

Book Review
for
Canadian Philosophical Reviews

Francis Jeffry Pelletier
Department of Computing Science
Department of Philosophy
University of Alberta
Edmonton, Alberta
Canada T6G 2E1

Kees van Deemter and **Stanley Peters** (eds.) *Semantic Ambiguity and Underspecification*.
(CSLI Publications: Stanford, CA). 272+xxiii.

This is a book of articles about a new theoretical underpinning for computational linguistics. Despite this aim, it contains much that is of interest to philosophers of mind, epistemologists, and philosophers of language, regardless of whether they also have an interest in computational implementations.

Computational linguistics has many subareas of research, such as voice recognition and machine translation and natural language interfaces to databases—to pick just three more or less at random. The area of computational linguistics that is under discussion in this book concerns the attempt to have a computer understand utterances in a natural language to the same extent that a person would understand them, starting with written (not spoken) natural language text, performing whatever syntactic and semantic measures are called for, and ending up with the ability to answer simple questions and draw some conclusions concerning the information contained in the input. This part of computational linguistics, which might be called natural language understanding (NLU), has generally followed a *representationalist* methodology. A sentence is presented as an ordered sequence of words; each word is looked up in a *lexicon* to find out its part of speech (noun, verb, adjective, etc.); these categorized words are then sent through a (syntactic) *parser* which produces a description of how the sequence of words is constructed into a sentence (which of the words form the noun phrase subject, which constitute the verb phrase, etc....the details depend upon what theoretical categories the syntactic grammar

presumes); this phrase structure description forms the input to a semantic module. This module uses the meanings of the individual words (as they are stored in the lexicon) plus information about how the words were combined by the parser to form longer phrases, to generate an overall interpretation of the sentence. There might then follow a process of generating some implications of this interpretation, especially for use in question-answering systems and for aid in the interpretation of any sentences that might follow the one just interpreted. All along through this procedure of generating a final interpretation out of a string of words there is a representation of the information as thus far processed. The lexicon contains a representation of the words of the natural language, including not only its part of speech but also its meaning. The result of the parsing procedure is to generate a syntactic representation of the string of words. And the result of the semantic interpretation is produce some representation of the sentence's meaning. A major trend in NLU has been the investigation of the formal powers of different types of representations. But we need not immerse ourselves in these details in order to appreciate the topic of this book; for, although the various authors adopt a variety of differing outlooks on many theoretical NLU issues, the underlying subject is quite universal and easy to understand.

NLU has had a rocky history, full of fits and starts, optimism and reversal of fortune. But one of the unexpected discoveries to emerge from the attempt to get computers "to talk just like you and me" is that psycholinguists and philosophers of mind/language had dramatically misunderstood humans' understanding of ambiguities. This was due mostly to a lack of appreciation of how much ambiguity "really" was in natural language. In formal accounts of meaning, even the most innocuous-seeming sentence can be wildly ambiguous. And that this is so was brought home very forcefully when the early computer programs were constructed and showed all the possible interpretations of the sentences. Consider *I do not like green eggs and ham*. Besides all the different meanings of the individual words (the *Webster's Student Dictionary* lists <5, 21, 1, 17, 16, 8, 4, 6> meanings for the words in the order given in the sentence), there is a question of whether I *dislike* them ("like them not") or merely *do not like*

them (maybe I am neutral), there is a question of whether one or the other of these attitudes applies to all ham or only to green ham, there is even the question of whether I have one of these attitudes towards green eggs, or only towards mixed ham and eggs that are green. *Three girls mailed four letters* is commonly thought to be ambiguous between there being four and twelve letters mailed; such an ambiguity is called a "scope ambiguity" (as is also the earlier *do not like them* vs. *like them not* ambiguity). The sentence *In most democratic countries most politicians can fool most of the people on almost every issue most of the time* has 120 different meanings for its different scope orderings and each of these meanings is true in distinct circumstances. Yet people (think they) can understand these sentences when they are uttered, and they do not seem to pause over the wide range of different interpretations possible.

Although the phenomenon of ambiguity has been recognized for many years in the study of language and language processing, the sheer amount of it and its presence in even very simple sentences has forced a reappraisal by workers in the cognitive sciences of their views as to how people process sentences. (Bar-Hillel was possibly the first to draw attention to this, in the late 1950s. But it wasn't until the late 1980s that it became fully appreciated). The discovery by NLU researchers that (psychologically) simple sentences were (formally) staggeringly ambiguous has forced others to give up these three views, each one of which described an attitude that researchers had entertained:

1. People construct mental representations for all of the meanings of a linguistic item and then use contextual information (both of the previous linguistic discourse and of the situation of utterance) to pare these down to some small number (perhaps one, in the normal case).
2. People work with a "most likely" meaning of the initial part of a sentence as it is being parsed, and then incrementally build on it by taking a "most likely" meaning of the next part and composing it in the "most likely" manner with the earlier part, until the entire sentence is represented with a "most likely" meaning.
3. The preceding linguistic context and the situational context conspire to make a particular mental representation of the sentence be determined. Pretty much any sentence that is

uttered that contains the same words, and pretty much in any order, will give rise to this unique representation, in that context.

