PREFACE

David Israel

In the early days of Artificial Intelligence, researchers talked confidently about
building fully intelligent entities, "robots", which could simulate (or emulate) any
natural cognitive human achievement. Among such achievements, of course, is
the (learning,) use and understanding of natural language. The early, semi-official
view seems to have been that the human linguistic capacity could be duplicated by
the application of perfectly general, non-domain specific operations to the
admittedly special, but largely then uncharted, domain of language. This was, at
the time, a perfectly reasonable position. It has ceased to be so and the essays in
this volume bespeak the recognition of that fact by the vast majority of natural
language researchers in Artificial Intelligence. They do so precisely by renouncing
even a residual obligation to speak to the issue. Whereof we should no longer
have to speak, we should pass over in silence.

Thus, it is no longer OK to "work in natural language" completely innocent of
current research in syntactic theory. It should, by the way, be less than perfectly
alright for people in syntactic theory to be as innocent of work in computational
linguistics, especially in the theory of parsing, as many still are. The same point
could be made about the inexcusability of ignorance of contemporary work in
semantics; but nowhere near so glibly. I shall return to this difference later.

With syntax and semantics "covered", can pragmatics be far behind? Indeed
not, but here the situation is very different. Certainly the source of the concepts
deployed by AI researchers is to be found in the works of philosophers (especially
Austin and Strawson, Grice and Searle). I think it can plausibly be maintained,
however, that those concepts would remain "blind" without having to
accommodate the "intuitions" that come from attempts to design computational
artifacts capable of becoming language-users. (My apologies to Immanuel Kant.)

Imagine, for instance, that one wanted to build a thing which could reasonably
be said to understand what you were talking about - at least when you were talking
about some previously delimited domain; and to act accordingly, again within the

limits of its sensors and effectors. (I shall limit myself to imagining a natural-language underlining understanding system, leaving aside the difficulties of generation. See David McDonald's piece in the volume for a pioneering effort in production.) Imagine, further, that you have been given - free of charge - a perfectly adequate parser for English.[1] Close your eyes even more tightly and suppose you find, under your bed, say, a tractable algorithm which assigns a correct semantic representation to each unambiguous sentence-under-analysis of English and the right n such to each n-ways ambiguous sentence-under-analysis. (This last bit of fantasizing might require large doses of mind-altering substances.)

How much work have you got left to do? A whole lot, and one would have to be an unregenerate syntactico-semantic imperialist to think otherwise.[2] But it was only with the disciplined attempts actually to build such devices that one could see just how much and what. And more: for it is only with such attempts that one could appreciate the extent of feed-back, and feed-forward, across the heretofore largely insurpassable boundaries of the **sentence** and, independently, of the **black box.**

For notice that, a priori, it seems natural to elaborate the little daydream hinted at above by supposing that our parser-cum-semantic interpreter works on individual sentences in glorious isolation one from another; as if even a monologue (directed, though, at another presumably intelligent agent) could be thought of as built up out of independently intelligible units of meaning. The fully modularized dream, surely, is that one runs the sentences first through the parser, then through the semantic interpreter, and then passes the output of that stage to a module (or two or three) which, inter alia, keeps usable track of the previous sentences (their syntactic analysis and their meaning?) and of the "current non-linguistic situation". (Ah, but which aspects of this, and in what form?)

Dreams die hard; but not many are left who insist that the processing story, for individual sentences, must go the route traced by the little arrows in diagrams of the linguistic module; that is, who insist on reading high-level structural diagrams as flow charts.[3] Still, there is the matter of the assumption that one can, and should, do the syntactic-semantic analyses of individual sentences in principled

---

1. Mind you, there are interesting questions afoot as to the appropriate conditions of adequacy, let alone "perfect adequacy", on parsers.
2. There are radicals who wonder whether this package from the gods is necessary. In particular, there are those who wonder how much beyond the lexicon one really requires from semantics; such people might be characterized as syntactico-pragmatic hegemonists. We shall ignore them for the duration.
3. We shall return to this point shortly.

ignorance of the uses to which those sentences, now considered as members of a coherent collective, might be put in the situation at hand. This assumption may not do. Much of the best work on discourse suggests that it will not do.

Considerations such as the above, fleshed out, of course, with the study of actual data and the analysis of hypothesized computational regimens may seem to lead researchers in discourse phenomena to a dilemma of sorts. Remember what was said above about the importance of contact with research in linguistics and philosophy. Much of the best work from those quarters is based on and argues for a very strict modularity as among syntax, semantics and, well, the rest. There is a tension, isn't there, between the import of this work and the interactions among modules hinted at above?[1] Yes and no.

