

Applications of Ontologies and Problem-Solving Methods

Asunción Gómez-Pérez and V. Richard Benjamins

■ The Workshop on Applications of Ontologies and Problem-Solving Methods (PSMs), held in conjunction with the Thirteenth Biennial European Conference on Artificial Intelligence (ECAI '98), was held on 24 to 25 August 1998. Twenty-six people participated, and 16 papers were presented. Participants included scientists and practitioners from both the ontology and PSM communities. The first day was devoted to paper presentations and discussions. The second (half) day, a joint session was held with two other workshops: (1) Building, Maintaining, and Using Organizational Memories and (2) Intelligent Information Integration. The reason for the joint session was that in all three workshops, ontologies play a prominent role, and the goal was to bring together researchers working on related issues in different communities. The workshop ended with a discussion about the added value of a combined ontologies-PSM workshop compared to separate workshops.

One of the main motivations underlying both ontologies and problem-solving methods (PSMs) is to enable sharing and reuse of knowledge and reasoning behavior across domains and tasks. PSMs and ontologies can be seen as complementary reusable components to construct knowledge systems from reusable components. Ontologies are concerned with static domain knowledge and PSMs with dynamic reasoning knowledge. To build full applications of information and knowledge systems from reusable components, both PSMs and ontologies are required in a tightly integrated way. The integration of ontologies and PSMs is a possible solution to the interaction

problem (Chandrasekaran 1998), which states that representing knowledge for the purpose of solving some problem is strongly affected by the nature of the problem and the inference strategy to be applied to the problem. Through ontologies and PSMs, this interaction can be made explicit and taken into consideration.

Ontologies

Ontologies aim at capturing domain knowledge in a generic way. An *ontology*, therefore, provides a commonly agreed understanding of a domain, which can be reused and shared across

Ontologies provide a common vocabulary of an area and define—with different levels of formality—the meaning of the terms and the relations between them.

applications and groups (Uschold and Grüninger 1996). Ontologies provide a common vocabulary of an area and define—with different levels of formality—the meaning of the terms and the relations between them. Ontologies are usually organized in taxonomies and typically contain modeling primitives such as classes, relations, functions, axioms, and instances (Gruber 1993).

Until now, few domain-independent methodologies have been report-

ed to build ontologies. Uschold's methodology (Uschold and Grüninger 1996), Grüninger and Fox's (1995) methodology, and METHONTOLOGY (Gómez-Pérez 1998; Fernández, Gómez-Pérez, and Juristo 1997) are the most representative ones, which have in common that they start from the identification of the purpose of the ontology and the need for domain knowledge acquisition. However, having acquired a significant amount of knowledge, Uschold proposes codification in a formal language, and METHONTOLOGY proposes expressing the ontology at the knowledge level as a set of intermediate representations based on tabular and graph notations.

Several languages can be used to formalize the content of an ontology at the symbol level. Usually, a language is attached to a given ontology server. The most representative languages are ONTOLINGUA (Gruber 1993), CYCL (Lenat and Guha 1990), and LOOM (MacGregor 1991). ONTOLINGUA is the language used by the ONTOLOGY SERVER (Farquhar et al. 1997). CYCL is the language used in the CYC Project, and LOOM is the language used by the server called ONTOSAURUS (Swartout et al. 1997).

Although ontologies can be used (Uschold and Grüninger 1996) to communicate between systems, people, and organizations, interoperate between systems, and support the design and development of knowledge-based and general software systems, the number of applications built that use ontologies to model the application knowledge is small. That is, many times such ontologies have been built just for a given application without special consideration for sharing and reuse. Several problems make it difficult to reuse existing ontologies in applications: Ontologies are dispersed over several servers; the formalization differs depending on the server on which the ontology is stored; ontologies on the same server are usually described with different levels of detail; and there is no common format for presenting relevant information about the ontologies so that users can decide which ontology best suits their purpose. These problems are probably the cause for the rel-

A PSM defines a way to achieve the goal of a task. It has input and output and can decompose a task into subtasks.

atively small number of known applications until now in areas such as knowledge management, ontology-based brokers, natural language generation, enterprise modeling, knowledge-based systems, and interoperability between systems.

