Knowledge representation and knowledge of words*

Richmond H. Thomason

Intelligent Systems Program
University of Pittsburgh
Pittsburgh, PA 15260

May 29, 1991

Abstract

This paper surveys some opportunities for cooperative research between linguists and computer scientists in lexical semantics. There are exciting possibilities and challenging problems.

1. Introduction

I will try in this short paper to present some general thoughts on knowledge representation and word meaning for the June, 1991 SIGLEX Workshop on Lexical Semantics and Knowledge Representation. I believe that the topic of this workshop is very timely, and as important and strategic as any in the cognitive sciences. That is the good news. The bad news is that it is very hard to feel confident about this topic, since progress in this area will have to overcome fundamental limitations of several of the sciences that are most closely involved: artificial intelligence, linguistics, and logic. The right emotions should be a combination of excitement and fear, or at least caution.

Difficult problems don't have quick and easy solutions. I don't promise to say anything that will really make a substantive contribution to the research problems. But I will try to explain why I believe the problems are hard and to provide some perspectives on the new area that is emerging here. This paper was written under time pressure. I received the abstracts of the papers that were accepted for the conference only a short time ago. This has made it possible (I hope) to make the paper relevant, but has not allowed much time for scholarship. I hope to prepare an enlarged version of the paper after the workshop, that will try to provide adequate references to the workshop papers, and to the rest of the literature.

2. Goals

*The author acknowledges the support of the National Science Foundation under grant IRI-9003165.

We need a theory of linguistic meaning that is well grounded in linguistic evidence, that is broad in its coverage of linguistic constructions and explanatory power, that can be integrated with appropriate reasoning procedures, and that provides applicable models for technology, such as machine translation, information retrieval, and word-oriented instructional software. How are we going to achieve these goals?

3. Background in logic

My own interest in this topic grew in part out of my work some years ago in Montague grammar. This field has developed into a healthy area of linguistics with many well developed research problems. But fairly early on, it seemed to me that a lot could be learned by concentrating instead on the limitations of the approach; some of these limitations are described in [Thomason 1987]. The shortcomings of a logicist approach to semantics are probably clearest in connection with word meaning.

Knowing such meanings involves access to a broad spectrum of relevant knowledge. Technical terms like 'myelofibrosis', make the point most vividly, but (as Minsky and others have often pointed out), it also is true of everyday terms like 'birthday party'.

A logic-based approach like Montague grammar uses *meaning postulates* to account for inferences like

(1) *Bill saw someone kiss Margaret*
So, someone kissed Margaret

In fact, the underlying logic provides a fairly powerful apparatus for writing these postulates. Lambda abstraction over variables of higher-order types enable the postulate writer to attach conditions to words (in the case of this example, to the word 'see') so that the right intersentential consequences will follow. (Roughly, 'see' has the property of expressing a relation such that if anyone is related to a state of affairs by this relation, then that state of affairs obtains. These things look horrible in English, but fine in Intensional Logic.)

This condition on 'see', though, is far from a characterization of its meaning; it doesn't distinguish it from a large class of similar terms, such as 'hear', 'learn', 'remember' and 'prove'. And the underlying logic doesn't deliver the capability of providing such characterizations, except in a few cases (like 'and') that are closely connected to the original purpose of the logic: explicating mathematical reasoning.

Mathematics provides deep chains of exceptionless reasoning, based on relatively few primitives. Thus, concepts can be connected through definitions. Most common sense domains provide relatively shallow patterns of defeasible reasoning, based on a large number of loosely connected concepts. It is difficult in many cases to separate what is primitive from what is derived. Given enough background knowledge it is possible to characterize the meanings of terms, but these characterizations seldom take the form of necessary and sufficient conditions. It is difficult to find reliable methods for articulating the background knowledge and general ways of applying such knowledge in characterizing meanings.

