– IBROW3 –
An Intelligent Brokering Service for Knowledge-Component Reuse on the World-Wide Web

V. Richard Benjamins 1,3 Enric Plaza 1 Enrico Motta 4
Dieter Fensel 2 Rudi Studer 2 Bob Wielinga 3
Guus Schreiber 3 Zdenek Zdrahal 4

1 Artificial Intelligence Research Institute (IIIA), Spanish Council for Scientific Research (CSIC), Campus UAB, 08193 Bellaterra, Barcelona, Spain, {richard,enric}@iiia.csic.es
2 University of Karlsruhe, Institute AIFB, 76128 Karlsruhe, Germany, {dfe,studer}@aifb.uni-karlsruhe.de,
3 Dept. of Social Science Informatics (SWI), University of Amsterdam, Roetersstraat 15, 1018 WB Amsterdam, The Netherlands, {richard,schreiber,wielinga}@swi.psy.uva.nl
4 Knowledge Media Institute, The Open University, Waltin Hall, Milton Keynes, United Kingdom, {E.Motta,Z.Zdrahal}@open.ac.uk

Abstract

The World-Wide Web is changing the nature of software development to a distributive plug & play process. This requires a new way of managing software by so-called intelligent software brokers. The aim of the European IBROW3 project is to develop an intelligent brokering service that enables third party knowledge-component reuse through the World-Wide Web. Suppliers provide libraries of knowledge components adhering to some standard, and customers can consult these libraries – through intelligent brokers – to configure a knowledge system suited to their needs by selection and adaptation. IBROW3 integrates research on heterogeneous databases, interoperability and web technology with knowledge-system technology and ontologies. The aim is to develop a broker that can handle web requests for classes of knowledge system (e.g. diagnostic systems) by accessing libraries of reusable problem-solving methods on the Web, and selecting, adapting and configuring these methods in accordance with the domain at hand.

The aim of this paper is to give a general overview of the project and to presents its main ideas and approach. IBROW3 has started on January 1, 1998 and thus we can only present preliminary results.

1 Introduction

The World-Wide Web has the potential to change the nature of software development to a distributive plug & play process. This requires a new way of managing software by so-called intelligent software brokers [Decker *et al.*, 1997, Fensel, 1997a, Fensel *et al.*, 1997a].

The aim of the IBROW3 project is to develop an intelligent brokering service that enables third party knowledge-component reuse through the World-Wide Web. Suppliers provide libraries of knowledge components adhering to some standard, and customers can consult these libraries – through intelligent brokers – to configure a knowledge system suited to their needs by selection and adaptation. A customer in this context is a person/company that wants to solve a reasoning problem, such as the design of a technical product; a supplier could be a commercial or non-commercial provider of knowledge components that implement reasoning components. The project, thus, aims at providing intelligent *reasoning* services on the Web, as opposed to the more common *data* services [Wiederhold, 1996].

The knowledge components we want to reuse are *problem-solving methods* (PSMs) [McDermott, 1988]. Today, many PSM repositories exists at different locations [Benjamins, 1993, Breuker, 1994, ten Teije, 1997, Barros, L. Nunes de *et al.*, 1997, Valente & Löckenhoff, 1993, Chandrasekaran, 1990, Motta & Zdrahal, 1996], but they are not accessible for outsiders nor are they compatible (interoperable). There are some proposals for standards for reuse in knowledge engineering [Genesereth & Fikes, 1992, Schreiber *et al.*, 1994b], and there are emerging standards in interoperability [Orfali *et al.*, 1996, JavaSoft, 1996]. The objective is to investigate whether it is possible to combine and elaborate both technologies to enable scaling up software reuse to a world-wide level, using the Internet.

2 Motivation

We think that reuse will play an essential role in industrial knowledge system development in the next century [Neches *et al.*, 1991]. The current technical state-of-the-art in knowledge systems is that existing PSM libraries provide *individually* substantial support for engineering knowledge systems (see [Speel & Aben, 1997] for an example in industry). However, there are two important outstanding problems that prohibit more widespread industrial take-up and employment of existing knowledge-system technology. First, the PSMs of the existing libraries are implemented in heterogeneous languages, and therefore they cannot *interoperate*. Second, the components (PSMs) of the libraries are not *accessible* for third parties who do not know the component's implementation language. In the database world, work is being performed on interoperability of distributed databases using so-called *mediators* [Wiederhold & Genesereth, 1997, Wiederhold, 1996, Arens *et al.*, 1993, Duschka & Genesereth, 1997]. We aim to reuse, adapt and extend this work to solve (parts of) the two outstanding problems mentioned above.

