Computer Assisted Ontology clustering for Knowledge sharing

Valentina A.M. Tamma[1], *Pepijn R.S. Visser*[2], *Donato Malerba*[3] *and Dean M. Jones*[4]

[1]Department of Computer Science, The University of Liverpool, Chadwick Building, Peach Street, Liverpool, L69 7ZF, UK, tel: +44-151-794 3706, fax: +44-151-794 3715, email: valli@csc.liv.ac.uk
[2]SAP AG Entw. HR - HRMS Netherlands, PO Box 1461, D-69190 Walldorf, Germany, tel: +49-6227-76 47 45, fax: +49-6227-77 47 45, email: Pepijn.Visser@sap.com
[3]Dipartimento di Informatica, Universita' degli Studi di Bari, Via Orabona, 4, 70126 Bari, Italy, tel: +39-080-544 32 69, fax: +39-080-544 31 96, email: malerba@di.uniba.it
[4]BT Adastral Park, Martlesham Heath, Ipswich, Suffolk, IP5 3RE, UK, tel: +44-1473-605 487 fax: +44-1473-605 487, email: dean.m.jones@bt.com

1. Introduction

The Internet has made it possible to access huge amounts of information from different places all over the world. This possibility has stimulated a growing demand for understanding how to integrate multiple and heterogeneous knowledge sources. This field of research aims to recognise and combine relevant knowledge so as to provide a richer understanding of a particular domain. The integration is particularly valuable if it enables the communication between different sources while allowing them to maintain their autonomy. The problem of integrating heterogeneous resources has been addressed in the literature although relatively few research efforts have focused on heterogeneity between resources other than databases (such as knowledge bases and constraint solvers). Dealing with knowledge sources of diverse nature requires the development of specific techniques. *Ontologies* have been developed to make knowledge sharing and reuse possible across domains and tasks (Gómez Pérez and Benjamins, 1999). Their role is to make (differences in) semantics explicit. An *ontology* is "a logical theory accounting for the intended meaning of a formal vocabulary, i.e. its ontological commitment to a particular conceptualisation of the world" (Guarino, 1998). Ontologies can provide a commonly agreed understanding of a domain that can be reused and shared across several applications. Among the various approaches for the integration of heterogeneous sources presented in the literature we focus our attention on the structure of multiple shared ontologies as presented in (Visser and Tamma, 1999). This architecture aggregates multiple shared ontologies into *clusters*, so as to obtain a structure that is able to reconcile different types of heterogeneity and is also intended to be more convenient to implement and give better prospects for maintenance and scaling. Furthermore, such a structure is thought to avoid information loss when performing translations between diverse resources. This paper discusses the feasibility of semi-automatic ontology clustering in order to obtain such a structure. More particularly, the paper investigates the possibility of adapting some *machine learning* (ML) techniques that have been developed for the unsupervised learning problem, such as conceptual clustering (Michalski and Stepp, 1983) in order to build both clusters of ontologies and the intermediate ontologies describing them in a semi-automatic fashion.

The paper is structured as follows: section 2 highlights different approaches for knowledge sharing in the literature and introduces our approach. Section 3 presents the ontology clustering architecture and section 4 describes the knowledge model used to represent ontologies. Section 5 presents an overview of the ML techniques while section 6 investigates the problems that arise when trying to cluster ontologies in a semi-automatic fashion. Finally, in section we 7 draw conclusions and present open problems and future work.

2. **Multiple shared ontologies for knowledge sharing**

All approaches for knowledge sharing are based on certain *mapping functions* performing the translations between the ontologies (shared or not). Concepts can be shared between different resources if an appropriate mapping function can be found that translates a concept understood by one resource into a concept that is understood by another resource. This is the minimal requirement for two resources to share knowledge.

The integration of heterogeneous sources can be accomplished without an intermediate ontology in a so-called ’one-to-one’ approach. In such an approach each ontology has an associated set of translating functions to allow the communication with the other ontologies. This technique is used in the OBSERVER system (Mena *et al.*, 1996), where translations of concepts from a source ontology into a destination ontology are accomplished by defining appropriate mapping functions. The approach seems feasible only if there are a few ontologies (resources) as it is not easy scalable.