We should remember that similar inadequacies were responsible for the failure of attempts (most notably, by Rudolph Carnap) to extend Frege's formalization of mathematical reasoning to the empirical sciences.[1] Carnap discovered that Frege's method of deriving definitions failed with color terms, and that terms like 'soluble' could not be given

[1] See [Carnap 36-37].

natural and correct definitions in terms of terms like 'dissolve'. The failure of logic-based methods to provide a means of formalizing the relevant background knowledge even in relatively scientific domains provoked a skeptical reaction against the possibility of extending these logical methods.[2]

Montague motivated his addition of possible worlds to the Fregean framework with a problem in derivational lexical semantics—that of providing a theory of events that would allow predicates like 'red' to be related to their nominalizations, like 'redness'.[3] Trying to account for derivational interconnections between word meanings (rather than providing a framework for making principled distinctions in meaning between arbitrary words) is a more modest goal, and much can be learned by extending a logic-based theory in this direction. But the work in lexical semantics that began in [Dowty 79] seems again to be limited in fundamental ways by the underlying logic. The definition that Dowty provides of agent causality in terms of event causality fails, for logical reasons, in a way that offers little hope of repairs. And, though the idea of *normalcy* that Dowty found to be needed in accounting for progressive aspect seems intuitively to sanction defeasible inferences, Intensional Logic provides no good way of accounting for the validity of examples that have exceptions, like

(2) *Harry is crossing the street.*
So Harry will cross the street.

There is a natural progression between examples like this, which are focused on inferential properties of telic constructions, to cases that draw more broadly on world knowledge (in this case, knowledge about the normal uses of artifacts), like

(3) *Alice used the match to light a fire.*
So Alice struck the match.

4. One relation between lexical semantics and knowledge representation

Linguistic logicist work in the semantics of words, then, is closely related to logicist work in knowledge representation. Though the relation has not been much exploited yet, it suggests a clear line of research that is likely to benefit both linguistics and AI.

I should add that I am thinking of long-term benefits here. I don't claim that this extension of the logicist inventory will provide a representation scheme for words that is nearly adequate. I do believe that such work is an essential part of any satisfactory solution to the problem of lexical representation. There is research in lexical semantics that is oriented towards applications but lacks a theoretical basis. The logical work, on the other hand, is limited in its applicability to lexical problems but provides an interface with sentence meaning; this approach is at its best in showing how meanings of phrases depend on meanings of their parts. Along with this, it provides a specification of correct reasoning that—though it may not be implementable—is general and precise, and can be essential at the design level in knowledge representation applications.

Part of the human language capacity is the ability to deal effectively with both words and sentences. Though we may not have a single computational approach that does both,

[2] See [Quine 60].
[3] See [Montague 69].

we can try to stretch partial approaches towards each other in the hope that together they'll cover what needs to be covered. This is why I am enthusiastic about extensions to the lexical coverage of the logicist approaches.

Logicist work in AI has generally recognized the need for augmenting the Fregean logical framework in order to deal with problems of common sense reasoning. The most generally accepted line of development is the incorporation of nonmonotonicity into the logic. And this feature, it turns out, is precisely what is needed to accommodate many of the problems that emerged in Montague-style lexical semantics. It is the defeasible nature of telicity, for instance, that makes it difficult to deal with (2) adequately within a standard logical framework. It is no surprise that lexical semantics is full of defeasible generalizations, and a general technique for expressing such generalizations would greatly extend the coverage of logicist theories of word meaning.

The available approaches to nonmonotonicity could readily be incorporated into the framework of Montague-style semantics without any changes to the undefeasible part of the logic.[4] Thus, the linguistic side has much to gain from the work in AI.

Work on common sense reasoning, on the other hand, would also gain much from cooperative applications to the study of derived word meanings. For one thing, the project of accounting for such meanings discloses a limited number of notions that are obviously of strategic importance for common sense reasoning.[5]

Moreover, the linguistic work uses a well developed methodology for marshaling evidence and testing theories. Given the difficulty of delineating common sense reasoning and deciding between competing theories, this methodology could be very useful to the AI community.