The difficulties computational linguistics programs encountered in being overwhelmed by ambiguity soon showed that people would have the same problems if sentence comprehension took place by means of representationism using any of the above three mechanisms. People simply do not maintain huge databanks in their heads consisting of all the ambiguous meanings, nor has anyone come up with a plausible mechanism by which people can compute the “most likely” meanings, nor is it true that people proceed by having “scripts” that merely need to have some keywords be uttered to fill in the blanks. In fact, people seem mostly to be *unaware* that there is range of possible meanings, and they go about their business only extracting what is called for from the information conveyed in the sentence.

This calls for a radical rethinking of the way people process semantic information, and has led to a new direction in NLU. The book under review presents works that display the new attitude that is emerging among a wide variety of computational linguists concerning this problem of overwhelming ambiguity. The main idea has come to be called *ambiguous logical form* or *underspecified semantic representation*, and is now a hotbed of active research in computational linguistics. The idea is to avoid the need for separate storage of all possible interpretations by instead storing only one “ambiguous” representation. The editors trace the interest in this sort of solution for the problem of ambiguous lexical items to Bronnenberg *et al* (1979) in the Philips NLU system and to Schubert & Pelletier (1982) where the idea was applied to scope ambiguities of quantifiers and other items. (See the book Introduction for full references). In these approaches only one ambiguous representation would be stored and disambiguation would happen at some later point in the processing of a sentence when more contextual information was available. But, continue the editors, these ideas did not really catch on in mainstream NLU until the late 80s when researchers at SRI came up with their underspecified “quasi logical forms”.

The move to ambiguous logical forms and underspecified semantic representations raises a number of conceptual questions, as well as practical issues of how to process this sort of representation. Here are some of the topics that are addressed by authors in this collection:

- How are different interpretations/meanings to be individuated? The fact that a given sentence can be used to truthfully describe two different situations does not necessarily make for ambiguity. It could instead be generality of a single meaning. How can the two be separated?
- What sorts of processing should be done with the underspecified representations? In particular, can they enter into inferences?
- When is disambiguation necessary?
- Is ambiguity a significant notion? After all, if an ambiguous expression can be processed in much the same way as a nonambiguous one, and used in the same mental activities, then what does the difference between them amount to?

The 10 papers in this collection, all previously-unpublished, deal with the above-mentioned questions, as well as with certain technical issues of formal semantics. The paper by Georgia Green concludes that disambiguation cannot in general be accomplished, because reasoning from pragmatic information can be arbitrarily complex. Philosophers will find this paper in particular to be a very rich source of examples and puzzles for “ordinary” philosophy of language. A paper by Jaap van der Does & Henk Verkuyl deals with the so-called distributive/ cumulative distinction (*The philosophers wrote the books*), attempting to show that such sentences require only one representation (which can be interpreted in two modes), despite the possibility of intermediate readings. Jerry Hobbs also addresses the question of how to individuate interpretations, arguing that most of the information conveyed by a sentence is not due to there being different semantic interpretations but rather is due to some “abduction-driven” method of interpretation of a single semantic interpretation. The papers by Hideyuki Nakashima & Yasunari Harada, by Sasa Buvac, and by Anne-Marie Mineur & Paul Buitelaar, are about the

issues of how to represent ambiguous information and how to disambiguate it. The first of the three is presented within the framework of Situation Semantics; the second makes use of the notion of “discourse context” and investigates how information can be shared by different contexts; and the third concerns the extent to which the computer science notion of “type coercion” can be employed in disambiguating lexical items. Hiyan Alshawi and Massimo Poesio each present papers that interpret ambiguous expressions by means of an “underspecified” representation, and then semantically interpret these representations as sets of truth values. Poesio’s paper will be of general interest to philosophers of language because it presents a survey and critical review of much of the literature on ambiguity, and makes many interesting comments about the psychological literature on the topic. Poesio also touches on the issue of how to reason using the underspecified representations, but the topic is taken up in much more detail in Kees van Deemter’s paper and in Uwe Reyle’s paper. The former paper proposes a many-sorted logic containing ambiguous constants and explores the different notions of logical consequence that can be defined in such a language. Reyle’s paper presupposes the formalism of Discourse Representation Theory but essentially amounts to treating the premises of an argument “universally” – that is, universally quantifying over all the many interpretations of the premises – and treating the conclusion of the argument “existentially”. An argument is valid in this framework when every interpretation of the premises (classically) entails some interpretation of the conclusion.

Much of the work in this book is, or could be, of interest to philosophers who are not already engaged in the enterprise of “getting computers to talk like you and me”. For one thing, there is a lot of information here relevant just to the question of what ambiguity is, and for another thing there are many provocative claims made concerning whether disambiguation is a good thing and how it might be done. Researchers interested in the distinction between “meaning vs. world knowledge” will find much to think about in Green’s article especially. And finally, logicians and formal semanticists could profitably turn some of their attention to the question of what reasoning would or should be in a language of ambiguous logical forms.

