INTRODUCTION

The problem of explicating the meaning of natural language discourse and how this meaning is obtained has been of central importance to the study of philosophy since the time of Aristotle. This problem has continued to perplex philosophers, linguists, psychologists, and other academicians right up to this very day, such perplexity manifesting itself in disputes such as nominalism *vs.* realism, behaviorism *vs.* mentalism, logical positivism *vs.* the "ordinary-language" approach, and structuralism *vs.* transformationalism. Basic to this perplexity has been the lack of an adequate set of tools with which to formalize meaning in natural language, resulting in a basic split between those who use existing tools to produce formalizations that are oversimplified and inadequate and those who attempt to account for the full complexity of natural language but, in doing so, abandon any attempt at formalization. Clearly, any move toward providing a more powerful set of tools with which to formalize the semantics of natural languages will have implications that will reverberate through a variety of academic disciplines.

A formalization of natural language semantics that is keyed to the computer will have a variety of very practical applications as well. Two decades of research in machine translation have failed to produce effective systems, chiefly because machines have not yet been able to duplicate the translator's function of understanding the material in the source language and then restating it in the target language. Once computers are capable of analyzing natural language to a depth where it becomes possible to mechanically determine equivalence of meaning in context of statements in two different languages, and capable of generating well-formed natural-language discourse from such an analysis, mechanical translation of scientific and other expository text will be straightforward. Information storage and retrieval will also be much simpler once it is possible to store and retrieve information by means of statements, questions, and commands in natural language--which requires, at least, that these natural-language inputs be interpretable as commands to the storage and retrieval mechanism. Computer-assisted instruction will not realize its full potential until student answers can be evaluated with respect to their meaningful content and this evaluation used to generate appropriate remedial instruction form a body of lesson information. And most importantly, a capability for meaningful analysis and generation of natural language will make computers accessible to people for an infinite variety of problem-solving applications in the way which is best suited to them as human beings.

The thesis of this paper is that, as a result of recent advances of the state of the art in linguistics, computation, and their interface, we are on the verge of being able to formalize the semantics of natural languages for the computer in a manner that will be philosophically interesting, linguistically and psychologically revealing, and computationally useful in the manners just suggested. Such a formalization, it shall be argued here, may be based on the notion that natural language is basically explicable as a method of programming a particular kind of computer--a computer with a relational associative memory structure that is purposive and goal-directed in its actions, in the manner of a human being.

THE PROBLEM

A. Requirements for a Semantic Theory

1. A Definition of Semantics

The term semantics is generally used to denote the system of relations between the expressions of a language and their meanings--in contrast to syntax, which describes the acceptable structural forms of linguistic expression, and pragmatics, which concerns itself with the effects of communications in a language upon the communicants. The term "meaning" here is to be analyzed in terms of its particular relevance to the act of communication. An act of communication includes a sender who encodes a message into a signal (usually a string of phonetic or alphanumeric symbols), a channel along which the signal is sent, and a receiver who decodes the signal into the original message and interprets the message. Meaning is the functional import of the message, which calls forth a response--cognitive, affective, and/or conative--from the receiver in the particular communication situation; it is a relation of the message to the receiver. Let us now make the simplifying assumption that this meaning is determined by the functional form of the message under some "standard" functional interpretation (this is basically an oversimplification, since the rules of functional interpretation, especially on the affective response dimension, will vary from receiver to receiver and even for the same receiver over time); we may then define semantics as the explication of the relationship between the surface forms of linguistic utterances (as spatial or temporal arrangements of morphological units, which are the minimal linguistic units that have a direct functional relationship to the determination of meaning) and the functional forms of the messages they express. For a semantics to be formalized, this relationship must be sufficiently well defined to enable the processes of encoding and decoding to at least be formally explicable in terms of it, if not totally formalized within it.

2. The Semantics of Formalized Languages

The most successful attempts to formalize semantics have been for the formalized languages of the deductive sciences. Tarski, in his classic paper [43], has set down a systematic method for formalizing the semantics of a formalized language, which is any language for which the set of meaningful sentences is defined by a formal syntactic grammar, all sentences are unambiguous, and there is a set of (syntactically defined) axioms and rules of inference from which theorems in the language may be derived. The method involves construction of a metalanguage containing expressions and axioms of a general logical kind, translations of the expressions and axioms of the language L to be characterized, and expressions and axioms which define the syntax of L. The notions of satisfaction and truth for the language L, which together constitute a semantic description of L, can be defined in the metalanguage M as follows: Given a domain $\mathcal{D}$ of individuals, a "semantic interpretation function" ϕ is defined in M which assigns to each individual constant and individual variable of L an individual of the domain $\mathcal{D}$, to each function letter of L of degree $n \geq 1$ a function from $\mathcal{D}^n$ (the set of all ordered n-tuples of elements of $\mathcal{D}$) into $\mathcal{D}$, to each predicate letter of L of degree $n \geq 1$ an n-ary relation on $\mathcal{D}$ (defined as a subset of $\mathcal{D}^n$), and to each phrase-forming rule of L a function τ_i, which determines the semantic interpretation of the phrase formed in terms of the semantic interpretations of the constituent phrases. An n-tuple $\langle a_1, \ldots, a_n \rangle$ of individual constants then satisfies the sentential function $\mathcal{F}$ of n free variables whenever $\phi(\mathcal{F})\ [\phi(a_1), \ldots, \phi(a_n)] = T$ (truth). The notion of truth is defined as a special case of the notion of satisfaction for sentential functions of zero free variables (i.e., sentences of L). Examples of the application of the Tarski approach are to be found in any standard logic textbook in the truth-table interpretations (standard and non-standard) of the propositional calculus and the standard set of semantical rules for the first-order predicate calculus.

