OPERATIONAL PROBLEMS OF MACHINE TRANSLATION:

A POSITION PAPER

Paul L. Garvin

State University of New York at Buffalo

The late Uriel Weinreich in the first of his lectures in the Trends in Linguistics lecture series at the Indiana University Linguistic Institute on July 13, 1964 referred to machine translation as "linguistics' most conspicuous and expensive failure."[1] Two years later the Automatic Language Processing Advisory Committee of the National Academy of Sciences, National Research Council in what has since become known as the ALPAC Report (1966: 24) stated that "No one can guarantee, of course, that we will not suddenly or at least quickly attain machine translation, but we feel that this is very unlikely."

In the light of these two highly authoritative statements of position, and in view of the abrupt reduction of funding for machine translation research, is it at all reasonable to discuss operational problems of machine translation these days? The answer is of course that if one is to talk about machine translation at all, it must be in terms of some reasonable operational objective, since research without such objectives will at best be related to machine translation only indirectly. The question as to whether or not such objectives are reasonable depends in this author's opinion upon the researcher's basic orientation: with a predominantly theoretical orientation, machine translation research will clearly be close to pointless; with an operational orientation, on the other hand, machine translation research will not only be interesting and valuable in its own right, but will also constitute one of the few available conclusive means of verification of the findings of linguistics (cf. Garvin 1962: 387).

This paper will attempt a survey of the major controversial issues in the field of machine translation, all of which, in the light of the above discussion, are considered operational. These issues are considered to fall into three basic categories: linguistic problems, design problems, and bread-and-butter problems.

LINGUISTIC PROBLEMS

The Role of Linguistics Theory in Machine Translation

It is a commonly held view among linguists, both the few who are interested in machine translation and the many who are not, that any application of linguistics—and in the linguist's view this certainly includes machine translation research—must be based primarily on a strong linguistic theory. (For a recent statement of this view see Bar-Hillel 1970). This is essentially a capsule view of the theoretical orientation. While nobody will deny that any applied work must have a sound theoretical basis, from an operational standpoint there are a number of things seriously wrong with an over-emphasis on theory.

(1) Machine translation is considered primarily an operational rather than a theoretical problem. Consequently, an application of sound linguistic research methods is more important than a further elaboration of linguistic theory.

(2) Most strong linguistic theories are essentially generative in nature. However, the basic problem in machine translation is not a generative but a recognition problem. Recent research in psycholinguistics has confirmed an opinion long held by this investigator, namely that a recognition problem cannot be resolved by simply reversing a generative system.

(3) The kind of questions that arise in machine translation research are not necessarily the kind of questions that are most popular among linguistic theorists. More specifically, the basic problems of ambiguity resolution are different in machine translation from the formulation they have received in

recent explanatory linguistic theories.

(4) It is an old operationalist adage that one can best learn by doing. This is particularly true in the case of machine translation where the machine manipulation of linguistic data forces the investigator to recognize a great many inaccuracies and intuitive shortcuts that are usually glossed over in theoretical linguistic research. Thus, rather than relying excessively on the contributions of linguistic theory to machine translation research one should expect significant contributions to linguistics from research on machine translation.

Nature of Models

Linguistic models can be categorized as strong or weak, depending on whether or not they have strong or weak formal pretensions. Current trends in linguistics favor strong models; this is of course based on an epistemological attitude that is oriented towards the elaboration of theory rather than of method. In line with the discussion in the preceding section, it is here considered that, particularly for purposes of an application such as machine translation, weak models are to be preferred to strong ones. The reason is that strong models are considered to prejudge the direction of research in a situation in which there are too many unforeseen and as yet insufficiently known factors. Clearly, there have to be grammars of both the source and target languages at the base of any machine translation system. Equally clearly, however, these grammars need not be formal grammars; as a matter of fact, in this author's opinion descriptive grammars are strongly preferable to formal grammars for purposes of machine translation because they are much better able to account for the indeterminacies of natural language structures which, as was so well stated by Charles F. Hockett recently, are essentially ill-defined systems (Hockett 1968: 44-45). Descriptive grammars can best be developed in a primarily method-oriented, rather than a primarily theory-oriented, frame of reference. As a matter of fact, in such a frame of reference conventional grammars may be used as a reasonable point of departure, with the necessary modifications introduced as the requirements of machine translation become apparent in the process of the development of experimental systems. Operationally oriented machine translation research both in the United States (cf. Garvin forthcoming) and in the Soviet Union (cf. Bel'skaja 1969) has done just that. This author has made strong claims on behalf of his proposed version of an operational machine translation system (Garvin 1967); it is not known how far along comparable Soviet versions have progressed.

