Ontotherapy, or how to stop worrying about what there is

Yorick Wilks
University of Sheffield
Department of Computer Science
Regent Court, 211 Portobello Street
Sheffield, S1 4DP, UK
y.wilks@dcs.shef.ac.uk

Abstract

The paper argues that Guarino is right that ontologies are different from thesauri and similar objects, but not in the ways he believes: they are distinguished from essentially linguistic objects like thesauri and hierarchies of conceptual relations because they unpack, ultimately, in terms of sets of objects and individuals. However this is a lonely status, and without much application outside strict scientific and engineering disciplines, and of no direct relevance to language processing (NLP). More interesting structures, of NLP relevance, that encode conceptual knowledge, cannot be subjected to the "cleaning up" techniques that Guarino advocates, because his conditions are too strict to be applicable, and because the terms used in such structures retain their language-like features of ambiguity and vagueness, and in a way that cannot be eliminated by reference to sets of objects, as it can be in ontologies in the narrow sense. Wordnet is a structure that remains useful to NLP, and has within it features of both types (ontologies and conceptual hierarchies) and its function and usefulness will remain, properly, resistant to Guarino's techniques, because those rest on a misunderstanding about concepts. The ultimate way out of such disputes can only come from automatic construction and evaluation procedures for conceptual and ontological structures from data, which is to say, corpora.

Key words: ontologies, thesauri, necessary and sufficient conditions, identity criteria, evaluation.

1. Introduction

Is there a problem about ontologies? Are they really distinct from semantic nets and graphs, thesauri, lexicons, taxonomies or are people just confused about any real or imagined differences? Does the word "ontology" have any single, clear, meaning when used by researchers in Artificial Intelligence (AI) and Natural Language Processing (NLP) and, if not, does it matter? And are we, assuming us to be researchers in AI/NLP, just muddled computer people who need therapy, that is, to have our thoughts firmed up, cleaned up, or sorted out by other, more philosophical, logical or linguistic, experts so we can do our job better? And, for those who read newspaper columns on how we are losing the meanings of key words, does it matter that what ontology means now, in our field at least, is somewhere between the notions above, and what it traditionally meant, namely the study of *what there actually is in the world*, such as classes or individuals. For that is what metaphysicians since Aristotle thought it meant.

These questions will be addressed, if not answered, in this paper; the last question will, of course, be answered in the negative, since philosophers have no monopoly over meanings, any more than

the rest of us. The central question above, the one that mentions AI, is a recurrent question in AI itself, and one to which I shall declare a practical, and negative, answer right at the start: namely, that decades of experience shows that for effective, performing, simulations of knowledge-based intelligence, it is rarely the case that enhanced representations, in the sense of those meeting strong, formal, criteria derived from logic, are of use in advancing those ends.

Since the topic is a metaphysical one, I declare my own hand at the beginning where these matters are concerned: my own PhD thesis (Wilks, 1968) was a computational study of metaphysical arguments, as contained in classic historical texts (or corpora, as we would now say), and my claim (barely supported by the exiguous computing capacity of those days) was that such argument proceeds and succeeds by methods quite different from the explicit, surface, logical argument structure proposed by its original philosophical authors, but rather by methods involving rhetorical shifts of the senses of key words and of which the author may not even be aware e.g. the whole rhetorical purpose of the philosophy of Spinoza, set out in the form of logical proofs, all of which are faulty, is actually to shift our sense for the word "Nature" (see Bosanquet, 1946).

It was this early investigation that alerted me to the possibility that representational structures are not always necessary where they were deployed, and that it is hard to be sure when representations are or are not adequately complex to express some important and complex knowledge, and that one should be very wary of the usefulness of logic-based formalisms where language is concerned. At one extreme, reference could be made here to Schvaneveldt's Pathfinder networks (1985) which were simple associative networks derived from word use that seem able, contrary to most intuition, to express the kinds of skills fighter pilots have. I recall also the dispute, originated by Fodor, as to whether or not connectionist networks could express recursive grammatical structures, an argument I believe he lost when Pollack (1990) produced his Recursive Auto-Associative Networks that could do exactly that. I myself once wrote a paper on machine translation (MT) that was later widely cited, to my surprise, where I argued (1994) that most MT systems do not in fact work with the formalisms their architects use to describe them.

The core issues here seem to me to be: first, whether any particular formalism can encode the requisite types of knowledge for some purpose and do so better than some simpler, more tractable, representation and, secondly, whether that representation supports whatever inference procedures are deemed necessary for that knowledge. If an ontology is taken in something like its classic sense, as a hierarchical structure of sets and individuals, then we know it will support simple Boolean/quantificational inference and we also know it will have difficulty representing phenomena that do not fall easily under set theory, namely intensional and continuous phenomena. But we also know that these can be accommodated by a range of extensions that make an ontology more like a general AI knowledge representation schema. In fact, it seems obvious to me that that is what has happened, at least in the more formal DAML/OIL parts of the ontology movement, and that ontology there is to a large degree AI renamed, with all the same representational difficulties and nothing much gained, in theoretical terms, beyond new and exciting application environments, except inn terms of the computational tractability of the formalism, a point we shall return to below.

But again, I would stress that appearances are not realities in this area, and much-used ontologies in the public domain do not have the structure they appear to have and would need if the inference constraint were to be taken seriously.