The Tarski approach runs into difficulty, however, in that it establishes only a single notion of truth, without reference to the different ways in which truth may be epistemologically established. This becomes critical when it comes to establishing rules of substitutability for "oblique" or "nonextensional" contexts such as quotation, indirect discourse, modal sentences, and belief sentences, in which the unrestricted substitutivity of equivalence does not preserve truth-value--one must here define types of equivalence stronger than identity of reference, which is the type of semantic equivalence defined in the Tarski approach. It was mainly to deal with this problem that Carnap [9] proposed a method of semantic analysis called "the method of extension and intension." In the framework of Tarski's formulation, the method can be stated as follows: Given a "model" of the language L, consisting of an individual domain $\mathcal{D}$ and a semantic interpretation function Φ for L over $\mathcal{D}$, the <u>extension</u> of any well-formed expression $\mathcal{E}$ in L is defined as the set of values for $\mathcal{E}$ of all semantic interpretation functions Φ' over $\mathcal{D}$ which differ from Φ at most on their assignments to the free variables of $\mathcal{E}$. Now if one considers the domain of possible models for L, the <u>intension</u> of any well-formed expression $\mathcal{E}$ in L may be defined as that function over models of L which yields as its value for any model the extension of $\mathcal{E}$ in that model. The notion of intension may be formalized by considering the metalanguage translations of well-formed expressions of L to be <u>intensional structures</u> for these expressions, which, since intensions are functions, will take the form of function definitions. The translations may be defined by a translation function θ which assigns to each individual constant, function letter, and predicate letter of L an appropriate function letter of the metalanguage M, to each variable of L a variable of M ranging over functions on models of L which map into the appropriate extensional range, and to each phrase-forming rule of L a function-definition operator σ_i (which could be functional composition, complement, union, intersection, iteration, transitive closure, summation, minimalization, etc.). Given this definition, one can recreate Tarski's definition of satisfaction and truth by noting that for any model M_i of L the corresponding semantic interpretation function ϕ_i is given by $\phi_i\ (E) = [\theta(E)]\ (M_i)$. And one can define, along with the ordinary notion of (extensional) equivalence, the notions of L-equivalence and intensional isomorphism as equivalence of intension and intensional structure respectively--and show, as Carnap does in [9], how these stronger types of equivalence permit the establishment of suitable substitutability criteria for "oblique" contexts.

To fit the semantics of formalized languages into our general definition of semantics, which presupposes use of the language for the purpose of communication, we must introduce one more thing into the metalanguage, namely, the performative operators of asserting, questioning, and commanding. We posit that the language L is being used to communicate between two information and control systems A and B, both of which possess incomplete and/or changing models of L[2] corresponding to knowledge of some environmental situation over which both A and B can exercise

[2] Note that the reference to disparate models of L (the only situation in which communication would make sense in this context) and changing models of L necessitates the use, at least implicitly, of an intensional semantics.

certain degrees of control. Then for any sentence S of L, an assertion 'S.' from A to B carries the functional import of instructing B to modify its model so as to make S evaluate to truth, a question 'S?' from A to B instructs B to evaluate S in its model and return the result to A, and a command 'S!' from A to B instructs B to modify its environment (if possible and if necessary) so that S will evaluate to truth in the model of the environment so changed. The metalanguage translations of the assertion, question, and command signs in L will be, of course, the corresponding performative operators. By identifying, now, the notion of message with metalanguage translation, the notion of arrangement of morphological units with syntactic description in the metalanguage, the notion of decoding with the translation function θ, and the notion of encoding with the inverse of θ, we show how the semantics of formalized languages meets our general requirements for a formalization of semantics.

3. Natural Language Semantics

We may now arrive at a set of specific requirements for a formal theory of natural language semantics by examining the crucial differences that are known to exist between natural languages and formalized languages and noting the revisions and extensions of the formalized-language paradigm that are required to take these differences into account. This approach is indicated by the fact that the semantics of formalized languages represents the most highly-developed point of departure from which to undertake a formal description of the real-world phenomenon of natural language semantics, and thus, if it indeed contains the potential of producing a description that fits the phenomenon, brings one much closer to that description than if one were to start with only the general definition of semantics given at the beginning of this section. What, then, does this approach indicate for the features of a revised paradigm under which the known properties of natural language as an instrument of communication may be subsumed?

Natural languages, first of all, are used for a vastly wider variety of communication acts than are formalized languages. The messages that are communicated in natural language relate to virtually every area of human activity and extend to nearly every purpose involving some kind of human interaction. As a result, there is a large inventory of different types of messages that are expressed in natural language, each in its own particular way or ways.[3] The three basic types of performative operators--assertion, questioning, and commanding--are subject to modification as to the manner of the request conveyed by the message (which indicates, among other things, the speaker's perceived or intended relation to the hearer[4]) and to functional combination, as in the case of threatening and warning, both of which combine commanding with asserting. A formal theory of natural language semantics must explicate these dimensions of the performative and also the relation of the performative to the notions of speaker, hearer, and context of utterance.

The explication of the non-performative parts of messages as intensional definitions requires some extension and elaboration

---

[3]Austin [1] has compiled what is perhaps the most extensive and systematized inventory of these message types.

[4]A very interesting discussion of this aspect of communication is contained in Watzlawick et al [46].

in order to be applicable to natural language. First, people carry in their memories not one model but many, corresponding to the many different situations that they have knowledge of. Thus, a message must refer either to a specific model, to a specific range of models, or generically to all models in which the specified intensions have nonempty extensions. Restricting the range of applicable models is accomplished in natural language through presuppositions, which indicate prior conditions that must be satisfied in a model for a given message to be applicable to it. Indicating generic vs. model-specific information, as well as "given" (for locating the appropriate model) vs. "new" (for adding to the model) information in the model-specific case, is accomplished through the subject-predicate division and through an extensive assortment of quantifiers. Furthermore, as Morris [31, 32] has pointed out, natural-language expressions not only designate but also appraise and prescribe--thus, natural-language intensions may take on as extensions not only objects, sets, and relations, but also values and actions. Natural-language intensions may also take on as extensions other intensions, thus giving natural language a "recursiveness" of logical order and a self-referential capability (which leads, naturally, to the classical logical paradoxes).

