[*Daedalus* 102; pp.217-230]

MARTIN KAY

Automatic Translation of Natural Languages

THE HISTORY OF man's attempt to build a translating machine for natural languages has not been illustrious. There has probably been no other scientific enterprise in which so much money has been spent on so many projects that promised so little. In the late fifties and early sixties, numerous people obtained, from one agency or another of the United States government, appreciable sums of money, in return for which they promised to deliver, in a very few years, a computer program or even an actual machine that could produce high-quality translations automatically. The events that brought these euphoric days to a sudden end are, by now, well known even to people who have no other knowledge of work in machine translation. Stimulated partly by the displeasure of some high-ranking civil servants and military officers at having received less than the best value for their money, the National Academy of Sciences in 1962 established the Automatic Language Processing Advisory Committee (ALPAC) and ordered it to investigate the entire matter of the federal sponsorship of research on machine translation.

In its report, delivered in 1963,[1] ALPAC was as kind to the designers of automatic-translation machines as it could possibly be, It concluded Out there was no possibility of producing a satisfactory translating machine in the foreseeable future and recommended that no further funds be spent on contracts that had such development in view. The committee did not, however, see the development of such machines as forever beyond the wit of man and, in fact, expressed support of the funding of research that aimed at hastening the day when it would be reasonable to let such a development contract.

Reactions to the report were predictable. For almost ten years, any application for financial support for a project involving language and computers, however modest or sound, could expect a swift and categorical refusal. None of the positive recommendations of the ALPAC report were acted upon, and a disservice may thereby have been done to many serious and inventive research workers as well as to the country. Nevertheless, although the number of research projects in computational linguistics has diminished, the discipline has attained far greater maturity. It required dedication to stay in a field that no longer had a ready source of money and whose center of interest had become an object of abuse. However, researchers were now free to look closely at the theoretical problems that stood in the way of successful machine translation. This is not to say that the profession has lost its lunatic fringe. It is not difficult to learn something about how computers are programmed, and many people know a foreign language. Those who know a little of both will always be susceptible to revelations about how a machine might be made to translate. What is to be feared is the predilection that some government agencies are apt to show for proposals that come from precisely this lunatic fringe.

The first machine-translation system to be put into full-scale operation was installed in 1964 at the Foreign Technology Division of the United States Air Force, where it remained in daily operation until 1970. It was a very ingenious machine called the Mark II translator, and it was one of the most interesting products of the early period of work on machine translation. Unfortunately, its ingenuity cannot be accounted sufficient to repay its prodigious cost. A study by Arthur D. Little, Inc. found its translations time-consuming, expensive, and of poor graphic quality; furthermore, they were not very accurate, even after human editing,

The machine made use of a so-called *photoscopic store* consisting of a glass disk, about ten inches in diameter, on which information was inscribed in concentric circles in much the same way as a movie's sound track is represented on the edge of the film. During the life of the system, a vast Russian-English dictionary of stems, prefixes, and suffixes was amassed and new disks were made periodically to incorporate the new information. The logical capabilities of the machine, however, were rudimentary. Each stem and affix on the disk was accompanied by a pair of codes indicating classes of stems and affixes that could occur before and after it. Thus, when a Russian word was sought in the dictionary, various alternative classes might be found, and the one chosen would be determined by the choice made for the item immediately preceding it.

In the heyday of machine translation, Leon Dostert at the University of Georgetown had three independent projects under his supervision. After the publication of the ALPAC report, two of these projects continued elsewhere, though less vigorously, and were eventually quietly buried. The third was delivered as an operational system to translate Russian materials into English to the Atomic Energy Commission at Oak Ridge and to the European Atomic Energy Community (EURATOM) in Ispra, northern Italy. This system, which is usually referred to simply as the "Georgetown program," was designed for use on a standard, general purpose computer, the IBM 7090. Its logical capabilities therefore far surpassed those of the Mark II translator, though the enhancement is not always apparent in the quality of the resulting translation.

The Georgetown program is very complicated. It consists of a large number of instructions that make use of several magnetic tapes on which various kinds of information are stored temporarily so as to make room in the main memory of the machine for other operations. In the course of translating a text, the program goes through a series of more or less well-defined steps called "dictionary lookup," "syntactic analysis," and so on. When this program was designed, work was just beginning on the formal properties of languages and the kinds of processors they might require, and what little was known was, in any case, largely ignored by the designers of this supposedly practical system. The absence of suitable formalisms is not to blame for the scarcity of impressive results from the Georgetown and other early systems, but it is to blame for their monstrous size and complexity.

