

NATURAL-LANGUAGE PROCESSING: CRUCIBLE FOR COMPUTATIONAL THEORIES OF COGNITION¹

Stan Rosenschein
Artificial Intelligence Center
SRI International
333 Ravenswood Ave.
Menlo Park, California 94025

Abstract

The intellectual challenges raised by natural-language processing (NLP) are serving as a focus for artificial intelligence research as a whole; consequently, progress in NLP will almost certainly have far-reaching scientific and practical effects. This survey of recent directions in natural-language processing explains the broad conception of language and thought that underlies work being done in this field today.

1. The Promise of Natural-Language Processing Research

Natural-language processing (NLP) has a history almost as long as any sector of AI and can boast more than its share of accomplishments. Over the years, NLP researchers have developed a solid core of language analysis techniques, have implemented numerous practical systems, and have made many significant theoretical contributions to the areas of knowledge representation, inference, and planning. This paper will contend that, by and large, NLP is on the right track in its fundamental conception of the mechanisms underlying language and thought, in the specific issues it has chosen to address, and in the computational methods it has adopted. NLP is therefore likely to play an increasingly central role in AI, serving as a focal point in the development of integrated computational models of cognition.

To appreciate the importance of progress in computational approaches to language, it is necessary to understand the range of concerns that motivate research in the field. Section 2 discusses three separate but equally important incentives: practical utility, scientific implications for the broader study of language, and significance for the development of AI in general. Section 3 delves into the issues dominating the attention of language researchers today, examining problems encountered in specific subareas of language research and assessing some of the major approaches being pursued. Because each of these subareas has seen substantial

progress in recent years, Section 4 argues that our theoretical base is now strong enough to support the development of a new generation of sophisticated natural-language (NL) systems. Some methodological suggestions are made for capitalizing on that progress and facilitating the further development of both theories and systems.

2. Three Good Reasons for Studying Natural-Language Processing

2.1. Practical Utility

It seems superfluous for an ivory-tower researcher to predict that natural language systems will have practical effects. After all, there are entrepreneurs, financiers, and customers who are risking substantial sums in the belief that NL systems *already* have practical utility. What used to be a cottage industry in natural-language interfaces to databases grows from year to year in the number of products available, the number of companies involved (with several major computer companies starting to show interest), and also, presumably, in the number of customers and the volume of sales. Still, these advances are easily dwarfed by the potential impact of language processing technology. The key lies in two areas of development: (1) the spread of small, fast, cheap processors optimized for symbolic processing, and (2) speech recognition. It is generally anticipated that the former will soon become reality (perhaps with some encouragement from Japan); the latter may take longer to materialize, but confidence in its eventual accomplishment remains strong.

If we assume that these two areas will indeed develop as predicted, we shall then certainly see the application of language processing on a very wide scale. Much has been written about how every device and appliance we interact with (from cars to typewriters, from stoves to sprinkler systems) will incorporate so-called smart processors. These intelligent devices are potentially more flexible in their behavior, which means we need a flexible way of *affecting* that behavior by giving them information and communicating our desires and intentions. This is exactly the function that natural languages fulfill so well. For example, envisage yourself saying the following: "Stove, I'm going to be late tonight. If the kids call and say they'll be home early, cook up the spaghetti. Otherwise, wait till about 7:00 before starting on the casserole; I'd like it to be done around twenty minutes or so after I arrive."

¹This work was partially supported by the Defense Advanced Research Projects Agency under Contract N00039-80-C-0575 with the Naval Electronic Systems Command and by the Office of Naval Research under Contract N000014-80-C-0296. The views and conclusion contained in this article are those of the author and should not be interpreted as representative of the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency, the Office of Naval Research, or the United States government.

Some experts have argued that formal languages should be used for interacting with these devices because formal languages are more "precise." Of course, one could just as easily claim on the same grounds that *people* should talk to one another in formal languages--indeed, sometimes they do (e.g., in mathematics). But more often, of course, they do not, and for good reason. To use a formal language requires effort and special training; why bother if English serves the need? This rationale applies equally to conversations with intelligent machines.

