Semantic Information Mediation among Multiple Product Ontologies

Alex G. Büchner*, Mervi Ranta§, John G. Hughes*, Martti Mäntylä§

* Northern Ireland Knowledge Engineering Laboratory, University of Ulster
§ Laboratory of Information Processing Science, Helsinki University of Technology
email: ag.buchner@ulst.ac.uk, mervi.ranta@hut.fi, jg.hughes@ulst.ac.uk, martti.mantyla@hut.fi

ABSTRACT

The concept of shared ontologies has been the basis for some substantial work in enterprise and product modelling environments. Albeit their success, there has almost always been problems with single shared product ontologies and the associated ontological commitment in real-world environments, especially when semantic conflicts have to be resolved. The main objective of this work is to tackle these issues through the development of a mediation mechanism that is capable of handling multiple related product ontologies, while keeping the autonomy of each participating entity.

The notion of product concept hierarchies is introduced, each of which represents a domain-related ontology of a participating partner in the product process. Mappings that supersede the capabilities of existing generic approaches are defined, which are centred around product modelling-specifics. Such specifics encompass leaf node—leaf node, non-leaf node—leaf node and non-leaf node—non-leaf node relationships across ontologies. These mappings are prerequisites for modelling generalisation, specialisation, part-of and is-a relationships among product ontologies. Furthermore, a product data management-specific relationship called alternative is introduced, which allows the specification of optional terms in related ontologies. Based on this fundament, mediation is introduced, which resolves semantic conflicts between different sites along the product process. The entire system (ontologies, mappings, mediator, and architecture) are accompanied by examples, which stem from an industrial scenario.

INTRODUCTION

Traditional product data management (PDM) systems have been tailored around a locally agreed notation, which usually mirrored a function at a stage in a product process. The relatively fast realisation of the limitations of these insular systems has led to more advanced approaches based on commonly shared product ontologies. The underlying philosophy of such a pre-defined collection of concepts and their interconnections to describe information entities, is the ontological commitment each participating site — organisational units in an enterprise or organisations in extended enterprises — have to be liable to. This restriction has been proven resolvable in small systems, but has shown almost infeasible in medium to large environments, let alone forthcoming Internet-based product development processes in virtual enterprises.

The objective of this work is to tackle this intrinsic obstacle in relaxing the restriction of a shared ontology through the support of multiple domain-related product ontologies, without sacrificing each partner's autonomy. The types of relationships that occur in PDM environments, which cause inconsistencies are

- leaf node — leaf node;
- non-leaf node — leaf node (and vice versa); and
- non-leaf node — non-leaf node.

These mappings are the prerequisite for modelling generalisation, specialisation, part-of (sub-type — super-type) and is-a relationships among product ontologies. The possible semantic conflicts among those relationships are dealt with in great detail through a customisable dynamic mediation facility.

The contribution of our approach is manifold. Firstly, hindering ontological commitment based on shared product ontologies and concept libraries is removed. Secondly, the autonomy of each participating site in those loosely coupled environments is much higher than in traditional tightly coupled systems. Thirdly, a PDM-specific relationship called alternative is introduced, which allows the specification of optional terms in related ontologies. And lastly, the resolving mechanisms of semantic conflicts allow a much higher degree of collaboration in heterogeneous product modelling environments.

The outline of the paper is as follows. First, the structure of concepts hierarchies and relationships within these product trees is introduced. Then, their restrictions are outlined and multiple product ontologies as well as possible semantic discrepancies are described. Next, a brief recapitulation of semantic information mediation in general is given, before the concepts are mapped onto the PDM context. The synergy of all provided components leads to a prototypical architecture. Finally, related work is evaluated, before conclusions are drawn and further work is sketched out. The entire paper is supported by a running example from the PDM domain, namely thermometer manufacturing and retail for domestic, that is non-experimental and non-laboratory usage. The accompanying example has been kept as simple as possible for didactical reasons.

PRODUCT CONCEPT HIERARCHIES

The most typical construct to describe PDM concepts is that of hierarchies, hence product concept hierarchies, which are also known as terminological ontologies. The advantage of tree-like representations is its familiarity among all participated experts, for instance, the product manager, the designer, the constructor, sub-contractors, and the retailer. Although some have argued that network-like structures are conceptually more powerful (Chen and Lynch, 1992), hierarchies have proven to be more appropriate. Formally, a concept hierarchy is defined as follows (Büchner, Bell and Hughes, 1998) [1].