Many architectures to integrate resources comprise a single shared ontology, an example is given by InfoSleuth, (Bayardo, *et al.*, 1997) and by the first KRAFT architecture (Gray *et al.*, 1997). Whether such an approach is feasible depends on the case at hand (Shave, 1997; Visser and Tamma, 1999; Visser and Cui, 1998). In contrast to an approach in which all resources share one ontology we here propose to locate shared concepts in multiple but smaller shared ontologies. This approach, which is thought to be more flexible and scalable, is referred to as ontology-based resource clustering, or shortly, *ontology clustering* (Shave, 1997). Resources no longer commit to one comprehensive ontology but they are clustered together on the basis of their similarities in conceptualising their domains. This idea has been used in a small experiment conducted with a telecommunications business process (Visser and Cui, 1998). More generally, multiple-ontology architectures are currently being investigated by our team.

3. **Ontology Clusters**

Ontology clustering is based on the similarities between the concepts known to the different resources, where each resource represents different aspects of the domain knowledge. Since our resources need to communicate in a sensible fashion they are all supposed to be familiar with some high level concepts. We group these concepts in an ontology rooted at the top of the hierarchy of ontologies. As it describes concepts that are specific to the domain and tasks at hand we refer to this ontology as the *application ontology* (following Van Heijst *et al.*, 1997). These concepts are reusable within the application but not necessarily outside the application. The concept definitions in the application ontology are chosen from an existing top-level ontology, which in our case is WordNet (Miller, 1990). The application ontology thus contains a relevant subset of WordNet concepts. For each concept one or more senses are selected, depending on the domain.

If some resources share concepts that are not shared by other resources then this leads to the creation of two (or more) sibling ontologies. Each sibling is a consistent extension of its parent ontology, but heterogeneous with respect to its peers.

A cluster is referred to as a group of consistent ontologies (possibly one) in our structure and is described by an ontology which is shared by those composing the cluster. Both ontology clusters and ontologies within each cluster are organised in a

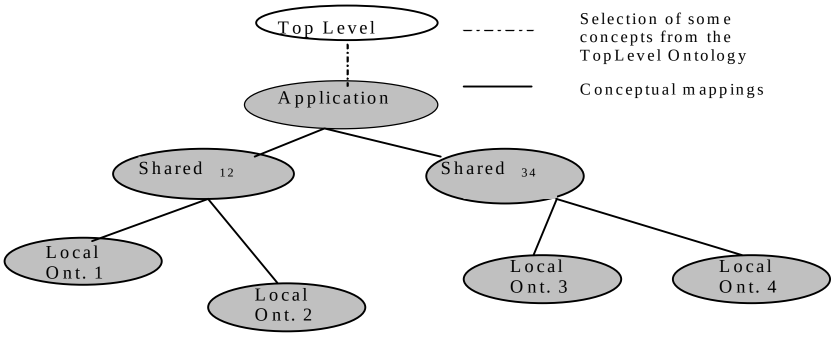

Figure 1 The structure of multiple shared ontologies

hierarchical fashion where each sibling cluster specialises the concepts that are in its parent cluster. The lower level clusters have more precise concept definitions than the higher levels, making the latter more abstract. Figure 1 illustrates this structure, where Local Ont. are the local ontologies, Shared_{12} is the ontology shared by the local ontologies 1 and 2 and analogously Shared_{34} is the ontology shared by the local ontologies 3 and 4. Shared_{12} and Application is an example of a cluster.

4. **The knowledge model**

The knowledge model used to represent the ontology plays a crucial role in the approach as it provides the elements that are needed for (semi-)automatic clustering. The knowledge model we propose is frame-based (Minsky, 1991) and it can be considered as an extension of OKBC (Chaudhri, 1998). Our model is based on *classes*, *slots*, and *facets*. *Classes* correspond to concepts and are collections of objects sharing the same properties. Classes are hierarchically organised and linked by *is-a* links. Members of the classes are described in terms of attributes, that is *slots*, which are attached to each class. *Slots* can either be sets or single values. A slot is described by a name, a domain, a value type and by a set of additional constraints, here called *facets*. *Facets* can contain the documentation for a slot, constrain the value type or the cardinality of a slot, and provide further information concerning the slot. More particularly, our framework has additional facets which facilitates the characterisation of attributes by describing both the domain and the attribute properties (for example whether there is an order relation on their domain) as well as the attribute' s behaviour (for example whether the attribute changes its value regularly

