The Future of MT is Now and Bar-Hillel was (almost entirely) Right

Elliott Macklovitch
Centre for Information Technology Innovation (CITI)
1575 Chomedey Boulevard
Laval, Quebec, Canada

1. Introduction

A seer, according to one dictionary definition, is a person who can predict events or developments; someone, in other words, who can foretell the future. As far as I know, Yehoshua Bar-Hillel never actually claimed to be a seer. He was first and foremost a scientist (a logician and philosopher of science), not a prophet. But in his many writings on machine translation in the 1950's and early 60's, Bar-Hillel did make a number of clear and important predictions about the future of MT; and based upon his assessment of what was possible in the field, he suggested certain priorities for future development. In this paper, I shall summarize some of the more striking of those early predictions and attempt to evaluate them in light of the current state of the art of MT, some thirty-five years later. As the title of my paper suggests, I believe that Bar-Hillel's predictions have been largely borne out; and yet his suggestions for more reasonable ways of putting the power of computers at the service of translators have gone by and large unheeded. In the final section of this paper, I will describe a system currently under development at the CITI called TransCheck, which is a novel kind of support tool for human translators. The approach to translation automation that underpins TransCheck is wholly consistent with Bar-Hillel's own call for a modest and judicious use of mechanical aids as an alternative to classical MT, and as such I am quite confident he would have endorsed it. At the CITI, we are convinced it is the way of the future.

2. On the Nonfeasibility of Fully Automatic, High Quality Translation (or FAHQT)

The celebrated argument that Bar-Hillel published in 1960 on the nonfeasibility of fully automatic, high quality machine translation is nearly as well-known as the acronym he coined for it, and certainly does not bear repeating in extenso. As Bar-Hillel himself pointed out, the argument does not even concern translation proper; rather, it demonstrates the inescapable need for extra-linguistic knowledge in order to determine the meaning - and *hence* the translation - of a polysemous word like "pen" in a perfectly innocuous sentence like "The box is in the pen". As it turns out, the knowledge required to disambiguate “pen” in this context concerns not just common-sense expectations about the relative sizes of writing instruments and children’s enclosures; also required is the ability to reason about this knowledge, viz. if an object X is *in* some other object Y, X is normally smaller than Y.[1] These kinds of reasoning capabilities and extra-linguistic knowledge were obviously not available to existing machine translation systems thirty-five years ago. More important for us today, however, is the manner in which Bar-Hillel reacted to the suggestion that such information might eventually be put at the disposal of later MT systems:

> “What such a suggestion amounts to, if taken seriously, is the requirement that a translation machine should not only be supplied with a dictionary but also with a universal encyclopedia. This is surely utterly chimerical and hardly deserves any further discussion.” (Bar-Hillel [2], p.176)

Still, Bar-Hillel did accord the suggestion some further attention, elaborating on why he considered the idea of equipping an MT system with a universal encyclopedia so preposterous.

> “The number of facts we human beings know is, in a certain very pregnant sense, infinite. Knowing, for instance, that at a certain moment there are exactly eight chairs in a certain room, we also know that there are more than five chairs, less than 9, 10, 11, 12 and so on, ad infinitum... We know all these additional facts by inferences we are able to perform, and it is clear that they are not, in any serious sense, stored in our memory.” (Bar-Hillel [2], p.177)

At the core of Bar-Hillel’s argument, therefore, is not just the fact that translation routinely requires encyclopedic knowledge; i.e. knowledge not about the properties of language but about the real world. Rather, the nub of the problem is the fact that humans can instantaneously access **infinite** amounts of such knowledge, as a result of their ability to infer. And while Bar-Hillel was able to envision a translating machine that might eventually perform certain inferences, he found it inconceivable that such a machine would be able to do so in the same spontaneous manner or under the same circumstances as any intelligent human can, and as translators unconsciously do all the time.

