Machine Translation and Philosophy of Language

Alan Melby

Brigham Young University at Provo, USA

1 An image with multiple interpretations

The other day I woke up with an image in my head, but I had no idea what it meant. The image was simple, a straight horizontal line with an oval above it that touched the line, and I was sure that it had meant something in my dream, but I had only a faint, rapidly fleeting, recollection of my dream. You have probably felt the frustration of trying to bring back a dream.

I sat up and tried to grasp what the image might signify. My first impression was that it could be a balloon filled with water sitting on a table and turned so that you cannot see the mouth of the balloon that has been tied up.

I had the feeling that the image had something to do with humans, so I looked beyond balloons. My next thought was that it might signify another person encountered along the way during a stroll through a park. You don't know what to make of the other person. You can try to squeeze each of them into a mould based on the way they are dressed, but you can't really get to know them unless you interact with them and let them come out of their mould. The line could represent the path you are walking along, and the oval the mould that you put that person into when you form a first impression.

That still wasn't quite right. Perhaps the oval was a knot-hole in a piece of wood, and the line was the lower edge of the wood. But what did that have to do with people? Ah yes, the knot-hole-wood image reminded me of a story about a visit to a tree house which had boards on all sides. It didn't matter that grown-ups who build houses generally put wood panels vertically. This tree house was more in the style of a log cabin. According to the story, a very imaginative boy, Harold, was invited into the tree house by Peter, the neighbour boy who had built the house. Peter closed the door and the small room became nearly dark, except for one beam of sunlight that shone through a knot-hole in one of the boards. That knot-hole was, of course, the oval in the image from my dream. Harold suggested playing a game in which he would pretend that he knew nothing about the outside world, that he had always lived in the tree house in the dark. Peter, who had come to expect crazy thought games from Harold and enjoyed playing them, pointed out the dust particles floating in the beam, and they both watched them for a few minutes. Then Peter moved over so that the beam shone onto his face and began describing to Harold some of the things that were happening outside the tree house. Harold would not at first believe that Peter could see anything outside the tree house, insisting that he could see the beam just as well as Peter could and that there was nothing to see but dust particles floating in the air. Then Peter helped Harold look along the beam instead of at it, and the outside world opened up.

There is no single correct interpretation of an image except, perhaps, within a domain. For example, on a Forest Service map, the oval might unambiguously mean a campsite and a small rectangle might mean that the campsite has a picnic table. In another domain, the same symbols may mean something entirely different. There is no well-defined limit to the number of possible domains in which an image could have particular domain-specific meaning and no limit to its possible interpretations in general language. The world is infinitely categorizable. But interpretations are not exactly random meanings either. Each is somehow motivated by the original image. The first interpretation of the image in my dream is motivated by a similarity with the shape of a balloon filled with water. The second suggests the unjustified oversimplification of

a human to a stereotype which supposedly allows us to predict how that person will behave (a kind of extreme racism). And the third interpretation of the line and oval suggested at a second level the story of the beam of light shining through a knot-hole into a tree house. One could say that the first interpretation is literal while the second is metaphorical and the third is both literal and metaphorical. But note that, contrary to a common assumption about metaphor, the metaphorical interpretations are not based on the literal balloon interpretation.[1] [Editors' note: The footnotes are all given at the end of the paper under 'Endnotes'.] However, the beam-of-light story has a particular significance to the philosophy of language related to the second metaphorical interpretation. We concur with the philosopher Emmanuel Levinas who claims that selfhood is based on a recognition of otherness in the sense that other people also possess selfhood. Without others, selfhood has no meaning. However, with particular individuals, we can attempt to deny them their agency. One way is to stereotype another person and claim that a label is all we need to know about the other (as in the second interpretation of the oval). Another way to attempt to deny the agency of others is to pretend that they are not relevant to us by putting up barriers around us to shut out everything but our own little world (as in the story of the tree house). When we stop resisting the otherness of others and grant them the same agency we possess, then a whole new world opens up (suggested by looking along the beam instead of at it), an exciting world which is fundamentally ambiguous yet grounded in the ethics and economics of human relations, a world which opens up the possibility of dynamic general language.

