5. Formal Ontology and the Needs of Ontological Semantics

In this chapter, we briefly discuss the philosophical and formal approaches to ontology as they relate to the needs of ontological semantics. We try to position ontological semantics within the field of formal ontology, though we do not attempt to catalogue all the existing ontology development projects. A comprehensive survey is difficult to accomplish, largely because few ontologies are accessible for comparison either in their entirety or in terms of their architecture and acquisition.

What we attempt to do here is four things. First, we place ontology in the context of the philosophical discipline of metaphysics. Metaphysics deals with the most basic categories, and it is a scary thought, in the spirit of the tool influencing the object under observation, that a different category choice may change one's entire picture of the world. We will see that some claims associated with metaphysics pertain to our concerns but many do not, and we discuss briefly to what extent that should concern us.

Second, we address a number of formal issues in ontology, as developed in the field and as they pertain to our needs. Third, we discuss the important distinction between ontology and natural language, primarily in relation to the phenomenon of ambiguity. And finally, we offer a wish list from ontological semantics to be considered as an extended agenda of formal ontology.

5.1 Ontology and Metaphysics

Guarino (1998a: 4) suggests a distinction between "Ontology" and "ontology." The former is an academic discipline within philosophy, and we will use for it a more appropriate name, metaphysics, a name that many scholars have been hesitating to use since the positivists "made it into a term of abuse [accusing it of] isolating statements about mental life from any possibility of verification or falsification in the public world" (Kenny, 1989: ix). We would like to restore the term to its legitimate domain. Like Kenny and other authors of recent works reinstating metaphysics (see, for instance, Jubien 1997, Loux 1998), we must avoid "the confusion that can be generated by bad metaphysics" and crave "the clarity which is impossible without good metaphysics" (Kenny 1989: ix).

Metaphysics is a traditional philosophical discipline, perhaps the most ancient one, that can be traced back at least to Aristotle. It "attempt[s] to provide an account of being qua being (Loux 1998: x). In accounting for being, metaphysics delineates "the categories of being" (ibid.). These categories form the basic philosophical concepts, and those, in turn,

> "underlie all genuine scientific inquiry because science cannot even begin in the absence of philosophical assumptions and presuppositions. These assumptions are generally not stated explicitly and so may not even be noticed by practicing scientists or students of science. But they are there.
>
> As an example, physics presupposes the following three things: (1) that there exists a physical reality independent of our mental states; (2) that the interactions of the stuff constituting this reality conform to certain general laws; and (3) that we are capable of grasping physical laws and obtaining evidence that favors or disfavors specific proposed laws.... The first two are metaphysical in nature while the third is epistemological.... [T]hey

are not at all self-evidently true...... [T]hey are not themselves part of the subject matter of physics." (Jubien 1997: 3-4)

Now, the list of basic categories proposed by metaphysics, the "'official' philosophical inventory of things that are... is usually called an ontology" (Loux 1998: 15; cf. also Bergman 1992, Grossman 1992, Chisholm 1996). This is, of course, Guarino's lower-case ontology; it is the sense also, in which our ontology, that of ontological semantics, exists.

The philosophical discipline of metaphysics faces a number of difficult, empirically unsolvable issues. The central one is the existence of properties, on which philosophers have always been divided into two basic camps (by now, with innumerable gradations), viz., the realists and naturalists/nominalists. The realists recognize the existence of two types of entities, individuals, which exist in time and space, and properties of individuals, which are abstract and, as such, atemporal and aspatial. The naturalists recognize the existence of just individuals.

Both camps have serious problems. The realists have to cope with two different kinds of existence, including the unobservable and directly unverifiable existence of abstract properties. Free of that concern, the naturalists have a hard time explaining away the similarity of two individuals in terms of purely physical existence. Over the centuries, the battle has seen many ingenious proposals on both sides, but the issue will not and possibly cannot go away.

How does this serious problem affect an ontology? What impact does it have on ontological semantics? Had it had sided with the naturalists, it would not have had any properties in the ontology but, of course, it does. Furthermore, the ontology in ontological semantics includes abstract and non-existent entities alongside physical entities. In fact, a very large branch of this ontology, in all its implementations, is devoted to mental objects and another, to mental processes. A close comparison of the ontological node for a typical mental entity and a typical physical entity will show that the fillers of the pertinent properties do reflect the distinction, e.g., a mental process will manipulate mental objects and the physical, only physical.

What ontological semantics aims to reflect is the use of concepts by humans as they see it, introspectively and speculatively; and people do talk about properties, fictional entities (unicorns or Sherlock Holmes), and abstract entities as existing. For us, however, the decision to include the abstract and fictional entities in the ontology is not motivated by the fact that these entities can be referred to in a natural language. Rather, we believe that languages can refer to them precisely because people have these concepts in their universe.

