International Workshop on Formal Ontology

in Conceptual Analysis and Knowledge Representation

17-19 March 1993
Padova, Italy

Edited by N. Guarino and R. Poli
Ladseb-CNR Internal Report 01/93

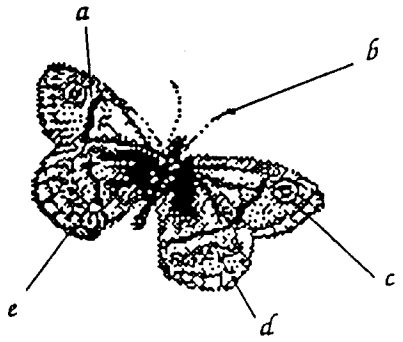

Organized by:

LADSEB-CNR - Institute for Systems Theory and Biomedical Engineering of the Italian National Research Council - Padova
Centro Studi per la Filosofia Mitteleuropea - Trento

Sponsored by:

AI*IA - Italian Association for Artificial Intelligence
CNR - Special Project on "Sistemi Informatici e Calcolo Parallelo"
University of Padova
University of Trento

Program Committee:

Giuseppe Attardi, Dip. di Informatica, Pisa
Amedeo Cappelli, ILC-CNR, Istituto di Linguistica Computazionale, Pisa
Pierdaniele Giaretta, Istituto di Storia della Filosofia, Padova
Nicola Guarino (co-chair), LADSEB-CNR, Padova
Diego Marconi, Dip. di Ermeneutica Filosofica, Vercelli
Roberto Poli (co-chair), Dip. Teoria Storia e Ricerca Sociale, Trento

A Linguistic Ontology

Kathleen Dahlgren

Intelligent Text Processing, Inc.
Santa Monica, CA 90403 kd@itp.socal.com

This conference addresses a spectrum of issues in the foundations of ontological representation. This paper defends the choice of a linguistically-based ontology for natural language processing and demonstrates that a single commonsense ontology produces plausible interpretations at all levels from parsing though reasoning. The paper explores some of the problems and tradeoffs for a method which has just one content ontology.

The AI application we are interested in would parse and interpret unrestricted text in unrestricted domains. It would select the one most plausible interpretation of each sentence and embed it in the discourse context, resulting in a fully-disambiguated cognitive model with all references resolved. Such a system needs to know something about everything people talk about. According to our critics, such as Fodor, it would have to have infinite knowledge. But our approach takes advantage of the psychological fact that only a very shallow layer of knowledge comes into play in initial linguistic interpretation. In fact, in writing interpretive modules, it has emerged that nearly 80% of the commonsense reasoning takes place with ontological knowledge alone.

Ontology is about "what there is". The world itself manifests an infinitude of objects, changes and relations among them. Moreover, ontologies can in principle be grounded in a number of different ways - in the world, in the mind, in the culture or in the language. Depending upon AI application, almost any of the possible perspectives on the real world is a potential way of classifying it. In practice, these perspectives are limited to those which are useful to some human institution which wants the AI application. The objective for world-relative ontologies is to approximate the characterization of the distinctions in a domain as closely as the relevant science permits.

Such world-relative grounding is very different from human-relative grounding. The latter investigates those classifications which have already emerged and have become codified either in mental wiring, cultural norms or language. Linguistic ontology is a reflection of culture, of the way people in a culture interact with the infinitude of objects, events and relations. The relative importance of the objects in a given culture is reflected in the type and extent of classification (Berlin 1972).

For linguistic ontology, a language reflects its culture's view of the world at several levels: morphology, syntax, formal semantics and lexical semantics. Some AI researchers in natural language have chosen to ignore these strictly linguistic indicators of ontological preferences in a language, and construct linguistically-independent

ontologies. However, there are a number of significant advantages to basing the ontology upon a particular language, its syntax and its word senses. First of all, these distinctions are the very ones which are needed in the interpretive tasks of syntactic and word sense disambiguation, formal semantic interpretation, temporal reasoning and discourse reasoning. Secondly, when the ontology is tied directly to the word senses (which are also syntactically determined), the resulting interpretation is far more precise than is possible when the knowledge base and lexicon are separated. Thirdly, the classes which a language countenances are time-worn, and likely to have sustained value. In particular, linguistic classification remains true across all domains (while perhaps elaborated or revised for domain terminologies). Thus a linguistically-based ontology is very transportable. The Intelligent Text Processing (ITP) ontology has easily moved from geography to finance to terrorism to human factors engineering. Changes took the form of additions. Fourthly, since a linguistically-based ontology corresponds to the way people think about objects, it is a useful way to predict their thinking about the knowledge in structured databases. Finally, inheritance of descriptive knowledge has a cognitive basis in such an ontology.

