

Book Reviews

Semantic Ambiguity and Underspecification

Kees van Deemter and Stanley Peters (editors)
(Institute for Perception Research and Stanford University)

Stanford, CA: Center for the Study of
Language and Information (CSLI
Lecture Notes 55), distributed by
Cambridge University Press, 1996,
xxiii+272 pp; hardbound, ISBN
1-57586-029-5, no price listed;
paperbound, ISBN 1-57586-028-7, \$24.95

Reviewed by
Peter J. Ludlow
State University of New York at Stony Brook

Consider the following two ways of thinking about ambiguity. According to the first way, let's call it the **orthodox view**, the ambiguity of a sentence is a function of there being more than one possible representation of the form of that sentence. So, on the orthodox view, an utterance of *Every man loves some woman* admits two distinct logical forms. For example, these:

- (1) [Every x : man(x)] [Some y : woman(y)] [x loves y]
- (2) [Some y : woman(y)] [Every x : man(x)] [x loves y]

Similarly, according to the orthodox view, the ambiguity of *John went to the bank* is due to there being two distinct lexical items for *bank*, and therefore to there being two different forms possible for the sentence *John went to the bank*, depending on which lexical entry for *bank* is inserted.

On the alternative way of thinking about ambiguity, let's call it the **heretical view**, matters are handled differently. Ambiguities may not depend on there being a plurality of possible forms so much as on the *underspecification* of form. If I say *John went to the bank*, I do not have two distinct lexical entries for the noun *bank*, but a single lexical entry which is underspecified as to whether a place of finance or a fluvial embankment is intended. A similar point can be made regarding sentential structure. Maybe there is only one syntactic representation for *All men love some woman*, yet different interpretations of that structure are possible in different environments. Or perhaps we can be radically heretical (indeed **apostate**) and say that the ambiguity can simply remain unresolved.

What turns on this distinction? There are many possible interpretations for the sentences we encounter on a regular basis, yet it is rare that we freeze up and ask "I wonder what she meant by that" (newspaper headlines are an obvious exception here). How is it that we are so good at disambiguating actual utterances or inscriptions given all the possible interpretations? How can we make natural language processing systems good at it too?

The orthodox view admits two ways of approaching these questions. One way (let's call it the **expensive** way) is to first generate all the possible forms that an utterance or inscription might have, and then determine which the speaker or writer intended. Another way (let's call it the **frugal** way) holds that (when things go right) we construct only one representation out of the available possibilities, using context (including our knowledge of speaker intentions) as a guide.

The heretical view suggests an altogether different approach to these questions. We bring a single underspecified representation to many different contexts, and let the context determine the intended meaning of that single representation. Or, taking the apostate view, perhaps the ambiguity is never resolved. Perhaps we routinely communicate and comprehend expressions that have underspecified meanings. That is, maybe we have thoughts that are ambiguous, and we communicate and reason with those ambiguous thoughts.

This excellent collection contains essays that explore the frugal orthodox view, the heretical view, and the apostate view of ambiguity, with a definite tilt towards the heretical and apostate theories. The collection has its origins in a workshop on ambiguity at Stanford in 1993, but the editors made the very sensible move of making a general call for papers on the topic. The result is a very strong collection, well focused, with contributions from a broad range of researchers.

The contributions are as follows:

"Ambiguity resolution and discourse interpretation" by Georgia Green

"Quantification and predication" by Jaap van der Does and Henk Verkuyl

"Monotone decreasing quantifiers in a scope-free logical form" by Jerry Hobbs

"Situated disambiguation with properly specified representation" by Hideyuki Nakashima and Yasunari Harada

"Resolving lexical ambiguity using a formal theory of context" by Saša Buvač

"A compositional treatment of polysemous arguments in categorial grammar" by Anne-Marie Mineur and Paul Buitelaar

"Underspecified first order logics" by Hiyan Alshawi

"Semantic ambiguity and perceived ambiguity" by Massimo Poesio

"Towards a logic of ambiguous expressions" by Kees van Deemter

"Co-indexed labeled DRSS to represent and reason with ambiguities" by Uwe Reyle.

There is an excellent introduction by van Deemter and Peters, which lays out the issues very nicely and provides a useful guide to the papers.¹ Along the way, there is some brief historical context-setting for the project, and the introduction concludes

¹ I note, however, that the handling of the references in the introduction is a bit careless. The volume and page numbers given for Schubert and Pelletier (1982) are completely wrong, and the reference for Bronnenberg et al. has both the date of publication and the title of the book wrong (actually, the series title is mistaken for the book title). In the latter case, this led to great difficulty in my tracking down the article.

with a brief survey of some of the approaches not covered in the collection (for example, statistical methods of disambiguation).