begin-ontology Italian-coffee
description “concepts used in preparing coffee in Italy”
...
concept: drink
subtype_of: liquid
slot: necessary-ingredient
value-type: substance **cardinality**: 1+
slot: preparation-method
value-type: preparation-method **cardinality**: 1
end-concept
concept: powdered-coffee
subtype_of: solid
end-concept
concept: espresso
subtype_of: drink
slot: ingredient
value-type: powdered-coffee **cardinality**: 1
slot: coffee-maker
value-type: espresso-maker **cardinality**: 1
end-concept
concept: espresso-maker
subtype_of: kitchen-tool
slot: produces
value-type: espresso **cardinality**: 1+
slot: water-reservoir
value-type: water-reservoir **cardinality**: 1
slot: filter-and-coffee-holder
value-type: filter-and-coffee-holder **cardinality**: 1
slot: coffee-reservoir **value-type**: coffee-reservoir
cardinality: 1
end-concept
...
end-ontology

begin-ontology Dutch-coffee
description “concepts used in preparing coffee in The Netherlands”
...
concept: drink
subtype_of: liquid
slot: ingredient
value-type: substance **cardinality**: 1+
slot: recipe
value-type: recipe **cardinality**: 1
end-concept
concept: ground-coffee
subtype_of: substance
end-concept
concept: Dutch_coffee_drink
subtype_of: Coffee_drink
slot: coffee-ingredient
value-type: ground-coffee
cardinality: 1
slot: coffee-maker
value-type: electric-coffee-maker
cardinality: 1
end-concept
concept: electric-coffee-maker
subtype_of: kitchen-appliance
slot: makes
value-type: Dutch_coffee_drink
cardinality: +
slot: water-reservoir
value-type: water-reservoir
cardinality: 1
slot: paper-filter-coffee-container
value-type: filter-and-coffee-holder
cardinality: 1
slot: jug
value-type: jug **cardinality**: 1
end-concept
...
end-ontology

Figure2: An example of similar ontologies described using the proposed knowledge model

in time or depending on the occurring of some event). Both these descriptions are used in the clustering. Two examples of ontologies described using this knowledge model are presented in Figure 2. The ontologies describe a fragment of the coffee-making domain in two countries, here Italy and The Netherlands (after Visser and Tamma, 1999). The fragment illustrates the semantic mismatch between the concept coffee in both cultures. Although in the example both concepts have been translated into English their semantic dissimilarity remains apparent.

5. **Conceptual clustering in ML**

Assuming that a feasible approach to knowledge sharing is the structure of ontology clusters described in section 3, here the question considered is whether some phases of the creation of such a structure can be automated. We assume the availability of a top-level ontology (WordNet) and individual resource ontologies and want to investigate means to automatically generate the intermediate ontologies and their clustering (*cf*. Fig 1). Our aim in this section is to describe problems arising when applying various machine learning techniques to create such structure (*cf*. fig 1). Presently, is not realistic to believe that such a task could be performed completely automatically. However, we try to analyse how other disciplines have dealt with the problem of obtaining automatic classifications and borrow some ideas in order to attempt the (partial) automation of the task. The motivation behind semi-automating the process of generating the intermediate ontologies is scalability: Using a structure of multiple shared ontologies to achieve knowledge sharing makes the architecture easy scalable. Resources can join and leave the system at any time and when a new resource joins the architecture it does so by referring to the shared ontology which is closer to the resource's way of conceptualising the world. This means that the shared ontologies have to respond dynamically to the joining or the leaving of a resource. This paper proposes to use a ML technique, more specifically conceptual clustering (Michalski and Stepp, 1983), to partially automate the generation of the intermediate ontologies.