Constructing an ontology should not be viewed as a task for metaphysics. Instead, this is a problem of representing knowledge and, thus, belongs in epistemology. That is, the object of study here is human knowledge about entities, not entities themselves. Inasmuch as humans know about unicorns as well as, say, goats, the respective concepts in the ontology have the same status. It is not, therefore, important for a constructed ontology that unicorns do not "exist" and goats do. The main criterion for inclusion is a consensus among the members of the community using the ontology (and generating and understanding documents to be processed with the help of the ontology) concerning the properties of a concept. The basic conclusion of this line of thought is that epistemology (and, therefore, any constructed ontology) is neutral with respect to the major metaphysical issue of existence.[58]

There is another claim routinely made about metaphysics that strikes us as importantly incorrect, even if it is made by those who, like ourselves, recognize the importance of metaphysics. "[M]etaphysics," Loux states, "is the most general of all disciplines; its aim is to identify the nature and structure of all that there is" (1998: x). But is it really? On the one hand, the very top levels of an ontology should contain the most basic—and general—categories that no particular area of research will claim as their own (Bateman refers to these levels as the "Upper Model" (see Bateman 1990, 1993; Hovy and Nirenburg 1992 reserves the term 'ontology' for the top levels only, and uses 'domain model' for the lower levels). On the other hand, metaphysics is not responsible for "the nature and structure" of, say, microbiological entities: microbiology, a part of biology, is. So, like all other disciplines, metaphysics has its specific and pretty limited domain: it is that of the universally shared basic categories.

The choice of categories for use by a science, while definitely influenced by metaphysics, whose categories are usually involved, is done by the philosophy of that science and not by metaphysics *per se*. It is also quite realistic to think, even if a particular example may be hard to come by, that a specific "low-level" discipline may stumble upon an important general property and claim a slot for it among the basic properties of metaphysics. Practically, it means that designing ontology is not a simple matter of putting the metaphysical categories on top and letting specific disciplines and domains add descendants.

5.2 Formal Ontology

Formal ontology is still a developing discipline, and a discipline which is clearly distinct from ontological semantics in its perspective, so we will not presume here to review or to analyze it in its entirety nor to anticipate or to prescribe the directions of its development. Instead, we will review briefly some of the more pertinent aspects of formal ontology, those bearing theoretically and practically on ontological semantics.

5.2.1 Formal Basis of Ontology

While metaphysics is an ancient discipline and ontology has been commented upon by every major philosopher of modernity, most influentially perhaps by Kant and Hegel, it is Husserl (1900-01) who is usually credited with founding formal ontology. He saw the field as parallel to formal logic: "[f]ormal logic deals with the interconnections of truths... [while] formal ontology deals with the interconnections of things, with objects and properties, parts and wholes, relations and collectives" (Smith 1998: 19).

Formal ontology is seen as being founded on the mathematical disciplines of mereology, which studies the relations between parts and wholes, theory of dependence, and topology. There is a body of work studying these disciplines, their relations to ontologies, and issues with their appli-

---

58. It is precisely because of this neutrality that text meaning (a form of recording knowledge) is neutral to its truth value. In other words, a sentence may be meaningful even if it does not have a truth value; that is, for instance, if it talks about unicorns or the present kings of France or even if it states that apples fall upwards. The above is a succinct refutation of truth-conditional semantics (see XREF).

cations to ontology (see, for instance, Simons 1987, Bochman 1990, Smith 1996, 1997, 1998, Varzi 1994, 1996, 1998[59]). Other scholars develop formal devices within mereo(topo)logy to accommodate such particular elements of ontology as space and time (Muller 1998), a particular kind of artifacts (Reicher 1998), deontic phenomena (Johannesson and Wohed 1998) or inheritance models (Schäfer 1998), among others. Complex ontological entities, such as patterns (Johansson 1998, Johannesson and Wohed 1998), stand out in this respect (cf. our own complex events in Carlson and Nirenburg 1990; Section 7.1.5 below). To all of this, one must add the study of inheritance (see, for instance, Horty *et al*. 1990).

More substantively, Guarino (1995: 5-6) sees "formal ontology... as the theory of a priori distinctions: among the entities of the world (physical objects, events, regions, quantities of matter...); among the meta-level categories used to model the world (concepts, properties, qualities, states, roles, parts...)" The two distinctions have—or should have—a different, hierarchical status in any formal theory: the higher-level distinctions are metaphysical, and the lower-level distinctions should be formulated in terms of those metaphysical distinctions or, at the very least, should be strongly determined by them.

The semantic aspect of formal ontology is, of course, a serious problem, as it is with any formal theory. In the matter of ontological definition, according to Guarino (1997: 298), "[t]he ultimate definition should make clear that it includes a structure, not just the taxonomy, that all the relations are given in terms of their meaning, and that there is a logical language that corresponds to the ontology, so that 'an ontology is an explicit, partial account of the intended models of a logical language.'"