With respect to the context-setting remarks, it is possible to complain that the editors underplay the extent to which the heretical view has been percolating in the literature for some time. They point to Bronnenberg et al. (1980) for an early underspecification treatment of lexical items, and to Schubert and Pelletier (1982) for a treatment of scope ambiguities, but claim that these approaches “did not gain great popularity until the late eighties” when **quasi-logical forms** were developed at SRI International (e.g., Alshawi 1990). Of course, whatever their popularity, such approaches were not uncommon throughout the 1980s (for example, one thinks of Hobbs and Shieber [1987], which appeared in this journal). If one looks to related fields, there has also been widespread flirtation with underspecification strategies.

In the linguistics literature, for example, Kempson and Cormack (1981) had proposed a single, surface-scope-preserving representation for multiply quantified sentences, and argued that pragmatics would yield the alternative scope readings. In a certain sense, the quantifier-storage mechanism of Cooper (1983) was in the same spirit. Even May (1985) ultimately opts for an underspecification approach by arguing that in *Every man loves some woman*, each quantified NP *c*-commands the other, and hence neither is in the scope of the other. For the resulting structure, scope is interpreted “arbitrarily.” These linguistic papers thus prefigured a number of essays in the volume (e.g., those by Hobbs and by van der Does and Verkuyl), which pursue more sophisticated forms of the same general strategy.

Stepping back a bit, it appears that these heresies have been whispered in many places and for some time now.² Of course it is one thing to commit heresy, and yet another to work out the details of an apparently heretical program, and that is where this volume departs from previous work (and leaves it far behind). I won’t survey each of the articles (the introduction by van Deemter and Peters does a fine job of that), but I do think that it might be useful to discuss some of the challenges that are

² Even philosophers have got into the act. See Atlas (1989) for an example.

We all tend to overstate our breaks with extant research (I’m as guilty of this as anyone), but some of the essays in this collection go too far in characterizing previous research as old guard. One example of this is the essay by Green, which reconstructs Hirst’s (1987) survey of the possible positions as follows (p. 3):

- (i) Mental representations (or neural structures) activated by linguistic processing of prior linguistic context determine a unique interpretation.
- (ii) All interpretations are accessed before context influences the determination of the meaning in that context.
- (iii) Interpretations are accessed in order of their “frequency of usage.”

Claiming that her view is “distinct from all three of these” (ibid.), Green offers the following (ibid.):

- (G) CONTEXT-ORDERED SEARCH: Interpretation of prior context biases the interpretive process before the ambiguity is encountered. The search for senses of an ambiguous word or phrase is determined by what the interpreter anticipates the speaker would have intended.

I have to confess that I find it difficult to distinguish (G) from (i). And (i) is just Green’s gloss of what Hirst said. If we compare (G) with what Hirst *actually* wrote (he was quoting Swinney and Hakes), my powers of discrimination fail me completely. From Hirst (1987):

- (i*) The PRIOR CHOICE hypothesis: “Prior context biases the interpretive process before the ambiguity itself is encountered . . . perhaps allowing only a single reading ever to be accessed” (Swinney and Hakes 1976, 683).

In fact, although not directly concerned with psycholinguistic issues, Hirst clearly says that the evidence favors the prior choice hypothesis (i*) in the case of structural ambiguity (if not lexical ambiguity).

taken on by the contributions in this volume (speaking to individual essays along the way).

One problem that many of the theories face is the heavy role that pragmatics must play in either the construction of the appropriate representation (on the frugal orthodox view) or the deduction of the intended meaning for the representation (on the heretical view). For example, in Hobbs's paper, we are told that the representation for a sentence such as *Few men work* will be a representation that literally says that there is a set containing few men, all of which must work. Pragmatics is supposed to allow us to deduce that this set contains all the men who work. Of course, that is a rather sophisticated pragmatic inference by Gricean lights. For example, Neale (1991, ch. 3) spells out the Gricean principles by which we can move from *some of the deaths were accidental* to the understanding that *some of the deaths were not accidental*. The inference is not trivial. If Hobbs's inference is to be carried out by abduction (as he proposes) then we shall want to see the details.

Indeed, the computational expense of using pragmatics to derive an intended meaning may well dwarf whatever expense there is in creating data structures for each possible scope representation, even if those possibilities number in the thousands. In effect, the "frugality" may amount to saving pennies on a billion-dollar budget. Still more, we may find in the end that having those representations handy will prove useful to any comprehensive effort to implement the pragmatic component (who can say, at this early point?).

Likewise for the heretical view (as the contribution from Poesio observes), there is no advantage to having a single underspecified representation if there is no constraint on the possible senses (interpretations) that must be computed. This point gives rise to an interesting conceptual question. Exactly what does it mean to say that one has an underspecified representation of an ambiguous sentence, if, in the end, the knowledge base (KB) has to represent the two interpretations differently? That is, let's suppose we have a natural language front-end to an intelligent knowledge base of some form. One can say that the natural language component admits only one representation, but, if the KB component still admits two, it seems like a rather hollow victory.