Problem-Solving Methods

PSMs describe the reasoning process of a knowledge-based system in an implementation- and domain-independent manner. A PSM defines a way to achieve the goal of a task. It has input and output and can decompose a task into subtasks. In addition, a PSM specifies the data flow between its subtasks. Control knowledge determines the execution order and iterations of the subtasks of a PSM. Control knowledge can be specified in advance, if known, or can be determined opportunistically at run time depending on the dynamic problem-solving situation (Benjamins 1995). PSMs can be used to efficiently achieve goals of tasks through the application of domain knowledge (Fensel and Straatman 1998). They can play several roles in the knowledge-engineering process, such as guiding the acquisition process of domain knowledge and facilitating knowledge-based system development through their reuse.

Work on PSMs covers different areas such as the identification of task-specific PSMs (for diagnosis, planning, assessment, and so on), the storing and indexing of PSMs in libraries, and the formalizing of PSMs. The difficulty of reusing PSMs is that one has to find the right PSM (that does—part of—the job), check whether it is applicable in the situation at hand, and modify it to fit the domain. To reuse PSMs successfully in a real-life application, one has to understand these processes. Recently, a few industrial applications have seen the light, which shows that the

reuse of PSMs is also interesting from an industrial point of view (for example, see four papers in Benjamins and Fensel [1998]).

Aim of This Joint Workshop on Ontologies and PSMs

In the past years, several separate workshops on ontologies (for example, ECAI'94, ECAI'96, IJCAI'93, IJCAI'95, FOIS'98) and PSMs (for example, KAW'96, KAW'98, IJCAI'97) have been organized that focused largely on theoretical aspects such as engineering, designing, building, maintaining, and using ontologies and PSMs and less on applications. As a result, there is now reasonable understanding of, and consensus on, the nature of ontologies and PSMs. Real applications built through ontologies and PSMs are, however, still rare. Moreover, the relation between ontologies and PSMs is an important issue. For example, Neches and his colleagues (1991) proposed an architecture for knowledge-based systems based on reusing generic reasoning modules and ontologies. Van Heijst, Schreiber, and Wielinga (1997) describe how to merge generic model components (task-model components and ontologies) to build the knowledge model of an application.

Apart from bringing together researchers from the two fields, this workshop had two aims: (1) promote a deep understanding of how ontologies and PSMs can be used in real applications and (2) pursue how they relate to each other.

Themes and the Papers in the Proceedings

The 16 papers accepted for this workshop deal with ontologies, PSMs, and their integration. Within each of these

categories, the following subdivisions can be made.

Ontologies

Classifying and characterizing ontologies: Uschold questions in his paper why there are no “killer applications” built using ontologies. He introduces a scheme for classifying ontology applications along a number of dimensions (purpose, representation languages and paradigms, meaning and formality, subject matter, development, conceptual architecture, mechanisms and techniques to use the ontology, and implementation platforms) and advocates that such schema be used when research applications results are reported.

In the paper “(ONTO)2AGENT: An Ontology-Based WWW Broker to Select Ontologies,” Julio Arpírez, Asunción Gómez-Pérez, Helena Pinto (all of Technical University of Madrid), and Adolfo Lozano (University of Extremadura) present three important problems that ontology users face when trying to reuse existing ontologies: No standardized identifying features characterize ontologies from the user's point of view; no web sites use the same logical organization to present relevant information about ontologies; and the search for appropriate ontologies is hard, time consuming, and usually fruitless. To solve these problems, they present a set of features to characterize ontologies from a user's point of view and (ONTO)2AGENT, a broker that uses a *reference ontology* (a living domain ontology about ontologies) as a source of its knowledge to retrieve descriptions of ontologies that satisfy a given set of constraints.

