

Why Can't a Computer Translate More Like a Person? **by Alan K. Melby**

Posters and Telescopes: an Introduction to Translation

Translation is difficult, even for people. To begin with, you have to know two languages intimately. And even if you speak two or more languages fluently, it is not a trivial matter to produce a good translation. When people start talking about the possibility of a computer replacing a human translator, someone will often bring up a sentence similar to the following:

** Time flies like an arrow.*

The person who brought it up will typically conclude by asserting that this sentence is an obvious example of a sentence that a computer could not translate. As a matter of fact, a computer could handle this sentence if it were programmed to handle just this sentence. The problem is getting a computer to deal adequately with sentences it has not been specifically programmed to handle. I will partially analyze this sentence and then give other superficially similar sentences that cannot all be translated in a parallel fashion.

This sample sentence about time flying is a figure of speech that combines a metaphor and a simile. Time does not really fly in the literal sense of a bird flying, but here we metaphorically say it does. Another example of a metaphor would be when we say that a ship ploughs the water, since there is no real plough and no dirt involved. The simile in this expression is the comparison between the metaphorical flight of time with the flight path of an arrow.

Now consider the following sentence, which is a rather dumb-sounding figure of speech modeled on the first one:

** Fruit flies like a baseball.*

Not all fruit, when thrown, would fly through the air like a baseball, except perhaps an apple, orange, or peach. But wait a minute. Suppose you substitute 'peach' for 'baseball' in the second sentence. All of a sudden, there is a new meaning. This time everything is literal. The 'fruit flies' are pesky little insects you can see crawling around on a juicy peach, having a feast. The 'peach' version of the sentence would be translated very differently from the 'arrow' version or the 'baseball' version.

The point of these sentences for human versus computer translation is that a human translator would know to handle the variation "Fruit flies like a peach" very differently from the baseball version while a computer would probably not even notice the difference and therefore could never replace a human translator. Why wouldn't a computer notice the difference? We will explore differences between humans and computer throughout this paper.

These sentences do show how words can shift in their usage. The word 'flies' shifts from signifying an action to signifying an insect, and in most languages it cannot be translated the same way in both usages. But we do not need anything nearly so exotic as these sentences in order to show that translation is full of pitfalls. Let me give you an example of a human translation of a simple poster, a translation that did not turn out very well.

This summer I attended a conference in Luxembourg and noticed in the train station a poster announcing a coming event. The announcement was in French with an English translation. I will refer to this announcement as the poster example. The English translation of the date and time of the event read as follows:

* *Saturday the 24 June 1995 to 17 o'clock*

Obviously, there are a number of problems in this translation. In English we say “the 24th of June, 1995” or “June 24, 1995,” rather than “the 24 June 1995.” Also, we say “5 o'clock” or “5 p.m.,” because in the United States we divide the day into two 12-hour periods, rather than one 24-hour period, except in the military. In England, the use of a 24-hour clock is more common but even there one would not say “17 o'clock.” Perhaps the most puzzling error in this translation is the use of the word ‘to’. At first glance, one would assume that the word ‘at’ was intended, so that the translation becomes, after all our changes:

* *Saturday, the 24th of June at 5 p.m.*

However, an examination of the French shows that this is incorrect. The French original used the word *vers*, which can mean either ‘in the direction of’ (as in a movement toward an object or to the left) or ‘at an approximate time’ (as in a promise to drop off a package around noon). Clearly, the second reading of *vers* is more likely here. Whoever translated the French probably used a French to English dictionary and just picked the first translation listed under the word *vers*, without thinking about whether it would work in this context. In the case of this poster, the translator did not have a sufficient knowledge of both languages, and the translation turned out not only awkward but just plain wrong.

This example of bad human translation is interesting because it was most likely done by a human yet in a manner similar to the way computers translate. (By the way, the conference I was attending in Luxembourg, where I saw the poster, was the fifth world summit on computer translation, which is usually called Machine Translation, hence the conference title: Machine Translation Summit V.)

Computers do not really think about what they are doing. They just mechanically pick a translation for each word of the source text, that is, the text being translated, without understanding what they are translating and without considering the context. An examination of the source text for our poster example will illustrate this.

* French source text: *le samedi 24 juin 1995 vers 17h00*

* Poster translation: *Saturday the 24 June 1995 to 17 o'clock*

* Better translation: *Saturday, the 24th of June, 1995, around 5 p.m.*

To give credit where it is due, the translator apparently knew enough about English dates to reposition the translation of the French article *le* to the other side of ‘Saturday.’ Other than the re-ordering of the article, the translation on the poster could be obtained using a simple word-for-word substitution technique by either a person or a computer looking up words in a dictionary. No real knowledge of either language would be required. Thus, people can easily translate like computers, that is, mechanically, usually with rather disappointing results. However, the opposite is not true. Computers cannot, in general, translate like people, at least

not like people who know both languages and are skilled translators. I have analyzed a poor quality human translation and provided an improved human translation. We will now look at a real-life example of machine translation. I will refer to it as the telescope example.

Last year I was at another conference on machine translation, this one being held at Cranfield University in the England. There were several major companies in the exhibit area demonstrating their commercial machine translation systems. On the way, I had picked up a French magazine similar to the American magazine *Air and Space*, and at the conference I fed a sentence from the magazine into one of the machine translation systems. Below is the French sentence that went in, followed by the English translation that came out of the computer.

* French source sentence: *L'atmosphère de la Terre rend un peu myopes mêmes les meilleurs de leur télescopes.*

* English machine translation: *The atmosphere of the Earth returns a little myopes same the best ones of their telescopes.*

Even without knowing French, one can see that the English translation is basically the result of a word-for-word substitution. In the poster example, the translation was awkward and somewhat misleading. This translation is perhaps even worse: it is practically incomprehensible. The context of the source text is an article from a French magazine discussing the problem of turbulence in the atmosphere. The magazine is addressed to a general audience rather than to professional astronomers. One possible human translation would be the following:

* *The earth's atmosphere makes even the best of their telescopes a little "near sighted"* (in the sense that distant objects are slightly blurred).