The project opens a novel perspective on building knowledge systems. It shifts the focus from centralized development of knowledge systems to configuring them by means of distributed third party reusable components. In current commercial practice, a knowledge system is either developed by a company's inhouse R&D department, or by contracting a software company. Successful completion of IBROW3 could contribute to changing this practice and to make knowledge system technology applicable on a much larger scale, while at the same time reducing its costs. The approach also allows the rapid generation of knowledge systems for exploratory prototyping purposes. This would enable industries to quickly and accurately assess whether knowledge-system technology is useful for them. In essence, intelligent brokers cooperate with, and take away much of the burdens an engineer is confronted with in the construction process. Associated benefits include the following.

- Knowledge-system technology will be brought closer to customers that could benefit from knowledge systems, but that in current practice have no access to this technology. Moreover, IBROW3 enables more widespread use of knowledge systems because the costs of the technology may decrease by an order of magnitude. This is in particular interesting for small and medium sized enterprises (SMEs) which currently lack the resources to exploit knowledge-system technology (as opposed to large enterprises).

- A new electronic market is established of knowledge-component providers and customers that converges with the telematics infrastructure and interoperability requirements of the future electronic marketplace [Garcia, 1997].

- Verification of reusable knowledge components will become worthwhile because this only has to be done once. The quality of systems built with such verified components will be, as a result, high [Wielinga *et al.*, 1993, van Harmelen & ten Teije, 1997].

- The technology that will be developed in the project can be applied directly to so-called *Intranets*. Intranets are organization-wide internal networks based on Internet technology.

3 Approach

The global approach to achieve IBROW3's objectives decomposes the problem in different, but related, issues. The overall picture is illustrated in Figure 1. The intelligent broker handles requests for reasoners from various customers. Based on these requests, it accesses different libraries available on the Web and searches them for candidate PSMs, which are adapted and configured into a knowledge system for the customer.

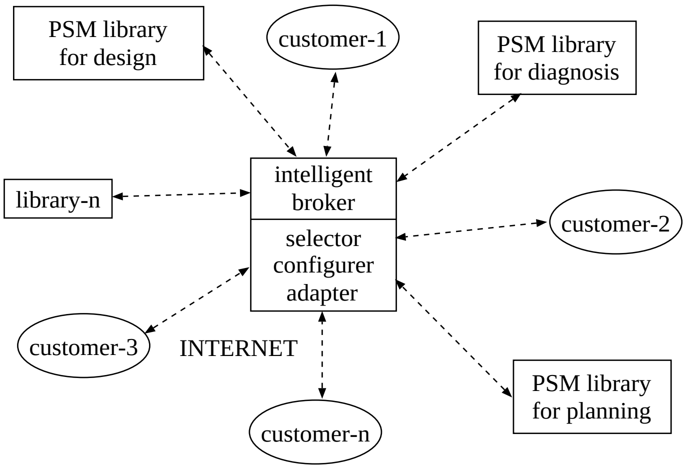

Figure 1: The IBROW3 broker on the Internet.

As a first step, the project will develop a prototype intelligent broker that is able to select, configure and adapt PSMs for a test case, using a PSM library for design [Motta & Zdrahal, 1996]. The approach of IBROW3 divides the objectives into the research topics described below.

PSM-description language For selecting PSMs from a library, the broker needs to reason *about* characteristics of PSMs, for example about their competence [Akkermans *et al.*, 1994, Van de Velde, 1988] and their requirements, aka assumptions [Benjamins *et al.*, 1996, Fensel & Benjamins, 1996, Benjamins & Pierret-Golbreich, 1996, Aamodt *et al.*, 1992]. But there might also be other characteristics relevant for selection, such as the cost of executing a PSM (in terms of execution time or required interaction with a human), its cooperation style (LIFO, LILO), empirical facts of the PSM (how often has it been selected and used successfully), etc. We will develop a PSM-description language for adequately capturing these characteristics: the Universal Problem-solving Method description Language (UPML). This language is independent of specific implementations of PSMs, as it is a meta-description defining the broker's view of the PSMs. In today's knowledge engineering research, there exist concrete proposals for many of the ingredients of such a language [Schreiber *et al.*, 1994a, Angele *et al.*, 1996, Gennari *et al.*, 1996] including the input and output descriptions of PSMs, the relation between their inputs and outputs, and their assumptions and domain requirements. Figure 2 illustrates the ingredients of PSM description language.

