

[MT Summit III, Washington, DC, USA, July 1991]

## Where Do Translators Fit Into MT?

### Panel Discussion

**Chair: Alex Gross, Cross-Cultural Research Projects, USA**

Claude Bedard, Traductix Translation Consulting, Canada

Harald Hille, United Nations

Martin Kay, XEROX and Stanford University, USA

Frederick Klein, USA

Sergei Nirenburg, Carnegie Mellon University, USA

#### 1 Introduction, Alex Gross

1. At the last MT Summit, Martin Kay stated that there should be “greater attention to empirical studies of translation so that computational linguists will have a better idea of what really goes on in translation and develop tools that will be more useful for the end user.” Does this mean that there has been insufficient input into MT processes by translators interested in MT? Does it mean that MT developers have failed to study what translating actually entails and how translators go about their task? If either of these is true, then to what extent and why? New answers and insights for the MT profession could arise from hearing what human translators with an interest in the development of MT have to say about these matters. It may well turn out that translators are the very people best qualified to determine what form their tools should take, since they are the end users.

2. Is there a specifically “human” component in the translation process which MT experts have overlooked? Is it reasonable for theoreticians to envision setting up predictable and generic vocabularies of clearly defined terms, or could they be overlooking a deep-seated human tendency towards some degree of ambiguity--indeed, in those many cases where not all the facts are known, an inescapably human reliance on it? Are there any viable MT approaches to duplicate what human translators can provide in such cases, namely the ability to bridge this ambiguity gap and improvise personalized, customized case-specific subtleties of vocabulary, depending on client or purpose? Could this in fact be a major element of the entire translation process? Alternately, are there some more boring “machine-like” aspects of translation where the computer can help the translator, such as style and consistency checking?

3. How can the knowledge of practicing translators best be integrated into current MT research and working systems? Is it to be assumed that they are best employed as prospective end-users working out the bugs in the system, or is there also a place for them during the initial

planning phases of such systems? Can they perhaps as users be the primary developers of the system?

4. Many human translators, when told of the quest to have machines take over all aspects of translation, immediately reply that this is impossible and start providing specific instances which they claim a machine system could never handle. Are such reactions merely the final nerve spasms of a doomed class of technicians awaiting superannuation, or are these translators in fact enunciating specific instances of a general law as yet not fully articulated? Since we now hear claims suggesting that FAHQT is creeping in again through the back door, it seems important to ask whether there has in fact ever been sufficient basic mathematical research, much less algorithmic underpinnings, by the MT Community to determine whether FAHQT, or anything close to it, can be achieved by any combination of electronic stratagems (transfer, AI, neural nets, Markov models, etc.). Must translators forever stand exposed on the firing line and present their minds and bodies to a broadside of claims that the next round of computer advances will annihilate them as a profession? Is this problem truly solvable in logical terms, or is it in fact an intractable, undecidable, or provably unsolvable question in terms of “Computable Numbers” as set out by Turing, based on the work of Hilbert and Gödel? A reasonable answer to this question could save boards of directors and/or government agencies a great deal of time and money.

#### 2 Claude Bedard’s Statement

##### 2.1 The translators’ absence

By suggesting that there should be “*greater* attention to empirical studies of translation” for MT purposes, Martin Kay is no doubt making a polite understatement. What he really means is that the world of MT should stop disregarding this topic.

MT researchers are widely known not to be very translation-conscious, I remember attending an MT con-

ference where one highlight turned out to be a presentation by a professional translator about the real-life problems of translation. Computational linguists present nearly fell off their chairs.

Although they remember Warren Weaver's mistake of seeing translation as a deciphering process, MT researchers still cannot help seeing translation as a rather straightforward, albeit difficult, process of linguistic transcription<sup>1</sup>; given enough linguistic and (since Bar-Hillel's the-box-is-in-the-pen example) cognitive information, they should get it right somehow.

Nor is translation-consciousness perceived as something particularly useful or meaningful by MT researchers. I remember again asking a panel of such people if they would find some use for a professional translator on their development teams. Very few could bring themselves to answer in the affirmative.

This being said, let's be fair: MT researchers rightly feel that there are so many difficult basic problems to address -- not the least being source-language parsing. Once these problems are reasonably solved, they would probably be quite willing to move on to more sophisticated issues involving translation know-how.

This may be a sensible attitude, especially if we consider that no one has really given those people a clue about the potential role of translators in MT research and development. In fact, if I were to ask the above question to a panel of translators, chances are that I wouldn't get any exciting answers either. By and large, the translator community has shied away from getting their minds into MT issues, and saying that they have been held at arm's length by computational linguists is a poor excuse indeed for staying comfortably aloof. I believe the time has come both for the MT community to genuinely acknowledge that translators have an important role to play, and for translators to recognize that as long as they are not part of the solution, they are part of the problem.

## 2.2 What kind of contribution?

The traditional contribution of translators so far has been mainly to work with MT systems as finished products. This is hardly trivial, mind you, because using such systems effectively in real-life translation situations involves developing a special kind of expertise which MT developers are not known to provide, I suspect that a new generation of "smart" MT users is on the rise; those users will have the ability to customize (more deeply than by mere dictionary manipulation) an MT system successfully for their own application, without losing control of the system's performance; this, in turn, will encourage designing systems to give the user more room for initiative. The smart MT user will also develop a flair for deciding about the applicability of MT to any particular situation.

