a time. Instead, the program continuously looks out for a given set of circumstances to which one or more of the rules might apply. When this happens, that particular rule 'fires' and is executed, which may in turn lead to other rules being fired.

Semantic nets are a form of graph structures used to represent knowledge, and it is possible to think of them as just a network of relationships between various items of knowledge. The reason they are called semantic nets, rather than simply nets or graphs, is because the individual linkages can have some meaning in their own right. Instead of an arc linking two nodes being merely an arc, it can be an arc that indicates the existence of a special kind of relationship between those two nodes. So, a node labelled 'table' may be linked to a node labelled 'furniture'. In this example, the relationship is that a table is a type of furniture — so we can call the linkage a 'type' linkage.

This kind of knowledge representation can be programmed in a conventional language. You might try it using BASIC and representing the various nodes and linkages by string variables. But, typically, one of the languages of artificial intelligence is more commonly used — such as LISP — because this makes it much easier to express these complex named relationships.

QUESTIONNAIRE

Frames are rather like a blank questionnaire that is specially designed for each type of situation a robot might encounter. The idea is very easy to understand and could be readily programmed in BASIC using simple two-dimensional string arrays — one dimension for the 'question' and another for the 'answers'.

In this approach, the robot is not considered to have complete knowledge of a situation until it has filled in every item on the questionnaire. Only then may it take the appropriate action. A refinement to this method is to have a large selection of different frames available to the robot. One of the tasks of the robot is then to select the appropriate frame for a given situation.

An important aspect of knowledge representation that we have touched on briefly is the part played by the programming languages themselves. LISP, for example, being a list processing language, is particulary appropriate for this kind of approach to the storage and retrieval of data. The choice of language will, therefore, make it easier to represent knowledge in particular ways.

In this instalment we have looked at a number of methods by which robots can be programmed with a fuller understanding of their environment. This is still the subject of considerable research in the computer science departments of universities throughout the world and the key elements are the use of sensors, feedback, learning and knowledge representation. In the next instalment of this series we will take a look at a different approach — computer simulation — in which computers are simply programmed to imitate real-life activities.

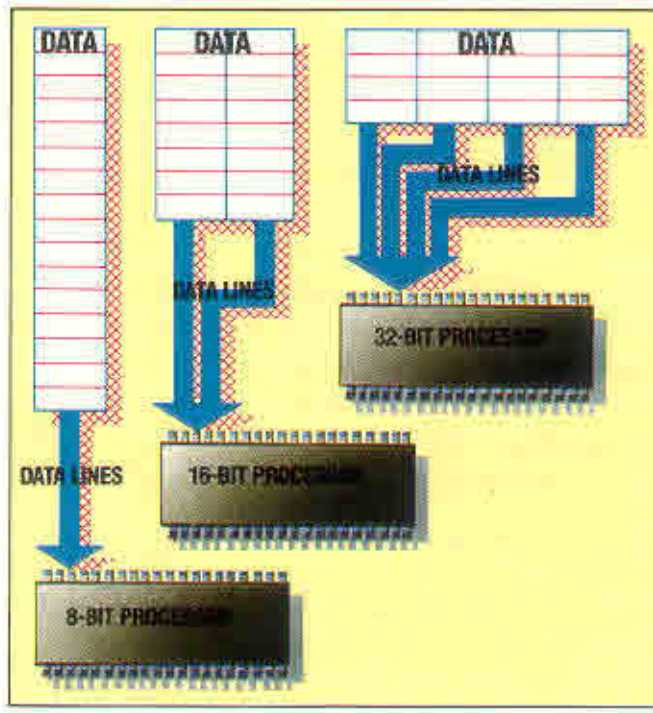

Parallel Lines

Data entering a robot's processing system through sensors must be processed sufficiently quickly to allow an acceptably speedy reaction to the incoming data.