The formal underpinning of the language are Abstract Data Types (ADTs). If we want to exchange our methods with other groups, then we have to build

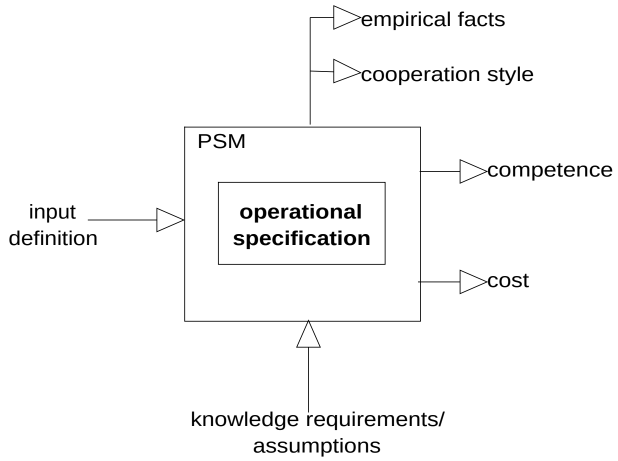

Figure 2: The ingredients of a universal PSM description language.

a wrapper. A KIF wrapper (Knowledge Interchange Format) [Gruber, 1993, Genesereth & Fikes, 1992] would be a good candidate wrapper because KIF is widely used and especially designed to serve as a standard "interlingua"

Ontologies An ontology is a shared and common understanding of some domain that can be communicated across people and computers [Gruber, 1993, van Heijst *et al.*, 1997, Gruber, 1993, Uschold & Gruninger, 1996]. Ontologies can therefore be shared and reused among different applications [Farquhar *et al.*, 1997], which is one of the main reasons why ontologies are popular nowadays. Most existing ontologies are *domain* ontologies, reflecting the fact that they capture (domain) knowledge about the world independently of its use [Guarino, 1995]. However, one can also view the world from a "reasoning" (i.e. use) perspective [Ikeda *et al.*, 1997, Fensel *et al.*, 1997b]. For instance, if we are concerned with diagnosis, we will talk about "hypotheses", "symptoms" and "observations". We say that those terms belong to the *task ontology* of diagnosis. Similarly, we can view the world from a problem-solving point of view. For example, propose & revise sees the world in terms of "states", "state transitions", "preferences" and "fixes" [Fensel *et al.*, 1997b]. These terms are part of the *method ontology* (or PSM ontology) of the propose & revise PSM. Several tasks and PSM ontologies are available including ontologies for design [Chandrasekaran, 1990, Motta & Zdrahal, 1996], diagnosis [Benjamins, 1993, ten Teije, 1997], planning [Barros, L. Nunes de *et al.*, 1996], propose & revise [Zdrahal & Motta, 1995], cover & differentiate [Eshelman, 1988], which can be reused and shared. Ontologies form an essential part of the intelligence of the broker needed for selecting, adapting and configuring PSMs. We refer to these processes as "ontological reasoning". Figure 3 illustrates the relations between the different ontologies.

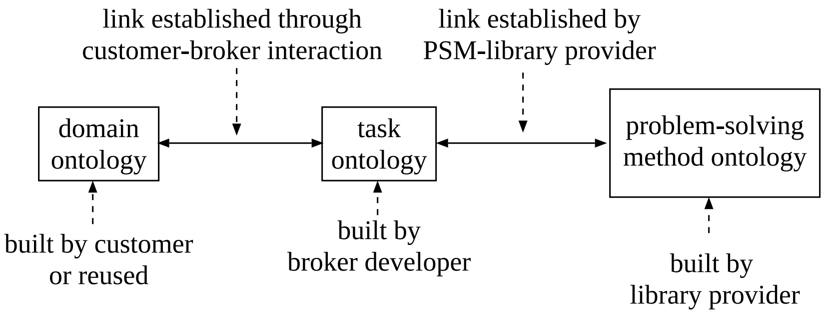

Figure 3: Relations between the different ontologies.

Selecting PSMs by relating task and method ontologies The PSM-description language offers the syntax in which to express the characteristics of the PSMs (e.g., competence and requirements). The broker "speaks" the PSM-description language and selects candidate PSMs for customer requests. Basicly, this is a matching process between a query of the broker (which originally came from the customer) and the characteristics of the PSMs. The first selection of relevant PSMs for the query considers the competence of the PSM. Several types of matches are possible between the query and the PSM [Schumann & Fischer, 1997], two of which are:

- PSM has weaker requirements and stronger competence than what the query requires (safe retrieval),
- PSM has stronger competence than the query asks, regardless of its requirements (unsafe retrieval).