Another theme I wish to introduce is that facts like this last plainly do not make structures useless, because these man-made objects (viz. ontologies, lexicons, thesauri) for classification of words and worlds contain more than they appear to, or their authors are aware of, which is why computational work continues to extract novelty from analyzing such objects as Websters 7th, LDOCE, Wordnet or Roget's Thesaurus. Margaret Masterman (2005) memorably claimed that the structure of Roget showed unconscious, as well as explicit structure, and it was the mental organization of a 19th century Anglican clergyman: above all an opposition between good and evil! If any of this is the case, then what structural objects that contain knowledge need is not so much conceptual clearing up but investigation of what they actually contain, and Wordnet has been subjected to a great deal of such analysis, e.g. (Peters, 2001)

As I noted already, those cursed with a memory of metaphysics are often irritated by modern AI/NLP where the word "ontology" is hardly ever used to mean what it used to, namely "the study of what there is, of being in general". Recent exceptions to this are discussions by Hobbs (1983) and others, but almost all modern usage refers to hierarchical structures of knowledge, whose authors never discuss what there is, but assume they know all that and just want to write down the relations between the parts/wholes and sets and individuals that undoubtedly exist.

To a large extent, and to avoid pointless controversy, this paper will go along with the usage while noting in passing that as a web search term, "ontology" locates two quite disjoint literatures with virtually no personnel in common: the world of formal ontology specification (e.g. Horrocks et al. 2002) and the world of ontologies for AI tasks to do with language (e.g. Nirenburg in Mahesh et al., 2000). Rare overlaps would be Lenat's CyC system (1995) which began life as an attempt to write down a great deal of world knowledge in predicate form but which was also claimed by Lenat as a possible knowledge form for use in language processing.

Before proceeding to the argument of the paper, it may be worthwhile to set out the basic terms being used, and in particular ontology, dictionary and thesaurus. The following are, in a sense, pure vanilla versions of these:

(1) is a bit of ontology for the concept denoted by the English word "pan":

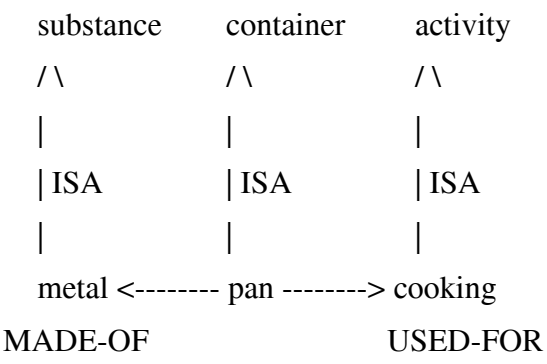

We shall be asking how this is this different from a dictionary entry for the term:

2) Pan (N): container, typically made of metal, used for cooking.

Or for a thesaurus entry for the same word, not in the form of Roget's thesaurus but something more up to date that marks SYN(ONYMS) and HYPER- and HYPONYMS explicitly:

3) pan: SYN pot,
HYPERNYM container, SYN holder
HYPONYM frying pan, skillet
steamer

WE shall say that, broadly, (1) is about things in the world, and (2) and (3) are about language, being respectively examples of a conventional dictionary entry and a thesaurus-style lexicon, and the difference will lie in the fact that all the terms in (1), except "pan" itself, are terms in a special defining language.

We could equally well consider:

	GENE	ALIAS	FUNCTION	ENZYME
4)	Pb33X	PBXalpha	activation	Transcriptaze
	Pb33-Xa			

The surface notation in (1) and (4) is different, but this does not matter, since we could easily redraw either in the style of the other. (4) is also about things in the world, genes, not words (other than with the rather special role of names and alias names), and this, in spite of its technical appearance is no more problematic than (1), or is it? We shall come back to genes at the end of the paper and ask whether we can be as sure of genes as we can about pans.

2. Ontologies and conceptual relations

The place to begin is one of the questions above about the conflation of ontologies (construed in the modern manner as hierarchical classification of things or entities) versus thesauri, or taxonomies, as hierarchical classification of words or lexical senses. There is a widespread belief that these are different constructs, as different (on another dimension) as encyclopaedias and dictionaries, and should be shown to be so. Others will admit that they are often mixed in together, as when Wordnet (Miller, 1995) is called an ontology, as parts of it undoubtedly are, on the above definition, but that this may not matter as regards its function as the most popular resource for NLP, any more than it matters that ordinary dictionaries contain many facts about the world like "a chrysanthemum is a flower that grows at Alpine elevations".

A random paper it would be unfair to cite, but which is typical of much recent work, recently offered an ontological coding scheme, made up of what it called Universal Words, and whose first example expression was:

(Drink > liquor)

which was intended to signify, via "universal words" that *"drink is a type of liquor"*. This seems at first sight to be the reverse of common sense: where liquors (distilled alcoholic drinks) are a type of drink. Again, if the text as written contains a misprint or misunderstanding of English and "liquid" was intended instead of "liquor" then the natural interpretation of the expression is true. There is probably no straightforward interpretation of the italicised words that can make

them true, but the situation is certainly made complex by the fact that "drink" has at least two relevant senses (potable liquid vs. alcoholic drink) and liquor two as well (distillate vs. potable distillate).

I bring up this, otherwise forgettable, example because issues very like this constantly arise in the interpretation of systems that claim to be ontologies, as opposed to systems using lexical concepts or items, and thus claiming to be using symbols that are not words in a language (usually English) but rather idealised, or arbitrary items, which only accidentally look like English words.

This situation, of would-be universal words, taken from English and whose unique sense has not been declared, is sometimes obscured by the problem of "foreigner English". By that I mean that, since English is now so widely known as a non-maternal language, it is often the case that the alternative senses of English words are simply not known to those using them as primitive terms. This can be observed in the interlinguas in use in some Japanese MT systems. This benign ignorance may aid the functioning of a system of "universal words" but it is obviously an unsatisfactory situation.

