory conceptualizations. From the early beginning heterogeneity is an essential requirement for this Ontology network. Means to deal with conflicting definitions and strong support in interweaving local theories are essential requirements for making this technology workable and scalable.

Take a Peer-to-Peer (P2P) network like Gnutella as an example (cf. [Oram, 2001]). Agents can dynamically enter and leave the network. Agents can communicate with a

¹³. Both, in a *historical* and in a *logical* sense.

local environment of other agents. This network is dynamically set up and collapsed according to the joined needs of a group of agents. Current work on Ontologies that focuses either on local domain theories or on principles, structures, and content of the right upper-layer Ontology are far way from supporting such a vision. What is needed is focus on:

- *linking local conceptualizations* dealing with heterogen definitions and personalized views,
- support in easy *configuration and re-configuration of such networks* according to the communication needs of agent coalitions, and
- methods and tools that help agents in *organizing consensus* allowing them to exchange meaning.

Ontologies ensure communication between various agents. They are “right” if they fulfill this purpose.

3.2 Development in time: Living Ontologies

Originally, an Ontology should reflect the “truth” of a certain aspect of reality. It was the holy task of a philosopher to find such truth. Nowadays Ontologies are used as means to exchange meaning between different agents. They can only provide this if they reflect an inter-subjectual consensus. Per definition they can only be the result of a social process. This gives ontologies a dual status for the exchange of meaning.

- Ontologies as *pre-requisite* for consensus: Agents can only exchange meaning when they have already agreed on a joined body of meaning reflecting a consensual point of view on the world.
- Ontologies as a *result* of consensus: Ontologies as consensual models of meaning can only arise as result of a process where agents agree on a certain world model and its interpretation.

In consequence, ontologies are as much a pre-requisite of consensus and information sharing as they are its results. Therefore, ontologies cannot be understood as a static model. An ontology is as much required for the exchange of meaning as the exchange of meaning may influence and modify an ontology. In consequence, *evolving* ontologies rather describe a process than a static model. Having protocols for the process of evolving ontologies is the real challenge. Evolving over time is an essential requirement for useful ontologies. As the daily practice constantly changes, Ontologies that mediate the information needs of these processes must have strong support in *versioning* and must be accompanied by *process models* that help to organize consensus.

Centralized process models have standardization bodies as central clearing unit. This central unit may soon become a bottleneck for the scalability of the entire process. Often

such standardization works slow and lead to mongrelized results. Decentralized process models for consensus achievement can be based on the natural consensus of working networks. Therefore, they can reflect true, proven useful, and broadly used consensus. In this context, one may want to take a look at P2P, where networks arise and are configured dynamically according to joined interests of loosely coupled groups.

4 Conclusions

Ontologies help to establish *consensual terminologies* that make sense to both sites. *Computers* are able to process information based on their machine-processable semantics. *Humans* are able to make sense of this information based on their connection to real-world semantics. Building up such ontologies that are pre-requisite and result of joined understanding of large user groups is far from being trivial. A model or “protocol” for driving the network that maintains the process of *evolving Ontologies* is the real challenge for making the *semantic* web reality.

Most work on Ontologies view ontologies as a isolated theory containing possible large number of concepts, relationships, and constraints that further detach formal semantics to them. In the paper we took a much broader view on ontologies. We view Ontologies as highly interwoven *networks* allowing to deal with heterogenic needs of the communication processes that should mediated by them. Second, these ontologies must shift over time as the processes they mediate based on consensual representation of meaning. It is the network and dynamic character of *Ontologies* that make further research work on them so exiting. The *glue*, that link together Ontology networks in space and time, is the actual challenge on current work on ontologies. It is the glue, stupid!

References

[Berners-Lee et al., 2001]

T. Berners-Lee, J. Hendler, and O. Lassila: The Semantic Web, *Scientific American*, May 2001.

[Fensel, 2001]

D. Fensel: *Ontologies: Silver Bullet for Knowledge Management and Electronic Commerce*, Springer-Verlag, Berlin, 2001.

[Fensel et al., 2000]

D. Fensel, F. van Harmelen, H. Akkermans, M. Klein, J. Broekstra, C. Fluyt, J. van der Meer, H.-P. Schnurr, R. Studer, J. Davies, J. Hughes, U. Krohn, R. Engels, B. Bremdahl, F. Ygge, U. Reimer, and I. Horrocks: OnToKnowledge: Ontology-based Tools for Knowledge Management. In *Proceedings of the eBusiness and eWork 2000 (EMMSEC 2000) Conference*,

Madrid, Spain, October 2000.

[Fensel et al., 2001]

D. Fensel, I. Horrocks, F. van Harmelen, D. McGuinness, and P. Patel-Schneider: OIL: Ontology Infrastructure to Enable the Semantic Web, *IEEE Intelligent Systems*, March/April 2001.

[Fensel et al., to appear (a)]

D. Fensel, J. Hendler, H. Lieberman, and W. Wahlster: The Semantic Web: Why, What, and How. In [Fensel et al., to appear (b)].

[Fensel et al., to appear (b)]

D. Fensel, J. Hendler, H. Lieberman, and W. Wahlster (eds.): *Semantic Web Technology*, MIT Press, Boston, to appear.

[Oram, 2001]

A. Oram (eds.): *Peer-to-Peer: Harnessing the Benefits of a Disruptive Technology*, O'Reilly, Sebastopol, 2001.