



*D*épartement
*I*ntelligence
*A*rtificielle
*e*t *M*édecine

DIAM Rapport Interne RI-95-148¹

August 19, 1995

7 pages

Methodological Principles for Structuring an “Ontology”

J. Bouaud, B. Bachimont, J. Charlet, and P. Zweigenbaum
{jb,bb,jc,pz}@biomath.jussieu.fr

*Département Intelligence Artificielle et Médecine.
Département de Biomathématiques – CHU Pitié-Salpêtrière PARIS 6
& Service d’Informatique Médicale de l’Assistance Publique-Hôpitaux de Paris*

*Adresse postale : SIM, 91, Boulevard de l’Hôpital, F-75634 Paris cedex 13. France
Tél. (+33)/0 1 40 77 96 17. Fax. (+33)/0 1 45 86 56 85.
E-mail. diam@biomath.jussieu.fr. Web. <http://www.biomath.jussieu.fr/>*

¹Presented at the IJCAI’95 Workshop on “Basic Ontological Issues in Knowledge Sharing”, August 19, 1995, Montréal, Canada.

Abstract

The knowledge used in most AI applications does not rely on a formal model of the domain. Therefore, it has to be normalized to ensure that the formal exploitation of its representation conforms to its meaning in the domain. Considering the intensional (non extensional) nature of concepts, which reflects the essences of the objects they denote, this normalization relies on a commitment on type definitions by necessary and sufficient conditions at the knowledge level. Our claim is that the taxonomic structure that accounts for the intensional nature of the ontology can be nothing but a tree. From this starting point, we derive methodological principles to constrain and justify the structuring of ontological types. Based on this methodology, we advocate understandability of an ontology rather than a putative reusability.

1 Introduction

The ontology is the heart of any knowledge description: knowledge is intimately related to the ontology, since it is necessarily expressed in terms of this ontology. Therefore, designing the ontology of a domain is a key issue for knowledge representation [Lenat and Guha, 1990; Gruber, 1993]. Designing an ontology corresponds to conceptual modelling; at least, the categories of objects that are considered to exist in a given domain have to be identified. These categories are structured by a subsumption relation and constitute a taxonomic hierarchy. But, when developing an ontology, one is faced with numerous problems and representational alternatives. One has to make choices, or "ontological commitments" [Gruber, 1993], for which very few guidelines exist in the literature [Brachman *et al.*, 1991].

In this paper, we present methodological principles to control the structuring and the use of an ontology in the context of knowledge-based systems [Bouaud *et al.*, 1994]. This work results from our experience in MENELAS, a natural language understanding system in the medical domain [Zweigenbaum *et al.*, 1995]. Medicine is a typical domain where knowledge is by nature descriptive, mainly expressed in natural language, and not formalized. Though a technical and restricted domain, there does not exist a set of conceptual descriptors from which every concept could be derived. Since a key issue in representing knowledge is to ensure that its formal exploitation at the symbol level by a program conforms to its meaning at the knowledge level for people, descriptive, informal knowledge has to be normalized. Concepts must be elicited and a commitment on their meaning is needed. Noting the intensional nature of types on the one hand, and the need for type definitions by necessary and sufficient conditions in order to be able to use a taxonomy on the other hand, we argue that this commitment has to take place at the knowledge level and is tied up to the essence of objects, thus to the

ontology. Assuming that types are defined through their essence, our claim is that the taxonomic structure which accounts for the intensional nature of the ontology can be nothing but a tree.

2 Taxonomies for Knowledge Representation

2.1 Conceptual type hierarchies

Within the framework of knowledge representation, types aim at representing the "categories of thought" and are called *conceptual types*. Types are organized in a taxonomy, often considered as the ontology, according to their levels of generality. How a taxonomy accounts for knowledge and can be used by a knowledge-based system is a knowledge acquisition issue. Understanding the very nature of the subsumption represented in a taxonomy is fundamental when using the taxonomy [Woods, 1991].