Intensional definitions are also more complex in their formal structure for natural languages than for formalized languages. Intensions may be defined by specifying the combination of tests and results that indicate which elements of any given model are to be included in their extensions--these tests may be on either "inherent" or "contextual" attributes of the element and the values of these attributes may be either countable sets or measurements on some continuous scale. For formalized languages, the identification function of an intension must distinguish clearly and unequivocally between exemplars and non-exemplars on the basis of a Boolean combination of the results fo the various tests. For natural languages, however, tests may be either criterial for identification of an exemplar or else have only a probabilistic bearing on identification; thus, the identification of exemplars of natural-language intensions is by no means clear-cut, but rather may resemble the assignment of degrees of confirmation to hypotheses (with a certain "level of confidence" being required for identification to take place). The use in English of generic determiners such as 'many', 'most', 'almost all', and 'few', and (corresponding) intensional adverbs such as 'commonly', 'usually', 'characteristically', and 'seldom', is indicative of the probabilistic nature of intensional definitions in natural language.[5]

The morphological structure of natural languages is also considerably more complex than that of formalized languages, as has been well recognized by contemporary linguists. The simple phrase-structure grammars that suffice to describe the syntax of formalized languages simply do not work for natural languages; to describe the surface syntactic structure of a natural language requires a system, such as a relational phrase-structure grammar (Bellert, [4]) or a complex-feature-symbol grammar, with the power of expressing the various relations of grammatical agreement among constituents. If the language to be analyzed is spoken language, the arrangements of morphological units are

---

[5]A full analysis of these generic determiners and adverbs, their logical interrelationships, and their relation to notions of probability is given in Celce and Schwarcz [13]. A capsule summary of this analysis is presented later in this paper.

not simply linear strings of symbols (connected by whatever grammatical relations) but are, rather, two-dimensional sequences consisting of both segmental and suprasegmental (stress, intonation, etc.) morphemes. Furthermore, the exact correspondence between grammatical sentences and semantically-interpretable sentences that obtains in formalized languages does not hold for natural languages, which permit of both "grammatical nonsense" and syntactically deviant utterances that make perfect sense--the first phenomenon requires a semantic theory to posit nonsyntactic conditions for semantic acceptability; the second, a procedure for syntactic error correction in decoding.

The above are only two of the phenomena that render explication of the process of encoding and decoding much more complex for natural languages than for formalized languages. In formalized languages all well-formed expressions are unambiguously interpretable in or out of context, their interpretations are determined in a straightforward compositional manner by function-definition operators in one-to-one correspondence with the syntactic formation rules of the language, and performatives are represented as single symbols preceding or following each sentence. For natural language none of these properties hold--indeed, semantic ambiguity and anomaly, discourse structure and other forms of context dependence, syntactic-semantic non-correspondence, idioms and figures of speech, and complex encodings of performatives are all common features of natural language. Their explication in a semantic theory requires, first, that the correspondence between syntactic form and semantic function be taken as many-to-many rather than as one-to-one; second, that intensional "well-formedness" relative to the particular domain of discourse and applicability to the model or range of models currently under consideration be taken as criteria for semantic acceptability in a discourse context; and third, that the theory specify the various alternative encodings of a message rather than a single encoding. There is also a need to incorporate analogical processes into the explications of encoding and decoding in order to account for the metaphorical use of language.

Finally, there are two inherent limitations that govern any attempt to formalize the semantics of natural languages: one formal, the other epistemological. The formal limitation is a consequence of Tarski's theorem [43], which states that any consistent and complete semantic theory of a language must be formulated in a metalanguage of higher order than the language being described. But since the set of theorems of any deductive system must be recursively enumerable, and since there are subsets of natural languages sufficiently powerful to define any recursively enumerable set, any formalization of natural language semantics using a deductive logic (including the logic of computation) as a metalanguage will be incomplete in the sense that there will be questions about the language[6] that are theoretically unanswerable in the metalanguage (one could, however, go to inductive logics and probabilistic metatheories as the basis for a metalanguage). The epistemological limitation derives from the fact that, while formalized languages are uniquely defined, no two speakers of a natural language have quite the same idea of what their language is. It is clearly impossible, then, to formulate a semantic theory that describes <u>all</u> the speakers of a natural language.

[6]Including, of course, any question as to the truth-value of a sentence expressing a logical paradox.

Neither is it practicable to attempt an "ideal speaker-hearer" theory that purports to explain how native speakers of a language "generally" assign meanings to utterances and express meanings through utterances, since validation of such a theory would be next to impossible. A more appropriate goal, especially in light of the fact that the data for any semantic theory must ultimately derive from the use of the language for communication, is to construct a theory of a "typical speaker-hearer" of the language in question, whose validity would then derive from the ability of a physical realization of the theory (e.g., as a program running on a digital computer) to engage in successful purposive communication with native speakers of the language.[7]

Let us enumerate, then, the requirements for a formal theory of natural language semantics that have been indicated here:

1. The theory shall be couched in a formal metalanguage.
2. The metalanguage shall contain models of possible discourse contexts, expressions representing extensions, expressions representing intensions, and axioms defining the relation of extension to intension for any given model.
3. The metalanguage shall contain expressions representing the messages communicated in the natural language, which will contain performatives specified as to type and manner, intensional definitions of both fixed and recursive logical order with criterial and/or noncriterial components on the descriptive, appraisive, and prescriptive dimensions, presuppositions, and both generic and specific quantifiers.

[7] Further reasons for preferring the "typical speaker-hearer" model to the "ideal speaker-hearer" model are given in Schwarcz [38].