In this regard, Guarino finds Gruber's (1993: 199) much quoted view of an ontology as an explicit specification of a conceptualization extensional and shallow because a conceptualization can be—and has been—easily confused with a state of affairs. Guarino wants to add intension to it, i.e., to assign meaning to the relationship (see Guarino 1997, and 1998a: 5). He is not entirely clear about the practical steps for doing that but we believe he is on target with that desire. Nor is it entirely clear whether, for him, intensions are supposed to capture uninstantiated events while he sees extensions as instantiations: this is necessary to do, but the intension-extension dichotomy may be rather a confusing tool for the distinction. We believe that conceptualization and instantiation belong to different stages in ontological work: the former takes place during ontology acquisition while the latter is associated with ontology use. A recurring state of affairs deserves to be conceptualized, and an appropriate concept should be added to the ontology. In the process of using an ontology, it is usually the instance of a concept that is created and manipulated. Guarino may also be mistaken about the extensionality of Gruber's definition[60]: it seems that the distinction was immaterial for Gruber's own, largely engineering, design-oriented focus on ontology (see below).

59. Schäfer 1998 even gushes about mereology "provid[ing] two of the elements which were distinguishing for Chomsky's methodological [?!] reform of linguistics: The possibility of rigorous formalisation and a cognitive interpretation of the results" (108). His naiveté in choosing the role model for formal ontology aside (cf. Ch. 1, fn. 17), the latter "element" is difficult to see as accomplished in Chomskian linguistics, and it remains a problem and a bone of some contention, for instance, between Gruber and Guarino, as we will see shortly.

Guarino's insistence on complete semantic interpretability of formal ontological statements is most laudable. In fact, he addresses a very sensitive point when he sympathetically quotes a remark by Woods (1975: 40-41) that

> "philosophers have generally stopped short of trying to actually specify the truth conditions of the basic atomic propositions, dealing mainly with the specification of the meaning of complex expressions in terms of the meanings of elementary ones. Researchers in artificial intelligence are faced with the need to specify the semantics of elementary propositions as well as complex ones."

Formal ontological statements in ontological semantics are, of course, TMR propositions (see Chapter 6), and they are fully semantic in nature.

5.2.2 Ontology as Engineering

While ontology has a crucially important philosophical aspect discussed above, Guarino (1998a: 4) is essentially correct in observing that

> "...in its most prevalent use in AI, an ontology refers to an *engineering artifact*, constituted by a specific *vocabulary* used to describe a certain reality, plus a set of explicit assumptions regarding the *intended meaning* of the vocabulary words[61]. This set of assumptions has usually the form of a first-order logical theory, where vocabulary words appear as unary or binary predicate names, respectively called concepts and relations. In the simplest case, an ontology describes a hierarchy of concepts related by subsumption relationships; in more sophisticated cases, suitable axioms are added in order to express other relationships between concepts and to constrain their intended interpretation."

Gruber (1995: 909), unnecessarily, takes the same idea away from philosophy and metaphysics, while coming up with useful engineering criteria for ontology design:

> "Formal ontologies are *designed*. When we choose how to represent something in an

60. Gruber himself borrows the notion of conceptualization from Genesereth and Nilsson (1987): A body of formally represented knowledge is based on a conceptualization: the objects, concepts, and other entities that are presumed to exist in some area of interest and the relationships that hold them. This is not so distinct from Guarino's own bases: Quine's view (1953; cf. Guarino 1997: 296) that a logical theory is committed to the entities it quantified over, Newell's definition of knowledge as "whatever can be ascribed to an agent, such that its behavior can be computed according to the principle of rationality" (Newell 1982; cf. Guarino 1995: 1-2), or Wielinga and Schreiber's (1993) similar statement that "[an AI] ontology is a theory of what entities can exist in the mind of a knowledgeable agent." But, Guarino (1998a: 5) claims,

    > "[t]he problem with Genesereth and Nilsson's notion of conceptualization is that it refers to ordinary mathematical relations on [a domain] D. i.e., *extensional* relations. These relations reflect a particular state of affairs: for instance, in the blocks world, they may reflect a particular arrangement of blocks on the table. We need instead to focus on the *meaning* of these relations, independently of a state of affairs: for instance, the meaning of the 'above' relation lies in the *way* it refers to certain couples of blocks according to their spatial arrangement. We need therefore to speak of *intensional* relations..."

61. It should be emphasized right away that these "words" are not words of a natural language, something about which Guarino himself and some other scholars are not always careful enough—see Section 5.3 for further discussion.

> ontology, we are making design decisions. To guide and evaluate our designs, we need objective criteria that are founded on the purpose of the resulting artifact, rather than based on *a priori* notions of naturalness or Truth. Here we propose a preliminary set of design criteria for ontologies whose purpose is knowledge sharing and interoperation among programs based on a shared conceptualization.
>
> 1. **Clarity**. An ontology should effectively communicate the intended meaning of defined terms. Definitions should be *objective*.... Wherever possible, a complete definition (a predicate defined by necessary and sufficient conditions) is preferred over a partial definition (defined by only necessary or sufficient conditions).
>
> 2. **Coherence**. An ontology should be coherent: that is, it should sanction inferences that are consistent with the definitions.... If a sentence that can be inferred from the axioms contradicts a definition or example given informally, then the ontology is incoherent.
>
> 3. **Extendibility** [sic]...[O]ne should be able to define new terms for special uses based on the existing vocabulary, in a way that does not require the revision of the existing definitions.
>
> 4. **Minimal encoding basis**.... Encoding bias should be minimized [to allow for various encoding options.]
>
> 5. **Minimal ontological commitment**.... An ontology should make as few claims as possible about the world being modeled, allowing the parties committed to the ontology freedom to specialize and instantiate the ontology as needed."

