

# Algorithm and Tool for Automated Ontology Merging and Alignment

Natalya Fridman Noy and Mark A. Musen

Stanford Medical Informatics, Stanford University, Stanford, CA 94305-5479  
{noy, musen}@smi.stanford.edu

## Abstract

Researchers in the ontology-design field have developed the content for ontologies in many domain areas. Recently, ontologies have become increasingly common on the World-Wide Web where they provide semantics for annotations in Web pages. This distributed nature of ontology development has led to a large number of ontologies covering overlapping domains. In order for these ontologies to be reused, they first need to be merged or aligned to one another. The processes of ontology alignment and merging are usually handled manually and often constitute a large and tedious portion of the sharing process. We have developed and implemented PROMPT, an algorithm that provides a semi-automatic approach to ontology merging and alignment. PROMPT performs some tasks automatically and guides the user in performing other tasks for which his intervention is required. PROMPT also determines possible inconsistencies in the state of the ontology, which result from the user's actions, and suggests ways to remedy these inconsistencies. PROMPT is based on an extremely general knowledge model and therefore can be applied across various platforms. Our formative evaluation showed that a human expert followed 90% of the suggestions that PROMPT generated and that 74% of the total knowledge-base operations invoked by the user were suggested by PROMPT.

## 1 Ontologies in AI and on the Web

Ontologies today are available in many different forms: as artifacts of a tedious knowledge-engineering process, as information that was extracted automatically from informal electronic sources, or as simple “light-weight” ontologies that specify semantic relationships among resources available on the World-Wide Web (Brickley and Guha 1999). But what does a user do when he finds several ontologies that he would like to use but that do not conform to one another? The user must establish correspondences among the source ontologies, and to determine the set of overlapping concepts, concepts that are similar in meaning but have different names or structure, concepts that are unique to each of the sources. This work must be done regardless of whether the ultimate goal is to create a single coherent ontology that includes the information from all the sources (**merging**) or if the sources must be made consistent and coherent with one another but kept separately (**alignment**).

Currently the work of mapping, merging, or aligning ontologies is performed mostly by hand, without any tools

to automate the process fully or partially (Fridman Noy and Musen 1999). Our participation in the ontology-alignment effort within DARPA's High-Performance Knowledge-Bases project (Cohen et al. 1999) was a strong motivation for developing semi-automated specialized tools for ontology merging and alignment. Several teams developed ontologies in the domain of military planning, which then needed to be aligned to one another. We found the experience of manually aligning the ontologies to be an extremely tedious and time-consuming process. At the same time we noticed many steps in the process that could be automated, many points where a tool could make reasonable suggestions, and many conflicts and constraint violations for which a tool could check.

We developed a formalism-independent algorithm for ontology merging and alignment—PROMPT (formerly SMART)—which automates the process as much as possible. Where an automatic decision is not possible, the algorithm guides the user to the places in the ontology where his intervention is necessary, suggests possible actions, and determines the conflicts in the ontology and proposes solutions for these conflicts. We implemented the algorithm in an interactive tool based on the Protégé-2000 knowledge-modeling environment (Fridman Noy et al. 2000). Protégé-2000 is an ontology-design and knowledge-acquisition tool with an OKBC-compatible (Chaudhri et al. 1998) knowledge model, which allows domain experts (and not necessarily knowledge engineers) to develop ontologies and perform knowledge acquisition. We have evaluated PROMPT, comparing its performance with the human-expert performance and with the performance of another ontology-merging tool.

## 2 Related Work

Researchers in various areas of computer science have worked on automatic or tool-supported merging of ontologies (or class hierarchies, or object-oriented schemas, or database schemas—the specific terminology varies depending on the field). However, both automatic merging of ontologies and creation of tools that would guide the user through the process and focus his attention on the likely points for actions are in early stages. In this section, we give an overview of some of the existing approaches to merging and alignment in the fields of ontology design, object-oriented programming, and heterogeneous databases.

## 2.1 Ontology Design and Integration

Researchers working on tools for ontology merging have expended the greatest amount of effort on finding mostly syntactic matches among concepts in the source ontologies. Such systems rely on dictionaries to determine synonyms, evaluate common substrings, consider concepts whose documentation share many uncommon words, and so on (see, for example, (Chapulsky et al. 1997) or the Scalable Knowledge Composition project (Wiederhold and Jannik 1999)). These approaches, however, do not take into account the semantics of concepts, the structure of the ontology, or the steps the user takes during merging.

