

AD-A087 250

UNIVERSITY OF SOUTHERN CALIFORNIA MARINA DEL REY INFO--ETC F/6 5/7
TOWARD A SPEECH ACT THEORY FOR NATURAL LANGUAGE PROCESSING.(U)

MAR 80 W C MANN

NSF-MC578-07#10

NL

UNCLASSIFIED

ISI/RR-79-75

AD
A087250

END
DATE
FILMED
9-80
DTIC

154
LEVEL II

12

ISI/RR-79-75

March 1980



William C. Mann

Toward a Speech Act Theory for Natural Language Processing

ADA 087250

DTIC
ECTE
JUL 29 1980
D
C

This document has been approved
for public release and sale; its
distribution is unlimited.

INFORMATION SCIENCES INSTITUTE

4676 Admiralty Way/Marina del Rey/California 90291
(213) 822-1511



UNIVERSITY OF SOUTHERN CALIFORNIA

80 7 25 187

UUC FILE COPY

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

DD FORM 1473
JAN 73

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

20. Abstract

This report explores the prospects for using concepts from Speech Act Theory in the design of processes that operate on natural language. The potential benefits of creating processes to identify the illocutionary force of utterances in text are particularly significant. These benefits include systematic derivation of implicit communication, identification of relations between text and prior text, and the possibility of applying existing action-oriented knowledge within AI to new natural language processing tasks.

The report not only explores these benefits and the potential uses of Speech Act Theory but also exhibits some simple processes for identifying illocutionary forces and for simulating the effects (including much implicit communication) of speech acts.

Accession For	
NTIS GRA&I	☒
DDC TAB	☐
Unannounced	☐
Justification	
By _____	
Distribution/	
Availability Codes	
Dist	Avail and/or special
A	

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

ISI/RR-79-75

March 1980



William C. Mann

Toward a Speech Act Theory for Natural Language Processing

UNIVERSITY OF SOUTHERN CALIFORNIA



INFORMATION SCIENCES INSTITUTE

4676 Admiralty Way/Marina del Rey/California 90291

(213) 822-1511

This material is based on work supported by the National Science Foundation under Grant No. MCS78-07410. Any opinions, findings, conclusions or recommendations expressed in this publication are those of the author and do not necessarily reflect the views of the National Science Foundation.

CONTENTS

Abstract	iv
Preface	v
1. Toward Incorporating Speech Act Theory Into Natural Language Processing	1
1.1 Some Persistent Sources of Difficulty in Natural Language Processing	1
1.2 Conceptual Similarities	3
2. Speech Act Theory	5
2.1 Illocutionary Forces	5
2.2 Conditions on Speech Acts	7
2.3 Syntax and Illocutionary Force	8
3. Uses of Speech Act Theory In Natural Language Processing	10
3.1 Convenient Concept for System Organization	11
3.2 Access to Useful Knowledge	11
3.3 Necessary Level of Abstraction	13
3.4 Component of a Theoretical Base for Discourse Processing	13
4. Gaps In the Theory	14
4.1 Enumerating the Effects of Performing a Speech Act	14
4.2 Recognizing Illocutionary Forces	15
5. Some Simplistic Approximations for Filling the Gaps	17
5.1 A Simple Method for Enumerating the Effects of a Speech Act	18
5.2 Toward a Simple Illocutionary Force Recognition Process	19
6. Conclusions	24
References	25

PRECEDING PAGE BLANK-NOT FILMED

ABSTRACT

This report explores the prospects for using concepts from Speech Act Theory in the design of processes that operate on natural language. The potential benefits of creating processes to identify the illocutionary force of utterances in text are particularly significant. These benefits include systematic derivation of implicit communication, identification of relations between text and prior text, and the possibility of applying existing action-oriented knowledge within AI to new natural language processing tasks.

The report not only explores these benefits and the potential uses of Speech Act Theory but also exhibits some simple processes for identifying illocutionary forces and for simulating the effects (including much implicit communication) of speech acts.

PREFACE

This is an introductory report, aimed at workers in Artificial Intelligence who are interested in but not sufficiently aware of the relevant work in philosophy. It is intended mainly to bring Speech Act Theory to the attention of the reader in an accessible way and to advocate its active pursuit. At the same time, the report serves to identify consequential divergences between methods of philosophy and AI, divergences that will require AI workers to supplement the existing base.

Since this report was written in mid-1978, some people, but not others, have found it useful. Several readers have seriously misunderstood its intent and have read it expecting content it was never intended to convey. The report is intentionally sketchy, in the hope that more substantive work might be stimulated. In particular, the section on processes (5) is intended to be suggestive rather than substantive, although it goes beyond many of the simple treatments of speech act phenomena in existing AI systems.

Since the early circulation of the report, some of the hoped for activity has been seen. While we cannot take credit for much of it, the trend is encouraging. However, the general status of the work is unchanged. Existing treatments of speech acts in AI systems are incomplete and mostly unproven. There is now an even greater opportunity for significant basic work on speech acts in the AI tradition. Recent work by Phillip Cohen, James Allen, and Gretchen Brown exemplifies the trend.

1. TOWARD INCORPORATING SPEECH ACT THEORY INTO NATURAL LANGUAGE PROCESSING

1.1 Some Persistent Sources of Difficulty in Natural Language Processing

The dominant natural language processing problem in computer science has been the design of person-computer interfaces. Robots, question-answering systems, instructional systems and expertise-delivery systems have been designed with natural-language interfaces, and there is a continuing interest in doing more. However, computer science actually addresses a much broader range of ambitions for processing natural language, including comprehension of unrestricted text, automatic indexing and retrieval of text, and machine translation. So far, accomplishments on these broader problems have been fragmentary and preliminary, mainly limited to exploratory attempts to subdivide the problems and to demonstrations of techniques for solving particular subproblems.