I also believe it is not sufficient to say, as someone like Nirenburg consistently does (see Nirenburg & Wilks, 2001), that ontological items simply seem like English words. My own view is firmly that items in ontologies and taxonomies are and remain words in natural languages -- the very ones they seem to be, in fact -- and that this fact places strong constraints on the degree of formalisation that can ever be achieved by the use of such structures. The word "drink" has many meanings (e.g. the sea) and attempts to restrict it, within structures, and by rules, constraints or the domain used, can only have limited success. Moreover, there is no way out here via non-linguistic symbols or numbers, for the reasons explored long ago in (McDermott 1976). Those who continue to maintain that "universal words" are not the English words they look most like, must at least tell us which of the senses of the real word closest to the "universal word" they intend it to bear under formalization.

One traditional move at this point, faced with the demand above, is to say that science does not require that kind of precision, imposed at all levels of a structure, but rather that "higher level" abstract terms in a theory gain their meaning from the theory as a whole. This is very much the view of the meaning of terms like "positron" adopted by Katz (1972) and, from a quite different position in the philosophy of science by writers like Braithwaite (1953), who argued that scientific terms (like "position") at the most abstract level of a scientific theory should be interpreted by a process of what he called "semantic ascent" from the interpretations of lower, more empirical, terms.

This argument is ingenious, but suffers from the defect that a hierarchical ontology or lexicon is not like a scientific theory (although both have the same top-bottom, abstract-concrete correlation) because the latter is not a classification of the world but a sequential proof from axiomatic forms.

However, what this "positron" analogy expresses, is very much in the spirit of Quine's later views (1951), namely that not all levels of a theory measure up to the world in the same way, even though there is no firm distinction between high and low levels. This in consistent with his well-known, but often misunderstood, view that "to be is to be the value of a bound variable", a notion we shall return to later, when we contrast this class of views with those of a writer like Guarino, who very much wants to argue that *all* levels of an ontology have their own "Identity

Conditions", which determine the other levels to which any given level can be linked by a relationship like ISA, normally taken to be set inclusion or type specification.

The most pragmatic possible response to this formalist criticism (like Guarino's which we shall discuss in detail below) is to argue that hierarchies, or any other structures, are only justified by their retrieval properties for some desired purpose, and that such evaluation overrides all other considerations. This is a tempting situation for anyone of a practical bent. However, as a position it has something of the defects of connectionism, to which Charniak famously responded that he would not work with a representation that could not be understood.

Thus even though we, in language engineering, are not doing science in any strong sense, we must, of course, aim for perspicuous, defensible representations, and never fall back on "it works, shut up about it" as a form of explanation. To the same end, I take the result of many years and rounds of the McDermott discussion (see above) to be that we cannot just say that representations are arbitrary, inessential, and could be unexamined English or even binary numbers. Given all this, one clear way of presenting the issue central to this paper is to ask the question: are thesauri and taxonomies really different in type from ontologies? Does one describe words and the other describe worlds?

The formal answer is that they are in principle different and should be seen to be so. This is very much the position of Guarino (Gangemi et al., 2001) which we will come to in a moment. However, it is hard to believe they are utterly different in practice, since the principle, whatever it is, is hard to state in clear terms. Facts about words and worlds are often all mixed together, as when the dictionary (about words) tells us that "edelweiss is a flower that grows at alpine altitudes", which is a plain fact about the world. If Quine is right that analytic and synthetic prepositions cannot be clearly discriminated (1951) then it follows straightforwardly that facts about words and the world cannot be separated either. Carnap (1946) proposed a dualism by which a sentence could be viewed in two modes, the material and formal, so as to express both possibilities, roughly as follows:

(F) "Caesar" is a symbol denoting a Roman Emperor.

(M) Caesar was a Roman Emperor.

Carnap's proposal was defective in many ways -- the sentences are not synonymous under translation out of English, for example -- but was implicated in the origin of Quine's later views, by providing an over-simple opposition, from his teacher, that he sought to overcome and refine.

The position of the present paper is the following: the persistent, and ultimately ineradicable, language-likeness of purported ontological terms (see Nirenburg and Wilks, op.cit.) means that we cannot ever have purely logical representations, purged of all language-like qualities. That enterprise is therefore ultimately doomed to failure , but should be pushed as far as it can be, consistently with the intelligibility of the representations we use (on this see also, Sowa 2004. As the history of Wordnet, and its popularity and productive deployment have shown, mixed, unrefined, representations can be useful, a fact formal critics find hard to understand or explain.

It is for this reason that data-mining research has been done for almost forty years (see Olney et al., 1966) on such human information structures as dictionaries, thesauri, ontologies and wordnets. Were they fully explicit, there would be little for such research to discover.

3. Ontological basics: a reminder

Before proceeding to the core of the paper, let us just remind ourselves again of the basic vocabulary for stating the problem: the initial issue in structuring an ontology is to decide what there is out there: are there basically individuals and sets of them (as Russell believed) or are there also types, concepts, universals etc., considered as separate entities from individuals, or perhaps as being among them? These are among the oldest intellectual questions known to mankind and there is no settling them. All one can do is make choices and take the consequences.

If one's task is roughly that of Linnaeus----- the taxonomy, classification or ontology of the natural species and genera in the world---then things are relatively simple, ontologically at least. You can assert canaries are birds with a branch structure:

Canary → ISA→ Bird

where the ISA link is set inclusion C, not set membership E, although one could write

Tweety → ISA → Canary → ISA → Bird

as long as the two ISAs are distinguished, thus:

Tweety → E → Canary → C → Bird

And this would give us the transitive inference that Tweety is a bird, and Linnaeus could have done this had he been interested in individuals and not only classes. In general, of course, inferences involving E are not transitive, as C is.

Part of the problem with describing an entity like Wordnet is that it does include Linnaean sub-nets taken straight from biology e.g.

Horse → C → Ungulate

If the world consisted only of sets and individuals, then the representational task is over, but the problem really arises when we move to the hierarchical relationships of concepts that may or may not reduce to sets and individuals, and above all, the problems that arise when we mix these, or refer to concepts in ways that do not reduce to claims about individuals and sets of them.

