

Formal Measures of Learning in a Secondary School Mobile Robotics Course

Illah R. Nourbakhsh, Emily Hamner
The Robotics Institute
Carnegie Mellon University
Pittsburgh, USA
{illah, etf}@andrew.cmu.edu

Kevin Crowley, Katie Wilkinson
Learning Research and Development Center
University of Pittsburgh
Pittsburgh, USA
{crowleyk+, kaw16+}@pitt.edu

Abstract—During the summer of 2002, Robotic Autonomy was taught to thirty students at Carnegie Mellon West in cooperation with NASA/Ames (Moffett Field, CA). The authors at Carnegie Mellon University's Robotics Institute and at the University of Pittsburgh's Learning Research and Development Center planned a methodology for formally evaluating the educational efficacy of Robotic Autonomy. This article describes the educational analysis methodology and the statistically significant results of our analysis, demonstrating the positive impact of Robotic Autonomy on student learning beyond the boundaries of specific technical concepts in robotics. Educational robotics is gaining traction at all levels of the educational system, however formal analysis of its effectiveness has been lacking. This research project seeks to address this weakness by presenting statistically significant evidence of broad learning.

Keywords—educational robotics, personal robotics, curriculum

I. INTRODUCTION

Robots have been playing an active role in education since the advent of the LOGO Turtle [4,7,8,18]. Both as project foci in laboratory coursework and as team challenges in national contests, the process of designing, building and programming robots has served to excite students across a broad age range. The current field of robotic educational endeavors is extremely large and diverse; see [9,10] for an overview.

We have had two primary goals in designing and executing a new robotics course. First, we planned to explicitly evaluate the educational impact of robotics on secondary level students. We were interested in quantifying lessons learned in service of robotics that are broadly applicable to learning. Second, we hoped to collect data covering a far longer span of time than can be afforded based on a weekend robotics course. Not only would the planned course need to fill a summer; but the students would be able to continue their explorations at home, even after course completion.

To enable our basic goal—the educational assessment of a long-term course of study in robotics—we embarked upon a complete design effort to create and execute a new class, Robotic Autonomy. A new educational robot was designed and fabricated in quantity for this course [11]. The robotics curriculum, inspired by earlier efforts in robotics education [16,17], was synchronized to the competencies of the new robot. Finally, educational assessment was planned into the course schedule and staff hiring needs.

II. COURSE OVERVIEW

A. Course Organization

Robotic Autonomy was taught over a seven-week period in the summer of 2002 at the Carnegie Mellon West campus, located within NASA/Ames Research Center. The top-level goal for this course was straightforward: to provide selected high school students with an immersive exploration of mobile robotics using leading-edge technologies. Course graduation was intended to mark, not the completion of these educational activities, but a launching point: every student would take home a robust, programmable mobile robot system for continued exploration for months and years. Although robotics would be the focus of this curriculum, we hoped that lessons learned would encompass important concepts reaching well beyond just robotics.

A sufficiently engaging mobile robot was not available commercially at a reasonable price for such long-term student robot interaction. Thus short-term robotic educational efforts often turn to Lego building blocks, usually designing curriculum both around robot morphology and construction as well as robot programming and interaction [20,21]. Another successful approach has been the integration of research robots and field robot prototypes into curriculum, where time with the robot is rare and therefore valuable [5,6,13]. In order to provide every graduate of Robotic Autonomy with a rich, programmable robot that would be robust to hundreds of hours of use, we chose instead to design and produce a new educational robot.

Dubbed the Trikebot, this robot includes on-board computer vision and distance sensing as well as a wireless 802.11b networking card (see Fig. 1). In preparation for the course, 30 Trikebot fast-build kits were designed and fabricated.

The Robotic Autonomy course was aimed at students entering their senior year of high school, and specified one prerequisite: the successful completion of any introductory programming course. Following the application and acceptance process, student composition ultimately included 18 students attending under full scholarship and 10 paying students. The scholarship students were from various underprivileged backgrounds, and were primarily Hispanic. The course was comprised of 8 girls and 20 boys.

