Teaching the Foundations in AI:
Mobile Robots and Symbolic Victories

David A. Shamma
dshamma@ai.uwf.edu
Institute for Human and Machine Cognition
The University of West Florida
11000 University Parkway
Pensacola, Florida 32514
1.850.474.3304

Carl W. Turner
carl_turner@chiinc.com
CHI Systems, Inc.
Gwynedd Office Park
716 N. Bethlehem Pike, Suite 300
Gwynedd, Pennsylvania 19002
1.215.542.1400

Abstract

Introductory AI courses present students with difficult theoretical issues related to the intelligence and behavior of agents operating in the world. We believe that mobile robot projects, when integrated with readings and discussions, help students bridge the gap between abstract AI theory and implementation. In this article we describe an introductory AI class designed around readings, discussion, and student-led mobile robot projects. We then review the outcome of a recent robot project. Finally, we report on the results of surveys taken during the semester to evaluate the success of the course in forming student opinions towards intelligence and agency.

Introduction to AI

One of the major complications with an introductory AI class is covering the foundational aspects; the lectures and readings exemplify the classical theoretical issues. However there are a number of current research issues that are not well addressed by the classical texts, e.g., autonomous agents acting within a dynamic environment. To address this problem, we have designed our Introduction to AI class that supplements readings on reasoning, planning, and action with a root project that asks the students to build intelligent artifacts to operate in the real world. In doing so, students are required to identify and carefully consider many of their underlying assumptions about biological systems, intelligence, and the possibility of intelligent machines.

Robot Projects Help Students Understand Difficult AI Concepts

Unlike many introductory AI courses, most of which focus solely on theory, Introduction to AI includes a substantial applied component: a semester-long robotics project is the central feature of the course. We believe that the design and implementation of autonomous robots provides experience with hard problems in artificial intelligence, experience that traditional classroom lectures and programming assignments cannot match. The projects are carefully designed to direct the student toward confronting one or more important topics in AI during the course of the semester-long project. The emphasis is on building a smart machine that can exploit and interact with a given environment. At semester end, the course culminates in a competition between the student teams.

Surveying Opinions on AI

Much has been written on how best to present course material to students in computer science and AI. However, there is little research on how a given teaching strategy actually changes students' understanding of difficult concepts in computer science and AI. As a part of last year's course, we surveyed students' attitudes about AI before and after the semester in an effort to determine the effectiveness of our instructional approach.

Taking a Stand: The PSS Hypothesis

We begin by stating our position at the outset: that machines are physical symbol systems, capable of intelligent thought and behavior (Newell 1980). Many of the objections to this hypothesis are made on a very abstract level, and serve — we believe — to confuse rather than enlighten. The aim of the robot project is to force the dialogue among students and teachers into the real world: to talk about things as they exist rather than as the way they should be.

By the physical symbol system (PSS) hypothesis, systems that show general intelligence demonstrate a number of defining characteristics, listed below (Newell 1980). Among other characteristics, intelligent agents a) can operate in real time, b) show effective adaptive behavior, c) store vast amounts of knowledge, d) are robust, e) have language, and so on. This is a rather different approach to understanding intelligence than the commonly-held view that humans must serve as the baseline for any definition of intelligence. Each robot project is designed such that students robot demonstrate one or more of these characteristics. Of course, the robots satisfy some of these constraints, and don't satisfy others. In order to compete successfully, the robots must demonstrate some degree of universality, operate in real time, and effectively manipulate symbols. They acquire limited capabilities and evolve over the course of the semester, with students as the engine of evolution. On the other hand, they have no language and cannot store vast amounts of knowledge. Just how robust they are varies widely from one robot to the next, and is sometimes the subject of debate.

Are the robots "intelligent?" The answer: in some ways, by definition, yes, in others, no. Reference to the PSS as a basis for judgment changes the question from "how much is the robot like a human?" to "how well does the agent fulfill a set of observable criteria?" a crucial distinction when trying to build intelligent agents. It is intended that students come to see AI as less an issue of building artificial humans than understanding and implementing agents that can exist and work in their own way on difficult problems in the real world. Here, students are able to grapple directly with real issues of knowledge, representation, and intelligence by working interactively with their robots.