Ontomorph (MacGregor et al. 1999) defines a set of transformation operators that can be applied to an ontology. A human expert then uses the initial list of matches and the source ontologies to define a set of operators that need to be applied to the source ontologies in order to resolve differences between them, and Ontomorph applies the operators. Therefore, aggregate operations can be performed in a single step. However, a human expert receives no guidance except for the initial list of matches.

Chimaera (McGuinness et al. 2000) is an interactive merging tool based on Ontolingua ontology editor (Farquhar et al. 1996). Chimaera allows a user to bring together ontologies developed in different formalisms. The user can request an analysis or guidance from Chimaera at any point during the merging process. The tool will then point him to the places in the ontology where his attention is required. In its suggestions, Chimaera mostly relies on which ontology the concepts came from and, for classes, on their names. For example, Chimaera will point a user to a class in the merged ontology that has two slots derived from different source ontologies, or that has two subclasses that originated in different ontologies. Chimaera leaves the decision of *what* to do entirely to the user and does not make any suggestions itself. The only taxonomic relation that Chimaera considers is the subclass–superclass relation. We discuss the differences between Chimaera and PROMPT in more detail when we discuss the results of our comparative evaluation in Section 7.

Medical vocabularies provide a rich field for testing of various ontology-merging paradigms. Not only is there a wide variety of large-scale sources, but also medicine is a field where standard vocabularies change constantly. Oliver and colleagues explored representation of change in medical terminologies using a frame-based knowledge-representation system. The authors compiled a list of change operations that are relevant for the domain of medical terminologies, and developed a tool to support these operations. However, the user has to do all the operations manually; there is no automated help or guidance.

## 2.2 Object-Oriented Programming

Subject-oriented programming (SOP) (Harrison and Ossher 1993)—an area of object-oriented programming—supports building object-oriented systems through composition of

subjects. **Subjects** are collections of classes that represent subjective views of, possibly, the same universe that need to be combined. The formal theory of subject-oriented composition defines a set of possible composition rules, these rules’ semantics, and how the rules work with one another. Interactive tools for subject-oriented composition are currently under development. However, the SOP approach relies more heavily on the operational methods associated with classes rather than on declarative relations among classes and slots. Alignment (as opposed to merging) is extremely uncommon in composition of object-oriented hierarchies, whereas it is common in ontology design.

## 2.3 Integration of Heterogeneous Databases

Developers of heterogeneous databases have dealt with issues of bringing various information sources together. These issues include merging or mediating between relational and object-oriented databases, varying formats from different database vendors, varying underlying schemas or basic assumptions. The common theme in the research on heterogeneous databases, however, is to bridge the gaps on demand by creating an extra mediation layer. The approaches include:

- Develop *mediators* —a facility for answering queries about an information source. Each source may have a *wrapper* that works as an interface between the mediator and the source itself. TSIMMIS is an example of such mediator-based systems.
- Define a *common data model* and then map the source and target to it .
- Specify a set of *matching rules* that directly translate between source and target .

Database are usually integrated at the syntactic rather than semantic level. And, in fact, ontologies are more semantically complex and are often larger than database schemas.

## 3 Knowledge Model

We now turn to the discussion of the PROMPT ontology-merging and alignment algorithm. We start with the description of the knowledge model underlying PROMPT. The knowledge model is frame-based and it is designed to be compatible with OKBC (Chaudhri et al. 1998). At the top level, there are classes, slots, facets, and instances:

- **Classes** are collections of objects that have similar properties. Classes are arranged into a subclass–superclass hierarchy with multiple inheritance. Each class has slots attached to it. Slots are inherited by the subclasses.
- **Slots** are named binary relations between a class and either another class or a primitive object (such as a string or a number). Slots attached to a class may be further constrained by facets.
- **Facets** are named ternary relations between a class, a slot, and either another class or a primitive object. Facets

may impose additional constraints on a slot attached to a class, such as the cardinality or value type of a slot.

- **Instances** are individual members of classes.

These definitions are the only restrictions that we impose on the input ontologies for PROMPT. Since this knowledge model is extremely general, and many existing knowledge-representation systems have knowledge models compatible with it, the solutions to merging and alignment produced by PROMPT can be applied over a variety of knowledge-representation systems.