Progress has been limited by four general difficulties:

1. Knowledge/Action--Natural language has been conceived as an obscure knowledge representation system rather than as a system of actions, in contrast to the action-oriented conceptions of programming language. This has severely limited the degree to which the knowledge of processes has been applied to understanding natural language.
2. Explicit/Implicit--Although text conveys a great deal of its meaning implicitly, natural language processing research has concentrated on the explicit portion of the meaning, especially the explicit meaning of individual sentences.
3. Content/Context--Although text derives a great deal of its meaning from the context in which it occurs, natural language processing research has concentrated on the portion of the meaning that can be identified in isolation from context.
4. Message/Communication--Although text is used for many other purposes besides conveying information, research has concentrated on text used for informing (and for requesting information), with little attention to the way that text affects its recipients.

The difficulties have often been traceable to a lack of suitable concepts--definite enough for implementation, broad enough to be significant, and plausible enough to try.

Given these problems, Speech Act Theory (SAT) is attractive as a source of help. It offers an action-oriented conceptual framework, specific ideas about implicit communication, some help in relating text to context, and significant steps toward a communication model of language interpretation. SAT can also make a direct contribution to language comprehension processing because the communicated content of text includes the knowledge of the speech acts performed by that text.

The purpose of this report is to explore the possibility that SAT could be genuinely helpful in addressing these problems (and to the general progress of natural language processing (NLP)). A second purpose is to facilitate SAT's use by filling some gaps--suggesting ways to meet some needs of NLP that have not yet been given sufficient attention in SAT.¹ (But see [6, 4, 14].)

We should recognize from the outset that NLP can only adapt SAT, not simply implement it. There are significant differences in the amount of detail that it is necessary to represent, and NLP criteria for judging the adequacy of a technical result are different from those in other disciplines. We (in NLP) put a high value on explicitness, on representation as a collection of processes, and on having tested the processes on a diversity of cases, but a low value on parsimony and on total freedom from counter-examples. A process producing approximately correct results is preferred to a perfect abstract description not complete enough to be translatable into process form. This difference in methodology, working by successive explicit approximations rather than by accumulating well-attested abstractions, leads to significant differences of interest and emphasis between SAT and NLP.

The first part (sections 2 and 3) of this report reviews some relevant themes in modern SAT and explores the uses that could be made of SAT in NLP. The second part (sections 4 and 5) examines two specific gaps in SAT and suggests approaches to the creation of processes that would fill them.

¹ In doing so we make no attempt here to supplement or correct SAT itself. Rather we take it as is, even to the extent of preserving a few errors of detail in the examples. (Previous readers of this report have found apparent errors in the quoted material by Searle and Habermas, but the value of their concepts does not rest on resolution of such errors.) The perspective of NLP does help one to see how SAT might be improved, but doing so is outside the scope of this report.

1.2 Conceptual Similarities

Speech Act Theory arose in philosophy, beginning with J. L. Austin in the 1950's [3]. It is currently a very active topic in the part of the field called Philosophy of Language. Major works on SAT include *Speech Acts* by John R. Searle and Volume 3 of *Syntax and Semantics* [21, 7].

SAT is particularly accessible to workers in NLP because many of its concepts closely resemble familiar concepts in computer science--in some cases concepts for which we already have algorithms and experience. The most important commonality is the *active* nature of the basic element of SAT, the *speech act*. A speaker performs a speech act by producing an *utterance*. A speech act can be decomposed into its *propositional part* and its *illocutionary force*.

- (1) Thomas Jefferson is buried in Virginia.
- (2) George Washington is buried in Virginia.
- (3) Is George Washington buried in Virginia?

Utterances 1 and 2 differ in proposition expressed but not in illocutionary force; 2 and 3 differ in illocutionary force but not in proposition expressed.

The illocutionary force is comparable to an operator in a programming language.² The speaker and hearer form a two-processor information system roughly analogous to the system consisting of the programmer and his computer.

The division between illocutionary force (the operator) and the proposition expressed (the operand) is important for the same reasons that the division between program and data is important. The operators are reused with high frequency and are part of the user's communicative competence, whereas the propositional parts are highly variable and seldom recur. (For this reason, the term "speech act" is often used informally to mean only the illocutionary force, without the accompanying propositional part.)

² There are various ways to take this notion of speech act as operator. One can regard an utterance U as encoding a recognizable force F and proposition P, and that speaking U is an event that causes a corresponding event, the speech act F(P). At another level of abstraction, much more closely related to a theory of communication, F(P) operates on a state of the hearer to produce a new systematically altered state. Use of SAT concepts is thus a possible part of a scheme of comprehension more oriented toward communication than present schemes.

For both programming languages and natural languages, a significant part of the skill of using the language is in knowing how to use combinations of operators to achieve a desired effect.

There is an interesting resemblance between the *felicity conditions* that are part of SAT and the preconditions and input assertions associated with computational processes.

If you see the statement

(4) $X \leftarrow X + 1$

In the middle of a program, knowing about the addition and assignment operators, $+$ and $\leftarrow$, you can draw several conclusions:

a. that X has a value,

b. that the value of X is suitable as an operand for $+$,

and so forth. These are conditions on the effective use of the statement. In some programming languages such implicit conditions are used to avoid redundant type-declarations and thus achieve brevity.

Similarly, if someone makes the following request to you:

(5) "Please give me five dollars."

you can conclude that

a. he believes you can give him five dollars,

b. he wants you to give him five dollars,

c. if he did not make this request, then you might not give him five dollars,

and so forth. These are likewise conditions on the effective use of the statement [21].

The forms of these particular conclusions come from the operator--requesting--not from the propositional part. People very frequently communicate such assertions implicitly, just by the illocutionary force of their speech acts. As with the implicit type-declarations in programming languages, it is a way of being brief.