2.2 Ontology Acquisition

The goal of the knowledge acquisition task is to obtain a formal representation of human knowledge that a computer can use at the symbol level in conformance with its meaning at the so-called knowledge level [Newell, 1982]. A same symbolic expression is given a computational interpretation at the symbol level like an instruction, and a semantic interpretation at the knowledge level by people like a piece of knowledge. Although we can "read" a representation, this does not imply that a program does the same, and a clear distinction between the two interpretation modes must be kept. The same matters hold for taxonomies and a great attention must be paid to this important issue since types are the basic building blocks from which more complex representations are elaborated.

As a result, there are two related key issues pertaining to knowledge acquisition in the context of a knowledge-based application. The first one is the problem of the acquisition of the conceptual type hierarchy. The second one concerns its use. If the classical logical interpretation of taxonomies is clear, there are no means to control that the semantic, non-formal interpretation of a type hierarchy by people is the same. Every knowledge representation relies on the strong assumption that they are. When building the ontology of a domain, the designer has to tackle carefully this key issue, which becomes crucial with atomic types.

2.3 Reusing Ontologies

The reusability and sharing of ontologies is the foundation of numerous projects. CYC [Lenat and Guha, 1990] is one of the most famous. The ontology is the central part of the CYC knowledge base. For the authors this ontology is rather universal than reusable insofar as it carves up the universe to represent enough common-sense and encyclopedic knowledge to support natural language capabilities. KIF [Genesereth and Fikes, 1992] is a logic-based, implementation-independent language for knowledge representation. It provides for the definition of objects of a conceptualization and therefore the

description of an ontology. Assuming these ontologies are reusable, Ontolingua [Gruber, 1993] makes these KIF ontologies portable since it translates them into several knowledge representation system idioms, thus enabling knowledge sharing.

As pointed out previously, most of the problems that arise when one tries to use an existing taxonomy deal with the semantics of subsumption and the interpretation of actual links. The formal semantics given by a program should not be contradictory with human interpretation. This is especially true with atomic types when the "natural language" labels bring its inherent ambiguity and variability in interpretation. "Ontological commitments" [Gruber, 1993; Guarino *et al.*, 1994] are needed to constrain the possible interpretations of an ontology so that they can be understandable and usable.

3 Modelling an Ontology

When building knowledge-based systems, the most important point is to ensure that the formal exploitation of the knowledge representation conforms to its meaning in the domain. On the one hand, the issue of knowledge representation is to design a formal representational system to account for the cognitive aspects of knowledge. On the other hand, the issue of knowledge acquisition is to constrain knowledge so that it can fit into the formal representational system. To achieve this goal, a *normalization* of knowledge is a necessary step during which the *ontology* of the domain must be elicited.

3.1 Fundamental Issues of Knowledge Representation

Generally speaking, AI consists in solving problems using human knowledge that is usually expressed in natural language. Since there is no other operational expression of this knowledge—*i.e.* the domain is not formal—we have to tackle the semantic richness of language and its multiple potential interpretations. For example, there is no formal theory for modelling common-sense knowledge or complex empirical medical knowledge. Knowledge people have about objects in a domain is mainly *descriptive*. It relies on the notion of prototype, whose logical status is not clear enough to be usable by a computer. People reason with this descriptive knowledge using the semantic laws of meaning such as analogy, metaphor, prototypical inference, etc. In non-formal domains, we have no complete definitions for the categories of objects accounted by descriptive knowledge.

In order to have a computer deal with this kind of knowledge, we need to normalize it: we need to agree on the notions it mobilizes and on their use. Knowledge must be considered as an objective notion, *i.e.* shared by people. Basically, we have to fix the meaning of terms so that we all understand the same thing when we use a term, and so that we all agree when a computer infers on it. Knowledge will be adequately described if we cater for explicitly rendering its objective and shared aspects. Normalizing is nothing but this. The use of terms will as a consequence no more rely on the full power of semantic interpretation, but on the manipulation of their explicit descriptions.

3.2 Normalizing by Necessary Conditions

A usual way of normalizing descriptive knowledge consists in stating the necessary relations between domain notions. Let us recall that types are categories of thought. Now, let us consider the notion of pain, which we all know and for which we want to confront our understanding. However, we have no means to access all pains in the real world in order to see objectively whether we are able to agree or not on what pain actually is. The basic reason is that such a notion is *intensional*. Objects are considered through the way we think of them, not through their very nature, since it is not accessible—except for God.

