

Interpretation without Semantics

Stephen Helmreich

*Computing Research Laboratory
Box 30001/3CRL
New Mexico State University
Las Cruces, NM 88003-0001
shelmrei@nmsu.edu*

Abstract

I suggest that Geoffrey Nunberg's work can make two contributions to the issue of the relationship of lexical semantics to knowledge representation. The first is the claim that Lexical Semantics and Knowledge Representation are the same. The second (which is both more controversial and more difficult to implement) is that the connection between this world knowledge and lexical items must be quite flexible, if not, in fact, non-existent and constructed anew for each interpretive act. I outline several arguments for this position and present a formal method of incorporating this within a standard Montague-grammar framework.

1 Introduction

In this paper, I attempt to relate the NLP issue of the relationship between lexical semantics and knowledge representation to the claims made by Geoffrey Nunberg in *The Pragmatics of Reference* [1]. His basic claim is that interpretation of lexical items in an utterance is almost entirely a pragmatic matter and not one related to language-specific semantics. He uses a theory of referring functions to account for the variety of uses to which any one lexical item may be put.

Though phrased in linguistic terminology, the corollary of this thesis for NLP is that purely linguistic lexical information is quite limited (perhaps only syntactic part of speech and phonological information), with encyclopedic and world-knowledge information constituting most of the knowledge base of the system. In addition, Nunberg claims that the connection between these linguistic symbols and the information in the knowledge base must be quite flexible (and, perhaps unspecified).

The first claim echoes those viewpoints which see no justification for differentiating lexical semantics and world knowledge. It is the second claim that differentiates what might be called a "radical semantics" position from a "radical pragmatics" approach. In both cases, lexical knowledge is treated identically with world knowledge. In the first case, however, this information is regarded as "semantics", knowledge about the meaning of the lexical items. (Langacker [2], Lakoff [3], Martin [4] and Pustejovsky [5] can be seen as typical of this approach.) In the second, it is regarded purely as information about the world and the connection between the *use* of lexical items in language and this world knowledge is the result of pragmatic interpretation.

Since it is this second claim that would require the most radical change in NLP systems, I shall concentrate on this aspect of Nunberg's approach. In the remainder of this position paper I will first restate some of the arguments of Nunberg for such a pragmatic approach to lexical interpretation. Then I shall set out in brief the approach I take in formalizing this approach within a typical Montague Grammar.

2 Argumentation

The first argument points to the many uses of lexical items that cannot be “stored” as part of the semantics of the item, but rather must be inferentially derived as part of the interpretation of the utterance. These include self-referential usages, *ad hoc* usages, and metonymical and metaphorical uses.

Self-referential uses. Example: *Rice starts with an “r”*. Since these types of usages (referring to the phonological, orthographical, phonetic or other characteristics of the lexical item itself) are universal for every lexical item in every language, it is clearly wasteful to store this idiosyncratically as part of the “meaning” of the lexical item.

Ad hoc uses. Example: *Your coat is on the refrigerator door*. In these cases, the interpretation depends on highly context-specific information. In the example sentence, the intent of the speaker was to refer to the dry-cleaner’s claim ticket which was attached to the refrigerator door by a magnet.

Metonymy and metaphor. More complex issues arise with various common “literary” usages such as metaphor or metonymy. There are claims on both sides for idiosyncratic lexicalization and also for free generation (Martin [6], Fass [7], and Pustejovsky [5]).

Taken together these types of usage are only suggestive. They require an extension, an analytic leap: if these types of usage are to be handled pragmatically, perhaps all types of usage should be so handled.

At this point, it should be pointed out that simply because particular information is (or is most conveniently) stored rather than derived is not an argument that this information is semantic and not pragmatic. As long as it *can* be so derived, it can be regarded as stored pragmatic information.

A second argument in favor of treating the relationship between lexical items and world knowledge as pragmatic rather than semantic is an Ockham’s razor argument: If one can exhibit a pragmatic system that performs lexical interpretation as well as an idiosyncratic language-specific semantic system can, then it would be best to opt for the pragmatic system. This is because the sub-systems needed by the pragmatic system (world-knowledge, context knowledge, and inferencing, at a minimum) are all types of information and processing that are needed for other independent and non-linguistic activities. Why duplicate them with a specific linguistic-semantic system? Note that this argument in particular requires a demonstration of the existence and workability of such a system.

Yet a third argument of Nunberg that I feel has not been sufficiently understood and internalized is that of the idealized speaker/hearer of the Chomsky model. Nunberg claims that this model is sufficient for phonology and syntax, but fails for semantics. Clearly at an individual level, the experiences on which world-knowledge is based varies widely from one to another, and in addition, (even given a relative commonality of experience) the way in which this knowledge is organized can be expected to vary widely from one person to another. Therefore, there seems to be no reason for lexical knowledge to be in any way identical or identically organized from person to person.

On yet another level, Nunberg claims that the way in which this lexical knowledge is utilized requires varying ways of looking at this knowledge. Thus the meaning that one attributes to a word in a particular context varies with the type of interaction, the social grouping in which one finds oneself, and the particular assumptions one makes about the interlocutors. Examples of this kind of varying usage are such terms as “classical music” or “jazz”. In one context, the term “classical music” might include Broadway musicals,

while in another it may refer only to the music of Mozart and early Beethoven.

