

On keeping logic in its place

Yorick Wilks
Computing Research Laboratory,
New Mexico State University,
Las Cruces, NM, 88003.

There need be no real dispute on this panel about what is meant, in the broadest terms, by formal semantics (FS) when opposed to common-sense semantics (CSS): after registering his complaints and worries, the opposition David Israel opts for in his paper is broadly the one adopted here, and the model-theoretic semanticists he mentions will do just fine for me, and I suspect for Karen Sparck-Jones in her characterisation of a "logicist" approach to natural language processing. As will appear below, though, I want to list a range of strengths of FS commitment, not all of which are model-theoretic. So, whatever we find to argue about on this panel, it needn't be those two terms.

Formal semantics (henceforth FS), at least as it relates to computational language understanding, is in one way rather like connectionism, though without the crucial prop Sejnowski's work (1986) is now widely believed to give to the latter: both are old doctrines returned, like the Bourbons having learned nothing and forgotten nothing, but FS has nothing to show as a showpiece success after all the intellectual groaning and effort. Here, I must register a small historical protest at Israel's claim that "until Montague, undeviatingly, the techniques of pure mathematical semantics were deployed for formal or artificial languages". It all depends what you mean by techniques, but Carnap in his *Meaning and Necessity* (1947) certainly thought he was applying Tarskian insights to natural language analysis. And the arguments surrounding that work, and others, were very like those we are having now. I need that point if the Bourbon analogy is to stick: FS applied to natural languages is anything but new.

But there have been recent changes in style and presentation in the purely computational area as a result of the return: many working in the computational semantics of natural language now choose to express their notations in ways more acceptable to FS than they would have bothered to do, say, ten years ago. That may be a gain for perspicuity or may be a waste of time in individual cases, but there are no clear examples, I suggest, of computational systems where a FS theory offers anything integral or fundamental to the success of the processes that could not have been achieved by those same processes described at a more common-sense level (what I am calling common-sense semantics, or CSS). However, I do not at all intend to define CSS by any particular type or level of notational description: by it I mean a primary commitment to the solution of the main problems of language processing, those problems that have not obstructed progress in the field for thirty years. That set I take to include: large scale lexical ambiguity (i.e. against realistically large sense ambiguity for lexical items of English), the problems of phrase, and other constituent, attachment, where those require meanings to fix, and the whole mass of problems that collect around the notions of expertise, plans, intentions, goals, common knowledge, reference and its relationship to topic assumption etc.

Hence, on these descriptions of FS and CSS they are not necessarily exclusive: it would be quite conceivable for an FS system to aid with the solution of a problem important to CSS. And it will always be possible for a successful CSS theory to be subsequently axiomatised---but that does not equate CSS and FS and more than axiomatising physics does away with experiments: theories come first, axiomatisations follow. My case is that that has not in fact happened in natural language processing and there is no reason to believe it will, because the origins and ultimate preoccupations of FS are always elsewhere. The examples Israel chooses to discuss in detail (Henkin and Kripke) make my point exactly. He says he could have taken others and that "the nature of the languages studied makes no essential difference". But, as he seems to concede elsewhere in the paper, those are the very areas where the techniques can be shown to work and that is why they are always chosen ("the choice of tools....should be grounded somehow in the nature of the phenomena to be analysed"). My case, to be set out in a little more detail below, is that the nature of the language chosen (natural versus artificial) makes as much difference as any choice could make, and that in the last quotation Israel is dead right.

A small concrete example may help: the choice between generating "a" and "the" is notoriously difficult in English, one that non-native speakers continually get wrong. Examples are sometimes hard to see in one's own language, and the choice between "des" and "les" in French is similarly crucial and notorious, though it is not the same distinction, yet it rests on the same kind of semantic criteria. It is not a problem with an arbitrary solution: French grammar books always claim to offer principles that underlie the choice.

It also seems, on the surface, to be a problem that FS or any logical approach to language structure ought to help with: it is certainly some form of idiosyncratic quantification. Those particles are exactly the kind that Montague grammar, say, offers large, complex structures for (just as, as Sparck-Jones notes in her paper, such systems offer such minimal, vacuous, codings for content words. This is probably the clearest quantitative distinction between FS and CSS). As far as I am aware no FS solutions have ever been offered, or would seem remotely plausible if they were; at best they would simply be a recoding in an alternative language of criteria satisfactorily expressed in other ways (see #1 below). Yet this is exactly the sort of place where FS should offer help with a concrete problem if it is to be of assistance to the NLP task at all, for it is to such items that its representational ingenuity is devoted.

