ON SPEAKING TERMS

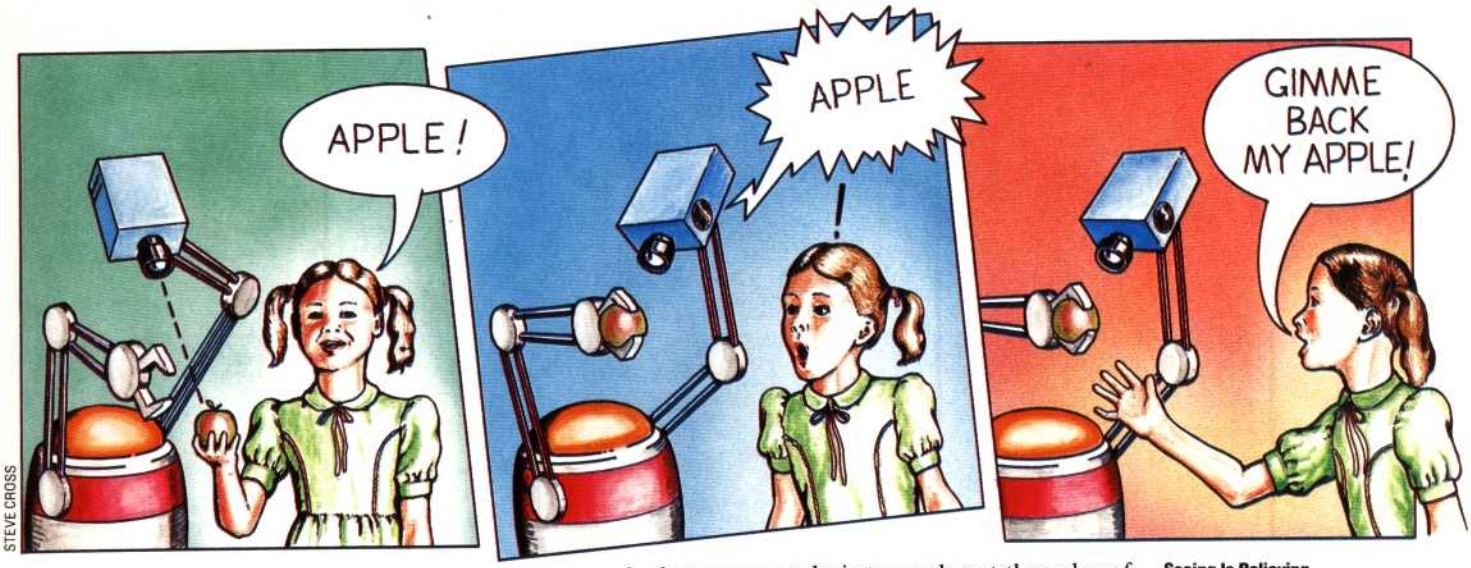

Seeing Is Believing
When a human sees an object, like an apple, and applies a name to it, there is an understanding of the meaning of 'apple'. The robot can visually recognise the object by matching what it sees with an internal image, and can repeat the sound pattern it has stored to go with the apple. But the robot has no understanding that the object is an edible fruit, nor, perhaps more importantly, that the apple actually 'belongs' to the human. This, of course, is something the human understands perfectly

Speech is one of the most difficult tasks for a robot to achieve and the reason for this is because the way in which humans learn to speak is not fully understood. In order to understand some of the problems associated with robot speech, therefore, it is necessary to discuss some of the important theories of language acquisition.

The study of human speech has produced two schools of thought: those who believe that language skills are innate — something that we are born with — and those who believe that language is acquired, or learnt. Those psychologists who argue that language is innate point out that man is the only creature to communicate by language. Those who believe it is acquired cite experiments with animals that have been taught to communicate successfully with humans by sign language.

If people learn speech simply by being exposed to it then it would make sense to look for a method of making robots do the same. After all, it would make life so much easier if a robot could learn the language just by listening to you speak it.

Certain limited attempts have been made to enable a computer to expand its knowledge of grammar by being given extra examples of grammatical sentence structures, while other experiments have tried to allow a robot to learn new words and morphemes (language elements) in any language simply by being shown them. But no system has yet been devised that has succeeded in teaching a robot to learn speech.

So, for all practical purposes, robot language skills are dependent on the assumption that language is innate, that the skills are not learned, and what we must do is to work out the rules of language and embed them permanently into the robot as if the robot had been born with them. In general, this consists of two distinct phases: syntactic analysis and semantic analysis.

Syntactic analysis is concerned with the grammar of what is being said and decodes the surface structure of the message or encodes the message into a grammatical form ready for transmission by the robot. The most common method of doing this is by means of a 'parsing tree' that gradually breaks down, or builds up, a sentence from the various parts of speech. It isn't an easy task — but it is a task that is gradually being tackled with some success.

Semantic analysis is much harder and involves working out the sense of the message (when the robot is listening to you speak); or working out what message needs to be conveyed (when it wants to speak to you). The problem with semantic analysis is that language is not context-free — its meaning depends upon the context in which it is spoken (and this does not apply to the spoken context alone, but to the entire context of the message). This context may encompass knowledge about the state of the world as one speaks, as well as the knowledge that each party has of the other.

This approach has been adopted in experiments conducted by the computer scientist Terry Winograd, who wrote a program that enabled a robot to understand what was said to it and to act on instructions. However, Winograd used a computer simulation of a robot that was only able to operate in a very closely-defined world. In this case, its world consisted of a number of building blocks that it was able to manipulate. Winograd's program, known as SHRDLU, was able to make a