While speakers of the same language use the same lexical items and the same syntactic structures, their internal representation of world knowledge and the way they connect these representational structures with linguistic items may vary greatly from person to person and from social group to social group. Speakers of a language must take this into account in trying to communicate. They cannot assume that all other speakers have the same shared knowledge, organized in the same way, and connected identically to the same lexical items.

The implications of this for NLP is that any language processing system should be able to work equally well with world-knowledge representations that may vary significantly in format, structure, and contents. If lexical items are strictly tied to elements in a particular knowledge base, this is not possible.

3 Implementation

It remains, however, to demonstrate that such a system is feasible. In the rest of this position paper, I shall outline the manner in which I have attempted this. Formalizations of Nunberg's referring functions have been made by Sag [8] (on a small scale) and by Fauconnier [9] (in a more cognitive context).

I have taken a typical Montague Grammar (truth-conditional/model theoretic) approach (Dowty, Wall, and Peters [10]) and eliminated the function which provides interpretations for each lexical item in the language. Instead, this is replaced by an Interpretive Component (which can be viewed in the large as related to Kamp's Discourse Representation Theory [11]).

The only requirement on the model is that it contain elements that map onto the lexical items of the language (thus allowing for self-reference). In fact, the only fixed "semantics" of the system associates lexical items in the language with these representative elements in the model.

In the Interpretive Component are Montague Grammar operators which function like Nunberg's referring functions. That is, they operate on items (and sets of items, and intensions of sets of items) and return other items, sets of items, and intensions of sets of items. These referring functions (MG operators) also contain usage/context and speaker indicators.

Certain referring functions operate on the lexical items in the model and return sets of things that particular speakers (in particular contexts) often refer to by using that particular lexical item. This is the closest to a "meaning" that the system contains. However, it is not part of the language, may vary greatly from speaker to speaker of the language, and is subject to certain rules of application before an interpretation can be generated.

These particular referring functions are not directly inserted into the sentence to provide its interpretation. Rather, it is a matter of co-determining a context and an interpretation. Part of the context is already given, while part is abductively added by the referring functions themselves.

I can show that this type of approach is useful for several purposes. First of all, it provides an existence proof for a semantic system that is both compositional and has no lexical semantics, only a lexical pragmatics—a system in which lexical semantics is highly constrained and the connection between the lexical items and the encyclopedic knowledge

is flexible. Second the concepts involved can be used to usefully distinguish between metonymy and metaphor (by determining whether or not the referring functions used have 1-1 or 1-many extensional equivalents). It also allows a definition of “same lexical item,” that is, two senses are the same lexical item if they are reachable through the set of given referring functions.

Although not an implementation of this approach, I would like to conclude by making a brief reference to the work that Eric Iverson and I have been doing on meta5/metallel—a program of word sense disambiguation that also identifies instances of metaphor and metonymy (Helmreich et al [12]). Although our research has not been directly reflective of the concerns of lexical knowledge versus world knowledge, we are working with a program which treats both types of knowledge as equivalent. That is, the structure of the lexicon contains both encyclopedic and what would be called strictly lexical information. In addition, metaphorical and metonymical usages are not directly stored in the lexicon, but are derived inferentially on the basis of the linguistic and world-knowledge context.

References

- [1] G. Nunberg, *The Pragmatics of Reference*, Indiana University Linguistics Club, Bloomington IN, 1978.
- [2] R. Langacker, *Foundations of Cognitive Grammar*, Stanford University Press, Stanford CA, 1987.
- [3] G. Lakoff, *Women, fire, and dangerous things*, University of Chicago Press, Chicago IL, 1987, University of Chicago Press.
- [4] J.H. Martin, “A unified approach to conventional non-literal language”, *Proceedings of the Fifth Rocky Mountain Conference on Artificial Intelligence (RMCAI-90)*, Las Cruces, NM, 1990, pp. 5–40.
- [5] J. Pustejovsky, “The semantic representation of lexical knowledge” *Proceedings of the First International Lexical Acquisition Workshop*, 1989.
- [6] J.H. Martin, *A Computational Model of Metaphor Interpretation*. Academic Press, Cambridge MA, 1990.
- [7] D. Fass, *Met*: A Method for Discriminating Metonymy and Metaphor by Computer*. CSS/LCCR TR 89-15, Centre for Systems Science: Burnaby, B.C., 1989.
- [8] I. Sag, “Formal Semantics and Extralinguistic Context”, *Radical Pragmatics*, P. Cole, ed., Academic Press, New York NY, 1981.
- [9] G. Fauconnier, *Mental Spaces: Aspects of Meaning Construction in Natural Language*. MIT Press, Cambridge MA, 1985.
- [10] D. Dowty, R. Wall, S. Peters, *Introduction to Montague Semantics*, Synthese Language Library 11, D. Reidel, Dordrecht Holland, 1981.
- [11] H. Kamp, “A Theory of Truth and Semantic Representation”, J.A.G. Groenendijk, T.M.V. Jansen, and M.B.J. Stokhof, eds., *Formal Methods in the Study of Language*, Mathematisch Centrum, Amsterdam, 1981.
- [12] S. Helmreich, E. Iverson, and F. Laroche, “Modular Meta5: Further Research in Collative Semantics”, *Memoranda in Computer and Cognitive Science*, MCCS-90-192, Computing Research Laboratory, New Mexico State University, Las Cruces NM, 1990.