Def. 1. A **concept hierarchy** o is an undirected, connected, acyclic graph which is defined as the tuple $o = (C, E)$, where $C = \{c_0, c_{1_1}, c_{1_2}, ..., c_{1_{|c_1|}}, c_{2_1}, c_{2_2}, ..., c_{2_{|c_2|}}, ..., c_{n_1}, c_{n_2}, ..., c_{n_{|c_n|}}\}$ and $E = \{e_1, e_2, e_3, ..., e_m\}$. Each e_k has the form $e_k = <c_i, c_j>$; $c_i, c_j \in C$, c_0 has indegree 0, $c_1 ... c_n$ have indegree 1. c_i is **subconcept** of c_j iff $c_i \subset c_j$; c_i is **superconcept** of c_j iff $c_j \subset c_i$. ♦

The connotation of the symbol o is that each product concept hierarchy represents an ontology. C represents the set of all concepts within o, E the set of all edges in the universe of discourse U. According to Gruber (1995), an ontology is an *"explicit specification of a conceptualization"*, which should fulfil the criteria of clarity, coherence, extendibility, a minimal encoding bias, and minimal ontological commitment. Fox, Chionglo and Fadel (1993) have proposed the partially overlapping criteria generality, efficiency, perspicuity, extensibility, transformability, granularity, scalability as well as competence.

The ontology M_1, which is used for the purpose of demonstration is depicted in Fig. 1. A hypothetical thermometer manufacturer uses a simple bill of material, which distinguishes between the case and the tube. Former is produced by the manufacturer itself, latter is purchased from a sub-contractor. The case is made of a specific material, which can be wood, plastic or metal; the supported scale is expressed in degree Celsius, Fahrenheit, or Kelvin.

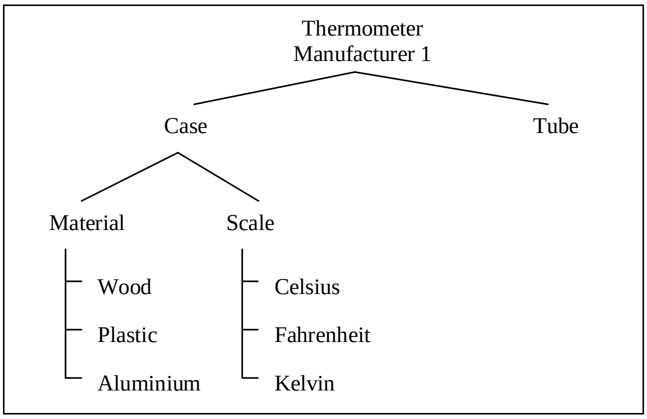

Fig. 1 An ontological representation of thermometers

Hierarchies also provide an implicit mechanism of describing various relationship types, which are relevant in PDM scenarios. The most often used types are is-a relationships, aggregation and its inverse generalisation, as well as part-of relationships. For instance, wood is-a material; aluminium, plastic and wood can be aggregated to materials; and a case and a tube are part-of a thermometer.

Additionally, it is often useful to package sub-trees in separate entities, which represent a product component. Each composite can be re-used in different contexts or can be unplugged, if not required. Further, when embedded in an object data model, mechanisms such as inheritance and overloading can be facilitated (Dubitzky, Büchner, Hughes and Bell, 1999). Formally, a product composite is defined as follows.

Def. 2. A composite p is defined as a subgraph of a concept hierarchy o, such that $p \subset o$, that is $\forall p_i \in$ p $\{p_I \in o\}$ and $p \neq o$.

The tube in the ontology depicted in Fig. 1, can be seen as a composite, which contains a bulb as well as a capillary. Obviously, other packages could have easily derived.

---

[1] The indegree of a node n is the number of arcs coming into n, the outdegree the number of arcs leaving n.

