Integrating Robotics Research with Undergraduate Education

Bruce A. Maxwell
Engineering
Swarthmore College
500 College Ave.
Swarthmore, PA 19081
(610)328-8081
maxwell@swarthmore.edu

Lisa A. Meeden
Computer Science
Swarthmore College
500 College Ave.
Swarthmore, PA 19081
(610)328-8565
meeden@cs.swarthmore.edu

Abstract

The field of robotics moves so quickly and encompasses such a wide range of disciplines and applications that education in robotics must be adaptive and incorporate a multidisciplinary approach. The key is to provide an education in the fundamentals--as best we can define them--while maintaining strong connections to current research results. To build these connections we are integrating current robotics research into both the classroom and summer research projects. This kind of integration brings benefits both to us as robotics researchers and the students as future roboticists. Our vehicle for motivating and directing the students during the summers of 1998 and 1999 has been a national robot competition. This goal has provided motivation, given us hard deadlines, and encouraged teamwork among the students. It has proven to be a successful way to balance research and implementation within the time frame of a summer project.

Keywords: robotics education, robot competition, undergraduate education

Integrating Robotics Research with Undergraduate Education

Abstract

The field of robotics moves so quickly and encompasses such a wide range of disciplines and applications that education in robotics must be adaptive and incorporate a multidisciplinary approach. The key is to provide an education in the fundamentals--as best we can define them--while maintaining strong connections to current research results. To build these connections we are integrating current robotics research into both the classroom and summer research projects. This kind of integration brings benefits both to us as robotics researchers and the students as future roboticists. Our vehicle for motivating and directing the students during the summers of 1998 and 1999 has been a national robot competition. This goal has provided motivation, given us hard deadlines, and encouraged teamwork among the students. It has proven to be a successful way to balance research and implementation within the time frame of a summer project.

Keywords: robotics education, robot competition, undergraduate education

1 Introduction

The field of robotics moves quickly and encompasses a wide range of disciplines and applications. Likewise, education in robotics must be adaptive and incorporate a multidisciplinary approach. The key is to provide an education in the fundamentals--as best we can define them--while maintaining strong connections to current research results. Therefore, it is essential to integrate current research into our undergraduate students' education.

At Swarthmore College we take a two-prong approach to undergraduate education, both of which integrate robotics research with educational goals. The first prong is to offer courses that support the field of robotics. Courses in artificial intelligence, computer vision, and robotics have all been offered at Swarthmore in the past two years. Either of the first two courses--AI or vision--serve as a prerequisite for the robotics course.

The second prong is to involve students in ongoing research projects--usually during the summer--as part of their undergraduate experience. Participation in a research project is required of honors students, and encouraged for students who choose not to participate in the honors program. This combination of courses and research provides these students with an excellent background in robotics and related areas.

This paper discusses both our current robotics course--offered spring 2000--and the robot projects we have undertaken in the summer of 1998 and 1999. Our emphasis is on demonstrating that integrating robotics research at the undergraduate level is not only feasible, but also successful in terms of both educational and research outcomes.

2 Undergraduate Robotics Course

We offered an upper-level robotics course for the first time at Swarthmore College in the spring semester of the 1999-2000 school year. The course was cross-listed in Engineering and Computer Science and was team-taught by the authors. About one-third of the 25 students enrolled were Engineering majors and the other two-thirds were Computer Science majors. Students were expected to have a strong background in programming and to have had at least one related upper-level course such as Artificial Intelligence or Computer Vision.

