PHICS
A Graduate Student Conference on Philosophy of Mind, Philosophy of Language, and Cognitive Science
Carleton University, September 28-30, 2001

Levels of Description: A Role for Robots in Cognitive Science Education

Terry Stewart[1] and Robert West[2]
[1]Department of Cognitive Science
[2]Department of Psychology
Carleton University

In this paper, we describe our use of robots to introduce first-year undergraduate students to the idea that understanding cognitive science requires one to view a 'mind' at multiple different levels. The importance of the multi-level description is unique to cognitive science, and so can be difficult to teach in a standard academic setting. Instead, we allow students to interact with a robot that has been pre-programmed with some behaviour. Through this interaction, they develop various descriptions of the robot (in effect, performing cognitive science in a microcosm). These various different descriptions can then be discussed and the benefits of each can be revealed. Furthermore, we found that it becomes evident that no one single description is 'best' or 'most useful'.

Introduction

As part of a first-year university seminar course in cognitive science, we decided to experiment with the use of robots in the classroom. Inspiration for this work came from the growing use of robots in educational situations, where they often provide an excellent hands-on approach to science and engineering. In most of these situations, the educational goal is to learn about robots themselves. For example, the MIT Media Lab has gained a strong reputation for creating remarkable machines. In our case, however, our major goal was for the robots to provide a concrete example of one of the more important tenets of cognitive science: the necessity of having multiple levels of description (sometimes known as the tri-level hypothesis [Dawson, 1998]). In this particular case, we wanted the students to discover for themselves various different ways of characterizing the behaviour of the robots. Once these different levels were uncovered, the students could then be led into a discussion of the usefulness of these levels.

We wanted to focus on this particular concept both because it is fundamental to cognitive science and because it is seldom studied or explicitly explored outside of graduate-level courses. As we discovered through contacts at the American Association for Artificial Intelligence's 2001 Spring Symposium on Robots in Education, very little real-world data is available in this sort of situation. Furthermore, we hoped to provide a clear example of what separates cognitive science from the traditional disciplines. As any teacher of a first-year introductory course in cognitive science can attest, it is difficult to explain cognitive science as something other than 'using psychology, philosophy, and computer science to analyse the mind'. We believe that the instructional scenario described herein can highlight part of what makes cognitive science unique.

Related Educational Robots

It is clear that robots have been used extensively in a field tightly related to cognitive science: artificial intelligence. Not only are robots one of the major products of artificial intelligence research, but they also provide very clear examples of 'emergence'. For example, in [Pfeifer, 1997], a number of robots are placed in an environment filled with randomly scattered blocks. After a while, all the blocks are pushed by the robots into a few piles (or against the walls of the enclosure). When students are asked 'what are the robots doing?', they generally answer that the robots are 'cleaning up'. However, a look at the control program for the robots indicates that they are merely programmed to generally avoid obstacles. They do this imperfectly and, due to their physical interactions within the environment, this leads to lone blocks being pushed into other blocks, forming clusters.

This example demonstrates the idea that the actual high-level behavioural description of a robot may be very different from that expected by a look at its program code. This concept of 'emergence' (high-level behaviours arising from the interaction of low-level rules) is a vital one for both artificial intelligence and cognitive science, and indeed is directly relevant for identifying the necessity and usefulness of having multiple levels of description. However, for a cognitive science audience, we need more than an example of the construction of an emergent behaviour. We need to look at how a science can effectively deal with emergence.

Educational Framework

In order to show how emergence is dealt with in cognitive science, we decided on the following classroom situation. First, we would implement some behaviour in a robot. Then, the students would be allowed to interact with the robot for some time. During this time, the students would be encouraged to try to characterize the behaviour of the robot. Afterwards, we would lead a discussion based on the students' ideas, and try to bring the various (hopefully diverse) descriptions of the robot into a common framework. In this way, we could show that the 'correct' description of the robot is not in the programmed rules, nor is it in the overall behaviour. Both of these descriptions are useful in different situations, and neither is better than the other.

We also wanted to frame this problem by explicitly asking the students to "do cognitive science on the robot's mind." Although significantly simpler, the problem of the students characterizing the robot's mind parallels the problem of cognitive scientists characterizing the human mind. Furthermore, the methods for doing so are fairly similar, including observing the robot in its 'natural' environment, developing hypothetical models of the robot's mind, and performing experiments on the robot to test these hypotheses.

The Trial Run

We used the above methodology as part of a first-year undergraduate seminar course on cognitive science at Carleton University. It was performed near the end of the two-semester course in a class with nineteen students present. These students divided themselves into four groups of four to five students each, and each group was given one of four identical robots. They were told to interact with the robots and to "characterize the robot's behaviour; be able to predict it and describe it". They were also told that after forty-five minutes of such interaction they would be asked to informally report their findings as a group to the class.