The distinction between *intension* and *extension* was pointed out long ago by philosophers. The denotation of a type is its extension, *i.e.* the set of objects characterized by the type. Basically, there is no explicit relation between the intension of a type, its meaning, and its extension. Its extension is not a characterization of the type. Moreover, there are no means to yield the extension from intension in non-formal domains, although this is not the case in artificial domains. Citing Sowa [Sowa, 1984]: "*The type lattice represents categories of thought and the lattice of sets and subsets represents collection of existing things. The two lattices are not isomorphic and the denotation operator that maps one into the other is neither one-to-one nor onto.*"

The famous Frege's example of *Morning star* and *Evening star* illustrates this. These two stars are in fact the same, and correspond to the planet Venus. Hence, these two concepts are extensionally equivalent. But, obviously they are not intensionally equivalent.

Because of the lack of an explicit relation between the intensional definition of a type and its extension, one must be careful not to confuse the properties that characterize extensions with those that concern intensions. There are a priori no definitions of types relying on universal necessary and sufficient conditions; it is not possible to relate biunivocally objects to concepts of them. Since only intensions are available, we cannot but agree only on intensions in the normalizing process.

For instance, pain can be considered as a human affection. If we accept this (partial) definition for pain, then we can all use this notion to understand each other. But this necessary condition is still intensional, since the necessary inference from pain to human affection is drawn only by considering the meaning of pain and not by studying every individual pain existing in the world. Such subsumption-based inferences do not rely on the inclusion of extensions but on the meanings of the notions. A first step of the normalization process is then to agree on and fix the necessary intensional conditions of types.

3.3 Normalizing by Necessary and Sufficient Conditions

Subsumption by the use of necessary conditions is the basis of all taxonomic organizations. The laws that rule intensional subsumption are compatible with the minimal logical interpretation of a taxonomic link in case

of such a partial definition. Such taxonomies, based on necessary conditions, are used for property inheritance.

But taxonomies only built from necessary conditions cannot allow classification. Their basic properties do not enable the addition of a new type while being sure there are no ambiguities. The defining conditions are not *sufficient* conditions: the differentiae between a type and its genus are only necessary conditions. There is no equivalence between the type and its definition. The result is that it is not possible to use the taxonomy either to look for a type through its definition or add a new type. Partial definitions are not enough to be sure that we talk about the same notions. We need sufficient conditions in order to agree that a notion corresponds to the same knowledge. Being aware of the type subsumption nature is fundamental when one tries to understand, use or extend a given taxonomy—especially with atomic types where differentiae are not explicit.

The knowledge normalization must be carried on in order to assign complete definitions to types. Since types are intrinsically intensional, there must be a trade off to determine which properties will be considered as *definitional*. Normalizing by complete definitions is then extracting the *essence* of the notions that are used—*i.e.* their basic meaning on which we agree and that we can use in the same unique way.

3.4 Characterizing the Essential Properties

By assigning complete definitions to notions, we aim at capturing their essence, thus at determining their *essential* properties. Capturing the essence of objects typically falls within the province of *Ontology*. The term "ontology" comes from philosophy and refers to "the study of being as such". An ontology may be defined as the set of objects that exist in a domain. So, building an ontology is to decide which objects one retains as existing in the studied domain. The notion of object here is intensional; this step corresponds to the conceptualization of the domain. Thus, defining types by deciding what are their essential characteristics amounts to building the ontology of the domain.

However, as pointed out previously, when the domain is not formalized, there are no methods that give the essential properties. One has to choose carefully the properties to consider as essential. A complete definition holds when no counter-examples occur. If we consider that an object belongs to a type but violates its definition, then the definition must be revised. As a consequence, the meaning of a type is the result of a negotiation process. In one case, a given property is essential and the other not, and in some other case, the contrary occurs. One must adopt a point of view, an interpretation, to be able to define essential properties. As a result, a careful distinction must be made between essential and *incidental*, *i.e.* non essential, properties [Brachman *et al.*, 1991]. This distinction is important to clarify what must be considered as the basic meaning of a type: its essence. The essential characteristics of objects are those such that, if a defined object loses only one of them, it no longer exists as such. These properties must be true by

intension as long as the object exists. These properties are definitional, in the sense that objects that bear them are recognized as members of the type in every possible world.

