LANGUAGE AND PROBLEMS OF KNOWLEDGE: BRIDGING NATURAL AND ARTIFICIAL INTELLIGENCE

Rosli Omar
Jabatan Kejuruteraan Elektrik
Universiti Malaya

Introduction

In society today, computers are very much part of our lives. They seem to be capable of doing so many "clever" things that we sometimes think they are intelligent. In certain popular books, authors talk about computers which will soon be more intelligent than man. They are actually dumb machines. Any perceived intelligence is due to their programmer A programmer in programming a computer to do a certain task will lay out all the steps necessary to perform the task. These steps constitute an algorithm. All the possible situations that can occur will be thought out by the programmer so that when the computer is faced with any one of them, the necessary step to take has already been specified.

This method of specifying a task is possible where there are clearly defined goals, constraints, and where all possible situations can be known in advance. But in the real world, the world we live in, it is not possible to determine beforehand all possibilities that can occur Let us assume that we wish to programme a robot, with a computer as its brain, to go to market. A tree may fall on its path; the road may be closed and a detour provided, a robber may snatch away its basket; the market it usually goes to may be closed. The possiblities are endless. It is just not possible to think ahead about all the foreseeable circumstances. Thus the robot must have intelligence to overcome these unthought-of situations. This is because the essence of intelligence according to Winograd (1987: 98), is the ability "to act appropriately when there is no simple pre-definition of the problem or the space of states in which to search for a solution".

The discipline which is involved in trying to make computers more intelligent is called Artificial Intelligence, popularly known as AI. According to Margaret Boden of Sussex University, "the least tendentious definition is Marvin Minsky's, 'Artificial intelligence is the science of making machines do things that would require inteligence if done by men'" (Boden, 1987 4). Some aspects of intelligence studied by AI are vision (how we understand what we see), reasoning, knowledge representation (how knowledge is stored and retrieved) and natural language processing and understanding. The term "natural" language is used here to distinguish it from the formal languages of computer science such as Fortran or Cobol. AI is multi-disciplinary is character and involves researches from computer science, mathematics and logic, psychology and philosophy of mind, neuroscience and linguistics among others. Understanding intelligence is no simple matter. Hence, the need for cooperation from various diciplines. Since human intelligence is the best intelligence (best in the manner defined above for intelligence), AI researchers seek to understand this and to simulate it in computers.

There are two approaches to AI according to Edward Feigenbaum (1985) of Stanford University. The first approach is through the study of the human mind, how human intelligence is achieved. From the insight gained, it is hoped that a model of intelligence might be simulated in the computer.

This approach relies heavily on the psychology and philosophy of mind and generally constitutes the *cognitive science approach* to AI. The second is what Feigenbaum calls the *engineering approach*. Here AI tries to produce programs that can solve problems intelligently. It does not matter if the way these programs solve them is not the way we would solve them. In these two approaches we can see that the computer is just a tool to implement our ideas.

What are the reasons for making computers understand language? For one, it would be easier to communicate with them using natural language rather than formal languages. This would bring about economic benefits - a major motivation in itself. The ability to communicate with computers using language would eliminate one sense of the "alienation of man from technology", especially high technology which one does not understand. Another reason of greater interest is that, language is at the root of intelligence. Whereas other attributes of intelligence, such as vision, are present in other creatures, the ability to speak and understand language is the sole prerogative of human intelligence, the ability of certain primates to understand a very limited number of words notwithstanding.

It is our ability to use language that has given us our present civilisation. According to King (1990), the genetic differences between man and the chimpanzee, the nearest to man in terms of intelligence, is less than one percent. But the difference in terms of the civilisation produced by the two is so vast. This is due to our language capability. The ability to pass knowledge from one person to another and from one generation to the next through written texts enormously speeds up civilisation [1]. It would be interesting to understand this unique language feature of ours by simulating it in the computer. Of course this is only one way of trying to understand the language phenomenon. Another way originates from the humanistic school. "understanding humanity is best done by studying humans". This is, of course, a very valid approach. The task of understanding the language phenomenon is so very difficult that any help from any quarter ought to be welcomed.

Language is at the root of intelligence. But to understand language, intelligence is required. Is this a chicken and egg situation? As will be elaborated, intelligence requires knowledge of the world we live in. Since language cannot be separated from intelligence, some general aspects of intelligence will have to be discussed in order to understand our language capability.

Natural Language Processing and Understanding

In the computer, the first stage to understanding language starts with structural analysis. The structure of a sentence is produced through the parsing process, which results in a parse tree. Parsing follows the grammar of a given language in building the parse tree. Grammar can be considered as knowledge about the language structure that needs to be known before any attempt is made to understand a particular language. Thus in English, a sentence S is made of a "noun phrase" NP and a "verb phrase" VP.

The NP and the VP can be broken down further into their respective components. This can be written as.

GRAMMAR			DICTIONARY
S	—>	NP VP	
NP	—>	proper-noun	Bill, Mary
NP	—>	pronoun	She, he
NP	—>	determiner NP2	
NP	—>	NP2	
NP2	—>	noun	roses
NP2	—>	adjective NP2	
NP2	—>	NP2 PP	
PP	—>	preposition NP	
VP	—>	verb	gave, saw
VP	—>	verb NP	
VP	—>	VP PP	
determiner			the that
adjective			red, beautiful
preposition			to, with

(Grammar adapted from [Winograd, 1983])

The dictionary of the computer stores all the words needed as its vocabulary. Words not in the dictionary, needless to say, will not be understood by the computer. For example, parsing the sentence

"He gave the beautiful red roses to Mary" will result in the parse tree shown in figure 1.