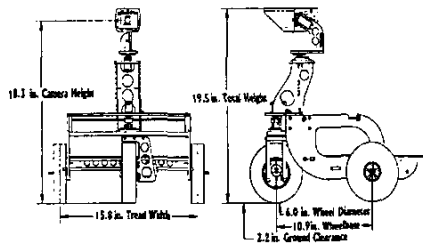


Figure 1. The Trikebot educational robot

The course structure depended primarily on teamwork. Principles governing effective teamwork were explicitly discussed. Students were divided into teams of three, with single-gender teams whenever possible. Based on previous experience teaching robotics courses at the undergraduate level, we felt that single-gender female teams would be more likely to encourage active participation by all members of the team, especially in the case of shy female students. Throughout the seven weeks, each team shared joint responsibility to meet course challenges, with all members of the team receiving the same grade on each week's activities. In order to tackle weekly assignments, the team used just one of their three robots in early weeks, but by the first month's end made use of all three robots in robot team exercises.

B. Curriculum Overview

Robotic Autonomy sought to present the complexity of mobile robotics incrementally and in a bottom-up fashion through a series of challenge-based exercises. The bottom-up approach, which maximizes exploration and self-discovery, is inspired by Constructionism [18]. Challenge-driven robot exercises have been popular in a variety of venues, from short-term robot contests [20] to undergraduate education [1,2,3,12,13,14,17].

The week-by-week progression of concepts was designed in view of these guiding principles while the intra-week structure of the course was designed to be consistent throughout. Monday and Tuesday were spent presenting new material and posing a new, open-ended challenge for each team to tackle. Wednesday was Challenge Day, including extensive testing of the challenge submissions of every team. In addition, a portion of this day was set aside for each team to document their weekly solutions, including source code, prose, pictures and videos to be placed on a specially configured team website. On Thursday morning, teams received the details of the end-of-week contest, which would apply the concepts learned for that week's challenge in an enjoyable and competitive format. Thus Thursday was spent preparing carefully for the next day's contest. Friday was Contest Day, with invited guests (parents, NASA administrators, and visitors) watching and cheering as team robots engaged in games such as line-following races, bomb defusing contests, musical chairs, etc.

In summary, new concepts were largely presented early in the week, with the most difficult bar set by the Wednesday challenge. Following this intellectual apex, the Friday contest offered a chance for students to reuse lessons learned that week in an enjoyable and playfully competitive atmosphere. In

addition to the direct lectures and challenges, weekly guest speakers were brought in on Mondays and Tuesdays to provide one-hour discussions on their areas of expertise. These speakers provided both an outside perspective on robotics and a window into the lifestyle of career roboticists.

Robotics learning topics included: kinematic analysis, fundamentals of electrical engineering, dead-reckoning, feedback control, reactive control, human-robot interfaces, proprioception, back-EMF based speed sensing, visual servoing, vision-based navigation, etc. For the complete curriculum of Robotic Autonomy as well as all student web site material see [15,19].

III. EDUCATIONAL ANALYSIS METHODOLOGY

We assessed impact of the course experience on two levels. First, we conducted a broad evaluation of all the students' experiences in Robotic Autonomy. This evaluation was intended to provide both formative and summative information about whether the course was connecting with students at the appropriate level and making progress toward the broad instructional goals. Second, we conducted an in-depth study of one week of the course. This study, focusing on the experience of two teams of students, was intended to identify some of the micro-genetic mechanisms of learning that might inform patterns of change described in the broader evaluation.

A. Data Collected: Whole Course Evaluation

At the broadest level, four classes of data were used to evaluate the educational effectiveness of the Robotic Autonomy class. First, students completed anonymous surveys about what they were learning throughout the course. On the first day of class, students completed an initial survey of 14 questions covering their technological backgrounds, their expectations for what they would learn in the course, and their plans for college and beyond. Each Monday throughout the course, students also completed a written survey asking them to reflect on the prior week's activities. During the last week of class, students completed a final survey that included similar content to the initial survey, but also asked specific questions about whether and how students had learned about the core themes and content of the course.

Second, in addition to the weekly written feedback, an on-site ethnographer conducted on-camera interviews with each team. The teams were asked about their progress on the assignments and whether anything particularly notable had occurred that week. A total of 9 hours of weekly team interviews were collected, with approximately 1 hour of interview time per team.