Ontologies and natural language generation: Marcel Fröhlich (University of Tübingen) and Reind van de Riet (Free University of Amsterdam) present in their paper “Using Multiple Ontologies in a Framework for Natural Language Generation” a framework that integrates domain-independent tools for natural language generation (KPML) and integration of existing ontologies in a natural language generation pipeline architecture.

In “ONTOGENERATION: Reusing Domain and Linguistic Ontologies for

Spanish Text Generation," Guadalupe Aguado, Alvaro Sánchez (both of Technical University of Madrid), John Bateman (University of Sterling), Gómez-Pérez, and colleagues propose a general approach to reuse domain (chemicals) and linguistic (generalized upper model) ontologies with natural language generation technology (KPML), describing a practical system for the generation of Spanish texts in the domain of chemical substances.

Integration of ontologies: Stuart Aitken (AIAI) describes in "Extending the HPKB-Upper-Level Ontology: Experiences and Observations" an experience of extending the HPKB upper-level ontology. The major claim is that reuse by extension is only possible if the ontology to be reused is understandable, and its design principles are made explicit.

Hercules Dalianis (KTH) and Eduard Hovy (University of Southern California) describe a semiautomated method to merge two STEP schemata. The paper shows how a set of simple metrics can be applied for revealing likely matches and nonmatches across concepts taken from different ontologies.

Pepijn Visser (University of Liverpool) and Zhan Cui (BT Laboratories) describe in "Heterogeneous Ontology Structures for Distributed Architectures" an experiment in the domain of business processes. They propose to have a hierarchical structure of heterogeneous ontologies to be used for the integration of heterogeneous and distributed information systems. Rather than trying to achieve one shared consensual domain ontology for all information systems, their starting point is to define a hierarchical ontology structure that comprises several smaller but heterogeneous shared ontologies. Each information system that is part of the distributed architecture maps its local ontology onto one of the ontologies in the hierarchy using a kind of mapping relation.

Applications that use ontologies: In his paper ("Spotting Ontological Lacunae through Spectrum Analysis of Retrieved Documents"), Edward Hoenkamp (NICI) describes how an ontology-based information filter can spot a potential lacuna in an ontology by analyzing retrieved documents.

The feedback of a spectrum analysis of the retrieved documents and terminological data in WORDNET helps to locate lacunae in an ontology that represents user-information need.

The paper by Zhi Jin, David Bell, George Willkie, and D. Leahy (all of University of Ulster), "Automatically Acquiring Requirements of Business Information Systems by Reusing Business Ontologies," proposes an approach for automatically acquiring requirements for business information systems by reusing the domain knowledge, which consists of a business ontology and a domain ontology.

In "Applying the Process Interchange Format (PIF) to a Supply Chain Process Interoperability Scenario," Steve Polyak (University of Edinburgh), Michael Grüninger (University of Toronto), J. Lee (University of Hawaii), and C. Menzel (Texas A&M University) describe the application of PIF in a knowledge-sharing effort to facilitate business-process reengineering of supply-chain activities. PIF acts as an interlingua between two separate tools used in the modeling and simulation of the proposed processes. The paper illustrates the benefits of using a process ontology to capture domain knowledge in a generic way so that it can be reused across applications and shared across groups.

Problem-Solving Methods

Building KBS through reuse of PSMs: In the paper "Experiments in Building Program Supervision Engines from Reusable Components," Monica Crubezy, Sabine Moisan (both of INRIA), and Mar Marcos (University of Jaume) describe a number of knowledge-based systems built from reusable problem-solving methods. It shows how the problem-solving methods from three application systems (PEGASE, PULSAR, and MEDIA) shared sub-methods and were configured from a common set of components. All these systems were used in the area of program supervision.