2.2. Scientific Synthesis in Related Disciplines

By related disciplines, I mean primarily linguistics, philosophy, and psychology. At present, each of these is quite "comfortable" operating within its own traditional boundaries, making occasional nodding references to the others. None of them has come up with a compelling framework for understanding the phenomenon of language in its totality; none has achieved a grand synthesis. This is not surprising; after all, language is the manifestation of a mental *process*, and an enormously complex one at that. The science of computation is the only science which studies complex processes with any degree of rigor: how to induce them, how to describe them, how to relate levels of description, etc. Without the concepts of computer science in their analytic toolbox, these other disciplines are severely handicapped. This is not to say that they lack insight. Still, their separate theories have no common framework, no "glue." The computational approach provides precisely that "glue," a perspective on syntax, meaning, and use of language that interlinks these concerns by relating everything back to formally and precisely described *processes*.

2.3. NLP Research: A Powerful Accelerator for AI

2.3.1. Explicit Models of Mind

The development of an integrated model of mind is a long-term goal for AI as a whole, but for language processing it is an immediate necessity. The purpose of the mind is to understand and affect the environment; the purpose of language is enable its users to understand and affect other minds. We can imagine a "mind" that has a great deal of knowledge about infectious disease without its having any knowledge of other *minds*. So when we build an expert system, say, for infectious disease, we do not feel compelled to formalize a theory of mind just to have a metadescription of the expert system; we feel quite satisfied with our informal models. In any case, the expert system is not even a full "mind" with complex goals and intentions, but rather a specialized reasoning engine. In the case of language-processing systems, however, we are forced to deal with difficult problems that can be safely ignored in other domains.

This is true even in seemingly mundane applications. Database interfaces, for example, are quite likely to encounter queries like the following:

User: I would like to know which products
the US imported from Japan last year.

There are at least three reasons a system able to respond appropriately to this query would have to have a model of mind. First, and perhaps most obviously, the sentence

contains the word "know," a verb of propositional attitude; any reasonable semantic analysis of this verb would presuppose at least a rudimentary model of mind. Second, the sentence is a declarative sentence, which literally only states something about a preference of the speaker. To understand the speaker's true intent, the hearer must conceive of the speaker as *having* intent; this too induces us to devise an explicit theory of mind.

Finally, the response generated by the system will be judged according to its effects on the mind of the hearer. He said he wanted to "know" something. In what form should that knowledge be imparted? Should every product be listed? Should the information be summarized by categories? If the system doesn't know how to answer the query directly, to what lengths should it go to satisfy the speaker's desire? Is the speaker really interested in the specific information he requested or does he actually want to know whether our imports are concentrated in high-technology items or steel? If the system can impart more useful general knowledge, should it? To respond appropriately, the system must model the effects of its response on the mind of the user. Such a requirement imposes upon the system-builder, as an essential component, an explicit concern for models of mind.

Even without the need for the language system to model other minds, there would be a trend in NLP to focus on integration and comprehensive theories of cognition. To fully understand language requires a broad range of abilities: high-level planning, low-level signal processing, the representation of complex knowledge, the integration of special-purpose algorithms (parsing, for example) with general-purpose reasoning systems, focus of attention (e.g., in a discourse), and belief revision in light of new information. Any application at which such diverse concerns converge will almost inevitably induce designers to take some position on overall mental "architecture."

2.3.2. Knowledge Representation and Acquisition

NLP has a special role to play in advancing knowledge-representation research. Artificial intelligence studies techniques for enabling a computational process embedded in the world to model (or, to use common AI terminology, "represent") this world. Natural language, in addition to its many other functions, is perhaps most importantly a device for describing aspects of the world in which we live. Its vocabulary provides complex elements of meaning, its syntax allows those elements to be combined in ever new and fascinating ways, and, by uttering sentences, we are able to transmit complex information about the world (with extremely fine control over the subtlest nuances of meaning) and to cause complex inferences to be drawn.