4. The metalanguage shall contain axioms characterizing the functional import of messages, sufficient to define both extensional and intensional equivalence, entailment, and contradiction among messages up to the limits of theoretical decidability.
5. The metalanguage shall contain expressions and axioms defining a "standard" syntax of the language at the level of surface arrangements of morphological units, in terms of both phrase structure and relations of grammatical agreement.
6. The metalanguage shall contain axioms defining the possible encodings of any message in any discourse context to which it is applicable.
7. The metalanguage shall contain axioms defining the possible decodings of arrangements of morphological units that are well-formed in the "standard" syntax or deviate from it by at most a tolerable degree and determining the intensional well-formedness and applicability to a given discourse context of these decodings.
8. The theory, to be validated as a description of a "typical speaker-hearer" of the language under consideration, must support a physical embodiment that is capable of engaging successfully in purposive communication with native speakers of the language.

B. Computational Avenues of Approach to a Semantic Theory

Since language is an instrument of communication and communication is essentially purposive, any semantic theory that one develops for the computer will of necessity be based, unless one is simply engaged in an academic exercise, on the purpose for which one wishes to communicate with the computer in natural language. In this section several such purposes and the sorts of semantic theory they have led to or are likely to lead to will be described.

The purpose of oldest vintage is, of course, translation by machine from one language to another. The problem here is, given a discourse in one language, to produce a discourse in a second language that has the same functional import with respect to a model of the domain of discourse as the first. Perhaps the reason that no efforts in this direction have achieved notable success to date is that the model of the domain of discourse and its functional interaction with the language have generally been ignored in the design of translation systems. The direction that will lead to a breakthrough here is that of developing domain-specific (rather than language-specific) translation systems for well-understood and formally structurable domains of discourse such as physics and mathematics--once a formal model of the subject matter and a canonical procedural language for communicating with that model are defined, efforts can be directed toward specifying the decodings of as much of the relevant natural-language subsets as possible into the procedural language and reasonable encodings of the procedural language into each of the natural languages.

Data management and information retrieval is another purpose of widespread interest. The domains of discourse to which these systems may apply may be arbitrarily broad or narrow; whatever the case, the requirements for formal structurability and a formal procedural language for storing and retrieving information in the data base are present. If the system is to do deductive question answering (or what Travis [45] has called "analytic information retrieval"), the system must be able to store and utilize the logical relationships among concepts and facts. The problem of specifying encodings and decodings here is simpler than for machine translation, since the user may make do with a fairly restricted natural-language subset for input, and natural-language output may be generated in a canonical form if it is in fact necessary at all. Thus it is possible here to get by with an oversimplified semantic theory, but for that very reason it can be expected that more progress can be made sooner with this than with any other approach (and this has, in fact, turned out to be the case).

Another purpose is the use of natural language to interface with pictorial information. Here the model is a set of logical statements describing the visual image, derived by the application of pattern-recognition operations to the visual image. The model, once derived, can then be either directly encoded into a set of natural-language sentences or else used as a data base for information retrieval. An alternative approach is to decode natural-language retrieval statements into search procedures on the visual image itself, performing the pattern-recognition operations, then, during the execution of these search procedures. If the visual image is what a robot sees in its environment, the robot may not only be asked about what it sees but also told to move about in its environment and to move parts of the environment about; thus, the intensional structure of the robot's message language will include a prescriptive as well as a descriptive dimension. As in the case of data management and information retrieval systems, the input language can be restricted and the output language can be minimal, thus obviating the need for sophisticated formulations of decoding and encoding.

The use of the computer to develop models of human thought processes is a purpose that can lend revealing insights into the nature of a semantic theory. Here one starts with hypotheses about the structure of human memory and the information processes that take place there, embeds these hypotheses in a computer program, and runs the program to determine the consequences of these hypotheses in terms of predictions of observable behavior. In terms of a semantic theory, the emphasis here is likely to focus on models, messages, and the pragmatic functions of messages on models; only limited attention is likely to be paid to the syntactic structure of the language, and encoding and decoding are likely to be formulated in a rough-and-ready heuristic fashion rather than in a way motivated by linguistic considerations.[8] Nevertheless, such models are an excellent way to test the workability of semantic ideas, for the models' linguistic poverty is compensated for in experimentation by their designers' linguistic flexibility--and once the innards are working right, they may serve as a basis for a more linguistically sophisticated formulation of decoding and encoding.

A purpose incorporating both analytic information retrieval and psychological modeling is computer-aided instruction with natural language.[9] The capabilities required here are to semantically analyze a student's natural-language response or question, to compare an analyzed response to a standard "correct response" to determine the logical difference if any, to generate remedial feedback in natural language by application of "tutorial decision rules" to the structure representing this difference,

---

[8]The one exception is models of linguistic performance, as discussed in Schwarcz [38]; there, of course, linguistic considerations are paramount from the beginning.

[9]This approach to CAI is described in Bennik, Schwarcz, and Silberman [6].

and to answer a student's analyzed question and generate a natural-language reply. For natural-language CAI all the components of a semantic theory, except perhaps for encoding, must be developed to their full extent with respect to the subject areas to be taught. The linguistic requirements are not quite so severe as for machine translation, since the capability of dynamic interaction enables students to put up with a certain amount of rigidity on the machine's part and since the machine will not be required to analyze or generate long coherent discourses, but the requirement of thorough and complete logical analysis is more demanding here than in any other application of a semantic theory.

Finally, there is the purpose of enabling people to program the computer in natural language. Messages here are statements in a general-purpose programming language which includes capabilities for defining both macros and closed subroutines; they will thus have both descriptive and prescriptive dimensions. Nouns, verbs, and adjectives will be decoded into either data items (if proper names, numbers, or truth values), primitive functions, macros, or closed subroutines, conjunctions and prepositions will decode into operators for combining program steps, adverbs will decode into designations of program sequencing, and quantifiers will decode into specifications for iterative loops. The decoding process will likely be some form of syntax-directed compiling, which exactly fits the decoding paradigm for formalized-language semantics, except in that the procedure may allow for a small degree of ambiguity. Encoding will either be completely standardized or else be defined in terms of a sublanguage of output specifications that may be associated in an arbitrary manner with computational procedures. All this assumes, however, that natural language is being used to program the computer in for small western cities?' produced a ten-item request, and the question 'For the smoggy high-income cities what is the age-income value-range?' produced twenty separate procedural requests. Although most of Kellogg's experimentation has been performed on a data base of census information, his system has also been successfully demonstrated with airline-schedule and educational data bases.