Third, students were required to open-source and document their challenge programs on the class website. The format included an explanation of what the program did and how to use it, an analysis of its performance and limitations, suggestions for future improvements, and photographs and videos of the robot performing *in situ*. Each team created seven open-source robotics websites to fulfill this requirement. Also associated with each weekly challenge was a grade assigned by the instructor using both quantitative and qualitative grading criteria. The student documentation and grades enabled us to

analyze the “output” of student learning over the span of the course.

Finally, after completion of the course, follow-on data were collected in the form of monthly online surveys. These surveys asked students about their attitudes toward robotics, science, and engineering; their activities with respect to robotics over the past month; and their future robotics and career plans. In the first 6 months following the end of class, monthly surveys were consistently collected from more than two-third of course graduates.

B. Data Collected: One week in-depth evaluation

In addition to the overall evaluation of the Robotics Autonomy class, an intensive, one-week study of two of the nine teams was conducted to develop a more detailed description of the learning and problem solving that occurred in the course on a minute-to-minute basis. Each team spent approximately four hours a day engaged in group work leading up to the contest and challenge problems. An ethnographer videotaped these problem-solving sessions. Due to space constraints one-week ethnography results are not detailed here.

C. Development of Learning Themes

In order to facilitate the evaluation of learning in the students, it was important to partition expected learning into a set of learning themes for which data would then be quantitatively coded. We hypothesized that six learning themes were particularly well suited to the learning taking place in an interdisciplinary program such as Robotic Autonomy. The themes chosen were: Mechanics, Programming, Teamwork, Problem Solving, Robot Point of View (Robot POV), and Self-Identification with Science and Technology (ID with Technology). The first two themes, Mechanics and Programming, encompass obvious lessons garnered from direct interaction with building and programming robots. The remaining four themes represent important additional opportunities for learning. These themes (Teamwork, Problem Solving, Robot POV, ID with Technology) represent the types of broader learning goals popular in curriculum design. Although popular as design goals, such broad categories rarely yield demonstrable gains, particularly in short-term programs such as Robotic Autonomy.

Mechanics. Mechanics embodies the interrelationship between various kinematics substructures of the robot and the kinematics of the overall robot. This includes an understanding of mechanical components and the manner in which all these components function together as a deterministic whole system. Basic mechanisms (servos, motors, chassis, suspension, bearings) and electronics (motor controllers, microprocessors, range-finding sensors, the vision system, the iPAQ) comprise this category. Because Robotic Autonomy students began the course by constructing the Trikebot rover using a fast-build kit, we hypothesized significant learning in the area of Mechanics, particularly in the early weeks of the course.

Programming. Programming includes learning how to write commands and scripts that control the robot using, in this case, the Java programming language. The programming skills learned extend well beyond robotics, encompassing code

generation / code writing, debugging, documenting, and commenting. Because the Robotic Autonomy challenges posed to the students were primarily challenges for the behavior of the Trikebot, we anticipated that a great deal of the direct learning with respect to overcoming daily challenges would fall in the category of Programming.

Teamwork. Learning how to work effectively in teams is a crucial ingredient for success in many endeavors. Specific skills within teamwork include generating and vetting new ideas; assigning roles and responsibilities; and co-constructing knowledge through observation, imitation, conversation and other socio-cognitive processes. Thus learning progress relative to teamwork would be an important focus of any educational evaluation. In Robotic Autonomy all students worked in teams of three on every phase of project completion. The Robotic Autonomy teams were formed in the first week and left intact throughout the seven-week curriculum.

Problem Solving. Robots such as the Trikebot are extremely complex machines. As such, the process of understanding and refining solutions using the Trikebot requires mastery of problem solving methodologies. Such skills include developing effective strategies for solving the problems that arose throughout the course: setting appropriate subgoals, using feedback from the robot to effectively identify weaknesses in current strategies, knowing when to abandon ineffective approaches, etc.