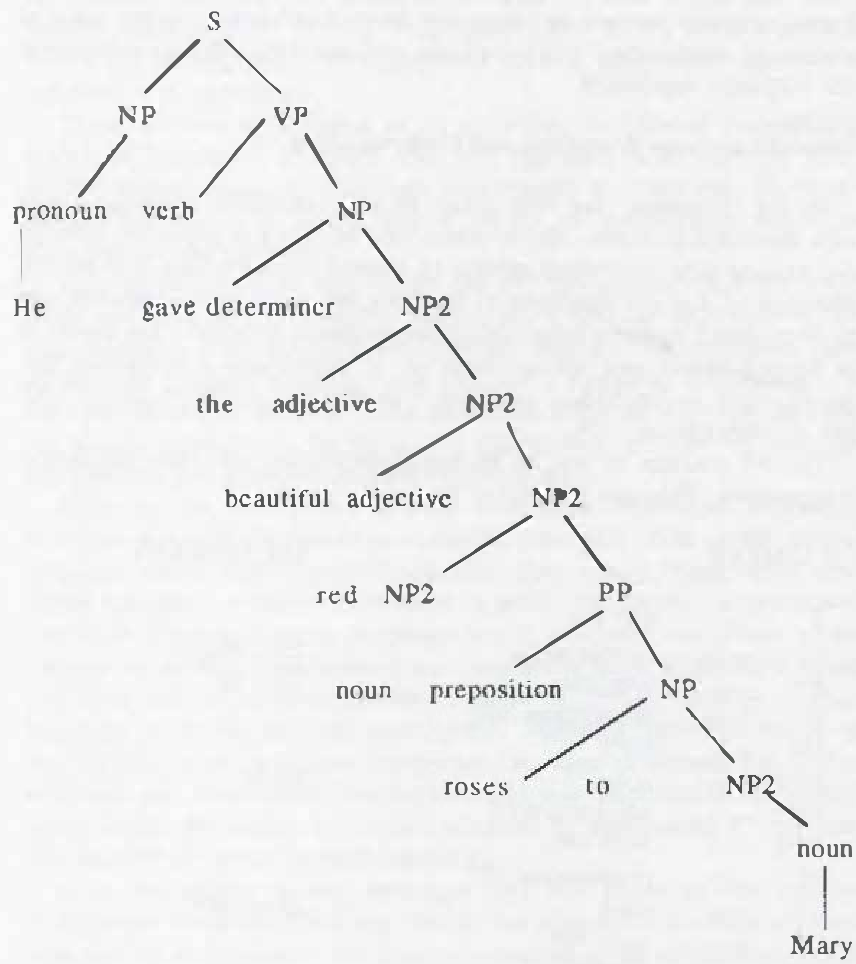

Figure I. The parse tree for the sentence. ...."He gave the beautiful red roses to Mary"

But syntactic analysis by itself is not enough to determine the meaning of a sentence. Consider the sentence: *I read an article about the Vietnam War in the newspaper*

Here, syntactic analysis has two choices:

1) "the Vietnam War in the newspaper" - war occurs in the newspaper

2) the prepositional phrase "in the newspaper" is modifier for "read" i.e. "read in the newspaper" [Charniak, 1985]

Only semantics can determine that the second meaning is the correct one. One of the classic AI programs to demonstrate language understanding through the use of semantics (and also that language understanding requires knowledge of the topic being discussed, apart from grammar knowledge) is Winograd's SHRDLU, a fragment of which is shown in figure 2 [Winograd, 1972]

Figure 2

1. Pick up a big red block. (Winograd's input)
   OK. (Computer response)
   The system answers "OK" when it carries out a command. In order to pick up the red block, it had to clear it off by finding a space for the green one and moving the green one away. The above figure shows the initial state of the scene. (Winograd's comments)

2. Grasp the pyramid.
   I DON'T UNDERSTAND WHICH PYRAMID YOU MEAN
   There are three pyramids in the scene, so the program does nothing.

3. find a block which is taller then the one you are holding and put it into the box.
   The system has to figure out what is meant by "one" and "it," as well as using deduction to decide which block it is holding
   BY "IT," I ASSUME YOU MEAN THE BLOCK WHICH IS TALLER
   THAN THE ONE I AM HOLDING
   Printed out to keep us informed.
   OK.

But even sementics will be unsure in this situation:

.....I read an article about the war in the newspaper

Here, was may mean a war between the editor and the reporters, in whcih case the war is happening in the newspaper

Thus, actual world knowledge is required in order to understand the sentence. It is the problem of integrating world knowledge in understanding language that is the main hurdle for computers in understanding language. Why this is so will be discussed below

Language Understanding and World Knowledge

AI's approach to language understanding and the cognitive science approach in general is largely based on the reductionist approach. The reductionist dictum says that "the meaning of the whole is the sum of the meaning of its parts" Thus to understand a phenomenon, the phenomenon is broken into its constituent parts. The total understanding of all the parts will be the understanding of the whole phenomenon. The basic assumption in the reductionist approach is that there is no cross-interaction between the parts. This approach has been successfully applied in the physical sciences where the cross-interaction is minimal (But note its failure in quantum mechanics - one of the most, if not the most, successful branch of modern physics) Reductionism's belief in its dictum springs from its belief that the world is objective; that knowledge is composed of context-free data, i.e. there is only one way of interpreting a phenomenon. In the physical world, this assumption is perhaps more valid than in the social world, the world of language and human interaction

In language understanding, the reductionist dictum means that the meaning of a sentence is merely the sum of the meaning of all the words in the sentence. Just the opposite is the holistic approach which says "The whole is greater than the sum of its parts", i.e. meaning of a sentence is more than just the meaning of its parts. To show that a sentence is more than just the sum of its parts, consider this fragment of conversation between two children (Papert and Minsky's example in Dreyfus, 1985):

Janet: "That isn't a very good ball you have. Give it to me and I'll give you my lollipop."