As long as we are comparing general language to a world, let us extend the metaphor to include domain-specific language. Start with our planet earth and its various languages/cultures focused in various geographical areas; then think of the various satellites orbiting the earth as artificially created domains. Some domains, such as the domain of the maintenance and repair manuals for a piece of machinery sold world-wide, will be almost completely shared across several languages, just as one satellite can transmit to several areas of the earth. Even then there will be minor variations such as the voltage and frequency a machine expects when fed electrical power and the type of plug placed on the outlet to obtain that power. Other domains are tied to one culture, sometimes even within the same language, just as a weather satellite and a military satellite may be in the same orbit yet be incompatible. For example, the domain of Law in the United States and the domain of Law in Great Britain are two domains, whether you count British English and US English as one language or two. Translating between incompatible domains can be as challenging as translating between general languages.[2]

Human translators are able to handle both general-language and domain-specific texts. As a starting point, a translator must be competent in two or more general languages. Then, for each new domain, the human translator must gain new expertise. The same requirement applies to a machine translation system in that the lexicons as well as any knowledge base the system may have must be updated in order to produce high-quality translations from a new domain. But here the similarity ends. Human translators can produce high-quality translations of general-language texts which are dynamic, that is, full of metaphor, allusions, and intentionally unusual usage. Current machine translation systems cannot. Current techniques in machine translation produce fully-automatic high-quality translation only when applied to a body of similar texts which are all restricted to the same domain. The texts must be static in that they do not contain new metaphors, allusions, or grammatical constructions. Sometimes this restriction occurs naturally and the texts form a sublanguage. More often, the restrictions must be enforced with the cooperation of authors, resulting in what is called controlled language. Many have noted that machine translation works better in a narrow domain. The reason in obvious: everything is better-defined and less ambiguous than in general language. What is less obvious is whether the machine translation techniques that work quite well within a domain can gradually be extended to apply equally well to general language. Or will one encounter a phenomenon of diminishing returns or even an unscalable wall. Terrance Hook, who has developed a domain-specific Dutch-English machine translation system, made a typical comment. He said that when restricted to a domain, the output of his system is good enough to be used, as it is, for some purposes. However, when, "on a rainy afternoon", he tries a passage from a newspaper, he gets gibberish. Is this a temporary limitation of domain-specific systems or will they gradually improve in their ability to

handle general language texts until they do as well on general language as they do on domain-specific texts? We claim that current techniques of machine translation will never be extended to handle general language texts.

2 Techniques that do not extend

A major shift has occurred in machine translation. John Hutchins, the acknowledged historian-in-residence of machine translation, has noted [3] that up to about ten years ago the assumption was that systems should be general; but now the assumption is that systems (at least systems aimed at high-quality output) should be domain-specific. The issue among professionals is no longer whether current techniques in machine translation work equally well in a domain and on general text. They do not. The issue is whether current techniques can ever be extended to handle general language effectively. I have proposed that they cannot be extended to dynamic general language. This claim is highly controversial.[4] How could I be so bold as to make such a claim? The reason is based on (1) the fact that current techniques depend on philosophical underpinnings called objectivism and (2) my belief that general language does not conform to the assumptions of objectivism, thus invalidating current techniques as applied to general language.

George Lakoff, a prominent linguist and early supporter of Chomsky, long ago broke off from the objectivist camp and has spent recent years developing a non-objectivist approach called experientialism. He summarizes objectivism as the belief that:

> rational thought consists in the manipulation of abstract symbols and that these symbols get their meaning via a correspondence with the world, objectively construed, that is, independent of any organism
> (Lakoff 1987:xii)

This view has many implications. It implies that the human mind is an abstract machine and that any machine, including a digital computer, which is properly programmed, is theoretically capable of thinking just as well as or even better than a human mind. Note that this view includes a strong form of mind/body dualism, which means that a human body is not at all necessary for human-like thought. Some researchers in Artificial Intelligence (AI) take what is known as the strong-AI position, which is that computers will someday be able to perform any intellectual task that humans can perform.[5] Marvin Minsky, a strong-AI proponent, recently wrote an article in which he estimates the knowledge a human acquires over a lifetime amounts to not more than the equivalent of about three gigabytes, which is approximately the amount of information that can be stored on one CD-ROM. He speaks of nanotechnology that places individual atoms in desired positions and that will allow us to produce much smaller and faster computer chips than we now can build. He then states, speaking of future robots as our virtual offspring, our **MIND-CHILDREN**:

> Once we know what we need to do, our nanotechnologies should enable us to construct replacement bodies and brains that will not constrain us to work at the crawling pace of "real time". The events in our computer chips already happen millions of times faster than those in brain cells. Hence, we could design our "mind-children" to think a million times faster than we do.
> (Minsky 1994:90)

Minsky also notes that many scholars from a variety of disciplines "firmly maintain that machines will never have thoughts like ours because, no matter how we build them, they will always lack some vital ingredient". Minsky says he has no patience with such arguments because they are all flawed by assuming, in one way or another, "the existence of some magical spark that has no detectable properties".