## 4 The PROMPT Algorithm

Figure 1 illustrates the PROMPT ontology-merging and ontology-alignment algorithm. PROMPT takes two ontologies as input and guides the user in the creation of one merged ontology as output. First PROMPT creates an initial list of matches based on class names. Then the following cycle happens: (1) the user triggers an operation by either selecting one of PROMPT’s suggestions from the list or by using an ontology-editing environment to specify the desired operation directly; and (2) PROMPT performs the operation, *automatically* executes additional changes based on the type of the operation, generates a list of *suggestions* for the user based on the structure of the ontology around the arguments to the last operation, and determines *conflicts* that the last operation introduced in the ontology and finds possible solutions for those conflicts.

Since there are several research groups working on methods for determining linguistic similarity among concept names (see Section 2.1), a specific implementation of the PROMPT algorithm will use whatever measure of linguistic similarity among concept names is appropriate. In our Protégé-based implementation (Section 5), we use Protégé component-based architecture to allow the user to plug in any term-matching algorithm. In PROMPT, we start with the linguistic-similarity matches for the initial comparison, but concentrate on finding clues based on the structure of the ontology and user’s actions.

The following is at the heart of our approach: We identify a set of knowledge-base operations for ontology merging or alignment. For each operation in this set, we




Figure 1. The flow of PROMPT algorithm. The gray boxes indicate the actions performed by PROMPT. The white box indicates the action performed by the user.

define (1) changes that PROMPT performs automatically, (2) new suggestions that PROMPT presents to the user, and (3) conflicts that the operation may introduce and that the user needs to resolve. When the user invokes an operation, PROMPT creates members of these three sets based on the arguments to the specific invocation of the operation.

The set of ontology-merging operations that we identified includes both the operations that are normally performed during traditional ontology editing and the operations specific to merging and alignment, such as:

- merge classes,
- merge slots,
- merge bindings between a slot and a class,
- perform a deep copy of a class from one ontology to another (includes copying all the parents of a class up to the root of the hierarchy and all the classes and slots it refers to),
- perform a shallow copy of a class (just the class itself, and not its parents or the classes and slots it refers to).

We identified the following conflicts that may appear in the merged ontology as the result of these operations:

- name conflicts (more than one frame with the same name),
- dangling references (a frame refers to another frame that does not exist),
- redundancy in the class hierarchy (more than one path from a class to a parent other than root),
- slot-value restrictions that violate class inheritance.

Both lists grow as we gain more experience.

For example, suppose the user is merging two ontologies and he performs a `merge-classes` operation for two classes A and B to create a new class M. PROMPT then performs the following actions:

- For each slot S that was attached to A and B in the original ontologies, attach the slot to M with the same value type and other facets. If S did not exist in the merged ontology, create S.
- For each superclass of A and B that has been previously copied into the merged ontology, make that copy a superclass of M (thus restoring the original relation). Do the same for subclasses.
- For each class C in the original ontologies to which A and B referred (that is, for each superclass, subclass, slot value, and class restricting a slot value of A and B), if C has not been copied to the merged ontology, suggest that it is copied to the merged ontology.
- For each class C that was a facet value for A or B and that has not been copied to the merged ontology, declare a dangling-reference conflict.
- For each pair of slots for M that have linguistically similar names, suggest that the slots are merged. Later, if the user chooses to merge the slots, suggest that the classes restricting the values of these slots, are merged as well.
- For each pair of superclasses and subclasses of M that have linguistically similar names, suggest that they are merged: these classes have similar names and, in addition, they were both superclasses (or subclasses) for A and B, which the user declared to be similar.

- Check for redundancy in the parent hierarchy for  $M$ : If there is more than one path to any parent of  $M$  (other than the root of the hierarchy), suggest that one of  $M$ 's parents is removed.

Note, that PROMPT bases most of the decisions in the preceding list on the internal structure of the concepts and their position in the ontology and not syntax.

When we describe the tool based on the algorithm in the next section, we outline a few other actions that can be taken after each operation, such as rearranging the current list of suggestions to maintain the user's focus.

## 5 Protégé-based PROMPT Tool

We implemented the PROMPT ontology-merging algorithm as an extension to Protégé-2000—the latest in a series of knowledge-acquisition tools developed in our laboratory (Grosso et al. 1999). Protégé-2000 is a knowledge-base development tool, which is designed to make it easier for domain experts to create and maintain knowledge bases. Protégé-2000 uses direct-manipulation techniques for ontology editing. The Protégé-2000 component-based architecture allows users to add new features by developing plug-ins—applications that use Protégé as the knowledge-base access layer and that use Protégé graphical user interface to create the knowledge base itself. PROMPT is such a plug-in for merging two source ontologies in Protégé-2000 into one coherent ontology. The tool and a sample interaction with it have been described in detail elsewhere (Fridman Noy and Musen 1999). We will outline a few of its features here.