There is a tangle of issues here; and, fortunately, this is not the place to attempt to untangle them. (For an heroic and insightful attempt, see Jerry Fodor's "The Modularity of Mind" [Fodor 1982]. For a more circumscribed study, see Berwick and Weinberg's "The Role of Grammars in Models of Language Use" [Berwick and Weinberg 1982]) A first point to note is that, qua designers of natural-language understanding computational artifacts, researchers in Artificial Intelligence need not be bound by current theories of human cognition, in particular of psycholinguistics; although they may be, perhaps should be, and as a number of the pieces in this volume make clear, have been, influenced by such work. (And a little more vice-versa, please.)

More concretely, we should bear in mind the happy compatibility of theories of static modularity with processing models incorporating high degrees of dynamic interaction. (I borrow the two phrases, and much else, from my colleague Rusty Bobrow, who has thought long and well about these issues.) More concretely still, processing hypotheses which stress rich interactions between syntax, semantics and even pragmatics are not ruled out, for example, by any sane version of the thesis of the Autonomy of Syntax.[2] The crucial constraint imposed by acceptance of the Autonomy thesis is that the syntactic module have, so to speak, a mind - and rules and representations - of its own, that it not speak or understand semanticese, for example. Just so, the semantic interpreter should be untainted by acquaintance with the rules of conversation. There is no prohibition, though, against the autonomously specified syntactic component being guided, as it goes about its business of assigning syntactic structures to input, by requested output from, e.g. the semantic component - as long as use of that output requires

---

1. Note that here the insurpassable barrier is not the sentence, but the boundaries of the black box.
2. For some relief on this oft tortured subject, see Noam Chomsky's "Questions of Form and Interpretation" (in [Chomsky 1977]).

no "knowledge" on syntax's part of semanticese. As long, that is, as the communication medium between the two components is highly constrained, constrained enough to be "neutral" as between the concepts proper to syntax and those within the purview of semantics alone. Note, by the way, that it makes no sense to imagine that the semantic interpreter understands no syntactic concepts; structures describable only by use of them are, after all, its input. So again, highly interactive models can capture the "asymmetry" between syntax and semantics - the first can remain ignorant of the other's terms; not so, the second. Neither, of course, need know anything about the other's internal workings.

Things become much more controversial, and much more interesting, when we venture beyond the confines of the "linguistic system proper"; as, it seems, we must when we venture upon the treatment of pragmatics and discourse phenomena generally.

Now is almost as bad a time as any to make some terminological distinctions; or to confess to having ignored them in the foregoing. There is pragmatics and there is pragmatics. When I have talked of pragmatics, I have not had in mind the work of Richard Montague, David Kaplan, and others on the logic of indexicals or token-reflexive elements, expressions whose denotata are a function of specifiable aspects of the contexts of their use. Rather, as suggested earlier, I have meant the theory of the use of language in communication. This delimitation is irremediably vague and open-ended; and therein lies the rub. Formal pragmatics (as the work of Montague et. al. [Montague 1974] might be called) is an annex of formal semantics; hence, part and parcel of the study of the linguistic system. Communication-theoretic accounts, on the other hand, can respect no such departmental boundaries. Some discourse phenomena *can* be traced directly to syntactic and semantic features of the constituent sentences, taken one at a time as it were; some, cannot. Again, the rub of open-ended interactionism.

With the rub comes the challenge. Everyone agrees that the actual use of language in communication involves all manner of cognitive modules acting together. Researchers in language, from Linguistics, Philosophy, and Artificial Intelligence, hold that this interaction must not be understood as evidence against a high degree of (static) modularization, especially with respect to the linguistic system, taken now as a whole, as against the rest of the mental apparatus. The challenge is to think clearly and in a theoretically well-motivated and disciplined way about modes and media of interaction. It's dirty work; but somebody's got to do it. Past this preface, the reader will find evidence that some, at least, are trying.

And now for an anticlimax. I warned the reader that we would have to return to the issue of formal semantics for natural languages. Forewarned is forearmed. The papers in this volume are by and large silent on specific grammatical issues; the researchers.in most cases assume the availability of parsers with usefully broad

syntactic coverage. (It is no small praise of work in Computational Linguistics to note that this assumption is a reasonably safe one.) There are, of course, open problems in both syntax and parsing; but we are by no means completely at sea. Oddly enough, a bit of the same can be said for the theory of pragmatics and discourse. Here, building on the work of Grice, Searle, and others, and extending it impressively across hefty chunks of talk, researchers in AI have developed, at the very least, a habitable framework for hypothesis and experimentation. (See, for instance, the papers in this volume.)

