CONSIDERATIONS FOR
COMPUTATIONAL THEORIES OF SPEAKING:
SEVEN THINGS SPEAKERS DO

John H. Clippinger, Jr.
Teleos
Cambridge MA 02138

I. INTRODUCTION

Fundamental progress has to do with
the reinterpretation of basic
ideas.

Whitehead

The opinion which is fated to be
ultimately agreed to by all who
investigate it is what we mean by
the truth and the object
represented in this opinion is the
real.

C. Pierce

Any discipline in the course of its
maturation experiences identity crises in
deciding what it is and where it is going.
Often an initial commonality of interests,
methods, and even disgruntlements hides an
eventual diversity of purposes and goals.
It is often at this time that self-conscious
decisions have to be made regarding the
overall scope and directions of its
enterprises. To a limited degree,
computational linguistics is beginning to
experience its own birth pains. Bedfellows
who did not seem so strange initially are
beginning to appear progressively more so.
Perhaps now is a time for stock taking and
some preliminary self reflection.

The principal intent of this paper is
to sort out what I feel to be some of the
more basic purposes of computational
linguistic research, and then to discuss
some of the explanatory requirements for one
aspect of such research, the modeling of the
human speaker. The bulk of the paper will
focus upon some of the more basic issues in
describing human speaking by a computational
model.

II. THE DUAL CHARACTER OF COMPUTATIONAL
LINGUISTIC RESEARCH

Superficially, what distinguishes
computational linguistics from other hybrid
forms of linguistic research
(sociolinguistics, psycholinguistics,
ethnolinguistics, generative linguistics,
etc.) is the concern to represent linguistic
descriptions in terms of computer programs.
The common bond in computational linguistic
research is the computer, or stated somewhat
more broadly, the information processing
approach. However, while this common
interest does serve to distinguish this form
of linguistic research from a growing number
of others, it does, I believe, masque some
more elemental differences over the purposes
of using computer programs to describe
language. There seem to me to be two
complementary but ultimately different
directions in computational linguistics; the

first is concerned primarily with improving
and developing the software technology
whereby computers can use and process
natural language; under this heading would
be machine translation, question and
answering systems, automated secretaries,
and assorted text processing systems; the
second is concerned with using the computer
as a means of developing an accurate and
empirically valid computational model of the
linguistic and cognitive behaviors of a
human speaker. To date these two interests
have been regarded as one; the rationale
being that human beings and machines would
use analogous mechanisms in their
comprehension and generation of language to
make the system work efficiently. But if
one's goals are descriptive, then I don't
feel that any such latitude is justified and
that the use of different models and
approaches should be justified with
reference to certain theoretical criteria
and empirical evidence--to the extent it is
possible. This of course, is not to say,
that the two branches of research, the
technological and the descriptive, will not
or should not share their results or
influence one another. Only that they each
have different objectives and hence
different methodological and theoretical
requirements.

Having made the distinction between
technological and descriptive forms of
computational linguistics, I am now going to
have to hedge a bit. Because, whereas both
forms of research do differ with respect to
specific research objectives, they do share
a common meta-language and methodology, and
are dependent upon that language for their
success. I am now referring to the
dependence of computational linguistics upon
abstract theories of programming and
computation. Yet even this distinction
between the abstract and the applied
modeling is not so clearcut, as practical
programming requirements, such as language
comprehension, affect the design of
programming languages, and they in turn
affect how specific problems are represented
and understood. However, I suspect that as
the field of computational linguistics
matures, the descriptive and the
technological forms of research are going to
exert different selective pressures on
computational theory to produce different
types of programming languages, one being
more specialized and efficient, and the
other being more generalized and probably
less "practical". Yet here too a
qualification is needed, for where the more
general and human-like problem solving
language is shown to have practical and
technological uses, then it too will have a
technological application.

Therefore, any advance in improving a
machine's facility with natural language
would analogously constitute an advance in
our understanding of human language
behavior. The hedge word here, of course,
is analogously, for it is not clear to me
whether a program designed for very
specialized and restricted purposes can be
any more than superficially analogous to a
human speaker. A simulated secretary for

example has far fewer concerns than an actual secretary and therefore is designed and programmed for more specialized purposes. Moreover, and perhaps most significantly, the simulated secretary has no inherent intentionality of its own; it doesn't care what others think of it, it does not need to please, it isn't distracted, and it doesn't deplore chauvinism and resist demeaning requests. Certainly there are analogies between what the simulated and the real secretary do, as both perform similar clerical and linguistic tasks, but these similarities are quite few when compared to their dissimilarities. Moreover, it is highly questionable whether the programming techniques and methods used to model this very restricted form of behavior can be easily generalized to other types of behavior. Most experience in artificial intelligence suggests the contrary. The point which I believe needs to be made is that while, yes, at some level, the information processing behavior of machines and programs is similar or analogous to that of human beings, that for the most part, this level is too general to be of much use in understanding human thought and speech at present. Hence, computational linguistics and artificial intelligence should begin to distinguish between those research goals that are technologically oriented and those that are intended to be descriptive of human action. For if one's goals are purely technological such as machine translation or a question and answering system, then one is certainly justified in using virtually any mechanism one wants.