There is much discussion in the speech act literature of the *intent* of the speaker, of what intentions accompany what actions, and of how each is related to meaning and to the consequences of speaking. There are direct parallels to the *goal-pursuit schemes in AI*, where means-ends analysis, formal goals, planning, and plan execution have been applied to robotry and other tasks. Of course the robots did not pursue their goals by speaking, but much of the knowledge is clearly incorporable in NLP.

Other resemblances, all imperfect but some rather close, are indicated in Figure 1.

SPEECH ACT	COMPUTER SCIENCE
illocutionary force	operator
proposition	operand
category of act	category of operator
felicity condition	precondition, input assertion
illocutionary effect	main effect
perlocutionary effect	side effect
indirect speech act	use of one operator to achieve the effect of another

Figure 1. Correspondences of SPEECH ACT concepts and COMPUTER SCIENCE concepts

Before we can see how to adapt SAT for NLP we must understand more of what a speech act is in the terms of Philosophy of Language.

2. SPEECH ACT THEORY

2.1 Illocutionary Forces

Searle introduces speech acts with this example:

- (6a) 1. Sam smokes habitually.
- (6b) 2. Does Sam smoke habitually?
- (6c) 3. Sam, smoke habitually!
- (6d) 4. Would that Sam smoked habitually.

Each performs a different kind of action, but all of them express in various ways a single proposition and make a particular set of references. They do not differ in

truth-value, but in the kind of action performed, so each is assigned a different illocutionary force [21].

Searle has written a taxonomy of illocutionary forces in which he identifies 12 dimensions along which illocutionary forces differ. He posits five basic categories, named Representatives, Directives, Commissives, Expressives, and Declarations; these are exemplified (for speaker S, hearer H and proposition P) in Figure 2.

Examples of Utterances in Searle's Five Primary Categories:

Representative--commits S to the truth of P

(7) "The Emperor has no clothes on."

Directive--S attempts to get H to do something

(8) "Begin the parade!"

Commissive--commits S to future course of action

(9) "I will march through the city."

Expressive--expresses psychological state of S regarding P

(10) "Thank you for your clever work."

Declaration--creates the condition expressed in P

(11) "You are sentenced to 1 year in jail."

Figure 2. Five categories of utterances

Within each category there is much diversity. Several other workers have produced taxonomies, with up to several hundred particular verbs, such as "boast" and "warn," indicating kinds of acts at the lowest level [10].

A speaker performs a speech act with a particular illocutionary force when he makes an utterance that conforms to a particular set of *constitutive rules*. The rules are

constitutive in the sense that they create the possibility of the action, just as the rules of chess create the possibility of making various moves in chess. The nature of these constitutive rules is widely debated, but they are generally seen as a further specification of Austin's claim that speech acts operate "by convention"³ [8, 3].

2.2 Conditions on Speech Acts

Associated with each illocutionary force *F* there are various conditions, sometimes called *felicity conditions*. These are propositions that must be satisfied if the utterance is actually to have the force *F*. For example, for a request by *S* that *H* perform action *A*, the felicity conditions would include

1. that *S* believes that *H* is able to perform *A*,
2. that it is not obvious to *S* that *H* is going to perform *A* without being requested to do so,
3. that *S* wants *H* to perform *A*.

Some of these conditions are testable or at least assessable by the hearer. Those that are not, which represent some part of the "psychological state" of *S*, such as 1 and 3 above, are called *sincerity conditions*. For a Representative *R(P)*, there is a sincerity condition that *S* believes *P*.

There are usually several of these conditions on the force of an utterance--typically 2 to 10. Each condition can give rise to some implicit communication, since by performing the act the speaker indicates implicitly that he believes that the conditions for performing that act hold. (So from the first condition on a request above, the request communicates that *S* believes that *H* is able to do *A*.) The result is that from this mechanism alone there is far more implicit communication than explicit.⁴

³The debate on the nature of the rules has also shown that many of the conditions associated with particular acts by Searle and others are in need of refinement. The present defects seem to be matters of detail rather than representing unsoundness in the conceptual structure. In addition to Searle's, there is an alternate line of development of SAT begun by Alston. The two lines differ far more at the interface between SAT and ethics than in the way they characterize illocutionary force [1, 2].

⁴This is an important consideration in building comprehension systems. If most of what is communicated arises implicitly rather than by direct interpretation of the symbols employed, then strong mechanisms for imputing implicit communication are essential to effective comprehension.

SEARLE'S DESCRIPTION OF A REQUEST: (Speech Acts, p. 66)

Speaker S requests hearer H to do action A by speaking utterance U.

Propositional content: Future act A of H

Preparatory conditions

1. H is able to do A. S believes H is able to do A.
2. It is not obvious to both S and H that H will do A in the normal course of events of his own accord.

Sincerity condition: S wants H to do A

Essential condition: Counts as an attempt to get H to do A.

2.3 Syntax and Illocutionary Force

How is illocutionary force related to the syntax of sentences?

We know that there may be syntactic correlates of the illocutionary force. The Directives include questions and commands, which have their own syntactic forms. Commissives may start out "I will..." or "I promise...."

We might therefore suspect that analysis of the syntactic form and vocabulary of an utterance would fully determine its illocutionary force, and that therefore illocutionary force is merely a relatively obscure syntactic property of an utterance. This is not the case, as we shall see in a moment.

Some utterances are used to perform what are called Indirect Speech Acts. For these, instead of or in addition to a most straightforward "direct" act, there is another act performed that differs in illocutionary force, proposition expressed or both, from the direct act. Declarative sentences such as

(12) "You are standing on my hand."

(13) (compare "Get off of my hand.")

and questions such as

(14) "Do you have a match?"