Aspects of Linguistic Structure

All linguists seem to agree that the system of language is hierarchically structured. That is, they all look upon the system of a language as having different levels, or strata, or components. From a machine translation standpoint, of course, it is most important to know which distinctions between different aspects of language are relevant for the development of machine translation systems. The least significant seems to be that of phonology and grammar, since no machine translation system to this author's knowledge is concerned with phonology at all. The most important is the distinction between grammar and lexicon, since all machine translation systems known to this author make some distinction somewhere between a dictionary lookup based on the lexicon and an algorithmic portion based in part on the grammar.

Linguistic approaches differ in regard to whether or not the lexicon is considered a part of the grammar or a dimension separate from it. In either case, the lexicon and the grammar (or the remainder of the grammar) are kept clearly separate by most linguists. The difficulty in machine translation is that the lexicon and the grammar cannot be hermetically sealed off from each other. The dictionary and the algorithmic portion correspond only roughly to lexicon and grammar respectively; the dictionary, after all, contains a grammar code which is based on the grammar, and the algorithmic portion serves to resolve not only grammatical but also lexical ambiguities. Nevertheless, an understanding of the differences between lexicon and grammar is essential for a proper operational assessment of all the variables that enter into the design of a machine translation system.

Related to the conception of levels or strata of language is the methodological problem of conducting the analysis "from the bottom up" or "from the top down". In the first case, the minimal units of language are considered as the input into the analysis and the output yields the maximum units which are, for all practical purposes, the sentences of the text. In the second case, the input are the sentences and the output is a decomposition of the sentences into their constituents. Clearly, since in machine translation the grammatical information is transmitted to the algorithmic portion from the dictionary by a lookup of individual textwords, and since therefore the initial input elements into the algorithmic portion are the "bottom units", a "bottom to top" approach is the most operationally efficient one for machine translation.

DESIGN PROBLEMS[2]

Sensing Units and Translation Units

This is one of the oldest and also most important problems faced in the design of machine translation systems. Sensing units are linguistic units which the computing equipment can read, that is, for all practical purposes, strings of letters separated by spaces and/or punctuation marks. Clearly, these correspond only partially to the translation units, that is to say, the grammatical and lexical units that must be manipulated in order to effect translation. The problem consists in providing the machine translation system with a capability for transforming the sensing units into appropriate translation units. In a sense, the entire recognition problem in machine translation is a consequence of this difference between sensing units and translation units. Were it not for that, the brute force conception that machine translation can be effected by a large enough dictionary with some adaptations to make room for syntactic and semantic differences between the two Languages, would indeed be adequate. And, needless to say, everyone who has had any experience with the field knows that this is not so.

Intermediate Language

There has been a good deal of discussion in the machine translation literature about the presumed advantages of an intermediate language (for an example of this, cf. Andreyev 1967). It has, for instance, been asserted that with the help of an intermediate language one could ultimately accomplish translation from any one of a number of different languages into any one of a number of other languages much more economically than by designing a corresponding series of translation systems for the different conceivable language pairs. Some have gone so far as to give a mathematical expression to this theoretical economy. From an operational standpoint, of course, the real problem is whether or not machine translation of any satisfactory quality can be achieved at all, rather than how it can be achieved most economically for an indefinite number of language pairs.