---

<sup>1</sup> I will use this term, which isn't meant to be offensive, but to draw attention to a certain reality.

But what about translator involvement in MT system development? So far, their contribution has been limited to helping build lexical and structural transfer rules, as exemplified in the TAUM Aviation project and also in the Systran project. But their work is still confined within a linguistic framework predefined by non-translators.

I strongly believe that translators can and should make a much more active contribution, because they bring in at least two vital ingredients.

First, translators are rather pragmatic, result-oriented people. This is particularly true in comparison to computational linguists, who are trained to think in terms of Chomsky-style formal elegance (a tremendous liability in the context of real-life linguistic applications), who will rather favor more scientifically exciting approaches, and who tend to get bored when they stop discovering. Translators can use their own practitioner's intuitions to help devise empirical mechanisms which, at the price of being somewhat inelegant and unworthy of scientific credibility, may do useful work. This is important if we want MT to somewhat pay for itself on the way to the Ultimate Solution.

Second, translators are the ones who are "translation-conscious." I keep saying that an in-depth awareness of real-life manual translation is essential to success in mechanizing translation; without this knowledge you may do fine with linguistic transcription in protected applications, but you're in trouble when it comes to real-life MT applications.

In fact it is translators' responsibility to provide answers to questions such as: What lies between linguistic-transcription-level MT output and the finished, professional-quality translation? Which ingredients? Which processes? How does the translator operate, and what does this suggest about how MT systems could possibly operate? What forms of automation are most useful to the translator in his work?

I feel it useful at this point to say more about translation consciousness. I would subdivide this concept into three distinct areas:

- Basic translation training. First of all, being translation-conscious means more than merely being bilingual. It takes several years of university training to fully realize the range of problems which confront the translator, and to learn about the solutions and techniques to solve them.
- Mundane translation experience. At the end of basic training, you still don't know about the reality of real-world, real-text, production-oriented translation -- which I will call "mundane translation" for short, as opposed to more highly rated forms of translation such as literary and academic translation exercises. After some exposure to the work place reality, you come to realize that both the purpose and circumstances of mundane translation differ from those of university experience. This level of expertise, in contrast to basic training, has been lit-

tle studied. My own efforts at shedding some light on mundane translation suggest that the closer we come to real-life considerations, the more difficult (and unattractive) the research — but the more useful its results.

- Exploring the inner mechanics of translation. In neither of the previous cases does the translator stop to analyze the process beyond the point necessary for his native intelligence to acquire the necessary skills. University research is massively oriented towards high-level problems (literary, transcultural issues, etc.) which start far beyond this point indeed. I may also add that the task is anything but trivial: thinking about the human processes involved in translation with a view to automating them proves to be a rather exhausting intellectual experience, maybe because translation is such a broad, zigzag-fashion, multi-level, open-ended task.

It is translators' responsibility to bring their various levels of consciousness into useful play. At this very early stage, I can only try to suggest a few areas of exploration, my own vision being still fuzzy,

Overall I see two distinct purposes for translator involvement: improving the quality of MT output and improving the usefulness of MT technology for translator use. These purposes are discussed below.

### 2.2.1 Improving MT output quality

The most obvious area of translator involvement lies in improving MT output quality — regardless of its end use (be it postediting by translators or direct use by the target-language audience).

One item of translation-consciousness which comes into play here is that each translation situation is specific. In fact, a common beginner's mistake is to quickly jump from the text to be translated to external sources of information, thereby failing to take advantage or heed of the text's own internal resources and peculiarities -- thus running the risk of introducing solutions which, though not bad in themselves, are ill-suited to the situation.

This can be applied to two areas, which I present here as examples: text-level clues for disambiguation, and handling the situational aspect of terminology.

**Text-level clues for disambiguation** Every practicing translator knows that any given text covers only a subset of the wide world, and can thus provide very useful disambiguation clues about itself -- which are perfectly in tune with the sub-world in question.

Translators routinely take advantage of this fact. In particular, the less they know about the text's technical domain, the more they resort to using low-level contextual clues — thus performing what is sometimes called "translation by radar," performed without real high-level understanding, but with a coarse knowledge of things which seem to be true about the text's sub-world.

For the purpose of automation, these clues could simply be a collection of basic syntactic relationships

(subject-verb, verb-object, modifier-noun, etc.) present in the text, which can be seen as answers to the question "does such a thing exist in the text's sub-world?"

Here is how it could work. The MT system would perform a first parsing pass on the text, and collect all syntactic relationships which happen to be unambiguous. Then a second pass would take place, with this collection of proven relationships being used as a disambiguation dictionary to decide on the parsing of ambiguous cases. For example, the phrase "liquid oxygen tank" is ambiguous locally, but if we observe in other areas of the same text (or even related texts) such occurrences as "liquid oxygen will be carried by truck," we find one good reason (there is such a thing as liquid oxygen in this text) for choosing "liquid oxygen" over "liquid tank."