To understand this, a lot of world knowledge and concepts are required:

Time, space, words, thoughts, talking (explaining, ordering, persuading, pretending), social relations (giving, buying, bargaining, begging, stealing), playing (real and unreal, pretending), eating (how does one compare the values of food with the values of toys?), owning (belongs to, master of), living (girl, awake, plays), intention (want, plot, goal), emotion (moods, dispositions). Those items in parenthesis are concepts in their own right and need to be further elaborated. The list seems endless, although Papert and Minsky think it is not, merely a large one [2] Berkeley's Dreyfus does not think so. He is especially sceptical of the approach of studying each "microworld" For example, the above microworld of "play", in relative isolation and then attaching it to other microworlds, that is, the reductionist approach

In the social world where language discourse takes place, every aspect of living interacts with others Language meaning then cannot be analysed in isolation that is, objectively, but rather, relative to a situation Winograd's (1987) simple example is very effective in showing this:

"Is there water in the fridge?"

Each one of the words in this sentence seems to have an objective meaning and the sentence itself is straightforward. But even this simple sentence cannot escape different interpretations. To a chemist who wants to keep his chemicals absolutely water free, water can mean H_2O molecules. But to a jogger, it can mean any drinking fluid. For a computer which cannot differentiate between the various situations, trying to find the correct interpretation will be a very difficult, if not impossible, task.

Knowledge Representation, Tacit Knowledge and Living- in-the-World

Knowledge can be classified as declarative and procedural knowledge Declarative knowledge is "knowing what or knowing that" - knowledge about facts in the world such as that the world is a sphere, two and two is four etc., whereas procedural knowledge is "knowing how", knowledge on the inter-relatedness of various types of declarative knowledge; knowledge of how to do things (Winograd, 1985) Since both types of knowledge of the world are necessary in understanding language and for general intelligence, it must be represented and made available, in the computer This will enable syntactic and semantic analysis to interact with world knowledge in interpreting sentences. Knowledge representation is at the heart of much AI research since intelligence is deemed to be the result of knowledge manipulation by the computer (Smith, 1985) This is consonant with the cognitivist idea that knowledge represented in the brain gives rise to our intelligence (Dennet, 1988). A major critic of this idea is Hubert Dreyfus of Berkely (1979, 1985, 1988)

How is world knowlege represented in the computer? There are various means: formal logic, semantic network, frames, rules etc. The knowledge is represented relatively context-free and will form the knowledge base of the computer The various means of knowledge representation is done by what is called the process of abstraction

[Brooks, 1986] The abstraction process abstracts only those facts that are considered to be relevant and pertinent to the task at hand. Abstracting only facts that are relevant is necessary because world knowledge is so vast and probably limitless. Furthermore, in the beginning, AI thought this was enough for understanding. But to Brooks, an important essence of intelligence is the ability to determine from the myriad facts of a given situation, which ones are relevant. If humans do the abstraction for computers, then the latter will never be intelligent since our ability to select relevant facts is very much associated with learning, which is another aspect of intelligence. Thus, to be truly intelligent, to learn what is relevant and what is not, computers need to do its own abstraction.

Humans also engage in abstraction. But we select relevant facts of a particular situation from a large knowledge base formed from a whole lifetime of living experience in the world (Dreyfus 1979) From this lifetime experience, some of the knowledge forms tacit knowledge (Wittgenstein, 1953; Polanyi 1967), knowledge that is so deep that we cannot even formulate it into words. As Polanyi puts it "We can know more than we can tell" This lifetime experience forms our broad knowledge base, described as an outer-horizon by the philosopher Husserl, a pre-understanding by Winograd, from which a background for understanding language is made possible. Just as explicit knowledge is necessary to understand language, so is tacit knowledge.

Tacit knowledge at the syntactic level, is shown by the fact that we cannot completely explain or provide rules on how we structure a sentence, either for understanding a given sentence or uttering. As to the effect of tacit knowledge on ordinary day-to-day situations, consider the classic example: the concept of "bachelor" (adapted from Dreyfus, 1985). Say that a computer is provided with the definition "an adult human male who has never married" - which seems reasonable enough. But in our everyday use of the word, the above definition is clearly not enough. Note this conversation:

Host to computer: "I'm having a party next weekend. Do you know any nice bachelors I could invite."

Suppose the computer's answer is:

1) Arthur.

(He has lived with Alice for several years now and has a child.)

2) Charlie.

(Charlie is 17, lives at home and is still at school.)

3) George.

(He is 17 and quit school at 13. He is now living on his own, and with a very successful business.)

4) Eli.

(Eli is a homosexual.)

5) Pope John Paul.

(No comment!)