Albeit the powerful mechanisms that are provided by product concept hierarchies, they come, as most other representation mechanisms, with an intrinsic drawback. If a product tree is used as shared product ontology, each participating site has to be committed to the notational usage of the agreed concepts as well as their arrangements. Unfortunately, this is a very unrealistic requirement in large organisations and particularly in extended enterprises or Internet-based product processes. Furthermore, very often domain-related ontologies already exist in companies, which decide to collaborate in some form or the other (for instance, merge in the most extreme case). Again, an agreement in this evolutionary scenario is even less likely than in ontologies which are being built revolutionary (Mizoguchi and Ikeda, 1996). A far more realistic and powerful approach is the explicit support of multiple related product ontologies, which is discussed in the sequel of this paper.

MULTIPLE PRODUCT ONTOLOGIES

As mentioned in the previous section, the main justification for the introduction of multiple related product ontologies in collaborative environments is the near impossibility of combining pre-existing ontologies in a sensible way and the restriction of the ontological commitment, once a shared / global ontology has been agreed upon. Furthermore, autonomy is an important feature in competitive manufacturing environments.

There are two key issues when dealing with multiple ontologies. Firstly, they have to be domain-related; domain in the context of PDM means that they have to represent the same or at least similar concept(s) differently. Secondly, there still has to be a minimal shared ontology, which acts as the mechanism that binds different ontologies semantically together. This global ontology does not have to be in hierarchical form; a host rules, a set of constraints, or whatever domain knowledge is appropriate, can be facilitated. This core ontology is domain-independent and unlikely to be modified over time.

For the sake of completeness, the set of (product) ontologies that are used in a heterogeneous PDM environment can be defined as follows.

Def. 3. An ontology space is spanning the set of ontologies $O = \{o_1, o_2, \ldots, o_n\}$. ¨

The global ontology *G*, which describes the nucleus of a mediator or brokering mechanism facilitated in collaborative PDM environments is defined at a later stage, when mappings among entities are introduced as part of the overall architecture.

For illustration purposes, a second product concept hierarchy (M_2) is added to the one depicted in Fig. 1, which represents the same domain from another idiosyncratic view, that is from a different thermometer manufacturer. The bill of material of the second manufacturer competing in the enterprise distinguishes between the board, the glass and the liquid in the tube. Latter can be either spirit (in form of red liquid) or mercury for high temperature measurements or highly precise thermometers. Boards are either made from metal, PVC, or wood, where latter is either oak or pine.

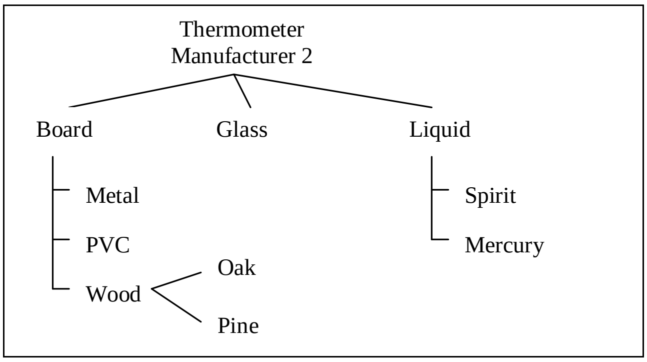

Fig. 2 An alternative ontological representation of thermometers

However, the introduction of multiple product concept hierarchies causes a host of challenges, which have to be tackled before the approach is of any feasible usage in collaborative environments in general and PDM scenarios in particular. The major problem is that of developing a flexible mechanism, which provides a vehicle that allows constructs, which have been expressed in one ontology can be mapped onto another. From a co-operative information systems point of view, this is handled as a semantic heterogeneity problem, as known from the multi-database world. The particular area of semantic heterogeneity we are confronted with has been coined by Kashyap and Sheth (1998) as interontology interoperation. The types of semantic conflicts that typically occur in heterogeneous environments are

- identical usage of terms;
- usage of a composite term for an atomic term (specialisation); and
- usage of an atomic term for a composite term (generalisation).

Kashyap and Sheth (1998) have used the relationships synonymous, hypernymous and hyponymous respectively to describe the three potential conflicting types. We have adopted their notation and formally specify these relationships as follows.