All in all, Bar-Hillel’s celebrated argument covers no more than a page in the original text; and yet there are at least two quite different ways in which it can be interpreted. On the first and narrower interpretation, it is a logically impeccable demonstration of the unattainability of FAHQT, seen not as a straw man but rather as the actual, though often unstated, goal of many of the groups working in MT at that time. For the purposes of this demonstration, it is not necessary for Bar-Hillel to quantify the frequency with which

---

1. Hence, given that boxes are generally larger than writing instruments, the box in question is most likely within a kind of enclosure, such as a playpen. Barring indications to the contrary, this at least is a more plausible interpretation of the sentence than the one in which “pen” is a writing instrument.

sentences like those of his simple example occur in various types of documents; as long as any instance of this kind of ambiguity appears, requiring extra-linguistic knowledge and/or reasoning for its resolution, it is sufficient to scuttle the fully automatic part of the FAHQT ideal. In other words, as soon as a human post-editor has to be called upon to resolve the ambiguity of one polysemous word like "pen", the MT system involved necessarily becomes less than fully automatic. More interestingly, perhaps, once the need for human intervention in the translation process is admitted and accepted, one can envisage a wide range of possible modes of cooperation, or divisions of labour, between man and machine. To find the most productive or cost-effective arrangement was, for Bar-Hillel, an empirical question that demanded careful study. Unfortunately, it was not a question that seems to have interested the majority of MT workers at the time, many of whom continued to confuse the aims and methods of MT as an area of fundamental research with those of MT as a practical endeavour. Bar-Hillel recognized the validity of both pursuits; what he deplored was their confusion. We return to this issue in section 6 below.

3. The Future of MT

There is another possible interpretation of Bar-Hillel's famous demonstration, one that is broader and more pessimistic, in which he can be seen as arguing for the unattainability of FAHQT not merely in the short term, but altogether. That this in fact was the real intention of his argument is suggested somewhat passingly in the original article, where Bar-Hillel states "that no existing or *imaginable* program will enable an electronic computer to determine the meaning of the word "pen" in the given sentence..." (p.175; emphasis added). But he reinforces this interpretation and makes it perfectly clear in a piece he published in 1962 in the Times Literary Supplement, entitled "The Future of Machine Translation." In it, Bar-Hillel declares that "MT has reached an impasse from which it is not likely to emerge without a radical change in the whole approach..."

> "...with all the progress made in hardware, programming techniques and linguistic insight, the quality of fully autonomous mechanical translation, even when restricted to scientific and technological material, will *never* approach that of qualified human translators and therefore MT will only under very exceptional circumstances be able to compete with human translation." (Bar-Hillel [3], p.182; again, my emphasis)

To justify this pessimistic prognostic, Bar-Hillel points to the syntactical (or scoping) ambiguity of the adjective in a simple phrase like "slow neutrons and protons." The example he chooses is different, but in essence, it is the same concise argument as he invoked in the celebrated 1960 article, based on the fact that human translators routinely make use of their vast background knowledge, no counterpart of which he says could conceivably stand at the disposal of computers. Bar-Hillel leaves it to the reader to generalize the argument to the "innumerable semantical ambiguities which nothing but plain, factual knowledge or considerations of truthfulness and consistency will resolve..." (ibid [3], p.182).

The final conclusions Bar-Hillel arrives at in "The Future of Machine Translation" must have appeared apocalyptic to MT workers at the time; those of the ALPAC report pale in comparison:

> "I would say that there is no prospect whatsoever that the employment of electronic digital computers in the field of translation will lead to any revolutionary changes. A complete automation of the activity is wholly utopian.... The quicker this is understood, the better are the chances that more attention will be paid to finding efficient ways of improving the status of scientific and technological translation –... including judicious and modest use of *mechanical aids*." (Bar-Hillel [3], p.183)

4. The Future is Now

It has now been over thirty years since Bar-Hillel published these provocative views on the future of machine translation. Of course, the future is by definition boundless. But suppose we decided to call in the bets *today*, and arbitrarily decreed that the future is now. How would Bar-Hillel's predictions fare in light of the current state of machine translation? In particular, some of the questions we would like to consider are the following: Was Bar-Hillel right in declaring that high quality and full automation are almost always mutually exclusive in machine translation?[2] Has time shown that he was correct, or just short of imagination, in asserting that no *imaginable* program would ever allow a computer to perform sense disambiguations on polysemous words like "pen" in underdetermined contexts like that of his famous example? Is it reasonable to characterize as exceptional those attested cases in which MT *has* been able to compete favourably with human translation? More generally, has machine translation yet emerged from the impasse in which Bar-Hillel saw it miring in 1962? And finally, has anything like a radical change occurred in the dominant approach to the whole problem of translation automation, like the one Bar-Hillel called for in the early sixties?