The pathological cases are well known, as in the apparently valid, but actually false, syllogism:

My car is-a (E) Ford.

Ford is-a (E) car company.

Therefore, my car is-a (E) car company.

The problem here can be seen in one of two ways: first, as the illegitimate chaining of ISAs (all of which are E, set membership), and therefore not valid, or, secondly, as the simple word sense ambiguity of the symbol "Ford", which refers to Ford cars as objects in the first line, and to the

car company with that name in the second. In the latter case, Ford is some form of conceptual individual, the Ford Car Company and which stands in a relationship to the set of cars it manufactures in a way that is hard to define, since it is not that of a concept to its corresponding extension set. The conceptual counterpart of a Ford car would presumably be "Ford" meaning something like "whatever it is that makes a car a Ford car". In this case, both premises above are true, but they refer to different objects under the same name, which also leads, by a different route, to an invalid conclusion.

The issues here are subtle and beyond those of the classic one of the duality of intensional concept and its extensional set, diagnosed by Woods in his famous "What's in link?" paper (Woods, 1975).

One can illustrate the core of the problem for ontologies, which is beyond the Woods duality and reference to fallacies like the one above, by considering a quite different area, treated roughly as follows in Wordnet. One can write down a (partial) taxonomy of religions, say, as follows:

Religion → Christianity v Islam….

Islam → Sunni v Shia v Ismaili….

Christianity → Catholic v Protestant v Orthodox…

Protestant → Anglican v Methodist v Baptist….

and so on.

This seems something like common sense but it is not at all like a Linnaean taxonomy or an ontology, because it is not clear how to interpret the implied ISA on each arrow link, nor exactly how any node is to be interpreted. If, for example, each node is a set of buildings occupied by a sect, or a set of believers, then the ISA links are set inclusion, provided we can assume the Linnaean disjunction: that no individual or class falls under two or more higher nodes. This is pretty safe in biology (since it is made up that way), though less so in, say, religion in Japan where many are both Buddhists and Shinto believers. But let us assume the disjunction is exclusive so as to continue.

The difficulty is that no such set theoretic interpretation (or model), about buildings or believers, is intended by the common sense interpretation of statements like the ones above, which is usually and loosely described as a taxonomy of concepts. This is sometimes expressed by interpreting the hierarchical links above as PART-OF, which can in turn be seen as either a relationship between concrete objects or conceptual ones or even both, and is (like C, but not E, transitive). Thus, one could write (now using → to mean “has as part”):

Body → Foot v Hand

Foot → Toe

Hand → Finger

to mean "a finger is part of a hand and a hand is part of a body and (so) a finger is part of a body". If we interpret the expressions that way we cannot at the same time be interpreting → as C since a set of fingers is not a subset of a set of hands. Nor can → be interpreted as E since a toe is not a member of a foot-set. So, the items indicated by the predicates must be "idealized

individuals", rather than particular ones, and that is a notion almost identical to that of a concept or intension or sense, namely, what it is to be an individual of a certain kind. Often, what it is that constitutes falling under a concept is the fact that the concept fits below its "mother node" above. Thus "being a finger" is, in large part, being part of a hand, and one could now reinterpret the earlier set of examples in this way, so that "being Catholicism" becomes being part of Christianity. Yet, much of what one wants to put into the definition of a concept does not come from hierarchical relations, of course. So, if we take:

"A Catholic priest is a male".

"A US President is over 35 years old".

These are both true statements about conceptual content that are only incidentally remarks about individuals falling under the description, i.e. the second is true of George W. Bush but not about him. The first is currently true, and true till now of the class of Catholic priests, but could change at any time.

The problem of conceptual inclusion, and how to interpret it in a way different from set inclusion or membership relations between objects covered by the concept, is the problem at the heart of the definition of an ontology. Such a structure is different from both a straightforward Linnean taxonomy/ontology (where relations are always set theoretic) on the one hand, and, on the other, the purely lexical thesaurus like Roget where a concept can be said to fall under another without any analysis or criteria being given for that inclusion.

A move frequently made at this point is to appeal to the notion of possible worlds, and to concepts as picking out, not just a set of individuals or sets in this world, but in some sub-set of all possible worlds. The chief appeal of this move is that it moves entities like the golden mountain to some set of worlds that does not include this one, and round squares to no world at all. Concepts that would once have been considered as expressing a 'necessary' or 'analytic' relationship, such as "animate cats" then appear in all words, or at least all worlds containing cats. It is often thought essential for formal reasons to constrain possible worlds to a single set of entities, whose (non-essential) properties may change from world to world. This seems an extraordinary constraint on possibility, namely that there is no possible world not containing, say, Tony Blair. Most people would have no difficulty imagining that at all.

This move will not be discussed further in this paper, as it is known to have no computational content, i.e. no process would correspond to searching among all possible worlds. Moreover, if Putnam's (1970) arguments have any force at all one cannot possibly know that there are no worlds in which cats are not animate.

In many conceptual ontologies the concepts are themselves considered individuals, so that set membership and inclusion relations can again be brought to bear. In the religions taxonomy above, we can, quite plausibly and in tune with common sense, consider Christianity and Islam as individual religions, members of the set Religions. If we then break the former into (sect-like) concepts of Protestantism and Catholicism then, if we wish to retain transitivity, Christianity and Islam will have to become sub-sets of religions and not members of the conceptual class above, at which point the problem returns as to what they are subsets of.

4. Is Wordnet an ontology?

Wordnet is certainly not an ontology in any ordinary sense, and about that Guarino and its other critics are surely right, but it does contain within it a whole range of relations including classically ontological ones such as set membership and inclusion (what we have called Linnaean). In what follows, I shall concentrate on Guarino's formal critique of WordNet, but he and his colleagues are intended only as examples of a school or style, and I could equally well have cited the work of (Smith, 1995).