Def. 4. Given two ontologies o_1 and o_2, a relationship $r_s(c_1, c_2)$ is defined as **synonym**ous iff $c_1 \in o_1 \wedge c_1$ has outdegree = 0 and $c_2 \in o_2 \wedge c_2$ has outdegree = 0. ¨

A synonymous example from the above PDM ontologies M_1 and M_2 are M1.Case.Material.Plastic and M2.Board.PVC.

Def. 5. Given two ontologies o_1 and o_2, a relationship $r_r(c_1, c_2)$ is defined as **hypernym**ous iff $c_1 \in o_1 \wedge c_2 \in o_2$ and c_1 has oudegree = 0 and c_2 has outdegree ≥ 1. ¨

The inverse relationship of Definition 5 is defined as follows.

Def. 6. Given two ontologies o_1 and o_2, a relationship $r_o(c_1, c_2)$ is defined as **hyponym**ous iff $c_1 \in o_1 \wedge c_2 \in o_2$ and c_1 has outdegree ≥ 1 and c_2 has outdegree = 0. ¨

Due to the fact that hypernymous and hyponymous relationships are symmetric to each other, a single example can illustrate both linkage types. M1.Case.Material.Wood is hyponymous to both, M2.Board.Wood.Oak and M2.Board.Wood.Pine, whereas M2.Board.Wood.Oak and M2.Board.Wood.Pine are hypernymous to M1.Case.Material.Wood, respectively.

In addition to these three generic relationships, a fourth type is provided, which we call *alternative* relationship. An alternative relationship allows the specification of two or more optional parts in either the same or alternatively in different ontologies. This is extremely useful in PDM scenarios, for instance, in which alternative components are required in case of an unforeseen process bottleneck situation.

Def. 7. Given an ontology space O, a relationship $r_a(c_1, c_2, \ldots, c_n)$ is defined as **alternative** iff $c_i \in o_i$ and $c_1 \equiv c_2 \vee c_1 \equiv c_3 \ldots \vee c_1 \equiv c_n$, where the symbol ≡ represents a synonymous, hypernymous or hyponymous relationship. ¨

To illustrate the usefulness of the newly introduced relationship, consider the wood branches of M_1 and M_2. In certain scenarios it is useful to declare a relationship that expresses that M1.Case.Wood can either be M2.Board.Wood.Oak *or* M2.Board.Wood.Pine. The equivalent to that example is to state that M1.Case.Wood is the same as M2.Board.Wood. Although this scenario can be described with a combination of the three relationships described above (synonyms, hypernyms and hyponyms), this becomes very opaque when more complex real-world ontologies have to be mapped onto each other.

The usage of alternative relationships across ontologies becomes clear when more participants are involved in a collaboration. An example has been given in (Ranta, Büchner, Mäntylä and Hughes, 1999), where a thermometer retailer is requesting information about available components. In case one manufacturer has either run out of a particular part or does not provide it at all, alternative components can be considered. Similarly, prices of interchangeable parts might vary from manufacturer to manufacturer, which requires the specification of alternative relationships.

SEMANTIC INFORMATION MEDIATION

Although mediation functions have been developed for more than two decades, mediators in information systems as such have been formalised for the first time by Wiederhold (1992). The pioneer has defined a mediator as a "… *module that exploits encoded knowledge about some sets or subsets of data to create information for a higher layer of applications. It should be small and simple, so that it can be maintained by one expert, or at most, a small and coherent group of experts*". Much work carried out in the area of mediation has originated from projects under the umbrella of ARPA's I^3 (Intelligent Integration of Information) initiative.

From an architectural point of view, a mediator is an independent application layer, located between a data source layer and a data source (Lee, Madnick and Siegel, 1996). Each mediator has an interface to each boundary, namely a resource access interface to the database (for example, a product catalogue) and a service interface to the application (for instance, production scheduler). The mediator itself contains domain-specific code which is based on a pre-defined terminology. Specialised add-ons have been added to this simplified architecture to overcome some initial problems, for instance, wrappers to deal with legacy data or integrators to combine resources.