There are obviously a number of problems in the machine translation. These problems stem from the ambiguity of word meanings. For example, the French verb *rend* can be translated as 'return' or 'make,' depending on the context. The French word *même* can be translated as 'same' or 'even,' again depending on the context. In both cases, the computer mechanically chose a translation and in both cases the poor thing got it dead wrong. It is hard to tell whether the computer couldn't find the word *myopes* in its dictionary and just passed it through unchanged or whether it found it and translated it inappropriately for the audience. A 'myope' is a technical term in English for someone who is myopic, that is, near-sighted. However, this is the wrong level of language to use in a publication intended for a general audience. Computers have no sense of audience; they just blindly follow rules. Another machine translation system, when given the same French sentence, did better in some ways but made other mistakes:

* *The atmosphere of the Earth renders a same myopic bit best of their telescopes.*

Professional human translators seldom, if ever, make errors like the ones we have seen in the poster example and the telescope example. Nevertheless, humans with nothing but a dictionary in hand can choose to stoop to the level of computers. In contrast, computers have not risen to the level of professional human translators. Why not? Why can't a computer translate more like a person?

It is interesting to observe how various persons who have not worked on machine translation react to the title question of this paper. Some believe that there is no fundamental difference between humans and machines. They assume that the quality of machine translation will

someday rival the quality of human translation in all respects. They point out that computers can do arithmetic much faster and more accurately than people. Then they remind us that math is harder than language for many students. Furthermore, they take it as obvious that the human brain is ultimately a type of computer. From this basis, they conclude that it is just a matter of time until we have a new kind of computer that will function like the brain, only faster and better, and will surpass the capabilities of humans in the area of language processing. Others take a contrary position. They believe that humans and computers are so entirely different in the way they work inside that computers will never approach the capabilities of human translators. Still others are puzzled by the question. They were under the impression that the problem of machine translation was solved years ago.

The fact of the matter is that machine translation is a problem that is far from solved. Experts in the field agree that computers do not yet translate like people. On some texts, particularly highly technical texts treating a very narrow topic in a rather dry and monotonous style, computers sometimes do quite well. (In the annex to this paper, I give a sentence of English and its computer-generated translation that was offered by a vendor as part of a showcase example of machine translation.) But with other texts, particularly with texts that are more general and more interesting to humans, computers are very likely to produce atrocious results. Professional human translators, on the other hand, can produce good translations of many kinds of text. People can handle a range of text types; computers cannot. Where the experts disagree is on the question of why computers are so limited in their ability to translate. I will present an answer to this controversial question, but only at the end of this paper. I will build up to it in the following stages:

- First, I will present a few more examples of why translation is difficult for both people and computers, even for people who know two languages and for computers that are carefully programmed to translate.
- Secondly, I will very briefly describe the mainstream approach to characterizing human language and point out how it fails to address the difficulties presented in the first stage.
- Thirdly, I will discuss a key factor that is missing in current theories of human language, a factor that I believe will be needed in computers for them to be able to translate more like people.

Some Difficulties in Translation

One difficulty in translation stems from the fact that most words have multiple meanings. Because of this fact, a translation based on a one-to-one substitution of words is seldom acceptable. We have already seen this in the poster example and the telescope example. Whether a translation is done by a human or a computer, meaning cannot be ignored. I will give some more examples as evidence of the need to distinguish between possible meanings of a word when translating.

A colleague from Holland recounted the following true experience. He was traveling in France and decided to get a haircut. He was a native speaker of Dutch and knew some French; however, he was stuck when it came to telling the female hairdresser that he wanted a part in his hair. He knew the Dutch word for a part in your hair and he knew one way that Dutch word could be translated into French. He wasn't sure whether that translation would work in this situation, but he tried it anyway. He concluded that the French word did not convey both meanings of the Dutch word when the hairdresser replied, "But, Monsieur, we are not even married!" It seems that the Dutch expression for a separation of your hair (a part) and a permanent separation of a couple (a divorce) are the same word. When you think about it, there is a logical connection, but

we are not conscious of it in English even though you can speak of a parting of your hair or a parting of ways between two people. In French, there is a strong separation of the two concepts. To translate the Dutch word for 'part' or 'divorce' a distinction must be made between these two meanings. We will refer to this incident as the haircut example. Some questions it raises are these: How does a human know when another use of the same word will be translated as a different word? And how would a computer deal with the same problem?

We expect a word with sharply differing meanings to have several different translations, depending on how the word is being used (Oxford English Dictionary, 1989, pp. 930-931). The financial sense evolved from the money changer's table or shelf, which was originally placed on a mound of dirt. Later the same word came to represent the institution that takes care of money for people. The river meaning has remained more closely tied to the original meaning of the word. Even though there is an historical connection between the two meanings of 'bank,' we do not expect their translation into another language to be the same, and it usually will not be the same. This example further demonstrates the need to take account of meaning in translation. A human will easily distinguish between the two uses of 'bank' and simply needs to learn how each meaning is translated. How would a computer make the distinction?

Another word which has evolved considerably over the years is the British word 'treacle,' which now means 'molasses.' It is derived from a word in Ancient Greek that referred to a wild animal. One might ask how in the world it has come to mean molasses. A colleague, Ian Kelly, supplied me with the following history of 'treacle'. The original word for a wild animal came to refer to the bite of a wild animal. Then the meaning broadened out to refer to any injury. It later shifted to refer to the medicine used to treat an injury. Still later, it shifted to refer to a sweet substance mixed with a medicine to make it more palatable. And finally, it narrowed down to one such substance, molasses. At each step along the way, the next shift in meaning was unpredictable, yet in hindsight each shift was motivated by the previous meaning. This illustrates a general principle of language. At any point in time, the next shift in meaning for a word is not entirely unlimited. We can be sure it will not shift in a way that is totally unconnected with its current meaning. But we cannot predict exactly which connection there will be between the current meaning and the next meaning. We cannot even make a list of all the possible connections. We only know there will be a logical connection, at least as analyzed in hindsight.

What are some implications of the "haircut", "bank", and "treacle" examples? To see their importance to translation, we must note that words do not develop along the same paths in all languages. Simply because there is a word in Dutch that means both 'part' and 'divorce' does not mean that there will be one word in French with both meanings. We do not expect the two meanings of 'bank' to have the same translation in another language. We do not assume that there is a word in Modern Greek that means 'molasses' and is derived from the Ancient Greek word for 'wild animal' just because there is such a word in British English. Each language follows its own path in the development of meanings of words. As a result, we end up with a mismatch between languages, and a word in one language can be translated several different ways, depending on the situation. With the extreme examples given so far, a human will easily sense that multiple translations are probably involved, even if a computer would have difficulty. What causes trouble in translation for humans is that even subtle differences in meaning may result in different translations. I will give a few examples.