The historical and intellectual source of FS lies in an alternative approach to what constitutes a proof: to meta-logical methods for the establishment of properties of whole systems, such as complete, consistent etc., and the employment of such properties as decidability to establish the validity of a theorem independently of normal proof-theoretic methods, i.e. by semantic methods. The applicability of this methodology has been perfectly clear in the case of programming languages, and to proofs of correctness of programs (even if the scope of the applications is still depressingly small) but in the application to natural language understanding its original aims have simply got lost.

From time to time, an application within the original methodology surfaces, such as Heidrich's proof (1975) of the equivalence of the methods of generative semantics and Montague grammar, but the result proved is then seen to be vacuous, in the terms of CSS at least, in that nothing was established whose absence had constituted any pre-existing problem. The usefulness or (in)adequacy of generative semantics was not anything that could be established or questioned by the sort of guaranteed

equivalences that the proof offered. The problems with generative semantics, whatever they were, lay elsewhere and were not alleviated by any such proof.

The heart of the issue is good old decidability, or whether or not the sentences of a language form a recursive set in any interesting sense. It is clear to me that they do not and I contrasted various sense in which they might back in 1971. Contributions like Israel's only makes sense on the assumption that sentences do form some such set, unless he is adopting only a "descriptive logical language position" (#1 below), and his whole position paper suggests he is doing far more than that.

The alternative position is that natural language is just not a phenomenon of the sort required and assumed by the various systems of logic under discussion; and that the interpretation of any sentence in any context is an approximative matter (including whether or not it has a plausible interpretation at all, and hence whether or not it is a sentence at all), one computed by taking a best-fit interpretation from among a number of competing candidates, and that is not a process reducible to a decidable one in any non-trivial manner. Indeed, I recall going further in (1971) and arguing (against Quine, I think it was) that that process of assigning an interpretation to a string, from among competing candidates, could be taken as a criterion of having a meaning: namely, having one from among a set of possible meanings. I mention this point for purely self-serving reasons: I do not want Israel to get away with identifying CSS with Schank and Chomsky as he does. I do not object to him joining them as bedfellows: that has been done before, all serious opponents are said to share premisses and, in the case of those two, the similarities become clearer as time goes on. Those are: a certain commitment to genetic claims, but above all a commitment to representations rather than procedures. My self-servingness is to point out that my own approach, preference semantics, was not about commitment to representations of a particular kind at all, but to certain procedures (based on coherence and connectedness of representations) as the right way to select interpretations from among competitors, for competitors there will always be. CSS can be about procedures as much as representations (Winograd has made a similar point in his (1985)). In his own commitment to representations, and dismissal of procedures ("semantics, even construed as a theory of language use, is not directly a theory of processing") Israel is actually in the same bed with that distinguished company. But not to worry, it is a big bed.

Preference semantics was, in a clear sense, procedural and had the advantage of declaring strings that did not admit the assignment of a single interpretation (i.e. remained ambiguous with respect to interpretation) as meaningless. Meaninglessness on that view was not having no meaning but having too many. I found, and still find that a satisfying position, one true to the process of language interpretation. Moreover, it offers an opportunity, it seems to me, for computation, processing, artificial intelligence, or what you will, to have something to say about fundamental questions like meaning. It is clearly an assumption of Israel, and all who think like him, that that cannot be: "real semantics", as he puts it, is being done elsewhere.

It is one of the advantages of the connectionist fun and games, from my point of view, that it has, against FS, brought implementation back to center stage from the wings; thus the new movement can be of enormous political interest and importance, whether or not one is a believer in it. At this workshop I know I shall feel like an anti-communist atheist in Poland, hugely enjoying the sermons of Bishops.

Let me try to separate levels (or may be just a continuum of positions or aspects) that the claims of FS make on natural language processing:

1) the use of a certain style of formal language of description of the basis of any system of natural language description or processing. That this claim can be pretty weak can be seen clearly if we remember McCarthy's insistence on first order forms of expression combined with his advocacy of heuristics (indeed his claim that they constitute the essence of AI). Hence, one can accede to this demand without giving any support to the central part of the FS claim about decidability.

2) Compositionality has been central to FS since Frege: from a computational point of view that doctrine is almost certainly trivial or false. I am sure this has been said many times; but I discovered it in (1984).

3) an emphasis on the particular role of quantifiers and the need of a field of distinguished entities to quantify over (this is quite independent of #1, and more central to FS). The set of distinguishable entities is easily given with Gensyms and doing that in no way concedes the FS claim. It must be admitted that it was a notable weakness in some early CSS systems (e.g. Conceptual dependency or Preference Semantics) that they did not offer a clear identification of individual entities, independent of intensional codings of concepts. But that lack was so easily remedied. The first demand can always be met by special quantifier procedures (e.g. Woods, Winograd, in their papers on procedural semantics). Nothing more is needed and demonstrations like Montague's stock example of radical quantifier ambiguity ("Every nice girl likes a sailor.") are effectively quantificational garden paths and no plausible natural language processing system is under any obligation to treat them. How could any system for the representation of natural language depend upon such cases, for they have no relation to crucial experiments in the sciences, but strong relations to the use methodological use Chomsky has made of incomprehensible infinitely-nested sentences?