The non essential properties of types are not part of the ontology. Knowledge associated with them must be declared separately in a knowledge base. The associated inference mechanisms are then specific.

3.5 Essence and Taxonomy

When the knowledge normalization process is completed, types are defined in terms of necessary and sufficient conditions expressed through essential, intensional, properties. By definition, an essence is unique, and therefore a type can only have a unique definition. Moreover, unicity of essence can be argued as follows. Let us consider a notion which has two essences: that means that there are two different sets of necessary and sufficient conditions to determine it. We argue that either these two definitions determine two different notions, or are intensionally equivalent. Let us assume first that there exists a possible world where the two essences denote different objects. If, in this context, we consider nevertheless that both meanings are identical, we rely on the assumption of a particular world where they denote the same objects. But, this property is contingent to this particular world and is not implied by the definitions: these definitions do not correspond to the same notion because there exists a possible world where their extensions are different. On the other hand, if these notions have an identical extension in every possible world, then they are intensionally equivalent. Thus, a first consequence of the unicity of essence is that the extensional equivalence (in a particular world: contingent equivalence) does not enforce intensional equivalence (in every possible world: essential equivalence).

Having stated that types are defined by necessary and sufficient conditions, one could hope that the relationships between types replicate the formal properties of subsumption hierarchy and that an ontology consists in a type lattice. However, such an expectation is not legitimate and can be the source of many misuses and misunderstandings of an ontology designed in an informal domain. Although surprising at a first glance, this fact is easy to understand. It relies on two observations. The first is that an ontology is a subsumption hierarchy: a type inherits the properties of the types that subsume it. Hence, a type consists of its specific properties added to the others deduced from its subsumers. The second fact is that there is no compositionality (in the usual meaning of formal semantics) of properties in a type. For example, the property "diuretic function" does not mean the same thing when it qualifies a drug (therapeutic function) or a molecule (biologic function). Hence, the type "drug with diuretic function and with molecule XYZ" will not have the same meaning as "drug with molecule XYZ and diuretic function". Moreover if properties are not compositional, the hierarchy cannot be a lattice, but just a tree: given two paths from the same node to the root, either they are identical, or they combine properties in a different order. In such a case, the end node

has a different meaning according to each path. Hence, there can be only one path between a node and the root. The main reason for this is that the intension of a given differentiating property, though expressed in the same terms, is not independent of its subsumers, which create an interpretation context for the term. The meaning of a property must be understood through its position in the ontology; the same property in different positions, *i.e.* definitions, does not have the same meaning. This fact comes from the natural language in which concepts in an informal domain are expressed at the knowledge level: natural language semantics is not compositional.

4 Methodological Principles for Ontology Structuring

Based on the previous section, we consider that the induced taxonomic structure for an ontology is a tree. We present here methodological principles to control the acquisition and structuring of the ontology [Bouaud *et al.*, 1994].

4.1 Principles of Structural Organization

A focal point for examining the issues of building taxonomies is the local system composed of a type and its subsumees, the *parent* and its *children*. We shall call such a system a *local tree*. By stating principles applicable to every local tree, we act on the whole structure.

Four principles are formulated as relations that must hold between types in a local tree. The first two principles are directly related to Aristotle's definition principles by genus and differentiae. The two others deal with the relations between siblings. These principles are exemplified by adding the type *mental state* as a subtype of *state*.

P1: Similarity Principle

This first "similarity" principle controls the conformance of a type to its genus. A child must share the type of its parent. This is the basic meaning of the taxonomic link. At the level of the whole local tree, this also means that all children share a common meaning. Being of the parent's type is a necessary condition for a child. For instance, a *state* consists in a description of the world; *e.g.*, it may be true or false. A *state* asserts something about objects in the world. A *mental state* consists in a description about *mental objects* which are objects of the world. As it shares a common meaning with *state*, type *mental state* is a subtype of *state*. Being a *state* is a necessary condition to be a *mental state*. The kind of object involved in the *state* is therefore a good candidate for providing the semantic axis along which children of *state* will be defined. Type *mental state* specializes *state* along this axis.