Though the Georgetown system purported to be concerned largely with syntax, it incorporated neither the notion of a grammatical rule nor the notion of a syntactic structure. The complexity of the syntactic part of the program was devoted to nothing more than resolving ambiguities in the assignment of words to grammatical classes. If a word to be translated could, in the abstract, be either an adjective or a noun, the process examined the word's context to determine in which capacity it functioned in the given sentence. The methods by which this was done were *ad hoc*, and they always provided a single answer to each problem regardless of genuine syntactic ambiguities in the sentence. Of course, an attempt was made to find the solution that would be correct in most cases. The grammatical classifications that were thus appended to the words in a text could be used later to determine which of a list of possible English alternatives would serve to translate the word and to help decide on the eventual order of the words in the second language. Such information about the structure of Russian and English as the program used was built into the very fabric of the program so that each attempt to modify or enhance the capabilities of the system was more difficult and more treacherous than the last. After a while, such a program becomes so complex that any further development is virtually impossible.

In the nearly ten years since the publication of the ALPAC report, much has been learned about linguistics and computer science, but few substantial inroads have been made into the basic problems that beset machine translation. Using the best knowledge that the profession has amassed, an automatic-translation system could be developed far more cheaply and easily today than was possible ten years ago, but there is little evidence that it would be able to produce translations of markedly higher quality.

It is generally agreed that any machine-translation system intended to produce results of high quality must carry out a syntactic analysis of every sentence in the text to be translated. The product of this analysis usually appears as a labeled tree representing the surface or preferably the deep structure of the sentence. Developing a structure of this kind has two important advantages. First, the function that a word or group of words fulfills in a sentence cannot usually be determined simply by examining neighboring words and phrases. It can be determined only by insuring that any function proposed for it is compatible with that proposed for every other word and phrase in the entire sentence. In other words, the most solid basis on which to assess whether a function has been correctly assigned is provided by a structural analysis of the sentence.

Tree structures are also valuable because they permit the definition of a simple but immensely powerful set of operations, known as transformations, in terms of which the structural changes that must be made to produce the sentence in another language can be stated. Suppose that a text is to be translated from a language like English in which the subject usually precedes the main verb and the object follows, into a language like Japanese in which the main verb invariably comes at the end of the sentence. The necessary adjustment in word order is easy to make if the syntactic analysis of the sentence identifies entities like subjects and objects in such a way that their relative positions can readily be altered.

Since there is no theoretical limit on the number of words that can constitute a subject or an object, the structure on which the rearrangement operations are carried out must have a way of connecting indefinitely many words into a group with a name so that it can be treated as a single item. Furthermore, subjects and objects can include other sentences with their own subjects and objects. Take the sentence, "Claims that John had passed the examination surprised the professor." The subject of the sentence is "Claims that John had passed the examination," which contains the second sentence, "John had passed the examination," which has its own subject, "John." The relationships of these various parts to one another can be conveniently represented in a tree diagram, as follows:

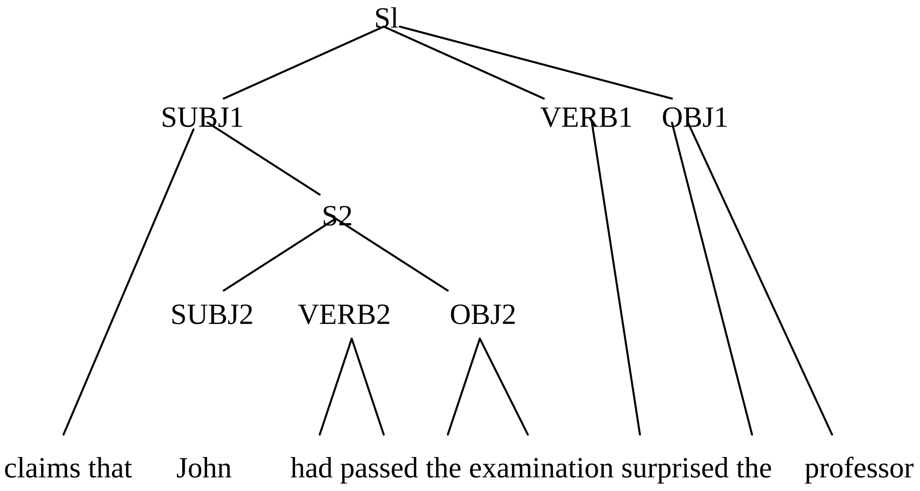

The labels S1 and S2 correspond to the first and second sentences respectively, and lines project down from each of these to labels representing the subject, verb, and object of the sentence.