The English word 'fish' can be used to refer to either a live fish swimming in a river (the one that got away), or a dead fish that has been cleaned and is ready for the frying pan. In a sense, English makes a similar distinction between fish and seafood, but 'fish' can be used in both cases. Spanish makes the distinction obligatory. For the swimming fish, one would use pez and

for the fish ready for the frying pan one would use pescado. It is not clear how a speaker of English is supposed to know to look for two translations of 'fish' into Spanish. The result is that an unknowledgable human may use the wrong translation until corrected.

The English expression 'thank you' is problematical going into Japanese. There are several translations that are not interchangeable and depend on factors such as whether the person being thanked was obligated to perform the service and how much effort was involved. In English, we make various distinctions, such as 'thanks a million' and 'what a friend,' but these distinctions are not stylized as in Japanese nor do they necessarily have the same boundaries. A human can learn these distinctions through substantial effort. It is not clear how to tell a computer how to make them.

Languages are certainly influenced by the culture they are part of. The variety of thanking words in Japanese is a reflection of the stylized intricacy of the politeness in their culture as observed by Westerners. The French make an unexpected distinction in the translation of the English word 'nudist.' Some time ago, I had a discussion with a colleague over its translation into French. We were reviewing a bilingual French and English dictionary for its coverage of American English versus British English, and this word was one of many that spawned discussion. My colleague, who had lived in France a number of years ago, thought the French word nudiste would be the best translation. I had also lived in France on several occasions, somewhat more recently than him, and had only heard the French word naturiste used to refer to nude beaches and such. Recently, I saw an article in a French news magazine that resolved the issue. The article described the conflict between the nudistes and the naturistes in France. There was even a section in the article that explained how to tell them apart. A nudiste places a high value on a good suntan, good wine, and high-fashion clothes when away from the nudist camp. A naturiste neither smokes nor drinks and often does yoga or transcendental meditation, prefers homeopathic medicine, supports environmental groups, wears simple rather than name-brand clothing when in public, and tends to look down on a nudiste. There is currently a fight in France over which nude beaches are designated naturiste and which are designated nudiste. Leave it to the French, bless their souls, to elevate immodesty to a nearly religious status. I trust my French colleagues will not take offense.

The verb 'to run' is another example of a word that causes a lot of trouble for translation. In a given language, the translation of 'run' as the next step up in speed from jogging will not necessarily be the same word as that used to translate the expression 'run a company' or 'run long' (when referring to a play or meeting) or 'run dry' (when referring to a river). A computer or an inexperienced human translator will often be insensitive to subtle differences in meaning that affect translation and will use a word inappropriately. Significantly, there is no set list of possible ways to use 'run' or other words of general vocabulary. Once you think you have a complete list, a new use will come up. In order to translate well, you must first be able to recognize a new use (a pretty tricky task for a computer) and then be able to come up with an acceptable translation that is not on the list.

The point of this discussion of various ways to translate 'fish,' 'thank you,' 'nudist,' and 'run' is that it is not enough to have a passing acquaintance with another language in order to produce good translations. You must have a thorough knowledge of both languages and an ability to deal with differences in meaning that appear insignificant until you cross over to the other language.¹ Indeed, you must be a native or near-native speaker of the language you are translating into and very strong in the language you are translating from. Being a native or near-native speaker involves more than just memorizing lots of facts about words. It includes having an understanding of the culture that is mixed with the language. It also includes an ability to deal with new situations appropriately. No dictionary can contain all the solutions since the problem

is always changing as people use words in usual ways. These usual uses of words happen all the time. Some only last for the life of a conversation or an editorial. Others catch on and become part of the language. Some native speakers develop a tremendous skill in dealing with the subtleties of translation. However, no computer is a native speaker of a human language. All computers start out with their own language and are 'taught' human language later on. They never truly know it the way a human native speaker knows a language with its many levels and intricacies. Does this mean that if we taught a computer a human language starting the instant it came off the assembly line, it could learn it perfectly? I don't think so. Computers do not learn in the same way we do. We could say that computers can't translate like humans because they do not learn like humans. Then we still have to explain why computers don't learn like humans. What is missing in a computer that is present in a human? Building on the examples given so far, I will describe three types of difficulty in translation that are intended to provide some further insight into what capabilities a computer would need in order to deal with human language the way humans do, but first I will make a distinction between two kinds of language.

Certainly, in order to produce an acceptable translation, you must find acceptable words in the other language. Here we will make a very important distinction between two kinds of language: general language and specialized terminology. In general language, it is undesirable to repeat the same word over and over unnecessarily. Variety is highly valued. However, in specialized terminology, consistency (which would be called monotony in the case of general language) is highly valued. Indeed, it is essential to repeat the same term over and over whenever it refers to the same object. It is frustrating and potentially dangerous to switch terms for the same object when describing how to maintain or repair a complex machine such as a commercial airplane. Now, returning to the question of acceptable translation, I said that to produce an acceptable translation, you must find acceptable words. In the case of specialized terminology, it should be the one and only term in the other language that has been designated as the term in a particular language for a particular object throughout a particular document or set of documents. In the case of general vocabulary, there may be many potential translations for a given word, and often more than one (but not all) of the potential translations will be acceptable on a given occasion in a given source text. What determines whether a given translation is one of the acceptable ones?

Now I return to the promised types of translation difficulty. The first type of translation difficulty is the most easily resolved. It is the case where a word can be either a word of general vocabulary or a specialized term. Consider the word 'bus.' When this word is used as an item of general vocabulary, it is understood by all native speakers of English to refer to a roadway vehicle for transporting groups of people. However, it can also be used as an item of specialized terminology. Specialized terminology is divided into areas of knowledge called domains. In the domain of computers, the term 'bus' refers to a component of a computer that has several slots into which cards can be placed. One card may control a CD-ROM drive. Another may contain a fax/modem. If you turn off the power to your desktop computer and open it up, you can probably see the 'bus' for yourself.

As always, there is a connection between the new meaning and the old. The new meaning involves carrying cards while the old one involves carrying people. In this case, the new meaning has not superseded the old one. They both continue to be used, but it would be dangerous, as we have already shown with several examples, to assume that both meanings will be translated the same way in another language. The way to overcome this difficulty, either for a human or for a computer, is to recognize whether we are using the word as an item of general vocabulary or as a specialized term.