Although over the years I have generally had little patience with Minsky and his outrageous claims [6], he has a good point here. In a post-religious society such as ours, it does little good to use an 'undetectable magical spark' as the basis for an academic claim. Instead I have decided to focus on what hurdles would have to be overcome by a machine before it would even have a chance of handling dynamic general language better than or on a par with humans. I do not claim that it will never be possible to build machines that can think like humans and, in particular, can handle dynamic general language as well as humans. Instead, I try to show that the current techniques of natural language processing (NLP) will never be extended to accomplish such tasks. Entirely new techniques will be needed. In particular, we will need techniques that avoid the assumptions of objectivism. We will see why in the next section.

3 Avoiding objectivism

Both mainstream philosophy and mainstream linguistics have built into them assumptions based on objectivism. Here are some of those assumptions:

(a) Words and fixed expressions such as multi-word terms are mapped to a short list of discrete senses, often to a single sense.

(b) Each sense exists independently of any particular word or sentence and has the properties of a mathematical set. For example, the sense of horse that corresponds to an animal is a set of objects in the real world. Any particular object is either in the set (if it is a horse) or is not in the set (if it is not a horse). There is nothing in between. Since these senses are independent of particular sentences and independent of people, they correspond to the way the world is, to the way the word objectively divides itself up.

(c) The meaning of a sentence treated in isolation can be obtained by combining word senses of the words of the sentence from the bottom up. If a word of the sentence is ambiguous then there may be multiple composite sentences for the sentence, unless all but one are weeded out by selectional restrictions.

These assumptions are embedded in the standard framework which divides language into syntax (including morphology), semantics and pragmatics, with emphasis on syntax and semantics. According to this framework, linguistics is a branch of individual rather than social psychology. To someone committed to the mainstream view, this framework is perfectly standard and obviously true. There are many flavours within Generative Grammar, but they all share this framework and most work in machine translation is explicitly or implicitly based on it.

However, dynamic general language violates all three of the basic assumptions listed as (a), (b) and (c). It violates assumption (a) in that new word senses, sometimes called nuances, can be generated dynamically as needed in speech or writing, often for the purposes of a single text. Indeed, this dynamic aspect of meaning is found in all interesting writing, not just in great literature. Only in a well-defined domain can the meanings of words be pinned down. And that is because we humans create a domain specifically so that the senses of a term will be limited and discrete, with the goal being one concept per term and one term per concept in each language.

Dynamic general language also violates assumption (b) in that its categories are not mathematical sets tied directly to the way the world divides itself up. Lakoff (1987) gives abundant evidence to

this effect from several disciplines. For example, he shows that categories of general language exhibit prototype effects in which some members are better members than others [7], a behaviour not allowed in mathematical sets. Again, in a domain, we divide up the world a certain way for a particular purpose. So from the point of view of the domain, the world can be seen as divided up into a neat ontology of domain concepts which are mathematical sets.

Assumption (c) is violated in that general language is always understood in a certain context. Martin Kay and his colleagues (Kay et al 1994) put it this way: "language is situated". When humans process general language, they do not delay consideration of pragmatic factors such as the situation. The syntax, then semantics, then pragmatics model is only applicable to domains in which the situation is constant and therefore implicitly taken into account at all levels.

So we see that dynamic general language violates all three assumptions on which most natural language processing is based. But controlled language restricted to a well-defined domain conforms to all three assumptions if we engineer it so. At a dinner speech, Martin Kay once put it something like this: "Success in NLP has been seen primarily in cases where natural language resembles formal language". That comment, although intended to be humorous, is on-target and has a serious side. The syntax/semantics/pragmatics model of bottom-up composition from well-defined concepts is essentially a description of a formal language such as a computer programming language. Formal languages conform to all three assumptions while dynamic general language conforms to none of them. Thus, NLP techniques that are based on these assumptions apply to domain-specific text inasmuch as it resembles formal language and inasmuch as it does not exhibit the dynamic possibilities of general language.