Suppose, now, that the sentence is to be translated into Japanese. Two kinds of modification must be made. First, the verb of every sentence must be placed at the end, and second, whenever a subject or an object includes a noun and verb that make a complete sentence, that sentence must be placed before the noun it modifies. Arranging the English words in their Japanese order, we obtain, "John the examination passed had claims the professor surprised." The tree diagram representing this sentence is as follows:

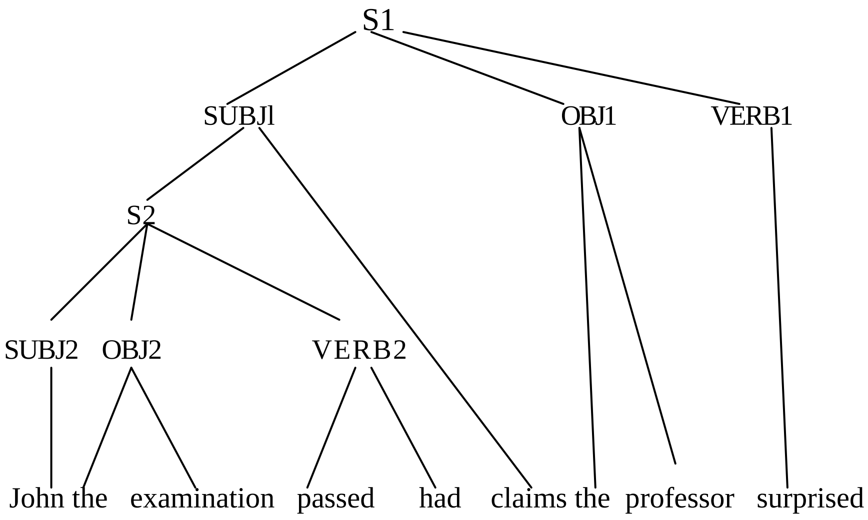

The new tree structure can be obtained from the original by treating the diagram as a mobile and changing the relative positions of the items that hang from particular places.

All the mechanical-translation systems that have been put into regular use are normally described as "machine-aided" translation systems. This is because the translations they produce are not, in general, comprehensible, but must be edited, often heavily, by a person who is familiar not only with the subject matter of the document but also with both languages involved. Therefore the production of a suitable translation by one of these systems can often be complicated, time-consuming, and expensive. All graphic material must first be removed from the text, leaving an indication of where it should be reinserted in the translation. If any of the graphic material contains matter in the foreign language, this must be specially translated and the appropriate amendment made to the tables, graphs, or pictures. The textual material must be represented in a form that the computer can read, and since optical character-recognition devices are still not equal to reading print, this must be done by a human operator at the keyboard. When the automatic translation has been done, a human editor must revise the translation, the graphic material must be reinserted, and a presentable copy must be produced.

In a letter published in *Science* on December 17, 1971, Dr. Wallace Sinaiko described some tentative results of an informal experiment he has been conducting. In 1964, the Foreign Technology Division agreed to have a Russian paper translated for him, using the Mark II translator then in service. The Russian paper was itself a good translation of an English paper, made by a professional translator. Without any detailed knowledge of Russian, Sinaiko was thereby enabled to assess the quality of the product of the mechanical system, allowance being made for the scarcity of data (the original English paper contained only 1085 words) and the possibility that error had been introduced by the professional translator Sinaiko was provided with the unedited output of the machine, making it easier to judge what its contribution to a satisfactory translation would have been,

Sinaiko had the same paper translated again in 1971 by the new system recently installed at the Foreign Technology Division, and he was given both the output of the machine and the final translation after human editing. In possession of two additional translations of the Russian text that he had obtained from professional linguists in 1964, Sinaiko was thus able to compare the raw output of the two translation systems, the final human-edited output of the present system, and the work of the two professional translators.