Nevertheless, the question of intermediate language deserves serious consideration, from an operational as well as from a theoretical standpoint. Some authors, such as Andreyev (ibid.), have talked about an intermediate language as if it were a real language such as English or Russian. That is, translation would be done from the source language into an intermediate language conceived of as a real language, and then back out of the intermediate language into the target language. If the intermediate language is so conceived, then this means encumbering the translation process by an additional step: instead of translating in a one-step process directly from the source language into the target language, it becomes a two-step process in which translation is first effected from the source language into the intermediate language, then from the intermediate language into the target language.

If on the other hand the intermediate language is not conceived of as a genuine language with all the appurtenances, then the conception of an intermediate language becomes much more rational from an operational standpoint, and also much more trivial from a theoretical point of view. In that case, the intermediate language is nothing more than a series of symbolic notations to record the output of the recognition routine and to serve as input into the command routine by which the text in the target language is to be generated. This, as was said, is operationally effective—it is also operationally necessary, because there must be some way in which the information gathered by the recognition routine is stored and transmitted out into the command routine. The use of the term intermediate language then becomes trivial, because this information store will certainly not have the language-like qualities which the term implies. It is further conceivable from an operational point of view, although certainly premature at the present state of machine translation research, that the same information store can be filled by a number of different recognition routines for different languages, and in turn feed into a number of different command routines for different target languages. The information store then will be combined with a kind of switchboard that will direct the appropriate recognition routine into the store and make sure that the output of the store is fed into the appropriate command routine. Thus, the theoretical efficiency talked about in the preceding paragraph is conceivable, but in a sense which for the current state of affairs is operationally trivial.

Total Accountability

Many linguistically oriented researchers in machine translation have claimed that in order for machine translation to be possible, it is necessary to account for all of the linguistic conditions that exist in a language. Some, such as Bar-Hillel (1970) have gone even further and claimed that not only linguistic conditions, but also pragmatic conditions have to be accounted for in order to make machine translation of the desired quality possible.

From an operational standpoint, this is an inappropriate identification of the aims of exhaustiveness in linguistic research with the aims of machine translation. Clearly, only those linguistic conditions which have a bearing on the translation process need be accommodated in a machine translation scheme. Thus, most of derivational morphology, although of great interest to the linguistic researcher, is essentially irrelevant to the translation process, since derived forms can be entered into the machine translation dictionary with their appropriate translations without going through the trouble of underlying analysis. Similarly, it is certainly not to be expected of a machine translation system any more than of a human translator to translate unambiguously passages which are inherently ambiguous in the source language. Likewise, no machine translation system should be expected to account in its entirety for those pragmatic factors which under ordinary circumstances would remain obscure to the human peruser of the source language text.

Morphological Analysis

Quite a few linguistically oriented machine translation researchers have given a great deal of attention to automatic morphological analysis as part of the machine translation process. This analysis has been primarily concerned with attempting to determine morpheme boundaries within printed words; some researchers have limited themselves to separating inflectional endings from the base portions of the words, while other researchers have gone further than that and also included the segmentation of derivational morphological material. One of the reasons given for this has been the requirement of total accountability which was discussed in the preceding section. Another, operationally more valid, reason has been that separating inflexional endings from base portions, while it may encumber dictionary lookup, saves a great deal of storage space in the dictionary portion of the program. The reason given for segmenting derivational material has been that it facilitates the recognition of neologisms. Clearly, the latter two reasons apply primarily to "highly inflected" languages such as Russian or German.

As far as the segmentation of inflectional morphemes is concerned, which some machine translation groups have called "stem-affixing", this is a perfectly reasonable space-saving procedure when it comes to high frequency regular inflectional patterns. In the case of the so-called exceptions, particularly when the irregularities involve changes in the base portions of the words, no operational gain is derived from the segmentation of inflectional morphemes from base portions.