It is no easy task to attempt to summarize the current state of the art in a field as diverse and ebullient as MT is today; and so inevitably, my personal assessment will appear to some people as incomplete or tendentious. Still, it seems to me that there are certain indisputable facts about the current state of machine translation which anyone, regardless of his or her particular bias, must necessarily accept. And one of these is that the growing use of computers over the last ten or fifteen years has not been accompanied by anything like the revolution in translation that early MT practitioners had hoped for. On this point at least, Bar-Hillel has turned out to be undeniably correct. While accurate data

2. "Those who are interested in MT as a primarily practical device must realize that full automation of the translation process is incompatible with high quality." (Bar-Hillel [5], p.167)

on the world translation market are notoriously difficult to obtain, I am aware of no studies or estimates that accord machine translation more than 5% of that market – and that figure appears to me to be generous. Why is it that in 1995 MT still occupies such a marginal role in professional translation worldwide? The unavoidable answer, in my view, is that in the vast majority of translation situations, currently available MT systems are simply not able to satisfy users' needs and/or expectations.[3] And in most cases, this is because a significant gap continues to exist between the quality of the "raw" machine output and the quality requirements of the end users, such that the cost of the post-editing necessary to make the machine output usable turns out to be prohibitive. Here too, it would seem that Bar-Hillel was quite correct in asserting a general incompatibility between high quality and fully automatic machine translation.

What about those cases where MT does prove cost-effective and so can compete favourably with human translation? These would appear to fall into one of two categories. In the first, the application domain is so narrow that the developers have been able to craft a specialised system capable of producing translations of acceptable quality, principally because the restrictions on the language used in the domain effectively reduce the full range of linguistic ambiguities to manageable proportions.[4] In the other class of successful MT applications, the end users agree to accept less than top quality translations – either because this is the only way of obtaining a translation at all, or as a cost saving measure for texts that will not receive wide or public distribution. Overall, however, the translation situations in which either of these two conditions obtains are relatively rare, or as Bar-Hillel qualified them, "exceptional".

5. MT and AI

Turning now to the more technical question of the capacity of current MT systems to perform lexico-semantical disambiguations like those that Bar-Hillel illustrated with his famous "box in the pen" example[5], let me first remark how appropriate it is that this question be considered at a symposium on the foundations of artificial intelligence. For the issue raised by such examples directly links machine translation to the broader fields of AI and natural language understanding. As mentioned above, the knowledge required to disambiguate "pen" in Bar-Hillel's example is not primarily linguistic; rather, it involves common-sense expectations about real-world objects, as well as the ability to reason and

3. Other explanations are possible, at least in principle. For example, adequate MT systems might exist but simply be too expensive for the majority of potential users. Currently, however, this is certainly not the case; on the contrary, inexpensive PC-based systems have made MT more widely available than ever. The problem lies in the deficiency of the technology, regardless of what one is prepared to pay for it.

4. Or, in the absence of naturally occurring sublanguages (like that exploited by Canada's METEO system), artificial restrictions may be imposed on the language in which technical documentation is drafted, in order to simplify the input to an MT system further downline. There has been an increase in the use of controlled language for MT in recent years.

5. Or, for that matter, syntactical disambiguations like the one he illustrated with the "slow neutrons and protons" example, since both require a capacity to reason over extra-linguistic knowledge.

draw inferences over this kind of extra-linguistic information. Here, it would be disingenuous not to acknowledge the substantial progress that has been achieved since the early 1960's. For one thing, the main thrust of Bar-Hillel's argument is now a generally accepted truth in the field: everyone involved in MT research and development today recognizes that in order to correctly translate unrestricted text, an MT system must be supplied not just with dictionaries and grammars, but with some portion of encyclopedic knowledge as well. Over the last three decades, numerous researchers have attempted to design MT systems in such a way as to accommodate this fundamental fact about translation. This is not the place for an exhaustive survey of these efforts; but we would like to attempt a general evaluation of their overall success.

