Report on the AAAI 2010 Robot Exhibition

Monica D. Anderson, Sonia Chernova, Zachary Dodds, Andrea L. Thomaz, and David S. Touretzky

■ *The 19th robotics program at the annual AAAI conference was held in Atlanta, Georgia, in July 2010. In this article we give a summary of three components of the exhibition: the Small-Scale Manipulation Challenge: Robotic Chess; the Learning by Demonstration Challenge; and the Education Track. In each section we detail the challenge task. We also describe the participating teams, highlight the research questions they tackled, and briefly describe the systems they demonstrated.*

This year marked the 19th edition of the Robotics Program at AAAI, cochaired by Monica D. Anderson and Andrea L. Thomaz. The program has a long tradition of demonstrating innovative research at the intersection of robotics and artificial intelligence.

In both the workshop and exhibition portions of the event, we strive to have the robotics program be a venue that pushes the science of embodied AI forward. Over the past few years, a central point of the event has been the discussion of common robot platforms and software, with the primary goal of focusing the research community's energy toward common "challenge" tasks.

This year, the Robotics Exhibition included two such robotics challenge problems: manipulation and learning by demonstration. In the Small-Scale Manipulation Challenge four teams demonstrated systems playing robotic chess. This exhibit was organized by David Touretzky and Monica D. Anderson. In the Learning by Demonstration Challenge, three teams demonstrated systems learning a block-sorting task. This exhibit was organized by Sonia Chernova. Additionally, this year marked another successful turnout for the Robotics Education Track, organized by Zachary Dodds, which highlights student- and educator-led robotics projects.

In this article we give a summary of these three components of the exhibition. In each section we detail the challenge task. Then we describe the participating teams, highlighting the research questions they tackled and resulting systems they demonstrated.

 ISSN 0738-4602

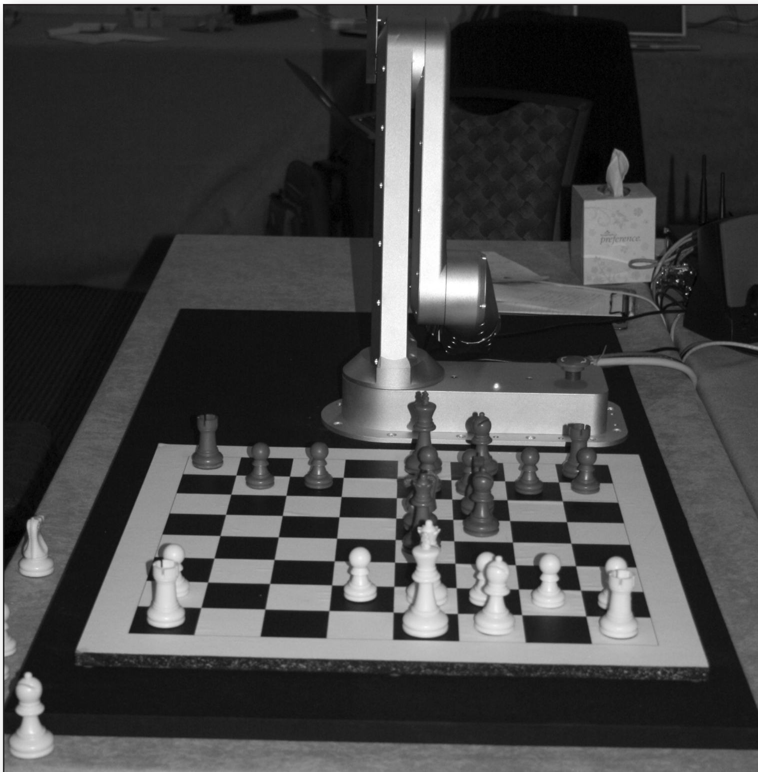

Figure 1. Gambit, the Intel/University of Washington Entry, and Winner of the Small-Scale Manipulation Chess Challenge.

Small Scale Manipulation Challenge: Robotic Chess

The 2010 Small Scale Manipulation Challenge was inspired by the IJCAI 2009 Mobile Manipulation Challenge organized by Willow Garage,[1] where human-scale robots competed on household tasks. Loading a dishwasher is an excellent problem for a robot like the Willow Garage PR2, but we wanted to demonstrate that smaller-scale robots, including low-cost models intended for educational use, could also exhibit manipulation skills. What we envisioned was mobile robots that run around on a tabletop, and so we sometimes referred to the challenge as "tabletop-scale manipulation," but in the interest of attracting a wide range of participants we opened the competition to all classes of platform. As a result, the four teams that met in Atlanta differed significantly in team makeup, budget, and technical approach. The Intel/University of Washington (UW) and Georgia Institute of Technology (Georgia Tech) teams used a medium or large fixed arm, while Carnegie Mellon University (CMU) and the University of Alabama used tabletop-scale mobile robots with small arms. All were able to compete successfully, but no one escaped incidents of hardware or software failure.

The task was to play chess on a real board using standard tournament-style (Staunton) chess pieces, recognizing the opponent's moves through computer vision. The Intel/UW and Carnegie Mellon teams opted to use blue and yellow pieces, which made color segmentation easier; the Georgia Tech and University of Alabama teams used more traditional tan and black or brown pieces. The board could be either standard (2.25 inch squares) or slightly larger than standard (2.75 inch squares). Only Carnegie Mellon chose to use the larger size board. Because of the differences in pieces and boards, and because two of the robots were not easily repositionable, each robot played on its own board, with human judges copying the move onto the opponent robot's board.

