

LET'S PUT THE AI BACK IN NLP

Lawrence Birnbaum

Yale University
Department of Computer Science
New Haven, Connecticut

Artificial intelligence is, or should be, at the heart of natural language processing research. After all, it is AI more than any other cognitive science that has made *processing* a central issue in the study of the mind. Yet, it seems to me that there has been a tendency recently on the part of many natural language researchers -- even, rather inexplicably, on the part of some within the AI community itself -- to view AI as playing a secondary role in the study of language, at best a useful engineering adjunct to the more important "theoretical" studies carried out elsewhere. One need not reach as far back as the blunderbuss attack of Dresher and Hornstein -- an attack whose ferocity in fact reflected a certain amount of healthy respect for AI, or at least anxiety about its success -- to find signs of this tendency. Consider, for example, the title of a recently published book, *Natural Language Parsing: Computational, Psychological, and Theoretical Perspectives*. It seems reasonably clear what the "computational" and "psychological" perspectives mentioned in the title are intended to refer to, but what does "theoretical" mean in this context? That becomes clearer, perhaps, when we observe that the book was edited by several linguists. The use of the term "theoretical" in place of "linguistic" reflects, I suspect, not only a desire to rhyme but also an unconscious assumption on the part of the editors that they amount to the same thing when it comes to the study of parsing.

Such an assumption may be pardonable as reflecting a natural pride in their field on the part of some linguists. What is much more surprising is evidence that this attitude exists within AI as well. For example, a recent monograph on natural language processing begins by propounding the following historical view: When NLP research first began, linguists were preoccupied with syntax, so AI researchers had no choice but to cobble together semantic theories as best they could. But now that the linguists (and philosophers) have turned their attention to semantics in a serious way, these *ad hoc* AI theories can and should be replaced by implementations of the far more rigorous products of our brethren sciences. This is only a slight exaggeration of an argument which seems quite seriously intended.

It should be obvious that in my view this attitude is detrimental to progress in NLP research. AI's unique contribution to the study of the mind stems from its dedication to the proposition that functional considerations arising from the need to perform realistic tasks, rather

than considerations of parsimonious empirical description, should be the primary constraints on cognitive theories. To take the view that AI's job is to "implement" the theories produced by other cognitive sciences is therefore to abandon what makes AI worth doing in the first place. Natural language processing may include computational linguistics, but there is a lot more to it than that.

One unfortunate consequence of the tendency to ignore what AI can genuinely contribute is that a great deal of effort gets devoted to implementing theories (primarily linguistic theories) that were never intended to be process models in the first place, while somewhat paradoxically attempting to stick as close to the original conception as possible. The results are generally uninteresting both from the perspective of linguistics -- since such an implementation is likely to be, at best, only a somewhat more rigorous reformulation of the original theory -- and from the view of AI -- since the original theory was not formulated with a view towards making any interesting functional claims.

To return to an old controversy, consider the case of AI models which draw their inspiration from linguistic theories of syntactic competence -- that is, theories which attempt to capture the content of our knowledge of language structures -- which are based on the assumption of syntactic autonomy. I do not question the substantial *empirical* contributions made in pursuit of these linguistic theories themselves. The question is, what additional contributions are made by the AI theories based on them?

The majority of the parsing models which are inspired by these linguistic theories, such as ATNs and Prolog-based parsers, depend quite explicitly on nondeterminism. The rules that they employ, and the representations that they build, seem for the most part taken over from pre-existing linguistic theory. Very few of these models seem to have anything to say about how the space of hypotheses generated by the grammars that they implement are to be searched, or how they are to be integrated into the understanding process as a whole, or how these factors might impinge on the rules and representations employed. None of them, that is, has much to say about the most specifically *AI* issues involved. The indiscriminate reliance on nondeterminism is particularly troubling in this respect. As process models, these theories simply fall back on a general model of symbolic computation -- namely, backward chaining with back-up.

On the other hand, Marcus's theory of deterministic syntactic analysis is a far more profound attempt to build an AI model of parsing based on linguistic theory. Marcus tries to provide a genuinely computational explanation for certain putative properties of English syntax,

by arguing that they are a natural consequence of functionally motivated aspects of his model. I happen to think he fails -- largely because his overall claim that autonomous, deterministic syntactic analysis is possible is seriously undercut by the failure to address such issues as lexical ambiguity or genuine structural ambiguity, particularly prepositional phrase attachment -- but at least some genuine claims are being made. Unfortunately, the more recent work of Marcus, Hindle, and Fleck marks a step backwards in this regard. In order to maintain the position of autonomous syntactic processing, Marcus *et al.* propose that the output of the parser be a somewhat vague description of the syntactic structure of the input sentence, capturing whatever structural information can be gleaned without semantics or nondeterminism. What claim is being made here? It is tautological that autonomous deterministic syntactic analysis is possible if the output is defined to be whatever can be yielded in such a fashion. In order to support a meaningful claim of syntactic autonomy, certain *functional* criteria must be met: It is necessary to show that such an output will be useful, and that it can be used without violating the assumption of syntactic autonomy -- i.e., that the rest of the language processing system will not need to employ syntactic knowledge. There is good reason to believe that this is not the case.

In the case of generation, the attempt to maintain syntactic autonomy is equally counterproductive. A generator must be able to relate semantic representations and pragmatic goals to the syntactic constructions that can be used to express the appropriate meanings and achieve those goals. If a model of generation deals only with purely syntactic rules and representations then it cannot, by definition, deal with such relations, and therefore it cannot address the interesting planning issues raised by generation -- that is, the specifically AI issues. Thus, for example, McDonald's model of autonomous syntactic generation makes virtually no decisions itself -- everything from what to say to which words to use has already been decided. If this theory is to make any meaningful processing claims, therefore, it must at least be shown that the rest of the language processing system can make all of these decisions without any recourse to knowledge of the syntactic options offered by the language. Particularly in the case of lexical selection, there is good reason to doubt this.

The real tragedy in all of this, of course, is that many genuine, and genuinely important, AI problems get lost in the shuffle. NLP is desperate for good methods whereby contextual constraints can be brought to bear in a timely fashion to help resolve such problems as lexical and structural ambiguity in language analysis. Lexical selection has received far too little attention in generation research. The problem of controlling search in conversational planning remains virtually untouched. Aside from Charniak's recent work and a few other attempts, the problem of controlling explanatory inference in understanding has largely been put aside since the heyday of script/frame theory, despite its centrality. What are the criteria (e.g., parsimony,

completeness, etc.) by which explanations are judged, compared, chosen? What kinds of constraints do they impose on knowledge representations? All of these issues are crucial to natural language processing, and AI is crucial to their solution. Let's put the AI back in NLP. We might even put some of the fun back in at the same time.