Robot Point of View. This relatively focused learning theme relates to a critical skill in the understanding of a robot’s operating sphere of influence. Robots are extremely limited, in that their sensory and effector systems are highly constrained relative to that of a human. By *robot point of view* we mean the ability to “see” through the robot’s eyes and thus understand the sensor limitations and action constraints under which the robot must operate. It is only by assuming an appropriate robot point of view that a robot designer can begin to discern the space of possible behaviors that are feasible from those that are impractically ambitious.

Self-Identification with Science and Technology. This extensive learning theme encompasses broad empowerment with respect to science and technology. This includes developing an interest in technology, confidence in one’s ability to work with technology, and interest in pursuing education and future careers in science and technology. In short, this theme considers students coming to see themselves as people who enjoy and are capable of technological explorations.

D. Theme Coding Process

Two reviewers collaborated to code the learning themes. Each of the six themes was divided into general and specific subcategories. For example, for the Programming learning theme, a response that simply said “programming” would be put in the General Programming subcategory, while a response that said “programming in Java” would be coded under the specific subcategory of Java or Other Programming Language.

Of the 452 responses coded in the Initial, Final, and Weekly surveys, only 5 did not fit into the learning themes. That 98.9%

of the responses fit the learning themes supports the validity of the coding scheme.

Once the themes were coded we calculated the proportion of times each student mentioned each specific category. The same proportions were calculated for the whole class using first a sum of the total mentions of a theme and then a count of the number of students who specified a theme. Analysis of variance (ANOVA) was conducted on data in each subcategory. Few differences were seen, so we analyzed theme totals (collapsing all categories) as well as specific theme vs. general theme comparisons. Below we report F values from the ANOVAs and the associated p values. We adopt the convention of considering p values less than .05 to indicate statistically significant findings.

IV. WHOLE COURSE EVALUATION FINDINGS

We now present analyses of the initial surveys, weekly surveys, and final surveys. The surveys were used in two ways: to track the success of the course, and to track what students thought they were learning about each of the six core themes.

A. Overall Success

Responses indicated that the course kept the students' interest and that the curriculum sequence was effective. Every week students were asked to anonymously rate how much they enjoyed the week on a scale of 1 to 5, with 5 being the highest. All weeks except for the fifth week were given a mean rating of 4 or above. Ratings for the fifth week, which was the week when autonomous navigation was presented, averaged 3.4. Consistent with the overall ratings of enjoyment, students found the contests and challenges to be increasingly motivating and engaging. On each weekly survey, students were asked whether the challenges and contests for that week were their favorite so far in the course. At least 33% of the students each week reported that it had been their favorite week thus far. As the course progressed, students consistently reported high mean levels of learning each week (3.7 and above).

On the final survey student responses also suggested that they had been engaged appropriately by the overall course experience. Students rated instructor effectiveness at a mean of 4.9 on a 5-point scale. Students thought the pacing of the course had been appropriate, rating pacing at 3.6 on a scale from 1 ("Too Slow") to 5 ("Too Fast"). The guest speakers were appreciated (4.7 out of 5) with every student agreeing that speakers should be included if the course is taught again.

B. Learning the Core Themes

We first asked the question of how students' understanding of their own learning changed from the beginning to the end of the course. Students' expectations for their learning of each of the six themes were coded from their responses to the initial survey question: *What do you expect to learn in this course?* On the final survey, students understanding of their learning of each of the themes was coded from their responses to a question that asked them to list the main things they had learned in the course.

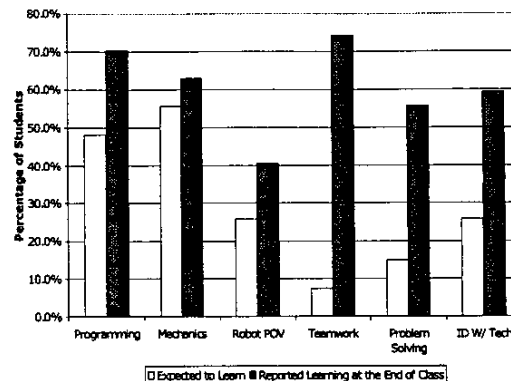


Figure 2. Student self-reports of learning opportunities and results for each of the core themes in the course.