The degree of commitment to an ontology in an information system may vary from zero to vague awareness to ontology-drivenness. "In some cases," Guarino writes (1998a: 3), "the term 'ontology' is just a fancy name denoting the result of familiar activities like conceptual analysis and domain modeling, carried out by means of standard methodologies." Ontology really comes into its own when its "own methodological and architectural peculiarities" *(ibid.)* come into play. In this case, the ontology becomes an integral component of the information system, "cooperating at run time towards the 'higher' overall goal" (*op. cit.*: 11). While definitely "ontology-driven," ontological semantics can, we believe, claim an even higher status: it is actually ontology-based, or ontology-centered.

5.2.3 Ontology Interchange

The last important issue of formal ontology we will touch upon briefly here is the movement to share and reuse ontologies. In fact, this is what Gruber's Criterion 5 above includes. Our own ontology has been shared by us with other groups and multiply reused both at CRL and elsewhere.

There are, however, two non-trivial issues with ontology interchange. One of them is the dichotomy, well-known in descriptive and computational linguistics, between a specific domain and a multidomain situation. When designing an ontology for a domain (or describing a sublanguage of natural language—see Section 9.3.6; cf. Raskin 1990), one can take full advantage of its limited nature and achieve a higher accuracy of description. This comes, however, at a price: the more domain-specific the description the less portable it is outside of the domain.

Furthermore, some scholars claim that no domain knowledge is or can be independent from a particular task for which it is developed and a particular method employed—this approach is, of course, well-known in physics. According to Bylander and Chandrasekaran (1988), ontology design cannot be free of “the so-called interaction problem: ‘Representing knowledge for the purpose of solving some problem is strongly affected by the nature of the problem and the inference strategy to be applied to the problem.’” Ignoring the long-standing debate on the issue in post-Bohr physics, Guarino (1997: 293) mounts his own defense, coupled with a reasonable plea:

> “I will defend here the thesis of the independence of domain knowledge. This thesis should not be intended in a rigid sense, since it is clear that—more or less—ontological commitments always reflect particular points of view (for instance, the same physical phenomenon may be described in different ways by an engineer, by a physicist or by a chemist); rather, what I would like to stress is the fact that reusability across multiple tasks or methods *can and should be systematically pursued*[62]...”

The ontological community has devoted a considerable effort to pursue this goal systematically. In a widely shared opinion, Gruber (1993: 200; cf. Nirenburg *et al*. 1995) stated correctly that

> “[k]nowledge-based systems and services are expensive to build, test, and maintain. A software engineering methodology based on formal specifications of shared resources, reusable components, and standard services is needed. We believe that specifications of shared vocabulary can play an important role in such a methodology.”

The second non-trivial issue in ontology interchange is developing formal tools for making the importation and interchange of ontologies possible. Gruber (1993) proceeded to define such a formal tool, Ontolingua, the best-known system for translating ontologies among notations. Ontolingua uses KIF, the Knowledge Interchange Format, designed by Genesereth and Fikes (1992):

> “KIF is intended as a language for the publication and communication of knowledge. It is intended to make the *epistemological-level* (McCarthy & Hayes, 1969) content clear to the reader, but not to support automatic reasoning in that form. It is very expressive, designed to accommodate the state of the art in knowledge representation. But it is not an implemented representation system” (Gruber 1993: 205).

All of this is quite appropriate for translating ontologies because the same can be said of ontologies themselves. Designing ontologies for portability means that

> “[e]xplicit specifications of... ontologies are essential for the development and use of

62. Outside of formal ontology proper, in the field of knowledge representation, Doyle and Patil (1991: 289) cast a vote of confidence in the possibility of “general purpose representation systems.” They “argue that general purpose knowledge representation systems should provide:
    - •Fully expressive languages,
    - •Tolerance of incomplete classification,
    - •Terminological classification over relevant contingent as well as definitional information,
    - •Nondeductive as well as deductive forms of recognition which permit “approximate” classification and “classification” of concepts involving defaults, and
    - •Rational management of inference tools” (op.cit.: 266).

    These additional properties seem to complement nicely the more external constraints by Gruber in Section 5.2.2 above.

intelligent systems as well as for the interoperation of heterogeneous systems.

> Ontology construction is difficult and time consuming. This large development cost is a major barrier to the building of large scale intelligent systems and to widespread knowledge-level interactions of computer-based agents. Since many conceptualizations are intended [or can be found] to be useful for a wide variety of tasks, an important means of removing this barrier is to encode ontologies in a reusable form so that large portions of an ontology for a given application can be assembled from existing ontologies in ontology repositories" (Farquhar et al. 1995: 1; cf. Farquhar et al. 1996, 1997).

