

A Situated Perspective on Natural-Language Processing

Susan U. Stucky
Stanford University

Use the word 'situated' these days and all anyone can think of is situation semantics.¹ That's unfortunate from my point of view, because I would like to explore an idea that follows from ideas handily referred to as "a situated perspective." (Not surprisingly, both situation semantics and its companion logic situation theory are by and large consistent with this point of view, but they aren't the only way of working out a theory compatible with it.) Viewing computation, language, and inference through this perspective, I will maintain, suggests a conception of natural-language processing both more complicated and more realistic than that underlying much of current practice. No system yet exists under this conception, though there is one under design (see note 4); therefore all I can reasonably hope to do is engage your interest and convince you that the overall conception merits consideration.

This "situated perspective" I have in mind is best thought of as a cluster of mutually reinforcing assumptions. As implied, these assumptions don't constitute a theory in and of themselves; they're merely indicative of one or several. Furthermore, a curious fact is that many of the underlying assumptions are unobjectionable. They're the sort of thing about which you find yourself thinking well, who could argue with *that*. Take the claim that language is efficient—Barwise and Perry's way of referring to the fact that you can use the same expression over and over again to talk about different things. That's as familiar a notion as the productivity of language (Barwise and Perry, 1983). Thus, to claim that language is efficient is to claim that interpretation depends on context: to know what the interpretation of the word 'you' is (i.e., who is being referred to) on a particular occasion of use depends on who you are talking to. (I will be using the word 'interpretation' to mean the actual stuff referred to, the properties it is said to have etc.) Yet, if you pair the claim about the context-dependency of interpretation of language with the equally reasonable claim that computational processes similarly have context-dependent interpretations, you are all set for the more complicated perspective on natural-language processing I mentioned.² At least, this is what I am going to argue.

Another central assumption of the "situated perspective" is the idea that meaning is relational. A kingpin of situation semantics, this assumption follows quite

reasonably from even a modest version of realism. Facts about the truth or falsity of an utterance are determined by the actual situation (used non-technically here), facts about the conversation and its participants as well as facts about the language. Thus, the meaning of the phrase 'near here' is a relation among the facts about the situation the phrase is used in, the phrase itself, and what it is being used to describe.

Though I won't catalog the assumptions that make up this "situated perspective" entirely (others will emerge as we go along), I will briefly discuss one that isn't necessary to the "situated perspective", as far as I'm concerned, though it figures centrally in situation semantics. I am not going to assume that meaning and interpretation can be defined in terms of observable behaviour (even broadly interpreted) without reference to internal architecture. Nor am I going to assume they can't. Similarly, I'm not going to assume you can or can't explain the structure of language in terms of external phenomena. A plausible assumption is that it may well take a pretty complicated story on both sides to provide an adequate theory of language use. What matters most to the situated perspective is the circumstantial dependence of interpretation, where circumstance is not restricted to phenomena external to the machine.

Now, let's put these three assumptions together—that interpretation for both language and computational process is context-dependent, that meaning is relational, and that a plausible explanation of language and language use may well make appeal to a generous supply of facts about both internal and external phenomena. Here's a way of thinking about natural-language processing in this situated perspective. We start with an agent to whom an expression u in some language has been put. The agent then processes u and arrives at some internal state m , where m is a state of the machine, defined at a particular level of description of the machine. Which level is that? The one that has interpretation that can be outside the machine and in virtue of which the machine understands u . Now there are a variety of relations that could hold between u and m , but one possible constraint is that u and m have the same interpretation; they describe the same state of the world. To put this more concretely, we obviously want the plane that the air controller refers to with his or her use of the phrase 'that plane' to be the same plane that the resulting m is about. (Of course, u might correspond to one or more m 's, and vice-versa.) An important point is that u and m can't have the same meaning: if we have adopted a relational account of meaning, then what u is related to (e.g., states of the world and m) and what m is related to (e.g., states of the world, other states of mind or machine, and u) are fairly likely not to be the same. This perspective rules out one familiar approach to natural language processing, namely, the one in which a representation of the syntax of u ($R_s u$) is first computed (e.g., by parsing), whereupon a representation of the meaning of u ($R_m u$) is said to be computed from $R_s u$, whereupon it is assumed that $R_m u$ is the same as $R_m m$

(a representation of the meaning of m). Well, all right, you say, but suppose that what is really computed from R_3u is R_iu , a representation of the interpretation of u . Then, you ask, can I assume $R_iu = R_im$? Well, you can. But then the chickens come home to roost. After all, R_im isn't m itself. And given the circumstantial dependence of m 's interpretation, getting from R_im to m may not be trivial. For instance, suppose that we've adopted a default such that the time at some point in the computation is taken to be whatever time it is when some particular bit of program is evaluated. Then R_im might be "15-OCT-86 16:30:17" where m is a kind of internal indexical having the force of 'now'. Thus, what we want, ideally, is a system that can go directly from 'now' in u to an internal state having the equivalent interpretation, in a theoretically principled fashion, and without invoking an intermediate representation in which what time it is is explicitly represented. And this seems right: after all, *we* can deal with the word 'now' without knowing what time it is.