This proposition is only a sample of what translator insights can suggest. Incidentally, it comes in line with a recent tendency to recognize texts as ready reservoirs of information which can be tapped somehow for various purposes; one of them could certainly be applying the results on the same text for self-disambiguation.

**The situationality of terminology** A commonly held view is that if we could connect a "complete" term bank to an MT system and give it domain and semantic indications, this would take care of the terminology issue.

Experienced translators, hearing such a statement, would gently shrug with a mild smile. They do appreciate the usefulness of a massive term database as an input to a discerning and resourceful human thought process, but not as a direct input to an MT system. The reason is that technical terms are not clear-cut, precision-made, unambiguous, standardized parts which fit nicely in any context (even though technically correct). Here are, briefly, some details:<sup>2</sup>

- The whole of technical terminology is not too coherent; but it can be adapted into a coherent body in specific situations, thanks to the efforts of the writer or translator.
- Some individually valid terms don't fit well with others (their combination in certain contexts may create ambiguities, for example).
- The abundance of synonyms is such that choices are called for all the time. Such choices cannot be made only by choosing the first item on the list (or even the most standardized one), mainly because the choice should rather be guided by the above-mentioned considerations.
- Terms are routinely made up by the author of the source text, so that the equivalents must be created from scratch by the translator.
- Every now and then, a source-language term calls for a translation by something other than a term, and conversely.

<sup>2</sup> For full details, see Bedard, Claude, *La traduction technique: principes et pratique*. Linguatex, Montreal, 1986.

To summarize, it is surprising how much adaptation work must be poured into a technical translation, especially terminology-wise. Many equivalents are valid only for the text at hand (or even for one specific context — this is what I could call “translator’s equivalents” as opposed to “terminological equivalents,” which have a more generic application). How can automation deal with that? Translators will have to find out.

### 2.2.2 The mystery of idiomaticity

Different languages not only have different lexicons and grammars; they also don’t say things the same way. The acute consciousness of this fact about the languages he handles (keeping each different “style” to the language it belongs to) is one of the main differences between a translator and a bilingual person.

I remember looking with some initial satisfaction at an MT output which appeared rather good (i.e. few parsing or syntax errors), but discovering after sitting down to post-edit it that it was almost unusable. With almost every sentence, I had the feeling that “it could not really be said this way” in the target language; the raw translation simply “did not sound right,” and had to be largely rewritten. Not that there was anything literary or particularly stylish about the text, which was merely a discussion about accounting principles.

The problem was of course idiomaticity — one of the deepest mysteries in translation. The problem extends far beyond the limited realm of recognized idiomatic expressions (to kill two birds with one stone, to have a bat in the belfry, etc.), covering as it does the omnipresent reality of casual situations where “a native speaker would (or would not) say it this way.”

How does the translator operate with regards to idiomaticity? I would suggest that when starting on a sentence, the translator first tries the easiest path, i.e. a rather literal translation. If it happens to be adequate, he goes no further; but if not, he has to find an alternate path, which brings him to reword the sentence to a varying degree.<sup>3</sup> This process suggests several very interesting questions: when does the translator decide that an easy path is not adequate, and why? Then, as he looks for an alternate path, to what constraints is he responding? What resources does he call up? What degree of rewording is considered 1) barely adequate, 2) fully adequate, 3) unnecessarily high? Also, how does the translator (or any native speaker) acquire his knowledge about what is idiomatic or not? Here are some suggestions about the considerations involved:

- At the local level, the translator has to satisfy contrastive constraints of a lexical, syntactic, and stylistic nature.

<sup>3</sup> Some insights about the degree of freedom available for rewording are provided in Bedard, op. cit., pp. 177-179, and in “La reformulation: quelques astuces d’un traducteur” in *Circuit*, Montreal, December 1985, pp. 21-22, and March 1986, pp. 16-17.

- He then has to fit together the various local solutions into a coherent whole at the overall sentence level; this may involve backtracking and discarding existing local solutions to find alternate ones.

Beyond idiomaticity itself, other factors also come into play:

- Detecting ambiguities which may crop up in the translated sentence, and rewording accordingly.
- At paragraph level, rewording for explicitness or clarity in the organization of message rendition.

As to how the translator acquires his idiomaticity skill, my answer will not come as a surprise: ingesting idiomatic sentences, i.e. reading. Hence, the translator who says about a raw output “It just isn’t said this way, but rather this way -- Why? -- I can’t really explain, I just know it” (a familiar situation in translation classes where the teacher tries in vain to comment on idiomaticity mistakes) may have something invaluable to offer to MT research: maybe indeed further advances in MT could be achieved if we could accept “I just know it” for an explanation. Maybe referral to raw accumulated evidence (a corpus) is the only way to solve certain problems,

All MT systems I know about don’t have any convincing approach to solving the problem of contrastive idiomaticity - except in the DLT-BKB project,<sup>4</sup> where the machine translation is pieced together from a collection of existing human translations; this MT design attempts to tap the translator’s know-how as expressed in the end-result of his work. One thing I like about it is that it does not try to explain or rationalize its process; it is based on a decision of the type “I have seen it work before, so it should work here in a related context.” It hints at the idea that translation may largely be an experience-based, organic process (just like the basic use of language), and that the solution may come from an organic source using information in its natural state, with no pre-processing based on human judgement of value. Translation research should try to look much more deeply into this issue.