Wordnet has at least the following relations:

Linnaean inclusion: ungulates ← horses

Simple subsets: shoes ← tennis shoes

Set membership: painters ← Picasso

Abstract membership: Carcompany ← Ford

Whole-part: Body ← hand ← finger

?Concept--component: Islam ← Sunni

?Concept-subconcept: Physics ← Gravity

As we noted above, Wordnet has many conceptual mother-daughters of the latter kind:

Religion → Islam → Sunni

Religion → Islam → Shia

Religion → Buddhism → Theravada

We also noted already that these cannot be interpreted or modelled by sets of or individual people (though WordNet actually tries this for one sense!), buildings, parts etc. It is not simply not plausible to interpret any of the above lines as set inclusion relations on, say, religious buildings or people. Because that is not what is meant by anyone who says "Sunni is a major form of Islam".

If therefore one takes the view that an ontology must at its core be a classification modelled by sets, as say the Rolls-Royce jet engine ontology is basically reducible to sets of components, or the SmithKlineGlaxo drug ontology to chemical components, then Wordnet is not one, both because it mixes such relations in with others and, most crucially, because its underlying relationship is synonymy, the relationship between members of a synset, and that is certainly not a relationship of sets of objects, and is not an ontological relationship even in the widest sense of that word. It is a logical triviality that one can refer to a concept as well as its instantiations, but this distinction is not well served if both cannot be ultimately unpacked in terms of sets of individuals in a domain. We noted earlier that the true statement:

The US President must be over 35 years old

does not refer to any particular president, but it is easy to state it in a conventional form so that it quantifies over all presidents. But, as we saw, simple quantification does not capture the intended meaning of one who says Sunni is a major branch of Islam, which is a claimed relation of concepts that goes well beyond saying that if anything a person and a Sunni they are a Muslim.

Wordnet clearly has ontological subtrees, often in the biology domain, yet it cannot be an ontology overall, nor is it a thesaurus, which we may take to mean words clustered by meaning relations, together with major upper level "meaning subsumption" classes e.g. words to do with motion, or games. It is often noted that Wordnet has no way of relating all its elements relevant to the notion of "game", sometimes referred to as the "game problem" in Wordnet (in that tennis, tennis shoe and tennis racquet, for example, are in widely separated parts of WordNet). Yet Wordnet's basic unit, the synset—a list of semi-synonymous words----is very much like that of a row of words in a classic thesaurus like Roget, which also has the "top level" heads and the possibility of cross referencing to link notions such as game.

Interestingly, Roget in his introduction to his Thesaurus gave as part of his motivation in constructing it the notion of biological hierarchies, though what he produced was in no way a set theoretic inclusion system. Efforts have been made over the years to provide formal structures that could combine, within a single structure, both set-theoretic inclusions (of an ontology) and the meaning inclusion relations of the kind that typify a thesaurus: the boldest of these was probably the thesaurus-lattice hypothesis of Masterman (2005) but that cannot be considered to have been generally accepted.

But the key distinction between Wordnet and an ontology is this: Wordnet has lexical items in different senses (i.e. that multiple appearance in Wordnet in fact defines the notion of different senses) which is the clear mark of a thesaurus. An ontology, by contrast, is normally associated with the claim that its symbols are not words but interlingual or language-free concept names with unique interpretation within the ontology. However, the position argued here is that, outside the most abstract domains, there is no effective mechanism for ensuring, or even knowing it to be the case, that the terms in an ontology are meaning unique.

This issue is venerable and much wider than the issue of ontologies: it is the issue of the interpretation of terms in all formal systems that appear to be the words of an natural language, but where their designers deny that they are. The issue is set out in full in (Nirenburg and Wilks, 2001) a paper in the form of a Greek dialogue where my character argued that, no matter what formalists say, the predicates in formalisms that *look like* English words, *remain* English words with all the risks of ambiguity and vagueness that that entails. This whole issue cannot be recapitulated here, but it turns on the point of what it is to know that two occurrences of a formal predicate "mean the same" in any representation. For example, it is generally agreed that, in the basic original forms of the LISP programming languages, the symbol "NIL" meant at least false and the empty list, though this ambiguity was not thought fatal by all observers. But it is exactly this possibility that is denied within an ontology (e.g. by Nirenburg, op. cit.) though there is no way, beyond referring to human effort and care, of knowing that it is the case or not.

Anecdotal evidence can be sought here in the representation language CyC, codings in which (1995) have been going on for nearly 30 years and where it is said that there is no way of knowing whether the basic predicates have changed their meanings over that time or not. There is simply no effective method for discussing the issue, as there now is for the ambiguity of word senses in text and their resolution.

5. Guarino's program for reforming Wordnet

Guarino (in Gangemi et al., 2001) has been a long term critic of Wordnet (Miller, 1995) and has proposed an alternative methodology for a formal ontology: OntoClean. I want to argue here that this is precisely the way in which WordNet cannot be cleaned up. The approach ignores the virtues of Wordnet and the many effective uses to which it has been put, computationally, and, given the argument above, there is no reason to believe that the kind of precision Guarino seeks is available for language terms of the kind WordNet (and by extension any ontology) consists of. There is a basic disagreement here about how far natural language can be "firmed up" in the way he proposes.

Guarino's critique of Wordnet begins from observations that it is a mixed bag of representations, a point conceded by its critics and defenders: so, for example, it mixes types for a given initial concept:

> apple given as fruit and food (only former is "necessary" for Guarino):
>
> window given as panel and opening (it cannot be both says Guarino):
>
> person given as living thing and causal agent (not "necessarily" the latter, says Guarino).