1. **Introduction to robots and robot motion**
   - Basic robot and sensor concepts
   - Maps and evidence grids
   - Configuration space
   - Path planning algorithms
2. **Navigation using behaviors and learning**
   - Subsumption architectures
   - Saphira architecture and fuzzy logic
   - Artificial neural networks
   - Reinforcement learning
   - Evolutionary computation
3. **Robot vision and sensor integration**
   - Real-time vision techniques
   - Edges and segmentation
   - Calibration and stereo
   - Motion detection and tracking
4. **Robot Control Architectures**
   - Multiple-layer designs
   - Cognitive models and concept development
5. **Human-Robot Interface Design**
   - Kinematic analysis
   - Speech recognition
   - Robot personalities and emotions

Figure 1 Major topics in the Robotics course

2.1 Course organization

Robotics is a wide ranging topic that is well served by having multiple instructors. In our case, Maxwell is an Engineer whose research focus is computer vision and Meeden is a Computer Scientist whose research focus is learning. Although our areas of expertise do overlap in some aspects, for the most part they are quite divergent which allowed us to cover more material than either one of us could have done alone.

We divided the topics into five sections of two to three weeks each. While one instructor led the lectures--three 1-hour lectures per week--the other instructor ran the labs--one 3-hour session. Our goal in selecting topics was to give students a grounding in some fundamental areas--geometry, machine learning, sensor analysis, visual sensing, robot software architectures, and robot-human interfaces--while at the same time providing examples of current research that extended and implemented variations of these concepts. A rough outline of the material covered is shown in Figure 1.

This organization loosely followed the structure of the main text chosen for the course, *Artificial Intelligence and Mobile Robots* [1]. This book contains an edited collection of case studies of successful robotic systems which is divided into three main parts: (1) Mapping and Navigation (topics 1-2 above), (2) Vision for Mobile Robots (topic 3 above), and (3) Mobile Robot Architectures (topic 4 above). While a number of the case studies are becoming dated, the concepts they present are appropriate as an introduction to the field.

We supplemented this text with more than 20 research articles. For example, to learn more about robot navigation, students read Latombe on path planning [2], Moravec on evidence grids [3], and Stentz on D* (a real-time replanning algorithm based on A*) [4]. For additional information about robot architectures, students read Sun on bottom-up learning and Cohen on developing human-like concepts in robots [5][6]. The complete list of citations for the supplemental reading materials are listed at http://www.palantir.swarthmore.edu/~maxwell/classes/e28/readings.htm. These supplemental readings were essential to presenting both the history of robotics and the current state of the art.

2.2 Laboratories

Optimally, a robotics class should provide students with frequent hands-on experience with physical robots. Practically, this is not always possible, depending on the class size and equipment availability. When equipment is limited, it is possible to use simulation as an initial test-bed for design and implementation. However, it is important that students have the opportunity to ultimately test their control code on an actual physical robot. These transfer tests often reveal problematic assumptions based on the simulation that do not hold in the noisy, dynamic, real world.

Our experience with moving from the simulator to the real robot indicated the following issues.

- The simulator coordinate systems or function calls did not always behave identically to the actual robot. This required small changes--e.g. query and replace--to code before it would work correctly.
- The simulator constants such as velocity and acceleration had to be adjusted when moving to the actual robots in order to maintain safety for both the robots and the students. The students tended to crank these values up in the simulation.
- The actual sensor values tended to be noisier, and possess biases, compared to the simulator.

We held the lab sessions in a lab with 25 Unix workstations. The students typically worked in teams of three or four (although a few students preferred working alone or in pairs). The teams tended to be fairly fluid, changing from one lab to the next, until the students found a group they felt compatible with. Students had access to two Nomadic Super Scout IIs, four Applied AI Systems Kheperas, and one Active Media Pioneer. From each workstation, the students could also access simulations of the Scout, the Khepera, neural networks, and evolutionary computation. In addition, students were encouraged to use a variety of analysis tools--e.g. hierarchical cluster analysis and principal components analysis, which are both useful in understanding neural network processing.

Lab assignments generally required three weeks of work. At the end of an assignment, we had each team provide a thorough report about what they had accomplished in lab. The lab descriptions are linked to the course home page at http://www.palantir.swarthmore.edu/~maxwell/classes/e28. We encouraged the students to create web pages for their report which included the text as well as any appropriate diagrams and figures to help describe their results. For examples of student reports, see: http://www.palantir.swarthmore.edu/~maxwell/classes/e28/reports.

2.2.1 Lab assignment descriptions

We developed four regular lab assignments and then had the students undertake a final project on a topic chosen by each team.