P2: Specificity Principle

The role of the second "specificity" principle is to justify the differentiae of a type with respect to its genus. This specific difference in the context of the parent expresses necessary and sufficient conditions for the child. This often takes the form of a distinctive property that the child bears, or (*e.g.*, for actions) as a specific thematic

role always associated to the child. The specific difference of *mental state* relatively to *state* is that it involves *mental objects*, which are specific objects of the world. Being a *state that involves mental objects* is the complete definition of a *mental state*.

P3: Opposition Principle

The third principle of "opposition" further characterizes the relations between siblings. This principle is fundamental since the tree structure derives from it. Siblings are organized in a system of oppositions. Each child of a type is opposed to the other children of the same type. All children must be pairly incompatible. The difference between a child and each of its siblings must then be stated. Keeping the same example again and assuming that, *e.g.*, a *physical state* is also a subtype of *state*, the difference between *mental state* and *physical state*, is that the former involves a *mental object* whereas the latter involves a *physical object*, both kinds of objects being distinct and incompatible in the ontology.

P4: Unique Semantic Axis Principle

The last "unique semantic axis" principle provides a criterion on which children are to be compared to their parent. In order to satisfy the two difference principles together (P2, P3), one can simply constrain all subtypes of a given concept type to differ from their parent on a common dimension or axis, each subtype having an exclusive value for this dimension. For each type, a semantic axis along which children are projected must be chosen. In the previous example, the decomposition of *state* is performed according to the kind of *objects* involved. There are three kinds of objects in the ontology: *ideal*, *physical*, and *mental* objects. *State* is then decomposed in three subtypes: *ideal state*, *physical state*, and *mental state*.

4.2 The Ontology Tree

The combined use of the previous methodological principles at the node level in the taxonomy ensure that the whole taxonomy is a tree, assuming a unique root like "entity". As the intensional type definitions are supposed to be complete, this taxonomy becomes a *decision tree*. It can be (re)used and understood: placing or searching a type in this ontological taxonomy can be done easily without ambiguities.

MENELAS knowledge bases are currently under development. The domain is related to coronary diseases. Basically, we mainly aim at representing *medical acts* and *pathological states* of a patient with many related notions. In the current state, the ontology contains about 1800 types and 250 relations with knowledge related to them.

A tool with a window interface (KMT) has been developed to help in the acquisition of an ontology. It follows the guidelines of the methodology in order to manage, with the designer, the construction of the domain ontology. Pieces of knowledge related to types are captured here. KMT is also a browser for the ontology and knowledge bases. Moreover, comprehensive documentation including the ontology itself and knowledge bases is generated as HTML pages for consultation.

5 Discussion

5.1 The structure of an ontology

Our main claim that the ontology has a tree structure relies on the hypothesis that the notion of type is intensional and on the difficulty to solve intensional equivalence. In our structuring methodology, the third opposition principle (P3) is fundamental here since the tree structure is a consequence of its application, forcing mutual exclusion between the sibling subtypes of a parent. This principle of opposition is imported from structural linguistics [de Saussure, 1985; Eco, 1988], where the meaning of a term consists in the differences it has with others. Roughly speaking, to express something, we choose a linguistic realization and, as a consequence, discard the others. The meaning of the selected term comes from the fact that we invest this term with something specific that other terms do not have. In this sense, language is a system of differences, and so is knowledge expressed by this medium. In a similar way, a formal system such as the type hierarchy constitutes a system of differences: a type acquires its meaning from its position with respect to other types.

The fourth unique semantic axis principle (P4) seems desirable, but less crucial than the others. However, it is complementary to the opposition principle (P3) and helpful in practice. P3 imposes that sibling types be opposed. But guaranteeing that sibling types are intensionally opposed is close to the issue we already discussed about intensional equivalence. Knowing that intensions are distinct does not mean that they are opposed. Thus, the only way to know whether two notions are opposite is to project them along the same semantic axis. In this respect, P4 constitutes an efficient principle to enforce the satisfaction of P3. Moreover, the choice of a unique semantic axis for decomposing all the proper subtypes of a type prevents the use of the same differentiating criterion further down. It is a mean to avoid the production of redundant definitions and temptations to confuse types. To a certain extent, satisfying P3 can be viewed as a side effect of the satisfaction of P4.