Course Description

Readings and Discussion

As with most introductory courses, our class includes a lengthy reading list and discussion section. Reading include of a thorough review of computer science theory from an AI perspective, a collection of foundational and philosophical AI papers, current AI work, and theoretical discussions (Turner *et al.* 1996b). The class requires extensive participation from the students. Often times, students will have to present the assigned readings, but discussion with the instructor and other classmates provides the most beneficial interaction.

The 6.270 Kits

The building of any beginning AI robot project is straightforward and requires few integral parts. To start, a small CPU is, of course, needed. There are many benefits of using a weak microprocessor. Students see that a high processor speed is not always an advantage or even a prerequisite for an AI implementation. Programming is also kept to a minimum. Interfacing the CPU should be as trivial a task as possible. Since the emphasis of the class is AI and not engineering. We find that the MIT 6.270 Robot Controller, or the Handy-Board, makes an excellent choice in combination with its Motorola 68HC11 CPU. Analog and digital sensors can be plugged directly into the board; an A/D converter onboard the 68HC11 handles most of the low-level hardware implementation.

We chose to use Interactive C for programming because beginning students tend to be more familiar with a procedural language as opposed to a functional one, such as LISP. Thresholding and interfacing is done through the use of many predefined C libraries. Again, the notion here is that the students should not be troubled with particulars of robot building and programming. It is also for this reason that Lego bricks and gears are supplied for chassis building (Turner *et al.* 1996c). Lego provides an excellent rapid prototyping tool for constructing a robot, yet exhibits robust stability for actual competition. Furthermore Lego does not, at first, intimidate the student.

Monitoring Student Performance

Monitoring each student's performance is the key to success for a robot project in Introduction to AI. Typically two milestones during the semester must be achieved. At each milestone the students must demonstrate their robot's proficiency for some small task in the actual competition. A good first milestone is to have students make a prototype geartrain for the robot, thus giving them a chassis on which to start building their robot. The second milestone usually consists of a demonstration of the perception sensors for the robot, making sure it is on its way towards properly examining its environment. Between milestones, the teaching assistant also makes sure that the students are taking the correct steps towards the next milestone. Often, a student may spend time on a function that does not help the robot's role or its task. The maze position is of little or no value to the robot. Furthermore determining the position is a difficult task. Nevertheless, students would spend considerable effort on this task.

It is also important that students are not easily and quickly overwhelmed by the many eclectic readings. During class lectures, parallels are drawn (by both the professor and the TA) between the the planning phase of the robot and complex reasoning within a dynamic environment. AI touches upon many fields, such as philosophy, mathematics, and cognitive psychology to name a few. Teaching classical theory is simply not enough. Students should be presented with the foundations, as well as current research (Wyatt 1996). Even textbooks can often be misleading due to biased criticism from other sources (Hayes & Ford 1995a). During the course of the semester two tests, a midterm and a final, are given to test the students on the reading material. These tests are mostly essay and focus both on the particulars of different theories and their impact and implications in various scientific fields.

Competition: the Spice of (Mechanical) Life

1996: The Autonomous Pac-Man

Our first graduate AI robot competitions featured tasks which, by their nature, required the robots to modify their environments. However, the environments were, for the most part, static; students were able to implement problem-solving plans with the knowledge that the environment would not change unless done so by the robot itself. For the 1996 competition, the contest organizers decided to make the environment more dynamic and unpredictable while keeping the emphasis on goal-driven behavior, strategy, and planning.

The 1996 competition took place on a large, rectangular, obstacle-filled playing field bounded by wooden walls. Small battery-powered lamps with mercury switches served as targets for the robots; the targets remained illuminated while in an upright position and turned off when knocked down. The competition proceeded in rounds, with two robots on the field during a round. One robot, the "runner," tried to find and knock down the illuminated targets in order to score points. The other robot, the "chaser," tried to discover and tag the runner. If the chaser tagged the runner, the runner lost the points collected during the round. Obviously, the contest rules were inspired by the game "Pac-man." Each robot was equipped with a modulated infrared light emitter, which permitted the robots to see each other. Each robot played a round as runner and chaser against the other robot contestants. In the past, each team conducted their own run and depended upon the static environment. Each robot had to constantly monitor its surroundings. This was not a trivial task, as constant detection of walls and obstacles occupied 20 – 30 percent of the 68HC11's time.