In each case we can see that the definition is not enough to capture the concept of "bachelor" As Winograd, (1987) puts it, 'The question "Is X a bachelor?" cannot be answered without providing the potential answers to "Why do you want to know?" This is what Gadamer means by "A question is behind each statement that first gives it meaning" In each of the above situations, further elaboration has to be provided for the definition. For humans, our abstraction mechanism will determine the relevant facts according to each situation. And these facts are abstracted from the large knowledge base of lifetime experience, a part of which is tacit knowledge. But for computer abstraction, we need to think ahead of all possible situations. But how are we to think of all the possible situations - and to explicitly express it in words? What about someone like Ali, who is allowed by his religion to have four wives? He now has two and is hoping to have a third one. Is he not a bachelor? It is the tacit knowledge that we have that enables us to understand a word, a concept, a sentence or a discourse in different situations. Johannessen succinctly establishes the relationship between language, which he calls propostional knowledge and tacit knowledge: "Propositional knowledge rests, we can conclude, on a substratum of tacit knowledge - without tacit knowledge there can be no propositional knowledge" (1988: 300)

Guha and Lenat who are trying to build "a large common sense knowledge base spanning human consensus knowledge" (Guha, 1990: p.33) say that most of what we need to know in the world is prescientific. This is knowledge that is too commonsensical, too obvious, to be included in reference books; Example of such knowledge include the following: animals live for a solid interval of time, nothing can be at two places at one time, string can be pulled but not pushed etc. How are we to present all these in the computer, facts that are so obvious that we do not think about them at all unless absolutely necessary (such as when trying to make computers intelligent!). We acquire these facts through our senses by virtue of the fact that we live in the world.

Again, some of the prescientific knowledge is so deep that it forms tacit knowledge.

One aspect which is said to characterise human intelligence is our intuition, the ability to perceive something usually as being "not-quite'right" or "I can't quite put my finger on it, but I think this is right' without being able to fully explain why it is so. Think again of our concept "bachelor" Here I'm not thinking of mystical or religious intuitions. It is my contention (intuition?) that intuition arises from tacit knowledge. This knowledge is too deep in the region of the subconscious for us to be able to express it in words. We just know it. How is intuition to work in computers? For this, computers definitely will need tacit knowledge. But computers by virtue of having explicit knowledge, through human abstraction, will not be able to have this. Even if it is possible (never mind the paradox of tacit knowledge in computers for a while,) how is this to be brought to the surface? This is because any question of bringing knowledge to the surface for a computer requires that the knowledge be explicitly known? How is it that man can know yet be unable to express what he knows or how is it that he knows but is not conscious that he knows? How does the subconscious relate to the conscious? These are some of the questions whose answers are pertinent if computers are to achieve intelligence.

Tacit knowledge results from experience gained by living in the world. Living in the world enables learning to distinguish between words and concepts and how they are used and provides opportunities for practice in their use. Practice makes perfect. This is true not just for language skills but for practical skills as well. To Dreyfus, (1985) bodily skills form a part of our intelligence. And bodily skills are only acquired by living and practising in the world. Consider bicycle riding and art and crafts. Furthermore, we cannot express our skills in words. Since computers do not have bodies and do not live in the world, the fact that these skills cannot be formulated as abstraction for the computer, means that they can never be truly intelligent, at least not like human beings. But having said that, what about robots? They can have bodies for mobility and other purposes, and sensors to interact with the environment. Can they "live in the world", and as such, achieve intelligence in the same way that we do? If so, does this refute Dreyfus?

Interpretation, Context and Expectation

Since language is not objective but relative to a given situation, interpretation is required before the correct meaning is arrived at. To arrive at this interpretation, the context of the particular situation has to be determined. Consider again the examples mentioned earlier of "Is there water in the fridge?", the concept of bachelor; and the Papert-Minsky example. The first example uses context to determine meaning on the basis of different needs. The second uses context to differentiate between various tacit knowledge subtleties and the, third to differentiate between various motivations. The same conversation among adults may show different motivation and thus different meaning

The sentences "to wreck a nice beach" and "to recognise speech" (from Winograd 1987) when spoken, sound very similar It is the context of the situation in which the sentence is uttered that will determine the meaning. For example, if the topic of conversation is about pollution and oil spills, then the former interpretation is logically consistent. However, if it is about language and speech understanding, then the latter interpretation is more likely

Language understanding is not only relative to different situations but also relative to different cultural practices that is, different social situations. History, geography, religion, children's stories, folklore, literature and technological level are some of the factors that will need to be taken into account. For example, to understand modern Hebrew fluently, it is necessary to understand the Bible in Hebrew since many words used have meanings related to biblical connotations (Hofstadter 1980: 377) Furthermore, to Hofstadter (the author of the classic *Godel, Escher, Bach*) language is relative not only to culture but also to subculture. To those in farm areas, the difference between a pickup and a truck is more pronounced than to those living in cities. Here, a knowledge of the cultural and social context is necessary for interpretation. Language practice within a culture forms a world view, a background from which we interpret language meaning. Computers which are not grounded in the world and thus not immersed in culture will not be able to distinguish subtle differences in meaning

Ultimately, we understand language because we are human. We know the meaning of hunger because we have a body which gets hungry And so it is for other bodily attributes. We may nicely put into a thousand words the concept of hunger but these words are just empty symbols to the computer Intrinsic understanding will forever elude it. For Weizenbaum,(1984) to be human is to live in a society with its associated values, objectives and interests. What does it mean for computers to have values, objectives and interests? Speaking of what it is to be human, Weizenbaum writes: "His life is full of risks, but risks he has the courage to accept, because, like the explorer, he learns to trust his own capacities to endure, to overcome. What could it mean to speak of risk, courage, trust, endurance, and overcoming when one speaks of machines?" (1984 280)

Even to understand a simple concept such as a chair, one needs knowledge of what it is to be human It presupposes certain facts about the human body (fatigue, comfort) and a network of other culturally determined equipment (tables, floors) and skills (writing, reading) on the basis of different needs. The second uses context to differentiate between various tacit knowledge subtleties and the, third to differentiate between various motivations. The same conversation among adults may show different motivation and thus different meaning.