In the Artificial Intelligence field, the classification problem has been widely studied in *inductive learning*, Categorisation (better known in machine learning as clustering) is a form of learning by observations that aims at grouping objects into classes (called *clusters*). In order to decide how clusters should be built, a measure to evaluate the *similarity* between objects has to be provided. Clusters are well formed if the intra-class similarity (similarity between objects belonging to the same class) is high and the inter-class similarity (similarity between objects belonging to different classes) is low. However, having an *aggregation* criterion is not enough to have a good categorisation or taxonomy. In fact classes should also have a simple conceptual description (*characterisation*) which makes them easy to interpret. This is the aim of that area of machine learning called *conceptual clustering* (Michalski & Stepp, 1983). Clusters have logically disjoint descriptions and optimise a given quality criterion. The process of building clusters of ontologies can benefit from the use techniques developed for conceptual clustering in order to semi-automate the two phases of the clustering, that is the *aggregation* phase when ontologies that are deemed similar are grouped together and the *characterisation* phase which permits to generate a description for the cluster that can become the intermediate ontology.

Many conceptual clustering systems have been reported in the literature. Most of them operate top-down (divisive approach), by recursively partitioning the initial set of observations: CLUSTER/2 (Michalski & Stepp, 1983), UNIMEM (Lebowitz, 1986), COBWEB (Fisher, 1987) and its extension CLASSIT (Gennari et al., 1989), LABIRYNTH (Thompson & Langley, 1991). Some bottom-up clustering algorithms have also been designed (aggregative approach), such as those implemented in the systems WITT (Hanson & Bauer, 1989), KBG (Bisson, 1992), and PYRAMID (Brito, 2000). All these systems work only for zero-order representations, that is vectors of attribute-value pairs. The only two exceptions are CLUSTER/S (Stepp & Michalski, 1986) and KBG, which was developed in the frame of the ESPRIT project (2154) Machine Learning Toolbox. In both cases observations are represented as conjunctions of ground atoms. In particular, the KBG representation of observations is based on *entities*, *predicates*, *values* and *types*. In KBG it is also possible to express some form of background knowledge, usually represented as *rules*. The background knowledge is used to saturate the initial observations: in this way information that was implicit in the example, given the background knowledge, is made explicit (Rouveirol, 1992).

Conceptual clustering algorithms working on first-order representations are important in order to solve the problem of ontology clustering. Although aggregating similar ontologies is an important aspect of the ontology clustering problem, it is not the only one. Indeed, the hierarchical organisation of clusters suggests inheritance mechanisms that enable the definition of common general concepts in the higher level clusters and the specification of local conceptualisations of resources in the lower level clusters. Therefore, it is equally important to *characterise* each cluster by associating a description of the properties that are common to all ontologies in it. Our interest for conceptual clustering methods is justified by their possibility of expressing such cluster characterisation. Unfortunately, the knowledge representation in the ontology case is more complex than the one considered in the classical conceptual clustering approach, meaning that many conceptual clustering systems cannot be used to cluster ontologies. First-order clustering systems seem to fit better the ontology clustering problem, although they still present some limitations, due to the fact that they have been designed to aggregate the extensions of a concept, that is the set of its instances, and not the intensional descriptions of concepts, as it is the case with ontology clustering.

6. **Some preliminary ideas on semi-automatic ontology clustering**

The lessons learned about conceptual clustering provide useful insights on the categorisation problem and can give hints on how to proceed in the attempt of clustering ontologies. In the remainder we point out some of the difficulties that emerge while trying to build ontology clusters and we try to highlight similarities with problems already studied in clustering in order to see whether some conceptual clustering technique can be adapted to solve those problems. Ontology clustering might be considered as a form of conceptual clustering, where the objects to classify are hierarchies of concepts. As for the classical conceptual clustering, similarity between objects should not be context free (that is, it should depend entirely on the properties of the objects whose similarity is being evaluated), but it should be context sensitive. In ontology clustering *context* plays a crucial role. Indeed, the main assumption behind clustering is that the ontologies to be clustered are all representing knowledge about different aspects of the same domain, therefore, there should be a way to permit communication between them, even if this communication occurs at a high level of abstraction. Ontologies are clustered on the basis of the similarity between concepts composing them. As concepts are hierarchically organised, we can assume that if there is a similarity between concepts higher in the hierarchy then the ontologies are similar. However, evaluating the similarity between concepts is not a trivial task, because concepts can be described by different attributes, therefore usual similarity measures, which are based on comparison of the values taken by the same attribute in two concept descriptions, cannot be used. Here we propose to evaluate the similarity between concepts in different ontologies by evaluating two aspects: the "*semantic"* aspect and the *"descriptive"* aspect. The semantic aspect evaluates the *semantic similarity* between concepts while the *descriptive aspect* bases the similarity evaluation on the attributes used to describe the concepts.