As shown in Fig. 2, students developed different ideas about learning opportunities from the beginning to the end of the course. First, consider what students reported about the three themes that are most specific to the technical aspects of robotics. At the beginning of the course, 56% of students expected to learn about Mechanics while, at the conclusion of the course, 63% reported Mechanics as one of the important things they learned. Similarly; 48% of students expected to learn about Programming and 70% reported that they had, in fact, done so. These findings do not strike us as remarkable; after all, a course about autonomous robots would certainly include the mechanical and programming aspects common to all robotics.

What are more interesting are the larger differences seen in self-reported learning of Teamwork, Problem Solving, and ID with Technology. While 7% of the students initially expected to learn about Teamwork, that theme turned out to be the most commonly reported learning outcome at the end—74% of the students listed it as something they had learned. Similarly, Problem Solving and ID with Technology were commonly reported as learning outcomes at the conclusion of the course, although they had been infrequently mentioned as possible outcomes at the beginning. These findings suggest that the course was successful at meeting the deeper goals of developing domain-general interest and skills that would prepare students for success in broader technology and science education in college.

In addition to coding whether students mentioned learning opportunities for each of the themes, we coded relevant questions from the initial and final surveys to track how much specific detail students reported when they described learning opportunities around specific themes. Although students mentioned Mechanics and Programming a similar number of times in the initial and final surveys, they provided significantly more specifics about each theme on the final survey. For instance, while students mentioned vague statements about "robot technology" on the initial survey, they were more likely to mention specific technologies such as "IR sensors" or "back-EMF" on the final survey, $F(1, 52) = 5.47, p < 0.05$. While they mentioned "learning to program" on the initial survey, they were more likely to talk about "states in programming" or "Java" on the final survey, $F(1, 52) = 8.61, p$

< 0.01. Thus, student descriptions of their own learning became more specific and grounded in the curriculum content.

How students talked about the themes of Teamwork and Problem Solving also changed to include more specifics by the end of the course. Students originally said they would “learn teamwork” or “work in teams of three”. In the final surveys comments like “Teamwork is hard especially with varying levels of skill and different personalities, [it] can be rewarding only through compromise” and “teamwork leads to victory” were more common, $F(1, 52) = 15.91, p < 0.001$. Similarly, from a few general statements about “learning how to solve problems” on the initial survey, student statements changed to specific observations such as learning to “really pay attention to what I am doing and try to solve it first before asking for help”, $F(1, 52) = 12.00, p < 0.001$.

We now turn to an analysis of the weekly surveys students completed each Monday. Two of the key questions on the survey asked students to reflect on whether they had, in the preceding week, made a breakthrough or discovery and whether they had struggled to understand anything. Responses for all weeks and students were summed for analysis. There was a possibility for 162 responses to each question, but not every student reported a struggle and breakthrough every week. For all six surveys given there were 51 reported struggles, between five and thirteen per week, and 87 breakthroughs, between nine and seventeen per week.

As shown in Fig. 3, student struggles were mostly around two themes: Programming and Mechanics. This is not surprising, because those topics are most directly tied to the challenges. Typical responses are shown below.

Our program had a bug which turned out to be a missing zero.

There were long time delays between commands.

Robots need to be tested in the same conditions as where they will perform.

In contrast, student breakthroughs occurred widely among the six themes. Mechanics and Programming were still mentioned most often, but breakthroughs coded as involving Teamwork, Problem Solving, Robot POV, and ID with Technology were also common (see Fig. 4). Representative responses follow.

Programming: New programming languages are easier to understand than I thought.

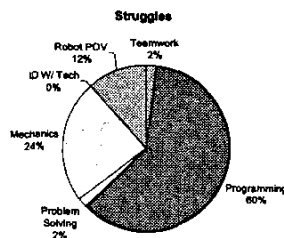


Figure 3. Percent of reported struggles by learning theme.

Mechanics: Understanding how sensors are so wonderful and yet so error prone.

Teamwork: The big discovery was that if I try hard, by working with my teammates, we could make a lot of things happen.

Problem Solving: Don't ever leave anything at the end or else you will be struggling to finish it on time.