4) reference and the world: the claim of FS to offer this is its weakest one, yet practitioners return to it repeatedly. If you don't adopt our methods, the claim is, you are trading in mere symbols, unrelated to the real world we, as plain men, know we share. I never cease to be amazed at the barefacedness of this: the classic statement of the position is David Lewis' attack on Fodor and Katz as peddlers of markerese (1972). But what else did he, or anyone else in FS, offer but symbol-to-symbol transformations? What else could they, in principle, ever offer by any conceivable formal methods? For symbols, and only symbols, are as much the trade and language of the denotational semanticist as of any computer modeller! Whatever the mysterious nature of the relation of symbols to things, it is not one on which FS could possibly throw light. Their solipsism is CSS's solipsism, and their position is metaphysically identical to ours.

Of course, arbitrarily named nodes identifying individuals are a handy, not to say essential, device, but no monopoly of FS. Worse yet, the proof procedures of FS demand such sets of entities to quantify over, but there is no formal way of guaranteeing that the entities established (so as to provide the guarantees that the proved theorems are true within the model that such entities form) are appropriate, in the sense of corresponding to any plausible real entities in the world. Any model set whatever that allows proved theorems to be true would suffice for the purposes of FS. Proofs of program correctness have faced this problem but FS applied to natural languages and common sense reasoning has not and cannot. Any claims to give access by such means to a plausible and appropriate world are not only false but utterly misleading as to the nature of FS.

What we reach by any formal or computational methods is always other symbols, and a "theory of meaning" for computational process over natural language should

recognise this fact (a suggestion along these lines was made in Wilks 1974).

An important additional claim of FS, certainly made in Israel's paper, is that semantics, whatever it is, must be separated off from knowledge of the world: "Lexical semantics does not yield an encyclopedia" and "Any plausible semantic account, then, will have to distinguish between analytic truths and world-knowledge". It is interesting to see that baldly and forthrightly stated without qualification, as if Quine were not still alive and well, but had never been. Practical experience with natural language processing always suggests the opposite: Sparck-Jones shows in her paper how this borderline has to be fudged by peculiar means in certain FS approach to practical problems. I am sure the Words and Worlds session of Tinlap3 will find no agreement whatever with the first quotation: real lexicons are such that information about meanings and the world are inextricably mixed, or are simply alternative formulations (at least Carnap, in opposing his "formal" and "material" modes of expression, got that right in 1947)

A final comment: we should notice the repeated offer, to our sloppy and heuristic discipline, of a Real Serious Theory for the field (remember Chomsky's for machine translation and language wellformedness; and more recently FS's for the semantics of selected AI systems). The chances always are that this prescription is unrelated to the disease; has Chomsky really helped language processing? We do indeed need a good theory but these are quack cures trading chiefly on the fears and inadequacies of practitioners and patients in the field.

One other consistent position is possible (I suspect it may be Dijkstra's, 1986): one can point out frequently, as he likes to do, the gap between the the interpretations normally given to logical entities (e.g. propositional implication or conjunction) and the interpretations usually given to apparently corresponding language items. One can also, at the same time, advocate the most stringent formal methods in computational applications to areas whose underlying structure or properties will support such methods. By such standards, natural language is not such an area and therefore one should not attempt this form of computational activity. That is a consistent position, but not one that most participants at Tinlap3 can take. It poses no problem for CSSers, but does, I believe, put a serious choice before FSers who want to remain in some way relevant to real language processing.

References:

- Carnap, R. (1947) *Meaning and Necessity*.
Dijkstra, E. (1986) Public lecture at NMSU.
Heidrich, C.M. (1975) Should generative semantics be related to intensional logic?, in Keenan (ed.) *The Formal Semantics of Natural Language*.
Lewis, D. (1972) *General Semantics*, in Davidson and Harman (eds.) *Semantics of Natural Language*.
Wilks, Y. (1971) *Decidability and Natural Language*. *Mind*.
Wilks, Y. (1974) *One small head: models and theories in linguistics*. *Foundations of Language*.
Wilks, Y. (1984) Is Frege's principle trivial or false? *Proc. Linguistic Assn. of GB, Annual Conference, University of Essex*.
Winograd, T. (1985) *Moving the semantic fulcrum*. *Linguistics & Philosophy*.