The mediator concept has been heavily employed in the database domain, in which almost all PDM applications are located. In such domains, the task of the information mediator is to negotiate requests from different clients (for instance the earlier mentioned thermometer retailer), based on site-specific information, available to the mediator. Techniques used for the implementation of the mediator are domain- as well as application-specific. The internal design of a mediator is not discussed in here, since it is out of the scope of this paper. The key is that of locating information about participating sites and negotiating solutions which satisfy the receiver and are accepted by the source. The structural aspects, that is the representation of semantic heterogeneity among sites in a collaborative product-based environment are discussed in the sequel.

A commonly accepted and widely used technique is that of mapping information about one site onto information about another site. The mapping have to provide linguistic support the the relationships which are modelled in the environment (Kashyap and Sheth, 1998) and thus acts as a global ontology. In our case, these are synonyms, hyponyms, hypernyms, and alternatives, which have proven sufficient in product-related domains. We now specify the mappings schematically, before PDM examples, based on the earlier specified ontologies M_1 and M_2 are given.

In order to specify synonym, hypernym, hyponym, and alternative mappings, we adapt Kashyap and Sheth's (1998) notation. A synonym requires a canonical form (for identification among more than two ontologies), an idiosyncratic term and the ontology to which the term belongs to. Synonyms between M_1 and M_2 are listed in Table 1. For simplicity, the dot notation for identifying entities in a product hierarchy have been rejected.

Table 1. Synonyms between M_1 and M_2

Canonical Form	Ontology M_1	Ontology M_2
Metal	Aluminium	Metal
Plastic	Plastic	PVC

Due to the fact that hyponyms and hypernyms are symmetric to each other only one type has to be supported. Each indegree-different relationship requires two terms and their ontologies where they are hypernym / hyponym of. Potential sub-type — super-type relationships across M_1 and M_2 are shown in Table 2. Obviously, more (or less) hypernym / hyponym relationships can be defined, depending on the context of each site.

Table 2. Hypernyms and Hyponyms between M_1 and M_2

Ontology M_1	Ontology M_2
Wood	Oak
Wood	Pine
Tube	Spirit
Tube	Mercury
Tube	Glass
Celsius	M2

Special attention should be drawn to the last entry in the above table. M_2 does not explicitly support different scales, whereas M_1 does. In order to map information from one ontology to the other, semantic information is required. In the given scenario, M2 only manufactures thermometers with a Celsius scale, which does not require the ontological representation. This fact can only be represented in the mapping, which then contains semantic information otherwise not available. The conflict resolution is being performed by the mediator.

Finally, alternative relationships have to be modelled schematically. From a term in an ontology, an (infinite) number of optional terms and the ontologies they belong to have to be given. As can be seen from the example in Table 3, they could theoretically be replaced by synonymous, hypernymous and hyponymous relationships. However, through the explicit specification of alternatives, it is possible to provide the mediator with context-sensitive negotiation directives. Furthermore, the notation reduces the complexity of mapping specifications, which is relevant when a new site is entering the enterprise.

Table 3. Alternatives between M_1 and M_2

Term	Ontology	Alternative(s)
Wood	M_1	Oak(M2) | Pine(M2)
Thermometer	R_1	M1 | M2 | M1.Case & M2. Glass & M2.Liquid

Two types of alternative are given. First provides an intra-ontology specification of wood in M_1, which can be either oak or pine in M_2, respectively. Second defines a mapping across ontologies, that is a thermometer from the retailer R_1 can either be fully purchased from M_1, M_2, or it can be the case from M_1 and the glass tube and the liquid from M_2. The & and | operator connate logical and and or operations, respectively.

A PROTOTYPICAL ARCHITETURE

We now propose a prototypical architecture in which collaboration in heterogeneous environments along product processes can be facilitated. Each participating site contains some local product-related data, usually stored in databases or data warehouses and a product-related ontology. These components are loosely coupled to a broker; the communication is performed through intelligent (transport) agents. The overall architecture is depicted in Fig. 3.