(15) (compare "Please give me a match.")

can be used to perform acts that are not just statements of information or requests for information respectively.

Given a particular sentence that can be used to perform an indirect speech act, such as (12) or (14) above, it appears that one can almost always construct a context in which that sentence has its direct illocutionary force rather than the more easily imagined indirect force. This being the case, we are guaranteed that the sequence of words of the sentence is insufficient for determining whether that sentence is used with direct or indirect force. So, *for the great majority of utterances, illocutionary force cannot be assessed on syntactic grounds*, nor on any other grounds that deal only with the sequence of words that constitute the utterance. This is not to say that syntactic form and speech act are not related, but the relation is far from determinative. Coyne [p. 17] says,

It is of crucial importance that form, logical or grammatical, does not seem to correlate nicely with illocutionary force except in the most artificial, pompous or ceremonial of instances.*

*This is not to deny the existence of certain typical relations: it is common to perform assertions by issuing utterances in the indicative mood, commands by issuing utterances in the imperative, etc. It does not take much examination to see, however, that such relations are merely typical, certainly not universal, and not in any straightforward way necessary.

Sentences like (12) and (14), which can have different illocutionary forces in different contexts, demonstrate the ambiguity of illocutionary force.⁵ (In addition to

⁵In this report the "context" of an utterance always means text accompanying that utterance, and "situation" always means the extralinguistic knowledge applicable in the interpretation of the text.

ambiguity, there are also utterances having more than one illocutionary force, and therefore performing more than one speech act.)

Similarly,

(16) "Do you need money?"

can be used as either a request for information or an offer.

Notice that questions such as

(17) "Do you know what time it is?"

are used just as predicates are often used in LISP--the user expects to get back either NIL or a useful value. As such they function as simple questions if the answer is negative, but also as requests if the answer is affirmative.

3. USES OF SPEECH ACT THEORY IN NATURAL LANGUAGE PROCESSING

Because the aims and methods of Philosophy of Language differ greatly from those of NLP, SAT must play a different role in NLP than it has played in Philosophy of Language. To function in NLP the theory must somehow facilitate building systems of information processes that produce or analyze natural language. There are several potential uses to consider, and it is important in this case to see more than one use in SAT.⁶

⁶ These are uses in the NLP discipline, not in the systems NLP produces. In language processing systems, the potential uses for SAT depend on the system task.

Whether or not SAT is represented in the organizing abstractions of a text-analysis system, the system may have to face illocutionary force discrimination issues.

(a) Your files are all composed of text.

(b) I will run out of tapes before Friday.

Sentence (a) can be either an assertion or a request for verification; (b) can be either a prediction or a promise. *The illocutionary force of an utterance is often one of the essential pieces of information about that utterance.*

Identifying the illocutionary force leads directly to a number of propositions about the speaker, his perception and knowledge of his own situation, and possibly his perception of the hearer. In a system these may function as premises for inference, as additional knowledge derived from the text or as useful sources of options in reference resolution and ellipsis filling. Since there are normally syntactic and semantic ambiguities to resolve, these propositions can also provide information on the reasonableness of various interpretations. For text where concepts of belief, doubt, argumentation, or credibility are important, knowledge of illocutionary force would seem to be particularly desirable for use in deriving the agent-to-knowledge relations.

For text generation, the various illocutionary forces can serve as the primitive actions of the system.

3.1 Convenient Concept for System Organization

The concepts of SAT can be used by system designers to impose conceptual order on complex phenomena. By recognizing certain aspects of language use as actions, the designer is better equipped conceptually to create the necessary processes. At the same time, opportunities to use action-oriented knowledge from AI--knowledge of planning, goal-pursuit, action selection, and action control--can be recognized. The theory can thus provide a convenient level of abstraction.

A similar use of the theory is to include in these systems processes that model the judgments of SAT, such as classification of an utterance according to its illocutionary force. As part of some larger task such as text comprehension, assigning these judgments to a special set of processes may provide a convenient factoring into tractable parts of the large and diverse set of judgments such a system must make.

Such a factoring is by no means necessary, but it may reduce the complexity of individual tasks and make the system easier to explain.

The decision to apply SAT in such a system is like the decision to implement a particular level in a Hearsay system [16].

A great many circumstances are relevant, and the choice depends on more than just technical necessities. There are so many clever ideas promising conveniences in implementation that convenience alone does not justify much attention to SAT.

3.2 Access to Useful Knowledge

A second use for SAT is as a common vocabulary for assimilating knowledge about language from outside NLP. There are substantive claims about language within linguistics, psychology, philosophy, computer science and elsewhere. While many of the specific claims of one discipline often seem incommensurate with those from another, statements about SAT constructs are exceptional; they often cross discipline boundaries. If speech acts are seen as actions, many questions are of common interest to the various disciplines, e.g.,

1. Under what conditions is this action performed?
2. What are its effects?

3. By what method is this action performed?
4. How can the action be recognized?
5. Can it be decomposed into simpler actions?
6. How do the various actions differ?

The various disciplines pose these questions in such a way that they overlap in their domains, as do the various answers, creating opportunities for results from abstract studies in other disciplines to be moved into concrete computer systems.

For example, Habermas (following Wunderlich) [12, 11, 18, 25] claims that under certain easily satisfied conditions there are four kinds of validity claims on each utterance, recognized by both speaker and hearer as a kind of background consensus:

- a. that the utterance is understandable,
- b. that its propositional content is true,
- c. that the speaker utters it sincerely, and
- d. that the speaker has a right to perform the speech act.

These clearly have consequences for the builder of processes to comprehend language, since they may be part of the substantive communication of the utterance and since they are subject to the same consistency conditions as the explicit content of the utterance.⁷

Notice too that the implicit communication specified by these claims is not the same as that specified by the conditions on particular speech acts, since it is constant across all illocutionary forces. Neither is it the same as the direct effect of any particular act. Thus Habermas' claims suggest several specifications on particular kinds of implicit communications occurring for all utterances and spanning all domains of knowledge.