Apart from the matching process, additional knowledge is needed to select suitable PSMs for the tasks of the customers (e.g. diagnosis, design) [Fensel *et al.*, 1997b]. Tasks and PSMs may use different terms which are respectively captured in task and method ontologies. Both task and method ontologies are expressed in the PSM-description language, and before the broker can match two expressions (one in the task ontology and one in the method ontology, the expressions have to be mapped onto each other. It is the responsibility of the library providers to link their method ontologies to task ontologies. The broker exploits then these links before it tries to match the expressions.

Usually, the quality of retrieval is evaluated by two parameters: recall and precision [Salton & McGill, 1983]. The recall gives the ratio between the number of relevant components retrieved and the total number of relevant components in the libraries considered. The precision gives the ratio between the number of relevant components retrieved and the number of retrieved components. If no components are retrieved at all, then either the query can be weakened or the component (the PSM) can be strengthened [Benjamins *et al.*, 1996].

Configuring PSMs by reasoning about task ontologies When the customer provides the broker with a request (formulated as a task to be realized), the broker has to select adequate PSMs to realize the task. In general, it will not be the case that there is one complete PSM that can be retrieved for this task. The broker thus might need to combine different PSMs that together solve the problem [Benjamins, 1995, Puerta *et al.*, 1992, Runkel & Birmingham, 1993, Klinker *et al.*, 1991]. One way of doing this is to decompose the task into subtasks and then look for PSMs to realize these subtasks, etc. Once the relevant PSMs are found, the same knowledge can then be used to compose the problem solver from its comprising submethods. In order to do so, the broker needs knowledge about the particular task under consideration that tells him how to decompose tasks, how the subtasks relate to each other (e.g. control regimes), whether PSMs for the subtasks are mutually exclusive or dependent, etc. Such knowledge is also captured in task ontologies [Benjamins, 1993] (besides the vocabulary of the task). In IBROW3, we build on research dealing with the configuration of PSMs, which can either be performed manually [Puerta *et al.*, 1992, Runkel & Birmingham, 1993] or semi-automatically [Benjamins, 1995, Barros, L. Nunes de *et al.*, 1997, ten Teije *et al.*, 1996, ten Teije, 1997].

Adapting PSMs by relating method and domain ontologies Since PSMs are generic components, they have to be adapted to the particular domain knowledge of a customer before we obtain a knowledge system [Gennari *et al.*, 1994]. This adaptation process basically consists of relating the method ontology to the domain ontology (cf. adapters in [Fensel, 1997b]). That is, it has to be explicated how PSM concepts map onto domain concepts. Today, a huge repository of domain ontologies exists [Farquhar *et al.*, 1997, Fridman-Noy & Hafner, 1997, Guarino & Poli, 1995], and we will reuse one for the test case. It is the responsibility of the customers to establish the links between method and domain ontologies, which will then be used by the broker.

Interoperability protocol The project aims at component reuse in a heterogeneous environment, which means that components implemented in different languages will have to cooperate with each other. Interoperability is thus an important issue. We recall that the objective of the broker is to configure a knowledge system for a customer, in other words, the broker provides the user with a reasoning service. There are different possibilities as where to execute the configured knowledge system, that is, how to provide the reasoning service.

- The broker retrieves the PSMs from the libraries and the customer downloads the problem-solver from the broker. The customer's domain knowledge stays at customer's site.

  - advantages: domain knowledge remains confidential; no network delays
  - disadvantages: customer needs the right implementation platform in order to execute the problem solver; traditional update problems

- The problem solver runs at the broker's site. The broker retrieves the PSMs from the libraries and takes care of interoperability problems (if PSMs are implemented in different languages).

  - advantages: execution is independent of hard/software of customer; no updates needed at the customer's site; there is a central working memory for the broker, so the content of the message remains relatively simple.
  - disadvantages: customer needs to upload his domain knowledge (confidentiality problems).

- The problem solver runs distributively. The broker does not physically retrieve the PSMs from the libraries, but they are executed in the libraries where they are found. Interoperability needs to be taken care of (e.g. CORBA [Orfali *et al.*, 1996]).

  - advantages: execution is independent of hard/software of customer and of broker; no updates needed at broker's site.
  - disadvantages: no central working memory available for the broker, so it will have to be sent by message passing; customer needs to upload his (confidential?) domain knowledge bit by bit (as the execution requires); execution speed depends on speed of network.