This is the "multiple appearance" of terms we noted above, and the way in which WordNet expresses sense ambiguity; it is one of its thesaurus, as opposed to ontology-like, features. Guarino's solution to the perceived problem is the provision of Identity Criteria (ICs) for all concepts: ICs make things what they are and allow their reidentification, which is to say, sufficient conditions for identity, later shifted to necessary conditions, which he thinks are easier to express. On his view, you cannot then hierarchically link concepts with different ICs e.g. ordered sets are NOT sets, a price he is prepared, apparently, to pay, since the identity criterion for being an ordered set is quite different from that for being a set. Thus, for him, the concept *person* cannot be subsumed by *physical object* because the IC for one is quite different from that for the other. Meeting the sufficient conditions for being a person are sufficient for being a living thing, but not sufficient for being a physical object since *disembodied persons are not impossible*. One sees immediately and from the last case that these issues will be tricky, even within Guarino's frame of reference, and that it may be very difficult to get general agreement, as on the last italicised statement.

But can we have a useful hierarchy if a person cannot be seen to be a physical object? Consider the interpretation of:

> Smith tripped and fell on the guinea pig and killed it.

Bergson (1900) in his classic text on humour defined jokes as being when human beings fell under physical laws, as in the standard banana skin scenario, a notion impossible even to state for Guarino since persons do not fall under physical laws, not being physical objects: "...person should not be subsumed by physical object (as a physical object, a body has persistence conditions different from a living being). Yet, these ISA links exist in WordNet" (Oltramari et al., 2002).

Guarino also turned to Nirenburg's Mikrokosmos ontology (Mahesh et al., 2000) as a target for reform, and objected to the following subsumptions:

social event -- communication event
mental event /
physical event -- perceptual event

and proposed the alternative:

social event -- communication event
mental event -- perceptual event
physical event /

where the lower figure has what he calls (vertical) ontological levels (Oltramari et al., 2002). One could still doubt whether any effective ontology could be set up consistently with his principles? Surely mental events and perceptual events do not have the same identity criteria, an observation made in most elementary philosophy classes, as when one notes that mental event are observable only by one person.

One must concede that Guarino may well be able to develop some less restrictive form of his principles, but here we are concerned only with their spirit which may be worse than what we have already shown: if, that is, the issue is really one of differing IC s at each hierarchical level. How can any two hierarchical levels have the same IC s, since, by definition they share features but necessarily have differentia, just in virtue of being different levels? Again, one will expect those differentia to be part of the IC for the appropriate level, in which case how can there be hierarchical relations at all? Canaries have ICs quite different from those of birds, namely *being yellow* among other things. If the IC is taken seriously a Canary cannot be a Bird.

Guarino does not, of course, intend that Canaries/Birds subsumptions are not in an ontology; my question is how, on his published claims, he can avoid the conclusion in a non-arbitrary way. There is a constant danger of higher nonsense in this area of formalization; consider:

"A piece of coal is an example of a singular whole. A lump of coal will still be a topological whole, but not a singular whole, since the pieces of coal merely touch each other, with no material connection. It will therefore be a plural whole" (Oltramari et al. p.4).

It may well be that they really intend "pile" here and this is no more than a slip in their understanding of English, but that is not much cause for comfort because it only serves to confirm my point of how tightly these issues are involved with the understanding and use of the natural language that is more or less isomorphic with the formal language. Yet this point is seldom noticed by the formalizers, it just seems to them an irritant, as opposed to a fundamental issue. On the view advanced here, they take as technical terms, on which to base a theory, words of a language (English) which will not and cannot bear the interpretations required. They cannot express or even admit the senses of the words they use, which is something WordNet, for all its faults, explicitly allows.

Perhaps this issue is very old indeed: "The Languages which are commonly used throughout the world are much more simple and easy, convenient and philosophical, than Wilkins' scheme for a real character, or indeed any other scheme that has been at any other times imagined or proposed for the purpose". This is Horne Tooke (quoted by Roget at the end of his 1862 Preface to his Thesaurus) attacking Wilkins, perhaps the first ontological formaliser in the 17th Century.

One way out of this impasse may be something explored by Pustejovsky (2001), namely a linking of Guarino's ontological levels to his own theory of regular polysemy. Thus Guarino's levels:

mental event

physical event

social event

might be extended to cover classic regular polysemy ranges like:

{Ford = company, car, ?management}

where, for certain concepts, there is a predictable set of functions they can have, and this can be considered as the predictable part of the word-sense ambiguity (alias polysemy) problem. Pustejovsky's approach is not one of multiple entries for a word, in the way a standard lexicon lists a word once for each sense, but a more "compressed" view of a single entry plus lexical rules for its "expansion". However, any link from phenomena like this to 'ontological levels' would require that the "multifunctional entity" (e.g. Ford) would still appear in more than one place in a structure, and with different interpretations, which would then make it, again, a thesaurus not an ontology.

There is a further issue, much broader than any touched on so far, and which relates to the thrust of much modern philosophy of meaning. Within the Anglo-Saxon world, there has been general acceptance of the arguments of Quine (1951) and Putnam (1970), refining as they did earlier positions of Wittgenstein, that it is not possible to continue the two millennia-old Aristotelean analysis stated in terms of what features a thing or class much have to be what it is: which is to say necessary properties are a dubious notion, and that the associated analytic-synthetic distinction among propositions, deriving from Kant, cannot be maintained.

But there is no hint in Guarino that he is aware of any of this, only that there may be practical difficulties---and he concedes explicitly that finding conditions, necessary or sufficient, for IC s may be very hard-----difficulties in detail, in specifying necessary and sufficient conditions for something's being in a certain class; he never allows that the whole form of analysis is dubious and outdated.