If Kellogg's system can be faulted as a semantic theory, other than in its lack of a nontrivial formulation of encoding, it is principally in its failure to deal with certain of the requirements specific to the semantics of natural languages. Chief in importance among these are noncriterial attributes of intensions (except those quantified by 'some'), recursiveness of logical order, the appraisive dimension of language, discourse structure recognition, disambiguation by discourse context, and deviations from standard syntax.[14] The logic of equivalence, entailment, and contradiction among messages, particularly on the intensional side, has also not been formalized to the extent that it could be. In all fairness, however, it must be pointed out that few if any of the other current approaches to semantics have dealt with any of these requirements (except the last, for predicate-calculus-based systems) in a formally satisfying way. Kellogg has succeeded in putting together the best of current knowledge in linguistics, formal semantics, and systems programming to develop an eminently usable formalization of English semantics for the computer.

Both the linguistic and the computational formalizations of natural language semantics, when looked at individually, can be seen to fall considerably short of the requirements for a

[14] This last item, as well as undefined words and semantic anomalies, is handled by Kellogg through appropriate feedback messages to the user.

semantic theory that is adequate for natural languages. When taken collectively, however, they contribute an enormous reservoir of ideas and experience upon which one may draw in undertaking the formulation of an adequate semantic theory. With the addition of recent advances in linguistic theory, programming languages, and artificial intelligence to this reservoir, we may draw from it the elements that will combine to make up an adequate approach. Let us now look at one possible such approach.

AN OPERATIONAL-MEANING APPROACH TO SEMANTICS

A. Methodological Basis

To arrive at a formal theory of natural language semantics, we must start from the set of requirements enumerated earlier--particularly those concerning the relation of a message to its functional import. Intensions are the principal components of messages, and they are classified according to their values along the descriptive, appraisive, and prescriptive dimensions. Prescriptive intensions have values which are actions of the communicating system, and therefore can be sensibly regarded only as <u>programs</u> for action. Appraisive intensions have values which are evaluations of one kind or another; the only sensible way to regard these, then, is as <u>evaluation functions</u>. Descriptive intensions have values which are objects, sets of objects, and relations among objects, where the objects may in turn be intensions. Here we adopt the operationist philosophy, in the formulation of Benjamin [5], and assert that descriptive intensions are <u>functional operations</u> on a data space which yield elements of knowledge as their result.

It is clear, then, that the most natural representation of intensions is as programs is some programming language. Since intensions are functions on models, the operations that constitute them will be performed on structures that represent models--and since intensions may be values of intensions, the structure of programs must be of the same form as the structures of models. Furthermore, the programming system which interprets the language must interpret it nondeterministically, for natural language may always specify alternative definitions of a concept, or alternative procedures for evaluation, or alternative ways to perform an action; with respect to the first, it is a fundamental premise of modern operationism that one can arrive at the same item of knowledge by means of different operational procedures, and that, in fact, the utility of a concept is largely as an expression of the generalization that a class of different operational procedures produce identical results.[15] Nondeterministic operation and the existence of evaluation functions characterizes a class of artificial-intelligence programs that have been written to do game playing, theorem proving, and general problem solving, all of which are based on the paradigm method of goal-directed heuristic search. A programming system based on this method of program operation, similar to the one that Pople [34] has recently implemented, would thus be indicated as the basis for an operational formalization of natural language semantics. Let us now turn to a sketch of how the semantics of natural languages might be formalized within such a system.

[15] This view is expressed clearly in Bridgman [8].

B. Models and Messages

There are two basic issues to be decided in the formulation of any model: what information is to be contained in the model, and in what form that information is to be represented. In the formulation of a message language for communicating with the model, it must furthermore be decided what computational processes are to be performed in the model. A semantic theory will rise or fall on the basis of the extent of information that can be represented in the model, the extent of information processing that can be formulated in the message language, and the ease with which translation algorithms can be formulated between the message language and the corresponding natural language subset.

Because of its close similarity to both formal logic and the attribute-value list structures and relational associative structures that have been employed in many artificial-intelligence programs, as well as its demonstrated advantages for linguistic formulations, the Fillmore case structure appears to be the most useful starting point for representing both models and messages. Additional specifications must be added in to represent the logical features which are lacking in Fillmore's formulation: logical connectives, quantification and quantificational ordering, other function-definition operators, the structure of the modality constituent, etc. The inventory of case relations must also be completely specified and, since case relations are all contextual, supplemented with a set of inherent relations that will enter into both extensional and intensional description--some of which, like the "spatially contains" relation, will be converses of the case relations themselves.

The formal content of both intensional and extensional description is still largely an open question, to which the various attempts to formalize natural language semantics can only suggest methods for solution. At the lowest level of semantic description, the actions that an operational semantic model will be able to perform will be computer actions and not human actions, and the evaluations that it will be able to make will almost certainly be pragmatic evaluations rather than aesthetic evaluations (this is not to assert, however, that no way will ever be found to program a computer to simulate a human being's appreciation of poetry, art, or music). It is the descriptive dimension that is most interesting, especially since both appraisive and prescriptive intensions are instances of second-order (or higher) descriptive intensions, with the consequence that the values and actions of human beings may be described in an operational semantic framework, and perhaps as a consequence also modeled by analogy though not applied directly. On the descriptive dimension Benjamin [5] lists the following types of operations, which we shall characterize as intensions with corresponding extensions:

Intension	Extension
1. Discriminating	Deictic references (present events)
2. Associating	
a. Co-occurrence	Co-occurrence classes and relations
b. Temporal succession	Temporal and causal relations; durable objects and states
c. Configurational	Part-whole relations; Gestalts

3.	Generalizing	Supersets and class-inclusion relations
4.	Ordering	Partial and total orderings (including enumerations)
5.	Measuring	Numerical values
6.	Analogizing and disanalogizing	Analogies, icons; models

All these operations may be combined, of course, by function-definition operators in defining intensions. Models may be defined in this context in terms of <u>situations</u>, which are hierarchically structured configurations of events where the elements of each level of the hierarchy are connected by relations of co-occurrence and temporal succession, and different levels are connected by part-whole relations. With each node of a situational hierarchy will be associated that extensional information which applies (inherently or contextually) to all events below it. Some situations will be associated with <u>goals</u> of the communicating system, which are intensional descriptions for which the communicating system seeks to transform the situation in order to satisfy. These goals form the basis for the operation of programs in the system.

The criteriality or noncriteriality of intensional attributes may be represented by associating with each attribution a quantifier; in an intensional definition these quantifiers represent levels of criteriality for attributions. In English and other natural languages there are five levels of criteriality for both positive and negative attributions that acquire lexical expression; these, as represented by generic determiners, adverbs of relative frequency and adjectives of possibility, are shown in the diagram below, along with their relations of implication and minimal mutual contradiction, and their relation to the absolute scale of probability (represented by the diagonal in the figure).

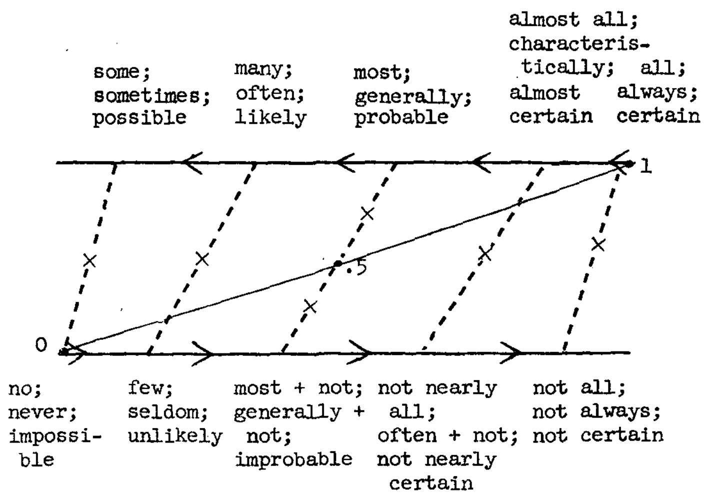

With a bit of intuition, patience, and attention to the requirements of commutativity and associativity, one may also construct a heuristic "multiplication table" that will define products of these levels of possibility, to handle conjunctive and disjunctive attributions.

The most general possible explication of messages, and probably the one that will prove to be necessary for a semantic theory, is that they be simply any programs in the system. Other than performative operations and intensional evaluation operations, the set of operations that constitute messages will include finding an instance of an intension, creating a new instance of an intension, finding or creating an intension similar to a given intension, inserting or deleting quantified relations between extensions and/or intensions, comparing extensions or intensions for equality, inclusion, or mutual exclusiveness, adding or deleting intensional definitions, modifying intensional definitions, rearranging and otherwise modifying situation structures, numerical computations, and the logical operations indicated by the function-definition operators. The specific form of the language in which all these operations may best be combined into programs is yet to be determined, but it may well turn out to be similar to Woods' procedural language, in which the basic statement form is a quantified "pattern-operation" rule.

C. Decoding and Encoding

The process of decoding natural language consists of three stages: syntax recognition, semantic translation, and application to the model representing the current discourse context. Syntax recognition includes recognition of both phrase structure and relations of grammatical agreement among the two or more constituents that are combined by a syntactic rule. Syntactic error correction might be handled by a method akin to Chomsky's [14] notion of "degrees of grammaticalness": relaxation of first grammatical agreement and then syntactic categorization conditions could be allowed until a parsing leading to a semantically-acceptable decoding was obtained. Semantic translation of the combination of constituents recognized into a functional form in the message language then proceeds by way of one or more interpretation functions associated with the rule of grammar. These interpretation functions will make tests for agreement among the inherent and contextual attributes of the intensions they combine; if any of these tests fails, metaphorical interpretation rules (which operate by analogizing and disanalogizing) might be invoked to attempt to resolve the conflict through appropriate construal of one or more of the constituents before the interpretation is rejected entirely. In the final step, application to the model representing the discourse context, antecedents of anaphoric and elliptical expressions are found through appropriate intensional evaluation, and further disambiguation may be achieved in case one or more of the alternative decodings contains presuppositions that are not satisfied in the model. Here additional rules of application may need to be introduced to add information to the model that is required by a presupposition but with respect to which the model is noncommittal.

The problem of encoding is that of transforming a message into a well-formed surface syntactic structure through lexical substitution and formation of syntactic constructions. Encoding may be formulated as a recursive top-down procedure, which operates from the outermost level of functional application in the message on inward to the point where each expression may be replaced by a lexical item, thereafter "unwinding" its way back outward, applying syntactic encodings followed optionally by syntactic transformations to each functional composition encountered on the way. There will, of course, be alternative paths that may be followed in the encoding of a message, because of the possibilities of alternative lexical substitutions, applications of alternative encoding rules, and optional application of transformations. Some of these paths may block because syntactic conditions on the application or output of the encoding rules are not satisfied, others because certain "performance-oriented" constraints, such as constraints on the level of certain types of embedding, are not met in the resulting surface structure.

The rules for lexical substitution can probably be formulated along the lines proposed by Gruber, those for encoding into nominal structures along lines suggested by Celce and Schwarcz [11, 12], and those for encoding into clauses and sentences along the lines suggested by Fillmore.