### 2.2.3 Improving man/machine cooperation

The foregoing suggestions and comments aim at improving the quality of MT output. But this factor, I believe, does not automatically bear a direct relationship with translator productivity; this issue deserves to be examined separately, which has seldom been done so far.

In many practical fields, popular wisdom has developed a well-tested saying that what is supposed to help

<sup>4</sup> See Sadler, Victor, *Working with Analogical Semantics*, Foris Publications, Dordrecht, 1989; and Bedard, Claude, “The BKB (Bilingual Knowledge Bank): a Promising New Approach to Machine Translation,” *Coming of Age — Proceedings of the 30th Annual conference of the ATA*, Learned Information Inc., Medford, NJ, 1989, pp. 437-441.

people and what actually helps them are often two different things. This "human factor principle" certainly applies to machine translation.

MT is put forward as a tool for translator productivity. But where, exactly, are the benefits of MT in the translator's work process? Let's break this question down as follows: First, what does the machine do for the translator? Second, what then remains to be done by the translator? Here is the gist of my answer:

What the machine does for the translator is, in the current state of commercial MT, 1) build sentences which are at best a successful, rather literal rendition, and 2) apply the contents of a dictionary. Not bad, maybe? But then, let's remember that both of these tasks are precisely what a translator would perform rapidly and easily.

What remains to be done, therefore, is the less-than-easy part of the job: understanding the source text (however obscure or badly written), making sure that the translation is accurate, making the translation idiomatic, finding the equivalents not already present in the machine's dictionary, adapting the vocabulary to the situation, conveying the message effectively, etc. If you watch a translator closely, you will realize that this not-so-easy part can be very time-consuming and take up a very large percentage of the overall translation time. For instance I recall typing a fine draft translation of a 25-word sentence in half a minute, then coming back to word 21 (on which much of the sentence's meaning depended) and spending 3 minutes figuring out what it really tried to say -- and how to render it. In the case of this isolated sentence, the easy part would represent 1/7 of the work, and the difficult part 6/7. Of course, over a full text these fractions would stabilize at much less dramatic values.

So what most people don't realize is that the difficult part of the work remains with the translator even though the machine did a decent job by its own (linguistic) standards. In other words, the surface area covered by an MT system does not remove anything from the depth of work required from translator. This tends to define a ceiling to the usefulness of MT as in productivity improvement claims.

On the other hand, there is a price (other than financial) to pay for this "help" from the machine: anything wrong with MT output will intrude upon the translator's natural flow of thinking. The translator is first puzzled by the raw translation, then detects what is wrong, goes to the source text to check what was to be translated in the first place, and comes back and fixes the raw translation. Depending on the raw translation quality, this can be quite a messy process, which forces the translator to shuffle back and forth between three versions of the text: the source, the raw (sometimes called the ugly), and his own final version. In such a situation, the translator is constantly confronted with things which don't make sense, and alerted to considerations which would normally go unnoticed; this puts an unnecessary burden on his mind and is likely to cause mental clutter.

The opposite situation would be the translator dictating, with his eyes constantly on the SL text and his mind on his translated version. All operations are then performed, so to say, "in RAM," at high speed, and to a large extent unconsciously. Also, throughout the work process, things always make sense.

These considerations, along with the high cost of sophisticated MT systems, led me to wonder whether there could be a more appropriate locus for interwork between the machine output and the translator than the current MT paradigm. I first decided that the sentence-building capability of MT systems might be done without, because 1) if it succeeds, it still doesn't save me much time since I could have done the same quickly and easily, and 2) if it fails, the result is likely to make me work harder than by starting from scratch.

What I was interested in was 1) the dictionary-lookup function (what I call the "recall function," especially useful when several translators are requested to work with the same terminology<sup>5</sup>) and 2) the ability to be instantly aware of what is to be translated (the source text).

This led me to develop a new approach which I have called Machine Pre-Translation (MPT).<sup>6</sup> It consists of a batch computer system performing only a partial translation of the text, in line with a number of simple rules:

- Only a (variable) percentage of the words are translated, with a priority on terminological words.
- All words whose translation would obscure the meaning of the source text are left untranslated.
- Multi-word target-language terms are highlighted by linking their component words with underlines.
- Words are not moved around, created, or deleted as a result of pretranslation,

The result looks pretty ugly at first sight. But surprisingly, the translator can easily read it as his source text and start at once turning the pretranslated text into his translation, working on his word-processor (or Dictaphone). Granted, the number of keyboard changes made is greater than with regular MT output; but it seems that the speed and mental comfort in making these changes amply compensate for their number. Since the machine output always makes sense and does not introduce disconcerting elements, the translator seldom has to backtrack; his mind is constantly moving forward -- not back and forth. After three years of experimentation in daily production work, MPT seems to me as a viable alternative to MT, at least in the current state of the technology.