Semantic similarity is a form of semantic relatedness using network representation, a problem that received much attention in the artificial intelligence field (Quillian, 1968), (Collins and Loftus 1975). Rada et *al.* (Rada *et al.*, 1989) suggest that similarity in semantic networks can be assessed solely on the basis of the *IS-A taxonomy*, without considering other types of links. One of the easiest way to evaluate semantic similarity in taxonomies is to measure the distance between the nodes corresponding to the items being compared, that is the shorter the path between the nodes, the more similar they are. This idea is inspired to some definitions of dissimilarities defined for cluster analysis, namely ultrametrics, tree distances and strong Robinsonian dissimilarities (Esposito *et al.*, 2000). More precisely, an ultrametric dissimilarity fulfils the ultrametric inequality:

$$d(a,b) \leq \max\{d(a,c),d(c,b)\}$$

where a,b and c are nodes of the taxonomy in our case. It can be shown that an ultrametric results from a hierarchical classification (dendrogram) of individuals, and conversely by putting d(a,b)="the lowest level at which the objects a and b meet in the dendrogram." This *bijection theorem* provides a unique characterisation of hierarchical classifications by ultrametrics. A tree distance, instead, characterises an additive tree, i.e. a tree T with *n* vertices and n-1 weighted edges: The dissimilarity d(a,b) is then reproduced as the sum of the weights of all edges of the (unique) path connecting two given vertices a,b in T. Additive trees are widely used in the reconstruction of phylogenetic evolutions, and might be useful in the context of ontology clustering only if a weight is associated to each concept inheritance in order to take into account the amount of inherited properties. Finally, Robinsonian dissimilarities characterise pyramidal classification of individuals. The pyramidal model generalises hierarchies by allowing non-disjoint classes at each given level, therefore Robinsonian dissimilarities should be considered only in the case of overlapping concepts in ontologies.

The concepts in the top-level ontology provide the bases for evaluating semantic similarity, indeed, the idea is to calculate the paths as those connecting the concepts in the resource ontologies to their nearest ancestor in the top-level ontology (Resnik, 1999). Unfortunately, we cannot completely rely on semantic similarity, as some terms can have more than one meaning, given rise to problems of ontological mismatches, as highlighted in (Visser *et al.*, 1998). For this reason the semantic aspect needs to be combined with the descriptive one. The evaluation of the semantic similarity provides a set of candidate similar concepts to which the descriptive aspect is evaluated. Once the candidate similar concepts have been determined, concept descriptions are taken into account in the attempt to establish whether the attributes can be related either on the basis of the semantic similarity or the domain analogy. For the semantic similarity evaluation concepts in the ontology are compared with concepts in the top-level ontology (which is on top of the application ontology and which provides the terminology to describe the concepts in the application ontology). For the domain similarity, which evaluates the descriptive aspect, all the slots' facets are compared. In summary, the similarity between concepts is a function of the semantic similarity between the concept names and the similarity between the attributes describing the concepts.

7. **Conclusion and future work**

This paper describes the so-called *ontology clustering* approach to knowledge sharing. In this approach ontologies are clustered together on the basis of the similarity they show in the way they conceptualise their domain. More in particular, the paper discusses preliminary ideas on how to obtain the ontology clusters in a semi-automatic fashion. We deem our approach to be novel for two reasons: the idea of using a set of heterogeneous and structured shared ontologies has not received much attention in the literature on knowledge sharing, and we are not aware of any other research effort which is trying to apply machine learning techniques to the knowledge sharing. The approach presented in this paper is still being investigated, therefore advantages and disadvantages of this approach can only be envisaged. The main advantage is that this approach should permit not only to obtain in a semi-automatic way the ontology clusters but also the shared ontologies characterising the clusters. One of the disadvantages is the heavy computational burden due to the evaluation of both semantic semantic and descriptive aspect. As this approach is still a proposal there are open problems to deal with, such as the evaluation of the similarity between concepts in heterogeneous ontologies depending on the semantic and the descriptive aspects. One of the main problems remains the discrepancy between ontology research, which is predominantly first-order-logic based, and the core of the conceptual clustering techniques, which are predominantly based on zero-order logic. In future research we aim to clarify more in detail how the semantic aspect and the descriptive aspects contribute to the evaluation of the similarity. To do so, a more precise definition of the descriptive aspect is required.