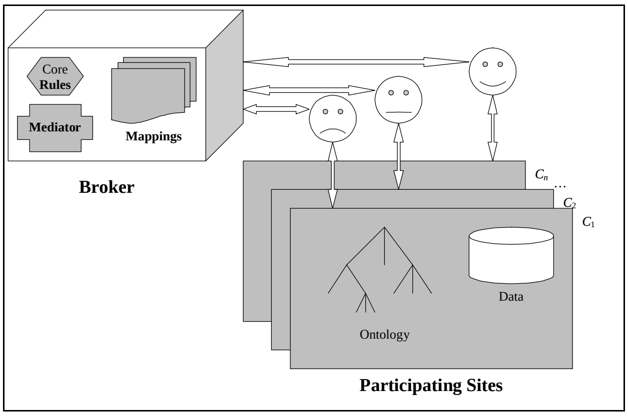

Fig. 3 A Prototypical Architecture

The broker contains a set of core rules, which contain system generic functionality, for instance, how to join a collaboration, how to specify mappings, et cetera. It also contains the mediator itself, which takes requests from the agents and negotiates among other sites based on the mappings, which are also located at the broker. The update of mappings is also performed through the agents. A more detailed description of the architecture and the interaction among the participating components can be found in (Ranta, Büchner, Mäntylä and Hughes, 1999).

The outlined architecture has partially been adopted from (Mena, Kashyap, Sheth and Illarramendi, 1996), and thus can be used as a generic vehicle for interoperation across ontologies, which is not limited to product data environments.

RELATED WORK

Various terminological relationships have been suggested, which provide a mechanism of mapping domain-related contents among ontologies. The most relevant to PDM are discussed in the following paragraphs.

Yu, Sun, Dao and Kerisey (1991) have facilitated common concepts, concept hierarchies and aggregate hierarchies in order to guarantee interoperability among heterogeneous databases. Attribute relationships are then used to specify semantically equivalent entities. The approach is similar to the synonym relationship in our approach; multi-level and inter-ontology relationships however are not supported.

Hammer and McLeod (1993) have proposed a set of relationship descriptors in order to capture relationships between terms across related ontologies. The relationships among objects have been divided into common concepts (identical, equivalent, compatible and incompatible) as well as related concepts (generalisation / specialisation and so called positive associations). The approach is similar to ours, but is limited in terms of mapping specification capabilities.

Michalski (1993) has defined a set of knowledge transmutation operators (generalisation, abstraction, similisation, generation, insertion, and replication). Although the intention of their approach is that of multi-strategy (machine) learning based on inference, the operations can be used for defining ontologies. However, no support is provided for mapping information across ontologies.

Kahng and McLeod (1998) have suggested four conceptual relationships (sub-concept, super-concept, overlap, and disjoint) in order to represent the commonalties between entities in related ontologies. The proposal is weaker in terms of modelling relationships per se, but has the advantage to support a certain degree of uncertainty through the support of a tolerance factor. This type of vagueness can either be caused by faulty information provision or by natural idiosyncratic interpretation of participating sites. We are handling uncertainty aspects as part of future work.

As mentioned previously, Kashyap and Sheth (1998) have created so called synonym, hyponym and hypernym relationships in order to allow inter-ontology interpretation, which has stemmed from Mena et. al.'s work (1996). We have adopted their methodology and applied it in the PDM context. We have also added a construct to their methodology, which allows the explicit allotting of optional mappings, namely alternative relationships.

CONCLUSIONS

A mechanism has been proposed that allows the collaboration of manifold idiosyncratic entities in a PDM environment. The underlying construct is that of multiple product-related ontologies, which are facilitated on each partners site. In order to allow inter-ontology interoperation, semantic information mediation has been introduced, which has proven to be able to solve typical collaboration conflicts among relationships (leaf node—leaf node, non-leaf node—leaf node, and non-leaf node—non-leaf node) successfully, which required to model PDM relationships such as generalisation, specialisation, is-a, part-of, et cetera. Additionally, a PDM-specific relationship, called alternative, has been introduced, which allows the specification of optional terms in related ontologies. Finally, the introduced components have been organised in a prototypical architecture, which will form the basis for further research.

Future work in that environment is threefold. Firstly, multiple ontologies on each participating site will be supported, which is useful if, for instance, a sub-contractor is joining a development process at two different stages, but uses different ontologies internally. Currently, the usefulness of multiple ontology brokers is also investigated. Secondly, uncertainty in the described mapping, which describes the membership degree of two parts on scale form 0 to 1 will be introduced. For example, wood in M_1 can be specified as oak in M_2 with a degree of 70% and pine with a degree of 20%, where the remaining 10% represent some form of ignorance. The resolution of conflicts has to be performed by the mediator, which has to be extended in order to be capable of handling uncertainty. Lastly, specifying the mapping between ontologies is a tedious and time-consuming task, which might cause potential partners not to join a collaboration. We currently investigate how knowledge discovery and data mining techniques (Anand and Büchner, 1998) can be exploited in order to semi-automate the process of mapping definition in PDM scenarios.