Humans have an amazing ability to distinguish between general and specialized uses of a word. Once it has been detected that a word is being used as a specialized term in a particular domain,

then it is often merely a matter of consulting a terminology database for that domain to find the standard translation of that term in that domain. Actually, it is not always as easy as I have described it. In fact, it is common for a translator to spend a third of the time needed to produce a translation on the task of finding translations for terms that do not yet appear in the terminology database being used. Where computers shine is in retrieving information about terms. They have a much better memory than humans. But computers are very bad at deciding which is the best translation to store in the database. This failing of computers confirms our claim that they are not native speakers of any human language in that they are unable to deal appropriately with new situations.

When the source text is restricted to one particular domain, such as computers, it has been quite effective to program a machine translation system to consult first a terminology database corresponding to the domain of the source text and only consult a general dictionary for words that are not used in that domain. Of course, this approach does have pitfalls. Suppose a text describes a very sophisticated public transportation vehicle that includes as standard equipment a computer. A text that describes the use of this computer may contain the word 'bus' used sometimes as general vocabulary and sometimes as a specialized term. A human translator would normally have no trouble keeping the two uses of 'bus' straight, but a typical machine translation system would be hopelessly confused. Recently, this type of difficulty was illustrated by an actual machine translation of a letter. The letter began "Dear Bill" and the machine, which was tuned into the domain of business terms, came up with the German translation *Liebe Rechnung*, which means something like "Beloved Invoice."

This first type of difficulty is the task of distinguishing between a use of a word as a specialized term and its use as a word of general vocabulary. One might think that if that distinction can be made, we are home free and the computer can produce an acceptable translation. Not so. The second type of difficulty is distinguishing between various uses of a word of general vocabulary. We have already seen with several examples ('fish', 'run,' etc.) that it is essential to distinguish between various general uses of a word in order to choose an appropriate translation. What we have not discussed is how that distinction is made by a human and how it could be made by a computer.

Already in 1960, an early machine translation researcher named Bar-Hillel provided a now classic example of the difficulty of machine translation. He gave the seemingly simple sentence "The box is in the pen." He pointed out that to decide whether the sentence is talking about a writing instrument pen or a child's play pen, it would be necessary for a computer to know about the relative sizes of objects in the real world . Of course, this two-way choice, as difficult as it is for a human, is a simplification of the problem, since 'pen' can have other meanings, such as a short form for 'penitentiary' or another name for a female swan. But restricting ourselves to just the writing instrument and play pen meanings, only an unusual size of box or writing instrument would allow an interpretation of 'pen' as other than an enclosure where a child plays. The related sentence, "the pen is in the box," is more ambiguous . Here one would assume that the pen is a writing instrument, unless the context is about unpacking a new play pen or packing up all the furniture in a room. The point is that accurate translation requires an understanding of the text, which includes an understanding of the situation and an enormous variety of facts about the world in which we live. For example, even if one can determine that, in a given situation, 'pen' is used as a writing instrument, the translation into Spanish varies depending on the Spanish-speaking country.

The third type of difficulty is the need to be sensitive to total context, including the intended audience of the translation. Meaning is not some abstract object that is independent of people and culture. We have already seen in examples such as the translation of 'thank you' in Japanese

a connection between culture and distinctions made in vocabulary. Several years ago, I translated a book on grammar from French to English. It was unfortunately not well received by English-speaking linguists. There were several reasons, but one factor was the general rhetorical style used by French-speaking linguists: they consider it an insult to the reader to reveal the main point of their argument too soon. From the point of view of an English-speaking linguist, the French linguist has forgotten to begin with a thesis statement and then back it up. Being sensitive to the audience also means using a level of language that is appropriate. Sometimes a misreading of the audience merely results in innocuous boredom. However, it can also have serious long-term effects.

A serious example of insensitivity to the total context and the audience is the translation of a remark made by Nikita Khrushchev in Moscow on November 19, 1956. Khrushchev was then the head of the Soviet Union and had just given a speech on the Suez Canal crisis. Nassar of Egypt threatened to deny passage through the canal. The United States and France moved to occupy the canal. Khrushchev complained loudly about the West. Then, after the speech, Khrushchev made an off-hand remark to a diplomat in the back room. That remark was translated “We will bury you” and was burned into the minds of my generation as a warning that the Russians would invade the United States and kill us all if they thought they had a chance of winning. Several months ago, I became curious to find out what Russian words were spoken by Khrushchev and whether they were translated appropriately. Actually, at the time I began my research I had the impression that the statement was made by Khrushchev at the United Nations at the same time he took off his shoe and pounded it on the table. After considerable effort by several people, most notably my daughter Yvette, along with the help of Grant Harris of the Library of Congress, Professor Sebastian Shaumyan, a Russian linguist, Professor Bill Sullivan of the University of Florida, Professor Don Jarvis of Brigham Young University, and Professor Sophia Lubensky of the State University of New York at Albany, I have been able to piece together more about what was actually said and intended.

The remark was not ever reported by the official Russian Press. Rather it was reported by a Russian-language newspaper called *Novoe Russkoe Slovo*, run by Russian emigres in the United States. It reported that along with the famous remark, Khrushchev said flippantly that “If we believed in God, He would be on our side.” In Soviet Communist rhetoric, it is common to claim that history is on the side of Communism, referring back to Marx who argued that Communism was historically inevitable. Khrushchev then added that Communism does not need to go to war to destroy Capitalism. Continuing with the thought that Communism is a superior system and that Capitalism will self-destruct, he said, rather than what was reported by the press, something along the lines of “Whether you like it or not, we will be present at your burial,” clearly meaning that he was predicting that Communism would outlast Capitalism. Although the words used by Khrushchev could be literally translated as “We will bury you,” (and, unfortunately, were translated that way) we have already seen that the context must be taken into consideration. The English translator who did not take into account the context of the remark, but instead assumed that the Russian word for “bury” could only be translated one way, unnecessarily raised tensions between the United States and the Soviet Union and perhaps needlessly prolonged the Cold War.

We have identified three types of translation difficulty: (1) distinguishing between general vocabulary and specialized terms, (2) distinguishing between various meanings of a word of general vocabulary, and (3) taking into account the total context, including the intended audience and important details such as regionalisms. We will now look at mainstream linguistic theory to see how well it addresses these three types of difficulty. If mainstream linguistic theory does not address them adequately, then machine translation developers must look elsewhere for help in programming computers to translate more like humans.

Does Mainstream Linguistic Theory Come to the Rescue?

