Integrating Robotics Research with Undergraduate Education

Bruce A. Maxwell and Lisa A. Meeden
Swarthmore College

TO KEEP UP WITH THE WIDE-RANGING, FAST-MOVING ROBOTICS FIELD, EDUCATION MUST BE ADAPTIVE AND MULTIDISCIPLINARY. THE AUTHORS' APPROACH BALANCES RESEARCH AND IMPLEMENTATION FUNDAMENTALS BY REINFORCING COURSE WORK WITH INTENSIVE SUMMER PROJECTS FOCUSED ON A ROBOT COMPETITION.

Robotics is a true multidisciplinary field that forces us to cross traditional disciplinary boundaries to develop working systems. In addition to the electromechanical systems that endow mobility, most autonomous robots also contain one or more computers and the software and hardware scaffolding necessary to support them. The field has evolved quickly over the past decade, largely because of the tremendous increase in computing power and the availability of an improved variety of sensors. We are now sending robots to Mars, to the depths of the ocean floor, and into hazardous nuclear reactors. We are giving robots the capability to learn, to act autonomously, and to interact with humans and their environment. These attributes are necessary if they are to successfully accomplish their tasks. Both graduate programs and industry need students who are ready for this challenge. The two most important things we can give students as undergraduates are a strong knowledge of the fundamentals and experience with real robot systems.

Swarthmore College takes a two-pronged approach to undergraduate education by integrating educational goals with robotics research. First, we offer courses in artificial intelligence, computer vision, and robotics. Both AI and computer vision serve as prerequisites for the robotics course. Second, we involve students in ongoing research projects as part of their undergraduate experience. Honors students in all disciplines must participate in a research project; other students are encouraged to do so. The research projects are team-oriented and have thus far focused on the design and building of a robot to do highly specific tasks, such as serving food.

In this article, we describe two undergraduate group projects we conducted, one from 1998 at the University of North Dakota advised by Maxwell, and one from 1999 at Swarthmore advised by both authors. The impetus for these projects was the American Association for Artificial Intelligence's annual robot competition. These experiences led to the development of a new robotics course at Swarthmore which we co-taught in the spring of 2000. This course focused on a set of fundamental concepts in mobile robotics that would prepare students for research projects in this area.

Robotics projects

In both 1998 and 1999 a team of seven undergraduates and two faculty members designed, built, and programmed a robot to serve hors d'oeuvres at the AAAI's Robot Competition's "Hors d'Oeuvres, Anyone?" event.

At this event, robots serve hors d'oeuvres to the conference attendees during the main conference dinner. The goal is to have robots interact with crowds of people in interesting ways while serving hors d'oeuvres in as large an area as possible, usually a convention hall. Other tasks can include the robot's returning to a refill station to replenish the robot's serving tray.

The major challenges that this event poses to students in programming the robot include sensing people's presence, navigating in a dynamic environment, interacting with people, and finding a specific location in a crowded environment. One advantage of this event is that it can be solved at different levels. It's not hard to build a robot that satisfies the task's basic elements, for example, but it is extremely difficult to develop a robot that serves hors d'oeuvres in a sophisticated manner. Such an open-ended problem can be a platform for both educational and research goals.

During the summer of 1998, Maxwell was principal investigator for a National Science Foundation Research Experiences for Undergraduates Site at the University of North Dakota Department of Computer Science. Under the supervision of Maxwell and Sven Anderson, seven of the 10 undergraduates designed, built, and programmed the robot nicknamed "Rusty" for the AAAI event. The team placed first. Rusty was an ActivMedia Pioneer I on which we placed a laptop, a table, camera, speakers, and microphone.

During the summer of 1999, after Maxwell moved to Swarthmore College, the authors coadvised seven undergraduates at Swarthmore for that year's AAAI hors d'oeuvres event. The team, with a robot named "Alfred," won first place and was also awarded Best Integrated Effort for all robots in all events.