Robot POV: Robots are babies.

ID with Technology: I made the discovery that building a robot could be very exciting instead of hard.

Finally, we analyzed the student self-report data for potential gender differences. Although we began the project with no particular expectations that girls and boys would have different experiences, we were sensitive to the historical problem that computer science has had in attracting girls to engage in advanced study. We were also acutely aware of the fact that the majority of the students were boys, all of the outside speakers were men, and that the instructor and all but one teaching assistant were men. As the robot course was one of the first intensive advanced technology experiences for most of the students, we were aware that it had the potential to work against or in support of existing stereotypes regarding girls and technology. Thus, we were particularly interested in whether the experience was successful and positive for the eight girls enrolled.

For most of our findings, there were no statistically significant differences between girls (sample size 8) and boys (sample size 20), suggesting that the course provided a supportive and interesting environment for both. We did observe three differences. First, on the weekly surveys girls were more likely to report having struggled with Programming, $F(1, 25) = 9.12, p < 0.01$. Second, girls also entered the class reporting less confidence with technology than boys, $F(1, 25) = 9.72, p < 0.01$. Third, girls' confidence with technology increased more than boys' by the end of the course, $F(1, 25) = 14.58, p < 0.001$. Thus, despite our initial concern, the course appeared to welcome and support the participation of girls.

In summary, findings on student reported learning suggest that the course was successful in meeting its specific instructional goals of teaching the technology of autonomy and also its general goals of supporting meaningful student engagement with technology to build general interest, skills, and confidence that could promote future success with technology education.

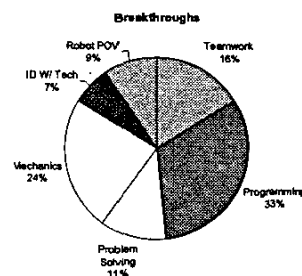


Figure 4. Percent of reported breakthroughs by learning theme.

V. CONCLUSIONS

The goal of this educational assessment was to characterize the impact of a hands-on robotics course using formal techniques. Our prior experiences with robotics education suggested that relatively broad forms of learning may be demonstrable, and this hypothesis has been validated. Learning about the coded themes of Mechanics and Programming is to be expected in a robotics course. Quantitative results based on self-report supported this expectation. More surprising were large jumps from expectation to reported learning along the themes of Problem Solving, Teamwork and ID with Technology. This suggests that a robotics course was able to meet deeper goals of developing domain-general interest and skills that can prepare students for broad success in technology and science education.

Coding for the level of detail in student comments regarding learning themes led to statistically significant increases in specificity. Significant trends were measured for "robot technology," Programming, Teamwork and Problem Solving. These results suggest that students learned concrete lessons for each theme, digging below the surface of abstract concepts to a functional level of detail.

The evaluation of self-reported struggles and breakthroughs supported the above conclusions. Student struggles were reported mainly around two themes: Programming and Mechanics. But, student breakthroughs were reported across a broad range of themes, including Teamwork, Problem Solving, Robot POV and ID with Technology. Once again the inclusion of non-technological themes reported as breakthroughs suggests that, during the course, learning extended beyond the content of technical challenges and into broader scientific and social lessons.

Finally, analysis of student self-report data for gender differences was intended to identify the effect of this advanced technology course on existing stereotypes regarding girls and technology. Thus a critical question would be the degree to which Robotic Autonomy was a positive and successful experience for the girls enrolled. Three statistically significant results summarize conclusions on this query. First, girls were more likely to struggle with Programming. Second, girls entered the course reporting less confidence with technology than boys. But third, girls' confidence in technology increased throughout the course significantly more quickly than the boys'. Thus the course appeared to support the participation of the girls and was able to compensate somewhat for the initial differences between girls' and boys' comfort with technology.

The Robotic Autonomy course lives on, taught again in 2004 by Prof. Mel Siegel, also at Carnegie Mellon West (NASA/Ames, Moffet Field CA). We hope that, as the case for the educational impact of robotics is strengthened by additional research, students in diverse age groups will benefit from robotics curriculum in secondary level education and beyond.