Mainstream linguistic theory emphasizes grammatical relations in a sentence. It is essentially a sophisticated form of sentence diagramming. Depending on when and where you went to high school, you may have encountered sentence diagramming or you may have missed it entirely. A sentence diagram shows all the words of a sentence and how they fit together. Mainstream linguistic theory has added a new dimension to sentence diagramming: Universal Grammar. According to Universal Grammar, there is only one method of diagramming sentences, this method applies to all the languages of the world, and it is universal because it is genetically encoded into the brain of every human child. This is a bold thesis and the large number of linguists around the world are working within this approach. Unfortunately, whether Universal Grammar is indeed universal or not, it says very little about the meaning of an individual word. It classifies words only according to the grammatical categories of nouns, verbs, adjectives, adverbs, and prepositions.

Not surprisingly, given the way it ignores word meanings, mainstream linguistics does not stack up very well when presented with the three types of translation difficulty we have discussed. It makes no mention of the distinction between general vocabulary and specialized terminology. This is because mainstream linguistics does not really deal with language in its entirety. It deals only with relatively uninteresting sentences that can be analyzed in isolation. Essentially, it deals with one very narrow slice of the pie of language that only appears to include general vocabulary and then calls that piece the whole pie. If it is true that mainstream linguistics does not really deal with the general vocabulary in all its richness, then it should be no shock to learn that it ignores the basic fact we have been exploring, namely, that a word can have several meanings, even within the same grammatical category. And if mainstream linguistics ignores the meanings of words, it has no need to take into account the context of a sentence. In fact, it has been a firm principle of mainstream linguistics for many years that the proper object of study is a single sentence in isolation, stripped of its context, its purpose, and its audience. This treatment of language on a local level (sentence by sentence) rather than on a global level has influenced the design of machine translation systems and we have seen the results in the telescope example.

It is a big job to take on the mainstream approach in any field. Actually, I am not saying that the mainstream is totally wrong. It does have many interesting things to say about grammar. Instead, I am saying that grammar, no matter how interesting it may be, is far from sufficient to teach a computer how to translate more like a person. Although none of the past three Barker lectures has dealt directly with translation, I detect in them considerable support for my thesis about the insufficiency of mainstream linguistics to deal with meaning, which I have shown to be highly relevant to translation. I trust my three colleagues would agree that mainstream linguistics does not treat meaning adequately.

Taking the past three Barker lectures in the order they were presented, we will begin with John Robertson, who warned us against the dangers of unwarranted reductionism. Robertson uses reductionism to describe an unwarranted oversimplification of a problem that leaves out an essential element. Reductionism, in a broader sense that I will use in this paper, is an approach commonly used in science. Reductionism, as suggested by the name, reduces a complex phenomenon to simple underlying components. It has in some areas been spectacularly successful, such as the reduction of visible light, infrared heat, radio waves, and x-rays to variations of a single phenomenon called electromagnetic radiation. But as Robertson points out, reductionism can go too far. In linguistics, the reduction of language to grammar separated from meaning is a highly unwarranted instance of reductionism. It may give the appearance of

allowing a scientific study of grammar, but ultimately it is a dead end approach that will not form a solid basis for studying other aspects of language beyond grammar and will not even allow a fully satisfying explanation of grammar.

My second colleague and Barker lecturer, Cheryl Brown, argued for the importance of words over grammar. Mainstream linguistics does not study real language as spoken by real people. Instead, it studies an “idealized, homogenous speaker-hearer community.” That is, it assumes that everyone has exactly the same internal grammar and vocabulary, that everyone is a carbon copy of everyone else. Brown ably shows through careful empirical studies that this idealization is not at all justified. She shows significant differences in the way men and women react to certain words. She gives examples of regional differences in the way certain words are used. And she shows that very advanced students of English in China are influenced by their culture in the connotations they give certain words. She illustrates the flexibility of humans in dealing with language, a flexibility which is not predicted by mainstream linguistics.

My third colleague and Barker lecturer, Jerry Larson, described the state of the art in regard to technology in language learning. He described many new developments that allow more sophisticated access to information, from text to sound to pictures to motion video. But he acknowledged that for a computer to evaluate the appropriateness of the speaking and writing of a student, when there is not just one predetermined response, we will need software that is “far more sophisticated than any currently available.” He rightly points out that such software would have to be able to recognize not just grammar but meaning and take into account the context of what is said and adjust for cultural factors.

This section of the paper was supposed to explore whether mainstream linguistics adequately addresses the types of translation difficulty I identified in the first section. I can now answer, with the support of my colleagues, in the negative. All three types of difficulty required a sensitivity to meaning, not just a mechanical attention to which words are used and how they are related grammatically. If mainstream linguistics cannot come to the rescue of those who want to program a computer to translate more like a person, then what kind of linguistics would it take? It is clear that it would take some approach to language that deals directly, not peripherally, with meaning. It is less clear what that approach should be. When you start working with meaning and try to pin it down so that it can be programmed into a computer, you begin to sympathize with the reluctance of mainstream linguistics to deal with meaning. And you come up against some pretty big issues in philosophy. For example, you eventually have to deal with the question of where meaning comes from. Are meanings already out there somewhere before we even make up words for them? Or do we create meanings out of nothing? How do we manage to communicate with others?

Some approaches to meaning assume that there is one basic set of universal concepts on which all other concepts are based. In this approach, which is sometimes called objectivism and dates back at least to Descartes, everyone, good or evil, must deal with these same starting concepts. I begin with the assumption that meanings are not absolutely imposed on us from the nature of the universe but that they are not entirely arbitrary either. Then where does meaning come from? I will now discuss a key factor that I believe to be missing from current theories of language. An approach to language that incorporates this factor should bring us closer to dealing adequately with meaning. Such an approach should guide us in the design of a computer that could translate like a person.

A Key Factor That Is Missing from Current Theories

That key factor which is missing from current theories is agency. By agency, I mean the capacity to make real choices by exercising our will, ethical choices for which we are

responsible. I will show a connection between agency and meaning. And since I have already shown that to translate we must consider meaning, I will then have shown that there is a connection between agency and translation. Any 'choice' that is a rigid and unavoidable consequence of the circumstances is not a real choice that could have gone either way and is thus not an example of agency. A computer has no real choice in what it will do next. Its next action is an unavoidable consequence of the machine language it is executing and the values of data presented to it. I am proposing that any approach to meaning that discounts agency will amount to no more than the mechanical manipulation of symbols such as words, that is, moving words around and linking them together in various ways instead of understanding them. Computers can already manipulate symbols. In fact, that is what they mostly do. But manipulating symbols does not give them agency and it will not let them handle language like humans. Symbol manipulation works only within a specific domain, and any attempt to move beyond a domain through symbol manipulation is doomed, for manipulation of symbols involves no true surprises, only the strict application of rules. General vocabulary, as we have seen, involves true surprises that could not have been predicted.