Research goals. The primary goal for both years' research projects was to build a system that could reasonably complete the competition's task. "Reasonably" was defined by a combination of the advisors' and students' vision and implementation ability. Ultimately, we wanted a robot that could locate people using computer vision, navigate to them using sonar, interact with them based on speech recognition and visual input, and travel to a refill station using vision and human interaction.

Our research goals fell into three categories:

- improving our institution's reputation in the field of robotics,
- developing working implementations of current research results, and
- generating publishable results.

We limited our research project goals for both summers, largely because of the undergraduates' background rather than their ability. The undergraduates had no prior robotics experience, nor had they taken a robotics course; only one student in 1998 had taken a computer vision course. Nevertheless, the project goal of assembling a working system—as we defined it—required us to implement and modify state-of-the-art algorithms and methods. This work supports our longer-term research goals because it provide algorithm implementations that we can later build on.

Educational goals. Our educational goals for both summers were as follows:

- Give students hands-on experience with real problems.

UNDER THE SUPERVISION OF MAXWELL AND SVEN ANDERSON, SEVEN OF THE 10 UNDERGRADUATES DESIGNED, BUILT, AND PROGRAMMED THE ROBOT NICKNAMED "RUSTY" FOR THE AAAI EVENT.

- Give students experience in understanding and implementing robotics principles from primary research literature.
- Give students confidence in their ability.
- Help students develop teamwork skills.

All four goals are essential if students intend to pursue graduate study and obtain technical positions within industry. Notably, the first and last goals are included in the Accreditation Board for Engineering and Technology criteria for engineering accreditation. The ABET criteria partly guide the Swarthmore engineering program's educational goals.[1]

The AAAI robot competition let us balance our educational and research goals for both groups of students. Both summers, the students built a real system in a fixed amount of time; they worked together to complete the project; and they implemented their respective portions of it on the basis of textbooks and primary literature. The open-ended nature of the competition afforded us significant flexibility in developing the robot designs.

Project organization

We addressed three major considerations in organizing the summer projects: scheduling, team organization, and the robot platform.

Scheduling. To meet the competition deadline, in both years we established target dates for completing the different task elements—vision, speech, navigation, and physical construction. The target dates focused the students' attention on the timing required to integrate the different subsystems. We scheduled approximately one-third (three to four weeks) of the 10-week project time for integration and testing, which meant that the individual elements had to be largely completed in six weeks. This difficult timeline reflects the constraints imposed by the time required to complete design, implementation, and integration.[2]

One of our project milestones was a complete test run at least a week before the competition, which gave us time to fine-tune the robot before the AAAI conference. It also meant we arrived at the AAAI competition with a working robot and could spend our time making additional improvements. The time we spent on improvements was a significant factor in our success both years.

Team organization. For both project years, the students' backgrounds varied considerably. The 1999 team included an art major and a theatre major in addition to the usual computer science and engineering majors. Such varied backgrounds gave a real boost to the project's creativity and resulted in more interesting visual and auditory interfaces than would have otherwise been possible. For example, the visual appearance of Alfred—the 1999 robot who looked like a penguin in a tuxedo—was significantly more sophisticated than Rusty, who was essentially a rolling table.

Although many engineering projects are rigid about organization,[2] in both years we let the students decide how the project team would evolve and also the area of the robot project on which they would work. As a result, two students worked on speech and interaction—including recognition and output—three on computer vision, and two on navigation and integration. Note that five of the seven students worked primarily on sensing or user interaction. This reflects the strong emphasis we had on the human interaction aspect of the hors d'oeuvres task, which is also its most challenging aspect.

Figure 1. Rusty was the 1998 University of North Dakota team's robot, left. Alfred was the 1999 Swarthmore team's robot, right. Rusty was an ActivMedia Pioneer I on top of which we built a table structure to hold a speaker, camera, and microphone. We were a bit more creative with Alfred, a Nomadic Super Scout II that we dressed up to look like a penguin butler. With Alfred we were able to hide the speakers inside, put the camera in his mouth, and have the microphone coming over his shoulder.

Students had ample opportunity to pursue a broad variety of interests and lend their creativity in different ways. For example, in 1999, two students—one each from the speech and vision groups—developed the physical robot structure, and one student from the navigation group was the robot's voice. The same dynamic occurred in 1998.