The sentences "to wreck a nice beach" and "to recognise speech" (from Winograd 1987) when spoken, sound very similar. It is the context of the situation in which the sentence is uttered that will determine the meaning. For example, if the topic of conversation is about pollution and oil spills, then the former interpretation is logically consistent. However, if it is about language and speech understanding, then the latter interpretation is more likely.

Language understanding is not only relative to different situations but also relative to different cultural practices that is, different social situations. History, geography, religion, children's stories, folklore, literature and technological level are some of the factors that will need to be taken into account. For example, to understand modern Hebrew fluently, it is necessary to understand the Bible in Hebrew since many words used have meanings related to biblical connotations (Hofstadter 1980: 377). Furthermore, to Hofstadter (the author of the classic *Godel, Escher, Bach*) language is relative not only to culture but also to subculture. To those in farm areas, the difference between a pickup and a truck is more pronounced than to those living in cities. Here, a knowledge of the cultural and social context is necessary for interpretation. Language practice within a culture forms a world view, a background from which we interpret language meaning. Computers which are not grounded in the world and thus not immersed in culture will not be able to distinguish subtle differences in meaning.

Ultimately, we understand language because we are human. We know the meaning of hunger because we have a body which gets hungry. And so it is for other bodily attributes. We may nicely put into a thousand words the concept of hunger but these words are just empty symbols to the computer. Intrinsic understanding will forever elude it. For Weizenbaum,(1984) to be human is to live in a society with its associated values, objectives and interests. What does it mean for computers to have values, objectives and interests? Speaking of what it is to be human, Weizenbaum writes: "His life is full of risks, but risks he has the courage to accept, because, like the explorer, he learns to trust his own capacities to endure, to overcome. What could it mean to speak of risk, courage, trust, endurance, and overcoming when one speaks of machines?" (1984: 280).

Even to understand a simple concept such as a chair, one needs knowledge of what it is to be human. It presupposes certain facts about the human body (fatigue, comfort) and a network of other culturally determined equipment (tables, floors) and skills (writing, reading)

There are all sorts of chairs (arm-chairs, dentist's chairs, beanbags). Chairs would not be equipment for us if our knees bend backwards, or if we have no tables as is the case in traditional Japan. Understanding chairs also includes social skills such as being able to sit appropriately (sedately, seductively, naturally, casually, provocatively) at dinners, interviews, concerts, in living rooms, courts and bars. A functional description as "something one can sit on" treated in a context-free manner will not even distinguish conventional chairs from saddles, thrones from toilet seats (Dreyfus, 1985. 83). Thus, when we speak a word, a whole world of related concepts is behind that word. In understanding then, "what is unspoken is as much a part of the meaning as what is spoken" (Winograd, 1987).

Another aspect of human intelligence is the importance of expectation in understanding language. To Husserl, intelligence is not based on passively receiving context-free facts into an already stored data but rather, it is a context-determined, goal- directed search of anticipated facts (Dreyfus, 1985). Consider again the sentences "to wreck a nice heach" and "to recognise speech". Context will determine the topic of conversation, and expectation will ensure that the interpreted meaning is consistent with the context. In a crowded and noisy situation, one may be near to two persons in a conversation. It may seem strange that one cannot make head or tail of what the conversation is all about, at least in the beginning, whereas the two participants seem to carry on the conversation may be easily enough. This is because, initially the context is not clear After a few related words are heard, the context can be determined and expectation can then function - and the conversation may be easily followed.

Expectation is not only necessary in language understanding but also in another aspect of intelligence, that of vision understanding: how to make sense of what we see. In an airport scene, we would expect to see aircraft, hangars and possibly helicopters. Looking at x-ray pictures, the layman cannot see the difference between those of a diseased person and a healthy person. To the expert, knowledge has provided expectation on what aspects to look out for (Chalmers, 1982). For a person who has been blind since birth, gaining vision after so many years of adjusting to a dark world, will be quite disorientating. There is no knowledge of what to expect. To Chalmers, even the physical sciences, which some of its practitioners proudly proclaim is composed of context-free knowledge, cannot escape from context-dependent data. Observations in the physical sciences are guided by theory which colours and provides a background of expectation for interpreting the observed data.

The fact that computers are not grounded in the world would make it difficult for them to determine context. For humans, because we live in the world, we are always in context or always-already-in-a-situation as Dreyfus (1985) puts it. This is a necessary outcome of that living. Living in the world enables us to develop our lifetime knowledge base and tacit knowledge, our outlook, background, pre-understanding, outer-horizon, world-view, for understanding and interpreting language meaning, and for intelligence in general. Living in the world provides opportunities for the learning and practice of language and bodily skills and context-determined expectations. Living in a society with its shared cultural practices gives us our values, interests and objectives that define what it is to be human, something that no mere knowledge representation as objective, context-free data in empty (representation) symbols will ever capture. The involvement of living is not there to provide causal, deep intentional semantics to the symbols.