The claim that agency must be included in an approach to meaning is perhaps unexpected. I will draw on five different sources to support this claim: (1) some work by Terry Warner (BYU Dept. of Philosophy); (2) some work by John Robertson (BYU Dept. of Linguistics); (3) some thoughts on physics by Jack Cohen, a biologist, and Ian Stewart (a mathematician); (4) some work on neural science by Antonio Damasio; and (5) an analysis of Shakespeare's Othello by Nancy Christiansen (BYU Dept. of English). So far as I know, these various parties have never collaborated, yet they are presenting various pieces of what is beginning to look like a coherent framework of support for the importance of agency in fully explaining human language.

Terry Warner has been working on the notions of agency and self-deception for many years. I remember studying his writings on the subject already back in 1982. But at the time, I did not see the connection with translation. Of course, there is a connection between translation and language. You need to know at least two languages in order to translate. So the key question is whether language depends on agency. If it does, then translation depends on agency, too, at least sensitive translation of general language texts. Then, a few years ago, the BYU Dept. of Philosophy organized a seminar on the philosopher Levinas. This brought Terry and me together in a new way and eventually resulted my seeing a connection between agency and language and in the writing of a joint paper, which has now been expanded into a book (Melby and Warner, 1995). I will not talk here about our general collaborative work, but only of the use we have made of Levinas. Levinas talks about otherness. Someone who is 'other' is outside of you and not under your control. A physical object can, of course, be outside of you yet totally under your control. A physical object can be under your control intellectually, if in no other way, in that you can include a representation of it in a system of ideas that enables you to label it exhaustively and predict its behavior.

We totalize objects in the physical world when we bring them totally under our control. Levinas points out that when we attempt to totalize other humans, we are treating them like objects rather than like humans. But we speak and listen only on the presumption that we are communicating with beings who are not objects but beings with an inner life of their own, just like ours, whose background and individuality we can take into account and who can take into account our background and individuality. That kind of language, not as idealistically represented as if it were a domain language, but in its dynamic reality, has ethics as at least part of its foundation. Note that we have not said that ethics is based on language; we have said that language is based on ethics, making ethics logically prior to language. We present this unusual notion in more detail in chapter 4 of our book and Levinas develops it at length in some of his writing. In order to make ethical choices we must have agency, that is, we must be agents.

Unless we regard others as agents, just like us, who in turn regard us as agents, then many key notions that are a basis for general vocabulary become meaningless. Without this interacting agency, there is no responsibility, no empathy or indifference, no blame, and no gratitude. So much becomes missing from language that what is left can be described as a technical domain and handled by a computer. Agency is not a layer on language; dynamic general language is a layer on agency. Without agency, we are reduced to the status of machines and there is no dynamic general language. Without dynamic general language, we would translate like computers and there would be no truly human translation as we now know it. Thus lack of agency is one factor that keeps computers from translating like people.

As I re-read John Robertson's Barker lecture, I noticed that on page 15 he points out that if language were just a machine that tells whether or not a sentence is grammatical, then language would not allow personal relations with God and other humans. He notes that there was a war in heaven a long time ago. This war is mentioned in the New Testament (Revelation 12:7). According to other ancient accounts in the Pearl of Great Price quoted by Robertson, the main issue of the war in heaven was whether or not people would have agency. Happily, the pro-agency team won out. Our agency is a prized possession. Neal Maxwell, at the October 1995 General Conference of our university's sponsoring institution, speaking of will, an essential element of agency as I have defined it, said, "Our will is our only real possession." The anti-agency team, lead by Lucifer, would have totalized all humans and there would have been no will, no agency, and thus no human language as we know it. We would be like computers sending meaningless data back and forth.

Robertson is exploring an approach to language which, unlike mainstream linguistics, is compatible with agency. Robertson has intensely studied the works of C.S. Peirce and finds in them an approach to language that is compatible with agency. Hopefully, I will be able to further explore Peirce with Robertson in the future to better see how it applies to translation. But initially it would appear the Robertson's approach to language is compatible with the Warner approach in that they both include agency as essential to fully human language.

We now turn from philosophy to physics. Bear with me while I attempt to make a connection between them. The issue I am concerned with is whether our current understanding of physics is compatible with agency. As a youth, I had the impression that physics viewed the world as entirely deterministic. In other words, what will happen next is supposedly determined exactly and precisely by the current state of the physical universe. In a deterministic view of physics, there is no room for human agency because we are part of the deterministic system. If there is no agency, then it should be possible to program a computer to do anything a human can. So it would be nice if physics allowed for agency.

The view of the brain as a deterministic machine is still held by very intelligent people. For example, Patricia Churchland, Professor of Philosophy at the University of California, San Diego, recently (October 12 and 13th, 1995) gave a series of invited lectures on our BYU campus. Two of her titles are revealing: "Understanding the Brain as a Neural Machine" and "Am I Responsible If My Brain Causes My Decisions?" From my own attendance at one of her lectures and from reports of a colleague, it is clear the Churchland holds the view that we have no real agency since our future decisions are completely determined by the current state of the machine we call a brain and by input our brain receives from the outside. However, as we will shortly see, the view of the universe as purely deterministic is out of date in physics.

In their book, *The Collapse of Chaos*, Cohen and Stewart (1994) take the reader on a tour of modern reductionist science. In the reductionist approach, as already mentioned in the report on John Robertson's Barker lecture, the complexity of the world around us is analyzed in terms of simpler constituents that are linked together by relatively simple laws—the laws of nature.

Typical examples of successful reductionism are the equations for electromagnetic phenomena already referred to in the summary of the Robertson Barker lecture and the equations predicting the motions of the planets using Newton's laws. As Robertson has pointed out, an unwise use of reductionism has been damaging in linguistics, but I had assumed that it had been uniformly successful in physics. However, in the past decade or so, some of the implications of chaos theory have begun to sink in. Now even classical physics is not seen as entirely deterministic, even if it is exact in analyzing past events and predicting many future events. There are systems such that even the tiniest differences in initial conditions can lead to large differences at some future time.