The concept of MPT has the merit, among other things, of challenging the old paradigm of MT output (a

---

<sup>5</sup> Claude Bedard, "Machine PreTranslation : A Further Step in Terminology Management," *The ATA Chronicle*, January 1991, pp. 19-20,

<sup>6</sup> For a full account, see Bedard, Claude, "La prétraduction automatique: un pas en arrière dans la bonne direction?," *Actes du Colloque international "Les industries de la langue -- Perspectives des années 1990,"* Montreal, 1991.

full translation) being THE material translators should work from. Where does this paradigm come from anyway? Here is my personal theory. Originally (in the pre-ALPAC era, when hopes were high), MT systems seem to have been designed for non-translators: linguistic processing (good or bad) had to be complete. After ALPAC and Bar-Hillel's assertion that FAHQT was unattainable, MT system output suddenly shifted to the translator clientele — but the basic paradigm remained the same: the output was 100% target-language. Someone forgot along the way that, although a target-language word (even inadequate) is more informative to a monolingual than a source-language (unknown) word, the opposite is probably true for a translator. This is why I like to call MPT “a step backwards in the right direction!”

This discussion about MPT simply aims at emphasizing how little thought has been devoted to the translator's mental processes, and also calls for a psychological, rather than linguistic, approach to what works best for translators.

### 2.3 Conclusion

This brief attempt at analysis is still superficial and piecemeal, but I believe that it hints at promising areas of investigation into the translator's mind. FAHQT, as we all know, is an awesome undertaking. But finding out exactly why seems an indispensable area of research if we want to eventually know how to solve the problem. The eighties have not brought any breakthrough within the framework defined almost exclusively by computational linguists. Without negating the usefulness of more traditional research, I believe that translators can provide ideas which would produce better results with the available technology, and maybe give MT research a “second wind” by pointing at alternate routes which challenge conventional linguistic approaches to MT. The MT community should open up to such fresh ideas — but on the other hand translators must roll up their sleeves and start to “think harder.” When both communities become together part of the solution, I am sure there will not be much room left for mistrust between them.

### 3 Harald Hille's Statement

Machine translation is, in my view, a natural and extremely interesting and challenging application of linguistics and computational techniques to a very old human activity, namely translation, which has been my life for most of my post-academic career. The past ten years of that career have been marked and enriched by an active interest in computers and applications in the humanities. My remarks on machine translation {MT} might strike some in the field as naive and hostile. In order to mitigate that impression somewhat let me say at the outset that I am very interested in MT, I read whatever I can about it and don't feel at all threatened by it. It is important and enriching for translators to participate in the development of what are ultimately

to become tools in the profession and for computational linguists involved in MT to know what translators think translation is all about. I hope that we will be working even more closely in the future. In his article “Human and Machine Translation” (in *Translation Perspectives II*, National Resource Center for Translation and Interpretation, SUNY-Binghamton, 1985) Wolfram Wilss described four different but complementary “memories” needed by a translator for his work: knowledge of the language system, knowledge of language usage, knowledge of the world and knowledge of the situation. Presumably, linguistic knowledge includes grammatical and lexical competence and the like; knowledge of language usage covers style, register and so forth; world knowledge implies understanding of what things are and can do and how and why things happen, and finally situational knowledge has to do with pragmatic and (psycho)logical understanding of the text or discourse and its context. Human translators have and use these four “memories,” whereas machine translation (MT) systems by and large include only the first, i.e. grammatical and lexical competence. This is a fair statement only if, for instance, one classifies the sort of basic semantic category coding in MT dictionaries (e.g. “city” is marked -animate, +location, +organization, etc.) as lexical and not as some form of world knowledge. Research into the problem of encoding world knowledge -- of the more sophisticated kind -- and programming it into a computer is being pursued at a number of centers, in particular Carnegie Mellon University, but as yet no production systems incorporate such knowledge. It is Wilss's view — and I agree with him — that of the various types of problems arising in translation, syntactic complexity, which draws mainly on the first type of knowledge, is not an insoluble problem for MT systems, but areas where major difficulties continue to exist (and, it would seem, will continue to exist for the near future) are those of semantic ambiguity and the indeterminacy of world knowledge. In a short example:

The stolen car was found by the police.

The stolen car was found by the building site,

Wilss points out the need for semantic rules (valence grammar) at an early stage in the analysis. All serious MT system dictionaries include features marking noun entries, the arguments of verbs and prepositions, etc. that would lead to the correct analysis of this ambiguity of *by* in passive sentences.

Situational knowledge, however, which Wilss doesn't elaborate much on, is equally important. For instance, in the case of:

The missing check was found by the welfare office.

Lexical information and even world knowledge won't solve the problem. The sentence is inherently ambiguous and its disambiguation requires understanding of the situation as it is being described by the text (i.e. does the “welfare office” function in this part of the text as an

agent — an organization that can act — or simply as an object, in this case a place). Here the human translator has the capacity (as yet unique to humans) of building up an understanding of the situation, of understanding the sentence not just as a linguistic object but as part of a communication. That understanding of pragmatic aspects of the text, such as relevance and intentionality, take the form of expectations that the human reader constantly builds and revises and, more importantly, uses to interpret the text as it unfolds.