Taking a leaf from the paper by Buvač, one might be able to keep an underspecified representation in the KB, and then disambiguate the representation when necessary by using "lifting rules" that draw on representations from a more general commonsense database.³ Of course, once we have disambiguated the KB representation, wouldn't we want to keep that disambiguated representation close at hand? Particularly if we might have to call on that information again? In short, we might save computational expense by only resolving ambiguities when we have to, but if we have to keep resolving the same ambiguities over and over, it is hard to see what is gained. In addition, as Buvač notes, we might not realize that the information we are giving to our database is ambiguous and that the ambiguity cannot be resolved even with the help of lifting rules. In that case, we might wish an intelligent database to be able to point out an unresolvable ambiguity and to query us for some disambiguating information. But if that is the case, then all ambiguities must be resolved when they enter the database. Why throw away that information?⁴

³ I'm not saying that Buvač is proposing to avoid disambiguated representations. I'm saying that his strategy might be applied in such an enterprise. For a similar strategy using "incomplete encodings," see Hirst and Ryan (1992).

⁴ The contribution by Mineur and Buitelaar allows underspecified representations to be disambiguated by relatively local information (thus minimizing the expense of recomputing the intended meaning), but, as van Deemter and Peters observe, this locality requirement prevents the proposal from being

Alternatively, following a path similar to that taken in the paper by Nakashima and Harada, one might try to avoid multiple KB representations by adopting some form of “situated inference” hypothesis (à la Barwise 1989), and hold that environmental conditions of some form will fix the content of the underspecified representations. That is, if the inference engine is sensitive to its external environment, the proper content for the underspecified representation might come for free. This might or might not work in the case of *bank* (where there are two different kinds of things in the environment), but it seems implausible when we encounter ambiguous expressions for abstract objects.

I have no quarrel with claims that we can be in meaningful epistemic relations with abstracta. The problem is that the only theories we have for cashing out *how* these epistemic relations are possible rely upon the mediation of syntactic forms (see, for example, Parsons [1990]). It is by our having finely individuated syntactic representations that we are able to “perceive” or be “attuned to” these abstracta. But if that is right, we will need fully specified representations *before* we can help ourselves to talk of abstract objects “in the environment.” And if that is the case, we cannot unburden our inference engine of multiple meaning representations.

This leads us to the apostate theories. These theories give up on the idea of disambiguation in the KB (situated or otherwise), and try to make sense of automated reasoning with genuinely ambiguous representations! Fortunately, several essays in the collection (e.g., the contributions by van Deemter, by Reyle, and by Alshawi) take up this fascinating idea.

Just to see how radical such an idea is, we need to bear in mind that it was concerns about rooting out fallacious reasoning that led the Stoic philosophers to invent the orthodox theory of ambiguity in the first place (see Atherton [1993]). The challenge for the Stoics (and for philosophers ever since) has been to ensure that we avoid the fallacy of equivocation—that is, reasoning fallaciously by trading on separate meanings for apparently identical terms in our arguments. For example, the following argument equivocates on two meanings of *end*:

Universal happiness is the end of ethics.
Everyone is happy.
Therefore it’s the end of ethics.

One wonders whether reasoning with underspecified representations will not have to smuggle in some form of the orthodox theory to be successful. That’s sort of how it looks in the gloss that van Deemter and Peters give to the paper by van Deemter (p. xix):

[I]n the inference $F(a) \models F(a)$, coherence requires both occurrences of the predicate F to have the same interpretation even if F is ambiguous, and this is, arguably, what keeps the inference valid in an ambiguous setting.

If some sort of bookkeeping is involved in ensuring that one is not shifting interpretations for F in a derivation, why can’t we think of the underspecified representations together with these bookkeeping mechanisms as just another way of talking about fully specified representations?

extended to handling a broad range of ambiguity phenomena (and clearly it isn’t intended to be a general solution to the problem of ambiguity).

Another concern about the apostate view is more philosophical. While we routinely say things that are ambiguous, it is always open to someone to come back and say “what did you intend by that: *A* or *A'*?” Perhaps this is the Fregean in me, but I can’t imagine saying that I was intending to express an ambiguous thought, for the simple reason that I don’t know what an ambiguous thought would be. To be sure, I might intentionally say something ambiguous, such as *I’m going to the bank*, to communicate that I am going out. In such a case, I might not care how *bank* is to be taken. But this is simply a case of my using an ambiguous expression to convey an unambiguous thought (“I’m going out”). There is no question that utterances can be willfully used in an ambiguous way, but it does not follow that we are communicating ambiguous thoughts.