A second dynamic that occurred was that, although we had not specified a leader upfront, in both years a primary programmer emerged to serve as a team leader in each of the vision, speech, and navigation groups. This person integrated the team members' work into a single program and made sure it worked, which simplified subsystem integration because it required the coordination of only three people.

Overall, for the first six weeks, the students worked primarily alone or in small groups. At integration time, about six weeks into the project, they began working in larger groups, with the navigation team coordinating communication. Once the students had completed the robot's program structure, the individual teams continued improving their respective components for later updating to the main program.

Robot platforms. For Rusty, 1998's robot, our primary equipment consisted of one ActivMedia 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's running over a person's foot. Figure 1 shows Rusty at the left.

For Alfred, 1999's robot, we had two Nomad Super Scout IIs with on-board computers. We used one in the competition and reserved the other for backup, which turned out to be essential because the hard drive on the primary robot died at the competition. The peripherals included a color camera, microphone, speakers, and a bumper down low to detect feet. Figure 1 shows Alfred at the right.

Project outcomes. As noted above, both teams won first place. Rusty set a standard for interaction and performance by being the first robot in the hors d'oeuvres event to integrate navigation, vision, and speech—both input and output—on a single platform. The innovations of the Swarthmore team's Alfred—which included a strong personality and a rudimentary ability to remember people—also set a new standard for the event.

Both robots performed their best in a preliminary round in which the judges and a small number of others participated. The robots were able to be heard, and to hear and recognize what other people said, which consisted mostly of yes/no and please/thank you. In the 1999 competition, for example, Alfred correctly recalled and commented on the fact that he had seen one judge twice. He also correctly commented on the politeness of the people with which he interacted. Fortunately, neither robot spilled anything on a judge.

Unfortunately, in the final rounds—when there are close to 1,000 people trying to watch and interact with the robots in a single noisy room—both robots had difficulty recognizing what was said. Partly as a result of this, people were not as willing to talk with the robot and get through a complete interaction. Nevertheless, the robots did serve many hors d'oeuvres and managed to run safely for over an hour in a difficult environment.

Research outcomes. We documented the teams' results in papers describing system implementation, the research of the 1998 team's vision system work,[3] and the 1999 team's robot system development have been published.[4]

Twice, we pushed the boundaries of what was possible with a low-cost mobile robot system. The total cost of the 1998 system was approximately $7,000, and the 1999 system cost roughly $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 (resembling Rusty) designed for human interaction and equipped with high-quality cameras, speakers, and microphone. Considering that two of the 1998 team members now work at Activmedia, this result is not necessarily surprising.

We also 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 improvement. It also sets the stage for innovations this coming year.

Finally, successful involvement in the AAAI competition garnered positive publicity, both for the University of North Dakota and for Swarthmore. The successes have also helped us become better known and more active in the robotics community, which will ultimately benefit the Swarthmore undergraduate research program.

Educational outcomes. We achieved good progress in all of the students' educational outcomes both from a pedagogical point of view and based on their own evaluations. The students 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 and computer vision literature and implemented at least a portion of an algorithm or concept from one or more of the papers. The need to collaborate fostered students' teamwork skills, while project completion and competition success greatly enhanced the students' self-confidence.

The robot project at Swarthmore, in particular, complemented the interdisciplinary nature of Swarthmore's engineering program and of the college as a whole. As other interdisciplinary engineering programs have noted, 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.[5-7]

To help evaluate the educational outcomes, we asked the 1999 students to discuss their experience with respect to their educational, personal, and professional goals. The "Robot project: Student participants' comments" sidebar lists some of the comments regarding the educational goals.

Student outcomes. Finally, the students themselves achieved several personal and professional outcomes from being part of the robot project. A number of these resulted directly from participation in the AAAI competition and would not necessarily have occurred otherwise.

- All students have a paper to build their credentials.
- Students met other roboticists.
- Students made robot industry contacts (students received job offers both years).
- Students built friendships with the other team members.

Below are representative student comments.