On the whole, then, this seems like a very natural and promising partnership.

5. Polysemy and context

It is encouraging to be able to point to an area of normal research at the interface of lexical semantics and knowledge representation, but at the same time it would be very misleading to imagine that all the problems of word meaning can be solved by nonmonotonic logic, or that the potential areas of partnership are all tidy and unproblematic.

In a number of papers published over the last twenty years, John McCarthy has claimed that a logical foundation for common sense reasoning should include not only a theory of nonmonotonic reasoning, but a theory of *context*.[6]

It is easy to see how an account of context is central in approaching reasoning tasks of great or even moderate complexity. It is essential to avoid being swamped in irrelevant detail. But if details are ignored, it is also essential to ignore them intelligently, so that the reasoning will retain appropriateness. Engaged reasoning is located in a local context which makes it focused and feasible, but nevertheless retains its applicability to the larger context of which the current context is part.

---

[4] For an example in which Intensional Logic is combined with Circumscription Theory, see [Thomason 90].

[5] At a symposium in the recent knowledge representation meeting in Cambridge, Massachusetts, Ray Reiter argued that common sense reasoning might not need to explicate causality; it may be as unimportant in the common sense world as it seems to be in modern physical theories. The ubiquitous presence of causal notions in processes of word formation is a strong argument against such a position.

[6] The need for a theory of context was mentioned in McCarthy's 1971 Turing Award address; see [McCarthy 87] for a revised version. A recent attack on the problem can be found in [McCarthy 89].

But there is a hierarchy here. Contextualization must also be controlled by reasoning processes, which themselves may well be located in contexts. Thus, contexts can have greater or lesser generality, and some representation of context must also be available to reasoners.

Though—if McCarthy is right—we may not yet have a satisfactory theory of context, which could be incorporated into a logicist framework, we do have many applications. Object oriented approaches to programming, in particular, achieve their power through catering to the human need for contextual organization of reasoning; they could equally well be called *context* oriented approaches.

Many of the most difficult problems of involving the meanings of words have to do with the variability of interpretation. In his experiments on the vagueness of terms, for instance, William Labov noticed that the distinction between 'cup' and 'bowl' was affected more by whether the interpreter was situated in a "coffee" or a "mashed potatoes" context than by factors such as the ratio of height to diameter of the artifact.[7]

To take another example, there is some reason to think that in a context where a bus is leaving for a banquet, 'go' can mean 'go on the bus to the banquet'. Of course, if someone says

(4) *I'm going.*

in such a context, it means 'I'm going on the bus to the banquet', but this effect could be attributed to the speaker meaning of the utterance, without assigning any special interpretation to 'go'. More telling is the fact that in this case it's possible to say

(5) *No, I'm not going; I'm taking my car.*

Some of the problems of polysemy that Sowa discusses in his contribution to this workshop and in other writings are best regarded, I think, as cases in which the procedures for interpreting words are adjusted to context. Unfortunately, this is an area in which we seem to have many alternative ways of accounting for the phenomena: vagueness, ambiguity, strategies of interpreting speaker meaning, and contextual effects. All these accounts are plausible, and each is best equipped to deal with some sorts of examples. But in many cases there is no clear way to pick the best account. Perhaps this problem should be solved not by treating the accounts as competitors and seeking more refined linguistic tests, but by providing bridges between one solution and the other; chunking, for instance, provides in many cases a plausible path from conversational implicature to a lexicalized word sense.

I have stressed the contextual approach to polysemy because it seems to me to offer more hope for progress than other ways of looking at the problem. It enables us to draw on a variety of computational approaches, such as object oriented programming, and it opens possibilities of collaboration with theoreticians who, influenced by McCarthy, are looking for formal ways of modeling contextuality. The ongoing work of theory development badly needs examples and intuitions; language in general and the lexicon in particular are probably the most promising source of these.