This potential application of Habermas' work shows how SAT can serve as a kind of "market language" for transferring into computer science results from other disciplines using SAT. For example, if we are using the Hearsay blackboard-of-hypotheses

⁷ Habermas also holds that if one of these four claims is called into question, then the conventional interpretive rules of the communication change in systematic, predictable ways. He claims further that the communications of scientific inquiry typically take place under the altered rules, not the rules of a "normal action context." Clearly such claims are relevant and suggestive to the designer of a text interpreter, both for knowledge representation (e.g., multiple sets of rules) and control structure (e.g., a stack for abandonment and restoration of particular interpretation methods).

architecture, specific implicit-communication mechanisms can be implemented by new, specialized knowledge-sources.⁸

This is not simply an interesting and useful result; it can be included in a system design because Habermas, a social scientist, has framed his claims in terms of SAT. We need to know and use SAT in order to do such importing.

3.3 Necessary Level of Abstraction

Beyond the "market language" role, use of SAT can become a necessary part of natural language processor design, essential in the same way that organic chemistry is essential between genetics and atomic physics. It may capture generalizations about the domain that can be stated in simple form only if this particular level of abstraction is recognized. (The sincerity condition on any utterance functioning as an assertion--that S believes the expressed proposition--might be such a generalization.) It is an open issue whether SAT can fill such a role. The evidence seems to me to be quite persuasive, but still inconclusive.

3.4 Component of a Theoretical Base for Discourse Processing

Processes that recognize or perform speech acts are essential to certain kinds of discourse processing based on discourse theories incorporating speech acts explicitly. A processor for any task that involves relating the details of text to motivations for producing the text might well be based on a theory that incorporates speech acts explicitly.

For example, an experimental system called DCS (Dialogue Comprehension System), currently under development, recognizes five principal kinds of objects in comprehending a dialogue transcript: syntactic subtrees, propositions, speech acts, dialogue-game uses and participants' goal structures. The theory is based on the premise that each utterance serves some overt goal of the speaker. The discourse structure it discovers has the connectivity of goal pursuit, with the primitive actions the speaker employs to pursue his goals being speech acts. Speech acts are suitable for embedding in plans

⁸ It is not necessary to have perfect agreement between the concept as implemented and the concept as previously described, nor to demonstrate empirically that they agree. Some deviation is inevitable, and the second draft may turn out better than the first.

produced for immediate purposes, and they also are seen as part of culturally transmitted shared plans.

The theoretical view is described in [17].

Which of these uses is sufficiently important to justify immediate attention in computer science? Certainly the convenience use is attractive, but many things look convenient. The "market language" use seems sufficient and more to the point, since NLP cannot afford simply to develop its own independent, idiosyncratic view of language. Whether or not SAT turns out to be "necessary," it can make an immediate contribution to NLP.

Beyond merely exploiting SAT as a source of convenient concepts for organizing familiar operations, NLP should begin to take advantage of the insights in other disciplines in which SAT has been accepted and used, including linguistics and cognitive psychology (in linguistics, [9, 10, 13, 15]; in psychology, [5, 19]). The explicitness of much of SAT, the similarity of some of its constructs to familiar computational constructs, and the broadening acceptance of SAT by other disciplines all make the effort required to incorporate it in NLP worthwhile. Nevertheless, there are difficulties.

4. GAPS IN THE THEORY

As in any other intellectual movement, speech act theorists have focused their attention on a small number of central concerns. In doing so, they have touched only lightly on some problems that would be prominent in attempts to use the theory in NLP.

Here we consider only two--enumerating the effects of a speech act and methods for recognizing the illocutionary force of an utterance.

4.1 Enumerating the Effects of Performing a Speech Act

Performing a speech act has particular effects, some of which arise directly from the fact that the act has been performed, others less directly. There has been serious effort to distinguish acts from effects, for example, to distinguish an act of Threatening from an effect of Intimidation. Much of the discussion centers on the terms

"Illocutionary" and "perlocutionary" introduced by Austin. Perlocutionary effects are often treated as a sort of residue class, comparable to the inconsequential side effects of a computation--a place to dismiss things from further consideration.

Speech act literature does not include a comprehensive enumeration of the effects of performing a speech act. For example, to my knowledge there has been no exploration of the contrast between informing and reminding, nor of the contrast between telling P to H1, who didn't know that P, and telling it to H2, who knew that NOT P. However, there are some interesting cases discussed, with suggestions about their specific effects.

On the positive side, there are some promising suggestions. Rogers, Stalnaker and others maintain that there are two kinds of effects, the first arising from the utterance itself and the second arising when it becomes evident that the utterance has been received without objection, its "uptake" [20, 24]. For an extreme example, one can say

(18) "Attack Hill 100."

and achieve certain effects thereby. An act of ordering has been performed. If the reply is

(19) "Why do you give me orders, Private?"

then the effects of having the order taken up do not occur, but the act of ordering is still recognized. The effects are directly parallel to the separate effects of bidding and accepting a Dialogue Game [17].

Although some Speech Act Theorists are considering effects, there is no reason to expect that SAT will ever provide an account at the level of detail required in NLP. Knowing the particular effects of a broad range of actions is peripheral for Philosophy of Language but central for NLP, so much of the required research must be done within NLP.

4.2 Recognizing Illocutionary Force

New work is also needed on recognition procedures. SAT cannot be utilized effectively in building language-reception processes until there are some reliable, explicit methods for identifying the illocutionary force (or forces) of a given utterance.