The problem is that in our ignorance of how higher order and unobservable human problem solving processes and mechanisms work, we have to rely upon experiments with abstract programming languages and simulated environments just to give us ideas about how such mechanisms might work. Certainly this is a reasonable strategy to pursue at this level of abstraction and experimentation. But when the intent then becomes the formulation of a computational model of some form of human activity, then I believe that there is a need to provide some theoretical or empirical justification for the model. For example, if the model is of a human interpreter of discourse, does the model make the same mistakes as people do, does it resolve ambiguity in the same way, can it make similar inferences from discourse as people, etc.? To answer these questions requires considerably more empirical evidence than we currently have.

The real issue here is accountability. For programming languages and technological applications of computational linguistics this is relatively easy to determine. But when the intent is descriptive, that is, when a model is presented as representing some human intellectual skill or process, then the issue of accountability becomes murky indeed. For it is not enough to say that it works, or that expert judges cannot distinguish the model from the real thing; Weizenbaum's Eliza program showed how easy it is to attribute powers to the computer it

doesn't possess and Colby's indistinguishability tests have similarly shown the gullibility of trained psychiatrists. Consequently there is a need on the part of descriptive computational linguistics to both specify the types of tasks a model of a speaker or listener should perform and in addition, to the extent it is possible, the manner in which they should be performed.

The problem then becomes knowing what it is one is wanting to describe and account for. And this is no easy matter, for as Wittgenstein noted (Wittgenstein, 1953), "the aspects of things that are most important for us are hidden because of their simplicity and familiarity (one is unable to notice something--because it is always before one's eyes)." Chomsky makes a similar point (Chomsky, 1970) "As native speakers we have a vast amount of data available to us. For this reason it is easy to fall into the trap of believing that there is nothing to be explained, that whatever organizing principles and underlying mechanisms may exist must be 'given' as the data is given." Traditionally linguists -- Chomsky among them -- have focussed upon the grammatical aspects of language use because it was so "given" and assumed that in accounting for the generation of grammatical utterances that they have in some way described at least some of the mechanisms involved in the generation of human speech. But as Max Black (Black, 1970), and a number of sociolinguists have pointed out (Hymes, 1972; Labov, 1974) to be able to characterize the grammaticality of a language with a finite set of rules does not mean that speakers of that language utilize such rules in their speaking. Therefore to my mind there is the real question of the value of such descriptions in accounting for either the individual or the collective use of language. Consequently when dealing even with the most obvious of linguistic "facts" there is a need to consider the more general role of language as a communicative, problem solving, and expressive medium. Similarly there is a need to characterize the type of explanation being sought; whether we want to explain language behavior in terms of reasons and intentions, or whether we want to describe it extensionally, in terms of causes.

This latter distinction is one which Toulmin (Toulmin, 1970), Radnitsky (Radnitsky, 1970), Dennett (Dennett, 1975) and Goffman (Goffman, 1974), make in their various discussions of the types of explanations appropriate to the human or social sciences. Taking Toulmin's argument for the moment, causes are essentially like physical laws, as they are beyond our control, whereas "actions done for reasons can be regarded as applications of procedures (methods of calculation, techniques, rituals, or other formalized modes of behavior) that we have learned during our life time." While ideally it may be possible to explain human language behavior in terms of causal descriptions, for example by neurological models, it is doubtful that this is the type of

description that computational linguistics is seeking. Rather the appeal of the computational approach is precisely in its capacity to characterize the symbolic procedures that we use and have learned. If this orientation is accepted, then I believe that the obvious conclusion to be made is that language behavior should be described and explained as a learned intentional activity. Going one step further, then, one of its goals should be the explanation of the various reasons and purposes for different types of linguistic behavior. However, to answer such questions satisfactorily requires a twofold description of the reasons for a procedural action. For taken from a historical or diachronic perspective, a given procedure is created to solve a particular set of problems at a particular time; hence its existence or reason for being is set in time. However, the same procedure may then again be used at some later time to achieve some other action perhaps unrelated to the first, and therefore has a reason independent of its derivation. Consequently, the reason a particular procedure was used can be explained in terms of the reason for its coming into being, or in terms of the immediate effect or result it was invoked to achieve. Psychotherapy is full of such cases: "Why do you smoke?" "Because I was bottle fed" or "Because I like the taste?" Similarly, for example, in understanding a speaker's use of intensifiers or dubitatives it is important to know whether they have a specific local meaning or whether they are a part of some standard discourse style -- either personal or cultural. For modeling purposes the differences are important, as they entail differences in representation.