The length of the incoming data queue is determined by the complexity of the algorithm interpreting the data and the processing speed. The volume of data generated by a sophisticated robot may be such that an eight-bit processor such as the Z80 or 6502 cannot cope, causing lengthy data queues to build up. The solution to this problem lies in the use of 16- or 32-bit processors that can process twice or four times as much data simultaneously

Pathways To Knowledge

SEMANTIC NET
Relationships Between Objects

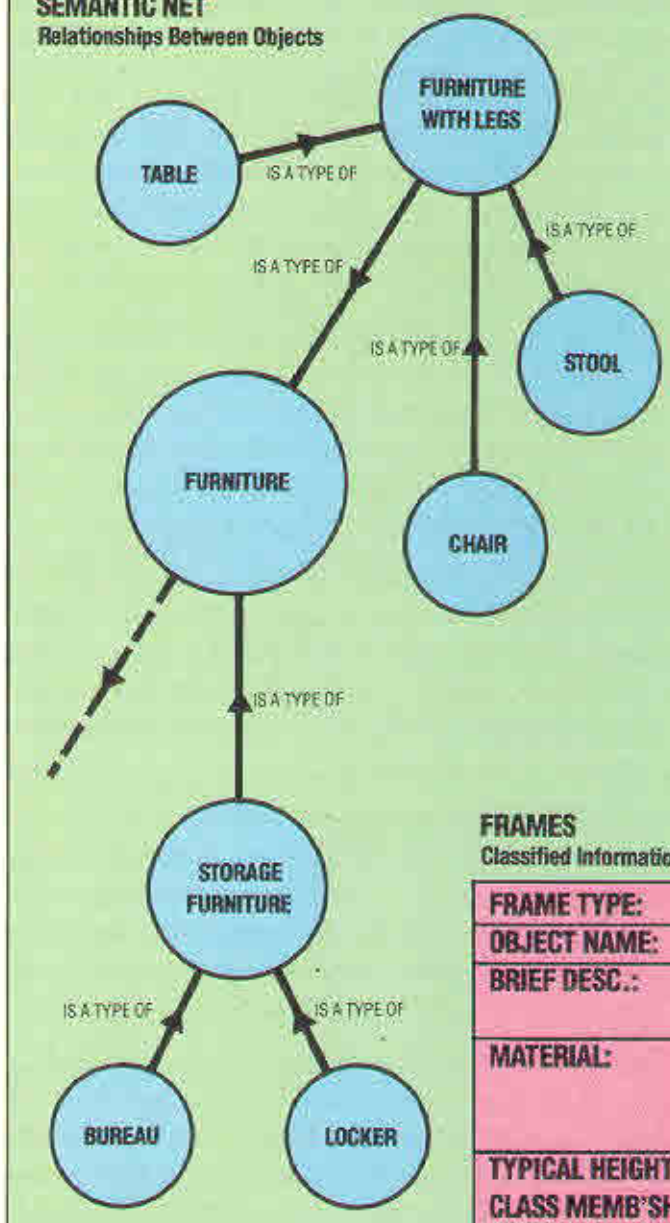

Knowledge can be defined as information in context. Computers are built to store information, but are not specially equipped to handle the relationships among data that turn information into knowledge. Appropriate methods for representing knowledge must, therefore, be invented. Among these are the semantic net, which can be stored as a linked list; frames, which are simply two-dimensional arrays; and production rules, which are lists of information and logical operators.

None of these methods is ideal for representing all knowledge, and combinations of methods are common. Here we use a semantic net to represent detailed knowledge. Notice that these are static representations, with no implications for action — in a robot that would be the outcome of the goal-seeking software's interaction with the knowledge base

FRAMES
Classified Information

FRAME TYPE:	1
OBJECT NAME:	STOOL
BRIEF DESC.:	3 LEGS SEAT
MATERIAL:	WOOD METAL PLASTIC
TYPICAL HEIGHT:	0.5-1m
CLASS MEMB'SHIP:	FURN.WITH LEGS SEATS

PRODUCTION RULES
Boolean Lists

IF STOOL **THEN** (3 LEGS **AND** SEAT) **AND** (WOOD **OR** METAL **OR** PLAST