It might seem that this type of ontology would fail in a machine translation task. However, tied as it is to psycholinguistics, to the extent that cultures are similar, and thereby have similar visions of the basic "joints" in the world, this ontology should provide a useful start in machine translation. Since word senses are individually attached to the ontology, problems of partial synonymy are eliminated in the method. However, it would involve the work of attaching all of the word senses of both languages.

ITP's Ontology

There are two different ontologies which any linguistic treatment requires-- the ontology of types required in the expression of truth conditions and the ontology of those types which organizes linguistic conceptualization - that is, lexical content. For the formal semantic types those justified in the DRT framework are assumed (Kamp, 1981, Asher, 1987). These types are individual, event, state, property and proposition. For the content ontology, the ITP team developed its own approach (Dahlgren and McDowell, 1986, Dahlgren, 1988). This approach is called "naive semantics" in deference to the fact that the classifications and descriptions in the lexicon reflect naive theories of the world, rather than scientific truths (Hayes, 1985). Given that semantic type and syntactic category are orthogonal, the naive semantic ontology handles nouns which name objects, nouns which name events or states, nouns which name properties, verbs which verbalize instruments and verbs which verbalize relations.

Knowledge of objects and events is of two types - ontological and generic, corresponding roughly to the difference between classification and description. However, it is always possible to treat any classification as a description, and vice versa. The principles we used to determine which distinctions would be found in the ontology as opposed to the

generic knowledge for words are of interest. First of all, people reason with classificatory knowledge monotonically. They think that something is a plant or it isn't, is a manmade object or isn't. On the other hand, properties of objects tend to be typical rather than necessary. While typical lemons are yellow, brown and green lemons are acceptable. Thus a potential ontological distinction must license monotonic reasoning. The most basic cuts in the ontology should reflect the significant joints in the world, as perceived by English-speakers, and as embodied in the English language. The first question to consider was what the different types of things were, and what was the evidence for them. The second question was what form the ontology should take. As for the first question, there were three criteria we settled upon: (1) the distinction plays a role in selectional restrictions; (2) psycholinguistic studies provide evidence of the psychological reality of the distinction, as in Rosch et al (1976); (3) the distinction is one sanctioned in Western philosophy. Subsequently new criteria were added which modified some of the nodes in relation to generalizations discovered in reasoning about disambiguation, coherence and anaphora resolution. For example, "route" and "measure" nodes under "abstract" were found to be required for reasoning with prepositional phrases.

The resulting ontology handles two major distinctions differently than other ontologies, sentient/non-sentient and social/natural. Philosophers have shown that these distinctions are fundamental to Western epistemology (Strawson, 1953, Quinton, 1975). Cognitive psychologists have demonstrated extremely early recognition of these distinctions in infants and small children (Gelman and Spelke 1981, Keil 1989) Interestingly, they also find their way into many layers of linguistic interpretation. Knowledge of sentience guides reasoning in parsing because of selectional restrictions of verbs and case roles of prepositional phrases. It is needed in anaphora resolution and coherence reasoning. Of any node in the ontology, the sentient one is by far the most frequently accessed. The ITP ontology has the sentient/non-sentient distinction high up, and sister to natural and social. This treatment means that people are classified in at least four ways (as human animals, as thinking persons, as members of institutions and as roles). That people are multiply classified has been amply demonstrated in the psycholinguistic literature (Cantor and Mischel, 1979). The advantage is that all sentients ("John", "secretary", "the United Nations", "the mob") are under a single node for satisfaction of the selectional restrictions of a verb like "say". Furthermore, all generic knowledge of humans qua humans ("has a heart", "breathes") is inherited under the "animal" node, while role knowledge ("works for someone") is under the sentient and social nodes.

The other philosophically-based distinction, that between natural and social, distinguishes those entities or events which only occur or arise in the context of human society from those which have natural origins. The distinction is particularly relevant in selectional restrictions for verb complements (objects,instruments, directions).