The techniques that Sinaiko used to compare these translations were simple and informal. The two characteristics he concentrated on were (1) untranslated words and (2) translated words that had two or more possible meanings indicated for them in the translation. The differences between the raw output of the two machine systems were insignificant. The earlier system left 1.2 percent of the words untranslated, whereas SYSTRAN failed to find English equivalents for 2.3 percent. The earlier system provided alternative translations for 6.3 percent of the words, whereas the later system provided alternatives for 5.3 percent. These types of error, if errors they are, would not be found in the work of a human translator.

A comparison of the raw output of the machine with the translation that resulted from editing showed that about 35 percent of the English words printed by the computer were altered by the editor. Every one of the approximately eighty English sentences had some editorial modifications most of them extensive. The most interesting statistic is the following: the manual translators worked at the rate of about 450 words per hour, whereas the editors working on the SYSTRAN output worked at the rate of about 400 words per hour.

Sinaiko was careful to point out that the results of this informal experi-ment ment are anything but conclusive. However, he observes, "It is apparent that little progress has been made during recent years. Moreover I do not know of any demonstrated advantages of MT over human translations."

Earlier I stressed that, while the last ten years have seen significant advances in the ease and elegance with which linguistic operations can be programmed as well as a bewildering array of new proposals in linguistic theory, no advance has been made that promises dramatically to improve the quality of machine translations. However, there may be ways that computer technology could serve translation other than those that have already been tried. At least two other ways have recently been suggested, one capitalizing on the recent development of machines that allow human intervention in the course of the computation, and one involving special artificial languages. If I seem unduly enthusiastic about the first of these, it must be remembered that I had some part in developing the idea.

The MIND system, developed at the Rand Corporation, is a package of computer programs that can be assembled in various ways to fill several linguistic functions. A version of the system was assembled in the latter half of 1971 that is intended to take over, as much as possible, the purely routine work involved in making a translation without ever attempting to solve problems for which it is not equipped. The program contains all the components that one would expect in a full-fledged translation program. There are facilities for analyzing the morphology of words, for obtaining their definitions, and for recording for each sentence all the information furnished by the dictionary about each of its constituent words. A thorough syntactic analysis of each sentence is performed that yields a *deep structure* (in the terms of modern transformational grammar) for each sentence. Transformational rules are applied to these deep structures to produce well-formed sentences in the second language. Finally, there is a component that provides the morphologically appropriate forms for each of the words printed out.

In addition, the system contains a component called a *disambiguator,* whose job is to mediate between the other components of the system with the help of a human consultant, to whom reference is made in all cases of difficulty or unresolved ambiguity. If a word has more than one meaning and the rules supplied to the system provide no basis for deciding which one applies in a particular context, the question will be referred to the consultant. If the rules allow more than one syntactic structure for a sentence, appropriate questions will be formulated to elicit the information necessary to decide among them. If it is necessary to know what a pronoun refers to before it can be correctly translated, the consultant will be provided with a list of possible referents and invited to choose the correct one.

These are the kinds of questions that cannot, as far as we know, be solved in a purely formal way. What is noteworthy about them is that they all arise in attempting to understand the original text rather than in attempting to compose a text in the second language. This suggests that a system of the kind just outlined might function very effectively with a human consultant who is familiar only with the language of the source document and its subject matter. If that is true, such a system might be made to produce creditable translations for technical documents without the services of a human translator or bilingual editor. Whether it can, in fact, do so still remains to be seen. The results of preliminary experiments in the translation of technical manuals from English into Korean are encouraging.