With the first lab our intention was to introduce students to the basic ideas of robot programming using the Nomad simulator. We had the students write control software for three tasks with increasing difficulty: a random walk through an empty environment, a random walk through obstacles, and navigating to a goal location through obstacles. In the first task, students could focus on understanding the basic movement commands without having to integrate sensing. In the second task, they could build on to what they had already created to use sonar sensors to respond to obstacles. Finally, they could tackle a much harder goal-based task.

We intended the second lab to provide students with experience in standard path planning techniques using the Nomad simulator and the actual Scout II robots. Their first task was to set up a state space for an environment containing obstacles. We gave them the choice of using a regular, two-dimensional grid (easiest), visibility graphs (harder), or Voronoi diagrams (hardest). Within this state space they could search for optimal paths in a known environment using the A* algorithm. This is a slow method and requires a complete recalculation of the plan whenever new information is discovered. Thus, in the next step we had them explore an enhancement to A* called D* which can efficiently re-plan. As the final task, we had them experiment with creating an evidence grid of an unknown environment based on sonar readings. Implementations of A*, D*, and evidence grids were provided, but students needed to modify the code for their needs and integrate them together.

Our goal in the third lab was to allow students to explore robot learning using the Khepera simulator and the Khepera robots. The controller for this lab was an artificial neural network with the parameters being learned through evolutionary computation. To conduct evolutionary experiments the students had to first determine a task to be learned and then design a fitness function that would measure success at the task. Artificial evolution is quite good at finding short cuts and loop holes in human-engineered fitness functions. Therefore, finding a robust fitness measure required extensive experimentation. Once a successful measure had been developed, the students were asked to analyze the final learned behavior in detail.

The final regular lab had students design a robot vision system for a particular task. The students were given several choices: tracking a target, following a sidewalk, or finding faces.

In terms of difficulty and time required, the first lab introducing basic robot programming and the final lab on vision were the most straight-forward. The second lab on path planning was by far the most difficult. In future offerings of the course, we would need to reduce the scope of this lab or break it into multiple assignments.

2.3 Discussion

The focus of our course was to give students a broad introduction to issues in robotics and robot sensing from both a symbolic and sub-symbolic point of view. Our lab assignments reflected this focus, asking students to use both styles of algorithms on traditional robot navigation and robotic sensing. Our final projects, rather than being focused on a competition, allowed students to explore an area of their interest.

This approach contrasts with that of Nourbakhsh, where the focus is on each lab group developing a single mobile robot system from low-level control to high-level planning [7]. Each lab group's robot then participates in a contest at the end of the semester that is open to the public. This model is also used at other universities [7].

Both of these approaches have merit. The strength of our approach is that students get to directly compare different approaches to the same task, and spend similar amounts of effort on each. The drawback, in comparison, is that students get less experience building a complex, integrated system. They do get some experience with system integration, however, as most lab groups developed working real robot systems for their final projects.

3 Undergraduate Robotics Projects

In addition to the curricular elements related to robotics education and research, both authors--Maxwell and Meeden--have organized and supervised extensive undergraduate robotics projects. The projects include both senior projects in computer science and engineering and summer research projects. This section focuses on the undergraduate summer research projects for the summers of 1998 and 1999. For both summers a team of seven undergraduates and two faculty members designed, built, and programmed a robot to serve hors d'oeuvres at the *Hors D'oeuvres*

Anyone? event in the American Association for Artificial Intelligence [AAAI] Robot Competition.

The goal of this event is to serve hors d'oeuvres to the conference attendees during the main conference dinner, which is generally held in a large convention hall. The main goal of this contest is to interact with people in interesting ways while serving hors d'oeuvres in as large an area as possible. Other suggested aspects of the task include returning to a refill station and food manipulation. The major challenges include: sensing people, navigating in a dynamic environment, interacting with people, and finding a specific location in a crowded environment. The nice thing about this event is that it can be solved at different levels; it is not difficult to put together a robot that will satisfy the basic aspects of the task, but it is extremely difficult to develop a robot that can serve hors d'oeuvres in a sophisticated manner. Thus, the problem is an open-ended one and can be a platform both for educational and research goals.