Analogy, Metaphor and other bits

Not only is word meaning deeply rooted in tacit knowledge, to further compound the problem for computers, word meaning changes with the use of analogy and metaphor. "To see analogically is to see one thing in another, not in the sense that one mistakes one for the other, but that she [3] conceives of the one in terms of the other" (Boden, 134). To Chomsky "language is a habit system, a system of dispositions to behaviour, acquired through training and conditioning. Any innovative aspects of this behaviour is due to 'analogy'" (1989: 137). According to Boden, there are two important creative uses of analogy. The first is the use of the familiar frame to prompt inquiry at developing the novel frame in an economical way. For example, the gas laws were explained in terms of billiard balls. Secondly, creative use of analogy enables one not merely to gather new factual knowledge about the novel phenomenon, but correlatively to understand and explain it by relating it to the concepts already accessible in the familiar frame (Boden, 326). Arthur Koestler talks of "bisociation", the merging of two different matrices or points of view in creating new insights, whether in science, humour or the arts. For example, in jest, Freud described the Christmas season as "the alcoholidays" - from the two matrices of alcohol and holiday (Martin, 1975).

Another important facet of human language and culture is that of poetry. Martin thinks that metaphor is the basis of poetry (1975: 209-210). The use of a word in a different context, as in the case of metaphor, brings with it certain connotations of that word to consciousness - or at least to the fringe of consciousness. For example, the use of the metaphor 'a rose in bloom' to describe the qualities of a loved one. Furthermore, Martin, in defense of poetry against charges (of some logicians, for example, Ayer in his *Language, Truth and Logic*) that it is imprecise, argues that the use of logic can at best, only *model* reality, whereas poetry relives and evokes images of reality

Consider Melville's Moby Dick: "And heaved and heaved, still unrestingly heaved the black sea, as if its vast tides were a conscience" To Martin. (1975:2) "If there is anything especially valuable about poetry, it is a value that belongs to the real world and is expressed in the speech of the unregenerate human animal. poetry's means are linguistic and semantic, and its subject-matter is experience" Complete understanding of language must therefore include this very important aspect of human culture.

To understand language then computers cannot just interpret words literally when they are used analogically or metaphorically Examples include dead metaphors such as "I see your point of view", to use Martin's example. The use of "see" here is accepted practice and thus the metaphor is "dead", that is not seen as novel or strange, and not even noticed that it is not used literally Dead metaphors and idioms may be collected and represented in the computer But words are always used analogically and metaphorically in new situations. Computers must be able to determine new meanings using some theory on how analogy and metaphor work, otherwise humans will always need to update new word meanings for them.

How does one choose the matrices Koestler talks about? Since analogy and metaphor work around certain relevant and common factors between two matrices, how are the relevant and common factors determined? How does the merging, the bisociation work? More needs to be studied before rules on how they work can be formulated. But the problems of tacit knowledge and living in the world again surface. It may not be possible to determine exactly how we do analogies and metaphor are produced. This can only be done through the experience of living

Words can also change their meaning through time as people use them in new ways. For example, the word 'nice' originally, in the thirteenth century, meant foolish. It has been variously used to mean ignorant, silly and wanton How it changes! It would not do to describe a respectable lady as nice when the word could mean wanton. A computer which has no facility to learn and make these changes will be limited in its usefulness if not outright dangerous to use.

New words are continually being coined, either formulated by linguists, experts through consensus or through being accepted by the populace through constant usage. Computers will need to have a mechanism to incorporate such words and their meanings.

Speech Acts, Searle and other bits

Apart from sentences that have a truth value there are others that are meant for actions to take place. An example of the first type "Stephen Hawking is the greatest particle physicist today " An example of the second is "Remove the above false sentence." The second type of sentences was studied by Austin and his student John Searle of Berkely Such sentences are called "speech acts". There are three types of speech acts. The first is *locutionary,* the act of building a sentence which obeys the syntactic and semantic rules, which can be said to be the aim of context-free-grammar The second is *illocutionary,* which is the purpose of the sentence, and thirdly *perlocutionary,* requiring an action to be taken.

The sentence "Would I abandon you?" is structurally interrogative, but might have the illocutionary force of a statement, and the perlocutionary force of reassuring (Ritchie, 1988). To interpret the various acts requires the context of utterance with its associated problems already mentioned.

Some other issues of language that need to be dealt with by the computer are "conversational implicature" and the "cooperative principle" of Grice and the "relevance theory" of Sperber and Wilson. How do we effectively represent these in the computer? How about the tone, timbre and body language of speech? It is said that only ten percent of meaning is derived from the literal meaning The rest comes from the other three factors. Consider the classic example, "I hate you" Coming from a wife to her husband it may mean the end of their road; coming from a mistress to her man, it may mean the very opposite. It all depends on the tone, timbre and body language. Can a computer ever hope to understand all these subtleties, all this deviousness of human language?