Terry Winograd is a prominent pioneer in artificial intelligence, whose SHRDLU system in the late 60's is still cited as providing a striking illustration of natural language understanding. In 1984, Winograd published an article in Scientific American entitled "Computer Software for Working with Language," in which he raised the following question:

> "Is there software that really deals with meaning – software that exhibits the kind of reasoning that a person would use in carrying out tasks such as translating, summarizing or answering a question? Such software has been the goal of research projects in artificial intelligence since the mid-1960's, when the necessary computer hardware and programming techniques began to appear even as the impracticability of machine translation was becoming apparent..." (Winograd [6], p.136)

Winograd briefly reviews some of the better known NLP projects in the 70's and 80's that attempted to encode knowledge of the world in a form that a program could use to draw inferences (e.g. Schank's scripts). The problem, he points out, is that all of these programs work only in highly limited, somewhat artificial domains, and it is not at all obvious how – or whether – they can be extended. In his conclusion, the reply he himself provides to his earlier question is certainly not very encouraging; in fact, it is quite reminiscent of the conclusions that Bar-Hillel had arrived at twenty-five years earlier.

> "The limitations on the formalization of contextual meaning make it impossible at present – and conceivably forever – to design computer programs that come close to full mimicry of human language understanding." (Winograd [6], p.142)

But does this necessarily entail that computers will never be able to translate adequately? Perhaps it might still be possible to design programs that would allow a machine to produce relatively high quality translations without having to simulate the full extent of human language understanding; if not all the time, at least often enough for human post-editing to be cost-effective. Actually, Bar-Hillel himself thought so for a time. Referring back to the expectations he first held for MT, he observed:

> "I knew then that nothing corresponding to items (3) and (4) [i.e. good general

> background knowledge and expertness in the field] could be expected of electronic computers but ... entertained some hopes that by exploiting the redundance of natural language texts better than human readers usually do, we should perhaps be in a position to enable the computers to overcome, at least partly, their lack of knowledge and understanding." (Bar-Hillel [4], p.213)

For Bar-Hillel, however, those hopes were short-lived. In 1962, he wrote:

> "Though it is undoubtedly the case that some reduction of ambiguity can be obtained through better attention to certain formal clues ... it should by now be perfectly clear that there are limits to what these refinements [of purely formal methods] can achieve, limits that definitely block the way to autonomous, high-quality, machine translation." (ibid, p.213)

Other AI researchers have been less pessimistic than Winograd and Bar-Hillel. Sergei Nirenburg, a longtime stalwart of the knowledge-based approach to machine translation, is one of these. He and Kenneth Goodman published a brave article [7] in 1990, in which they systematically take up many of the criticisms that are frequently directed against meaning-based (or interlingual) MT, and expose the double standards and many of the inconsistencies in the arguments that are often employed to repudiate and even disparage the interlingual paradigm. But in the face of widespread and persistent scepticism about the overall feasibility of developing a complete set of language-independent meaning representations for all possible linguistic expressions in all human languages, Nirenburg and Goodman can do little more than protest that "We are making inroads into these and other difficult areas."(p.184) In the end, they are forced to recognize that the only real way to silence the critics and convince the sceptics is to demonstrate the practical utility of the approach by actually building a production system prototype.

Until such time as we see the results of such a prototype, however, it seems to me we have every reason to remain sceptical. For as Winograd and others have been careful to point out, all the interlingual systems that have been documented or demonstrated to date have been more or less toy systems.[6] To the extent that they have been able to simulate some measure of inferencing in order to arrive at a correct translation, these AI-based systems may perhaps have shown Bar-Hillel to be wrong, at least in a narrow sense.[7] Nonetheless, I doubt that Bar-Hillel would have been very impressed with such demonstrations. For though an accomplished theorist in other domains, he never remained solely a theorist in matters of translation, but always exhibited a genuine