> 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. … I did feel a sense of ownership and responsibility.

Curriculum

We offered an upper-level robotics course for the first time at Swarthmore College in the spring semester of the 1999–2000 academic year. The course was cross-listed under Engineering and Computer Science and team-taught by us. About one-third of the 25 students enrolled were engineering majors; the rest 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. We developed this course, in part, to prepare students for summer research projects in robotics like those described above.

Course organization. Robotics is well served by having multiple instructors handle the teaching. In our case, Maxwell is an engi-

THE NEED TO COLLABORATE FOSTERED STUDENTS' TEAMWORK SKILLS, WHILE PROJECT COMPLETION AND COMPETITION SUCCESS GREATLY ENHANCED THE STUDENTS' SELF-CONFIDENCE

neer whose research focus is computer vision, and Meeden is a computer scientist whose research focus is learning. Team-teaching let us cover more material than either of us could have done alone.

The course outline follows:

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

We taught the topics as five two-to-three-week sections. While one instructor led the lectures—three one-hour lectures per week—the other instructor ran the labs: one three-hour session weekly. Our goal in selecting topics was, first, to ground students in the concepts of several areas—geometry, machine learning, sensor analysis, visual sensing, robot software architectures, and robot–human interfaces. At the same time, our goal was to expose students to current research that variously extended and implemented these fundamental concepts.

Our course work organization loosely followed the structure of our chosen text, *Artificial Intelligence and Mobile Robots*.[8] This book is an edited collection of successful robotic-systems case studies divided into three main categories. The studies correlate with the topics in the above list as follows: "Mapping and Navigation" (topics 1 and 2); "Vision for Mobile Robots" (topic 3), and "Mobile Robot Architectures" (topic 4). Although some of the book's case studies are now dated, the underlying concepts remain appropriate as an introduction to the field. Topic 5 was based on primary literature in the field.

To combine the history of robotics with current research, we supplemented this text with more than 20 research articles. For example, to learn more about robot navigation, students read about path planning,[9] evidence grids,[10] and D* (a real-time replanning algorithm based on A*).[11] For additional information about robot architectures, students read about bottom-up learning[12] and developing human-like concepts in robots.[13] The complete list of supplemental reading materials is at www.palantir.swarthmore.edu/~maxwell/classes/e28/readings.htm.

Laboratories. Ideally, students in a robotics class have frequent hands-on laboratory experience with physical robots. When class size and equipment availability constraints prevent hands-on work, simulation serves as an initial testbed for design and implementation. Ultimately, however, students must be able to test their control code on an actual robot. Such *transfer tests* often reveal simulation assumptions to be problematic because they might not apply in the real world.

Our experience with moving from the simulator to the robot highlighted the following issues:

- The simulator coordinate systems or function calls did not always behave identically to the actual robot. This required students to make small changes—for example, query and replace—to the code before it would work correctly.
- The simulator constants such as velocity and acceleration required adjustments for safety reasons when students began working with the actual robots. The students tended to "crank up" these values in the simulation.
- The actual sensor values tended to be noisier, and possess biases, compared to those in the simulator.

We conducted lab sessions 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 worked with two Nomadic Super Scout IIs, four Applied AI Systems Kheperas, and one ActivMedia Pioneer. From each workstation, the students could also access simulations of the Scout, the Khepera, neural networks, and evolutionary computation. Also, students were encouraged to use tools—for example, hierarchical cluster analysis and principal components analysis, which are both useful in understanding neural network processing.

Lab assignments. We developed four regular lab assignments and then had the students undertake a final project on a topic chosen by each team. Lab assignments generally required three weeks' work. At the assignment's end, each team provided a thorough report about what they had accomplished. The lab descriptions are linked to the course home page at 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 www.palantir.swarthmore.edu/~maxwell/classes/e28/reports.

Goals. The goal of our first lab was to introduce students, with the Nomad simulator, to robot programming's basic concepts. The students wrote control software for three tasks of 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 what they had already created to use sonar sensors to respond to obstacles. Finally, they could tackle a much harder goal-based task.