Thus, researchers in representation for AI systems would do well to study natural language as a rich, multifaceted source of inspiration and ideas. For instance, they might learn much about which abstractions of the everyday world are useful by examining natural language semantics, both the general semantics of various phrase and sentence types and the detailed semantics of the lexicon. A systematic investigation of this sort

specifically aimed at knowledge-representation questions might greatly accelerate our progress in so-called common-sense theories (cf. [8]). In addition, we might ask what inference-drawing mechanism seems implied by the way semantic content is packaged in natural language. This is likely to give us a new perspective on inference—perhaps somewhat richer than “structured objects” and inheritance, yet more computationally specialized than general-purpose logical deduction.

Research in this vein can also serve practical needs in the more application-oriented areas of AI. For instance, NL systems could fill a practical need as a tool in knowledge acquisition, frequently cited as a major bottleneck for expert systems. The expert's knowledge is often informal and most easily expressed in natural language, while *formalized* knowledge is required to implement an expert system. The passage from informal to formal knowledge is often difficult, especially for concepts that are built deeply into our language and thought. Any system that could mechanically extract the intuitive vocabulary and content from text provided by an expert, then transform it into a formal representation, could help bridge this gap.

A knowledge-acquisition system that is based on natural language could do much to uncover buried structural relationships implicit in language and often missed by human “knowledge engineers.” These include temporal relations as expressed by tense, causation as expressed in causative verbs, and notions of possibility and likelihood as expressed by modals and certain modifiers. There is already enough known about the semantics of these constructs and how they might be represented in a computer to make a knowledge-acquisition aid conceivable.

3. Where We Are in NLP Research

3.1. Levels of Language

It has been conventional practice to divide the study of language into *levels*, a convenient but admittedly gross and inadequate classification. The lowest level (if we agree to ignore phonology) is *syntax*, which is largely the study of grammaticality and phrase structure—the factors that determine which words can go in which places, and how larger phrases are built up out of smaller subphrases. *Semantics*, the next level, is concerned with the meaning of the various syntactic units (notably the sentence), where “meaning” has to do with the conditions in the world that the words are *about*. *Pragmatics*, the highest level, is concerned with the uses to which language is put and how speakers' goals are achieved by uttering sentences in context.

Progress in these areas comes in spurts, both in linguistics and natural-language processing. Modern linguistics began with a major spurt in syntax with Chomsky's theory of *generative grammar*; first proposed in the fifties, it was developed further in the sixties into a large body of work on *transformational grammar* and is continuing in modified form today. During the last several years there has also been an upsurge of nontransformational theories, generalized phrase structure

grammar [6] and lexical-functional grammar [4], to mention two examples.

In semantics the linguists have been greatly influenced by the logicians, especially Montague [12] and, more recently, the theory of situation semantics developed by Barwise and Perry [3]. Pragmatics is still largely uncharted territory to students of language, though the works of Searle [22], Austin [2], Kamp [9], and others, have been influential.

In natural-language processing, too, syntax was the first area investigated. The early sixties saw the growth of automated syntactic analysis of natural language, with much of its impetus derived from the study of parsing for computer languages. Semantic issues began to be addressed in the late sixties and early seventies, primarily at MIT [23] and Stanford [20]. Although even the early question-answering systems of the sixties can be said to have *embodied* a pragmatic component, it was not until the late seventies with the work of Grosz on discourse [7] and Perrault and his students on computational approaches to speech acts [15] that concern with pragmatics became generally evident and serious attempts were made to give the pragmatics of NLP some theoretical grounding.

The relative recency of research in pragmatics is somewhat surprising since the fundamental reason for the *existence* of language is pragmatic—the need to understand and control the complex cognitive states of others. Indeed, from the standpoint of functional utility and, plausibly, order of evolutionary development, pragmatics precedes semantics—which in turn precedes syntax. Yet, we have seen that, in the modern study of language, formal theories of syntax (Chomsky) preceded formal theories of semantics (Montague); we are still waiting for a formal theory of pragmatics and discourse. Thus, the growth in our formal understanding seems to be exactly the reverse of functional importance and development. Perhaps the theoretical progression from syntax to semantics to pragmatics can be viewed as “retracing” our developmental steps.