Thus we can conclude that current NLP techniques will never be extended to handle dynamic general language, since to do so they would at least have to abandon the three basic assumptions of this section. Any set of techniques which truly abandoned these principles would look so different from current techniques that it would be inappropriate to call them an extension of current techniques. But what can we say about how these new techniques would look?

4 What is needed

Please recall that I am not saying that there are no techniques which can handle dynamic general language. I am saying that current techniques are insufficient. So what would be sufficient? First, the new techniques would allow for fundamental ambiguity. Fundamental ambiguity goes beyond superficial ambiguity in that it entails both an indeterminate list of possible senses for a word and an indeterminate relation between the senses and the real world. Most people in NLP to whom I pose the question of whether they believe in a universal set of concepts determined by the structure of the universe will respond that they do not believe in any such thing. Yet their techniques are based on this assumption. Again, within a domain, we can act as if there were no fundamental ambiguity so long as we have a group of people who have come to a shared understanding of the concepts of the domain. This shared understanding comes about through human experts interacting in a mixture of general language and specialised terms. General language provides the metalanguage for arriving at a common understanding. But this approach falls apart when applied to general language, because there is no metalanguage in which to discuss general language. Yorick Wilks [8] has pointed out this problem when he asked how one can know whether everyone in a co-operative effort has the same understanding of the primitive concepts of an interlingua. This leads to the philosophical problem of the given. How do we obtain the atomistic concepts that are used to build up more complex concepts? What gives them to us? Chomsky would say that they are genetically hardwired. Philosophers would say that if they are not hardwired and we do not have them as children then we cannot get them through direct experience since concepts are required to interpret our experience.

A satisfactory solution must overcome the problem of the given. Chomsky's solution is unsatisfactory since it does not allow for fundamental ambiguity. One criterion that a satisfactory solution must pass is the test of dynamic metaphor. Current NLP techniques can easily handle frozen metaphor. We simply put a fixed expression in the dictionary. Although even there we run up against resolving ambiguities such as the English request to go jump in the lake or the French request to go cook oneself an egg, which may be literal requests to perform a specific task or idiomatic requests to just leave and not come back, depending on the situation. Dynamic metaphor is much more challenging than frozen metaphor. Dynamic metaphor is created for the purposes of one text or even one sentence. Understanding dynamic metaphor involves taking into account the entire situation and those aspects of general knowledge that are relevant to the situation. It is ultra context-sensitive and thus contrasts with the objectivist processing which assumes that the meaning of a sentence can be built up without taking into consideration the context at all. Some dynamic metaphor is so clever or poignant that it is frozen and preserved for future use. The prevailing wisdom is that metaphor is a secondary aspect of language that can and should be ignored until other problems are solved. Lakoff has shown that it is a pervasive aspect of language that needs to be solved up-front. Certainly, for general language, we cannot afford to ignore it. An interesting aspect of metaphor is that, although one cannot prepare in advance a list of all possible metaphorical uses of a word and although once a dynamic metaphor is created one cannot predict how it could be appropriately translated, every metaphorical usage is in retrospect motivated rather than random.

Ian Kelly supplied me with an interesting sense history of the word ‘treacle’ over the past two thousand years. At each change in sense, there was dynamic metaphor at play and each change is motivated though some are surprising. The ancestor in Ancient Greek of the word ‘treacle’ was a wild animal. It then metonymically became the bite of a wild animal. This sense then broadened to become a general injury and later shifted to the medicine used to treat such an injury. Later still it narrowed to the substance put into a medicine in order to make it more palatable and finally, in British English, to one such substance, molasses. Each step is logical and motivated for a human, but it would be asking too much of a machine based on objectivist assumptions to figure out the new meaning at any stage of the transition from wild animal to molasses. Some NLP projects have worked on understanding dynamic metaphor. They should not be expected to achieve human levels of performance unless they truly abandon their objectivist assumptions. But at least it should be possible to measure their performance in such tasks as translating texts containing dynamic metaphor.