6. As Hutchins [15] puts it in the summary of his chapter on AI-based MT systems: "It needs to be stressed, however, that none of the AI workers are expecting their work to result in the near future in 'operational' MT systems."(p.284)
7. Though even this is not entirely obvious. As mentioned above, Bar-Hillel was able to envision a machine capable of inferencing; what he found more difficult to imagine was "a scheme which would make a machine perform such inferences in the same or similar circumstances under which an intelligent human being would perform them." (Bar-Hillel [2], p.177) And of course, he also discounted any ad hoc procedure, mounted solely for the case at hand, "whose futility would show itself in the next example." (ibid, p.174)

concern with the practical problems of translation in the real world. And here, demonstrations of the possibility of machine reasoning within toy domains are of little consequence. The major obstacle to fully automatic, cost-effective machine translation today remains exactly the same as it was 35 years ago, to wit, the vast and unpredictable **range** of the knowledge that is required to allow a machine to achieve an understanding of a source text that is sufficient for the purposes of translation. As Hutchins and Somers put it in their Introduction to Machine Translation:

> "The problem for MT systems is that it is at present impossible in practice to code and incorporate all the potential (real world) knowledge that might be required to resolve all possible ambiguities in a particular system, even in systems restricted to relatively narrow ranges of contexts and applications. Despite advances in Artificial Intelligence and in computing technology, the situation is unlikely to improve in the near future: the sheer complexity and intractibility of real world knowledge are the principal impediments to quick solutions." (Hutchins & Somers [8], p.93)

Hence, it seems fairly safe to say that Bar-Hillel would not have substantially modified his views on the feasibility of FAHQT, had he lived beyond 1975 to witness some of the impressive successes in artificial intelligence. Without wanting to diminish the import of the advances of the last twenty years, it does appear to be generally true that AI and its daughter discipline MT are similar, in that their most impressive applications have been achieved in relatively restricted domains; in both fields, depth of understanding and breadth of coverage remain by and large mutually exclusive. Moreover, Bar-Hillel was already familiar with some of the early successes of AI research, e.g. the work on pattern recognition (or perceptrons) and programs to play checkers. Given his remarkable foresight, we would be rash to disregard the warning he formulated in 1962: "it would be disastrous to extrapolate from these primitive exhibitions of artificial intelligence to something like translation." (Bar-Hillel [4], p.214)

6. A Radical Change of Approach

In summary: an objective assessment of the state of the art in MT would seem to suggest that the field is still in fact mired in the same impasse that Bar-Hillel described in the early 60's. The question we now want to consider is whether we have begun to see anything like a radical change in the dominant approach to the whole problem of translation automation that Bar-Hillel called for back then. My own inclination is to answer in the negative (with an important qualification, to be specified below). Although they may not publicly admit it, the researchers on most current machine translation projects are still striving to achieve FAHQT, just as they were in Bar-Hillel's time. Granted, the techniques we now employ may have evolved; new approaches, or perhaps whole new paradigms, have emerged over the last decade, which appear at first glance to be radically innovative, e.g. Statistical Machine Translation and Example-Based Machine Translation.[8] But the goal remains essentially the same now as it was then: to develop a fully automatic translating machine capable of producing target texts of a quality comparable to that of a human translator – in other words, a translating robot. No one can object to this as an entirely worthwhile pursuit for a long-term research programme. The problem is that many of those working toward this objective continue to maintain the unspoken assumption that their research efforts – even if they do eventually fall short of their ultimate goal – will nevertheless prove useful in providing short-term solutions to the practical problems besetting a translation profession that is overwhelmed and unable to meet the demand for its services. This is far from a self-evident truth, however, as we will argue below. In his pre-ALPAC articles, Bar-Hillel deplored just this confusion between the aims and methods of MT as an area of fundamental research and MT as a practical endeavour. Perhaps it is time we finally learned the lessons of history and accepted the fact that classical MT – and by this I mean **any** approach in which the initiative in the translation process is given over to the machine so that it can autonomously produce a target version of the source text – will never contribute more than marginally to satisfying the ever-growing demand for translation, or at least not for many, many years to come.