The robots in question were built using the Lego Mindstorms Robotics Invention System. The physical form of the robots was identical to the 'Torbot', whose construction details are provided with the Mindstorms kit [LEGO, 2000]. This robot has two motors (one for each tread), and two sensors, which are touch-sensitive 'feelers' in front of the robot. These feelers are designed to be able to detect both walls and cliffs.

The behaviour of the robots was based on Braitenberg's vehicle 2B [Braitenberg, 1984]. In normal operation, both treads are set to run forwards. If either 'feeler' is active, then the opposite tread is reversed until the feeler is no longer active, plus a random amount of time from 0.1 seconds to 1 second. This gives the robot a robust obstacle-avoidance behaviour, and can be easily programmed using the software provided with the kit.

Importantly, the students were not informed of the details of the construction of the robot. This decision was made so as to maintain the parallel with normal cognitive science. When studying actual minds, we do not have the luxury of knowing exactly what the key components are.

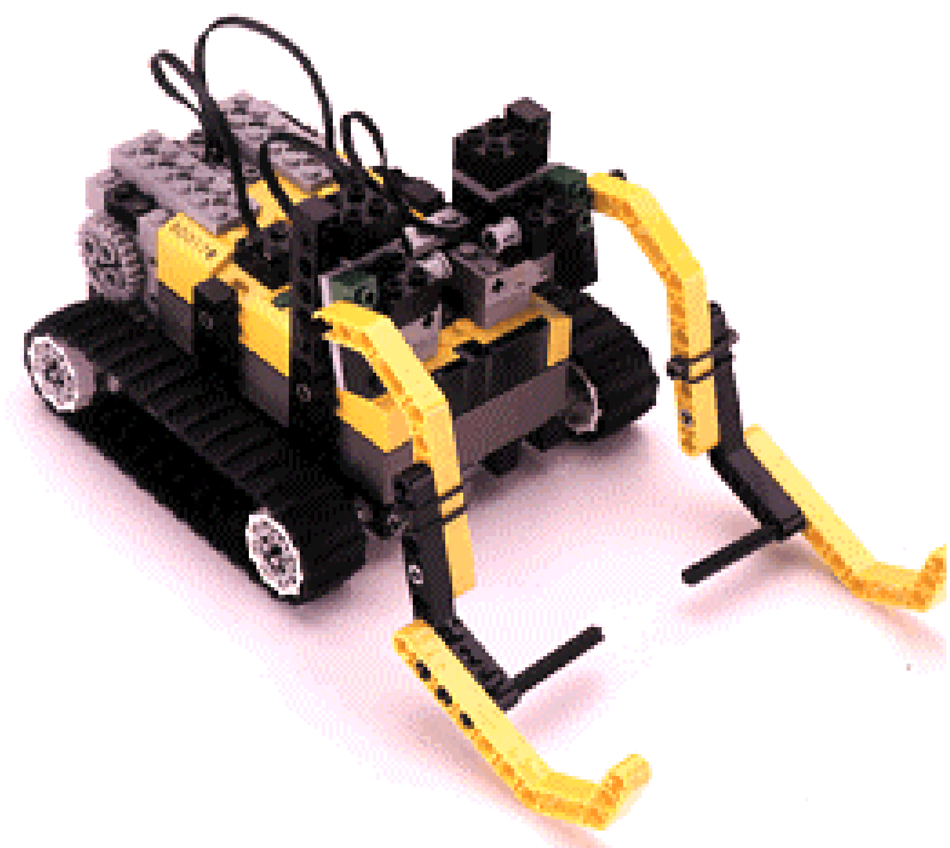

Figure 1: The Mindstorms 'Torbot' robot body.

Results

When the four groups of students were asked to present their findings, there were two distinct types of reports. Two of the groups described the behaviour in terms of low-level components. Their reports could be summarized as "if the left sensor is hit, it backs up and turns right, and vice versa." They correctly identified the basic mechanism the robot used to sense its environment (the touch sensor), and found a simple low-level rule that described the behaviour they observed. The other two groups, however, described its behaviour at a much higher level. The robot was described as "trying to avoid objects" and it was said that it "does not like hitting objects". Furthermore, the robot could experience 'frustration', as evidenced by a few attempts at running into an obstacle before becoming 'frustrated' and trying a new direction.

The rest of the class consisted of an attempt to reconcile these two views. One of our first goals was to counteract the impulse to say that the low-level description was the 'real' explanation of the robot's behaviour. The first way this was done was to reveal the existence of an even lower level of description: the implementation details. In the actual implementation, there is no 'turning left' or 'turning right', there is merely the reversal of signals to the motors. 'Turning' can be seen as being an emergent property of the state of the motors and the physical configuration of the robot. Indeed, we can then describe an even lower level of detail by discussing the compiled program code within the microprocessor itself (although this was not done in the trial run). This situation allowed us to point out the problem of infinite regress: there is no lowest level of description.

