Robots in an Interdisciplinary Course in the Liberal Arts

Ellen Walker*   and Lee Braver†

*Computer Science Department
†Philosophy Department
Hiram College
Hiram, OH 44234
{walkerel, braverlj}@hiram.edu

Abstract

This paper describes the use of robots within an interdisciplinary course at a liberal arts college. The course, entitled "The Nature of Intelligence" covered the major paradigms of Artificial Intelligence and their application to robotics, including contemporary objections to those paradigms. The project required students to implement their own reactive robots, and to put those robots into an environment where intelligent behavior could be exhibited.

Introduction

Hiram College is a small residential liberal arts college, typical of many across the United States. One of the strong values of our college is the importance of interdisciplinary experiences, and while robotics is often considered interdisciplinary in and of itself, the interdisciplinarity that is most valued is across "division" lines. To this end, we developed an interdisciplinary course, "The Nature of Intelligence," team-taught by a faculty member from Computer Science with research interests in Artificial Intelligence, and a faculty member from Philosophy with research interests in Philosophy of Mind.

The Course

Our course, INTD 341 was titled "The Nature of Intelligence", and had the following catalog description:

> What is consciousness? What is the difference between an intelligent response and a simple reaction? Can machines think? This course will explore the issues surrounding the topic of human and non-human intelligence, drawing on Computer Science, Robotics, Psychology, and Philosophy of Mind. Students will program robots to perform simple actions and debate whether this constitutes intelligence or not. We will also read various philosophers and psychologists' analyses of intelligence and attempt to apply them to real-world agents. No previous knowledge of computer programming, robots, philosophy, or psychology will be assumed.

During the course, we juxtaposed readings from *Mind Design II* (Haugeland, 1997) and readings from *Introduction to AI Robotics* (Murphy, 1999). Readings based on "good old fashioned AI", such as Newell and Simon's Turing Award talk were related to the hierarchical paradigm of robotics, and later readings were related to the reactive paradigm, which the students implemented on Lego Mindstorms robots. Additional topics that the course covered included biological models of intelligence (neural networks and genetic algorithms) and the idea of emergence, which related back to the reactive paradigm.

For each paradigm, we read original articles promoting the paradigm, contemporary articles criticizing the paradigm, and the relevant sections of Murphy's book relating the paradigm to robotics. In a sense, our course recapitulated the history of artificial intelligence and robotics. In light of the liberal arts nature of the course, much time was allocated to the examination of each system studied for its strengths and weaknesses, as well as critical discussion of the criticisms themselves.

The Robot Assignment

The primary goal of the robot assignment was to give students a hands-on experience with the reactive paradigm, experiencing its benefits and limitations directly. Secondary goals were for students to experience both the challenges and excitement of working with robots in the "real world," and to have the opportunity to critically evaluate their own systems.

The robot assignment required teams of students to build and program robot vehicles that had a set of specified behaviors. The detailed assignment is attached as an appendix, but the behaviors themselves are:

1. FORWARD: When the robot is turned on, it should move forward, going straight enough to cover at least 3 straight road segments without diverting from the road. No sensing is needed for this behavior.

2. FOLLOW_THE_ROAD: This behavior is made up of three parallel SENSE-ACT connections:

Figure 1: Road Following Demonstration

Figure 2: Three Robots on the Course

BLACK→FORWARD, AQUA→RIGHT[1], WHITE→LEFT. Your program will need to distinguish the three colors using values returned by the light sensor.

3. AVOID-OBSTACLE: This behavior is a simple sense and a more complicated action: TOUCHING→(STOP, LEFT, FORWARD). The idea is that the robot will move to the left of the obstacle, then continue its existing plan.

When all behaviors were correctly implemented, robots could be run on a previously unseen course consisting of a thick black "road" surrounding an aqua "lake" in a white "field", and the vehicles would follow the road, passing obstacles (including slower vehicles) as needed. Figure 1 shows one of the robots navigating the course. Additional photos and videos of the robot performances from Spring 2006 are available on our class's website, http://cs.hiram.edu/~walkerel/intd341/. After the demonstrations, students were engaged in a discussion as to the intelligence exhibited by their creations.