With respect to semantics, it's sadly a horse of a different color. As one reads the essays, one may get the impression that there has been a principled decision to opt for extensional, first-order languages (usually sortalized) as generating the logical forms of English utterances and hence as the vehicles of semantic representation. My own view is that this decision is best seen as purely a tactical or pedagogical one; at something less than worse, it might be understood as forced on one, faute de mieux. With the tactics and/or the pedagogy, I have no argument. For the wistful longing for a better way, I have only sympathy. As for the principle, I deny it. I even have something of an argument.

AN ARGUMENT:

(a) All (both) programs for formulating formal semantic accounts for significant fragments of natural languages (fragments containing "interesting" semantic constructions) exploit formal languages quite different than the language of first-order logic. We have in mind Montague-style semantics and the Situation Semantics of Barwise and Perry [1982].

(b) In other words: to our knowledge, no significant fragment of any natural language has ever been semantically analyzed by way of a (systematic) translation into a standard first-order language. Indeed, again to our knowledge, no one has ever even seriously attempted it.

(c) The foregoing facts don't seem to be accidents of history and there do seem to be good methodological reasons for the history. The main consideration is the ad hoc and unsystematic character of attempts at semantic analysis of particular sentences of English by way of paraphrase into standard first-order languages. Crucial here is the source of this ad hoc character. Any account of the semantics of natural languages which exploits a formal language has a choice about where to be risk-taking. (Such choices go with the territory of deep and general unsolved problems.) It can be novel and daring in the specification of the formal language and its semantics. This is the route taken by Montague and Barwise-Perry.

Montague's language (the language IL in "The Proper Treatment of Quantification in Ordinary English") is a "throw in everything but the kitchen sink" omega-order intensional logic, with free use of lambda-abstraction and intensional operators of all degrees. To put it crudely, the strategy seems to have been to devise a construct in the formal language for each construct in the natural language. Moreover, Montague was at all times motivated by concerns for generality; indeed, his approach might be called "Pentagon Semantics". Look at the worst-case context in the fragment in which, e.g., noun phrases occur, and assign a semantic type to noun phrases accordingly. Despite this, the semantic account of IL is, in a sense, standard; it consists in extending Tarski-style treatments of quantificational languages of arbitrary order to modal and intensional languages, an extension pioneered by - inter alia - Montague.

Barwise and Perry's formal language ALIASS, on the other hand, is designed to be much closer in its syntax to the surface syntax of English. *Exactly* what its semantics looks like is not yet fully clear in that no treatment of a significant fragment has yet been (widely) published. (There is an underground literature; there is even a bumper sticker: <u>Another Family for Situation Semantics.</u>) In some thoroughly uninteresting sense, it too will be more of the same; that is, more set theory. But this may be quite misleading, as there are hints at a more properly recursion-theoretic treatment. Thus, special constraints may be put on the kinds of sets, and operations thereon, to be allowed. Classical Tarski-style semantics places no such constraints.

To return to our theme, then, the route taken by Montague and by Barwise-Perry has the advantage that the inventiveness is confined to an area susceptible to precise mathematico-logical treatment. It also allows (in theory at least) for simple (recursive) translation procedures between English and the target formal language.

The other, standard first-order, route focuses on the procedures for paraphrase (translation) from the natural to the formal language. It leaves such procedures, however, in just the state that they assume in introductory logic texts: imprecise, non-formalizable rules-of-thumb; heuristics based on appeals to intuition. The contrast is illuminating, and, on grounds of good scientific methodology, highly unfavorable, we think, to the "conservative" strategy.

(d) Once the set-theoretic semantics of a new, non-standard, formal language has been given, one can see *in general* how to "compile" that language into a first-order language, typically with a bloated ontology. Such compilation may make great practical sense if, for instance, one has a powerful first-order logic machine sitting on one's desk. Still, keep in mind that it just may turn out that the detour through the weird and wonderful is necessary; that it may just be

(psychologically?) impossible to translate directly and systematically from a natural language into the language of first-order logic.

So much by way of argument; needless to say, more needs to be said. In the meantime, it's a pleasure to welcome the reader to this volume, to taste some of the first fruits of a happy blending of the sciences of language.

COMPUTATIONAL MODELS OF DISCOURSE