Unfortunately, the size of the project overwhelmed the two robot teams. The rigid guidelines for building and modulating an IR tower on each of the robots consumed the students with electrical engineering factors (Turner, Shamma, & Dobbs 1998). Interestingly, each team was able to implement each of the many required skills their robot needed in order to perform at a high level, but were unable to integrate all of the skills in any given version of a robot.

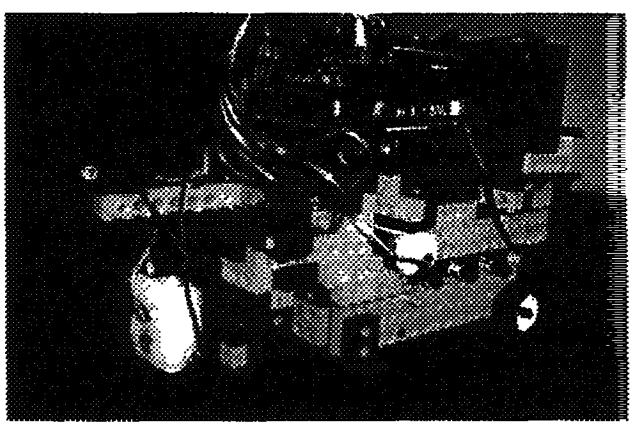

Figure 1: 1997 Robot Contender: Buff.

1997: Mars Mini-Rover

Last year's contest was intended to be a smaller scale version of the Mars Sojourner Rover. The competition involved a single robot in a playing field (see Figure 1). The robot's primary task was to locate and determine several small objects of varying size and color. From the size and color description, the robot would know to know if it was allowed to touch the object for further examination. In addition, the robot had to avoid certain traps in the environment. This competition had an environment that changed during a robot's run and yielded more successful results that the previous year's attempt.

Planning was an important part of the robot's task. This is a point of emphasis for our robot competition. The competition is designed to reward robots that exhibit intelligence within a specified domain. By incorporating planning into the overall function of the robot and an integral part of the competition, a team of students cannot implement an ad hoc design that will do well in the competition.

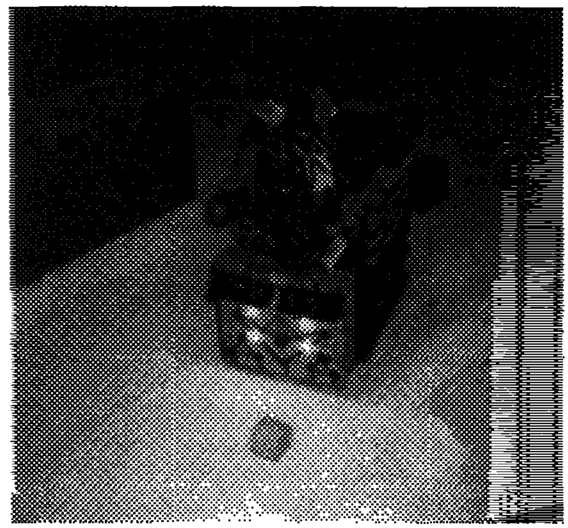

Figure 2: Winney the Brick examining a block.

Several good designs were demonstrated by the students. The Mini Mars Rover that outperformed the others was Winney the Brick (see Figure 2). Winney's success was due in part to the 4"x4"x1" piece of balsa wood used to mount the sensors. The contest organizers had to determine the legality of such an idea. The usage of Lego bricks is a limitation for the students and a feature of the contest, but it was determined that the balsa was being used only to simplify physical engineering so more work and time could be devoted towards AI programming.

Student Attitudes Towards AI

We have claimed in the past that the use of mobile robot projects in the Introduction to AI course has helped students reconsider some of the difficult theoretical problems that are raised when discussing the possibility of creating intelligent machines (Turner *et al.* 1996a) (Turner *et al.* 1996c). This opinion was based on informal feedback from students at the conclusion of the course. However, we have never put this claim to an empirical test. Therefore, we attempted to devise a measure that would demonstrate the effectiveness of the current approach to teaching an introductory AI course.