**Setting the preferred ontology.** It often happens, that the source ontologies are not equally important or stable, and that the user would like to resolve all the conflicts in favor of one of the source ontologies. We allow the user to designate one of the ontologies as **preferred**. When there is a conflict between values, instead of presenting the conflict to the user for resolution, the system resolves the conflict automatically.

**Maintaining the user's focus.** Suppose a user is merging two large ontologies and is currently working in one content area of the ontology. We believe that the system's suggestions that the user sees first should relate to the frames in the same area of the ontology in which the user is working and should not force him to change focus to a different part of the ontology. PROMPT maintains the user's focus by rearranging its lists of suggestions and conflicts and presenting first the items that include frames related to the arguments of the latest operations.

**Providing feedback to the user.** For each of its suggestions, PROMPT presents a series of explanations, starting with why it suggested the operation in the first place. If PROMPT later changes the operation placement in the suggestions list, it augments the explanation with the information on why it moved the operation.

**Logging and reapplying the operations.** The process of ontology merging or alignment is not a one-time exercise. After the user has merged or aligned ontologies

and perhaps has even developed an application based on the result, the source ontologies may change. This scenario is particularly likely for distributed ontologies developed by independent users. Ideally, reapplication of the merging or alignment process to the changed ontologies should be almost automatic. PROMPT logs *knowledge-level* ontology-merging and editing operations. If the original ontologies change, the user only needs to make adjustments to the operations in the log that explicitly refer to the changed frames and the system can then reapply the operations automatically to merge the original ontologies again.

## 6 Evaluation

We have evaluated PROMPT formatively. We measured the quality of its suggestions, measured its utility, and compared it to another ontology-merging tool. We performed three controlled experiments, in which human experts merged two ontologies using Protégé-2000 with PROMPT, generic Protégé-2000, and Chimaera (see Section 2.1).

The two source ontologies were the same for all the experiments: (1) the ontology for the unified problem-solving method development language (Gennari et al. 1998; Fensel et al. 1999) and (2) the ontology for the method-description language (Gennari et al. 1998). Both ontologies describe reusable problem-solving methods, and merging them was a real-life task in our laboratory. The source ontologies contained the total of 134 class and slot frames, and the resulting merged ontology had 117 frames.

All the testers had unlimited time to complete their tasks and we did not compare the actual rate at which the experts performed the tasks: Each merging process was performed by one expert, and the impact of the individual differences on any rate data is extremely significant. These test are preliminary and we plan to perform extensive testing with more subjects using the same sources and protocols later.

### 6.1 Quality of PROMPT's Suggestions

In the first experiment, human experts who were initially unfamiliar with PROMPT, used Protégé-2000 augmented with PROMPT to merge the two source ontologies. We evaluated the quality of PROMPT's suggestions by measuring the following: (1) how many of the PROMPT's suggestions the human experts decided to follow when merging the source ontologies, and (2) how many of the conflict-resolution strategies that PROMPT proposed were followed by the experts. We present the average value among the experts.

Our results showed that the human experts followed 90% of PROMPT's suggestions. The experts followed 75% of the conflict-resolution strategies that PROMPT proposed. During the merging process, PROMPT suggested 74% of the total knowledge-base operations that the users invoked.

## 6.2 PROMPT versus Generic Protégé-2000

We performed an ablation experiment to determine the value that PROMPT adds to a generic knowledge-editing environment. Two experts initially unfamiliar with PROMPT merged the same source ontologies: One expert used Protégé-2000 augmented with PROMPT and the other used the generic version of the Protégé-2000 ontology editor. We compared the contents of the resulting merged ontologies and the number of explicit knowledge-base operations that each user had to specify.

The merged ontologies, which the two experts produced, were quite similar: there was only one difference in class hierarchy, and a number of minor differences in slot names and their types. The user who was using the generic Protégé-2000 was able to find all the classes that should have been merged. However, using the generic Protégé required performing and explicitly specifying 60 knowledge-base operations. Since PROMPT generated and suggested most of the necessary knowledge-base operations, the PROMPT user needed to specify explicitly only 16 operations.