In the realm of language translation, one further line of investigation seems worthy of mention. Largely because of its sheer simplicity, it has usually been ignored or ridiculed in the past. We start from the premise that there are large numbers of people who need to read documents in some foreign language, Russian for example, but who have no knowledge of the language and no desire to learn it. Furthermore, we assume that many of the Russian documents would be read by such a small number of English-speaking people that it would be very difficult to justify the cost of making a translation. Let us further suppose that, though these people. are unwilling to invest the amount of time required to learn Russian, they might be prepared to spend a tenth, or possibly a quarter, of this time to learn a skill of equivalent utility, They might be willing to learn a much simpler language into which, for one reason or another, it proved very simple to translate Russian. If, for example, there were some language into which Russian texts could be mechanically translated in a simple but entirely reliable way, and if this language were also very easy for native English speakers to learn, then these people would have ready access to the foreign materials they needed,

No language with the properties just described in fact exists. But there is good reason to suppose that one could be created. If a dictionary were made that provided a counterpart for each Russian word, prefix, and suffix, and if the process of translation consisted simply of replacing the Russian words and affixes by the counterparts listed for them in the dictionary, a new language would have been created with the grammar of Russian but with a different vocabulary. If the vocabulary were such that each item in it corresponded to one and only one Russian item, the translation process would be completely reversible, capable of reconstituting the original text exactly. Thus, no information from the original text would ever be lost, a property that no other kind of translation has.

Suppose, now, that the items used us counterparts for Russian words were chosen, wherever possible, to be English words, or English-like words, with meanings suggestive of the meanings of the Russian words. Though it is impossible to find English words with the same meaning as some Russian words, that difficulty is encountered less in technical documents where precise equivalents are usually abundant. This method would leave it to the human reader to learn the idiosyncrasies of the most common words with the widest ranges of meanings. In return, it would relieve the human reader of his most time-consuming task, that of finding equivalents for the precise words, which, though they individually occur relatively rarely, comprise the bulk of the vocabulary encountered in technical documents.

Lest what is being proposed here be confused with some early and notoriously unsuccessful experiments in machine translation, it must be stressed that we do not expect native English speakers to be able, without training, to read texts in the curious Anglo-Russian that would emerge from this translation process. We do, however, expect that this language could be learned in much less time than Russian or any other natural foreign language. The production of these translations would be entirely mechanical, and the algorithm required is trivial, so that the cost could be extremely low. In my view, the products of a simple system of this kind would fill the needs of the Foreign Technology Division at least as well as their present system does. Furthermore, the steps that would have to be taken to extend the system to other languages are straightforward, simple and cheap.

At present, linguists are devoting more and more attention to problems of meaning. This was, of course, the principal center of interest in linguistic studies until the end of the nineteenth century when there was a temporary shift of attention to the origin and development of language. One of the most vexing aspects of the study of meaning is that there is very little agreement on the question of what the problems arc that need to be solved. Since almost anything that can be thought can be said, linguists have sometimes sought to exclude meaning from their field of study lest that field become too broad and amorphous. However, it is not clear that the study of meaning entails a study of everything that can be meant any more than that the study of logic entails an examination of every true and false argument. Some students of meaning have undertaken to provide a universally valid scheme for classifying words according to their meanings as Roget did in his well-known *Thesaurus of English Words and Phrases*. Such a categorization, for all that it is purely taxonomic, might be thought of as some kind of map of the territory over which the human spirit roams, or as the basis of a universal vocabulary into which the sentences of any language could be translated. To some scholars, the study of meaning has been effectively identified with the study of informal logic. Depending on how much rigor is introduced into this kind of study, it tends to take the form of an enriched, or corrupted, version of standard logical formalisms.

One of the principal points of contention among students of meaning concerns the question of whether there is, in fact, something that can eventually be captured and examined which is the meaning of a word or sentence. Every attempt to capture such an object leads, at best, to other words and expressions, possibly in some formal notation. Presumably the best that can be said is that the new set of words and expressions provides a more transparent representation of the meaning and shows the contributions of various components explicitly. But it cannot be claimed that anything set down on paper actually is a meaning. Some scholars have reacted to this situation by noting that the fact that words and sentences are meaningful is not grounds for assuming that there must be something which is their meaning.

The meaning, as Wittgenstein said, is the use. The meaning of a word or sentence is the total set of relations that it contracts with other words and sentences. When I learn a new word or a new fact about the world, the result is to change, however imperceptibly, the meanings of all other words and sentences in my language. While this view does not broaden the scope of linguistics so that it embraces the whole of science, it does claim for it much of the territory that was previously thought of as belonging to psychology and philosophy. In this view, a person's knowledge of the world is defined by his ability to describe that knowledge in language.

By what criteria should a theory of meaning in ordinary language be judged? Each theorist, of course, has his own answer. However, many people are prepared to concede that an ultimate test of a theory of meaning would be to incorporate it in the design of a machine, thereby enabling the machine to demonstrate the same kind of linguistic competence as a human being.