Sometimes, when someone asks me what I meant, I did not see that my remarks were ambiguous. Perhaps I did not realize that one of my terms had a second meaning. In these cases we have *unintentionally* used an ambiguous expression to convey an unambiguous thought. Still other times, I may use an ambiguous term having neither of the actual meanings in mind. Rather I might have some “smeared” understanding of the two. In this case I am perhaps having a disjunctive thought involving both *A* and *A'*, or perhaps I am having a third thought involving the smeared *A + A'* concept. But what would it mean to say that my thought itself was ambiguous?

This complaint is philosophical, but it mirrors a dilemma about applying the apostate theory to KBs. If a KB can resolve the ambiguity of a representation by appeal to the environment or a broader commonsense database, then it doesn’t appear that the representation is really underspecified after all. (This is clearly what van Deemter (p. 203) is committed to in view of his remarks on Barwise and Perry 1983.) On the other hand, if there is no way to disambiguate the representation, then that representation will support inferences involving both interpretations *A* and *A'* so long as they are kept separate. But this appears to be functionally equivalent to reasoning with a disjunctive representation for *A* and *A'*. Is there a third possibility? I’m not clear on what it would be.

All of these possible (and undeveloped) criticisms are not intended to put potential readers off from this excellent collection. To the contrary, I hope that (by marking out the many controversies spawned by these papers) I have piqued interest in the project. Without a doubt, this collection is required reading for all computational linguists, semanticists, and technically competent philosophers of language. All of the above should hurry out, buy this book (or borrow it), and read it; I can’t wait for the debates to begin!

Acknowledgment

Thanks to Graeme Hirst for comments on earlier drafts of this review, and for help in tracking down some of the (botched) references in van Deemter and Peters’s bibliography.

References

Alshawi, Hiyan. 1990. Resolving quasi logical forms. *Computational Linguistics*, 16(3): 133–144.
 Atherton, Catherine. 1993. *The Stoics on Ambiguity*. Cambridge University Press, Cambridge, England.
 Atlas, Jay D. 1989. *Philosophy without*

Ambiguity. Oxford University Press, Oxford.

Barwise, Jon. 1989. *The Situation in Logic*. Stanford: Center for the Study of Language and Information.

Barwise, Jon and John Perry. 1983. *Situations and Attitudes*. The MIT Press, Cambridge, MA.

Bronnenberg, W.J.H.J., Harry C. Bunt, S.P. Jan Landsbergen, Remko J.H. Scha, W.J. Schoenmakers, and E.P.C. van Utteren. 1980. The question answering system PHLIQA1. In Leonard Bolc (editor), *Natural Language Question Answering Systems*. Carl Hanser Verlag, Munich, pages 217–305.

- Cooper, Robin. 1983. *Quantification and Syntactic Theory*. D. Reidel, Dordrecht.
- Hirst, Graeme. 1987. *Semantic Interpretation and the Resolution of Ambiguity*. Cambridge University Press, Cambridge, England.
- Hirst, Graeme and Mark Ryan. 1992. Mixed-depth representations for natural language text. In Paul S. Jacobs (editor) *Text-Based Intelligent Systems*. Lawrence Erlbaum Associates, Hillsdale, NJ, pages 59–82.
- Hobbs, Jerry and Stuart Shieber. 1987. An algorithm for generating quantifier scopings. *Computational Linguistics* 13(1): 47–63.
- Kempson, Ruth and Annabel Cormack. 1981. Ambiguity and quantification. *Linguistics and Philosophy* 4, 259–309.
- May, Robert. 1985. *Logical Form: Its Structure and Derivation*. The MIT Press, Cambridge, MA.
- Neale, Stephen. 1991. *Descriptions*. The MIT Press, Cambridge, MA.
- Parsons, Charles 1990. Mathematical intuition. In Palle Yourgrau (editor), *Demonstratives*. Oxford University Press, Oxford.
- Schubert, Lenhart K. and Francis Jeffrey Pelletier. 1982. From English to logic: Context-free computation of 'conventional' logical translations. *American Journal of Computational Linguistics*, 8(1): 27–44.
- Swinney, David A. and David T. Hakes. 1976. Effects of prior context upon lexical access during sentence comprehension. *Journal of Verbal Learning and Verbal Behavior*, 15(6): 681–689.

Peter J. Ludlow is Associate Professor of Philosophy and a member of the Stony Brook Group in Logic and Formal Semantics. He is editor of *High Noon on the Electronic Frontier: Conceptual Issues in Cyberspace* (The MIT Press, 1996) and *Readings in the Philosophy of Language* (The MIT Press, 1997). Ludlow's address is Department of Philosophy, SUNY at Stony Brook, Stony Brook, NY 11794; e-mail: pludlow@ic.sunysb.edu