These distinctions are best represented as an acyclic directed graph rather than a tree. This form permits cross-classification. The

root of the tree is "entity". Initially this node crossclassified as

entity -> (abstract v real) & (individual v collective)

real -> (physical v temporal v sentient) & (natural v social)

(Recently, the collective distinction has been removed from the ontology for reasons explained below). This type of crossclassification is found elsewhere in the ontology, as well. The distinctions under real are illustrated below:

natural	social	
physical	rock	knife
sentient	man	programmer
temporal	earthquake	party

Cross-classification streamlines the reasoning, because it captures similarity and difference simultaneously. Thus "secretary" is like "person" in sentience, and unlike "person" in being essentially social. Thus when the time comes for satisfaction of verb selectional restrictions, both "secretary" and "person" can be subject of "say", but in nominals, only "secretary" can occur in the context "___to sentient". In fact, only roles under the node "assistant" and "agent" are allowed in the latter context. Representationally inside the system, this ontology is represented as first order logic as suggested by Hayes (1985).

It is important to note that the resulting ontology is not the only one which could be developed for use in natural language processing. It would in fact be preferable to run several different ontologies in parallel, because that is the way objects are classified cognitively. For example, people tend to classify plants by shape - tree, bush, grass, vine, but at the same time they understand some of the biological classification (such as deciduous vs evergreen) which puts plants into a different system. Here the possibility exists to create a cross-classification, if both systems of classification are important enough for linguistic reasoning. Another example turns on the natural/social distinction. For any given food, that distinction is simplistic. In fact, foods undergo stage after stage of social intervention from domestication and cultivation ("wild boar" vs "pig", "rye grass" vs "rye"), to distribution ("pork","sack of rye") to processing ("bacon","rye flour", "rye bread") to purchasing and meals. Foods in the store have a "store" ontology (meat, deli, produce, canned goods); foods in meals have a "meal" ontology ("appetizer", "main dish", "side dish", "dessert "). Humans classify objects multiply because classification is human-relative, and humans have a variety of needs and interests when interacting with objects. At times humans interact with objects as the grower, the distributor, the cook or the consumer of food. Similarly with all other objects, such as tools, and all types of commodities (people build them, buy them and use them). The way events are classified is even more complex. We chose to stay with just one naive semantic ontology, and represent information about other ontologies as generic knowledge, mainly because of processing time

considerations.

A Single Linguistic Ontology for Many Levels of Linguistic Processing

The ITP ontology has been empirically tested both by expanding it to thousands of nodes for all of the senses of thousands of nouns and verbs, and by using it in several levels of linguistic interpretation from parsing through reasoning. In this section each use of the ontology is illustrated.

Word Sense Disambiguation: Each lexicalized sense of each word is represented separately and given an ontological attachment. (Lexicalized metaphorical extension is recognized as separate word sense). Often (though not always), there are syntactic distinctions as well as semantic ones between senses. Consider some of the senses of the word "symphony" , which have ontological attachments as in (1).

(1) social event - the playing of the symphony
hard copy - the paper version
discourse - a structure of musical notes
institution - the players structured as a symphony company

Just on ontological grounds with selectional restrictions or reasoning rules, the correct sense of "symphony" in the sentences of (2) is selected by the ITP disambiguation module.

(2) Fred saw Mary at the symphony. (social event)
Fred went to Greece with the symphony. (institution)
Fred picked up the symphony. (hard copy)
The symphony lasted one hour. (social event)
The symphony formed last year. (institution)
Fred wrote the symphony last year. (discourse)

Notice that these words often disambiguate each other - because "symphony" has no "human" sense, "pick up" cannot have the reading where it means meet socially. Conceivably it could also mean transport, and that is determined by discourse context. Without further context, the module selects the "physically with the hands" reading of "pick up" because it is more frequent in English than the transport reading.

Syntactic disambiguation: Disambiguation takes place during parsing, pruning early to cut down on the combinatorial explosion.

For example, the sentences (3-5) illustrate three attachment points for a prepositional phrase: verb-modifying, verb-phrase modifying and sentence-modifying, respectively.

(3) Sara compensated Mary with a bonus.

(4) Sara compensated Mary with alacrity.

(5) Sara compensated Mary with certainty.

The verb-modifying attachment in (3) is triggered both by a syntactic feature of the sense of "compensate" meaning "pay", namely that it subcategorizes for a "with"-phrase, and by the fact that the object of "with" is ontologically an asset. The verb-phrase modifying attachment in (4) is determined by the "manner" prepositional phrase rule, which requires certain prepositions, and certain abstractions such as emotional states like "alacrity". The sentence-modifying attachment in (5) is found by noting that "certainty" is an epistemic abstraction. Often syntactic and ontological knowledge interact in syntactic disambiguation, underlining the value of tying the ontology to the lexicon.