To this effect, the Fikes group has actually implemented a website for ontology importation and integration (http://WWW-KSL-SVC.stanford.edu:5915/&service=frame-editor). Our own experience in augmenting and modifying existing ontologies shows that, while often simpler than acquiring ontologies from scratch, it is still a labor-intensive effort that can be facilitated in many ways, of which the specification of a reusable format may not even be the most important one. The development of dedicated semi-automatic ontology acquisition methodologies and tools is, in our estimation, much more useful, specifically because it concentrates on content, not format.

5.3 Ontology and Natural Language

In the preceding two sections, we discussed well-explored areas of philosophical and formal ontology, primarily as they pertain to ontological semantics. In this section, we are venturing into the difficult and underexplored part of formal ontology, namely, the relations between ontology and natural language.

5.3.1 A Quick and Dirty Distinction Between Ontology and Natural Language

Guarino (1998b) criticizes several existing ontologies, including the Mikrokosmos implementation of ontological semantics, for allowing ambiguity in certain ontological nodes, e.g., treating the node WINDOW as both an artifact and a place, thus effectively postulating what, from his point of view, is a non-existing concept that subsumes the properties of both. Similarly, he objects to a link from the COMMUNICATION-EVENT node to both SOCIAL-EVENT and MENTAL-EVENT as parents.

This criticism can be appropriate only if ambiguity in ontological concepts is not allowed in any form. In that case, the distinction between natural language and ontology is simple and clear-cut: words can be ambiguous; concepts, cannot.[63] To accommodate the absence-of-ambiguity principle, an ontology should have different nodes in different places for the concepts WINDOW-ARTIFACT and WINDOW-PLACE, or for MENTAL-COMMUNICATION-EVENT and SOCIAL-COMMUNICATION-EVENT, even if it chooses to use the same English word in their labels.

If ambiguity so clearly demarcates words from concepts, then it is rather surprising that Guarino considers linguistics a participant in the development of an ontology: "[o]n the methodological

63. Whether lexical ambiguity is accidental or systematic, has nothing to do with ontological concepts per se. For example, in the AQUILEX and CORELEX projects much effort has been expended on treating as many instances of ambiguity as possible in a systematic way. Then, Buitelaar (1998), a CORELEX contribution, does not seem to belong in the volume on ontology edited by Guarino.

side, the main peculiarity is the adoption of a *highly interdisciplinary approach*, where philosophy and linguistics play a fundamental role in analyzing the structure of a given reality at a high level of generality and in formulating a clear and vigorous vocabulary" (1998a: 3). In what sense, does—or can—linguistics contribute to this enterprise? Bateman (1993) provides a reasonably clear explanation: "ontology construction [should be based] on an understanding of natural language." Hovy and Nirenburg (1992) is more cautious and circumspect: the knowledge we obtain from our understanding of a particular natural language should be integrated with and into a language-neutral ontology, presumably by combining the material from different languages. It should be stressed, however, that both of the above opinions relate more to the ability of people to perceive and manipulate knowledge through language than to the formal discipline of linguistics and its legitimate purview.

Moreover, the use of the term 'vocabulary' in the initial quote above licenses mixing ontological nomenclature with units of the dictionary of a natural language, thus further contributing to the unjustified fusion of the metalanguage of ontological description with natural language. As we have argued elsewhere (see Nirenburg *et al*. 1995, Nirenburg and Raskin 1996: 18-20, and Sections 2.6.2.2 and 4.3.2), some scholars persist in this natural-language fallacy positively, as it were, by insisting on using natural-language words instead of ontological concepts to represent natural-language meanings, and others persist in it negatively by trying to expose an ontology as camouflaged natural language. The former has a long history, if not a legitimate place in natural language semantics; the latter is rather easily refuted by indicating that ontological concepts have no ambiguity.

The confusion has a deep philosophical origin, going back at least 50 years, namely, the so-called 'linguistic turn' in philosophy (cf. Footnote 5 in Chapter 2): the move away from world phenomena or even their representations in human concepts[64] to the analysis of the meaning of propositions about these phenomena or concepts. Kenny (1989: viii) gives this move a pretty fair shake:

> "In the last half-century many people have described themselves as adherents of, and many people have described themselves as enemies of, linguistic philosophy. Neither adherence nor opposition is a very useful stance unless one makes clear what one means by calling a particular style of philosophy 'linguistic.'
>
> 'Philosophy is linguistic' may mean at least six different things. (1) The study of language is a useful philosophical tool. (2) It is the only philosophical tool. (3) Language is the only subject-matter of philosophy. (4) Necessary things are established by linguistic convention. (5) Man is fundamentally a language-using animal. (6) Everyday language has a status of privilege over technical and formal systems. These six propositions are independent of each other. (1) has been accepted in practice by every philosopher since Plato. Concerning the other five, philosophers have been and are divided, including philosophy within the analytic tradition. In my opinion, (1) and (5) are true, and the other four false."