For those who remain concerned with the practical problems of working translators, does Bar-Hillel indicate the direction he thought that a genuinely radical change of approach should take? The citation reproduced at the end of section 3 above does provide one clue, where Bar-Hillel mentions the possibility of a "judicious and modest use of *mechanical aids*." (Here the emphasis is in the original.) In his "Aims and Methods in Machine Translation" [5], Bar-Hillel is even more explicit:

> "The only reasonable aim, then, for short-range research into MT seems to be that of finding some machine–post-editor partnership that would be commercially competitive with existing human translation, and then try to improve the commercial competitiveness of this partnership by improving the programming in order to delegate to the machine more and more operations in the total translation process which it can perform more effectively than the human post-editor." (p.172)

A partnership between machine and human translator/post-editor that takes as its starting point a judicious and modest use of mechanical aids: this would seem to point to what is now generally called machine-aided human translation (or MAHT), as distinct from classical MT or human-aided machine translation[9]. MAHT has not been a popular avenue of research over the last thirty years. To be sure, it has had its isolated champions; most notably, perhaps, Martin Kay (cf Kay [11]). But generally speaking, MAHT has not

8. The standard reference for SMT is Brown et al. [9], and for EBMT, Sato and Nagao [10]. The former approach, of course, does have historical antecedents, on which Bar-Hillel also had some very interesting things to say. See especially Bar-Hillel [5], p.171.

9. Again, the distinction between MAHT and HAMT may be framed in terms of which of the two – man or machine – retains the initiative in the translation process. In MAHT, it is the human translator, and the machine is viewed as a tool that may be called upon to amplify human capabilities, but only on such tasks that can be automated reliably.

succeeded in attracting anything like the attention and funding that has been poured into classical MT, no doubt because, as Bar-Hillel suggested, the latter is perceived as an intellectually more challenging pursuit. Only in the last few years, as more and more researchers have started exploring corpus-based approaches to NLP, has MAHT begun to receive the attention it deserves.[10] MAHT is the philosophical cornerstone of the CITI's machine-aided translation program. In this final section, I would like to illustrate our approach to MAHT and briefly discuss why we consider it to be a radical departure from traditional responses to the problem of translation automation, by describing one of the translator support tools we are currently developing.

The project in question is called TransCheck, and it is documented more fully in Macklovitch [13]. As its name suggests, TransCheck is intended to be used as a translation checker, somewhat like a spelling checker. But where the latter verifies certain (orthographic) properties of a single monolingual text, TransCheck is designed to validate certain properties that normally must hold between *two* texts that are in a translation relation. To do so, the system incorporates an alignment algorithm that automatically links segments (currently sentences) in the target text to their corresponding segments in the source text. Once the draft translation is completed, the translator submits the two language files to TransCheck, and the system then verifies the aligned segments to ensure that they do not contain any deceptive cognates, calques, illicit borrowings, or certain other commonly occurring translation errors. When it does detect an error described in its database of prohibited translations, TransCheck flags it and provides the user with information on the correct target language form that should be used. Preliminary tests of the first prototype have produced encouraging results, confirming the general viability of a translation checker based on a sentence alignment program and a part-of-speech tagger; again, see [13] for further discussion. Ultimately, of course, it is hoped that TransCheck will be able to detect more subtle types of translation errors, and that end users will be able to modify the contents of the database so that it reflects their own translation norms. Here, however, I would like to focus on another projected extension to TransCheck which has not yet been fully implemented, but which illustrates, I think particularly well, the interest of what Bar-Hillel called a judicious and modest use of mechanical aids.

One of the most boring tasks for a translation reviser (who may be the translator himself) is to verify that all numerical expressions in a source text have been correctly rendered in the target. Texts in domains such as economics or statistics can be packed full of such expressions, and the smallest error in one digit is tantamount to a serious mistranslation: not only is it extremely embarrassing for the translator, but it can

---

10. For more on this paradigm shift, see Isabelle [12], where the author outlines some of the reasons why the classical rule-based approach to MT has produced so few useful results in the way of translator support tools, and why, on the contrary, the corpus-based approach seems to lend itself so well to MAHT.