Allan Turing suggested that we could claim to understand the basis of human intelligence only when we could build a machine with which human beings could communicate and which resisted every attempt on the part of an interlocutor to determine whether it was, in fact, a machine There is a growing number of students of language, most of them, to be sure, not claiming to be linguists, for whom the adequacy of a theory of meaning must be assessed in just this way. They would claim that the studies of meaning and of intelligence are all one.

The value of this approach to the study of meaning does not depend on the validity of the specific projects that have hitherto been based on it or on how readily we expect to be able to develop machines whose performance approaches the ideal. It does depend, at least to some extent on such fundamental epistemological questions as whether it is ultimately possible to judge the grasp of meaning that a machine or organism has attained purely on the basis of its behavior, What would it be like to have a machine that not only could tell me that it was sorry I had a cold but could also be sorry? Is it possible to understand the meaning of a word like "sorry" without being able to experience the emotion? To put the question somewhat differently, what conclusions would we be justified in drawing about the human faculty of language from a machine that had been enabled, by various kinds of cunning and trickery, to masquerade as a human being? Clearly there would be no necessary connection between the components of the machine and the components of human psychology. But this is to say nothing that cannot be said with equal justice of any linguistic theory that has been proposed. The test of a scientific theory must be behavioral. We cannot expect scientific models to operate for the same reasons or by the same processes as reality, but only to operate in a manner sufficiently analogous to enable us to extrapolate about reality from the behavior of the machine. Because of this ignorance of motive, the scientific value of a talking machine cannot be assessed objectively, but only on the basis of such subjective criteria as the parsimony and elegance of its structure.

The attempt to build machines that mimic human behavior belongs to a field that has come to be known as *artificial intelligence.* A contribution to that field that has recently attracted a great deal of attention is a computer program designed by Terry Winograd of M.I.T. This program enters into a conversation with its human interlocutor about a very carefully restricted domain of discourse. The program causes a picture to be displayed on a television screen depicting a table top on which a number of simple objects—cubes, balls, pyramids and boxes of various sizes and colors—are distributed. The machine can be instructed to move these objects about on the table top and it does this using its single "hand," a depiction of which can be seen entering the display from the top of the screen. It can, therefore, move only one object at once. It is possible to imagine instructions that require some ingenuity to carry out. Suppose, for example, that there are three blocks on the table and that the machine is told to stack them on top of one another. It may be that some of the blocks are initially supporting other objects which must first be removed. Obstructions must be removed from the upper face of at least two of the blocks before the stacking can begin.

Winograd's program may have to design quite a complex strategy in order to carry out a particular instruction, but, according to the view on which this work is based, it can only be said to understand an instruction fully if it can respond in this positive way. The program can also be asked questions about the disposition of the objects on the table and about its reasons for making particular moves. It may, for example, be asked, "Why did you put the green block on the red one?" to which the answer might be something like "Because you told me to stack up three blocks so that I had first to stack up two blocks."

Students of artificial intelligence have worked with very diverse models from robots that use a television camera for an eye and can move from place to place negotiating obstacles to programs that prove mathematical theorems and play chess. Hitherto, few of these efforts have involved a determined attack on obviously linguistic problems. Interaction with machine has typically been through the medium of .specially designed languages but, to the extent that a wider view is taken of problems of meaning, these projects can be seen as contributing to our understanding of natural language. For Winograd, it is a matter of the first importance that his program communicate in English and he describes his work as contributing to *procedural semantics*, an explicitly linguistic enterprise. For him, the meaning of a sentence is the procedure that it sets off in the head of the hearer and he takes it as his task to replicate that process in a machine.

Any machine that processes textual data in nontrivial ways must have certain basic capabilities. It must be able to recognize words, making due allowance for the ways in which their forms vary with number, person mood, and the like. For each word, it must be able to retrieve information about its syntactic and semantic properties from a dictionary. It must be able to distinguish the correct syntactic structure from among the several possibilities in a grammatically ambiguous sentence. The details of how these processes are carried out depends on the theoretical stance of the designer. For some purposes, a strategy that is expensive in terms of computer resources may be preferred because it is considered a better model of the human strategy or because it is more perspicuous. On the other hand, if large amounts of text are to be treated, efficiency may be a prime consideration, For one purpose, it may be necessary to have all possible analyses of every sentence whereas for another it may be desirable to seek the analysis which is, in some sense, most probably correct.