The previous methodological principles aim at controlling that the taxonomy built reflects the ontology of a domain. Whereas these principles are necessary to structure the ontology, they are not sufficient to complete its acquisition. They do not provide instructions, or at least strong guidelines, for eliciting an ontology as they are more constraining than constructive.

Nevertheless, many other works concerned with ontologies do not impose such a constraint on them. Arbitrary tangled hierarchies allowing multiple parents for a type are usually considered. This is consistent with the model-theoretic semantics of predicate logic based on a universe of discourse. In this condition of an extensional interpretation of types, the type structure is equivalent to the inclusion lattice of extensions. This method is clearly adequate in formal domains where the intension gives the extension. But as pointed out previously, in non formal domains where only human knowledge is available, this is not the case.

Though, CYC, for instance, does not consider this and

has a practical engineering approach to ontology. While they want to represent knowledge, the authors' position is clearly extensional: "...there appear to be several categories of interest in the world, several worthwhile ways of dividing up the set of IndividualObjects." [Lenat and Guha, 1990].

Nevertheless, and not so surprisingly, the tree structure which can be seen as a constraint can be found in several practical as well as theoretical works in "knowledge representation". In their work on formalizing ontological commitments, Guarino *et al.*'s [Guarino *et al.*, 1994] definition of *well-founded ontological commitment* requires that the subsumption hierarchy of "substantial sortal types" is a tree. More practically, in applications using "concept forming languages" like KL-ONE [Brachman *et al.*, 1991] or Conceptual Graphs [Sowa, 1984], the hierarchy of atomic types, from which more complex descriptions are build, is often a tree, even if defined descriptions constitute an arbitrary hierarchy (cf. GALEN [Rector *et al.*, 1992]). In the same way, representational primitives are always considered as mutually exclusive. Considering the historical background, zoological or botanical taxonomies are always trees.

5.2 Reusability of an ontology

Claiming that an ontology is reusable entails that it can be used to express knowledge for other tasks than the one for which it was designed. Noting the previous discussion, we advocate that the essential properties that are characterized result from a point of view, an interpretation, of the domain notions. The consequence is that the ontology is dependent on this point of view. In non-formal domains, like medicine, this point of view is likely dependent on the application.

When Gruber [Gruber, 1993] provides examples of KIF "reusable" ontologies, they are developed in formalized, mathematical domains. But, in such domains for which there exist formal models, concepts are defined analytically and only hold by/through their definitions. Since they are not subject to interpretation, such ontologies can therefore be considered as reusable. Within CYC [Lenat and Guha, 1990], the use of common-sense knowledge has to go through the constitution of more specific "microtheories", *i.e.* a set of sentences which are true in a certain context. So, the CYC microtheories constitute so many points of view about a domain which can be mutually inconsistent. This highlights the dependency between an ontology and the context of its design. For instance, in non-formal domains like medicine, the fact that there could be multiple ontologies of a domain questions their reusability. This gives rise to the characterization of the domain limits. Does the domain include the task?

5.3 Understandability of an ontology

As explained, the goal of knowledge representation is to obtain the best possible correspondence between the computational properties induced by the structuring of the knowledge representation and the properties that human interpretation assigns to this knowledge. In the absence of precise guidelines, humans may load concepts

and their structuring with more meaning than is formally defined to the system; this is why ontological commitments are needed.

In our context, since the ontological types rest on complete definitions, the ontology can be used for classification. The tree structure enforces stronger constraints. It constitutes a decision tree: there are no ambiguities as there is a unique path to describe a type which excludes the others. Considering an existing ontology built this way, it becomes easy to understand it, use it and augment it, because each type is justified and opposed to others. Comments associated to each of the four principles provide useful documentation for each type. They correspond to the commitment made on the type meaning, which can be shared by others. Citing Guarino [Guarino, 1994], "... a rigorous ontological foundation for knowledge representation can improve the quality of the knowledge engineering process, making it easier to build at least understandable (if not reusable) knowledge base.", we advocate understandability of an ontology rather than a putative reusability.