We will identify the interoperability requirements needed for the project and we will evaluate existing proposals, like CORBA [Orfali *et al.*, 1996] as an object broker architecture and JavaBeans [JavaSoft, 1996] as a component architecture. Whereas CORBA and JavaBeans aim at *components*, IBROW3 deals with *knowledge components*. We will select an appropriate proposal (with identified deficiencies for dealing with knowledge components), adapt it to our requirements (yielding a first rudimentary version) and apply it to the components of the PSM libraries (see [Gennari *et al.*, 1996], for an example of how to use CORBA for PSMs).

4 Discussion

PSMs play an important role in knowledge-system development based on reusable knowledge components [Benjamins & Fensel, 1998]. However, the current status of building knowledge systems is a centralized process involving domain experts, domain ontologies, PSM libraries and knowledge engineers. In the project described in this paper, we explore the situation of distributed reuse through *customers*, *intelligent brokers* and *providers*. Our hope is that the intelligent broker – which resides on the WWW – takes away much of the burden a knowledge engineer is nowadays faced with. This objective introduces new issues like distributed sharing and reuse, heterogeneous knowledge components, and other problems related to large-scale web-based reuse.

In the USA, a new DARPA[1] project called High-Performance Knowledge Bases (HPKB) has been started in the spring of 1997. There is a significant intersection between the objectives of HPKB and IBROW3, including scaling up knowledge-component reuse, interoperability, Internet.

Acknowledgment

IBROW3 is funded by the European Commission under Esprit, Open LTR, proposal number: 27169. The consortium consists of the Intelligent Systems Lab Amsterdam (ISLA) of the University of Amsterdam, the Netherlands; the Artificial Intelligence Research Institute (IIIA) of the Spanish Council for Scientific Research, Spain; the Knowledge Media Institute (KMI) of The Open University, United Kingdom; and the AIFB institute of the University of Karlsruhe, Germany. IBROW3 is endorsed by an industrial advisory committee whose current members are Daimler-Benz, Germany and the Artificial Intelligence Applications Institute (United Kingdom).

Richard Benjamins is supported by the Netherlands Computer Science Research Foundation with financial support from the Netherlands Organisation for Scientific Research (NWO), and by the European Commission through a Marie Curie Research Grant (TMR).

References

[Aamodt *et al.*, 1992] Aamodt, A., Benus, B., Duursma, C., Tomlinson, C., Schrooten, R., & Van de Velde, W. (1992). Task features and their use in commonkads. Technical Report KADS-II/T1.5/VUB/TR/014/1.0, Free University of Brussels & University of Amsterdam & Lloyd's Register.

[Akkermans *et al.*, 1994] Akkermans, J. M., Wielinga, B. J., & Schreiber, A. T. (1994). Steps in constructing problem-solving methods. In Gaines, B. R. & Musen, M. A., (Eds.), *Proceedings of the 8th Banff Knowledge Acquisition for Knowledge-Based Systems Workshop. Volume 2: Shareable and Reusable Problem-Solving Methods*, pp. 29–1 – 29–21, Alberta, Canada. SRDG Publications, University of Calgary.

[Angele *et al.*, 1996] Angele, J., Decker, S., Perkuhn, R., & Studer, R. (1996). Modeling problem-solving methods in new karl. In Gaines, B. R. & Musen, M. A.,

[1]DARPA URL: http://www.teknowledge.com:80/HPKB/.

(Eds.), *Proceedings of the 10th Banff Knowledge Acquisition for Knowledge-Based Systems Workshop*, pp. 1.1–1.18, Alberta, Canada. SRDG Publications, University of Calgary. http://ksi.cpsc.ucalgary.ca:80/KAW/KAW96/KAW96Proc.html.

[Arens *et al.*, 1993] ARENS, Y., CHEE, C. Y., HSU, C., & KNOBLOCK, G. A. (1993). Retrieving and integrating data from multiple information sources. *International Journal of Intelligent and Cooperative Information Systems*, 2(2):127–158.

[Barros, L. Nunes de *et al.*, 1997] BARROS, L. NUNES DE, HENDLER, J., & BENJAMINS, V. R. (1997). Par-KAP: a knowledge acquisition tool for building practical planning system. In *Proc. of the 15th IJCAI*, pp. 1246–1251, Japan. International Joint Conference on Artificial Intelligence, Morgan Kaufmann Publishers, Inc. Also published in Proceedings of the Ninth Dutch Conference on Artificial Intelligence, NAIC'97, K. van Marcke, W. Daelemans (eds), University of Antwerp, Belgium, pages 137–148.