Is there anything else that would be needed in a viable approach for handling dynamic general language? Yes, it would be important to avoid falling into radical relativism when allowing for fundamental ambiguity. Radical relativism, typified by the Deconstructionist movement in literary theory, recognises the problem of the given and solves it by saying that nothing at all is given. Concepts are not genetic, neither are they built into the structure of the universe. Everything is relative. The problem with this approach is that it does not explain how we can communicate. How do we know that our concepts have anything to do with the concepts in the head of the person we are talking to? A series of distinguished philosophers, including Heidegger and, in his later work, Wittgenstein, have struggled with this problem. They have concluded that our concepts are grounded in our social interactions. This is a promising direction. Note that it implies that general-language linguistics is a branch of social rather than individual psychology.

Often it is said that a computer that could translate anything would have to understand what it is translating. But how do you tell if a computer understands? John Searle proposed a puzzle ("Minds, Brains, and Programs" John R. Searle, from The Behavioural and Brain Sciences, vol. 3, (c) 1980 Cambridge University Press) in which it is assumed that techniques are somehow developed which allow a person sitting in a box who speaks only English to answer questions about a story by mechanically following a set of rules. The catch is that the story, questions, and the answers are all in Chinese and the person is English monolingual. Within a domain, say the domain of Chinese weather bulletins, this could probably be done if someone who speaks only

English could follow rules similar to those used by the Meteo system to translate weather bulletins between English and French. Of course, it may take quite a while for the person in the box to produce an answer, but let us ignore that problem. The question is whether the ability to mechanically produce acceptable answers would constitute a demonstration that the person understands Chinese. Most people would say the answer is obviously no, while strong-AI people would say the answer is obviously yes.

Searle is on the side of those who think the answer is obviously no. He points out that if he were the person following the mechanical rules, he would get out of the box without knowing any Chinese. He would still know English and understand questions posed to him in English, but he would not understand Chinese. He points out that some people have suggested that an adding machine **UNDERSTANDS** arithmetic and that a door that opens automatically when someone approaches it and breaks a beam of light **UNDERSTANDS** the instructions of the photocell. He points out that this sense of 'understand' is not at all the same as the sense in which we note a person understands Chinese. Searle then goes through several types of replies he has received to his argument from strong-AI types. One type of reply is that perhaps a person that blindly follows the rules sitting closed up in a box does not understand Chinese, but if the rules were programmed into a small computer that was put into a robot, then the robot, thanks to its ability to move about and see things, would understand. Searle replies that this implies that understanding is solely a matter of formal symbol manipulation, which is one of the tenets of objectivism. Searle counters the symbol-manipulation theory by noting that an essential element of understanding is conscious intentionality. Most people would accept this. The problem is how to detect whether a machine intends to do something or merely follows a series of instructions. Strong-AI proponents must logically accept a form of mind-brain dualism, namely that the mind, including its intentionality, can be successfully implemented in a digital computer or in a human brain or, presumably, in 'Cartesian mental substance' if we ever run across any of that stuff, whatever it is. Searle takes delight in pointing out, however, that the AI literature contains "frequent fulminations against 'dualism'". Searle rejects this form of dualism and expresses his belief that intentionality is a biological phenomenon. If this is so, he points out, we should no more expect a computer program to have intentionality than a computer software simulation of photosynthesis to produce sugar. The problem is that unless we can somehow detect intentionality and prove that it is a biological phenomenon, we have a stand-off between Searle and Minsky. They may agree that a computer needs understanding and that understanding entails intentionality, but that leaves unanswered the question of whether a computer can have intentionality. In line with my previous stance of attempting to identify specifically what would be needed for a computer to handle general language rather than just saying that it would need an undetectable spark, we should perhaps look for indirect ways to detect understanding and intentionality.

Marvin Minsky, in the same article where he pontificates about artificial brains, says something with which I agree, namely, that one thing which separates current machines from humans is the flexibility of the human mind. When a computer program encounters a situation for which it has not been explicitly programmed, it stops or produces meaningless results. When humans encounter a new situation, they are able to try various solutions until something works. This applies to Searle's Chinese Box puzzle. Flexibility is a detectable aspect of understanding and intentionality. Even Meteo has occasional problems with a sentence, usually due to a typographical error or noise on the transmission lines. A human reviser handles these situations because they cannot all be systematized and therefore require the flexibility of the human mind. The human blindly following the instructions of Meteo would exhibit no more flexibility or robustness than a computer. Therefore, neither a computer, nor a human following instructions mechanically, truly understands. We have now made an additional requirement of a machine that might handle natural language. It must exhibit flexibility in handling new situations. This flexibility would probably be related to the ability to handle dynamic metaphor. It seems that ways of testing flexibility could be devised.