Jose Sierra and Martin Molina (both of Technical University of Madrid) describe in their paper, "Terminological Importation for Adapting Reusable Knowledge Representation Components in the KSM Environment," an

adaptation approach for reusable knowledge representation components based on a particular form of ontological mapping called *terminological importation*. Terminological importation is a mechanism to populate a representation terminology associated with a knowledge representation component with concepts taken from another terminology about a concrete domain. The approach is used in the KSM environment for adapting basic reusable, knowledge representation software components in different domains.

Integration of Ontologies and PSMs

B. Chandrasekaran, John Josephson (both of The Ohio State University), and Richard Benjamins (University of Amsterdam) ("Ontology of Tasks and Methods") claim that ontologies for problem-solving knowledge are just as important as for domain knowledge. Traditionally, the field of ontology has been concerned with static domain knowledge. The authors analyze problem-solving knowledge and provide a list of primitive components that, according to their analysis, should be identified when describing any problem-solving knowledge. In addition to these components, they define how PSM components get connected with domain knowledge, and they suggest making explicit the assumptions about how problem-solving knowledge expects factual knowledge to be structured.

The paper by Mitsuru Ikeda, Kazuhisa Seta, Riichiro Mizoguchi (all of Osaka University), and Osama Kakusho (Hyogo University), "An Ontology for Building a Conceptual Problem-Solving Model," is about the use of task ontologies and their role in forming a bridge between a user's vocabulary and the system's terminology. The authors pay considerable attention to the execution of a conceptual model.

In "IBROW3, An Intelligent Brokering Service for Knowledge-Component Reuse on the World Wide Web," the authors provide a global overview of an Esprit Project for the plug and play of PSMs. A web broker mediates between customers and PSM providers

to configure a customized knowledge system for solving problems of the customers. The different worlds of the customers, the broker, and the providers are modeled through ontologies.

In his paper "DESIRE, an Interoperative Environment for Distributed Expert Systems," Takahira Yamaguchi (Shizuoka University) proposed using an ontology to facilitate communication among expert systems to improve their performance in a distributed environment.

Joint Session

Besides presentations of the papers accepted for the workshop, there was also a joint session with two other workshops: (1) Building, Maintaining, and Using Organizational Memories and (2) Intelligent Information Integration. The common theme of the three workshops was ontologies. In knowledge management, ontologies can play the role of an organizational memory. In information integration, ontologies can be used to integrate heterogeneous information sources. This session had three invited talks: First, James Hendler of the University of Maryland talked about, among other things, the role of ontologies for knowledge representation on the web. Second, Gio Wiederhold of Stanford University discussed, among other things, the role of ontologies in large-scale information systems. Third, Ulrich Reimer of Swiss Life talked about the role of ontologies in knowledge management.

Conclusions

This workshop provided a platform to present, discuss, and evaluate applications of ontologies and PSMs in areas such as knowledge management and enterprise modeling, communication between people and organizations, interoperability between systems, natural language generation, and integration of ontologies and PSMs in applications. Overall, the combination of ontologies and PSMs in one workshop was evaluated positively—not that there were so many papers about their integration (although there were some), but it was felt that being aware

of the other work was certainly worthwhile. After all, a knowledge system contains both domain knowledge and reasoning knowledge. The workshop papers are available at //delicias.dia.fi.upm.es/WORKSHOP/ECAI98/index.html.

Acknowledgments

Richard Benjamins was supported by the Netherlands Computer Science Research Foundation with financial support from the Netherlands Organisation for Scientific Research (NWO).