[Barros, L. Nunes de *et al.*, 1996] BARROS, L. NUNES DE, VALENTE, A., & BENJAMINS, V. R. (1996). Modeling planning tasks. In *Third International Conference on Artificial Intelligence Planning Systems, AIPS-96*, pp. 11–18. American Association of Artificial Intelligence (AAAI).

[Benjamins, 1993] BENJAMINS, V. R. (1993). *Problem Solving Methods for Diagnosis.* PhD thesis, University of Amsterdam, Amsterdam, The Netherlands.

[Benjamins, 1995] 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 & Fensel, 1998] BENJAMINS, V. R. & FENSEL, D. (1998). Problem-solving methods. *International Journal of Human-Computer Studies*, to appear. Special issue on Problem-Solving Methods.

[Benjamins *et al.*, 1996] BENJAMINS, V. R., FENSEL, D., & STRAATMAN, R. (1996). Assumptions of problem-solving methods and their role in knowledge engineering. In Wahlster, W., (Ed.), *Proc. ECAI–96*, pp. 408–412. J. Wiley & Sons, Ltd.

[Benjamins & Pierret-Golbreich, 1996] BENJAMINS, V. R. & PIERRET-GOLBREICH, C. (1996). Assumptions of problem-solving methods. In Shadbolt, N., O'Hara, K., & Schreiber, G., (Eds.), *Lecture Notes in Artificial Intelligence, 1076, 9th European Knowledge Acquisition Workshop, EKAW-96*, pp. 1–16, Berlin. Springer-Verlag.

[Breuker, 1994] BREUKER, J. (1994). Components of problem solving and types of problems. In Steels, L., Schreiber, A. T., & van de Velde, W., (Eds.), *Lecture Notes in Artificial Intelligence, 867, 8th European Knowledge Acquisition Workshop, EKAW-94*, pp. 118–136, Berlin, Germany. Springer-Verlag.

[Chandrasekaran, 1990] CHANDRASEKARAN, B. (1990). Design problem solving: A task analysis. *AI Magazine*, 11:59–71.

[Decker *et al.*, 1997] DECKER, K., SYCARA, K., & WILLIAMSON, M. (1997). Middle-agents for the internet. In *Proc. of the 15th IJCAI*, pp. 578–583, Japan. International Joint Conference on Artificial Intelligence, Morgan Kaufmann Publishers, Inc.

[Duschka & Genesereth, 1997] Duschka, O. M. & Genesereth, M. R. (1997). Infomaster – an information integration tool. In *Proceedings of the International Workshop on Intelligent Information Integration*, Freiburg, Germany.

[Eshelman, 1988] Eshelman, L. (1988). MOLE: A knowledge-acquisition tool for cover-and-differentiate systems. In Marcus, S., (Ed.), *Automating Knowledge Acquisition for Expert Systems*, pp. 37–80. Boston, Kluwer.

[Farquhar *et al.*, 1997] Farquhar, A., Fikes, R., & Rice, J. (1997). The ontolingua server: a tool for collaborative ontology construction. *International Journal of Human-Computer Studies*, 46(May):707–728.

[Fensel, 1997a] Fensel, D. (1997a). An ontology-based broker: making problem-solving method reuse work. In *Proceedings of IJCAI workshop on Problem-Solving Methods*.

[Fensel, 1997b] Fensel, D. (1997b). The tower-of-adapters method for developing and reusing problem-solving methods. In Plaza, E. & Benjamins, V. R., (Eds.), *Knowledge Acquisition, Modeling and Management*, pp. 97–112. Springer-Verlag.

[Fensel & Benjamins, 1996] Fensel, D. & Benjamins, V. R. (1996). Assumptions in model-based diagnosis. In Gaines, B. R. & Musen, M. A., (Eds.), *Proceedings of the 10th Banff Knowledge Acquisition for Knowledge-Based Systems Workshop*, pp. 5.1–5.18, Alberta, Canada. SRDG Publications, University of Calgary. http://ksi.cpsc.ucalgary.ca:80/KAW/KAW96/KAW96Proc.html.

[Fensel *et al.*, 1997a] Fensel, D., Erdmann, M., & Studer, R. (1997a). Ontology groups: Semantically enriched subnets of the www. In *Proceedings of the International Workshop Intelligent Information Integration during the 21st German Annual Conference on Artificial Intelligence*, Freiburg.

[Fensel *et al.*, 1997b] Fensel, D., Motta, E., Decker, S., & Zdrahal, Z. (1997b). Using ontologies for defining tasks, problem-solving methods and their mappings. In Plaza, E. & Benjamins, V. R., (Eds.), *Knowledge Acquisition, Modeling and Management*, pp. 113–128. Springer-Verlag.