3.2. Fundamental Issues and Methodologies

Every area of language study can be viewed as having several distinct components. One is the choice of framework and, quite often, the choice of formalism. Another is the detailed knowledge, usually expressed in the chosen formalism. Third, perhaps, is the computational realization of this knowledge. Of course, the rate of progress in exploring these components differs, but some overall judgments are still possible. If percentages had to be assigned to the three major subareas, we might say that syntax is 80% solved, semantics 50%, and pragmatics 20%.

Syntacticians might object that 80% is a gross overestimate, and indeed computational syntax is a vigorous research area. Nevertheless, compared to the other subareas, we can say that syntax is a largely solved problem because there is broad agreement on the general question of what syntactic facts are, what kind of grammars can be written to account for them, how these grammars can be used to analyze language

automatically, and so on. Research has consequently moved on to second-order questions: elegance and generality in syntactic descriptions, improvements in parsing, ways of handling ungrammatical input, ways of tying syntax to semantics and pragmatics, disambiguation, scoping, discourse features signaled by syntax, and of course there is still much interest in the construction of larger, more accurate grammars [16, 19].

In semantics and pragmatics, on the other hand, there is still much disagreement about fundamental issues, stemming largely from divergent views about knowledge representation. For a large segment of the research community, the main conceptual tool here is logic, and nobody who is ignorant of logic can really do credible work in semantics, pragmatics, or in knowledge representation generally. This is admittedly a controversial point that I will not attempt to defend here, but it might be mentioned that in linguistics too semantics took a long time to develop into a serious area of research; it wasn't until a logician (Montague) offered a set of formal conceptual tools (as Chomsky had done with syntax) that the field began to make rapid progress. Now linguists can correlate syntax and semantics, and study fine-grained semantic phenomena with as much analytic precision as they study syntax.

There is a general recognition in AI of the difference between logic as a *conceptual framework* and general-purpose deduction as a *processing strategy*; one can embrace the one while rejecting the other, although both have been used to advantage in NLP. However, the depth of feeling generated by the question of logic's role in AI seems to indicate something deeper than the mere choice of conceptual or programming tools. It seems that, on the one hand, there are many researchers who feel intellectually uncomfortable with logic and for this reason try to denigrate its relevance as a tool. On the other hand, those who *do* feel comfortable with logic often use it as a weapon to assert their intellectual superiority over those who do not. This is a continuing problem for NLP, as it has been for AI as a whole.

In any event, even those who believe in logic as a conceptual framework for semantics and pragmatics do not see it as the entire answer, because logic in itself is neutral with respect to processing. In this regard, Schank and Winograd had more insight than Montague, who was developing his theories around the same time. Although their semantic representations could only be judged crude as compared with Montague's, they realized instinctively that, to explain human language use, representations of meaning had to be treated as computational objects, not just formal objects (e.g. functions from possible worlds to truth values.) Likewise, world knowledge implied more than just the requirement that admissible interpretations satisfy meaning postulates. Model-theoretic abstractions have their place, but they beg particularly interesting questions, e.g. how does the knowledge of all those possible worlds fit in our finite heads? What is the computational *mechanism*? Of course, in the early years Schank, at least, probably went too far in denying the legitimacy of formal semantics (and formal **syntax**, for that matter). But his concern, shared with Winograd, for language as a **total** process was exactly what was needed.

3.3. Why Is Language Processing So Difficult?

3.3.1. Pragmatics

As we argued above, it is quite likely that language emerged against a background of a fairly well-developed intelligence. People who already knew how to infer complex facts about their environments, make judgments on the basis of uncertain evidence, learn, reason, and the like could apply these faculties in trying to decipher one another's intended messages. The relative advantage gained from the acquisition of language was that it reduced the degree of uncertainty in people's minds regarding one another's internal cognitive state and increased mutual control over cognitive state, and hence over one another's behavior. Language functions pragmatically as a substitute for ignorance and wild conjecture.