Ontology and Case-role: Certain prepositions in combination with certain ontological classes form case-role generalizations. An example is that any propositional head modified by a PP with "about", "on", "concerning" or "regarding" has case-role of "theme", regardless of the ontological class of the object of the preposition.

Complement clause disambiguation: Post-verbal infinitivals with certain verbs are syntactically ambiguous, as illustrated in (6-8).

(6) The woman needed the bicycle to ride.

(7) The woman need the man to ride.

(8) The woman needed the bicycle to succeed.

To select the correct attachment point and argument structure for the infinitival, the ontology is inspected along with selectional restrictions. Thus, "bicycle" is a vehicle, and therefore a possible object of ride, so that in (6), the attachment is to the NP headed by "bicycle", while in (7) since "man" as human is a plausible subject of "ride", and not a plausible object, "man" is chosen as subject, and the infinitival is attached to the verb-phrase as complement. On the other hand in (8), "bicycle" as vehicle is not selected as subject of "succeed" so succeed attaches to the sentence as a purpose clause with "woman" as subject.

Nominal Disambiguation. When the head of a noun phrase is a nominal, the modifying prepositional phrases fill the theta grid of the corresponding verb (MacDonald, 1982). In (9), "treason" is a crime which fills an object role of one reading of "charge", while in (10), Jones is sentient, and so the plausible subject of a number of different readings of "charge.

(9) Charges of treason

(10) Charges of Jones

Noun compounds fall into a wide variety of predictable patterns with certain ontological classes. For example, in a phrase consisting of a noun attached to "material" and a following noun which is a manmade object, often the relationship is one of material, as in "leather jacket", "wood box". Another type of nominal is agentive. In "tapestry

runner", the fact that "tapestry" is a type of fabric disambiguates "runner", and in "long-distance runner" the fact that "long-distance" is a measure selects another reading of "runner".

Formal Semantics: The relational ontology for verbs and some eventive nouns reflects Vendlerian aspectual classes, so that every verb is marked as an event or state, and events are broken into activities, accomplishments and achievements (Vendler, 1967). In addition verbs are cross-classified as social or natural, and as either mental, emotional or nonmental. These distinctions enable the formal semantic module to assign semantic types to verbs and nominals. Events are marked with an event type reference marker, states with a state type reference marker, and so on. In addition, aspectual classes can be used in a treatment of tense, aspect and truth conditions. Formerly the individual/collective distinction fed the quantificational rules, but as will be described below, this distinction turned out to be inadequate and has been replaced with other syntactic and semantic lexical knowledge. Obviously treatments of tense and aspect would draw heavily upon a Vendlerian ontology.

Anaphora Resolution: The use of the ontology in pronoun resolution is well-known and standard. In addition, the ontology informs rules which guess antecedents of definite anaphors, distinguishing among event, state and individual types. For example, ontological knowledge informs the correct choices of antecedents in (11-13).

(11) The secretary bought a brand-new Mazerati. The woman ... (12) The secretary bought a brand-new Mazerati. The vehicle .. (13) The secretary bought a brand-new Mazerati. The move ...

Illustrating event types, knowledge of relational ontological classes makes possible the correct choice of antecedent in (14).

(14) After gathering in the neighboring villages, the guerrillas charged the base. The attack took place at dawn.

Since the nominal "attack" is an achievement, and "gather" is an activity, "gather" is rejected as the antecedent of "the attack". Ontological knowledge is only one of the types of knowledge required for anaphora resolution. In a text with a number of achievements accessible to a definite anaphor, additional generic knowledge of event implications would be required to find the most plausible antecedent.

Methodological Issues.

The drawbacks of a single ontology of this type arise from the multiplicity of world views people adopt, and from the complexity of the linguistic/conceptual scheme interaction. The first problem was described above in relation to foods and commodities, and there are other examples. A surprising one has to do with the role ontology. The ITP ontology classifies roles according to social/educational class: professional, worker, businessman, official, etc. A perfectly plausible alternative would be to categorize roles in terms of activity or industry

domains, so that "doctor" and "nurse" would fall under "medical personnel" and "lawyer" under "legal personnel". Regardless of which way this decision goes, neither treatment handles certain role terms which cut across all of the social classes and activity domains. These are the terms "president", "vice-president", "secretary" and so on. ITP invented a special "hierarchical role" node, but this wasn't a satisfying solution.