During the summer of 1998, Maxwell was PI for a National Science Foundation Research Experiences for Undergraduates [REU] Site at the University of North Dakota [UND] Department of Computer Science. Under the supervision of Maxwell and Sven Anderson, seven of the 10 undergraduates involved in the 1998 program participated in a designing, building, and programming the robot Rusty to serve hors d'oeuvres at the *Hors D'oeuvres Anyone?* event. The UND team placed first in this event.

The following summer, the authors co-advised seven undergraduates at Swarthmore College for the same event in the 1999 AAAI competition. The Swarthmore team won first in the *Hors D'oeuvres Anyone?* event and was given the award for "Best Integrated Effort" of all robots in all events.

3.1 Research and Educational Goals

Our primary goal in undertaking these research projects was not to win the competitions--although during the second year we had a better idea of the probability that it could happen. Instead, our primary goal for both years was to build a system that could complete the task in a reasonable manner. It was a combination of the advisors' and students' vision and implementation ability that drove the definition of "reasonable". Overall, our goal was to have a robot that could find people using computer vision, navigate to them using sonar, interact with them based on speech recognition and visual input, and get back to a refill station using vision and human interaction.

Our research goals fell into three categories:

- Improving our institution's image in the field of robotics,
- Developing working implementations of current research results, and
- Generating a publishable paper.

Our research goals for both summers were limited, largely because of the background of the undergraduates rather than their ability. None of the undergraduates involved in either summer had experience with robotics; none had taken a robotics course; and only one student in 1998 had taken a computer vision course. Nevertheless, the project goal of putting together a working system--as we defined it--required us to implement and modify state-of-the-art algorithms and methods. In some sense this supports longer-term research goals as it gives us algorithm implementations that we can later build upon.

Our educational goals for both summers were to:

- Give students hands-on experience with real problems,
- Give them experience understanding and implementing primary literature,
- Give them confidence in their ability, and
- Help them develop teamwork skills.

All four of these goals are essential to both further graduate study and technical positions within industry. Note that the first and last are also part of the ABET criteria for engineering accreditation, which partly guides the educational goals of the Swarthmore engineering program [8].

Having the robot contest as a goal turned out to nicely balance our educational and research goals for this group of students. We had to build a real system in a fixed amount of time; the students had to work together in order to complete the project; and each student had to implement their section of it based on textbooks and primary literature. The open-ended nature of the contest allowed us to push the envelope when possible given the students' background and limited time.

3.2 Project Organization

3.2.1 Scheduling

To ensure that we met the contest deadline, we set up specific target dates for completing the different aspects of the task--vision, speech, navigation, and physical construction. The main purpose of the target dates was to focus the students' attention on when integration of the different subsystems would take place. Overall, we scheduled approximately one-third--3-4 weeks--of the 10 week project time for integration and testing, which meant that the individual components had to be largely completed in 6 weeks. This was a difficult time-line, but reflects the tension that exists between the amount of time spent on design, implementation, and integration [9].

One of the later targets--during both summers--was a complete test run at least a week before the contest. This give us time to make final improvements to the system before the AAAI conference. It also meant we arrived at the AAAI contest with a working robot and could spend our time at the conference making improvements rather than implementing basic capabilities. Both years of the contest this was a significant factor in our success.

3.2.2 Team Organization

For both 1998 and 1999 we had seven students working directly on the robot project. As noted above, almost none of the students had prior experience with robotics, computer vision, or speech recognition. The students had a wide variety of backgrounds--especially in 1999, which included an art major and a theatre major in addition to computer science and engineering majors. The variety of backgrounds turned out to be a real boost to the creativity of the project and resulted in more interesting visual and auditory interfaces than would have otherwise been possible.

Unlike many engineering projects, we were not rigid about organization or specifying who was in charge at the beginning of the project [9]. We let students select projects with particular areas, and ended up with two working on speech and interaction--both recognition and output--three on computer vision, and two on navigation and integration. It is important to note that five of the seven people worked primarily on sensing or user interaction.