Some work has been done in SAT on the relation between syntax and illocutionary force, but (as we have seen) illocutionary force is not a syntactic phenomenon.

NLP requires that its methods be very explicit, but it allows some approximation. SAT, on the other hand, is unwilling to admit approximation, preferring to lose generality and have a few unassailable cases. So far, SAT has relied almost entirely on the reader as the recognizer of illocutionary force. However, to produce clear cases, some of the analysis has worked with sentences that contain an "illocutionary force indicating device" (Ifid), as, for example, in "I promise to pay \$1000 to the ABC Roof Company," where "promise" is the Ifid. I state that such sentences are relatively uncommon in ordinary text. Even this device is not infallible, as is neatly illustrated in the following widely-circulated example. Consider the sentence

(20) "I promise I will come to your party."

which seems necessarily to be a promise. But in conversation between A and B,

(21) A: "Why do I decide to give a party every weekend?"

(20) B: "I promise I will come to your party."

even the presence of the Ifid "I promise" is insufficient to produce a promise. However, the existence of the Ifid serves speech act theorists as a way to create examinable, publishable cases without having to address the general force-recognition problem.

The reliance on the reader as recognizer is especially clear with indirect speech acts. The reader must recognize that

(22) "Can you pass the salt?"

is often not a question about ability, and that

(23) "What do I know?"

functions as an assertion. For both direct and indirect speech acts, the more prominent question has been

"What kinds of speech acts are there?"

rather than

"How do I know that this speech act is of such and such a type?"⁹

There are several reasons to believe that SAT will not create an account of speech act recognition sufficiently complete and detailed for NLP. Although the definition of the term "illocutionary force" is part of SAT's domain, there has been little effort to make the definition explicit and operational by specifying a recognition procedure. The method that uses the reader as the recognizer of illocutionary force is satisfactory for present purposes and gains acceptance where any more explicit method would be subject to disputes.

More basic is the lack of any motivation toward comprehensiveness. Philosophy of Language specializes in a particular kind of facts and issues. If language comprehension requires more, then that is seen as being in another territory, most often linguistics, sociology, psychology, literary criticism, or rhetoric. As a result of this predisposition, SAT is not even testing the degree of coverage of its theories, even though they supposedly apply to "every utterance."

Finally, the notions of "context" are not sufficiently explicit, and no explicit use of "context" has been made in attempting to state how one recognizes a particular speech act.

NLP must therefore do its own development of speech act recognition procedures. In both of these cases we see that NLP research should include in its scope development of SAT, contributing to and extending it as well as adapting existing theory.

6. SOME SIMPLISTIC APPROXIMATIONS FOR FILLING THE GAPS

Before presenting particular processes to address the two gaps identified above, the role of these processes must be clear. In the tradition of all of AI, they are approximations to an unrealized ideal. As approximations, they are worthwhile to the degree that *they can be elaborated to yield satisfactory processes*. To be satisfactory, a process should accurately reflect the intuitions of native speakers and be capable of performing a useful role in a text processing task.

⁹ But see Heringer [13], Searle [22] pp. 60-64.

Although the processes sketched below are very rudimentary, we expect that more satisfactory algorithms would include similar kinds of processes as subparts. In the case of effects enumeration, we see the effects as a loosely coupled collection that would include the effects suggested here, along with others. In the case of illocutionary force recognition, we expect that the kind of analysis sketched here would be supplemented by methods for rejecting hypotheses of illocutionary force and by methods for generating alternative hypotheses after such rejections.

5.1 A Simple Method for Enumerating the Effects of a Speech Act

How can the effects of performing a speech act be enumerated simply? Consider a process that uses a knowledge base B of propositions representing H's knowledge about the speaker S and his world. Approximating the effect of a particular act F(P) performed by uttering U might include the following:

Adding to knowledge base B

1. that S believes (or otherwise conforms to) the sincerity, felicity and preparatory conditions of F;
2. that S claims the understandability of U, the truth of P, S's sincerity in uttering U, and S's right to perform F(P);
3. that S has performed F(P) by uttering U.

Such a process would correctly identify a significant portion of the implicit communication in ordinary utterances.

However, this approximation lacks some important features. It is unresponsive to the distinction between utterances that are "taken up" and those that are not. It also fails to add to B many of the effects of the hearer's acceptance of the utterance. Also, there is no provision at all for performing indirect speech acts.

Each of these faults would lead to frequent, significant errors on many kinds of text, including dialogue. The lack of mechanisms for producing the effects of handling indirect speech acts is particularly crucial, since people seem unable to limit themselves to using only direct speech acts.

Notice that these defects have a feature in common with those of the simple speech act recognizer: they cannot be overcome as long as only single sentences are analyzed. We can easily imagine corrections, but they all involve responding to the context (and not just the prior context) of the utterance, deciding how this particular utterance relates to the remainder of the discourse. For an alternative method in which response to context is used, see [17].

5.2 Toward a Simple Illocutionary Force Recognition Process

Creating an interesting simple process for recognizing illocutionary force is much more difficult, in spite of the fact that much more theoretical work has been done on illocutionary force than on effects enumeration per se. The source of difficulty seems to be simply that the indicators of illocutionary force are sometimes infrequent and usually unreliable. There are no syntactic attributes or particular usages of words or phrases to guarantee that a particular illocutionary force is present. Clearly the performative verbs are relevant when they occur, but they are infrequent, and their occurrence serves merely to constrain the selection of illocutionary force, not to determine it. Similarly, there are many syntactic attributes of utterances that, when present, serve to constrain the choice of illocutionary force.

We take this notion of *constraining the choice of illocutionary force* as the organizing idea below. The design of the process sketched below is based on formulating such constraints as rejection rules.¹⁰

There are several reasons for casting the illocutionary force recognizer in the form of a rejection rule processor rather than as a more conventional processor of positive evidence.