Joseph Weizenbaum is well-known for having written a computer program called Eliza that simulates a psychoanalyst. When it was first installed on a computer at a university, some people would **TALK** to it for hours on end through a computer terminal, exposing their darkest secrets and actually believing that it was a human psychoanalyst or at least that it really **UNDERSTOOD** them. Weizenbaum was appalled.[9] He knew that the computer program didn't understand a thing they were saying. It simply looked for key words and put together minor variations on stock replies. For example, if a person said: "My parents are divorced", Eliza would reply something like: "Tell me more about your family", using a table that lists 'parent' as a 'family' word. How was Eliza so successful in fooling intelligent people? First, it was dealing with a domain, the domain of the detached psychoanalyst gathering data. Eliza never said anything substantive, even mundane things. It could not even answer a question like "How many days are there in a week?" It just asked questions to keep the person talking, and who doesn't like to talk when **SOMEONE** will listen? Eliza clearly fails the flexibility test of being able to handle a new situation.

It is instructive at this point to look at one other person who has written about mechanistic approaches to language. Roy Harris (1987) in his book *The Language Machine* traces the history of the idea that human language can be put into a machine, going back to Gulliver's Travels in which there is a section about a machine which randomly produces sequences of words. Young men are employed for the purpose of sifting through the random sequences for ones that have meaning and putting the sequences together into books. This satire on a wrong way to create literature is surprisingly not too different from the deadly serious way that a Chomsky-style grammar randomly produces sentences in isolation to supposedly generate a human language, except that it is semantic rules that eliminate the millions of sequences that do not make sense instead of a room full of humans, a process euphemistically called 'overgeneration' and 'selection'. Along the way we find Saussure who posited a language machine in the brain in order to distinguish linguistics from language teaching. For him, the language machine was automatic and so no one had control. Thus there was no need to teach the inner workings of the language machine to humans and no danger of language teachers taking over part of linguistics. What is missing from Saussure is any mention of bilingual humans or social class differences in dialect. They were erased by the idealized *langue*. Saussure spoke out against prescriptivism, but, ironically, it was during Saussure's lifetime that the idea of a standardized national language arose, a triumph of prescriptivism, with theoretical support from Saussure's idealization of language. A national language is a creation which gives a false idea of uniformity and contributes to the view of language as a machine. Then Chomsky completed the project by making language into a machine that functions completely without human intervention. For Harris, the view of language as a machine has contributed to the exclusion of a moral dimension from language and a devaluing of a search for solid truth and knowledge, resulting in radical relativism. Another bizarre consequence of the language machine view is that communication is only an incidental aspect of language instead of the core aspect.

There is a contradiction between the model of language as a machine that is independent of social interaction and the deepest yearnings of the authors of these models. Chomsky, in a documentary on his life and work, stated that, although he has sought a connection between his linguistic theory and his political activism, which centers on manipulation of public opinion by the press, he has found none.[10] Shouldn't that lack of connection be worrisome? And Minsky, in his article about artificial brains, makes the rash claim that "No popular ethical system yet, be it humanist or religion-based, has shown itself able to face the challenges that already confront us." He is clearly concerned about the meaning of life for himself and others. He even ends his article with a sermon-like plea: "Our job is to see that all this work shall not end up in meaningless waste". It seems that a good place to start would be to place social interaction at the core of language and give some of the long-established ethical systems a chance to work instead of undercutting them.

Now we can put together the previous points of flexibility and social interaction to avoid both the problem of the given and radical relativism. We need a flexible grounding of language that allows for social interaction at the core of language. This leads to the work of Levinas on the questions of