6. Empirical issues

Much of this discussion is moot in the sense that, whatever its faults, WordNet continues to be useful in a large range of experimental activity, and defies systematic "cleaning up". Other structures such as Open Directory continue to be used as ontologies for search, although (as we noted earlier) they can be shown to contain contradictions (in the sense of hierarchies with forms like A→B→A) and so can have no effective models of any kind. More seriously, there is a long tradition of work starting in the 80s of extracting hierarchies empirically from machine readable

dictionaries, which appear to have the form of ontologies, at least in part, e.g. physical object→artefact→instrument→violin, and heuristics were sometimes applied during such construction so that an ambiguous term like instrument could be tagged as "musical instrument" within that particular sequent of set inclusions. This and much related work on deriving proto-ontologies from linguistic objects like dictionary is described in (Wilks et al., 1996); although fairly complete hierarchies were obtained by such methods, they do not provide anything that could reasonably be called an ontology, when the sequent's contain items, ambiguous or tagged, like religion→sect→Shi'ism, where, as we noted earlier, the relationship is purely conceptual (hyponymy in linguistic terms) and not one of real world objects, nor even of conceptual "objects" that can be equated to non-conceptual objects, in the way automobile→Ford→Mustang can be seen as a relation of concepts (the Mustang model) but equally well as set inclusion of classes of real cars.

Much of the research on ontologies as possible products from machine-readable dictionaries concentrated on how to establish reliable and generally acceptable forms at the very top levels of such a hierarchy, where relationship to real world objects is quite remote, and with inclusions like entities→living things. It is these very top levels that the Wordnet creators avoid constructing, perhaps wisely, but a great deal of work has gone on (e.g. Wilks et al., 1996, Hovy 1998, Hovy & Nirenburg, 1992) in trying to get agreed relationships between very abstract English terms at these highest levels. Roget's thesaurus avoided this, but in a different way from Wordnet, by simply declaring about 1000 abstract terms as the upper levels of the thesaurus (e.g. Motion) but which were not themselves put into hierarchical relations, thus avoiding the "top level" problem, and that of the single node under which everything falls.

Much recent work has turned to the population of ontologies empirically from text corpora (e.g. Brewster et al., 2005) and the attempt to use text extraction techniques either to create an ontology from scratch or to augment and update one constructed a priori. This is a separate topic deserving its own treatment, but it can be argued even here and in brief that one could hold the position that the only way to avoid ontologies being always criticised as creative inventions of individuals (as linguistic rules used to be) is by providing a general methodology that allows for their creation and maintenance by automatic means. If texts, in the broadest sense such as the World Wide Web, are our repository of explicit knowledge, then it must in principle be possible to extract that knowledge and convert it to a compressed ontological form. The objection always comes that much human knowledge is in fact implicit, learned as a child, and never stated, rather in the way that instructions in a phone box tell you about the functioning of phones where the box is, but never that you need to place them to your ear—this is implicit and taken for granted. Conversely, many stated "facts" in corpora will in fact be untrue: just as some obvious X-ISA-Y's are never stated, some X-ISA-Ys are untrue, or could be said to reflect individual or wrongly-shared ontologies e.g. Tomato→Vegetable or Whale→Fish.

7. What there is

At the beginning of this paper, we distinguished senses of "ontology" and resile ourselves from the discussion of the more traditional ontological enquiry, so as to get to where the arguments currently go on, or should go on. But ancient questions have a habit of returning to bite one at the end, even though, in this paper, we have taken as robust position, in the spirit of Quine (op. cit.) that whatever we put into our representations --- concepts, sets, etc. --- have existence, at least as a polite convention, so we can continue with the discussion. It is interesting to note at the

end that our initial scientific examples of genes are by no means as straightforward as we pretended.

Suppose we ask again, what are the ontological "objects" in genetics, say in the classic Drosophila data base FlyBase (Morgan et al., 2003)? FlyBase ultimately grounds its gene identifiers --- the formal gene names --- in the sequenced Drosophila genome and associates nucleotide sequences parsed into introns, exons, regulatory regions etc with gene ids. However, these sequences often need modifying on the basis of new discoveries in the literature (e.g. new regulatory regions upstream from the gene sequence are quite frequently identified, as understanding of how genes get expressed in various biological processes increases). Thus the "referent" of the gene id. changes and with it information about the role of the `gene'. However, for most biologists the `gene' is still the organising concept around which knowledge is clustered so they will continue to say the gene `rutabaga' does so-and-so quite happily even if they are aware that the referent of rutabaga has changed several times and in significant ways over the last decade. The curators and biologists are, for the most part, happy with this, though the argument that the Drosophila community has been cavalier with gene naming has been made from within it. This situation, assuming this non-expert description is broadly correct, is of interest here because it shows there are still ontological issues in the original sense of that word: i.e. as to what there actually IS in the world. More precisely, it calls into question Putnam's optimistic theory (1970, cited elsewhere in this paper) that meaning can ultimately be grounded in science, because scientists know the true criteria for selecting the referents of terms. The Drosophica case shows this is not so, and in some cases the geneticists have no more than a hunch, sometimes false in practice, that there are lower level objects unambiguously corresponding to a gene id., in the way that an elementary molecular structure, say, corresponds to an element name from Mendelev's table.

8. Conclusions

Ontological and lexical resources are resistant to logical “fixing” and can themselves be objects of research, as well as serving practical functions within NLP and AI programs. Attempts to clean up resources may make them impossible to retain, but an imperfect linguistic resource is better than none. Ontological and lexical resources do not differ in content, only in principle, and the fact WordNet is of mixed type does not disqualify it from practical use. Attempts to clean up resources should actually be done by automating their construction as best we can, rather than by post hoc logical devices. It may however be intellectually tidier to segregate scientific ontologies of the Linnaean type (interpretable under set inclusion) from those with irreducibly conceptual content, whose distinction from thesauri remains unclear.