[Fridman-Noy & Hafner, 1997] Fridman-Noy, N. & Hafner, C. D. (1997). The state of the art in ontology design. *AI Magazine*, 18(3):53–74.

[Garcia, 1997] Garcia, D. L. (1997). *Electronic enterprises: looking to the future*. PhD thesis, University of Amsterdam, Amsterdam, The Netherlands.

[Genesereth & Fikes, 1992] Genesereth, M. R. & Fikes, R. E. (1992). Knowledge interchange format, version 3.0, reference manual. Technical report, Logic-92-1, Computer Science Dept., Stanford University. http://www.cs.umbc.edu/kse/.

[Gennari *et al.*, 1996] Gennari, J. H., Stein, A. R., & Musen, M. A. (1996). Reuse for knowledge-based systems and CORBA components. In Gaines, B. R. & Musen, M. A., (Eds.), *Proceedings of the 10th Banff Knowledge Acquisition for Knowledge-Based Systems Workshop*, pp. 46.1–46.16, Alberta, Canada. SRDG Publications, University of Calgary. http://ksi.cpsc.ucalgary.ca:80/KAW/KAW96/KAW96Proc.html.

[Gennari *et al.*, 1994] Gennari, J. H., Tu, S. W., Rotenfluh, T. E., & Musen, M. A. (1994). Mapping domains to methods in support of reuse. *International Journal of Human-Computer Studies*, 41:399–424.

[Gruber, 1993] Gruber, T. R. (1993). A translation approach to portable ontology specifications. *Knowledge Acquisition*, 5:199–220.

[Guarino, 1995] Guarino, N. (1995). Formal ontology, conceptual analysis and knowledge representation. *International Journal of Human-Computer Studies*, 43(2/3):625–640.

[Guarino & Poli, 1995] Guarino, N. & Poli, R. (1995). The role of ontology in the information technology. *International Journal of Human-Computer Studies*, 43(5/6):623–965. Special issue on ontology.

[Ikeda *et al.*, 1997] Ikeda, M., Seta, K., & Mizoguchi, R. (1997). Task ontology makes it easier to use authoring tools. In *Proc. of the 15th IJCAI*, pp. 342–347, Japan. International Joint Conference on Artificial Intelligence, Morgan Kaufmann Publishers, Inc.

[JavaSoft, 1996] JavaSoft (1996). The Javabeans^{TM} 1.0 API specification. Technical Report http://java.sun.com/beans, Sun Microsystems.

[Klinker *et al.*, 1991] Klinker, G., Bhola, C., Dallemagne, G., Marques, D., & McDermott, J. (1991). Usable and reusable programming constructs. *Knowledge Acquisition*, 3:117–136.

[McDermott, 1988] McDermott, J. (1988). Preliminary steps towards a taxonomy of problem-solving methods. In Marcus, S., (Ed.), *Automating Knowledge Acquisition for Expert Systems*, pp. 225–255. Boston, Kluwer.

[Motta & Zdrahal, 1996] Motta, E. & Zdrahal, Z. (1996). Parametric design problem solving. In Gaines, B. R. & Musen, M. A., (Eds.), *Proceedings of the 10th Banff Knowledge Acquisition for Knowledge-Based Systems Workshop*, pp. 9.1–9.20, Alberta, Canada. SRDG Publications, University of Calgary. http://ksi.cpsc.ucalgary.ca:80/KAW/KAW96/KAW96Proc.html.

[Neches *et al.*, 1991] Neches, R., Fikes, R. E., Finin, T., Gruber, T. R., Senator, T., & Swartout, W. R. (1991). Enabling technology for knowledge sharing. *AI Magazine*, 12(3):36–56.

[Orfali *et al.*, 1996] Orfali, R., Harkey, D., & Edwards, J., (Eds.) (1996). *The Essential Distributed Objects Survival Guide*. New York, John Wiley & Sons.

[Puerta *et al.*, 1992] Puerta, A., Egar, J., Tu, S., & Musen, M. (1992). A multiple-method shell for the automatic generation of knowledge acquisition tools. *Knowledge Acquisition*, 4:171–196.

[Runkel & Birmingham, 1993] Runkel, J. T. & Birmingham, W. P. (1993). Knowledge acquisition in the small: building knowledge-acquisition tools from pieces. *Knowledge Acquisition*, 5:221–243.

[Salton & McGill, 1983] Salton, G. & McGill, M. J. (1983). *Introduction to Modern Information Retrieval.* New York, McGraw-Hill.