In our opinion, (5) does not contribute to the issue at hand, and the rest is false. Language is a tool of philosophy to not a larger extent than it is for any other human endeavor. Studying language

64. The "slippage" from studying categories of the world in metaphysics to studying human concepts reflecting (or even replacing) these "objective" categories is usually attributed to Kant (1787—cf. Loux 1998: 2)

takes away from philosophy like studying the screwdriver takes away from driving in a screw. We are actually with Chisholm (1996: 8), when he writes:

> "Aristotle says that in discussing the categories, he is concerned in part with our ordinary language. And he says this often enough to provide encouragement to those contemporary philosophers who believe that the statements of metaphysicians, to the extent that they are not completely empty, tell us something about our language. One of our principal concerns, however, is that of finding the *ontological presuppositions* of statements about language.
>
> Where some readers of this book may expect to find discussions of language, they will find discussions of thinking and intentionality instead."

We would like to take it a little further still by claiming that Chisholm's ontological presuppositions are ontological content, ontological meaning, and it is separate from natural-language meaning.

5.3.2 The Real Distinction Between Ontology and Natural Language

In this section, we question the premise that ambiguity is what distinguishes natural language and ontology. In particular, we will explore whether an ontology really must be unambiguous and whether this ideal is at all attainable. Next, we will argue that the real distinction is that languages emerge and are used by people, while ontologies are constructed for computers.

Is the objection of formal ontology to having a single concept WINDOW with the properties of both opening and artifact justified? The objection is predicated on the premise that there must be no ambiguity in ontology. What does this premise mean in reality? In a formal logical system, no two concepts can have the same name, and, conversely, no single concept can be referred to by more than one name. Obviously, no such blatant violation of formality can be expected in any practical ontology. So, what was, then, criticized by Guarino in the Mikrokosmos decision to use a single concept for WINDOW? It was precisely the decision to declare WINDOW a single concept carrying no ambiguity.

Why would one prefer to split WINDOW into two different concepts? A claim that the English word *window* has two distinct senses would have no bearing on this ontological decision because it should not be expected that there will be a one-to-one relationship between the space of word senses in a natural language (or, more accurately, the union of all word senses in all natural languages) and that of ontological concepts. On the contrary, it seems more important that there is apparently no natural language in the world in which the word for WINDOW does not realize both the opening and the artifact senses of the word. This semantic universal is probably the strongest evidence we may have that people seem to conflate the two concepts.

This phenomenon (known variously as regular polysemy, vagueness, underspecification, sense permeability, etc.) is pervasive: *book* (or *newspaper* or even *poem*) in all languages refers both to a physical object and its informational content; *say* (or *smile* or *wink*) to a physical phenomenon and to conveying a message; *bank* (or *school* or *shop*) to an organization and a building housing it, etc. Will formal ontology require that each of the concepts corresponding to these word senses be duplicated in the manner suggested?

In the same vein, should there be different concepts for *eat* corresponding to eating with a spoon, with a fork, with chopsticks or one's fingers? After all, each of the above are distinct processes.[65] Would it matter if in some language there were, in fact, different words for some of these processes? In general, languages do not use isomorphic sets of lexeme names. This has given rise to the widespread study of cross-language mismatches and translation divergences (see, for instance, Viegas et al. 1999:190-195), as in the well-known example of English *wall* vs. Italian *muro* 'outside wall' and *parete* 'inside wall.'

If an ontology is constructed according to Bateman, that is, based on an understanding of a natural language, and that language happens to be Italian, then the ontology will have two separate concepts for the inside and outside wall. Using such an ontology in NLP and defining lexical senses in ontological terms as it is done, for instance, in the Mikrokosmos implementation of ontological semantics, the two Italian words will be directly mapped into these concepts. Using the same ontology to support the acquisition of the English lexicon, the entry for *wall* will have a connection to two ontological concepts, and this is the definition of polysemy in ontological semantics. In other words, *wall* would have one sense more than if the ontology contained just one concept for WALL. It might be counterintuitive for an English speaker to consider that *wall* has two senses corresponding to the inside and the outside walls.

If, on the other hand, ontology construction is based on English, there will be a single concept for *wall* in it. From the point of view of the Italian speaker, this concept would be seen as a non-terminal node in the ontological hierarchy which would have the concepts for *muro* and *parete* as its children. This concept could be made a terminal node, thus becoming a conflation of the putative child concepts, very similarly to what Guarino sees happening in the case of using a single concept for WINDOW which conflates the notions of opening and artifact. There must be a reason to include or not to include the children in the ontology. We have just seen that decisions based on Italian and English are incompatible; which means that Bateman's principle of "basing ontology construction on an understanding of natural language" is not feasible.

We can think of a practical, descriptive reason for not splitting the concept of WALL. Let us assume that Italian is the only language with makes the above distinction. An ontology with two separate concepts will add a sense to the entries of the word for *wall* in all the other languages, which will result in many extra senses in the universal lexicon. If, on the other hand, the concept is not split, the only price to pay is to add a disambiguation constraint to the lexical entries for *muro* and *parete*.