undermine the credibility of the entire text. The reason why human revisers find this task so boring is that the "translation" of numerical expressions is so straightforward that it requires almost no intellectual effort (although between certain languages, there may be some minor differences of syntax); and yet every number must still be checked for the possibility of an error of transcription that does occasionally occur. Is this not exactly the kind of mechanical operation for which computers are better suited than humans? On the basis of the aligned sentences it has paired, a system like TransCheck should be able to verify that for every source segment containing a numerical expression, the corresponding target segment contains the equivalent numerical expression; and where it doesn't, that pair should be brought to the reviser's attention. In actual fact, the problem is not as trivial as I have suggested here[11]; nonetheless, we are convinced that even a rudimentary numerical component within TransCheck should be able to validate a large proportion of the numerical expressions in most texts. Moreover, an important characteristic of this approach to translation validation is that even if the system is less than fully exhaustive, whatever errors it does detect will still contribute to improving the quality of the final text; just as the errors detected by a spelling checker improve the final product, even though none of those systems is fully exhaustive either. Notice, however, that this is not generally the case with classical MT systems: the kind of partially correct translations generated by such systems do not always result in a reduction of the translator's workload, as many disenchanted MT users have testified over the years.

7. Conclusion

TransCheck is one of a new generation of translation support tools currently being developed at the CITI, all of which are based on the concept of translation analysis, as opposed to the classical MT approach of translation generation.[12] In our view, translation analysis does constitute a radical departure from traditional responses to the whole problem of how best to automate the translation process, although it is not as yet anywhere close to becoming the dominant approach in the field. In its favour, it has allowed for the development, within a remarkably short period of time, of promising new types of translation support tools, which are wholly consistent with the modest and practical strategy that Bar-Hillel advocated over 30 years ago. Whether these tools will actually live up to their promise and prove more useful to translators than classical MT has to date, only the future will tell. And no one but a seer can predict the future.

---

11. To illustrate just a few of the complications frequently encountered, there is the obvious problem of numerical expressions that are written according to different standards, e.g. "7 p.m." vs. "19 h"; which is why all such expressions will have to translated into a normalized form before being compared. Less obviously, the text in one language may use a numeral, e.g. the date "1994", where the translation properly refers to the same period by means of a non-numerical noun phrase like "last year".

12. See Isabelle et al. [14] for a description of some of the CITI's other projects, and a fuller discussion of the differences between translation analysis and translation generation.

References

[1] Bar-Hillel, Yehoshua. Language and Information: Selected Essays on their Theory and Application, Addison-Wesley Publishing Company, Inc., 1964.

[2] Bar-Hillel, Y. “A Demonstration of the Nonfeasibility of Fully Automatic High Quality Translation,” reproduced in [1], pp. 174-179.

[3] Bar-Hillel, Y. “The Future of Machine Translation,” reproduced in [1], pp. 180-184.

[4] Bar-Hillel, Y. “Four Lectures on Algebraic Linguistics and Machine Translation,” in [1], pp. 185-218.

[5] Bar-Hillel, Y. “Aims and Methods in Machine Translation,” in [1], pp. 166-173.

[6] Winograd, Terry. “Computer Software for Working with Language,” Scientific American, no.251(3), September 1984.

[7] Nirenburg, Sergei & Kenneth Goodman. “Treatment of Meaning in MT Systems,” Proceedings of the Third International Conference on Theoretical and Methodological Issues in Machine Translation, Austin, Texas, June 1990.

[8] Hutchins, W. John & Harold L. Somers. An Introduction to Machine Translation, Academic Press, 1992.

[9] Brown, Peter et al. “A Statistical Approach to Machine Translation,” Computational Linguistics, no.16:2, pp.79-85, 1990.

[10] Sato, Satoshi & Makoto Nagao. “Toward Memory-Based Translation,” Proceedings of COLING-90, vol.3, pp.247-252, 1990.

[11] Kay, Martin. “The Proper Place of Men and Machines in Language Translation,” CSL-80-11, Xerox PARC, 1980.

[12] Isabelle, Pierre, “Machine-Aided Human Translation and the Paradigm Shift,” Proceedings of the Fourth Machine Translation Summit, Kobe, Japan, July 1993.

[13] Macklovitch, Elliott. “Using Bi-textual Alignment for Translation Validation: the TransCheck system,” Proceedings of the First Conference of the Association for Machine Translation in the Americas, Columbia, Maryland, pp.157-168, October 1994.

[14] Isabelle, Pierre et al. “Translation Analysis and Translation Automation,” Proceedings of the Fifth International Conference on Theoretical and Methodological Issues in Machine Translation, Kyoto, Japan, 1993.

[15] Hutchins, W. John. Machine Translation: Past, Present, Future, Ellis Horwood, Chichester, 1986.