As far as the segmentation of derivational elements is concerned, the advantages derived from the facilitation of the recognition of neologisms have to be weighed against the disadvantages of introducing an additional elaborate systems task into the design. In this author's opinion, the segmentation of compounds into their components may well be extremely useful in the recognition of neologisms. On the other hand, the segmentation of derivational morphemes from the remainder of the base portions of the words is both operationally more cumbersome than the segmentation of compounds, and less likely to yield results in the correct recognition of neologisms. It is, after all, well known that the lexical meanings of derived words, particularly in the Slavic languages, are often not predictable from the sum of the meanings of the derivational morpheme or morphemes and the remainder of the base portion.

Grammar Code and Algorithmic Portion

Most workers in the field of machine translation agree that grammatical information is stored in the form of grammar codes in the dictionary of the system; the term grammatical information is here used loosely to include whatever lexical and other semantic information *is* available to the program. This information is then called by the algorithmic portion of the system for further processing to effect the required recognition of the source language input and subsequent generation of the target language output. This raises the question as to how much information is to be stored in the grammar code, and how much of the recognition and subsequent generation task therefore is to be left to the algorithmic portion. The current trend in much of linguistic theorizing has been to emphasize the significance of rules; this means, from a machine translation standpoint, that a great deal of the recognition burden is placed on the algorithmic portion, with only as much contained in the grammar code as is considered theoretically desirable. Since, however, a table-lookup operation is significantly easier to perform than an algorithmic one, there is a distinct operational advantage in maximizing the grammar-coded information and minimizing the role of the algorithmic portion. This does not of course mean that under any circumstances the role of the algorithmic portion becomes as trivial as some early machine translation researchers have assumed it to be. The other question, namely, whether or not the algorithmic portion should contain a separate table of rules, will be discussed in a subsequent section on bipartite and tripartite design in machine translation.

Content of the Grammar Code

One of the lessons learned by all machine translation research groups has been that the amount and type of grammatical information contained in conventional grammars is wholly inadequate for machine translation purposes. At least the following types of information have been found essential for most full words in a language such as Russian: (1) Word class information. This includes not only indication of the conventional word class, but requires a reformulation of word class distinctions in terms of syntactic functioning of Russian words. Thus for instance, instead of including participles under verbs as one of their forms, as is done in many traditional grammars, participles are considered a special type of attributives with the particular function of having the potential for governing dependent structures. Thus participles are included in a category of governing attributives together with a number of adjectives that are functionally equivalent to them, and are not included in the same class with verbs. Similarly, infinitives are considered a separate word class because, unlike finite verbs, they do not ordinarily have subjects, and unlike gerunds which likewise do not have subjects, they may serve themselves as subjects of Russian clauses. (2) Agreement information. This has to include information with regard to modifier-head-type agreement, as well as information as to whether or not the agreement is of the ordinary type (as for Russian adjectives) or of the exceptional type (as for Russian numerals). (3) Modifiability information. This information has to state what dependent words a given word may be modified by. That is, in the case of nouns, what adjectives may modify a given noun; in the case of adjectives, what adverbs may modify a given adjective; etc. (4) Complementation information. This information concerns the type of complements which may be associated with a given verb, noun, or other Russian word that may have a complement in association with it. Thus, verbs or nouns of location may have certain types of complements of place accompanying them; verbs or nouns of time may have certain complements of time accompanying them, etc. (5) Governor class information. For those words which may be governed by other words, information in regard to the particular kinds of words which may govern them: for instance, in the case of adjectives, the kind of nouns to which they may be modifiers; in the case of adverbs, the kind of adjectives to which they may be modifiers. (6) Government information. Those words which govern dependent structures, the kind of dependent structures which they may govern. For instance, the kind of case a verb or noun may govern, whether or not more than one dependent structure may be governed and in which case each of the possible dependents will stand, whether or not there is prepositional government (which preposition and demanding which case), etc. (7) Subject class information For verbs, the class of subjects which a given verb may take, such as animate, inanimate, human, etc. (8) Object class information. The same type of information as for subject class, except of course, concerning the object which a given verb or a given governing attributive may take.

The above includes only a part of the kind of information required for a complete grammar code. Much of this information is commonly considered semantic rather than grammatical; much of it has to do with not only the syntactic recognition of the sentence but also with the recognition of semantic compatibilities. A great deal more information is needed if in addition to this type of recognition correct choices are to be made in the case of multiple meaning.