A criterion for deciding when to stop splitting ontological concepts and, more generally, how to demarcate them, deserves to be one of the cornerstones of formal ontology. It seems to have to do with Gruber's (1995: 909—see also Section 5.2.2 above) general ontology-design criterion of 'minimal ontological commitment,' and it needs further elaboration, along with other topics in formal ontology (see Section 5.4 below). In linguistic semantics, it is hard to establish a similar principle for limiting the polysemy of a lexical item, and typically, monolingual dictionaries prefer to multiply the senses, the number of which can be radically reduced without loss both for the-

65. This example was used by Weinreich (1966) in his critique of Katz and Fodor (1963). Weinreich's point was that the semantic theory proposed by Katz and Fodor established no limits for polysemy.

oretical and practical purposes (see Nirenburg *et al*. 1995, Raskin and Nirenburg 1998: 192-199; Section 9.3.5 below). In multilingual dictionaries, an important motivation for distinguishing word senses in one language is the presence in another language of different words for realizing these senses. In an early contribution to universal lexical semantics, Hjelmslev (1958) proposed to format multilingual lexical semantic descriptions as follows:

<table>
<tr><th>French</th><th>German</th><th>Danish</th></tr>
<tr><td>arbre ‘tree’</td><td>Baum</td><td rowspan="2">troe</td></tr>
<tr><td rowspan="3">bois
1) ‘wood’ material
2) ‘wood’ part of landscape</td><td rowspan="2">Holz</td></tr>
<tr><td rowspan="3">skov</td></tr>
<tr><td rowspan="2">Wald</td></tr>
<tr><td>forêt ‘forest</td></tr>
</table>

In the above table, the French column features two sets of synonyms—*arbre/*$bois_1$ and $bois_2$*/forêt* and one polysemous word, *bois*. The Danish column features two polysemous words, each with two senses and one set of synonyms, $troe_2$/$skov_1$. The German column features three single-sense words (at least, for the senses illustrated in the table).

What Hjelmslev implied here is a method of crosslinguistic lexical semantic analysis. Dolgopol’sky (1962) implemented this method on the material of 28 languages. This method is based on a geometric metaphor: should one choose to extend all horizontal lines across the entire table, the resulting rows will correspond to what Hjelmslev called ‘values,’ i.e., relative, differential meanings. One would think that, in ontological terms, these values would correspond to the most detailed, atomistic conceptual representations, excluding any possibility of concept ambiguity or conflation. In reality, some of these extensions will be hard to interpret: thus, the extension to the left of the line between *troe* and *skov* would split the concept of wood as material into two unmotivated concepts.

As any dictionary of French, German or Danish will demonstrate, the words in the table realize three distinct senses—those of plant, material and landscape feature. It is an accident that these three senses correspond to the German words. One should expect that, in other cases of crosslinguistic description, it will be another language that will turn out to be nonpolysemous.[66] And yet in other cases, no language may be found to provide nonpolysemous coverage of the senses involved. The most appropriate ontology for representing the senses of the words in the table should contain three concepts corresponding to the German word senses. As will be discussed in the next section, we see it as a task for formal ontology to explore why it is so and what criteria can be discovered for making such decisions. We have a very strong intuition that these three concepts must be represented in the ontology to the exclusion of all alternatives. Such a decision strikes us as "natural" and obvious. As we have mentioned elsewhere (Nirenburg et al. 1995), this feeling is shared by all the members of our research team, which makes such decisions reproducible. It is not entirely clear to us what this certainty is based on, and we believe that it is the task of formal and philosophical ontology to address this issue.

Ambiguity is pervasive in language and must be handled. We have seen that it cannot be expunged from any specific implementation of ontology, because of no purely conceptual limits on grain size. Moreover, an effort to split ontological concepts into the ever smaller unambiguous units leads to a sharp increase in polysemy and, therefore, makes the task of disambiguation so much more difficult. As we will argue in the next section, no ontology exists in a vacuum. It interacts with other resources, such as knowledge representation languages, lexicons, analyzers, generators, general reasoning modules, etc. It is safe to assume that the overall amount of ambiguity to be addressed in any application is fairly constant at a given grain size. The differences among the various approaches to the treatment of ambiguity may be articulated as differences in the distribution across the above resources of the knowledge necessary for resolving the ambiguity. So, if an ontology is made less ambiguous, it only means that the ambiguity will have to be treated increasingly elsewhere.

The confusion about ambiguity creeps into the important issue of the distinction between ontology and natural language in yet another way. Wilks (Wilks et al. 1996: 59, Nirenburg and Wilks 1997—see also Section 2.6.2.2 and references there) has eloquently argued that the very fact of using English words as labels for ontological concepts smuggles natural language ambiguity into ontology, and thus there is no basic difference between the representation language of ontology and a natural language:

> "YW: The first feature of language that should concern us in this discussion is as follows: can predicates of a representational language avoid ending up ambiguous as to sense? The negative answer to this question would make RLs NL-like. It will also mean that understanding a representation involves knowing what sense a symbol is being used in. If NLs are necessarily extensible as to sense—and words get new senses all the time—then can RLs that use NL symbols avoid this fate?" (Nirenburg and Wilks 1997: 4—from Wilks' contribution to the dialog-format article).