References

- Benjamins, V. R. 1995. Problem-Solving Methods for Diagnosis and Their Role in Knowledge Acquisition. *International Journal of Expert Systems: Research and Applications* 8(2): 93–120.
- Benjamins, V. R., and Fensel, D. 1998. *International Journal of Human-Computer Studies* (Special Issue on Problem-Solving Methods) 49(4): 305–649.
- Chandrasekaran, B. 1998. Generic Tasks in Knowledge-Based Reasoning: The Right Level of Abstraction for Knowledge Acquisition. In *Knowledge Acquisition for Knowledge-Based Systems*, eds. B. R. Gaines and J. H. Boose, 65–77. San Diego, Calif.: Academic.
- Farquhar, A.; Fikes, R.; Pratt, W.; and Rice, J. 1997. The ONTOLOGIA Server: A Tool for Collaborative Ontology Construction. *International Journal of Human-Computer Studies* 46(6): 707–728.
- Fensel, D., and Straatman, R. 1998. The Essence of Problem-Solving Methods: Making Assumptions to Gain Efficiency. *International Journal of Human-Computer Studies* 48(2): 181–215.
- Fernández, M.; Gómez-Pérez, A.; and Juristo, N. 1997. METHONTOLOGY: From Ontological Art toward Ontological Engineering. Paper presented at the Spring Symposium Series on Ontological Engineering, 24–26 March, Stanford, California.
- Gómez-Pérez, A. 1998. Knowledge Sharing and Reuse. In *The Handbook of Applied Expert Systems*, ed. J. Liebowitz, 10-1–10-36. Northwest Boca Raton, Fla.: CRC.
- Gruber, T. 1993. A Translation Approach to Portable Ontology Specifications. *Knowledge Acquisition* 5(2): 199–220.
- Grüniger, M., and Fox, M. S. 1995. Methodology for the Design and Evaluation of Ontologies. Paper presented at the IJCAI Workshop on Basic Ontological Issues in Knowledge Sharing, 19–21 August, Montreal, Quebec, Canada.
- Lenat, D., and Guha, R. 1990. *Building Large*

Knowledge-Based Systems: Representation and Inference in cyc Project. Reading, Mass.: Addison-Wesley.

MacGregor, R. 1991. Inside LOOM Classifier. *SIGART Bulletin* 2(3): 70–76.

Neches, R.; Fikes, R.; Finin, T.; Gruber, T.; Patil, R.; Senator, T.; and Swartout, R. 1991. Enabling Technology for Knowledge Sharing and Reuse. *AI Magazine* 12(3): 36–56.

Swartout, B.; Patil, R.; Knight, K.; and Russ, T. 1997. Toward Distributed Use of Large-Scale Ontologies. Paper presented at the Spring Symposium on Ontological Engineering, 24–26 March, Stanford, California.

Uschold, M., and Grüniger, M. 1996. Ontologies: Principles, Methods, and Applications. *Knowledge Engineering Review* 11(2): 93–155.

van Heijst, G.; Schreiber, A.; and Wielinga, B. 1997. Using Explicit Ontologies in KBS Development. *International Journal of Human-Computer Studies* (Special Issue on Using Explicit Ontologies in Knowledge-Based System Development) 46(2–3): 183–292.

Asunción Gómez-Pérez received a B.A. in computer science (1990), an MS.C. in knowledge engineering (1991), and a Ph.D. in computer science (1993) from the Universidad Politécnica de Madrid, Spain. She also has an MS.C. in business administration (1994) from the Universidad Pontificia de Comillas, Spain. She was a visitor in the Knowledge Systems Laboratory at Stanford University from 1994 to 1995. She is associate professor at the Computer Science School at Universidad Politécnica de Madrid, Spain. She was also the executive director of the Artificial Intelligence Laboratory from 1995 to 1998. Her current research activities include theoretical ontological foundations, methodologies for building ontologies, ontological reengineering, evaluation of ontologies, uses of ontologies in applications, and knowledge management on the web. Her e-mail address is asun@fi.upm.es.

Richard Benjamins is a senior researcher and lecturer in the Department of Social Science Informatics at the University of Amsterdam, where he graduated cum laude in 1988 and obtained his Ph.D. in 1993 in cognitive science. He worked at several research institutes, including the University of São Paulo, Brazil, the University of Paris-South, and the Spanish Artificial Intelligence Research Institute in Barcelona. His research interests include knowledge engineering, problem-solving methods and ontologies, diagnosis and planning, and AI and the web. His e-mail address is richard@swi.psy.uva.nl.