Bipartite versus Tripartite Design[3]

The issue here is whether or not the rules of the grammar of the source language should be contained in *a* table to be called by a parsing algorithm, or whether these rules should be written into a more elaborate algorithm of which they become an organic portion. In the first case, the machine translation program would essentially consist of three portions: a dictionary, a parsing algorithm, and a table of rules—hence, the term tripartite. In the second case, the machine translation program will consist of only two portions: a dictionary, and a translating algorithm—hence, the term bipartite.

The main arguments in favor of a tripartite design are: (1) that it allows the processing by one and the same parsing algorithm of more than one table of rules; thus, if any corrections in the grammar are to be made, this involves only a relatively simple updating of *a* given rule table, and does not require any revision of the algorithm itself; (2) the labor of the programmer who is responsible for the parsing algorithm can be kept separate from the labor of the linguist who is responsible for the table of rules. In theory, these two advantages appear to be overwhelming. In practice, it turns out that the fundamental problem in the automatic recognition of grammatical structure of text is the correct sequencing of the application of the rules of the grammar which are supposed to effect the recognition. In this author's opinion, such a sequencing of the application of different grammatical rules can be effected only by making the rules of the grammar an organic part of the algorithm; this is the only way to insure that a given rule will be called only after all the conditions that are necessary for its operation have been previously recognized by other rules of the program, and that such a recognition has been effected in the correct order.

This requirement of sequencing of rule application is based not only on the recognition that the grammar of a language is hierarchically structured, that is, that there are levels to be gone through. It is also based on the recognition that in addition to the levels of the language, there is also an operational order in which grammatical and other information becomes available to the program. Thus, once again, it is apparent that the operational requirement does not parallel the theoretical desiderata.

Recognition Strategy

As was stated above, a bipartite machine translation design is considered operationally preferable to a tripartite one. This means that the algorithmic portion of a machine translation program operates on the basis of something like a pattern recognition strategy, rather than a parsing strategy. This means that the algorithmic portion will in essence carry out a number of context searches to recover the conditions necessary to effect recognition and subsequent translation.

This conception of recognition strategy raises the basic problem of the proper organization of the searches in a bipartite system. Two basic types of search patterns have been proposed in the literature: one is predictive analysis, devised by Ida Rhodes of the National Bureau of Standards and soon thereafter adopted by the then Machine Translation Research Group at Harvard University (Oettinger and Sherry 1961); the second is the author's fulcrum approach (Garvin 1968). Predictive analysis employs essentially a straightforward left to right search pattern. The algorithm looks at the grammar code of each textword, as it has been looked up in the dictionary, from left to right. For each textword it records its grammatical potential in the form of predictions and notes the extent to which predictions of previously noted words are fulfilled by each current word. If this succession of predictions and their fulfillments does not result in an appropriate recognition of the syntactic structure of the sentence, then unfulfilled predictions and unused fulfillment possibilities are retested in a program component called hindsight. The basic difficulty of this approach is that the more complex a sentence, the greater the burden placed upon the hindsight; from an operational standpoint, the greatest weakness of this approach has been that the hindsight has never properly been worked out. In the fulcrum approach, on the other hand, searches are designed to use words in order of their grammatical significance, rather than in the linear order of their appearance in text. Thus the searches are directed first at those words which contain the most grammatical information from the standpoint of the recognition of a particular structure (the so called fulcra), then they branch out from these pivot words in order to encompass the remainder of the structures in question. Since not all grammatical information is retrievable in a single pass, the fulcrum approach uses a succession of passes for the retrieval of the grammatical information contained in each sentence.

The reasons for which the fulcrum approach is considered operationally preferable to predictive analysis are the same for which a bipartite system is considered operationally preferable to a tripartite one: the need for the appropriate sequencing of the application of grammatical rules to the elements of the text.