1. Positive indicators of illocutionary force (including both the syntactic form of the utterance and the occurrence of performative verbs) usually admit a range of interpretation rather than a unique one.

¹⁰ Note that this formulation is equivalent to one which recognizes an illocutionary force by *conformity to a set of rules*, by taking all of the rules capable of rejecting a particular illocutionary force as a set, and conformity as inapplicability. Alston uses such a formulation in his definition of illocutionary force [2, 8].

2. In extending the process to recognize indirect illocutionary force correctly, the crucial information is in

- a. whether or not every direct illocutionary force was ruled out, and
- b. how the direct illocutionary forces were ruled out.

Let U be the utterance to be analyzed and P the proposition expressed by U . Let $F = f_1, \dots, f_n$ be the set of distinct illocutionary forces; recognition of the illocutionary force of U is selection of one member of F . Let $T = t_1, \dots, t_m$ be a set of tests, i.e., a set of predicates applicable to U , its context and its situation.

We can define a *force rejection rule* $C_{ij}(T, U)$ as a logical function of T and U such that

$$\begin{aligned} C_{ij}(T, U) &= \text{true iff the values of } T_1(U), T_2(U), \dots, T_m(U) \text{ prohibit} \\ &\quad \text{interpretation of } U \text{ as having direct illocutionary force } f_j, \text{ and} \\ &= \text{false otherwise.} \end{aligned}$$

Then the recognition of illocutionary force yields as a set of plausible illocutionary forces of U the set of f_j such that

$$C_{ij}(T, U) = \text{false.}$$

The adequacy of the approach rests on being able to find enough valid constraints and of course on being able to determine the features on which they depend. Both this process and the previous one presume the availability of a particular hypothetical P , the expressed proposition, for determining values of relevant features.

In the examples below we will use a simple system of four illocutionary forces, called Representative, Request, Query, and Commissive. The set is not complete, and no claims are made for its empirical adequacy. The Representatives correspond to the Representatives in Searle's taxonomy, the Commissives (which include promises) to his Commissives, and Queries and Requests are the subsets of his Directives that direct verbal and nonverbal actions respectively [23].

Examples of tests T would include such things as

- t_1 : U is first-person.
- t_2 : P is the occurrence of a future event.

- t3: U has interrogative syntactic form.
- t4: U has imperative syntactic form.
- t5: U has declarative syntactic form.
- t6: P is the occurrence of a verbal act.

Consider the following sentences (labeled with forces on the basis of ordinary but imaginary contexts and situations).

- (24) Request: "Compute the square of 11."
- (25) Query: "Tell me your phone number."
- (26) Query: "What is your phone number?"
- (27) Representative: "I was at your party."
- (28) Representative or Commissive (ambiguous): "I will be at the picnic."

For each example, we would like to reject all of the forces except those indicated by the labels. We can do so with the following constraints:

<u>Rule</u>	<u>Test Expression</u>	<u>Forces Rejected</u>
C1	t3	Representative, Commissive
C2	t5	Request, Query
C3	t4	Representative, Commissive
C4	Not t1	Commissive
C5	Not t2	Commissive
C6	Not t3 and not t6	Query
C7	t3 or t6	Request

The tests succeed on the given sentences as follows:

<u>SENTENCE</u>	<u>TESTS</u>
(24)	t4, Not t2
(25)	t4, t6, Not t2
(26)	t3
(27)	t1, t5
(28)	t1, t2, t5

We can therefore form a table of the forces and the rules that reject them as follows:

	<u>Representv.</u>	<u>Request</u>	<u>Query</u>	<u>Commissive</u>
(24)	C3	-	C6	C3, C4, C5
(25)	C3	C4	-	C3, C4, C5
(26)	C1	C7	-	C1, C4, C5
(27)	-	C2	C2, C6	C5
(28)	-	C2	C2, C6	-

As indicated by the rejection table, these constraints are sufficient to select the forces indicated above, and the tests on which they rely are tests of factors often cited in the speech act literature.¹¹

In considering extension of the scheme so that it can recognize the forces of indirect speech acts, other sorts of constraints are of interest. For recognizing indirect uses of questions, it is useful to test whether the speaker knows the answer. Similarly, for irony it is useful to test whether the speaker is known to believe not P. These tests can be used in constraints. The applicability of such constraints can be used to trigger creation of indirect speech act hypotheses. A scheme for doing so is described in [17].

The quality of such a procedure can be assessed on several grounds:

1. The frequency with which the derived set of plausible illocutionary force includes the correct one(s).¹²
2. The frequency with which the derived set of plausible illocutionary force includes incorrect ones.
3. The frequency with which the correct illocutionary force is constrained away.

¹¹ For homogeneity and ease of explanation, all of the rules are formulated as constraints above; it would be more efficient, and perhaps more realistic from a psychological point of view, to recast part of the scheme (perhaps a syntactic part) as a generator of hypotheses to which the remaining constraints were applied.

¹² Correctness here is simply agreement with an independently determined strong consensus of native speakers.

In the case of an indirect speech act, we expect that *all* of the direct illocutionary forces will be constrained away. In this case, we can judge the appropriateness of the manner in which the direct force was constrained away.¹³

The force recognition process described above does not make any special use of the "performative verbs" so prominent in SAT but relatively infrequent in ordinary speech. How difficult would it be to extend the process to deal with such cases [10]?

There is an intimate relationship between use of performative verbs and the use of embedded sentences. If one says in certain circumstances,

(29) "I warn you, that bull is dangerous."

the speech act of the utterance is the same as the speech act of saying

(30) "That bull is dangerous."

The phrase "I warn you" may serve to clarify or emphasize the act of warning, but warning is performed in either case. A recognition scheme could not rely on the presence of "warn" and similar verbs to identify cases of warning.