Acknowledgements

The paper is indebted to an anonymous referee for the stimulus for the elementary explication at the end of the Introduction, and to Professor E.W. Briscoe for the discussion of FlyBase genes. Many discussions with AKT colleagues (EPSRC Interdisciplinary Research Centre Grant, 2000-2006) have contributed to this paper especially with Christopher Brewster, but the mistakes are, as ever, my own.

References

Bergson, H., 1900/1924. Le rire: essai sur la signification du comique, Paris, Editions Alcan.

Bosanquet, R., 1946. Remarks on Spinoza's Ethics, In Mind, 55.

Braithwaite, R., 1953. Scientific Explanation, Cambridge, Cambridge University Press.

Brewster, C., Iria, J., Ciravegna, F. & Wilks, Y., 2005. The Ontology: Chimaera or Pegasus, In Proc. Dagstuhl Seminar Machine Learning for the Semantic Web.

Carnap, R., 1946. Semantics and Necessity, Chicago, University of Chicago Press.

Gangemi, A., Guarino, N. & Oltramari A., 2001. Conceptual Analysis of Lexical

Taxonomies: The Case of WordNet Top-Level. In Proc. of FOIS2001.

Hirst, G., 2004. Ontology and the Lexicon. In (Staab and Studer, eds.) Handbook on Ontologies, Springer Verlag, Berlin.

Hobbs, J., 1983. Ontological Promiscuity, In Proc. 23rd Annual Meeting of the Assn. Comp. Ling.

Horrocks, I., Patel-Schneider, P. & van Harmelen, F., 2002. Reviewing the design of DAML+OIL: An ontology language for the semantic web. In Proc. of the 18th Nat. Conf. on Artificial Intelligence (AAAI'02),AAAI Press.

Hovy, E., 1998. Combining and standardizing large-scale, practical, ontologies for machine translation and other uses. In Proc. 1st International Conf. on Languages, Resources and Evaluation, Granada, Spain.

Hovy, E. & Nirenburg, S., 1992. Approximating and Interlingua in a principled way, In Proc. DARPA Speech and Language Workshop, Hawthorne, NY.

Katz, J., 1972. Semantic Theory, Harper and Row, New York, NY.

Lenat, D., 1995. CyC: a large scale investment in knowledge infrastructure, In Comm. ACM, 38.

Mahesh, K., Nirenburg, S. & Beale, 2000. Toward full-text ontology-based word sense disambiguation. In (Nicolov and Mitkov, eds.) Recent Advances in Natural Language Processing II, John Benjamins, Amsterdam.

Masterman, M., 2005. Language, Cohesion and Form: selected papers of Margaret Masterman, In (Wilks, ed.), Cambridge, Cambridge University Press.

McDermott, D., 1976. Artificial Intelligence meets Natural Stupidity, reprinted in (Haugeland ed.) 1981, Mind Design, Cambridge, MA, Bradford Books/MIT Press.

Miller, G., 1995. Wordnet: a lexical database for English, In Comm. ACM 38.

Morgan, A., Hirschmann, L., Yeh, A. & Colosimo, M., 2003. Gene Name Extraction Using FlyBase Resources, In ACL Workshop on Language Processing in Biomedicine, Sapporo, Japan.

Nirenburg, S. & Wilks, Y., 2001. What's in a symbol: ontology, representation and language. Journal of Experimental and Theoretical Artificial Intelligence (JETAI).

Olney, J., Revard, C. & Ziff, P., 1966. Some monsters in Noah's ark. Research.Memorandum, Systems Development Corp., Santa Monica. CA.

Oltramari, A., Gangemi, A., Guarino, N. & Masolo, C., 2002. Restructuring WordNet's top-level: the OntoClean approach, In Proc. LREC02, Spain.

Peters, W., 2001. An Exploration of Figurative Language Use in WordNet , In Proc. Corpus Linguistics 2001, University of Birmingham.

Pollack, J.B., 1990. Recursive Distributed Representations, In Artificial Intelligence 46. Representations, Artificial Intelligence 46(1), 77-105.

Pustejovsky, J., 2001. Type Construction and the Logic of Concepts, In (Bouillon& Busa, eds.) The Syntax of Word Meaning, Cambridge, Cambridge University Press.

Putnam, H., 1970. Is Semantics Possible? In Metaphilosophy.

Quine, W.V.O., 1951. Two Dogmas of Empiricism in The Philosophical Review.

Schvaneveldt, R.W., Durso, F.T., Goldsmith, T.E., Breen, T.J., Cooke, N.M., Tucker, R.G. & DeMaio, J.C., 1985. Measuring the structure of expertise. International Journal of Man-Machine Studies, 23, 699-728.

Smith, B., 1995. Formal ontology, common sense and cognitive science. In Int. Journal on Human-Computer. Studies, Smith 43(5-6).

Sowa, J., 2004. The challenge of knowledge soup. In Proc. First International WordNet Conference, Goa, India.

Wilks, Y., 1968. Argument and Proof in Metaphysics, from an empirical point of view. Unpublished PhD thesis (advisor: Professor R.B. Braithwaite), University of Cambridge.

Wilks, Y., 1994. Stone Soup and the French Room: the empiricist-rationalist debate about machine translation. In (Zampolli, Calzolari and Palmer, eds.) Current Issues in Computational Linguistics: in honour of Don Walker. Kluwer: Berlin.

Wilks, Y., Slator, B. & Guthrie, L., 1996. Electric Words: dictionaries, computers and meanings. Cambridge, MA: MIT Press.

Woods, W., 1975. What's in a Link: Foundations for Semantic Networks. In (Bobrow and Collins eds.), Representation and Understanding: Studies in Cognitive Science, New York: Academic Press.