Acknowledgements

Valentina Tamma' s research at Liverpool University is funded by BT plc, United Kingdom.

References

1. Bayardo, R.J., Bohrer, W., Brice, R., Cichocki, A., Fowler, G., Helal, A., Kashyap, V., Ksiezyk, T., Martin, G., Nodine, M., Rashid, M., Rusinkiewicz, M., Shea, R., Unnikrishnan, C., Unruh, A., and Woelk, D. (1997): InfoSleuth: Agent-Based Semantic Integration of Information in Open and Dynamic Environments. In *ACM SIGMOD Record Vol. 26, No. 2 (June 1997), SIGMOD '97. Proceedings ACM SIGMOD International Conference on Management of Data*, Tucson, Arizona, USA, pp. 195-206.
2. Bisson, 1992. Conceptual clustering in a first-order logic representation, *Proceedings of the 10th ECAI*, Vienna, pp. 458-462.
3. Brito, P. (2000). Hierarchical and pyramidal clustering with complete symbolic objects. Chapter 8.3 in *Analysis of Symbolic Data. Exploratory methods for extracting statistical information from complex data, Series: Studies in Classification, Data Analysis, and Knowledge Organization,* Bock, H.-H. and Diday E. (Eds.) vol. 15, 312-324, Springer-Verlag, Berlin.
4. 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),* AAAI Press/The MIT Press, Madison, Wisconsin, pp. 600-607.
5. Collins, A., and Loftus, E. (1975): A spreading activation theory on semantic processing. In *Psychological Review*, vol.82, pp. 407-428.
6. Esposito, F., Malerba, D., Tamma, V., & Bock, H-H. (2000). Classical resemblance measures. In *Analysis of Symbolic Data. Exploratory methods for extracting statistical information from complex data*, Bock, H-H. and Diday, E. (Eds.), Series: Studies in Classification, Data Analysis, and Knowledge Organization, vol. 15, Springer-Verlag:Berlin, 139-152.
7. Fisher, D.H. (1987). Knowledge acquisition via incremental conceptual clustering. *Machine Learning Journal*, vol. 2, pp. 139-172.
8. Gennari, J., Langley, P. and Fisher, D. (1989). Models of incremental concept formation. *Artificial Intelligence Journal*, vol. 40, pp. 11-61.