But to reduce uncertainty is not to eliminate it. In spite of our complex system of linguistic conventions, pragmatics has to contend with the problems of interpreting events, revising beliefs, and predicting the effects of actions under conditions of extreme uncertainty. In the interpretation of any complex reality, there are usually numerous plausible theories as to what is actually going on. This presents enormous problems to system-builders attempting to incorporate pragmatic processing. On the comprehension side, they must stipulate in great detail which theory or interpretation is to be preferred. On the generation side, they must encode strategies for choosing from an enormous set of potential linguistic actions an utterance which will influence the world in just the desired manner. In general, this sort of detailed stipulation of pragmatic knowledge is exceptionally difficult; we are much too close to the process ourselves and often have difficulty seeing the forest for the trees.

Discourse presents special problems, since the level at which the mind must be modeled to explain discourse phenomena (e.g., focus) is more detailed than the level of just "having beliefs," which appears largely adequate for defining the speech-act sector of pragmatics. For example, we need to know how objects can become highlighted, how objects meeting certain descriptions can be, or fail to be, retrieved. This level of modeling the effects of a sentential utterance introduces many problems when we try to connect it with the "logical" meaning of the sentence. For instance, it is usual to suppose that definite references ("the big red ball") pick out objects that are somehow in focus in the hearer's mental discourse representation. In certain quantified cases in which a universal quantifier (e.g., "every," "each") has wider scope than a definite determiner (such as "the"), as in "the ticket bought by each man," there can be no single semantic object in focus, since the phrase refers to a *different* ticket for each man. Some of Kamp's work seeks a formal framework for addressing such problems as these, making the connection to model theory as well. However, computational approaches have inspired much of the work in this area and are likely to be a key component in the solution to the discourse state puzzle.

3.3.2. Semantics

In semantics, though the problems are far from solved, the situation is a little simpler. Here logic has provided a rather formal notion explicating the attachment of "meaning" to a syntactic language; this notion has served linguists well over the last decade. Although NLP has been split over the role of logic, there are large areas of agreement. For example, the computational semanticist who uses Montague's intensional logic as a meaning representation language may have more in common with a Schankian using conceptual dependency than he does with a theoretical linguist engaged in model-theoretic semantics. The linguist is only interested in abstract truth conditions for sentences, whereas the computational linguists are primarily interested in translating natural language into some formal language that lends itself to inferential manipulation.

Despite the controversy about what that language ought to be, there is a great deal of common focus on specific problems. This stems from the fact that *any* choice of target language to represent meaning carries with it the obligation of deciding how each of the various problematic phenomena of natural language is to be treated. These phenomena include tense, modalities, quantifiers (including nonstandard ones), sentential complements, intensional adjectives, and so on. To simplify our discussion considerably we can affirm that what makes semantics difficult is (1) comprehending what meaning is in the first place, (2) determining the most natural computational correlates of meaning, and (3) understanding the detailed structure employed by linguistic forms to carry these meanings.

3.3.3. Syntax

Although theoretical linguists are still much concerned with questions of framework in syntactic theory, the difficulties with syntax in NLP would exist even with consensus on framework. Indeed, in practice there is a good deal of consensus among the system-builders. The syntactic component of almost every major natural-language system uses some form of phrase structure tree (with features on the nodes) as the representation of syntactic structure and some form of grammar with a context-free skeleton and manipulation of the features as the specification of syntactic knowledge. Some systems keep the grammar "declarative" and quite separate from the parser, while others mix parser and grammar, giving the linguistic knowledge a more procedural flavor. Notable examples of this latter style include Winograd's SHRDLU program [23], various systems inspired by Schank and Riesbeck, and Marcus's parser [11], in which the grammar "rulesets" are, in effect, procedural routines. Our knowledge of syntax is perhaps the most precise form of our linguistic knowledge, but the breadth of our knowledge is sometimes actually an impediment. In particular, we generally know how to assign more syntactic readings to a sentence than were intended, some of which can be eliminated by a more careful writing of the grammar, while others require integrated understanding to resolve. But the process of refining our detailed syntactic knowledge has caused our grammars to grow, which in turn has led to a renewed

interest in grammatical (and metagrammatical) formalisms as we struggle to organize our grammatical knowledge more effectively.