III. CONSIDERATIONS FOR A COMPUTATIONAL THEORY OF SPEAKING

So far we have established that that segment of computational linguistic research concerned with describing human speakers should take some initial steps towards saying more concretely what it is trying to account for. Certainly if we are going to evaluate a piece of research in this field, we should have some preliminary consensus as to what a successful model of speaking should be able to do. Likewise, there should be some common understanding as to the types of descriptions and explanations that are being sought.

Taking the second question for the moment, I think that one of the unique contributions of a computationally based methodology to the human sciences is its capacity to give formal and teleological descriptions to complex forms of symbolic behavior. Therefore one of the principal, if not inviolate, ingredients of a computational theory of speaking should be a computationally based methodology. This requirement would to my mind disqualify statistically based, and to a lesser extent, predicate calculus based, models. Furthermore, I feel that it is incumbent upon a modeler of human discourse to make

informed use of more advanced programming concepts, as problems resolved at this level eventually contribute a richer and more complete representation at the purely descriptive level. For example, although notions such as frames, mini-worlds, demons, actors, pattern matching and the like are essentially programming concepts, they do represent solutions or partial solutions to programming complex symbolic environments which in all likelihood are far simpler than those encountered by the everyday speaker. Consequently a computational model of speaking requires these techniques and methods both technically and theoretically.

Now finally what a computational theory of speaking should in part, at least, account for. To date we seem to know very little about the types of mechanisms involved in the creation, formation, and regulation of speaking (here also, is included discourse and conversation). And while it would certainly be reassuring to compartmentalize speaking as a clearly bounded activity, both sociolinguists and phenomenologists have effectively dashed that hope.

Since we can't say what speaking is -- if indeed it is any single thing (language does a disservice here), we can point out some of those things that people appear to do while speaking. It follows then that a model of human speaking should be able to do these too.

Seven Things Speakers Do

1. People normally initiate a statement or a discourse out of a desire to be understood. They have some goal, or effect they want to achieve which they think speaking might facilitate.
2. Speakers alter what they are saying according to the physical and social context in which they are speaking. This in turn affects what is said and how it is said.
3. Speaker's have models of to whom they are speaking. They shape their remarks according to how they feel they are going to be interpreted. They apparently monitor what they are going to say, often making mistakes and changing their minds as they are speaking.
4. Speakers make effective use of the thematic organization of their conversations to direct interpretation, to specify role relationships, to qualify, intensify, and to amuse.
5. Speaking has a tone to it. Sometimes it is angry, other times placating. It can also be erudite, reverent, direct, evasive, sardonic, etc.
6. Speaking can also be stylized, having cultural and sub-cultural

constraints on what is said and how it is said; for example, Black narrative styles.

7. Speakers make effective use of idioms, "buzz words", vernacular expressions, slang words to express themselves.

Very likely most all these aspects of speaking appear as obvious, but to paraphrase the previous Wittgenstein quotation, it is often the obvious that escapes us. The argument can be made that yes, indeed, we will get to these things but let us first write a program that first speaks grammatically. I think that misses the point. Grammaticality is only one aspect of speaking behavior, and one which I doubt is as focal as it is made out to be. Grammar is just one of the means that speakers use to communicate their thoughts, intentions, and feelings to others. It should not be that element around which all subsequent investigations should be built. I feel that future work done with empirical data on speaking behavior will bear this point out and perhaps encourage the development of more comprehensive computational theories of speaking as an intentional and expressive human activity. If these theories are going to have a descriptive and explanatory value, they will at least have to be able to do these seven things that human speakers "obviously" do.

REFERENCES

- Chomsky, N., "Problems of Explanation in Linguistics," in Explanation in the Behavioral Sciences, Borger, R., Cioffi, F., eds., Cambridge University Press, p. 425-470, 1970.
- Dennett, D., Why the Law of Effect will not go Away, forthcoming, 1975.
- Goffman, I., Frame Analysis, Harper and Row, 1974.
- Hymes, D., Towards Communicative Competence, University of Pennsylvania Press, 1972.
- Labov, W., personal communication.
- Radnitsky, G., Contemporary Schools of Metascience, Scandinavian University Books, 1970.
- Toulmin, S., "Reasons and Causes" in Explanation in the Behavioral Sciences, Borger, R., Cioffi, F., eds., Cambridge University Press, p. 1-26, 1970.
- Wittgenstein, L., Philosophical Investigations, Blackwell, 1953.