The human translator can also revise the translation of earlier portions of the text in the light of new information. It is probably safe to say that MT systems have very few expectations about texts and, as far as I know, no MT system will go back and revise translations of earlier sentences. A similar problem arises in the case of anaphora within and across sentences of the text. As far as I know, none of the available MT systems can handle pronoun reference across sentence boundaries. In translating from English to French, for instance, it is crucial to be able to identify the referent of an English “it,” as all nominals in French have grammatical gender, i.e. they are either masculine or feminine, unlike English where that choice is relevant only for pronouns referring to humans and some animals and everything else is referred to as “it.” To paraphrase some remarks on translation by Roman Jakobson, languages differ not so much by what they can express but rather by what they must express. The information to solve the problem may well not be in the sentence at hand; previous sentences (situational knowledge) and world knowledge provide the clues. Translating in the reverse direction can also be tricky as French *il* has to be analyzed to see whether “he” or “it” is appropriate. Even worse, a French *elle* referring to “la personne” in the generic sense would probably want to be translated as “he” in English. Here again the human translator, who can remember what has been said up to the present and who understands the dynamics of the situation, can find the clues to resolve the reference problem, whereas MT systems, which are typically sentence-bound, have very limited if any capacity to forage around outside the sentence in the text for such information.

A similar problem arises when one language makes distinctions in its lexicon that others don’t make and there is no general word that allows one to fudge. An example of this is Spanish *pared* and *muro* for English “wall,” where the Spanish distinguishes between interior and exterior walls. German and Italian make similar distinctions. In translating from English into these languages one has to make a choice and usually the sentence alone will not provide sufficient information. Again, one has to explore the context. This sort of problem is encountered in all parts of the lexicon and grammar (look at verbal aspect in Russian) and can pose real difficulties, I would think, for interlingual MT systems, which try to capture all information in a language-independent abstract representation before going on to synthesize the target text. If your analysis has to anticipate all possible distinctions

in a wide range of target languages the overhead would seem to be tremendous and much of it unnecessary for a given language pair.

These are some areas where MT systems would seem to need more work. Of course, none of this is new to MT developers, although the literature and demonstrations still seem to focus on individual sentences, while human translators find themselves translating texts. While MT developers work on solutions to such problems, it would be useful for translators with some understanding of MT and natural language processing to analyze texts to determine the frequency of such problems in texts of various kinds and to try to evaluate the nuisance factor of such problems in post-editing. One could then set up a set of priorities for further development. A problem may be found to be relatively infrequent in actual texts and, unless it is judged to be particularly annoying to correct in post-editing, could be assigned low priority.

#### 4 Fred Klein’s Statement

I am very grateful for this opportunity of speaking to you today. I have enjoyed close to 30 years of translation experience on three continents. I haven’t had a formal higher education. This was not through my choice but due to the practical realities of the European cultural and political world at the time I was growing up.

Still, I have been highly successful in my chosen profession and can proudly boast of my many publications, prizes, and conference panel participations.

That does not mean that college or higher degrees are unnecessary, of course, but it does mean that it is possible to do without them. And I have been able to work up to seven times faster and much more efficiently based on my own resources than with the support of two different well-known MT systems. I am not implying that MT is useless; on the contrary, it is in my opinion simply limited to certain fields.

I am very happy to be here for this conference, and I hope to have an opportunity to talk with many of you individually while I am here. Although I enjoy a certain reputation as an outspoken person, I will do my best in any speaking or writing I do during this conference to behave myself. I promise not to criticize certain individual MT systems I have used, nor even to mention their names. I will, however, talk about advantages and disadvantages of systems in general and try to indicate solutions where these are applicable.

Let me first say something in favor of MT for information purposes or raw output, that it is an excellent technology, which cannot be matched by humans. It represents limited but useful output, although it may not be entirely legitimate to call this translation.

It is my belief that MT developers have failed to study what translating entails. They have started from wrong assumptions back in the ’Fifties, without understanding that word-for-word translation is useless.

Many professionals in the field of computational linguistics seem to view the computer with an almost reli-

gious fervor. For myself, I consider the computer to be the best and most indispensable tool I have ever owned. But a tool, no more than this. I certainly believe that MT is possible, but not necessarily in a cost-efficient manner. I would of course love to see a variety of different and more promising results.

The key issue remains: is MT truly viable at all? Yes, it can be cost-efficient and viable, but only under specific, limited, relatively rare conditions, as with the TAUM weather forecasts in Canada; or with the SPANAM environment, which seems to be the only logical result of the original Georgetown project, carried out under almost ideal (and thus unrealistic) circumstances; or in the case of abstracts, or under other such rigidly controlled linguistic conditions, for example, the Controlled English system used by Xerox.

The growth of MT has been slow even in the giant market of the European Community and the ready-made bilingual culture of Canada. Translation users have no inherent bias against MT. If it were workable, they would use it. As a result of slow progress in MT, a gulf has developed, and initial failures of MT have left translators, along with many others, skeptical about any real possibilities.

We do not know how we think. We do not know how we translate. Therefore the best hard- and software, the fastest system, the vastest memory, the most perfect programs have failed year after year.