4. The Future of NLP

4.1. The NLP Research Community: Sharing Ideas and Code

The NLP research community has collectively accumulated a great deal of intellectual "property": large, detailed grammars, knowledge bases, parsing techniques, representation languages, and considerable data in the form of interesting texts and dialogues. But we have here a clear case of the whole being less than the sum of its parts! Part of the reason, of course, is human nature—an innate reluctance to invest time and effort into learning what others have already done. But the willingness to expend effort is surely a function of perceived relevance; so far researchers have not been able to describe the results of their work in a way which allows it to be easily integrated with the results achieved by others. Common notations are the exception rather than the rule, crisply posed problems are rare, and easily understood formal descriptions of systems *at a level of abstraction useful to other researchers* have been even rarer. Most of our arguments are in prose, most of our articles are highly anecdotal. Given this state of affairs (this too is true for most of AI), it is not surprising that the cumulative effects of research have been less than could be hoped for.

We can think of language research as being arranged along three orthogonal dimensions: vertically (in terms of conceptual levels of language), horizontally (in terms of different research groups and paradigms), and along two planes in the third dimension (theoretical and system-building). There are signs that a process of increasing cohesion may be underway along each dimension.

Let us first consider the vertical dimension along the theoretical plane. For a linguistic theory to be vertically integrated, there must be smooth links in the theory between syntax and semantics, and between semantics and pragmatics. As regards the first connection, there has always been a tendency among theoreticians and grammar writers alike to strive for rule-to-rule correspondence between syntax and semantics. That is, each syntactic rule in a grammar that tells how to build more complex syntactic units from simpler ones is associated with a semantic rule indicating how to build complex "meanings" out of simpler ones. During the last few years, systems have been written [5, 18] demonstrating that the distance between "meanings," as conceived of both by linguistic theorists and their computational cousins, is not as great as had previously been supposed. In particular, the logics that formal semanticists regard as mere conveniences in specifying truth conditions, can often be pressed into service as target representations in computational systems. Often such logics can be manipulated deductively with far greater ease than their complex model theory would suggest.

The conceptual boundary between semantics and

pragmatics has also been considerably "smoothed" over the past few years. This is largely because of the increased use of formal languages for describing both the semantics of NL and the cognitive state of agents. Logic-based meaning representation languages have been developed that are (1) rich enough to describe the semantics of English constructs [13], (2) convenient to mechanically translate English into [21, 18, 5], and (3) formal enough to embed in an outer, "pragmatic" language. When this outer language contains operators such as "believe" and "want," which take propositional arguments, it is quite straightforward to express the idea that an agent *believes* [pragmatic] the *content* [semantic] of a sentence. For instance, if we use a logical language that allows us to quantify over the objects of belief, we can write

$$\exists p.\text{meaning}(\text{sent},p) \wedge \text{believe}(\text{agt},p,t)$$

The use of such notations makes it possible to describe the connections between *meanings* and *cognitive states* (cf. [17]) with a high degree of precision and rigor. This is a very positive accomplishment and one that augurs well for further developments in the field.

As far as the horizontal dimension is concerned, we are beginning to see a process of sifting through linguistic knowledge for abstracted kernels that can be shared. One example of this is in the area of syntax, where there is incipient recognition of the commonality of approach to the study of Lexical-Functional Grammars, Unification Grammars, Definite Clause Grammars, and other feature-propagating grammatical formalisms [14, 4, 10, 18]. This may ultimately lead to more effective sharing of the sort of detailed grammatical knowledge now being painstakingly developed in parallel for a number of separate formalisms (and implemented systems). If we understand the connections well enough, we can share pieces of grammars either by writing them in a common language or by passing them through some grammar translator or notational preprocessor.