Single versus Multiple Interpretation of Sentences

Many approaches to automatic syntactic analysis, whether connected with machine translation or not, favor the outputting of as many parsings of each sentence as is conceivable in terms of the given grammar code. The reason for this preference is the theoretical interest of showing the variety of conceivable analyses based on a given grammar code. From an operational standpoint, this is clearly undesirable, since the operational aim of machine translation is not to show the variety of conceivable interpretations of each given sentence, but to arrive at some reasonable form of translation with the minimum of waste motion. Thus, in an operational approach to machine translation priority must be given in each case to the most likely interpretation of any given sentence in the hope that this will indeed turn out to be the interpretation applicable in the particular case. As the machine translation system is refined, provisions can be included for superseding this most likely interpretation in favor of a less likely one, if the latter turns out to be the one applicable to the particular case.

Filter versus Heuristics

This question is closely related to the one treated in the preceding section.

A program component called filter has been used in some of the Soviet approaches to operational machine translation (cf. Mel'čuk 1964, Iordanskaja 1967). The intention of these systems is essentially to produce if not all, then a number of, the different possible syntactic interpretations of each sentence, by means of an algorithm which incorporates numerous table lookups and is essentially based on a variant of dependency grammar (cf. Hays 1964). It is not known to what extent this approach has been operationally successful; *it* is known, however, that the Mel'čuk group has since turned its attention to other problems of a more theoretical nature (cf. Mel'čuk and Žolkovskij 1970).

A machine translation design which gives a preferred single interpretation to each sentence obviously does not need a filter for the selection of one alternative from among many. What it does need is a capability for the revision of the one selected single alternative, in case overriding conditions in the grammatical makeup of the sentence require that it be superseded by another interpretation. The mechanism for overriding previously made determinations as to the interpretation of sentences is given by the inclusion of a heuristic capability in the machine translation design. The initial preferred interpretation of a sentence is given on the basis of information derived early in the syntactic processing. This information may have to be overridden on the basis of more powerful information obtained at later stages in processing. Consequently, the heuristic component must both recognize which interpretation; may be subject to later revisions, as well as identify the conditions on the basis of which any prior interpretation is subject to such a revision. Usually, the original interpretation is arrived at on the basis of the immediate context, and whatever revisions may be necessary arise from the inclusion of a broader, usually clause-wide, context. The advantage of combining a single preferred interpretation with a capability for revision based on heuristics is essentially that in most cases the original preferred decision, precisely because it is based on greater likelihood, may be allowed to stand. Thus a great deal of the processing involved in the use of filters can be avoided. (For a detailed discussion of the use of heuristics in the fulcrum approach, see Garvin 1968: 172-81).

BREAD-AND-BUTTER PROBLEMS

Quality of Translation

A great deal of discussion in the machine translation literature has been devoted to the feasibility or non-feasibility of high-quality machine translation. Much of this discussion has been quite unrelated to reality, because it has been based on an A Priori abstract conception of what constitutes high quality translation. Clearly, the question of the quality of translation has to be related to user need: the greater the need, the more it is possible to compromise with quality. This has recently been recognized even by Bar-Hillel (1970). For many purposes, machine translation output will be only casually scanned rather than carefully read; from a great mass of documents so perused a few may then be selected for later, more careful, human translation. Another factor to be considered is the speed with which machine translation can be effected, as compared to the time required to produce good quality translation by human labor. This has, of course, been used as an excuse for the perpetuation of operating, though operationally unviable, machine translation systems. Nevertheless, it is one of the practical problems deserving more careful consideration than has been afforded them in the past.

Input Preparation

In the view of most observers, the greatest practical handicap in the use of machine translation has been the high cost of key-punching the original document for input into the computing system. Clearly, the only way of overcoming this handicap is by the use of automatic character recognition.

Recent claims to the effect that character recognition is now feasible for a sufficient number of fonts to be practical seem to have some validity. Undoubtedly, this will have a great effect on the evaluation of the economics of machine translation in the future, provided the question can be approached with sufficient detachment from the mistakes of the past.

Staffing

This is the most complex practical problem in both machine translation research and in the maintenance and updating of the machine translation systems of the future, if indeed such systems will ever become practical.

