Delivering Configurable Problem-Solving Services to Web Users

Monica Crubézy and Mark A. Musen
Stanford Medical Informatics
MSOB X-215, 251 Campus Drive
Stanford, CA 94305-5479 USA
+1 650 723-6979
{crubezy,musen}@smi.stanford.edu

The World-Wide Web is suitable not only for distributing static data resources but also for providing users with dynamic services. Currently, a growing number of online services are appearing, offering functionality such as booking a performance ticket, organizing a meeting, or comparing experimental genomic data against reference databases of gene sequences. Existing services, however, tend to be integral: The underlying reasoning process that these services perform is neither modifiable, nor configurable for different domains, nor can it be integrated with other services to produce new functionalities. As a result, the level of functionality that users can expect from Web services is limited to performing simple, predefined tasks. Such tasks typically do not handle large amounts of knowledge. Meanwhile, the current evolution of the Web into a Semantic Web encourages communities of users to create and publish shared domain knowledge conceptualizations—ontologies—with which they describe the data that they exchange. We anticipate that user communities soon also will need to exchange the processes and services that operate on their shared data and knowledge to solve complex tasks in their domain. For instance, developers and users of image-processing techniques already try to share either software modules or experience with specific tools. In the context of the Semantic Web, the goal of our research is to develop the technologies and frameworks needed to provide sophisticated online reasoning facilities, configurable for different domains and applications. For instance, we can imagine making available online a generic heuristic classification technology, which different communities (e.g., paleontologists, geologists, biologists) or individual users can adapt for specific applications. We believe that the Semantic Web will provide the grounding infrastructure to deliver *intelligent problem-solving services* to users willing to achieve knowledge intensive tasks in their domain.

At Stanford Medical Informatics[1], we have been developing knowledge systems for different application domains and purposes for many years. Our approach consists in building intelligent systems by assembling reusable knowledge components, namely *domain ontologies* and *problem-solving methods*. Problem-solving methods are domain-independent problem-solving strategies, that provide standard ways of addressing stereotypical problems, or *generic tasks*, such as diagnosis, design, and classification. Our major contribution in the area of knowledge engineering is incarnated by the Protégé[2] suite of tools that supports developers in building a running application system by creating, configuring and assembling reusable components. Protégé offers a generic, full-fledged knowledge-modeling and user-interface support, extensible and customizable for different purposes and users. In particular, Protégé supports the process of selecting a problem-solving method suitable for the user problem, and configuring the method by integrating its generic components with domain-specific information through the use of explicit architectural elements—declarative *mapping relations*.

Delivering problem-solving services to Web users involves developing the necessary infrastructure and technologies that can enable service identification and interoperability on the Web. In

[1] http://www.smi.stanford.edu/

[2] http://protege.stanford.edu/

particular, the competence of each Web service needs to be specified and advertized so that brokering agents can then reason about these competence descriptions to match services to user needs and to combine services together. Ongoing work in the IBROW project[3] is addressing these issues and aims to produce a brokering agent that can locate and configure problem-solving services on the Web to solve specific tasks requested by users. In particular, the IBROW consortium has developed the Uniform Problem-solving Method development Language (UPML[4])—a modeling framework and a markup language (in the form of an RDF Schema[5]) for specifying problem-solving components on the Web. UPML provides a comprehensive ontology to specify the properties of each kind of component involved in a knowledge system. For instance, problem-solving methods define a set of input and output roles, as well as set of logical formulas, which provide a functional specification of their competence. UPML also provides explicit ways of specifying the assembly of components in a knowledge system. In the context of IBROW, we have developed a special-purpose editor, based on Protégé, that enable developers to model and advertize problem-solving services on the Web, as resources marked-up in UPML.

Jointly with the Knowledge Media Institute[6] at the Open University (UK), we also have developed the IBROW Internet Reasoning Service (IRS), a tool that enables developers to prototype knowledge systems quickly by configuring reusable reasoning components from online structured libraries. The process of configuring problem-solving components into a running application may involve several activities: mapping generic tasks and methods to a domain model (e.g., mapping a generic classification framework to a database of archeological artifacts to produce an artifact-classification application), mapping methods to tasks (e.g., selecting a particular abstraction method for performing a data-abstraction task), or, in general, refining existing components (e.g., specializing a data-to-solutions matching component by introducing fuzziness in the matching process). Based on the knowledge-level descriptions of the problem-solving components in UPML, the IRS provides different levels of user support, from interactive browsing and navigation facilities, which enable manual selection and configuration of reasoning components, to intelligent, semi-automated assistance in building an executable application. We have implemented a version of the IRS as an extension to Protégé, that interfaces domain knowledge bases to UPML-compliant libraries of problem-solving components.

The ultimate aim of our work is to make sophisticated problem-solving technology available to a wider audience and provide the level of intelligent support needed to allow rapid generation of web-based reasoning services. A side-effect of achieving this goal will be simply to make much artificial intelligence technology available online, thus making its use more widespread. The Semantic Web is a vehicle that will allow us to perform a large evaluation experiment of our approach, by delivering problem-solving services directly to users. At the same time, the heterogeneous, distributed and versatile nature of the Web will challenge us to incorporate more automated support, with simplified procedures and customized interactions, so that less-experienced users also can benefit from advanced distributed problem-solving services. We anticipate a future where the Internet will make libraries of problem solvers available to any person with a Web browser. The challenge is to describe and index these problem solvers meaningfully, and to develop the infrastructure that allows users to locate these problem solvers easily and to link them to specific data sets and knowledge bases in simple, intuitive ways. Our research is exactly a step towards a Semantic Web in which users can access not only knowledge but also intelligent problem solvers.

---

[3] http://www.swi.psy.uva.nl/projects/ibrow/home.html

[4] http://www.cs.vu.nl/~upml/

[5] http://www.w3.org/TR/rdf-schema/

[6] http://kmi.open.ac.uk/