## 6.3 PROMPT versus Chimaera

In the third experiment, we compared the performance of PROMPT and Chimaera, which is the tool closest to PROMPT in the limited set of existing ontology-merging tools (see Section 2.1).<sup>1</sup> We used PROMPT and Chimaera to merge the same source ontologies and we executed exactly the same sequence of merging steps in each of the tools. The executed operations included merging both slots and classes. After each step, we compared the set of new suggestions that the two systems generated. We used a human expert who has previously merged the two ontologies manually to judge whether the suggestions that the systems produced were correct or not.

PROMPT had 30% more correct suggestions than Chimaera did. Suggestions from Chimaera constituted a proper subset of PROMPT's suggestions. 20% of Chimaera's correct suggestions were roughly the same as PROMPT's. The remaining 80% were much less specific than the corresponding PROMPT's suggestions: Chimaera pointed to the class in the ontology where action was required, and PROMPT suggested a specific action (or alternative actions), which were required for that frame.

## 7 Discussion

Our results in the first two experiments demonstrated that a human expert agreed with a very large fraction of both the suggestions and the conflict-resolution strategies that PROMPT produced. PROMPT was able to perform a large number of merging operations on its own (or with simple "approval" of a human expert), thus saving the expert time and effort.

<sup>1</sup> Chimaera software is available at <http://www.ksl.Stanford.EDU/software/chimaera/>

We intentionally chose the source ontologies that were relatively small and uncontroversial to merge in order to factor out subjective differences in the experts' opinions. However, choosing the source ontologies in this way made it almost inevitable that the two resulting ontologies in the experiment that compared Protégé-2000 augmented with PROMPT and generic Protégé-2000 would be similar. This approach allowed us to compare the number of operations that needed to be generated but did not allow us to compare the quality of results. We believe that for larger ontologies, when it is harder for a human expert to track down all the frames that need merging, the results would be different. We plan to perform experiments to test that in the future.

There is a significant difference in the way Chimaera and PROMPT guide the user in the merging process: PROMPT presents a list of *specific* operations that it suggests to the user. Chimaera points to the classes where it determines that the user needs to do something, but it does not specify *what* exactly the user needs to do there. For instance, in the example in Figure 2, Chimaera suggested that the class `Ontology` requires the user's attention because some of the slots of that class came from different source ontologies. It did not specify which slots the user needs to consider and what he needs to do. In the corresponding case, PROMPT suggested that the `axioms` slots, which came from two different ontologies, must be merged. For classes with large number of slots, PROMPT's suggestion, which lists the slots, would be more helpful than simply suggesting that there are slots in this large list that require attention. On the other hand, the user may get overwhelmed if there are too many specific suggestions and in that case Chimaera's approach may be better.

We now plan to continue experimenting with ontology-merging and alignment to define more heuristics that will allow us to automate a larger part of the merging and alignment process. We will also extend our approach to




Figure 2. Differences in suggestions between (a) PROMPT and (b) Chimaera. PROMPT suggests that the user merges two specific slots (*axioms*) of the *Ontology* class. Chimaera points the user to the *Ontology* class that has two slots originating from different ontologies without specifying which slots the user needs to consider

consider standard OKBC facets, such as maximum and minimum cardinality, inverse, and so on. We will include instances of classes, as well as axioms that define additional constraints on frames in the ontology.

## 8 Conclusions

We have described a general approach to ontology merging and alignment. We presented PROMPT—an algorithm for semi-automatic merging and alignment. We discussed strategies that PROMPT uses to guide a user to the next possible point of merging or alignment, to suggest what operations should be performed there, and to perform certain operations automatically. The strategies and algorithms described in this paper are based on a general OKBC-compliant knowledge model. Therefore, these results are applicable to a wide range of knowledge-representation and ontology-development systems. We extended Protégé knowledge-modeling environment with a tool based on the algorithm and performed an empirical evaluation of the tool. Our results showed that PROMPT was very effective in providing suggestions: A human expert followed 90% of PROMPT's suggestions. PROMPT was also very effective in its coverage: 74% of the knowledge-based operations invoked by the user were suggested initially by PROMPT.

## Acknowledgments

We are extremely grateful to Monica Crubézy, Ray Ferguson and Samson Tu for participating in the evaluation experiments and for the valuable feedback on the earlier drafts of this paper. This work was supported in part by the grants 5T16 LM0733 and 892154 from the National Library of Medicine, by a grant from Spawar, and by a grant from FastTrack Systems, Inc.