The reason this is such an extremely difficult problem is because both the development and the maintenance of machine translation systems require the cooperation of personnel with two sets of qualities that are very rarely found in the same individuals. On the one hand, work in machine translation requires great originality, expertise, intuitive brilliance, and all the other qualities that make for good researchers. On the other hand, machine translation research also requires extreme intellectual discipline, patience, persistence, and willingness to give up one's individual original ideas in favor of the established parameters of the system.

One of the more easily resolvable problems of staffing is the decision as to whether the work of linguists and programmers should be combined in the same persons, or whether the two competencies should be kept separate. In this author's experience, no linguist will ever become a good enough programmer, and conversely. Therefore, in order to maintain the highest possible level of professional competence in the research staff, the two competencies should be kept separate but should learn to work in close coordination. This again is an extremely difficult objective to achieve in practice, although it is much talked about in theory.

Footnotes

[1] This statement was not included in the published version of his lecture (Weinreich 1966).

[2] For a discussion of the design of the author's proposed machine-translation system, see Garvin 1968.

[3] For a detailed discussion of the author's views of this issue, see Garvin 1966.

References

Andreyev, N. D. 1967. "The Intermediate Language as the Focal Point of Machine Translation." Machine Translation, ed. by A. D. Booth, Amsterdam, North Holland, pp. 1-27.

Automatic Language Processing Advisory Committee, 1966. Language and Machines: Computers in Translation and Linguistics. Publication 1416. Washington, D. C., National Academy of Sciences, National Research Council.

Bar-Hillel, Yehoshua 1970. "Position Paper on Machine Translation in 1970." University of Texas at Austin, Multilith.

Bel'skaja, I. K. 1969. Jazyk čeloveka i mašina. 2 vols. Moscow, Publishing House of Moscow University.

Garvin, Paul L. 1962. "Computer Participation in Linguistic Research." Language 38.385-9. Reprinted in Garvin forthcoming.

Garvin, Paul L. 1966. "Some Comments on Algorithm and Grammar in the Automatic Parsing of Natural Languages." Mechanical Translation 9.2-3. Reprinted in Garvin forthcoming.

Garvin, Paul L. 1967. Machine Translation - Fact or Fancy. Datamation. April 1967.

Garvin, Paul L. 1968. "The Place of Heuristics in the Fulcrum Approach to Machine Translation." Lingua 21.162-82. Reprinted in Garvin forthcoming.

Garvin, Paul L. Forthcoming. On Machine Translation. Selected Papers, The Hague, Mouton.

Hays, David G. 1964. "Dependency Theory: A Formalism and Some Observations." Language 40.511-25.

Hockett, Charles F. 1968. The State of the Art. The Hague, Mouton.

Iordanskaja, L. N. 1967. Avtomatičeskij sintaksiceskij analiz. Vol. 2. Mezsegmentnyj sintaksiceskij analiz. Novosibirsk. "Nauka" Publishing House. Siberian Branch.

Mel'čuk, I. A. 1964. Avtomaticeskij sintaksičeskij analiz. Vol.1 Obščije principy. Vnutrisegmentnyj sintaksičeskij analiz. Novosibirsk. Editorial and Publishing Division of the Siberian Branch of the Academy of Sciences of the USSR.

Mel'čuk, I. A. and A. K. Žolkovskij 1970. "Towards a Functioning 'Meaning-Text' Model of Language." Linguistics 57.10-57. The Hague, Mouton.

Oettinger, Anthony G. and Murray E. Sherry 1961. "Current Research on Automatic Translation and Predictive Syntactic Analysis." Proceedings of the National Symposium on Machine Translation. Held at UCLA, February 2-5, 1960, ed. by H. P. Edmundson, Englewood Cliffs, N.J., Prentice-Hall, 173-82.

Weinreich, Uriel 1966. "Explorations in Semantic Theory." Trends in Linguistics, ed. by Thomas A. Sebeok, Vol. 3: Theoretical Foundations. The Hague, Mouton. pp. 395-477.