Cohen and Stewart, who challenge the assumption that reductionism is sufficient even in physics, ask the intriguing question, *if complexity is explained by reductionism, then what explains the simplicity we see around us?* As one example, they consider crystals. The structure of a crystal is not readily explained in terms of the detailed vibrations of individual atoms. However, the structure of a crystal is probably influenced by the tendency to minimize energy, and this tendency is contextual rather than reductionistic. They go so far as to state, "We are surrounded by evidence that complicated systems possess features that can't be traced back [solely] to individual components." (pp. 426-427) In other words, reductionism is insufficient to explain the physical universe. William E. Evenson of the BYU Physics Department puts it this way (personal communication): "You have to make sure that the individual components are self-consistently adjusted for the context." You can't blindly build everything up from individual components without some notion of the big picture. That is quite an admission for a scientists and mathematicians. Cohen and Stewart do not claim that we must believe in God. Indeed, the claim that belief in God is a necessary consequence of science would be incompatible with agency, since there would be no room for faith.² But they point out that modern physics is not incompatible with a belief in God. They even refer to an interpretation of physics that leaves room for human agency. Cohen and Stewart, along with many others, discount some of the speculations of Roger Penrose (1989 and 1994), a mathematician who thinks that consciousness comes from quantum effects in a certain part of the brain. But they agree that the question of consciousness is an important one.

The interpretation they support comes from Freeman Dyson (Cohen and Stewart, 1994, p. 272) and does not depend on the details of the brain. In quantum mechanics, it is well known that you cannot measure the exact position and speed of a subatomic particle without influencing the position and speed of the particle by the process of measuring. This introduces a truly random element into the physical world, which means that the future is not absolutely determined by the past.³ Dyson says that quantum mechanics describes what a system might do in the future, while classical mechanics describes what the system ended up doing in the past. He suggests that our consciousness may be at the moving boundary between future and past, that is, the present. This interpretation of physics says that the future cannot be computed exactly even though the past can be analyzed exactly, leaving open the possibility of free will and thus agency through choices in future action. Hopefully, I have now made a convincing connection between physics and the philosophy of agency.

Dyson's explanation is reminiscent of the way word meanings shift. They are unpredictable in advance, as in the treacle example, but they can always be analyzed when in the past and a motivation can be established in retrospect. Cohen, Stewart, and Dyson have opened up to me a new view of physics. This new view is compatible with both modern physics and with a linguistics based on agency rather than deterministic generation of sentences.

Until recently, I assumed that the highest levels of translation would require a personal understanding of emotions, but I did not see any connection between emotions and other mental

functions needed for human-like translation. From brain science comes surprising support for a connection between emotion and human reasoning. Human reasoning is an essential aspect of agency. What good does it do to have the ability to make choices if one cannot use even common sense reasoning in making decisions. Now, on the basis of recent studies, the need for emotions is not a separate requirement for human-like translation. Agency and human reasoning ability imply feeling emotions, because without emotions, human reasoning is impaired. Antonio Damasio, a well-respected neurologist, has published an intriguing book, *Descartes' Error* (Damasio 1994), which challenges the claim made by Descartes that reason and emotion should be kept separate. Damasio draws on case studies of unfortunate people who are the victims of damage to a certain area of the brain, damage that robs them of the ability to feel emotions. Damasio shows conclusively that the inability to feel emotions hinders their ability to reason normally and make common sense decisions. For one thing, they become insensitive to punishment. In a way they may lose some part of their agency, since they can no longer feel emotions. Considerably more work is needed, probably in the form of masters and doctoral theses, before firm conclusions can be drawn. But in these cases of brain damage, along with a loss of common sense, I predict there will be a detectable loss of ability to produce sensitive translations of certain kinds of texts well, unless the patient's memory of having felt emotions is sufficient to maintain a full capability in language. This discussion should be continued as more evidence accumulates.

Finally, we turn back to Shakespeare and find that he may have understood the connection between language and agency all along. Nancy Christiansen (1994) points out that Othello is trapped because he can only see one interpretation of events at a time. We could say that he loses some of his agency by getting trapped in a domain. Iago, on the other hand, is acutely aware that multiple interpretations of the same facts are possible, but denies some part of his agency by denying any connection between ethics and choices. Shakespeare, meanwhile, sits back and sees all sides. He recognizes agency (which is ethics in action) as the basis of language. I wonder what would have happened if Shakespeare had been chosen as the linguist of his day. Perhaps everyone would be convinced that agency is needed for human-like translation. This effort to balance agency and determinism would then be much ado about nothing. Shakespeare saw the wave of determinism that has engulfed our generation, and saw beyond it. Great literature was never taken in. Further dialogue with Christiansen on agency and language is in order.

These five sources fit together in that they all are compatible with the claim that agency is essential to the richness of normal human language, as opposed to machine-like domain language. Warner speaks of both language and agency being based on ethics. Robertson claims that agency is essential to the development of relationships. Cohen, Stewart, and Dyson show that agency is compatible with modern physics. Damasio shows that fully human reason, which is essential to agency, is tied to emotions. And Christiansen shows us that Shakespeare understood the connection between ethics, language, and agency long before I started thinking about it.

Our concepts are not based on some absolute self-categorization of the physical universe. They are based in part on the ethical dimension of our relationships with others. Our agency, which includes both emotion and reason and the ability to choose how we will respond to demands placed on us by others, is the basis for human language as opposed to machine-like language.

Finally, we can answer the question of this paper. A computer cannot translate more like a person because it lacks, among other things, agency. It won't suffice to store massive amounts of information. Without agency, information is meaningless. So a computer that is to handle language like a human must first be given agency. But we should be careful, because if we give

agency to a computer it may be hard to get it back and the computer, even if it chooses to learn a second language, may exercise its agency and refuse to translate for us. Douglas Robinson (1992) puts it well. He asks whether a machine translation system that can equal the work of a human might not “wake up some morning feeling more like watching a Charlie Chaplin movie than translating a weather report or a business letter.”