It should be recognized that there may be limits to the organizing of human knowledge. FAHQT is an impossible dream and a complete waste of time and resources. Universal, context-independent, concept-neutral “knowledge domains” are absurd, unrealistic assumptions of people who have no knowledge of the real world of language or of translation. If we do not establish reasonable priorities and eliminate costly dreams, then we will all fail again. It should be added that non-native speakers of a language are unable to judge the quality of translation, even if they are famous computational linguists, gurus or sheer geniuses in other undertakings. You can postulate MT, but the global market will buy or reject the results. In a sense the global market is the true end user.

For MT to progress, computational linguists must in my opinion come to accept translators as indispensable partners in their work. Translators often work on an almost unconscious level: is it at all possible to program for this? But the fact of the matter has been that translators have been left out from MT for decades.

The knowledge of practicing translators can be best integrated into MT through the use of expert systems, which can do everything but translate or compile glossaries, since these glossaries are the cornerstone of MT. Publicity and public relations directed towards translators and of course their actual involvement is mandatory if MT is to succeed.

Here are some advance suggestions for promoting MT:

1. Publish a multilingual glossary of key MT terms

with simple definitions in at least five languages (include some from Prescott’s new survey for the EC definitions by MT developers, etc.). Try to use language which the average translator can understand. Do not aim too high.

2. Make such a glossary available in electronic form at the Summit. Provide on-line versions as well, if possible.
3. Consider an international MT newsletter. With ads by developers, but independent articles and papers, Down-to-earth. User-friendly. Translator-friendly. Exclude high-sounding theories and research. Researchers and developers must stop simply talking among themselves. They must talk in a non-condescending manner to the real users, who are very frequently translators. Ideally such a newsletter should be multilingual. And of course electronic.
4. Take some form of effective international action to limit false and misleading claims in this field. These include not only those overtly false claims that black boxes will automatically and flawlessly translate all possible materials to and from all languages, but also a failure to contradict a tendency among journalists and the general public to believe that such myths can be true. If MT Summit III has not yet set up a *hersteller-neutrale* (manufacturer-independent) work group to eliminate such claims, then I urge you to do so immediately. Such misleading publicity is not only false but represents a disservice to the real living translators, who—as you are perhaps just beginning to realize—may be absolutely crucial to the work you propose to do.
5. Provide up-to-date records of viable MT systems on a global scale. Do any such records exist, and if not, why not? For our purposes, forget the experimental ones completely.

Finally, a closing remark. I sincerely hope that you not be offended in any way by this notion. It seems to me, from the point of view of the future of machine translation, that this conference has perhaps been misnamed. Is a “summit” really the best means of communicating the potential this field may truly possess? Might it not be more appropriate to descend from this summit and begin to discuss what MT is and can become in terms closer to ordinary human utterance rather than the ones in which this subject is all too often couched? It is my great honor to be here with you, and I thank you for your patience in listening.

## 5 Sergei Nirenburg’s Statement

Human translators today almost universally produce better quality results than machine translation systems (which are currently valued mostly for their potential for translating large volumes of text cheaply). It therefore seems natural to strive to make MT systems behave more like human translators. This can be done either through the judicious use of texts already translated by



humans in the process of MT or by somehow encoding human expertise in an automatic translation system. Thus, Martin Kay's oft-quoted opinion is that MT developers do not pay sufficient attention to the work of translators. But the criticism may have been directed not at the absence of translators in the R&D loop but rather at the lack of use of past human translations in MT environments. This state of affairs is actually being remedied by extensive use in some MT systems of collocational information gathered in monolingual and especially bilingual corpora and, in MAT environments, by facilitating retrieval of past translations as templates for new translations.

Translators are the obvious source of insight into the mechanics of the translation process. But in reality practicing translators are seldom invited to take an active part in research and development of MT technology. This state of affairs is considered untenable, and it is a standard item of criticism against the MT research and development community. It may be interesting to analyze the reasons for this alleged aloofness.

For the sake of this discussion, let us assume that we are talking of MT R&D projects whose genuine aim is to develop a useful MT system rather than test a linguistic theory or a computational method. Let us also assume that there is a consensus of opinion among the translators with respect to the knowledge that they can impart to an MT system. The above assumptions are certainly nontrivial, as too many MT systems are method- rather than task-oriented, on the one hand, and too many human translations are judged inadequate by other translators, be it for real reasons or interpersonal differences in taste.

The crucial point is as follows; the lack of the translators' involvement in MT is criticized but no serious constructive suggestions are made as to how this should be remedied! The underlying reason for this absence is as follows. In order for direct human input to be felt in MT system design, the systems must recreate the decision processes of humans. Since we don't know how we think and how we make decisions (psychologists have only *theories* about that), it is difficult to formulate a set of appropriate rules for a system to follow. Moreover, since computers and humans are very different entities, it is not clear whether imitating humans is indeed the best way to go. (It has been amply shown that things that people find difficult, like high-volume number crunching, are easy for machines, whereas things that are easy for people, like speaking or text comprehension, are extremely difficult for the computer.) This point has been discussed at length in connection with the foundations of the discipline of artificial intelligence, in response to the criticism that using computers (at least, today's computers) to perform tasks that previously could be performed only by humans does not necessarily mean recreating human intelligence; it, in fact, means only recreating the results of human intelligence operation, by whatever means necessary. A very good example is the success of the computer chess program Deep Thought, based essen-

tially on fast calculations, in games against top human competition.