interiority and totalization. Levinas has shown that to be an interiority, that is, to have selfhood and agency, which is also an essential part of consciousness and understanding, one must acknowledge the existence of other interiorities that cannot be totalized. Put in more familiar words, to be a living, thinking person, one must acknowledge the existence of other thinking, living persons who are peers and agents themselves and whose actions and motives cannot be perfectly controlled or predicted. Even an attempt to control other people is an implicit acknowledgment of their agency which you wish to destroy. Totalization involves bringing something into your world and gaining complete control over it. Totalization in an NLP system that interacts with people in any way would involve making a model of the person with which the system is interacting and incorporating that model into the algorithm of the system, so that the computer actually interacts with the model, which is part of itself, not with the person.[11] An algorithm is a finite set of instructions such that each decision is binary (i.e. yes or no) and the process terminates in a finite number of steps. All computer programs (except those stuck in an **INFINITE LOOP**) are algorithms. Once we accept that, although totalization of the physical world is desirable and largely possible, totalization of other people is neither possible nor ethical, then we can draw the startling conclusion that an approach to dealing with natural language that truly allows for social interaction could not be a totalized system and therefore could not be algorithmic! For entirely different reasons, a prominent physicist Roger Penrose (1989) has suggested that the brain may operate non-algorithmically on the basis of faster-than-light processes of quantum mechanics.

This brings us back to the image at the beginning of the paper. The oval represented a knot-hole in a tree house. The tree house could stand for a domain-specific approach. Useful work can be and is accomplished in machine translation with a domain-oriented approach based on the assumptions of objectivism. However, before computers will have a chance of performing as well as humans on dynamic general language, they will at least have to avoid the assumptions of objectivism, allow for fundamental ambiguity, handle dynamic metaphor, become much more flexible, and become an agent, recognizing other people as agents (which involves being based on a non-algorithmic approach). The final step of becoming an agent and seeing others as agents that regard it as an agent, thus permitting social interaction, is suggested by the step of looking along the beam of light (in the tree house story told earlier) instead of at it. Until you do it, it is impossible to know what the result will be. Once you do it, a whole new world opens up.

5 Implications

The implications of this philosophical discussion are simple. Machine translation is headed in the right direction. Domain-specific approaches using controlled language should be continued and the controlled languages should be made to conform to all the assumptions of objectivism so far as possible. Dialogue-based machine translation can guide the user into writing in a controlled language. Low-quality indicative translation for information only is unarguable since many find it useful. But further work on fully-automatic high-quality machine translation of unrestricted text is a waste of time and money unless the issues in this paper are carefully addressed. If we ever reach a breakthrough in natural language processing which allows for the handling of dynamic general language, it will not be based on any extension of current techniques in machine translation. The electric light bulb did not result from research and development on the candle (personal communication from Roger Harris). Fully-automatic high-quality machine translation of unrestricted text will be a truly surprising, unpredictable breakthrough and therefore is not expected in the foreseeable future, even though it may come at any time.

We should not complain about the heavy requirements I have imposed on an approach that could handle general language at human levels of performance. In 1984, many of us reviewed the vision of the world presented by George Orwell in his novel Nineteen Eighty-Four and were thankful that things were not as bad as he had predicted, at least outside the Soviet Bloc in the Free World. I had occasional contact with people on the other side of the Iron Curtain and heard horror stories of oppression heaped upon those who dared think on their own in a way that opposed the government then in power. In Orwell's world, the Party had invented Newspeak, a deliberately restricted language in which it was impossible to think thoughts that were not approved by the

Party. Now ten years later, we have seen the Iron Curtain fall. If all language suddenly could be treated like domain-specific language, then a new and far worse Iron Curtain would, in Orwellian fashion, forever keep us from thinking truly new thoughts and we would become machines trapped in the prison of objectivism.

Endnotes

1. Terry Winograd (1987) provides an additional example of the fact that meaning is not always neatly divided up into a literal base meaning and figurative extensions. Suppose one asks the question "Is there any water in the refrigerator?". In the context of a typical American family, this would be a question about whether there is a pitcher in the family refrigerator containing enough cold water (above zero degrees Celsius but probably below ten degrees) to pour into a glass and have a good drink. However, a scientist asking another scientist this same question may be asking whether there is any substance in the laboratory refrigerator containing some H_2O that might interfere with an experiment using microwaves. Which is the literal meaning? If one tries to list all the possible meanings in all conceivable contexts, this is an admission that meaning is indeed dependent on context. If one argues that the literal meaning is the one that is most likely in a normal context, then this is also an admission that meaning is dependent on context. In this case, the context we have called the Utterly Boring World. There really is no meaning that is independent of all context.

2. There are, of course, even variations in the legal system between states in the United States and between England, Wales, and Scotland in Great Britain. Further complications arise when considering US territories and extra-British members of the United Kingdom, such as Northern Ireland, the Isle of Man and the Channel Islands.