Until recently, it was thought that each set of requirements demanded a new program and that there was no end to the designing of essentially different algorithms for basic linguistic processes, While there is, of course, no way of knowing what tomorrow's revelations may bring, it now seems likely that the best algorithms will turn out to be variants of a single overall strategy. Three strategies have been proposed for obtaining so called deep structures for arbitrary sentences. By "deep structure," I mean the kind of structure assigned to a sentence by some variant of transformational grammar. It is an attempt to make explicit the underlying logical relations among words rather than simply to label subjects, objects and the like. There has been rivalry among the proponents of three strategies, which were thought to be fundamentally different. However, it has recently become clear that the similarities are more striking than the differences. There appears to be a common core of operations that must be part of any algorithm for syntactic analysis.

The oldest of these strategies was the subject of Stanley Petrick's doctoral thesis at M.I.T.[2] It is a complicated procedure divided into several different stages and drawing heavily on the details of Chomsky's formalization of transformational grammar. The other two proposals make no direct reference to this formalism. William Wood's Augmented Transition-Network Parser[3] is inspired by parts of automata theory and, in particular, by the notions of automata theory with finite numbers of discrete states and of push-down stores. Kay's chart parser[4] capitalizes on the notion of general rewriting rules. It is at least in principle, possible to write equivalent grammars for programs that follow each of these three strategies. In other words, grammars can be written which would cause the three programs to deliver identical analyses of the same sentences. However, the grammars would be written in entirely different notations; furthermore, they would cause quite a different sequence of events to occur in the machine. From this point of view, grammatical formalisms take on the aspect of high-level programming languages, each of which requires a compiler to translate it into the language of a particular machine. The difference is that in this case, the machine is not simply a general purpose digital computer, but a special machine which might be called a *syntactic processor* It is not necessary to construct instances of this special machine out of pieces of hardware because a general purpose computer can be made to stimulate it by supplying it with the appropriate algorithm in a suitable programming language.

That it is possible to design a single machine with reference to which grammatical formalisms appear as high-level programming languages is, theoretically, not surprising. Indeed, it is not difficult to prove that, if the formalism is adequate for syntactic analysis at all, then it must be possible to solve the problem in this way. What is interesting is that the proposed syntactic processor turns out to have a simple and elegant design and that this approach to the problem of syntactic analysis is efficient and practical. The difference between the syntactic processor and the general purpose computer is the difference between the theoretically adequate machines that are the object of mathematical study and the machines that are manufactured by engineers.

It will take time to discover the cash value of something like the syntactic processor. At best, it will be shown to incorporate important components of the human faculty of language. At worst, it will be a useful piece of engineering. In any case, it belongs to the field of computational linguistics.

The strategy of syntactic analysis is a real problem on which some modest headway has been made. But it is not a problem that belongs obviously either to linguistics or to computer science and it would probably never have arisen in the normal course of work in either of these disciplines. The same can be said of many problems in semantics. The computational linguist, however, sees problems of meaning in a different light from other linguists. To him, the meaning of a sentence is, as I have said, a process—a program that will be carried out in the head of the hearer. The computational linguist is, above all, a specialist in the processes of language and he is coming more and more to see semantics as the field in which his main contribution will be made.

REFERENCES

1. Automatic Language Processing Advisory Committee, National Academy of Sciences, National Research Council, *Languages and Machines* (Washington, D.C.: U.S. Government Printing Office, 1966).
2. S. R. Petrick, *A Recognition Procedure for Transformational Grammars* (Cambridge, Mass.: M.I.T. Press, 1965).
3. W. A. Woods, "Transition Network Grammars for Natural Language Analysis" *Communications of the ACM*, XIII (1976); W. A. Woods and R. M. Kaplan, *The Lunar Sciences Natural Language Information System,* Report No. 2265 (Cambridge, Mass.: Bolt, Beranek and Newman, 1971).
4. M. Kay, *Experiments with a Powerful Parser* (Santa Monica, California: Rand Corporation, 1967), RM 5452-PR; R. M. Kaplan, *A General Syntactic Processor* (forthcoming).