Vertical integration is also starting to take place along the system building plane. At least one logic-based computational system has emerged [1] that tries to combine detailed reasoning about cognitive states with lower-level linguistic knowledge. And, of course, systems built at Yale have *always* been vertically integrated; their problem has not been lack of integration, but lack of connection with more broadly based theories of language.

Refinement has also been taking place in our notion as to how a language processing system should be organized. There has always been tension between the desire to have the modules of the language system correspond to the conceptual levels of language (roughly, syntax, semantics, and pragmatics) and the need to commingle the levels. The former affords conceptual clarity at the expense of artificial boundaries that tempt us to ignore interactions. The latter allows interactions, but threatens to swamp the system-builder with complexity. This tension is far from resolved, but the clearer our understanding of the interactions, the better we are able to parameterize the effects and control complexity. For example, by determining exactly which syntactic phenomena are correlated with discourse

phenomena we can incorporate the information into the feature system we need for syntax in any case, thus encapsulating the interaction and making it manageable. Another example of this would be in the design of intermediate languages encoding aspects of user intent that are derivable from syntactic form (or, in a speech-understanding system, intonation), but which are not strictly part of the *meaning* of an utterance.

4.2. The Broader Study of Language, Cognition, and Computation

The various language-related disciplines are starting to establish the requisite connections. Models of mind that are now emerging contain a componential analysis of mind with a clear relation to models of pragmatics and discourse. Use of the language and concepts of computation can facilitate communication across disciplinary boundaries. As computational linguists start constructing systems that bear recognizable links to theories of mind and language developing in the broader research community, a more unified conception begins to coalesce. There is growing recognition that while it is legitimate to study mental structures and capabilities in the abstract, these will eventually have to be explained in computational terms.

Once the principle is accepted that descriptions of mind and language are descriptions of a computational process *at some level of abstraction*, we have at least the potential for integrating our cognitive theories and systems. We can view theories emerging from linguistics, philosophy, and psychology as abstract descriptions of systems we could build, in which the connection between the code that implements the system and the abstract theory that describes it could, in principle, be made as rigorous as we like. As we mentioned earlier, the activity of establishing formal connections between levels of description of processes is one of the best-developed parts of computer science, virtually a specialty of the field.

This, in turn, offers the possibility of more systematic development of intelligent systems instead of current handcrafting. Instead of talking vaguely of giving our systems "knowledge," we shall understand in a deeper, more technical way how the "having knowledge" level of description is systematically related to more conventional levels of description. This will allow us to build "plug-compatible" world knowledge modules, knowledge compilers, and grammars seen as a specialized form of knowledge. It will enable us to integrate the special-purpose form of the parser/grammar theoretically within the larger notion of compiled knowledge. (Similar insights will apply to other forms of compiled knowledge, e.g., the knowledge of physics compiled into our visual sensors.) Psychological theories will be "abstract predicates" that will or will not apply to our systems. In short, the passage from theory to code and back again will be conceptually smooth. This is the goal toward which NLP research is currently moving.

Acknowledgments

I have benefited from discussions with Jane Robinson, Bob Moore, Barbara Grosz, Jerry Hobbs, Kurt Konolige and Stuart Shieber. These colleagues, along with many others, have contributed to the creative and stimulating intellectual climate at the SRI AI Center. Of course, I bear sole responsibility for any errors or misconceptions.