What Wilks seems to ignore here is that the meanings of the ontological labels are constructed—in the sense of being formally defined for the use by the computer. The computer perceives these labels straightforwardly; no comparison or reference is made to any words of any natural language with which these labels can be homographous. Wilks is right that a human reading such a

66. Throughout this discussion we have followed Hjelmslev and others in assuming that having a single word in a language to express a certain sense is somehow a privileged state of affairs compared to expressing this sense with a phrase. We have argued elsewhere (Raskin and Nirenburg 1995, 1998) for the 'principle of practical effability' among languages which removes any distinction between these two alternative ways of expressing a meaning. It was convincingly argued (Zvegintzev 1960) that this priority of single words over phrases is a fallacy and that this fallacy is the cornerstone of the Sapir-Whorf hypothesis. We are not sure that any neurolinguistic or psychological evidence exists for the primacy of the single word expression in the mind. What we do know is that a computational dictionary can include both words and phrases as entry heads.

label is likely to slip into an alternative sense of the homographous word not intended by the creator of the label. However, the computer has no capability of doing so.

It appears then that the crucial distinction between ontology and natural language does not lie exactly in the nonambiguity of the one and ambiguity of the other. This distinction is in the constructed and overtly defined nature of ontological concepts and labels on which no human background knowledge can operate unintentionally to introduce any ambiguity, as opposed to pervasive uncontrolled ambiguity in natural language. The entire enterprise of natural language processing is about designing knowledge structures for the computer to use. We have not yet achieved that goal, so we cannot suspect that a computer would be able to confuse an ontological concept or its label with its homographous, and possibly polysemous, lexeme of a natural language.

5.4 A Wish List for Formal Ontology from Ontological Semantics

Practical ontology building expects assistance from formal and philosophical ontology. In this section, we compile a wish list of issues that practical ontology builders would want to be tackled and solved in a principled way. The issues relate to

- the status of ontology vis-à-vis other knowledge resources in an application;
- the choice of what concepts to acquire;
- the choice of what content to assign to each concept; and
- the evaluation of the quality of an ontology using both the glass-box and black-box evaluation paradigms.

In practical applications, ontologies seldom, if ever, are used as the only knowledge resources. In the representative application of knowledge-based MT, for example, the ontology is used

- to supply the language for explaining lexical meanings, which are recorded in the lexicons of particular languages;
- to provide the contentful building blocks of a text meaning representation language;
- to provide the heuristic knowledge for the dynamic knowledge resources such as semantic analyzers and generators.

Formal ontology must help ontology builders to constrain the relationships between ontological concepts, structures that represent text meaning and lexicon entries. In particular, an account must be available of the difference between ontological concepts as entity types and text meaning elements (and the semantic components of lexicon entries) as entity instances. What we believe it means for formal ontology is the necessity to define the status, beyond knowledge representation format, of ontological instances. The latter come in several kinds, the most important for our discussion here being: instances of ontological concepts used for defining lexical meaning in lexicon entries; and facts, representations that result from compositional combination of meanings of individual words and/or phrases in the input into the meaning specification of a text.

A crucial concern for an ontology builder is the decision on what concepts to introduce and how to represent each of them. A good ontology will have a good coverage and be reasonably homogeneous.[67] While coverage is determined by the domain and the nature of the application, formal ontology can help to decide how to organize the concepts that must be included in the ontology, for instance, how to organize the most economical concept hierarchy and how to define the non-terminal nodes in it. Formal ontology will be much more useful in practice if it agreed not only to put forward desired properties of ontologies but also to offer criteria for the processes of ontology construction and judgments about sufficient depth and breadth of coverage. In other words, we are suggesting that formal ontology, as a theory, must be supplemented by a methodology (see Sections 2.4.2 and 2.5).

As we have just mentioned, formal ontology effectively concentrates on evaluating the quality of ontologies. In fact, even in that endeavor, we would benefit from a broadening in the scope of such evaluations. At present, formal ontology is concerned with the inherent properties of ontologies, considered independently of any concrete application. This is done by examining the content of the ontologies in search of potential contradictions, ambiguities and omissions. This type of evaluation is often called glass-box evaluation, as the internal workings of a resource are transparent to the investigator. Practical ontologists would benefit from extending the purview of the evaluation into a glass-box evaluation of an ontology under construction as well as into a black-box evaluation of both existing and nascent ontologies when the ontology itself is opaque to the investigator, and the quality of the system is judged by the quality of output of an application based on the ontology (see also Figure 16 in Chapter 2). In fact, Mahesh *et al.*, (1996) is a good example of such an evaluation: it had to develop the principles and criteria on the fly. It would have been better to take them off the formal ontology shelf.

67. This is an important desideratum. In fact, a major ontological enterprise was criticized in one evaluation for a lack of strategic direction for achieving a more or less uniform depth and breadth of knowledge coverage—see Mahesh et al. 1996.