6. Linguistic work

[7] See [Labov 73].

Of course, most of the recent linguistic research on word meaning has been done by nonlogicists. See [Levin 85], for instance, for a useful survey of work in the Government-Binding framework.

There is no substitute for the broad empirical work being done by linguists in this area. But as Levin's survey makes clear, it is very difficult to develop a theoretical apparatus that is well grounded in linguistic evidence in this area. Despite the efforts of many well trained linguists to devise good general tests for important notions like agency, the connection of these concepts to the evidence remains very problematic.

Despite difficulties with the high level concepts, the linguistic work has uncovered much taxonomic information that is relatively general across languages, and that evidently classifies words not only into categories that pattern similarly, but that share important semantic features.

This, too, seems to be an area in which cooperation between linguists and the AI community might be fruitful. The classification schemes that come from linguistics are not only well motivated, but should be very useful in organizing lexical information on inheritance principles. Moreover, it might well be useful for linguists who are grappling with methodological difficulties to learn to think of their problems along knowledge engineering lines rather than syntactic ones.

7. Linguistics and knowledge representation

Representation is crucial in contemporary linguistics, and is found in all the areas where linguistic structure is important. But syntax seems to be the primary source of representational ideas and methods for justifying them. For over thirty years, syntacticians have proposed formalisms (which in general are variations on labeled trees, representing phrase structure), along with rules for systematically generating them. They have also developed methods for justifying these formalisms, based mainly on introspective evidence about grammaticality, and an extremely rich battery of techniques for bringing this evidence to bear on hypotheses.

Though (except in some cases where natural language processing systems are integrated with the formalism), these representation systems are tested by introspective evidence, and their connection to experiments and to cognitive psychology is in fact tenuous and problematic, many linguists make cognitive claims for their representations.

The hope seems to be that eventually the structures that are well supported by the introspective methods will be eventually be validated by a larger psychological theory of processing that is well supported by experimental evidence.

Whether or not such a theory is eventually forthcoming, the current methods used to support different syntactic theories often seem to leave no way of settling even quite major issues. And when these methods are extended to semantics, they definitely seem to leave theoretical alternatives underconstrained by the available methodology of linguistic argumentation. Intuitions about meaning are even more problematic than those about grammaticality. Even though grammaticality is a fairly refined notion, and subject to contextual factors that are difficult to determine, it seems to be easier to agree about grammaticality judgments than about, for instance, judgments about ambiguity.

The criteria that have emerged in knowledge representation seem to me to be well worth considering in this respect. Here are some considerations.

1. The criteria are stringent—so stringent, in fact, that, in view of conflict between desirable features such as expressivity and tractability, there really are no general-purpose knowledge representation schemes meeting them all.

2. The criteria of knowledge representation can be added without much violence to the ones already imposed by linguistic theorists. In fact, the need for usability—assuming that the users are linguists—would require the use of representations that make linguistic sense. No special cognitive claims need to be made. The point is that, though it can be debated whether a generally accepted linguistic formalism is adequate as a representation of human cognition, there is no doubt—if it's generally accepted—that it is a useful way of displaying linguists' insights into linguistic structure.

3. It often is necessary in linguistics to represent large amounts of information. As lexicography becomes computerized, and the need is felt to connect these computerized linguistic knowledge bases to areas of linguistic theory such as syntax, a novel criterion emerges—does the theory allow a workable way of organizing large amounts of lexical information?

4. The need to associate knowledge management procedures with representations also provides new constraints, and—if the procedures can be implemented—may also help to automate the testing process. It is hard to see, for instance, whether a semantic theory can be tested as a mere theory of representation. Since the main purpose of semantic representation is to provide a level at which sound inference can take place, an explicit specification of the associated inference procedures is needed before we can begin to test the theory.