We also did not restrict the students to one particular aspect of the problem. For example, in 1999 two of the students--one each from the speech and vision groups--developed the physical robot structure, and one of the students from the navigation group was the voice of the robot. The same

Figure 2 Rusty, the 1998 UND robot (left), and Alfred, the 1999 Swarthmore robot (right).

dynamic occurred in 1998. It gave students a chance to pursue a broader variety of interests and lend their creativity in different ways.

A second dynamic that occurred was that in both years a primary programmer emerged in the vision, speech, and navigation groups. This person was the one who integrated the other team members' work into a single program and made sure it all worked together. Thus, although we did not specify a team leader up front, one emerged. This, in turn, made the integration of the three parts a simpler process because it only required the coordination of three people.

For the first six weeks, the students worked primarily in small groups or by themselves. At the integration time--about 6 weeks into the project--they began working in larger groups, with the navigation group coordinating communication. Once the structure of the robot's complete program was finished, the individual groups were able to continue improving their components on their own and could then upgrade the main program to use the new routines.

3.2.3 Robot Platforms

For the summer of 1998 our primary equipment consisted of one Pioneer 1 robot and a Toshiba Tecra laptop. The robot's peripherals included a color QuickCam, a microphone, speakers, and feelers over the wheels to prevent the robot from running over a person's foot. Rusty is shown in the left picture of Figure 2.

For the summer of 1999 we had two Nomad Super Scout II's with on-board computers. We used one in the contest and had the other for backup--which turned out to be essential as the hard drive on the primary robot died at the competition. The peripherals for 1999 were similar to 1998: color CCD camera, microphone, speakers, and a bumper down low to detect feet. Alfred is shown in the right picture of Figure 2.

3.3 Project Outcomes

As noted above, the 1998 UND Team won first place in the hors d'oeuvres event with the robot Rusty and set a standard for interaction and performance. The 1999 Swarthmore team also won first place in the hors d'oeuvres event as well as best overall integrated effort. The innovations of the Swarthmore team--which included a strong personality and a rudimentary ability to remember people--set a new standard for the event.

3.3.1 Research Outcomes

In both years we generated publishable quality work in the form of system implementation papers. The vision system of the 1998 team appeared in [10]. The system development paper for the 1999 robot appeared in [11], and has also been submitted to the Journal of Autonomous Robots.

In both years we pushed the boundaries of what was possible with a low-cost mobile robot system. The total cost of the 1998 system was around $7000, while the total cost of the 1999 system was around $10,000 (not including the backup robot).

These robots have also impacted commercially available robots. Activmedia now has bumpers capable of detecting feet and sells a robot--that looks similar to Rusty--designed for human interaction and equipped with high-quality cameras, speakers, and microphone. Since two of the 1998 team members are now employed by Activmedia, this result is not necessarily surprising.

In terms of long-term research goals, we now have a suite of algorithms that provide a variety of capabilities. Our experience with these algorithms lets us pick and choose what works and what needs to be improved. It also sets the stage for new innovations this coming year.

Finally, successful involvement in the contest gave us positive publicity, both at UND and at Swarthmore. For the authors, it has helped us to make connections within the field and become more active in the robotics community. This will, in the end, benefit our undergraduate research program at Swarthmore.

3.3.2 Educational Outcomes

We were able to make good progress in all of the educational outcomes for the students. They had a hands-on experience with a real problem and a hard deadline. Most of the students read one or more papers from the robotics/computer vision literature and implemented at least a portion of an algorithm or concept from one or more of the papers. The need to work together built teamwork skills; and completing the project and being successful in the contest greatly enhanced their confidence in their abilities.

In addition, the robot project fit well with the interdisciplinary nature of Swarthmore’s engineering program and the college as a whole. As noted by other interdisciplinary engineering programs, robot projects provide a nice mixture of both engineering and computer science topics, challenging the students to integrate and apply their knowledge in a coherent project [12][13][14].