ACKNOWLEDGMENTS

The Robotic Autonomy program was funded in part by NASA/Ames Research, National Hispanic University and Intel Corporation. Thanks to the following individuals who made

significant contributions to Robotic Autonomy: Ajinkya Bhawe, Adriana Cardenas, Daniel Clancy, John D'Ignazio, Maylene Duenas, Thomas Hsiu, Neeti Malhotra, Andres Perez-Bergquist, Raj Reddy and Steve Richards.

REFERENCES

- [1] Beer, R., Chiel, H., Drushel, R. Using autonomous robots to teach science and engineering. *Communications of the ACM*, June 1999.
- [2] Billard, A. Robota, clever toy and educational tool, Special Issue on Socially Interactive Robots, Robotics and Autonomous Systems, 2003.
- [3] Billard, A. and Hayes, G. Learning to communicate through imitation in autonomous robots, in: Proceedings of the 7th International Conference on Artificial Neural Networks, 1997.
- [4] Cooper, M. et al. Robots in the classroom—tools for accessible education, in: Proceedings of the 5th European Conference for the Advancement of Assistive Technology, 1999.
- [5] Coppin, P. et al. EventScope: a telepresence interface for internet-based education, in: Proceedings of the Workshop on Telepresence for Education, IEEE International Conference on Robotics and Automation, 2002.
- [6] Coppin, P. et al. Big Signal: information interaction for public telerobotic exploration, in: Proceedings of the Workshop on Current Challenges in Internet Robotics, IEEE International Conference on Robotics and Automation, 1999.
- [7] Druin, A. and Hendler, J. Robots for kids: exploring new technologies for learning, The Morgan Kaufmann Series in Interactive Technologies, Morgan Kaufmann, 2000.
- [8] Falcone, E., Gockley, R., Porter, E. and Nourbakhsh, I. The personal rover project, Special Issue on Socially Interactive Robots, Robotics and Autonomous Systems, 2003.
- [9] Fong, T., Nourbakhsh, I. and Dautenhahn, K. A survey of socially interactive robots. *Robotics and Autonomous Systems*, Special issue on Socially Interactive Robots 42 (3-4), pp 143-166, 2003.
- [10] Fong, T., Nourbakhsh, I. and Dautenhahn, K. A survey of socially interactive robots: Concepts, Design and Applications. Technical Report CMU-RI-TR-02-29, Carnegie Mellon University, Pittsburgh, PA USA, December 2002.
- [11] Hsiu, T., Richards, S., Bhawe, A., Perez-Bergquist, A. and Nourbakhsh, I. Designing a Low-cost, Expressive Educational Robot. In *Proceedings of IROS 2003*. Las Vegas, USA, 2003.
- [12] Kumar, D. & Meeden, L. A robot laboratory for teaching artificial intelligence. In *Proc. of 29th SIGCSE Symposium on Computer Science Education*, 1998.
- [13] Maxwell, B. and Meeden, L. Integrating Robotics Research with Undergraduate Education, IEEE Intelligent Systems November/December 2000.
- [14] Murphy, R. *Introduction to AI Robotics*. MIT Press, 2000.
- [15] Nourbakhsh, I., Crowley, K., Wilkinson, K., Hamner, E. The Educational Impact of the Robotic Autonomy Mobile Robotics Course. CMU Technical Report CMU-RI-TR-03-29. August 2003.
- [16] Nourbakhsh, I. When students meet robots. *Essay in IEEE Intelligent Systems and Their Applications*, 15(6), p15. 2000.
- [17] Nourbakhsh, I. Robotics and education in the classroom and in the museum: On the study of robots, and robots for study. In *Proceedings Workshop for Personal Robotics for Education*. IEEE ICRA 2000.
- [18] Papert, S. and Harel, I. *Situating Constructionism*, in: Constructionism, Ablex Publishing Corp., 1991.
- [19] RASC 2003. Web reference: <http://www.cs.cmu.edu/~rasc>
- [20] Stein, C. Botball: Autonomous students engineering autonomous robots, in: Proceedings of the ASEE Conference, 2002.
- [21] Wolz, U. Teaching design and project management with Lego RCX robots. In *Proc. SIGCSE Conference 2000*.