Many of these verbs are bizarre in such roles.

(31) "I boast that I know a lot about stocks."

(32) "I insult you that you are ugly."

The class of such verbs as a whole is better treated as a set of descriptive verbs applicable to utterances. As such they are applied both to parts of the utterances in which they occur and to separate utterances.

A necessary precondition for any sensible use of descriptions of utterances is a capability for representing propositions about utterances. This is a general requirement for language understanding; it goes far beyond the case of performative verbs, and includes the representation of believing, doubting, suggesting, regretting, hoping, denying, and claiming. It includes suitable treatment of embedded sentences and of quoted text.

¹³ If all of the direct forces are constrained away, but the utterance is not interpreted as an indirect speech act, then it is simply not comprehended. SAT includes some discussion of "defective" speech acts, but their status is not clear.

Design of processes for analyzing the effects of performative verbs is best left for a processing environment in which the methods for dealing with utterance descriptions are suitably elaborated. In such an environment, treatment of performative verbs might not present special problems at all.

6. CONCLUSIONS

The concepts of Speech Act Theory could contribute significantly to the progress of natural language processing. Although the prevailing style in Philosophy of Language is very different from that in AI, the action-oriented conceptual framework of Speech Act Theory is reasonably compatible with actively developing AI concepts. Natural language processing stands to benefit not only from the work of speech act theorists but also from other disciplines using the same concepts.

Processes that identify the illocutionary forces in text can be used in deriving implicitly communicated information and in relating text to its context and situational information. Speech act concepts are particularly attractive as a means for relating knowledge from text to knowledge about the speaker.

To take advantage of the opportunity, natural language processing must familiarize itself with Speech Act Theory, rerepresent the knowledge in process-oriented terms, and actively work to fill in some of the gaps. The results could include both a new version of the theory and a substantial expansion of natural language processing capabilities.

REFERENCES

1. Alston, William P., "Linguistic acts," *American Philosophical Quarterly* 1 (2), April 1964, 138-146.
2. Alston, William P., *Philosophy of Language*, Prentice-Hall, 1964 (in the *Prentice-Hall Foundations of Philosophy* series, Elisabeth and Monroe Beardsley (eds.)).
3. Austin, J. L., *How to Do Things with Words*, Harvard University Press, 1962.
4. Bullwinkle, Candace, "Levels of complexity in discourse for anaphora disambiguation and speech act interpretation," in *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, August 1977.
5. Clark, H. H., and P. Lucy, "Understanding what is meant from what is said: A study in conversationally conveyed requests," *Journal of Verbal Learning and Verbal Behavior* 14, 1975, 56-72.
6. Cohen, Philip R., and Raymond C. Perrault, "Overview of 'planning speech acts'," in *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, August 1977.
7. Cole, P., and J. L. Morgan (eds.), *Syntax and Semantics. Volume 3: Speech Acts*, Academic Press, 1975.
8. Coyne, Margaret U., *A Critical Study of Two Models for Speech Act Analysis*, Ph.D. thesis, Northwestern University, August 1975.
9. Davison, Alice, "Indirect speech acts and what to do with them," in Cole, P., and J. L. Morgan (eds.), *Syntax and Semantics. Volume 3: Speech Acts*, pp. 143-185, Academic Press, 1975.
10. Fraser, Bruce, "Hedged performatives," in Cole, P., and J. L. Morgan (eds.), *Syntax and Semantics. Volume 3: Speech Acts*, pp. 187-210, Academic Press, 1975.
11. Habermas, Jurgen, "Towards a theory of communicative competence," *Inquiry* 13, 1970, 360-375.
12. Habermas, Jurgen, *Knowledge and Human Interests [Erkenntnis und Interesse]*, translated by J. Shapiro, Beacon Press, Boston, 1971.
13. Heringer, J., "Some grammatical correlates of felicity conditions and presuppositions," *Working Papers in Linguistics* No. 11, Columbus: The Ohio State University, Department of Linguistics, 1972.

14. Horrigan, Mary Katherine, "Modelling simple dialogs," in *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, August 1977.
15. Lee, Patricia Ann, *Impositive Speech Acts*, Ph.D. thesis, Ohio State University, 1974.
16. Lesser, V. R., and L. D. Erman, "A retrospective view of the Hearsay-II architecture," in *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, pp. 790-800, 1977.
17. Mann, William C., A characterization of language comprehension (in preparation).
18. McCarthy, T. A., "A theory of communicative competence," *Philosophy of the Social Sciences* 3, 1973, 135-156.
19. Munro, Allen, *Speech Act Understanding in Context*, Ph.D. thesis, University of California, San Diego, 1977.
20. Rogers, Andy, "Remarks on the analysis of assertion and the conversational role of speech acts," in *Proceedings of the Fourth Annual Meeting of the Berkeley Linguistics Society*, Berkeley, Calif., 1978.
21. Searle, J. R., *Speech Acts: An Essay in the Philosophy of Language*, Cambridge University Press, 1969.
22. Searle, J. R., "Indirect speech acts," in Cole, P., and J. L. Morgan (eds.), *Syntax and Semantics*. Volume 3: *Speech Acts*, pp. 59-82, Academic Press, 1975.
23. Searle, J. R., "Taxonomy of illocutionary acts," in Gunderson, K. (ed.), *Language, Mind and Knowledge*, pp. 344-369, University of Minnesota Press, Minneapolis, 1975.
24. Stalnaker, Robert C., "Assertion," in Cole, P. (ed.), *Syntax and Semantics*. Volume 9: *Pragmatics*, pp. 315-332, Academic Press, 1978.
25. Wunderlich, D., "Pragmatik, Sprechsituation und Deixis," in *Beiträge zur Literaturwissenschaft und Linguistik*, Bad Homburg, 1971.