As promised, this situated perspective does seem to have complicated things. You can't just take your favorite grammar formalism, code it up, implement a parser, derive a semantic representation and be done with it (as if that were an easy thing to do). You can't, on this view, design a language front-end for a system, unless you know the structure of the relevant level of description of the computation. And we've demonstrated that that level can't be analyzed as being equivalent to the representation of the meaning or interpretation of the expressions of the language being processed.

On the other hand, there's a positive side to all this. It may well be possible to get from u to m in much more direct ways than we have so far imagined. And, having theories of the various kinds of contextual dependency and how they interrelate should allow for more realistic (if you will pardon the pun) systems. As I said in the beginning, all this context-dependence isn't a surprise. Anyone who has taken natural-language processing seriously has had to come to grips with that property of language from the very beginning. It's just that many theories of language and computation haven't.

Imagine, if you will, a system that understands a situated language, for which there is a well understood description of the level of computation relevant to interpretation of the sort we have been discussing *and* manifests what we might call situated inference. (Note that if interpretation of internal structures is context dependent, then inference is, *de facto*, situated.) Moreover, if inference is really going to be situated, then we won't be needing to flesh out (or even necessarily disambiguate) absolutely everything upon internalization, in principle anyway. For instance, we might expect a situated robot, upon discovering a note on Jones' door saying "I'm at lunch" to infer directly that Jones was not there then and so not deliver the cup of tea it was carrying; and do this *without* using a sort of logical form that has the import of

"Jones is not in his office at 12:00 p.m. October 15, 1986". In other words, we would expect our situated inference engine to do situated inference. And, we expect this because of the overlap in the (temporal) circumstances of the situation of inference and the situation being reasoned about: it's being in the stuff it reasons about is precisely what makes it situated. Of course, if Jones later complains to me that my robot failed to deliver the tea he had ordered, I will also expect that the robot will have the capability of rendering explicit more information about the time of the failed delivery, but it need not do so initially.³ One way of looking at this is to see that it will help keep our robot from drowning in information.⁴

I said at the beginning that the situated perspective was compatible with situation semantics and its companion logical theory. Can we expect the latter to help in the design of such situated systems? The answer is yes, I think, but not in the way you (or Barwise and Perry) might have expected. Because situation semantics and situation theory are designed to account for the circumstantial relativity of interpretation, the language of situation theory is a good vehicle for a theorist to use in giving an account of the full interplay of language, inference, and computation on the view sketched here. Similarly, situation semantics isn't a bad way to go about giving an account of the external significance of language, as surely we must. On the other hand, coding up some situation semantics or replacing the semantic representation in a current system with a representation of situation semantics won't do justice, it seems, either to situation semantics or to the machine.

Notes

¹Thanks are due to the Situated Inference Engine project members at CSLI for clarification of many the ideas discussed in this paper. The research reported on was made possible by a gift from the System Development Foundation to CSLI. I would like to thank the foundation for helping to create an environment in which multi-disciplinary research projects such as the SIE are both encouraged and supported.

²I say "equally reasonable" as if it were a) obvious that there are such dependencies and b) easy to say what these dependencies are. It does seem obvious (though it is not often acknowledged) that a lisp expression can be on a particular occasion of use *about* a particular airplane, say the one just now landing at San Fransisco International. On the other hand, it seems far from obvious that it will be easy to say what these dependencies are. As in the natural language case, the information carried (by the execution of a program, for instance) is complex; dependencies arise from both the internal machine environment and the state of the external world. Delimiting the kinds of context and finding appropriate ways to

characterize the complex of relations has only just begun. For work relevant to the situated perspective laid out here, see in particular (Smith, 1986a).

³Of course it is a long way from expectation to reality. For a characterization of the internal structures that are causally responsible for an agent's or system's actions, those that have interpretations in the sense I have been using them here, see Smith (1986b).

⁴Such a system is not entirely fantasy. The Situated Inference (SIE) project at CSLI is a project to design and build a computational system that engages in situated inference. However, the point is not just that the language the SIE uses will be situated (that much is true of current natural language systems). Or even that internal structure depends likewise on circumstance for interpretation (that much is true of current systems). Rather the interest lies in the SIE's being designed with two additional purposes in mind: (i) all three, inference, internal structures, and language will be situated in compatible ways, and (ii) there is a commitment to develop a common theoretical framework in terms of which to understand the full interplay among language, content, and the internal structures etc. Progress reports on the SIE appear from time to time in the CSLI Monthly, a publication of The Center for the Study of Language and Information, Ventura Hall, Stanford University, Stanford CA. 94305.

References

- Barwise, K. J. and Perry J. 1983. *Situations and Attitudes*. Cambridge, Massachusetts: The MIT Press.
- Smith, B. C. 1986a. The Correspondence Continuum. In *Proceedings of the Sixth Canadian AI Conference*.
- Smith, B. C. 1986b. Varieties of Self-Reference. In J. Halpern (Ed.) *Proceedings of the 1986 Conference on Theoretical Aspects of Reasoning about Knowledge*. Los Altos, Calif.: Morgan Kaufmann, 19-43. Revised version to appear in *Artificial Intelligence*.