The second approach to reconciling these two levels of description was to point out the usefulness of each. Does either description include the other? Can one practically derive one level of description from the other? Here, we were able to make use of the observation about 'frustration'. We posed the question to the class as to whether or not the low-level description would lead you to seeing that this sort of behaviour could happen. The students concluded that it was certainly not immediately obvious that such behaviour would emerge, and so it was useful to have this higher level of description. We were thus able to persuasively demonstrate that both of these levels of description are required for the sort of 'understanding' that cognitive science aims to achieve.

After the discussion, the students indicated that they 'got' what we were trying to get across. Whether or not this is actually true is not established here, but they all seemed to appreciate that if the situation can get this complicated with such a simple creature, then dealing with the human mind is an extremely ambitious task, and not one prone to the sort of cut-and-dry answers that traditional fields of science offer. Furthermore, the robots provided an excellent example for later discussions on the topic of consciousness and the problem of other minds. For example, their comments that the robot was 'trying' to avoid obstacles opens up the issue of attributing mental states: is the robot 'trying' to avoid obstacles in the same sense of the word as when humans 'try' to do something?

Considerations for Future Classes

We identified two issues which should be addressed before we use these robots in another class. To begin with a purely physical issue, the robots require a more robust design. During the class, we were forced to spend some time quickly repairing the robots after physical mishaps such as falling off a table. These repairs were easy to perform, but they did disrupt the experimental situation.

The second issue is the importance of having a 'natural environment' for the robots. Most of the interaction with the robots by the students took place on a table-top, where make-shift walls were formed with hands or textbooks. The lack of a static environment led to behaviour characterizations that were more directly stimulus-response based than would otherwise be expected. When these same robots are placed in a static environment with minimal human interaction (such as in a typical office), the behaviour seems much more exploratory. This sort of description did not arise in the classroom situation.

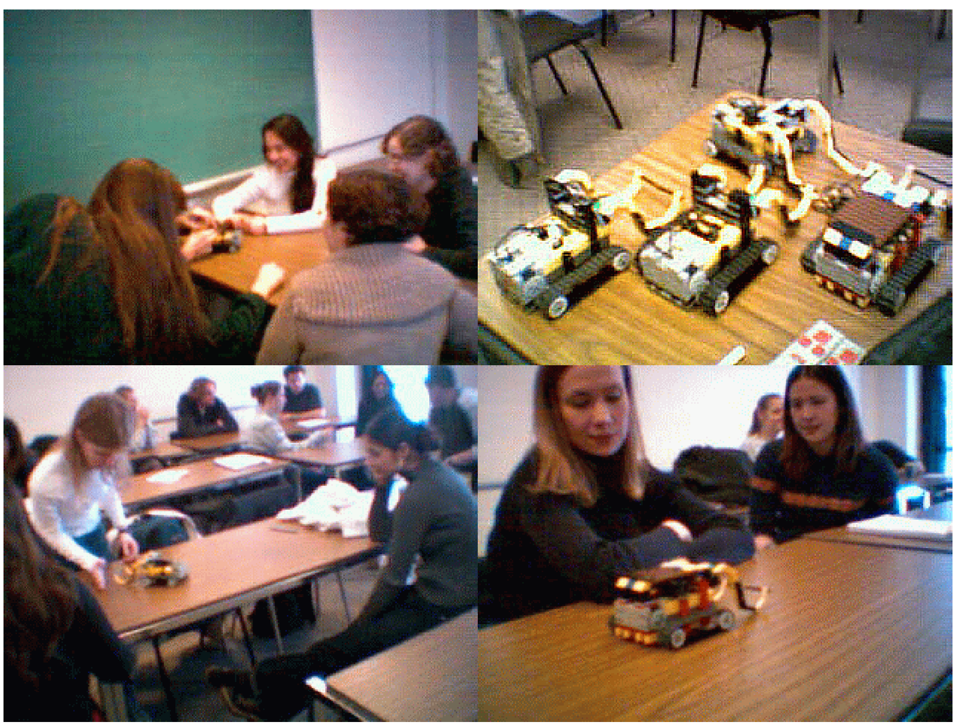

Figure 2: Students interacting with the robots

Conclusion

Although robots are receiving attention for education within certain fields, very little research has been seen within cognitive science education. However, as we have seen in our single trial, they do seem to be effective at introducing traditionally difficult concepts. Aside from providing a welcome change from the typical lecture situation, they can elucidate and provide the basis for discussion on one of the most important questions in cognitive science: What does it mean to understand how a mind works? Approaching this philosophical question with the concrete example of these robots can make this topic more tractable for students and educators.

Bibliography

Braitenberg, V. 1984. *Vehicles: Experiments in Synthetic Psychology*. Cambridge, Mass: MIT Press.

Dawson, M. 1998. *Understanding Cognitive Science*. Malden, Mass: Blackwell Publishers.

The LEGO Group. 2000. *LEGO Mindstorms*. Available at http://mindstorms.lego.com

Pfeifer, R. 1997. "Teaching powerful ideas with autonomous mobile robots." *Journal of Computer Science Education*, 7, 161-186.