3. John Hutchins commented on the shift from general machine translation systems to domain-specific systems at the 1994 Cranfield conference. At that same conference, Peter Wheeler, who in the past ten years has gone from working at the European Commission with Systran, to working for Logos (a machine translation developer), then to being an independent consultant, confirmed the accuracy of the remarks made by Hutchins.

4. At the first conference of the Association for Machine Translation in the Americas, a member society of the International Association for Machine Translation, held in Columbia, Maryland, October 6-8, 1994, a panel discussion treated the topic of the future of machine translation. Several panel members expressed their belief that current systems would gradually be extended to handle general language.

5. At the 1994 Cranfield conference, I took a straw poll during a debate on the limits of machine translation in which professionals from all over the world were participants. About ten percent of the participants indicated that they take the strong-AI position.

6. An outrageous observation at this point would be that there seems to be something about people whose names end in 'sky' (pronounced 'skee') that leads them off the deep end.

7. One of the best-known examples of prototype effects is that a Robin and a Penguin are both birds but a Robin is a better example of a bird.

8. Wilks made this point at the Cranfield conference which has been mentioned several times in these endnotes.

9. I learned how chagrined Weizenbaum was when I heard him give a lecture on the topic in the 1970s.

10. This statement was made in the film "Manufacturing Consent", a documentary on the life of Chomsky which has been shown on university campuses and art film theaters around the country.

11. This suggests another way to detect understanding. Ask someone to make friends with the computer program. Have them ask the computer for advice and try to determine whether the computer program really cares about the person or is just “going through the motions”. This would test both flexibility and interiority.

1. Terry Winograd (1987) provides an additional example of the fact that meaning is not always neatly divided up into a literal base meaning and figurative extensions. Suppose one asks the question “Is there any water in the refrigerator?”. In the context of a typical American family, this would be a question about whether there is a pitcher in the family refrigerator containing enough cold water (above zero degrees Celsius but probably below ten degrees) to pour into a glass and have a good drink. However, a scientist asking another scientist this same question may be asking whether there is any substance in the laboratory refrigerator containing some H2O that might interfere with an experiment using microwaves. Which is the literal meaning? If one tries to list all the possible meanings in all conceivable contexts, this is an admission that meaning is indeed dependent on context. If one argues that the literal meaning is the one that is most likely in a normal context, then this is also an admission that meaning is dependent on context¾in this case, the context we have called the Utterly Boring World. There really is no meaning that is independent of all context.

2. There are, of course, even variations in the legal system between states in the United States and between England, Wales, and Scotland in Great Britain. Further complications arise when considering US territories and extra-British members of the United Kingdom, such as Northern Ireland, the Isle of Mann, and the Channel Islands.

3. John Hutchins commented on the shift from general machine translation systems to domain-specific systems at the 1994 Cranfield conference. At that same conference, Peter Wheeler¾who in the past ten years has gone from working at the European Commission with Systran, to working for Logos (a machine translation developer), to being an independent consultant¾confirmed the accuracy of the remarks made by Hutchins.

4. At the first conference of the Association for Machine Translation in the Americas, a member society of the International Association for Machine Translation, held in Columbia, Maryland, October 6-8, 1994, a panel discussion treated the topic of the future of machine translation. Several panel members expressed their belief that current systems would gradually be extended to handle general language.

5. At the 1994 Cranfield conference, I took a straw poll during a debate on the limits of machine translation in which professionals from all over the world were participants. About ten percent of the participants indicated that they take the strong-AI position.

6. A outrageous observation at this point would be that there seems to be something about people whose names end in ‘sky’ (pronounced ‘skee’) that leads them off the deep end.

7. One of the best-known examples of prototype effects is that a Robin and a Penguin are both birds but a Robin is a better example of a bird.

8. Wilks made this point in an unpublished paper presented at the Cranfield conference which has been mentioned several times in these endnotes.

9. I learned how chagrined Weizenbaum when I heard him give a lecture on the topic in the 1970s.

10. This statement was made in the film "Manufacturing Consent", a documentary on the life of Chomsky which has been shown on university campuses and art film theaters around the country.

11. This suggests another way to detect understanding. Ask someone to make friends with the computer program. Have them ask the computer for advice and try to determine whether the computer program really cares about the person or is just “going through the motions”. This would test both flexibility and interiority.