Our Experiences

The course was taught in Spring 2006 with 18 students. Of the 18 students in the course, 11 were computer science (CS) majors, 1 was a biology major, and the remaining 6 were undeclared first-year students. There were 13 men and 5 women in the class, with 4 of the 5 women being non-CS majors.

When creating teams, we attempted balance prior experience, which meant that every team had at least one non-CS major. Each team was given a standard Lego Mindstorms kit and the RoboLab software for programming. The use of RoboLab was an attempt to level the playing field between the computer science majors and the other students. Unfortunately, even with this attempt, some teams designated a CS-major "programmer" for the team. On reflection, we plan to try the opposite model (teams of all CS majors or all non-CS majors) when the course is next taught.

We gave the students a design from Baum's book to use as a starting point; this was quite successful, in that every team had a reasonable vehicle and forward motion by the end of the first day. Students were quite creative in modifying the designs to their needs, so eventually there were six distinct vehicles, as can be seen in the pictures.

Aspects of the assignment were modified as the course went on. The original idea of using yellow and white boundaries to simulate the actual colors on the sides of U.S. roads wasn't viable due to the limitations of the Lego sensors. The idea of using road segments taped together succumbed to the abuse of vehicles running over paper boundaries. At the request of nearly the entire class, the final course was printed using our media center's poster printer. Unfortunately, we ran into a problem with that, as well; the "aqua" color from the poster printer was sufficiently different from the same color printed on our lab's printer that most teams had to reprogram at the last minute. The taped together course and the color difference are visible in Figure 2.

Since the robots needed to use "dead reckoning" to decide when to turn back to the road, the size and shape of the object to be passed was needed. We specified a tape dispenser (visible in Figure 2), which was easily available, relatively heavy so it wouldn't move when bumped, and roughly the size of the vehicles that we really wanted to pass. This proved to be a reasonable model, although the course itself was too small for multiple vehicles on opposite sides of the lake to pass at once.

The results of the experiment were that all teams were able to exhibit reasonable road-following behavior, on the taped-together road, if not on the poster road. By the

[1] The original plan was to have YELLOW→RIGHT, but the sensors could not sufficiently distinguish yellow from white.

faculty's expectations, the vehicles were successful in showing (some) intelligence. When the students discussed their results, however, we found that nearly all of the students failed to assign any intelligence at all to their vehicles! The most common explanation was that if the vehicle were intelligent, then it wouldn't have to be reprogrammed for different road colors.

Evaluation and Future Plans

The robot assignment served its purpose in that it gave all the students a taste of some of the problems and issues faced when dealing with the real world of robotics, while being simple enough that all teams could succeed to some level. All students were required to struggle with imperfect sensors and actuators, non-standard test conditions and the other issues that are faced by anyone trying to embed a system in the real world.

Despite these difficulties, students were highly engaged by the assignment -- putting longer hours into their robot than any of their papers. The lab was open and busy every night during the course of the assignment, as well as the 2-3 hours allowed for robot work during class. Students were excited enough about the project that we weren't able to provide enough lab hours to satisfy them. The robot demonstrations and ensuing discussion led to the liveliest class of the semester.

We were surprised that the students did not attribute intelligence to their systems, but wonder if it was at least partially due to their intimate familiarity with the programs. In the future, we can show videos of prior work to prime the discussion before the students know the details of how the robots were programmed, or how fragile those programs are.

As mentioned earlier, each team designated a CS major their "chief programmer", and on some teams the division of labor was so strong that the non-majors on the team never touched the computer. We plan to force non-CS majors to program next time by the simple expedient (if possible) of using segregated teams. As the robot project was graded based on the report, and not on the specific success of the robot, this will not put those teams at a disadvantage. If we continue to use Lego robots, we will continue to use a visual programming option, as we believe it is less intimidating to the students with no prior programming experience. If we can get access to more powerful robots, however, we would like to extend this assignment to a system that can show a stronger emergent intelligence; perhaps even adding a follow-on assignment where different robots communicate and cooperate.

References

Haugeland, J., ed. 1997. *Mind Design II*. Cambridge, MA: MIT Press.

Murphy, R. 2000. *An Introduction to AI Robotics*. Cambridge, MA: MIT Press.