9. Gómez Pérez, A., and Benjamins, V.R. (1999): Overview of Knowledge Sharing and Reuse Components: Ontologies and Problem-Solving Methods. In *Proceedings of the IJCAI'99 Workshop on Ontology and Problem-Solving Methods: Lesson learned and Future Trends*, Benjamins, V.R., (Ed.), CEUR Publications, Amsterdam, vol.18, pp.:1.1-1.15.
10. Gray, P.D.M., Preece, A., Fiddian, N.J., Gray, W.A., Bench-Capon, T.J.M., Shave, M.J.R., Azarmi, N., Wiegand, M., Ashwell, M., Beer, M., Cui, Z., Diaz, B., Embury, S.M., Hui, K., Jones, A.C., Jones, D.M., Kemp, G.J.L., Lawson, E.W., Lunn, K., Marti, P., Shao, J., and Visser, P.R.S. (1997). KRAFT: Knowledge Fusion from Distributed Databases and Knowledge Bases, *Database and Expert System Applications (DEXA' 97)*, Toulouse, France.
11. Guarino, N. (1998): Formal Ontology and Information Systems. In *Proceedings of Conference on Formal Ontology (FOIS'98)*, Guarino N. (Ed.), Trento, Italy.
12. Hanson, S., & Bauer, M., (1989). Conceptual clustering, categorization and polymorphy. *Machine Learning Journal*, vol. 3, pp. 343-372.
13. Lebowitz, M. (1987). Experiments with incremental concept formation: UNIMEM. *Machine Learning Journal*, vol. 2, pp. 103-138.
14. Mena, E., Kashyap, V., Shet, A., and Illarramendi, A. (1996): OBSERVER: An Approach for Query Processing in Global Information Systems based on Interoperation across Pre-existing Ontologies, in *Proceedings First IFCIS International Conference on Cooperative Information Systems (CoopIS '96)*, Brussels, Belgium, IEEE Computer Society Press, pp. 14-25.
15. Michalski, R.S. (1983). A theory and methodology of inductive learning. In *Machine Learning: An Artificial Intelligence Approach*, Michalski, R.S., Carbonell, J.G. and Mitchell, T.M. (Eds.), Tioga Publishing Co., Palo Alto, CA, pp. 83-134.
16. Michalski, R.S. & Stepp, R.E. (1983). Learning from observations: Conceptual Clustering. In *Machine Learning: An Artificial Intelligence Approach*, Michalski, R.S., Carbonell, J.G. and Mitchell, T.M. (Eds.), Tioga Publishing Co., Palo Alto, CA, pp. 83-134.
17. Miller, G.A. (1990): Nouns in WordNet: a lexical inheritance system. In *International Journal of Lexicography*, vol. 3, no. 4, pp. 245-264,
18. Minsky, M. (1992): A Framework for representing knowledge. In *Cognitive Science*, Collins A. and Smith, E.H. (Eds.), Morgan Kaufmann, Los Altos, CA, pp. 191-215.
19. Quillian, M.R. (1968): Semantic Memory. In *Semantic Information Processing*, Minski, M. (Ed.), MITT Press, Cambridge, MA.
20. Rada, R., Mili, H., Bicknell, E., and Blettner, M. (1989): Development and application of a metric on semantic nets. In *IEEE Transaction on Systems, Man and Cybernetics*, vol. 19, no. 1, pp. 17-30.
21. Resnik, P. (1999): Semantic Similarity in a Taxonomy: An Information-Based Measure and its Application to Problems of Ambiguity in Natural Language. In *Journal of Artificial Intelligence Research*, vol. 11, pp. 95-130.
22. Rouveirol, C. (1994): Flattening and saturation: Two representation changes for generalization. In *Machine Learning Journal*, vol. 14, pp. 219-232.
23. Shave, M.J.R. (1997): Ontological Structures for Knowledge Sharing. In *The New Review of Information Networking*, Vol 3, pp. 125-134.
24. Stepp, E. and Michalski, R.S. (1986). Conceptual clustering of structured objects: a goal-oriented approach. In *Artificial Intelligence*, vol. 28, pp.43-69.
25. Thompson, K. & Langley, P. (1991): Concept formation in structured domains. In *Concept formation: Knowledge and experience in unsupervised learning*, Fisher, D.H., Pazzani, M.J., & Langley, P. (Eds.) San Francisco, CA: Morgan Kaufmann
26. Van Heijst, G., Schreiber, T., and Wielinga, B. (1997): Using Explicit Ontologies in KBS. In *International Journal of Human-Computer Studies*, vol. 46, nr. 2/3, pp. 183-292.
27. Visser, P.R.S., Jones, D.M., Bench-Capon, T.J.M., and Shave M.J.R. (1998). Assessing Heterogeneity by Classifying Ontology Mismatches, *Proceedings Conference on Formal Ontology* (FOIS' 98),Guarino, N., (Ed), Trento, Italy, pp. 148-162. An earlier version of this article appeared at the AAAI 1997 Spring Symposium on Ontological Engineering, Stanford, CA, USA.
29. Visser, P.R.S., and Cui Z. (1998). On Accepting Heterogeneous Ontologies in Distributed Architectures, *ECAI'98 workshop on Applications of Ontologies and Problem-solving methods*, Brighton, UK.
30. Visser, P.R.S., and Tamma, V.A.M. (1999): An Experience with Ontology-based Agent Clustering. In *Proceedings of the IJCAI'99 Workshop on Ontology and Problem-Solving Methods: Lesson learned and Future Trends*, Benjamins, V.R., (Ed.), CEUR Publications, Amsterdam, vol.18, pp.:12.1-12.13.