Both decoding and encoding may be formulated most neatly as non-deterministic procedures employing heuristic search and evaluation. The rules employed in both, furthermore, are of the pattern-operation type, the syntactic structures they operate on are of the form of situation structures, and the semantic structures they operate on are, of course, components of models and messages. Therefore both procedures and rules for decoding and encoding should be formulatable, and perhaps formulatable most elegantly, in the message language, since it is a general-purpose programming language containing all these features. Such a formulation would have the further advantages of parsimony with respect to computer implementation of the semantic theory and easy modifi-ability through the ability to use natural-language statements to effect changes in these procedures and rules.

D. Implications

The approach described above, though it has not yet been imple-mented, can be regarded as a sincere attempt to meet the require-ments set forth for a formal theory of natural language semantics in the first part of this paper--an achievement that no other approach advanced to date can claim, despite the many valuable ideas these approaches have produced. Evaluation of this attempt as a semantic theory must, of course, await the satisfaction of the final requirement: that the proposed system, when programmed, engage successfully in purposive communication with speakers of a natural language. Simply as an approach that holds the promise of adequacy as a semantic theory, however, it can provide a unifying direction for research in a number of areas, including linguistics, lexicology, logic, theory of computing, and artificial intelligence. The unification of such a diversity of directions of exploration, along with the rigorous test that the approach implies for an operationist theory of knowledge and meaning, should render the approach an interesting and fruitful one for philosophical study and exploration.

Adopting an approach such as this, or any approach satisfying the requirements for a semantic theory, as a metatheoretical basis would also help greatly to resolve the confusion that exists today in linguistic theory. This state of disarray is a result of the fact that, with a very few exceptions, linguists have basically ignored the fundamental fact of language as being <u>a tool for communicating something to somebody</u>. They have almost without exception ignored the interface between language and the speaker's or hearer's model of the universe of discourse. Operating in this sort of vacuum, linguists are under too few empirical constraints to determine any theory of grammar, let alone one that is meaningful. Only a semantic metatheory that takes the communicational significance of language explicitly into account can provide a satisfactorily sound basis for a theory of grammar.

The exploration of the approach offered here would bear very much upon the interests of cognitive psychologists, too, in that it offers a unified framework for a theory of language and cognitive processing. The heuristic-search-and-evaluation mode of operation is the paradigm that has emerged from an extensive amount of empirical research on human thinking and problem solving; its successful extension to explicating the understanding and production of language would lend support to the view that the mechanisms employed in language processing are the same as those employed in human thinking in general. The specific forms of model and message structures would, conversely, provide a basis for a formal theory of cognition that would receive support from the linguistic side as well.

Finally, the formulation of the approach as a general-purpose programming system implies that it would be usable, in principle at least, for any application of computers to linguistic and semantic information processing, including all the ones mentioned earlier in the paper. Availability of suitable computer hardware and operating systems would, of course, be essential to any application of the approach on a realistic scale. A more demanding requirement, however, is that of encoding the definitions of the thousands of different words that make up any natural language into an appropriately structured lexicon. A standard dictionary is one possible source of this information, but it will obviously not contain enough to define every word of a language operationally. The work of Olney _et al_ [33] should be very helpful, however, in determining what _can_ be gleaned from a standard dictionary and whether this information can be appropriately supplemented to yield an adequate operational lexicon, or whether a major new lexicographic effort, more rigorous in its requirements than any that have gone before, will have to be undertaken. Whatever the case, the operational lexicon, once created, would be usable for all varieties of applications--and its construction, as well as the programming of applications, would be made easier by the capability implied by this approach to program the system in a natural language once it had been supplied with sufficient information to define the semantics of a suitable "base" subset of the language.

This approach, of course, is only one of many that might be taken to formalizing the semantics of natural languages. Like all other approaches that have been attempted or proposed so far, it will surely reveal its limitations somewhere along the way. But at a time when linguistics, semantics, and computational linguistics are all anxiously searching for a paradigm to follow, this may well be a fruitful one to try.

REFERENCES

1. Austin, John L., How to do things with words. Oxford, Clarendon Press, 1962.

2. Bar-Hillel, Yehoshua, The outlook for computational semantics. Proceedings of the Conference on Computer-Related Semantic Analysis (Las Vegas, Nevada, December 3-5, 1965), Detroit, Wayne State University Press, 1965.

3. Bar Hillel, Yehoshua, Dictionaries and meaning rules. Foundations of Language 3, November 1967, 409-414.

4. Bellert, Irena, Relational phrase structure grammar and its tentative applications. Information and Control 8, October 1965, 503-530.

5. Benjamin, A. Cornelius, Operationism. Springfield, Illinois, Charles C. Thomas, 1955.

6. Bennik, F. D., R. M. Schwarcz and H. F. Silberman, Linguistic and tutorial modeling for natural language CAI. System Development Corporation SP-3266, Santa Monica, California, December 1968.

7. Bohnert, H. G. and P. O. Backer, Automatic English-to-logic translation in a simplified model. IBM Watson Research Center RC-1744, Yorktown Heights, New York, January 1967.

8. Bridgman, P. W., The way things are. Cambridge, Massachusetts, Harvard University Press, 1959.

18. Craig, J. A., S. C. Berezner, H. C. Carney and C. R. Longyear, DEACON: Direct English Access and CONtrol. Proceedings 1968 AFIPS Fall Joint Computer Conference 29, 365-380.

19. Darlington, Jared L., Machine methods for proving logical arguments expressed in English. Mechanical Translation 8, July 1965, 41-67.

20. Fillmore, Charles J., The case for case. In E. Bach and R. Harms (eds.), Universals in linguistic theory, New York, Holt Rinehart and Winston, 1968.

21. Fillmore, Charles J., Types of lexical information. In C. J. Fillmore and I. Lehiste, Working papers in linguistics No. 2, CISRC Technical Report 68-3, Ohio State University, Columbus, Ohio, November 1968, 65-103.