ACKNOWLEDGEMENTS

This research has partly been funded by the IiMB Esprit project Nº 21108, which is greatly acknowledged.

REFERENCES

Anand, S. S and Büchner, A. G. "Decision Support Using Data Mining", Financial Times Pitman Publishers, 1998.

Büchner, A. G., Bell, D. A. and Hughes, J. G., 1998, "A Contextualised Object Data Model Based on Semantic Values", Proc. *11th Int'l Conf. on Parallel and Distributed Computing Systems*, pp. 171-176.

Chen, H and Lynch, K. J., 1992, "Automatic Construction of Networks of Concepts Characterizing Document Databases, *IEEE Trans. on Systems, Man, and Cybernetics*, Vol. 22, No. 5, pp. 885-902.

Dubitzky, W., Büchner, A. G., Hughes, J. G. and Bell, D. A., 1999, "Towards Concept-Oriented Databases", *Journal on Data and Knowledge Engineering*, forthcoming.

Fox, M. S., Chionglo, J. F. and Fadel, F. G. 1993, "Towards Common Sense Modelling of an Enterprise", *Proc. 2nd Industrial Engineering Research Conf., pp.425-429.*

Gruber, T. R., 1995, "Towards Principles for the Design of Ontologies Used for Knowledge Sharing, *Int'l Journal of Human Computer Studies*, Vol. 43, No. 5/6, pp. 908-928.

Hammer, J. and McLeod, D., 1993, "An Approach to resolving Semantic Heterogeneity in the Federation of Autonomous, Heterogeneous, Database Systems, *Int'l Journal of Intelligent and Cooperative Information Systems*, Vol. 2, No. 1, pp. 51-83.

Kahng, J. and McLeod, D., 1998, "Dynamic Classificational Ontologies: Mediation of Information Sharing in Cooperative Federated Database Systems", *M.P. Papazoglou, G. Schlageter (eds). Cooperative Information Systems*, pp. 179-203.

Kashyap, V. and Sheth, A., 1998, "Semantic Heterogeneity in Global Information Systems: the Role of Metadata, Context and Ontologies", *M.P. Papazoglou, G. Schlageter (eds). Cooperative Information Systems*, pp. 139-178.

Lee, J. L., Madnick, S. E. and Siegel, M. D., 1996, "Conceptualizing Semantic Interoperability: a Perspective from the Knowledge Level", *Int'l Journal of Cooperative Information Systems*, Vol. 5, No. 4, pp. 367-394.

Mena, E., Kashyap, V, Sheth, A. and Illarramendi, A, 1996, "OBSERVER: an Approach for Query Processing in Global Information Systems based on Interoperation across pre-existing Ontologies", *Proc. 1st Int'l Conf. on Cooperative Information Systems*, pp. 14-25.

Michalski, R., 1993, "Inferential Theory of Learning as Conceptual Basis for Multi-strategy Learning, *Machine Learning*, Vol. 11, No. 2/3, pp. 111-151.

Mizohuchi, R. and Ikeda, M., 1996, "Towards Ontology Engineering", Technical Report AI-TR-96-1, The Institute of Scientific and Industrial Research, Osaka University, Japan.

Ranta, M., Büchner, A. G., Mäntylä, M. and Hughes, J. G., 1999, "Collaboration in Heterogeneous Environments along Product Processes", *Proc. 4th World Conf. on Integrated Design & Process Technology*.

Wiederhold, G., 1992, "Mediators in the Architecture of Future Information Systems", *IEEE Computer*, Vol. 25, No. 3, pp. 38-49.

Yu, C., Sun, W., Dao, S. and Keirsey, D, 1991, "Determining Relationships Among Attributes for Interoperability of Multi-Database Systems", *Proc. 1st Int'l Workshop on Interoperability in Multidatabase Systems*, pp. 251-257.