Another problem area concerns the states of objects. Some verb selectional restrictions require certain types of plants or minerals, so that a commonsense ontology of objects and stuffs was settled upon. However, other verbs require these in their solid or liquid form. Worse, practically anything physical in the proper perspective can be viewed as a liquid: Fred poured the sand through the funnel; The people poured out of the building; The ships poured out of the harbor. (Example due to Vic Liptak).

Parts require a different ontology. Parts pose special difficulties because they share unpredictable properties of the objects they are parts of. For example, a human hand is living, but it isn't human. It is self-moving, but it isn't a vertebrate. And of course, disembodied, can it really be called living? Similarly, a carburetor is part of an engine, but it isn't an engine. And engines are parts of vehicles, but are not self-moving. Furthermore, many parts occur in commodity contexts, and in those contexts are separable conceptually from the objects they become part of, as for example nails, circuits, handles, and so on. But the key test for being a part and not a whole is conceptual inseparability from some whole.

The complexity of the linguistic/conceptual scheme interaction is best exemplified with the problem of collectives. First of all, early on it become clear that masses, sets and structures had different semantic and syntactic properties, so that the collective node was broken into those three. But the node collective still had problems. First of all, syntactic mass, which licenses singular agreement for determinerless nouns, is not the same as semantic mass: "skiing" is not a mass in the sense that "gold" is. Furthermore, not all semantic masses can occur in syntactic constructions for masses: "some gold", but not, *"some tidewater". So the semantic mass in the ontology wouldn't work to handle the syntactic constraints. Secondly, it would have been very nice if all nodes under collective had licensed "they" anaphora, but that was not true. Not masses, and only some sets and some structures, license "they" anaphora. "the school" is a structure which licenses "they" anaphora, but not "the ship" (the physical structure). "the chessmen" is a set which licenses "they" anaphora, but not "the forest". It would have been even better if a single node had characterized the possible subjects and objects of collective verbs like "gather" and "assemble". Not even a test like "collective and self-moving" works for "gather", because there are expressions like "the dust gathered" which don't satisfy it. Another set of collective verbs, such as "merge", requires a plurally marked subject, so that "the files merged" is good, but *"the herd merged" is bad. A big drawback of having a parallel collective ontology is that there is duplication of the entire complex ontology, but collectives are very

sparse. As a result ITP has chosen to handle collectives with syntactic and generic features. The syntactic features handle properties like singular agreement with no determiner, and the semantic features control quantification, and verb agreement.

Acknowledgements: This ontology was developed with the assistance of Claire Chyi, Vic Liptak, Carol Lord, Joyce McDowell, Susan Mordechai, Hajime Wada and Karen Wallace.

Bibliography

Asher, N. 1987. A Typology for Attitude Verbs and Their Anaphoric Properties. Linguistics and Philosophy. 10:125-198.

Berlin, B. Speculations on the Growth of Ethnobotanical nomenclature. Language in Society 1:51-86.

Cantor N. and W. Mischel. 1979. Prototypes in Person Perception. Advances in Experimental Social Psychology. 12:3-52.

Dahlgren, K. and J. McDowell. 1986. Kind Types in Knowledge Representation. Proceedings COLING86.

Dahlgren, K. 1988. Naive Semantics for Natural Language Understanding. Boston: Kluwer Academic Press.

Gelman, R. and E. Spelke. Thoughts about Animate and Inanimate

Objects. in J.H. Flavell and L. Ross, eds., Social Cognitive Development. Cambridge, England: Cambridge University Press.

Hayes, P. J. 1985. The Second Naive Physics Manifesto. In J. R. Hobbs and R. C. Moore, eds., Formal Theories of the Commonsense World, Norwood, NJ: Ablex.

Kamp, H. 1981. A Theory of Truth and Semantic Representation. In J. Groenendijk, Th. Janssen, and M. Stokhof, eds., Formal Methods in the Study of Language, Amsterdam: Mathematisch Centrum, 277-322.

Keil, F. 1989. Concepts, Kinds and Categories. Cambridge, Massachusetts: Harvard University Press.

MacDonald, D. 1982. Understanding Noun Compounds. Carnegie-Mellon

Doctoral Dissertation.

Quinton, A. 1975. Social Objects. Proc. Aristotelian Society. 75:1-27.

Rosch, E., C. B. Mervis, W. D. Gray, D. M. Johnson, and

P. Boyes-Braem. 1976. Basic Objects in Natural Categories. Cognitive Psychology. 8:382-439.

Strawson, P. C. 1953. Individuals. London: Methuen.