22. Green, C. Cordell and Bertram Raphael, The use of theorem-proving techniques in question-answering systems. Proceedings 1968 ACM National Conference, 169-182.

23. Gruber, Jeffrey, Functions of the lexicon in formal descriptive grammars. System Development Corporation TM-3770, Santa Monica, California, December 1967.

24. Katz, Jerrold J., Analyticity and contradiction in natural language. In Fodor and Katz, 519-544.

25. Katz, Jerrold J., Recent issues in semantic theory. Foundations of Language 3, May 1967, 124-194.

26. Katz, Jerrold J. and Jerry A. Fodor, The structure of a semantic theory. In Fodor and Katz, 479-518.

27. Katz, Jerrold J. and Paul M. Postal, An integrated theory of linguistic descriptions. Cambridge, Massachusetts, M.I.T. Press, 1964.

28. Kellogg, Charles H., Designing artificial languages for information storage and retrieval. In H. Borko (ed.), Automated language processing: the state of the art, New York, Wiley, 1967, 325-368.

29. Kellogg, Charles H., A natural language compiler for on-line data management. Proceedings 1968 AFIPS Fall Joint Computer Conference 33, 473-492.

30. Kirsch, Russell A., Computer interpretation of English text and picture patterns. IEEE Transactions on Electronic Computers EC-13, 1964, 363-376.

31. Morris, Charles, Signs, language, and behavior. Englewood Cliffs, New Jersey, Prentice-Hall, 1946.

32. Morris, Charles, Signification and significance. Cambridge, Massachusetts, M.I.T. Press, 1964.

33. Olney, John, Carter Revard and Paul Ziff, Toward the development of computational aids for obtaining a formal semantic description of English. System Development Corporation SP-2766/001/00, October 1968.

34. Pople, Harry E. Jr., A goal-oriented language for the computer. Ph.D. thesis, Carnegie-Mellon University, Pittsburgh, Pennsylvania, January 1969.

9. Carnap, Rudolf, Meaning and necessity. 2nd ed., Chicago, University of Chicago Press, 1956.

10. Celce, Marianne and Robert M. Schwarcz, A note on two basic semantic distinctions. System Development Corporation SP-3353, Santa Monica, California, April 1969.

11. Celce, Marianne and Robert M. Schwarcz, Counting, collecting, measuring, and quantifying in English. System Development Corporation SP-3378, Santa Monica, California, May 1969.

12. Celce, Marianne and Robert M. Schwarcz, Predicate nominals: lexically simple and lexically complex. System Development Corporation SP-3379, Santa Monica, California, May 1969.

13. Celce, Marianne and Robert M. Schwarcz, Nominal quantification in English: a semantic description. System Development Corporation SP-3380, Santa Monica, California, May 1969.

14. Chomsky, Noam, Degrees of grammaticalness. In J. A. Fodor and J. J. Katz (eds.), The structure of language: readings in the philosophy of language, Englewood Cliffs, New Jersey, Prentice-Hall, 1964, 384-389.

15. Chomsky, Noam, Aspects of the theory of syntax. Cambridge, Massachusetts, M.I.T. Press, 1965.

16. Cohen, L. Jonathan, Review of J. J. Katz, The philosophy of language. Journal of Linguistics 3, April 1967, 163-165.

17. Coles, L. Stephen, An on-line question-answering system with natural language and pictorial input. Proceedings 1968 ACM National Conference, 157-168.

35. Quillian, M. Ross, Semantic memory. Ph.D. thesis, Carnegie Institute of Technology, Pittsburgh, Pennsylvania, February 1966.

36. Quillian, M. Ross, Word concepts: a theory and simulation of some basic semantic capabilities. Behavioral Science 12, September 1967, 410-430.

37. Quillian, M. Ross, The Teachable Language Comprehender: a program to understand English. Report No. 1693, Bolt Beranek and Newman Inc., Cambridge, Massachusetts, March 1968. (To be published in Communications of the ACM)

38. Schwarcz, Robert M., Steps toward a model of linguistic performance: a preliminary sketch. Mechanical Translation and Computational Linguistics 10, September and December 1967, 39-52.

39. Schwarcz, R. M., J. F. Burger and R. F. Simmons, A deductive question answerer for natural-language inference. System Development Corporation SP-3272, Santa Monica, California, November 1968. (Submitted for publication to Communications of the ACM)

40. Simmons, Robert F., Answering English questions by computer: a survey. Communications of the ACM 8, January 1965, 53-70.

41. Simmons, Robert F., Natural language question answering systems: 1969. Computation Center report TNN-87, University of Texas, Austin, Texas, January 1969. (To be published in Communications of the ACM).

42. Simmons, R. F., J. F. Burger and R. M. Schwarcz, A computational model of verbal understanding. Proceedings 1968 AFIPS Fall Joint Computer Conference 33, 441-456.

43. Tarski, Alfred, The concept of truth in formalized languages. In A. Tarski, Logic, semantics, and metamathematics (trans. by J. H. Woodger), Oxford, Clarendon Press, 1956, 152-278.

44. Thompson, Frederick B., English for the computer. Proceedings 1966 AFIPS Fall Joint Computer Conference 29, 349-356.

45. Travis, Larry E., Analytic information retrieval. In P. L. Garvin (ed.), Natural language and the computer, New York, McGraw-Hill, 1963, 310-353.

46. Watzlawick, P., J. H. Beavin and D. D. Jackson, Pragmatics of human communication. New York, Norton, 1967.

47. Weinreich, Uriel, Explorations in semantic theory. In T. Sebeok (ed.), Current issues in linguistics, The Hague, Mouton and Co., 1966, Vol. 3, 395-477.

48. Wilks, Yorick, Computable semantic derivations. System Development Corporation SP-3017, Santa Monica, California, January 1968.

49. Woods, William A., Procedural semantics for a question-answering machine. Proceedings 1968 AFIPS Fall Joint Computer Conference 33, 457-471.