References

1. Appelt, Douglas E. Planning Natural-Language Utterances to Satisfy Multiple Goals. Technical Note 259, SRI International, Menlo Park, Ca., 1982.
2. Austin, J.L.. *How to Do Things with Words*. Oxford University Press, Oxford, 1962.
3. Barwise, Jon and John Perry. Semantic Innocence and Uncompromising Situations. In *Midwest Studies in Philosophy VI*, P. French, T. Uehling and H. Wettstein, Eds., University of Minnesota Press, Minneapolis, Minn., 1981, pp. 387-404.
4. Bresnan, Joan W. and Ronald Kaplan. Lexical-Functional Grammar: A Formal System for Grammatical Representation. In *The Mental Representation of Grammatical Relations*, Bresnan, J., Ed., MIT Press, Cambridge, Mass., 1980.
5. Gawron, Jean Mark *et al.*. Processing English with a Generalized Phrase Structure Grammar. Proceedings of the 20th Annual Meeting of the Association for Computational Linguistics, June, 1982.
6. Gazdar, Gerald and G. K. Pullum. Generalized Phrase Structure Grammar: A Theoretical Synopsis. Indiana University Linguistics Club, Bloomington, Indiana, 1982.
7. Grosz, Barbara. *The Representation and Use of Focus in Dialogue Understanding*. Ph.D. Th., University of California, Berkeley, 1977.
8. Hayes, Patrick J. The Naive Physics Manifesto. In *Expert Systems in the Micro-electronic Age*, Donald Michie, Ed., Edinburgh University Press, Edinburgh, 1979.
9. Kamp, H. A Theory of Truth and Semantic Representation. In *Mathematical Centre Tracts 135: Formal Methods in the Study of Language*, J. A. G. Groenendijk, T. M. V. Janssen and M. B. J. Stokhof, Eds., Mathematisch Centrum, Amsterdam, 1981.
10. Kay, M. Functional Grammar. Proc. of the Fifth Annual Meeting of the Berkeley Linguistic Society, Berkeley Linguistic Society, Berkeley, California, February 17-19, 1979, pp. 142-158.
11. Marcus, Mitchell. A Computational Account of Some Constraints on Language. In *Elements of Discourse Understanding*, A.K. Joshi, B. Webber, and I. Sag, Eds., Cambridge University Press, Edinburgh, 1981.
12. Montague, Richard. The Proper Treatment of Quantification in Ordinary English. In *Formal Philosophy*, Richmond H. Thomason, Ed., Yale University Press, New Haven, 1974.
13. Moore, Robert C. Problems in Logical Form. Proceedings of the 19th Annual Meeting of the Association for Computational Linguistics, June, 1981.
14. F. C. N. Pereira and D. H. D. Warren. "Definite Clause Grammars for Language Analysis--a Survey of the Formalism and a Comparison with Augmented Transition Networks." *Artificial Intelligence* 13 (1980), 231-278.
15. Perrault, C. Raymond and J. F. Allen. "A Plan-Based Analysis of Indirect Speech Acts." *American Journal of Computational Linguistics* 6, 3 (1980), 167-182.
16. Robinson, Jane. "Diagram: a Grammar for Dialogues." *Communications of the ACM* 25, 1 (1982), 27-47.
17. Rosenschein, Stanley J. Abstract Theories of Discourse and the Formal Specification of Programs that Converse. In *Elements of Discourse Understanding*, A.K. Joshi, B. Webber, and I. Sag, Eds., Cambridge University Press, Edinburgh, 1981.
18. Rosenschein, Stanley J. and Stuart M. Shieber. Translating English into Logical Form. Proceedings of the 20th Annual Meeting of the Association for Computational Linguistics, June, 1982.
19. Sager, Naomi. *Natural Language Information Processing*. Addison-Wesley, Reading, Mass., 1981.
20. Schank, Roger C. and Kenneth M. Colby, eds.. *Computer Models of Thought and Language*. Freeman, San Francisco, 1973.
21. Schubert, Lenhart K. and F.J. Pelletier. "From English to Logic: Context-Free Computation of 'Conventional' Logical Translation." *American Journal of Computational Linguistics* 8, 1 (1982), 27-44.
22. Searle, John R.. *Speech Acts*. Cambridge University Press, Cambridge, 1969.
23. Winograd, Terry. *Understanding Natural Language*. Academic Press, New York, 1972.