The second lab provided students with experience in standard path-planning techniques via the Nomad simulator and the actual Scout II robots. Their first task was setting up a state space for an obstacle-strewn environment. Three design choices of increasing difficulty faced the students: a regular, 2D grid; visibility graphs; or Voronoi diagrams. Within this state space, students searched for optimal paths in a known environment using the A* algorithm. This algorithm is slow and requires a complete recalculation of the plan whenever new information is discovered.

THE STRENGTH OF OUR APPROACH IS THAT STUDENTS DIRECTLY COMPARE ONE OR MORE DIFFERENT APPROACHES TO THE SAME TASK AND EXPEND EFFORT ON EACH.

In the next step, as a result of A*'s slowness, we had them explore an enhancement called D*, which can efficiently replan. As the final task, the students experimented with creating an evidence grid of an unknown environment based on sonar readings. Implementations of A*, D*, and evidence grids were available, but students had to modify the code for their needs and integrate the different techniques.

In the third lab, students explored robot learning with the Khepera simulator and the Khepera robots. The controller for this lab was an artificial neural network, with the students learning the network parameters through evolutionary computation. To conduct evolutionary experiments, the students had to determine a task to be learned and design a fitness function that measured task success. Artificial evolution is quite good at finding shortcuts and loopholes in human-engineered fitness functions. Therefore, finding a robust fitness measure required extensive experimentation. Once they had developed a successful measure, the students analyzed the final learned behavior in detail.

In the fourth lab, students designed a robot vision system for a particular task. The students' task choices were to track a target, follow a sidewalk, or find faces.

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

Discussion. We wanted to give students a broad introduction to robotics and robot-sensing issues from both a symbolic and sub-symbolic point of view. The lab assignments reflected this focus, requiring students to use both styles of algorithms on traditional robot navigation and robotic sensing. To provide flexibility we let the students pick their own final project, which usually involves picking task and programming the robot to solve it.

This broad approach that compares different techniques contrasts with that of Illah Nourbakhsh, who has students develop a single mobile robot system from low-level control to high-level planning.[14] In Nourbakhsh's approach, also used by other universities, each lab group's robot then participates in a public contest at the semester's end.[14]

The strength of our approach is that students directly compare one or more different approaches to the same task and expend effort on each. The drawback is that students get less experience building a complex system, although they do gain system integration exposure because most lab groups develop working robots for their final projects.

The 1998 and 1999 projects inspired us to develop our robotics course. This in turn inspired us to return to the AAAI competition in the summer of 2000. The course generated both interest and qualified students to work on it. Most of the students on the 2000

Swarthmore team took the course outlined above. The students' preparation paid off in terms of achievement in research projects completed during the summer of 2000. Swarthmore's robot team finished first in two events held by the AAAI: the "Urban Search and Rescue" and the ever-popular "Hors d'Oeuvres, Anyone?" events.

Throughout the project we maintained our focus on combining textbooks and primary robotics research literature, which contributed to the 2000 projects' success. So, our integration of robotics research with the undergraduate courses clearly smoothes the transition to an intensive robotics project and reduces the learning curve.

From the project outcomes, student comments, and introspection, several lessons have emerged. First, participation in the competition as a primary goal gave the students an appreciation for the difficulties of integration and system development. The drawback of making competition participation the primary goal was that we were unable to do as much pure research as we would have liked.

Another lesson is to begin the integration of the system components as early as possible. This gives the students a greater sense of how their project responsibilities fit together. Related to this is the need for regular group meetings to enhance communication and the sense of how each student's work fits.

We have successfully integrated robotics research into the undergraduate curriculum. Although the pure research results are not necessarily as strong as those that might result from a graduate program, the benefits to the students have been tremendous. Because the students are the primary focus at Swarthmore, this balance is just right for us. ■

Robot project: Student participants' comments

The following are selected student comments evaluating their participation in the robot projects we supervised:

> 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 …. I think that understanding of how learning can take place is probably the biggest thing that I took away from the summer which will benefit me in all my classes and further research efforts in the future .… 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.

> 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.