Searle and the Chinese Room

Searle in his celebrated article "Minds, Brains and Program" (1980, 1990), argues through his thought experiment "The Chinese Room" that computers by these very nature cannot achieve human intelligence and understanding. His argument runs as follows: Suppose someone or something, let it be Searle, were to be in a room, opaque to the outside world. Suppose that he were given all the rules of grammar of a language which he does not understand, in this case, Chinese, because Searle does not understand Chinese, and also a basketful of Chinese symbols. (Of course, the rules of grammar and how to use the symbols will be in a known language, say, English.) The Chinese symbols will be identified by the rules according to their shapes and no understanding of the symbols is required. Imagine that people outside the room, who understand Chinese, send in statements in Chinese symbols into the room. By manipulating the symbols, Searle can send out the correct response. To the outside world, Searle (or whatever is inside the room) understands Chinese.

But he is just manipulating symbols without any understanding. This is just like the computer· the "rules" are "programs", "people" are programmers and Searle the computer Searle is trying to show that the Turing test for understanding is not enough. The Turing test of understanding is passed if, to people outside the room, the response is indistinguishable, whether it is from a human or otherwise. The test is named after Alan Turing, one of the fathers of AI (Turing, 1950). The Turing test is essentially a behaviouristic test for intelligence, which to Searle is not enough since there is no real understanding by computers. Searle states this as "computer programs are formal (syntactic). Human minds have mental contents (semantic)" Thus, according to Searle, computers, by their very nature, inherently cannot achieve human level understanding

His arguments, published in 1980, have for a decade generated a heated debate, especially within the AI community, the community described as the "artificial intelligentsia" by Weizenbaum.

Conclusion

It does seem that the computer has a real problem on its hand, so to speak, in trying to understand language. To us, langnage is such a natural thing that we do not attribute any intelligence to it. We marvel at robots which can do many mechanical feats, whether in fiction or in reality, but we think nothing of the ability of HAL the robot (from the 2001 Space Odyssey fame) to converse with his human counterparts. But AI researchers, trying to make computers literate, are involved in a never-ending regression in understanding the philosophical, psychological and biological foundations of intelligence, which, in turn, are the foundations of langnage understanding. In trying to nnderstand language, one is led to an understanding of intelligence, which, in turn, leads to an understanding of knowledge, its representation, tacit knowledge (how to represent tacit knowledge - is that not a contradiction in terms?), and a whole philosophical debate between objectivity and relativity/subjectivity (of knowledge, of the world); between holism and reductionism, and questions on the nature of consciousness and the nature of being Until these and other deep issues are better understood, until we understand our own intelligence better, computer intelligence and language understanding will still be at a very rudimentary stage. Some have compared computer intelligence to that of insects. Still, this might be an insult to insects, which at least can negotiate and survive in their environment.

It also seems that the areas which AI finds so difficult to make computers understand occur in our everyday knowledge, our common sense ability to get along well in the world. Remember our market-going robot? Apart from language, this inclndes vision nnderstanding, how we understand what we see. Again, just as in language understanding, it is precisely world knowledge that has to be incorporated in interpreting the image seen. That is the root problem. AI has been relatively successful in formalising the intelligence of experts in the so-called knowledge-based systems, or expert systems, how they go about diagnosing, solving problems. This is so because their knowledge is highly formalisable due to the precision of the knowledge. It is not so of our everyday knowledge. This is precisely the problem, how to represent imprecise knowledge, myriad inter-related facts, facts which we may not even know exist in our mind/brain, tacit knowledge. Thus, there is a paradox of intelligence between men and computers. Computers find it easy to achieve expert intelligence but do poorly on mundane, everyday, common sense intelligence. Following Boden, (1987:478) Minsky's defintion of AI, in the light of these problems, has to be modified: It is "the study of how to build and program machines that can do the sort of things which human minds can do" The study of AI, far from dehumanising man, as some have complained, has in fact, produced a profound respect for the human mind/brain.

The problems mentioned above, tacit and prescientific knowledge, abstraction, holistic and relativistic interpretation, bodily skills, context, expectation, etc. - are taken care of by us because we live in the world. Computers, in order to overcome the infinite regress of knowledge representation, and those other problems, need to "live in the world" They need to have sensors and a body to interact with the environment. They need to have an automatic capability to acquire their own knowledge. A major research area in AI is knowledge acquisition, which is, therefore, prior to knowledge representation. Only in this way will there be a chance to have the large knowledge base of lifetime experience, and the possibility of tacit knowledge. From here, abstraction proper can be attempted. Knowledge acquisition and abstraction require learning of what is deemed as useful and relevant. Machine learning is another major research area. Living in the world will provide the always-already-in-a-situation, the context, necessary for expectation and, consequently, interpretation of sentences.

Nevertheless, will computer understanding be the same as ours? May be not. Computers may have a body - but it is not a body that feels hunger, that feels pain [4] Can our values, objectives and interests be internalised by the computer? Internalisation seems to require consciousness, or more precisely, self-consciousness of one self-on what it is to be human, to be a part of society. Again, self-consciousness seems to be the crucial factor behind the Chinese Room argument against computer intelligence. It provides the causal power, semantics, and the intentions (beliefs, desires etc) to otherwise empty symbols of representation [5]

Still, computer intelligence without deep understanding, is useful as long as what is manifested to the outside world, is no different from our own manifested intelligeucc. Anyhow, it would uot do to have machines which are conscious. A whole new problem area of defining a legal person, of law and responsibility, questions of ethics, of a human-machine relationship, will have to be thought out. But, having said that, without consciousness, how different will computer intelligence be? As T Edelson of Georgetown University puts it, "Can a system be intelligent if it never gives a damn?"