The only way, then, to try to use direct translator expertise in MT R&D is to treat the task as a knowledge acquisition task in the field of expert systems, where typically a specially trained developer conducts extensive interviews with an expert (in our case, this would be a translator) recording the latter's solutions to test problems. The next step, then, is to reformulate the translator's opinions in such a way that a computer program could use this expert knowledge in a practical MT system. This methodology has been used, with some success in a number of fields, usually, to help solve some kind of a diagnosis problem (in medicine, equipment maintenance or mineral prospecting, for instance). In translation, we deal with considerably greater complexity, since language understanding is centrally involved, in addition to any knowledge of a particular field or, orthogonally, of professional translator expertise. As a person who has been trained as both a translator and a computational linguist and who has been, at different times, a practicing translator and a practicing MT researcher, I can attest that a) I am not confident in my own perceptions about my thought processes during translation and b) it is difficult to suggest a detailed scheme for "debriefing" translators on the subject of their professional knowledge. The above does not imply that the expert system way is a wrong one, merely that the complexity of this approach seems to be similar to building systems which do not attempt a direct representation of human knowledge about translation.

In my opinion, there are two types of automation-related activities in which a translator can take part today. For immediate use, translators and developers have to join forces in developing advanced MAT tools. In a longer-term perspective, a two-fold development must take place. First, translators must realize that advanced research and development on MT must continue, as the expected benefits for the translation profession are quite high. Even if FAHQMT is a dream, the results of asymptotic convergence toward that goal can be sufficiently dramatic. Second, MT developers must realize that the need for quality and coverage breakthroughs in MT warrant a diligent study of translation techniques and day-to-day contacts with translators for the purposes at least of expertise acquisition and evaluation.

## 6 Summary, Alex Gross

The primary question for this panel is of course embodied in its title, "Where Do Translators Fit Into MT?" This question was then broken down into four sub-questions, which may be briefly summarized as follows:

1. Is it, true that translators have been left out of the MT development process, and if so, why?
2. Do human beings bring something to the translation process which machines cannot duplicate?

3. How could human translators (or their knowledge) best be integrated into MT processes? and
4. How are translators likely to view renewed efforts towards some form of FAHQT, and how might this in turn influence their feelings about MT?

Although these questions cover a broad scope, the range of answers provided by the panelists in their written remarks is broader still. Bedard holds that “MT researchers are widely known not to be too translation-conscious,” and that human translators should be much more active in this field because they are both “translation-conscious” and also rather pragmatic people. He further refers to a process used by human translators, even when the domain of a text may be unfamiliar to them, which he calls following “low-level contextual clues” or “translation by radar.” He also maintains that each translation situation is specific and that terminology, assumed by some MT developers to be invariable boiler-plate, is often situation-dependent. He is deeply concerned that “different languages ... don’t say things the same way” and feels that this poses a problem for most MT approaches. He furthermore observes that MT output often compels the translator “to stop for considering items normally disregarded in his normal thought process .. forcing the translator to work in zigzag” between not two but three texts, and he wonders if this “messy process” truly saves him any time or effort. To solve these various problems, Bedard proposes a list of specific steps for building up an improved human-machine interface, which he calls “machine pre-translation.”

Harald Hille bases his text upon a published paper by Wolfram Wilss, who discussed four different but complementary “memories” translators require for their work: “knowledge of the language system, knowledge of language usage, knowledge of the world, and knowledge of the situation.” He points out that generally speaking the computer, despite various pending projects, can so far be truly said to “possess” only the first of these. He also refers to the ability of humans to disambiguate inherently ambiguous sentences, resolve anaphora, and bring context or real-world knowledge to bear upon a problem. Cases where different languages make divergent distinctions about words and concepts comprise yet another class of instances where, Hille feels, “MT systems would seem to need more work.”

Far more outspoken in his criticisms of MT is Fred Klein, who feels that many computational linguists “seem to view the computer with an almost religious fervor.” Klein insists that MT is viable “only under specific, limited, relatively rare conditions,” a factor which, he believes, may explain why its “growth ... has been slow.” He sees FAHQT as “an impossible dream” and “universal ...concept-neutral ‘knowledge domains’” as “absurd.” For MT to progress, says Klein, “computational linguists must ...come to accept translators as indispensable partners in their work.” His paper concludes with a multi-point list of suggestions for promot-

ing MT through publications, enhanced publicity, and “some form of effective international action to limit false and misleading claims in this field.”

Sergei Nirenburg begins his paper with a statement that human translation output today is “almost universally ... better” than MT translation but thereafter follows a distinctly different line from the other panelists. The point of MT, he argues, is not to have it work exactly like humans or recreate intelligence but for it to evolve its own methods of producing superior output. He furthermore states that the main thrust of Martin Kay’s comment, cited above, may have not been to involve translators more closely in MT processes but rather to integrate past HT work into MT systems so as to improve overall performance. He concludes that “translators must realize that ... research and development on MT must continue” but also that MT developers should recognize that “a diligent study of translation techniques and day-to-day contacts with translators” is warranted.