Many students come to the course with strongly held opinions about AI, opinions which may reflect the way in which AI is commonly portrayed in traditional CS classes and the popular media: Few students in introductory AI classes typically claim to have experience working on AI-related projects. Many of the preconceptions held by students in Introduction to AI reflect the wider opinion of AI, e.g., computers lack flexibility and carry out pre-defined instructions only, consciousness and awareness are central to any definition of intelligence, inference and reasoning are difficult (and thus can be considered intelligent), but movement and perception are simple. The attitudes expressed could be summarized briefly as such: Humans are the only yardstick by which intelligence can be defined and measured. It is against this baseline that we tried to measure changes in students' attitude towards AI by the conclusion of the course.

Surveying Opinions on AI

A questionnaire consisted of three sections. Section 1 consisted of 11 statements that were intended to elicit opinions about current controversies in AI. Students were asked whether they agreed or disagreed with the statements. Responses were made on a scale from 1 (agree strongly) to 7 (disagree strongly) with 4 (neither agree nor disagree) as the midpoint. Section 2 consisted of two parts. Part 1 consisted of a single question that asked simply, "How much intelligence would you attribute to a person who could do the following:" There followed six scenarios, e.g., "Play master-level chess." Responses were made on a scale from 1 (a great deal) to 5 (very little). Part 2 of Section 2 posed the same question and scenarios about "a computer" rather than "a person." Section 3 asked for a short response to the question "What would a machine have to be able to do in order to be considered intelligent?"

Methods Questionnaires were distributed to students prior to the beginning of the first class of the semester (the pre-test) and again at the end of the last class (the post-test). Six students took both questionnaires and provided the data used in the analysis. The primary analysis was on: a) within-subject changes in magnitude between pre- and post-test of opinions in Sections 1 and 2, b) within-subject changes in direction of opinion in Sections 1 and 2, c) differences in attitudes toward human and machine intelligence expressed in Section 2, and d) changes in the content and depth of opinions expressed in Section 3. The findings should be considered highly tentative, owing to the small number of students who participated, the lack of appropriate control groups, and the subjective nature of evaluating answers in Section 3.

Results As expected, students expressed strong opinions towards AI before the course, as measured by the absolute value of the mean difference from 4 (no opinion) in Section 1 during pre-test, 1.73. Post-test, there was no difference in the expressed strength of their opinions, 1.71. However, the direction of students' opinions towards the questions in Section 1 changed substantially on several items, as seen in Table 1.

Question	Pre-Test	Post-Test
Computers can be thought of as formal systems	2	1.5
Consciousness is a pre-requisite for intelligence	5.33	4.67
Computers can do only what they are instructed to do	1.33	6
Artificial intelligence is the enterprise of constructing a machine that can reliably pass the Turing Test	5.5	4.83
Computer architecture will have to become more like a human brain before computers will be able to think	4	5.17
Neural networks are superior to serial computers because they can compute functions that serial computers cannot	3.83	5
"Intelligence is in the eye of the beholder." In other words, there is no objective definition of intelligence	3.33	3.17
Intelligence consists of a collection of separate abilities rather than a single ability	3.33	2.17
Expert systems are so successful that they should no longer be considered a part of artificial intelligence	5.17	5.5
Artificial intelligence is the business of using computation to make machines act more intelligently	2.67	4.17
Something either is intelligent or it isn't: there are no degrees of intelligence	5.83	5.33

Table 1: Mean pre- and post-test responses to questionnaire Section 1. Responses range from 1 (agree strongly) to 7 (disagree strongly).

Pre-test, students substantially agreed with the statement "Computers can do only what they are instructed to do," 1.33. Post-test, they disagreed strongly, 6.0. Pre-test, they expressed little opinion on the statements "computer architecture will have to become more like a human brain before computers will be able to think," 4.0, and "neural networks are superior to serial computers because they can compute functions that serial computers cannot," 3.83. Post-test, students somewhat disagreed with those statements, 5.17 and 5.0 respectively.

Students also agreed strongly pre-test, "AI is the

Question	Pre-Test	Post-Test
Converse in English on topics of general interest	2.83	2.33
Play master-level chess	2.33	1.83
Identify a leaf by touch	2.67	3.8
Drive a car in city traffic	3.33	3.17
Solve differential equations	3.17	2
Diagnose a patient with heart disease	2.33	2.17

Table 2: Mean pre- and post-test reaction to human intelligence: Section 2, Part 1. Responses range from 1 (a great deal) to 5 (very little).

business of using computation to make machines act more intelligently," 2.67. Unexpectedly, the post-test response was 4.17, close to no opinion, and makes one wonder what students think AI is about.