Notes

[1] But this is also the civilisation that produces the Hiroshimas and the Nagasakis.

[2] P.J Hayes (1985) thinks that the number of concepts required to represent common-sense, every day knowledge, is in the order of 10,000 to 100,000. According to Minsky, "a machine will quite critically need to acquire on the order of a hundred thousand elements of knowledge in order to behave with reasonable sensibility in ordinary situations. A million, if properly organised, should be enough for a very great intelligence." (Dennett: 45) Dennett thinks our knowledge is very much larger, but other pieces of knowledge could be generated by mind/brain inferences from the main knowledge body

[3] Boden uses 'she' in place of the usual 'he' to show the stereotypical representation of 'he' as representing both sexes - but even she would not try using 'woman' in place of 'man' to represent' 'mankind(!) To see the powerful effect-largely unconscious-of this stereotyping, see Hofstadter (1987), and Smullyan (1988)

[4] See Dennett's chapter "Why you can't make computers that feel pain"

[5] Paul Churchland (1988), UCSD"s psychologist and J.R. Lucas (1961), an Oxford philosopher, however think that self-consciousness is just a matter of the complexity of the brain Beyond a certain level of complexity, self-consciousness automatically appears

References

Boden, M., (1987) *Artificial Intelligence and Natural Man*, London: MIT Press.

Brooks, R.A., (1986) "Achieving Artifial Intelligence Through Building Robots", MIT AI Memo 899

Chalmers, A.F (1982) *What is This Thing Called Science?*, St. Lucia: University of Queensland Press.

Charniak, E. and D. McDermott (1985) *Introduction to Artificial Intelligence*, Reading, Ma: Addison-Wesley.
Churchland, P.M., (1988) *Matter and Consciousness*, London: Bradford Books.
Dennett, D.C., (1988) *Brainstorms*, London: Bradford Books.
Dreyfus, H.L., (1979) *What Computers Can't Do: Critique of Artificial Reason*, New York: Harper & Row (2nd. Edition).
Dreyfus, H.L., (1985) "From Micro-Worlds to Knowledge Representation: AI at an Impasse" in R.J. Brachman and H.L. Levesque (eds.), *Readings in Knowledge Representation* , Los Altos. Morgan Kaufman.
Dreyfus, H.L., (1988) "The Socratic and Platonic basis of Cognitivism", in *AI & Society, The Journal of Human and Machine Intelligence*, vol. 2 no. 2, Springer International.
Feigenbaum, E., (1985) in E.A. Torrero (ed.) (1985)*Computers, Next-Generation* New York: Spectrum Series, IEEE Press.
Guha, R.V., and D.B Lenat (1990) "Cyc: A Mid-Term Report" in FAI Magazine, Fall., vol. 11 No. 3
Hayes, P.J (1985) "The Second Naive Physics Manifesto" in R.J Brachman and H. J Levesque (eds), *Reading in Knowledge Representation*, Los Altos: Morgan Kaufman.
Hofstadter, D.R. (1980) *Godel, Escher Bach. An Eternal Golden Braid*, Harmondsworth: Penguin Books.
Hofstadter, D.R., (1987) *Metamagical Themas*, London: Penguin Books.
Johannessen, K.S., (1988) "Rule Following and Tacit Knowledge" *in AI & Society, The Journal of Human and Machine intelligence* vol 2 no. 4, Dec., Springer Internationall
King, M., (1990) "Genes of War" by D Noonan in *Discover*, Oct.
Lucan, J.R., (1961) Minds, Machines and Godel:, *Philosophy 36* Also in Andersons "Minds and Machines" "
Martin, G.D., (1975) *Language, Truth and Poetry*, Edinburgh: The University Press.
Polanyi, M., (1967) *The Tacit Dimension*, New York: Doubleday, Anchor
Ritchie, G., (1988) "Natural Language Processing", MSc Notes, Edinburgh University
Searle, J., (1980) "Minds, Brains and Program", in *The Behavioural and Brain Science* 3 Sept.
Searle, J., (1990) "Is the Brain's Mind a Computer Program", *Scientific American*, Jan. Vol 262 No. 1 New York.
Smith, B.C., (1985) "Prologue to 'Reflection and Semantics in a Procedural Language" in R.J Brachman and H. J Levesque (eds), *Readings in Knowledge Representation*, Los Altos: Morgan Kaufman.
Smullyan, R., (1988) *What is the Name of This Book*, Harmondworth: Penguin Books.
Turing, A.M., (1950) "Computing Machinery and Inteligence", in *Mind*, vol. 59, Reprinted in A. Anderson. (ed.)(1964) *Minds and Machines*, Englewoods Cliffs, N.J Prentice-Hall.

Weizenbaum, J., (1984) *Computer Power and Human Reason*, Harmondworth: Penguin Books.

Winograd, T., (1972) *Understanding Natural Language*, New York: Academic Press.

Winograd, T., (1983) *Language as a Cognitive Process*, Reading, Ma: Addison Wesley

Winograd, T., (1985) "Frame Representations and the Declarative/Procedural Controversy" in R.J. Brachman and H. J Levesque (eds), *Readings in Knowledge Representation*, Los Altos: Morgan Kaufman.

Winograd, T., and F. Flores (1987) *Understanding Computers and Cognition* Reading & New York: Addison-Wesley.

Wittgenstein, L., (1953) *Philosophical Investigations*, Oxford: Blackwell.