To help directly evaluate the educational outcomes, we asked the 1999 students to discuss their experience with respect to their educational, personal, and professional goals. Below are some of the comments regarding the educational goals. They are, perhaps, the strongest argument for integrating this kind of a project into an undergraduate education.

> *I think that I learned a lot about working with a team and how groups have to interact in order to make a good, finished product. I learned a fair amount about how research can be done, as in searching on the web or through published journals and finding successful ideas and then applying those ideas to help solve your own problem. Previously, I think I was very stuck in Pned* [illegible]

I took away from the summer which will benefit me in all my classes and further research efforts in the future... I guess that I have built a little bit of confidence in myself that I have the ability to do interesting things in the field of robotics or computer science. Like all the people that now come up to me and praise me for undertaking something so foreign, I once thought that these problems were ones best left to be solved by those much smarter than me. I feel that now I can read a wide variety of papers and identify with the struggles and understand the relevant issues. At the same time, I have learned a lot about perspective, as in what it is that I can really do and more importantly, what I can't do. I could very easily be working on these problems for a long time without progressing much except to begin to understand the complexity of all the issues involved in attempting to create artificially intelligent systems. I guess that I just obtained a better sense of where I fit in and what I can contribute to help solve these bigger problems and unanswered questions.

Working on the robotics projects greatly enhanced my computer programming skills. More importantly though, I learned to work and to interact with other students as part of a team... The most important thing I learned from this project was teamwork. It may sound cheesy, but I always assumed that the word "teamwork" was simply a word and had no bearing. However after being a part of the robot project I began to appreciate the responsibilities and patience required for a successful interaction with other people.

I got to be a better C programmer. It was also the first time I dealt with an API, and I feel like I'm better equipped to deal with that kind of environment in the future; and it was nice to have to integrate my code with the code of others... [I gained] confidence, confidence in the power of Swatties especially, but also in my own ability to help create.

It introduced me to Computer Vision which has been intriguing enough for me to embrace as a Honors preparation. It also showed me how various engineering / computer science principles are integrated and put into practice. This time, it really has to work! No theories on paper can save you... I thought the process of integration and the work days at the convention center were the most rewarding. At that point it was really about how different parts of the team communicate with each other both within the robot and on a human level; I learned a lot about competing and how to be ready.

3.3.3 Student Outcomes

Finally, the student's themselves achieved several personal and professional outcomes from being part of the robot project. A number of these were a direct result of participation in the AAAI contest, and would not necessarily have occurred otherwise.

- Each student has a paper to build their credentials,
- They met other roboticists,
- They made robot industry contacts (both years students received job offers), and
- They built friendships with the other members of the team.

Below are selected excerpts from the student's comments.

The experience of working in a creative team and the opportunity to work with my peers is definitely enriching. Some of the working relationships have turned into friendships, including with my professors.

I wanted to know whether my two majors [Theatre and Computer Science] would ever intersect. They did here.

I think that I finally participated in a project that I felt was shaped and directed by me as much as it was by any of the other team members. Although, you two [Maxwell & Meeden] certainly had more say than the students, I did feel a sense of ownership and responsibility that helped me to feel good about what I was doing and like I made a difference in how things turned out.

4 Discussion and Conclusions

For the first time in the summer of 2000 we had a pool of students at Swarthmore who had taken an advanced robotics course, and possibly a computer vision course. By putting together a curriculum that includes courses in artificial intelligence, computer vision, and robotics we have generated both interest and qualified students to work in the field of robotics. The student's preparation paid off in terms of achievement in the student's research proejcts during the summer of 2000. Swarthmore's robot team finished first in both the Urban Search and Rescue and the Hors d'Oeuvres Anyone events. This success occurred in part because we maintained our focus on combining textbooks and primary robotics research literature in the course as well as in the research projects. Thus, our integration of robotics research with the undergraduate courses makes the transition to a robotics project smoother and reduces the steepness of the learning curve.