Endnotes

1. An example of the need to be sensitive to cultural factors is the translation of descriptions of items on a menu in a restaurant. Last year while in Paris, I passed by a billboard outside a well-known restaurant. The billboard advertised a dish called *steak tartare*. The description in English mentioned that it included fresh ground beef and egg, but failed to mention that the ground beef is served completely raw. In fact, this dish typically consists of raw ground beef mixed with spices and topped with a raw egg and bits of raw onion. For an American tourist passing by, there was not a clue that the meat was served raw. For a British tourist, this may be common knowledge, and for a Frenchman, it is no big deal that the meat is raw. Their ‘well done’ is more like our ‘rare,’ and they sometimes order a steak bleu, literally, ‘blue,’ meaning barely warmed over a flame but not cooked in the American sense of meat preparation. An American could easily think that the word tartare refers to tartar sauce and order the dish thinking that it would be strange to serve tartar sauce with beef instead of fish but certainly expect the meat to be cooked. This example shows the importance of being aware of the differences between cultures when translating.
2. Douglas Adams, in *The Hitchhiker's Guide to the Galaxy*, gives us a whimsical account of why science should not attempt to prove the existence of God. The account is particularly appropriate for this paper since it involves translation.

In Adams’ novel (pp. 58-60), there is a small, yellow fish called the Babel fish that feeds on brainwave energy. If you place a Babel fish in your ear, you can understand anything said to you in any language. The Babel fish is thus extraordinarily useful, especially for someone hitchhiking across the galaxy. However, the story continues, some thinkers have used the Babel fish as a proof of the non-existence of God. The argument goes like this. It would be such a bizarrely improbable coincidence that anything so useful as a Babel fish could have evolved by chance, that we can conclude it did not evolve by chance. God refuses to allow a proof of his existence, since that would deny faith. But since the Babel fish could not evolved by chance, it must have been created by God. But God would not allow a proof of his existence. Therefore, there could be no God.

The silliness of the above argument is intended, I believe, to show the futility of trying to prove the existence of God, through physics or any other route. Belief in God is a starting point, not a conclusion. If it were a conclusion, then that conclusion would have to be based on something else that is firmer than our belief in God. If that something else forces everyone to believe in God, then faith is denied. If that something else does not force us to believe in God, then it may not be a sufficiently solid foundation for our belief.

Adams may also be saying something about translation and the nature of language. I can speculate on what Adams had in mind to say about translation when he dreamed up the Babel fish. My own bias would have him saying indirectly that there could be no such fish since there is no universal set of thought patterns underlying all languages. Even with direct brain to brain communication, we would still need shared concepts in order to

communicate. Words do not really fail us. If two people share a concept, they can eventually agree on a word to express it. Ineffable experiences are those that are not shared by others.

3. There is a famous thought experiment in quantum mechanics devised by Erwin Schrödinger. A cat is placed in an impenetrable box, along with a radioactive atom and a device that detects whether or not the atom has decayed, releasing poison gas when it does decay. According to one interpretation of quantum mechanics, the cat is in a state of superimposed life and death until some measurement is made. Many people have written about this thought experiment, which seems so counterintuitive. Cohen and Stewart do not challenge the evidence that quantum effects introduce true randomness, but they do challenge the assumption that the cat can be both alive and dead at the same time. They discuss what it means to make a measurement, and they suggest that the cat itself knows what is happening, invoking T. S. Eliot's poem "The Naming of Cats" (found, among other places, in *The Norton Book of Light Verse*, edited by Russell Baker, (c) 1986, W. W. Norton & Company: New York).

Annex: A showcase example of machine translation

The following English sentence was taken from a source text chosen by a major machine translation vendor. The source text was translated by computer into French, German, and Spanish and the output was offered as an example of what machine translation is like when things go well. Even here, note the different ways the abbreviation ATP was handled. In this English text, it obviously stands for Advanced Technology Program. However, in the French text, it was expanded into a French chemical term for the organic chemical that is called 'adenosine triphosphate' in English. This compound, which is broken down from more complex substances such as sugars, is the immediate source of energy to the cells of our body. Clearly, this is a serious translation error. The compound adenosine triphosphate, abbreviated ATP, has nothing to do with the Advanced Technology Program, also abbreviated ATP, except that it has the same abbreviation. The computer mechanically substituted the full French form for the chemical use of ATP, demonstrating a lack of understanding of what is being translated. In the German translation, something strange is going on as well. The abbreviation has been reduced to lower case, except for the first letter. This is probably because it has been treated as a normal German noun, and all German nouns are capitalized. In the Spanish translation, the acronym was left untouched. This is a mistake as well, since the full form was translated and the Spanish version should be abbreviated as PTA.

* English source text: *The purpose of the Advanced Technology Program (ATP) is to assist United States businesses to carry out research and development on pre-competitive generic technologies.*

* French machine translation: *Le but du programme de technologie de pointe (triphosphate d'adénosine) est aider des entreprises des Etats-Unis d'effectuer la recherche et le développement sur des technologies génériques précompétitives.*

* German machine translation: *Der Zweck des Programms der neuen Technologie (Atp) ist, Staatgeschäfte zu unterstützen, Forschung und Entwicklung auf vorwettbewerblichen generischen Technologien durchzuführen.*

* Spanish machine translation: *el propósito del programa de la tecnología avanzada (ATP) es asistir a los negocios de Estados Unidos realizar la investigación y el desarrollo en tecnologías genéricas precompetitivas.*

References

- Adams, Douglas. (1979) *The Hitchhiker's guide to the galaxy*. New York: Harmony.
- Christiansen, Nancy Lee. (1994) *Shakespeare's paideia: Rhetoric as ethics in the Tudor inheritance and in "Othello."* University of California Los Angeles PhD thesis, University Microfilms International order number 9500680.
- Cohen, Jack and Ian Stewart. (1994) *The collapse of chaos: discovering simplicity in a complex world*. New York: Penguin Books.
- Damasio, Antonio R. (1994) *Descartes' error: emotion, reason, and the human brain*. New York: Grosset/Putnam.
- Dreyfus, Hubert L. (1992) *What computers still can't do*. Cambridge, MA: The MIT Press.
- Melby, Alan and C. Terry Warner. (1995) *The possibility of language: a discussion of the nature of language, with implications for human and machine translation*. (in press) Amsterdam: John Benjamins Publishing Company.
- Penrose, Roger. (1989) *The emperor's new mind concerning computers, minds, and the laws of physics*. Oxford: Oxford University Press.
- Penrose, Roger. (1994) *Shadows of the mind: a search for the missing science of consciousness*. Oxford: Oxford University Press.
- Robinson, Douglas. (1992) "Neural networks, AI, and MT" in *The ATA Chronicle* volume XXI, Number 9 (October 1992). The American Translators Association.