THE KNOWLEDGE

MIKE BRIGWAN: DW

At Your Service . . .
We have created an imaginary robot butler that must synthesise a variety of sensory inputs to ply its trade. Its greatest problem is the fact that the guests are moving constantly, which means it must update its internal model of the space accordingly

In previous instalments of this series we have discussed in depth the individual senses that contribute towards a robot's 'intelligence'. Here we look at how these senses may be combined to give the robot a fuller understanding of its environment.

While examining the sensors that a robot may use to gain some knowledge of the world in which it moves, we have considered each type of sensory input (sight, sound, touch) as if used in isolation. This is a fair assumption if the robot has a single sensor only but, in practice, the better robots will have several. To understand its environment, the robot must be able to integrate these sensory inputs by using each one as a check upon the others in order to build up a complete internal model of its world.

This is hardly surprising, because this – it seems – is how humans work, too. Our senses do not exist in isolation: we are constantly using the input from one sense as a check on that from another sense and the result is that we build up a very complete picture of our environment. The best example of this concerns the case studies of people who were blind from birth but who have been given their sight through surgery. These patients often surprise their surgeons with the speed at which they can make full use of their vision. This is because blind people have a very good knowledge of the world as a result of being able to touch objects, to move around in the world and to hear descriptions of the environment. Once they can see, therefore, they are able to use this knowledge and apply it to interpreting things they see.

If we are to get the best out of robots we must allow for the interaction of their senses in just the same way as a person's. For example, a robot that is designed to pick up objects may be able to pick them up 'blind', but it could do much better if it had a vision system because it could then locate the objects even if they were slightly misplaced, or at an angle other than the one the arm was programmed to expect. To do this, the robot must build up some kind of internal model of its environment by using the inputs from all its sensors. It must be able to look at the object and recognise it, it must then position its end effector and make the necessary calculations to pick the object up.

The simplest illustration of this internal model is