5. There are many similarities of detail that make it easy to build smooth bridges between linguistic formalisms and ones from knowledge representation.

8. Conclusion

Let's be clear about the problems.

The field of knowledge representation began with a strong emphasis on applications in natural language understanding, but shifted its emphasis as it developed. This happened in part because the opportunities for productive research in the area are concentrated in relatively small scale, domain specific systems. It is hard to see how to build larger systems without sacrificing a clear understanding of what one is doing, and any hope of reliable performance. Thus, in returning to natural language understanding, we are straining the capabilities of what is known about representing knowledge. Since there is much interest in larger systems, and some hope of help from existing knowledge sources and from what linguists have learned about word meaning, lexical semantics might be a promising area for research in scaling up knowledge representation. But we have to remember that we are trying to extend the field in ways that are pretty fundamental.

Linguists have created a successful science by systematically ignoring cases where there are strong interactions between linguistic knowledge and broadly based world knowledge. They have developed a research methodology that works well for phonology, syntax, morphology, and some limited areas of semantics, but that breaks down in other areas of semantics and in pragmatics. They are comfortable with arguments that test representation systems for linguistic correctness, but not with ones that depend on engineering considerations like usability and transportability. Fairly radical departures from linguistic methodology are needed, I suspect, in establishing a unified theory of lexical semantics. To try to separate this project from the task of building large scale knowledge bases is to settle for a partial solution, which may well turn out to be incompatible with systems providing the world knowledge that ultimately needs to be used in natural language processing applications.

To integrate a computational semantics of words with knowledge representation techniques, we need to remember that representations can't be separated from reasoning. It is all too easy for any representation system to seem adequate until it is put to use in applications such as planning, that call for intensive reasoning. This requirement is probably going to be extremely difficult to observe in practice, but I think that we have to bear it in mind if we are going to have confidence in the representation systems that emerge from this work.

References

[Carnap 36-37] Rudolph Carnap. "Testability and meaning." *Philosophy of Science 3*, 1936, pages 419–471 and *Philosophy of Science 4*, 1937, pp. 1–40.

[Dowty 79] David Dowty. *Word Meaning and Montague Grammar.* D. Reidel, Dordrecht, 1979.

[Labov 73] William Labov. "The boundaries of words and their meanings." In *New ways of analyzing variation in English.* C.-J. Bailey and R. Shuy, eds., Georgetown University Press, Washington DC, 1973, pp. 340—373.

[Levin 85] Beth Levin. "Lexical semantics in review: an introduction" In *Lexical semantics in review,* B. Levin, ed., *Lexicon Project Working Papers 1,* MIT Center for Cognitive Science, Cambridge MA, 1985, pp. 1–62.

[McCarthy 87] John McCarthy. "Generality in artificial intelligence." *Communications of the ACM 30* (1987), pp. 1030–1035.

[McCarthy 89] John McCarthy. "Artificial intelligence and logic." In R. Thomason (ed.) *Philosophical Logic and Artificial Intelligence,* Kluwer Publishing Co., Dordrecht, 1989, pp. 161–190.

[Montague 69] Richard Montague. On the nature of certain philosophical entities. *The Monist 53* (1969), pages 159–194.

[Quine 60] Willard Quine. *Word and Object.* MIT Press and John Wiley, Cambridge MA and London, 1960.

[Sowa 91] John Sowa. "Logical structures in the lexicon." This volume.

[Thomason 1987] Richmond Thomason. "Remarks on linguistic semantics." In *Mathematics of language,* A. Manaster-Ramer, ed., John Benjamins, Amsterdam, 1987, pp. 374–388.

[Thomason 90] Richmond Thomason. Propagating epistemic coordination through mutual defaults I. In Rohit Parikh (ed.) *Proceedings of the Third Conference on Theoretical Aspects of Reasoning about Knowledge.* Morgan Kaufmann, San Mateo CA, 1990, pages 29–39.