[Schreiber *et al.*, 1994a] Schreiber, A. T., Wielinga, B. J., Akkermans, J. M., Van de Velde, W., & Anjewierden, A. (1994a). CML: The CommonKADS conceptual modelling language. In Steels, L., Schreiber, A. T., & Van de Velde, W., (Eds.), *A Future for Knowledge Acquisition. Proceedings of the 8th European Knowledge Acquisition Workshop EKAW'94*, pp. 1–25, Berlin/Heidelberg. Springer-Verlag.

[Schreiber *et al.*, 1994b] Schreiber, A. T., Wielinga, B. J., de Hoog, R., Akkermans, J. M., & Van de Velde, W. (1994b). CommonKADS: A comprehensive methodology for KBS development. *IEEE Expert*, 9(6):28–37.

[Schumann & Fischer, 1997] Schumann, J. & Fischer, B. (1997). NORA/HAMMR making deduction-based software component retrieval practical. In *12th IEEE International Conference on Automated Software Engineering*, pp. 246–254. IEEE Computer Society.

[Speel & Aben, 1997] Speel, P.-H. & Aben, M. (1997). Applying a library of problem solving methods on a real-life task. *International Journal of Human-Computer Studies*, 46(May):627–652.

[ten Teije, 1997] ten Teije, A. (1997). *Automated Configuration of Problem Solving Methods in Diagnosis.* PhD thesis, University of Amsterdam, Amsterdam, The Netherlands.

[ten Teije *et al.*, 1996] ten Teije, A., van Harmelen, F., Schreiber, A. T., & Wielinga, B. (1996). Construction of problem-solving methods as parametric design. In Gaines, B. R. & Musen, M. A., (Eds.), *Proceedings of the 10th Banff Knowledge Acquisition for Knowledge-Based Systems Workshop*, pp. 12.1–12.20, Alberta, Canada. SRDG Publications, University of Calgary. http://ksi.cpsc.ucalgary.ca:80/KAW/KAW96/KAW96Proc.html.

[Uschold & Gruninger, 1996] Uschold, M. & Gruninger, M. (1996). Ontologies: principles, methods, and applications. *Knowledge Engineering Review*, 11(2):93–155.

[Valente & Löckenhoff, 1993] Valente, A. & Löckenhoff, C. (1993). Organization as guidance: A library of assessment models. In *Proceedings of the Seventh European Knowledge Acquisition Workshop (EKAW'93)*, pp. 243–262.

[Van de Velde, 1988] Van de Velde, W. (1988). Inference stucture as a basis for problem solving. In Kodratoff, Y., (Ed.), *Proceedings of the 8th European Conference on Artificial Intelligence*, pp. 202–207, London. Pitman.

[van Harmelen & ten Teije, 1997] van Harmelen, F. & ten Teije, A. (1997). Validation and verification of conceptual models of diagnosis. In vanthienen, J. & van Harmelen, F., (Eds.), *EUROVAV'97: Proceedings of the Fourth European Symposium on the Validation and Verification of Knowledge Based Systems*, Leuven, Belgium.

[van Heijst *et al.*, 1997] van Heijst, G., Schreiber, A. T., & Wielinga, B. J. (1997). Using explicit ontologies in KBS development. *International Journal of Human-Computer Studies*, 46(2/3):183–292.

[Wiederhold, 1996] Wiederhold, G. (1996). Intelligent integration of information. *Journal of Intelligent Information Systems*. Special issue on Intelligent Integration of Information.

[Wiederhold & Genesereth, 1997] Wiederhold, G. & Genesereth, M. (1997). The conceptual basis for mediation services. *IEEE Intelligent Systems*, 12:38–47.

[Wielinga *et al.*, 1993] Wielinga, B. J., Akkermans, J. M., & Schreiber, A. T. (1993). Validation and verification of knowledge models. In Cardenosa, J. & Meseguer, P., (Eds.), *EUROVAV'93: Proceedings of the European Symposium on the Validation and Verification of Knowledge Based Systems*, pp. 29–50, Palma de Mallorca, Spain. Universidad Politecnica de Madrid.

[Zdrahal & Motta, 1995] Zdrahal, Z. & Motta, E. (1995). An in-depth analysis of propose & revise problem solving methods. In Gaines, B. R. & Musen, M. A., (Eds.), *Proceedings of the 9th Banff Knowledge Acquisition for Knowledge-Based Systems Workshop*, Alberta, Canada. SRDG Publications, University of Calgary.