From the project outcomes, student comments, and introspection, there are several lessons that emerged from the project experience. First, having participation in the contest as a primary goal provided many benefits, especially from an engineering education standpoint. It gave students hands-on experience with a real system on a real problem; it enhanced team development because people had to work together to achieve the goal; and it gave the students an appreciation for the difficulties of integration and system development. The drawback of making participation in the contest the primary goal was that we were not able to do as much pure research.

A second lesson from the project experience is that a leadership structure is helpful to the process of integration, whether it is specified up front or emerges from the group dynamics.

A third lesson is that we should begin the integration of the system components even earlier in order to give the students a greater sense of how their aspects of the project fit together. Related to this is a fourth lesson, which we are taking to heart in the upcoming year. This lesson is that we need more large group meetings to enhance communication and the sense of where each student's work fits. This was the most common comment in the 1999 students' evaluations.

Finally, a generally applicable lesson is that encouraging students to be involved in multiple aspects of the project enhances their interest, creativity, and learning. We saw this during both years of the project, and the result was always positive.

Overall, we feel like we have been successful in integrating robotics research into the undergraduate curriculum. While the pure research results are not necessarily as strong as those of a graduate program, the benefits to the students have been tremendous. Since the students are the primary focus at Swarthmore, this balance is just right for us.

5 References

[1] Kortenkamp, D, Bonasso, P, and Murphy, R (1998) *Artificial Intelligence and Mobile Robots: Case Studies of Successful Robot Systems*, AAAI Press, Menlo Park, CA, 1998.

[2] Latombe, J-C (1991) *Robot Motion Planning*, Kluwer Academic Publishers, Boston.

[3] Martin, M, Moravec, P (1996) "Robot Evidence Grids", CMU Robotics Institute Technical Report, CMU-RI-TR-96-06.

[4] Stentz, A (1994) "Optimal and Efficient Path Planning for Partially-Known Environments", in *Proc. of IEEE Int'l Conf. on Robotics and Automation*, May 1994.

[5] Sun, R, Merrill, E, and Peterson, T (in press) "From Implicit Skills to Explicit Knowledge: A Bottom-up Model of Skill Learning", *Cognitive Science*.

[6] Oates, T, Schmill, M, Cohen, P (in press) "A Method for Clustering the Experiences of a Mobile Robot that Accords with Human Judgements", *Proceedings of the Seventeenth National Conference on Artificial Intelligence*.

[7] Nourbakhsh, Illah (2000) "Robots and Education in the classroom and in the museum: On the study of robots, and obots for study", in *Proceedings, Workshop for Personal Robotics for Education*, IEEE ICRA 2000.

[8] *ABET/EAC Engineering Criteria 2000: Criteria for Accrediting Engineering Programs*, Accreditation Board for Engineering and Technology, Inc., Baltimore, USA.

[9] D'Andrea, R (2000) "Robot Soccer: A Platform for Systems Engineering", Computers in Education Journal, 10(1), pp 57-61.

[10] Maxwell, B, Anderson, S, Gomez-Ibanez, D, Gordon, E., Reese, B., Lafary, M, Thompson, T, Trosen, M, Tomson, A (1999) "Using Vision to Guide an Hors d'Oeuvres Serving Robot", *Proceedings of IEEE Workshop on Perception for Mobile Agents*, June 1999.

[11] Maxwell, B, Meeden, L, Addo, N-S, Brown, L, Dickson, P, Ng, J, Olshfski, S, Silk, E, Wales, J (1999) "Alfred: The Robot Waiter Who Remembers You," in Proceedings of *AAAI Workshop on Robotics*, July, 1999.

[12] DeVault, J (2000) "Robot Stories: Interdisciplinary Design with Autonomous Mobile Robots", Computers in Education Journal, 10(1), pp 21-27.

[13] Giolma, P, Aminian, F, and Ibaroudene, D (2000) "An Autonomous Robot-the Ideal Design Project", Computers in Education Journal, 10(1), pp 62-67.

[14] Heller, M (2000) "Mobile Robots and Interdisciplinary Design - Mobots", Computers in Education Journal, 10(1), 73-76.