6 Conclusion

Methodological principles to structure the ontology of a non-formal domain have been presented. These principles have been used in the design of a real ontology of more than 1800 types and 250 relations. Moreover this ontology has been effectively used and validated through a medical application. Our claim is that, in non-formal domains like medicine, knowledge is rather descriptive than formal and therefore needs to be normalized in order to be used by a program while keeping the same meaning. This normalization consists in finding an agreement on the meaning of domain notions and in assigning complete definitions to such intensional types. Elaborating such definitions amounts to characterizing the essences of types, *i.e.* to eliciting the ontology of the domain. Assuming the unicity of the essence of an intensional type, we argue that the ontology must be a subsumption tree.

The characterization of essential, definitional, properties is not a deterministic process. As a consequence, one must adopt a point of view on the domain notions. In our opinion, this fact has a serious impact on the possibility to reuse directly existing ontologies. However, because it has been built according to structuring principles, we think that our ontology can be understood without ambiguities and shared: the meaning of each type relies on a unique complete definition justified with respect to all the other types in the taxonomy through a system of differences.

References

- [Bouaud *et al.*, 1994] Jacques Bouaud, Bruno Bachimont, Jean Charlet, and Pierre Zweigenbaum. Acquisition and structuring of an ontology within conceptual graphs. In *Proceedings of ICCS'94 Workshop on Knowledge Acquisition using Conceptual Graph Theory*, pages 1–25, University of Maryland, College Park, MD, August 19 1994.
- [Brachman *et al.*, 1991] R. J. Brachman, D. L. McGuinness, P. F. Patel-Schneider, L. A. Resnik, and A. Borgida. Living with Classic: When and how to use a KL-ONE-like language. In Sowa [1991], chapter 14, pages 401–456.
- [de Saussure, 1985] Ferdinand de Saussure. *Cours de Linguistique Générale*. Payot, Paris, 1985.
- [Eco, 1988] Umberto Eco. *Le signe; Histoire et analyse d'un concept*. Editions Labor, Bruxelles, 1988.
- [Genesereth and Fikes, 1992] M. R. Genesereth and R. E. Fikes. Knowledge Interchange Format version 3.0 reference manual. Technical Report Logic-92-1, Computer Science Department, Stanford University, June 1992.
- [Gruber, 1993] T. R. Gruber. A translation approach to portable ontology specifications. *Knowledge Acquisition*, 5:199–220, 1993.
- [Guarino *et al.*, 1994] N. Guarino, M. Carrara, and P. Giaretta. Formalizing ontological commitments. In *Proceedings of the 12th National Conference on Artificial Intelligence*, pages 560–567, Seattle, WA, July 31–August 4, 1994. American Association for Artificial Intelligence.
- [Guarino, 1994] N. Guarino. The ontological level. In R. Casati, B. Smith, and G. White, editors, *Philosophy and the Cognitive Sciences*. Hölder-Pichler-Tempsky, Vienna, 1994. *To appear*.
- [Lenat and Guha, 1990] D. B. Lenat and R. V. Guha. *Building Large Knowledge-Based Systems, Representation and inference in the Cyc Project*. Addison Wesley, 1990.
- [Newell, 1982] A. Newell. The knowledge level. *Artificial Intelligence*, 18:87–127, 1982.
- [Rector *et al.*, 1992] A. L. Rector, W. A. Nowlan, and S. Kay. Conceptual knowledge: the core of medical information systems. In K C Lun, P Degoulet, T Piemme, and O Rienhoff, editors, *Proceedings of the 7th World Congress on Medical Informatics*, pages 1420–1426, Geneva, Switzerland, 1992. North Holland.
- [Sowa, 1984] J. F. Sowa. *Conceptual Structures: Information Processing in Mind and Machine*. Addison-Wesley, London, 1984.
- [Sowa, 1991] J. F. Sowa, editor. *Principles of Semantic Networks*. Morgan Kaufmann Publishers, 1991.
- [Woods, 1991] W. A. Woods. Understanding subsumption and taxonomy: A framework for progress. In Sowa [1991], chapter 1, pages 45–94.
- [Zweigenbaum *et al.*, 1995] Pierre Zweigenbaum, Bruno Bachimont, Jacques Bouaud, Jean Charlet, and Jean-François Boisvieux. Issues in the structuring and acquisition of an ontology for medical language understanding. *Methods of Information in Medicine*, 34(1), 1995.