> I got to be a better C programmer. It was also the first time I dealt with an [application programming interface], 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 … in my own ability to help create.

> It introduced me to computer vision, which has been intriguing enough for me to embrace as an honors' preparation. It also showed me how various engineering/computer science principles are integrated and put into practice ... the process of integration and the workdays 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.

References

1. Accreditation Board for Eng. and Technology Inc., "ABET/EAC Engineering Criteria 2000: Criteria for Accrediting Engineering Programs," Baltimore, 2000.
2. R. D'Andrea, "Robot Soccer: A Platform for Systems Engineering," *Computers in Education J.*, Vol. 10, No. 1, 2000, pp. 57–61.
3. B. Maxwell et al., "Using Vision to Guide an Hors d'Oeuvres Serving Robot," *Proc. IEEE Workshop on Perception for Mobile Agents*, IEEE Press, Piscataway, N.J., 1999.
4. B. Maxwell et al., "Alfred: The Robot Waiter Who Remembers You," *Proc. AAAI Workshop on Robotics*, AAAI Press, Menlo Park, Calif., 1999.
5. J. DeVault, "Robot Stories: Interdisciplinary Design with Autonomous Mobile Robots," *Computers in Education J.*, Vol. 10, No. 1, 2000, pp. 21–27.
6. P. Giolma, F. Aminian, and D. Ibaroudene, "An Autonomous Robot—The Ideal Design Project," *Computers in Education J.*, Vol. 10, No. 1, 2000, pp. 62–67.
7. M. Heller, "Mobile Robots and Interdisciplinary Design—Mobots," *Computers in Education J.*, Vol. 10, No. 1, 2000, pp. 73–76.
8. D. Kortenkamp, P. Bonasso, and R. Murphy, *Artificial Intelligence and Mobile Robots: Case Studies of Successful Robot Systems*, AAAI Press, Menlo Park, Calif, 1998.
9. J-C. Latombe, *Robot Motion Planning*, Kluwer Academic, Boston, 1991.
10. M. Martin and P. Moravec, *Robot Evidence Grids*, Tech. Report CMU-RI-TR-96-06, Robotics Inst., Carnegie Mellon Univ., Pittsburgh, 1996.
11. A. Stentz, "Optimal and Efficient Path Planning for Partially-Known Environments," *Proc. IEEE Int'l Conf. Robotics and Automation*, IEEE Press, Piscataway, N.J., 1994.
12. R. Sun, E. Merrill, and T. Peterson, "From Implicit Skills to Explicit Knowledge: A Bottom-up Model of Skill Learning," *Cognitive Science*, to be published.
13. T. Oates, M. Schmill, and P. Cohen, "A Method for Clustering the Experiences of a Mobile Robot that Accords with Human Judgements," *Proc. 17th Nat'l Conf. Artificial Intelligence*, AAAI Press, Menlo Park, Calif., 2000.
14. I. Nourbakhsh, "Robots and Education in the Classroom and in the Museum: On the Study of Robots, and Obots for Study," *Proc. Workshop for Personal Robotics for Education*, IEEE Int'l Conf. on Robots and Automation, IEEE Press, Piscataway, N.J., 2000.

Lisa Meeden is an associate professor of computer science at Swarthmore College. Her primary research interests are evolutionary robotics and robot learning. She received her BA degree in mathematics from Grinnell College and her PhD degree in computer science with a minor in cognitive science from Indiana University. Contact her at Computer Science Dept., Swarthmore College, 500 College Ave., Swarthmore, PA 19081; meeden@cs.swarthmore.edu.

Bruce Maxwell is an assistant professor of engineering at Swarthmore College. His primary research interests are computer vision, robotics, computer graphics, data mining, and visualization. He received a BA degree in political science and a BS degree in engineering with a concentration in computer science from Swarthmore College, an MPhil degree from Cambridge University, and a PhD in robotics from the Robotics Institute at Carnegie Mellon University. Contact him at Engineering Dept., Swarthmore College, 500 College Ave., Swarthmore, PA 19081; meeden@cs.swarthmore.edu.