The hypothesis that students attribute intelligence to humans and machines based on different criteria was confirmed by their responses in Section 2. Responses to queries about human intelligence appear in Table 2; responses to machine intelligence appear in Table 3. The mean absolute difference between corresponding questions in Part 1 and Part 2, was 1.0, pre-test. This view did not change as a result of having taken the course, as the post-test difference was 1.1. Students in the pre-test tended to attribute intelligence to humans if they could "play chess" (2.33) and "diagnose patients with heart disease" (2.33) and attribute rather less intelligence to machines for the same behaviors, 3.33 and 2.83, respectively. Of course, it was known to all that a machine had defeated the human chess champion a few months earlier, which may have influenced this response. Responses to the post-test questions were largely the same.

Question	Pre-Test	Post-Test
Converse in English on topics of general interest	1.67	1.17
Play master-level chess	3.33	2.5
Identify a leaf by touch	2.67	2.17
Drive a car in city traffic	2.17	2
Solve differential equations	4	3.17
Diagnose a patient with heart disease	2.83	3

Table 3: Mean pre- and post-test reaction to computer intelligence:Section 2, Part 2. Responses range from 1 (a great deal) to 5 (very little).

The pre-test responses to Section 3 were short, and could be classified mainly in one of two ways. Several answers referred to behaviors which, it was argued during the course of the semester, may have already been demonstrated by intelligent machines. The other class of responses dealt with unobservables, such as "self-awareness" and "intuition." Post-test, students generally wrote at greater length and with more insight into what machines currently can and cannot do and tended to phrase their answers in terms of measurable behaviors, with one explicit reference to the behavior of a robot. Of interest is that two students declared that "passing the Turing Test" would be a sufficient condition for judging a machine to be intelligent, in marked contrast to others' opinion of the Turing Test's validity (Hayes & Ford 1995b).

Summary

We have argued that an AI course which combines robot building projects with readings and discussion can provide students with insights into difficult theoretical issues in a way that traditional AI courses do not. Building machines which must reason, explore, and work in the real world tends to bring arguments over intelligence (broadly construed) into the real world and away from mere philosophizing. The present study suggests that student attitudes towards AI changed significantly in several ways as a result of having taken the course, but it is not possible to ascribe these changes only to the robot project. The course and its associated robot project will continue to evolve in the coming years, as will efforts to assess the effectiveness of the general approach to AI education.

References

Hayes, P. J., and Ford, K. M. 1995a. Intellectual archeology. *SIGART Bulletin* 6:19–21.

Hayes, P. J., and Ford, K. M. 1995b. Turing test considered harmful. In *International Joing Conference on Artificial Intelligence*, 972–977.

Newell, A. 1980. Physical symbol systems. *Cognitive Science* 4:135–183.

Turner, C. W.; Ford, K. M.; Dobbs, S.; Suri, N.; and Hayes, P. J. 1996a. Robots in the classroom. In Stewman, J. H., ed., *Proceedings of the 9th Florida Artificial Intelligence Research Symposium*, 497–500. Florida AI Research Society.

Turner, C. W.; Ford, K. M.; Dobbs, S.; Suri, N.; and Hayes, P. J. 1996b. Robots in the classroom: What smart machines can teach smart students. *Computer Science Education* 7:187–197.

Turner, C. W.; Ford, K. M.; Hayes, P. J.; Shamma, D.; and Manseur, R. 1996c. Robots in education. In *Proceedings of the US-Japan Graduate Student Forum in Robotics*.

Turner, C.; Shamma, D.; and Dobbs, S. 1998. Things that help us think: Mobile robots in a.i. education. *IEEE Transactions on Education*, CD-ROM Supplement.

Wyatt, R. 1996. Design considerations for an AI course for the typical CS student. In Stewman, J. H., ed., *Proceedings of the 9th Florida Artificial Intelligence Research Symposium*, 501 – 504. Florida AI Research Society.