## References

- Abiteboul, S., Cluet, S. and Milo, T. (1997). Correspondence and Translation for Heterogeneous Data. In: *Proceedings of the International Conference on Database Theory (ICDT)*.
- Brickley, D. and Guha, R.V. (1999). Resource Description Framework (RDF) Schema Specification. Proposed Recommendation, World Wide Web Consortium: <http://www.w3.org/TR/PR-rdf-schema>.
- Chapulsky, H., Hovy, E. and Russ, T. (1997). Progress on an Automatic Ontology Alignment Methodology.
- Chaudhri, V.K., Farquhar, A., Fikes, R., Karp, P.D. and Rice, J.P. (1998). OKBC: A Programmatic Foundation for Knowledge Base Interoperability. In: *Proceedings of the Fifteenth National Conference on Artificial Intelligence (AAAI-98)*, Madison, Wisconsin, AAAI Press.
- Cohen, P., Schrag, R., Jones, E., Pease, A., Lin, A., Starr, B., Gunning, D. and Burke, M. (1999). The DARPA High-Performance Knowledge Bases Project. *AI Magazine* **19**(4): 25-49.
- Farquhar, A., Fikes, R. and Rice, J. (1996). The Ontolingua Server: a Tool for Collaborative Ontology Construction. In: *Proceedings of the Tenth Knowledge Acquisition for Knowledge-Based Systems Workshop*, Banff, Canada.
- Fensel, D., Benjamins, V.R., Motta, E. and Wielinga, R. (1999). UPML: A Framework for knowledge system reuse. In: *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI-99)*, Stockholm, Sweden.
- Fridman Noy, N., Ferguson, R.W. and Musen, M.A. (2000). The knowledge model of Protégé-2000: combining interoperability and flexibility. Stanford Medical Informatics Technical Report, Stanford University.
- Fridman Noy, N. and Musen, M.A. (1999). SMART: Automated Support for Ontology Merging and Alignment. In: *Proceedings of the Twelfth Banff Workshop on Knowledge Acquisition, Modeling, and Management*, Banff, Alberta.
- Garcia-Molina, H., Papakonstantinou, Y., Quass, D., Rajaraman, A., Sagiv, Y., Ullman, J., Vassalos, V. and Widom, J. (1997). The TSIMMIS approach to mediation: Data models and Languages. *Journal of Intelligent Information Systems*.
- Gennari, J.H., Grosso, W. and Musen, M.A. (1998). A method-description language: An initial ontology with examples. In: *Proceedings of the Eleventh Banff Knowledge Acquisition for Knowledge-Bases Systems Workshop*, Banff, Canada.
- Harrison, W. and Ossher, H. (1993). Subject-Oriented Programming (A Critique of Pure Objects). In: *Proceedings of the Conference on Object-Oriented Programming: Systems, Languages, and Applications (OOPSLA'93)*, Washington, DC, ACM Press.
- MacGregor, R., Chalupsky, H., Moriarty, D. and Valente, A. (1999). Ontology Merging with OntoMorph. <http://reliant.teknowledge.com/HPKB/meetings/meet040799/Chalupsky/index.htm>
- McGuinness, D.L., Fikes, R., Rice, J. and Wilder, S. (2000). An Environment for Merging and Testing Large Ontologies. In: *Proceedings of the Seventh International Conference on Principles of Knowledge Representation and Reasoning (KR2000)*, Breckenridge, Colorado.
- Milo, T. and Zohar, S. (1998). Using Schema Matching to Simplify Heterogeneous Data Translation. In: *Proceedings of the 24th International Conference on Very Large Data Bases*, New York City, Morgan Kaufmann.
- Oliver, D.E., Shahar, Y., Shortliffe, E.H. and Musen, M.A. (1999). Representation of Change in controlled medical terminologies. *Artificial Intelligence in Medicine* **15**: 53-76.
- Ossher, H., Kaplan, M., Katz, A., Harrison, W. and Kruskal, V. (1996). Specifying Subject-Oriented Composition. *Theory and Practice of Object Systems* **2**(3): 179-202.
- Wiederhold, G. (1992). Mediators in the architecture of future information systems. *IEEE Computer* **25**(3): 38-49.
- Wiederhold, G. and Jannik, J. (1999). Composing Diverse Ontologies. In: *Proceedings of the IFIP Working Group on Database, 8th Working Conference on Database Semantics (DS-8)*, Rotorua, New Zealand.