In the interest of time, games were limited to 10 moves per side. Since we wanted to emphasize dexterity and accuracy over speed or chess prowess, 150 points were awarded for any legal move, with 25 extra points for captures or castling operations due to their greater complexity. If a move took more than 2 minutes to plan and execute, 10 points were subtracted for each extra minute, with forfeiture of the move after 5 minutes. There was also a 10-point penalty if a piece wasn't placed fully within the boundaries of its square, or if another piece was disturbed and had to be repositioned by a human.

We also defined several bonus tasks that could earn an extra 100 points each either during or outside of an official match. The only three tasks that were successfully completed were audibly announcing the robot's move (all four teams), recognizing and audibly announcing the opponent's moves (three of the four teams), and recognizing and audibly announcing a misplaced piece (one team). The other tasks that were not completed will be discussed later.

With four teams, we decided to have each play the other three, for a total of six official matches, held two at a time. At other times during the two days of the competition there were informal exhibition matches, and the robots also played against human spectators.

Intel/University of Washington

This team, led by Professor Dieter Fox, was composed of graduate and undergraduate students from the University of Washington and researchers and engineers from Intel Labs Seattle. The team members were Cynthia Matuszek, Brian Mayton, Louis Legrand, Robert Chu, Michael Kung, Liefeng Bo, Marc Deisenroth, and Joshua Smith. In collaboration with Roberto Aimi of Alium Labs, they developed a custom 6 degrees of freedom (DOF) arm called Gambit, (figure 1) which featured a spherical wrist, parallel-jaw gripper, and two cameras. Instead of rigid opposing surfaces, the gripper paddles consisted of an open frame covered in a compliant rubber material. This "opposing trampoline" structure proved to be a real advance in reliable chess piece manipulation.

The initial prototype of Gambit was constructed from three-dimensional printed plastic, but the final version demonstrated in Atlanta had a Computer Numerical Control (CNC)–milled aluminum housing. Gambit's shoulder-mounted depth camera looked down on the board and used a combination of color and depth information to track the board pose and game state in real time. Gambit's second camera, a tiny image sensor integrated into its gripper, provided close up, overhead images of chess pieces for piece recognition and visual servoing. Unfortunately, a hardware problem prevented demonstration of this capability during the competition.

Georgia Institute of Technology

The "Golem Chesster" team consisted of graduate students Hun-Soo Yi, Evan Seguin, Pushkar Kolhe, and Tobias Kunz, advised by Professor Mike Stilman. They employed a large, modular commercial arm, the 7-DOF Schunk LWA3, and a 7-DOF Schunk SDH2 dextrous hand with pressure sensor arrays in the fingers. A time-of-flight sensor was mounted directly above the board for an unobstructed view of the pieces. In contrast to a camera, the Swiss Ranger (SR4000) reported direct depth information by actively supplying infrared light and measuring the phase difference in the reflected signal.

Golem Chesster (figure 2) faced a completely different set of challenges from other teams because the Schunk arm was designed for larger and heavier manipulation tasks than chess. While the translational workspace of the arm was greater than the other robots, the distance between the wrist and the finger-tips required the LWA3 to completely change the configuration of its joints to achieve correct orientations with respect to pieces. In fact, at the edges of the board, the robot could only reach a subset of orientations, which meant that pieces had to be grasped at an angle. The Georgia Tech controls team handled these challenges by developing algorithms that maximized the workspace of the arm during motion and ensured proper repositioning despite grasp angle.

Schunk has recently designed a new PowerBall module that combines the last two wrist joints into a single device. This significantly shortens the finger-tip distance, leading to a greater workspace for orientation. The new module will be featured in future implementations of Golem Chesster.

Carnegie Mellon University

The Carnegie Mellon University entry was developed by Jonathan Coens as his computer science masters thesis. Jonathan used a hexapod robot called the Chiara, created in CMU's Tekkotsu Lab by his advisor, Professor David Touretzky, with a custom gripper for manipulating chess pieces. Ethan Tira-Thompson, the principal architect of the Tekkotsu software framework, also provided some technical assistance. Atlanta was the debut of the new delta series Chiara, which uses more powerful leg servos than the previous model and a faster processor based on an Intel Atom chip.

Figure 2. Golem Chesster, the Georgia Tech entry, and Second Place Finisher in the Small-Scale Manipulation Chess Challenge.

Vision for the Chiara was more challenging than for the other competitors because of the team's insistence on an "organic" solution in keeping with the Chiara's insectlike appearance. Instead of mounting the camera up high, it was mounted on a neck proportionate in size to the robot's small body. As a consequence, all views of the board were partially occluded due to the density of pieces and the camera's shallow viewing angle. Coens invested considerable effort in algorithms for compensating for occlusions by shifting the robot's body to take advantage of parallax, combining information from multiple views, and using chess knowledge to resolve ambiguities (Coens 2010).

Manipulation was also a challenge because the Chiara has a planar arm; translation along the z-axis is accomplished using the legs. To pick up a piece, the robot would use its legs to raise up its body so that the arm was above the plane of the pieces. It would then extend the arm until the gripper was above the target piece, use the wrist pitch servo to swing the fingers down on either side of the piece, and then close the foam-covered fingers to compliantly grasp the piece.

The short length of the arm required the Chiara to walk sideways to position itself properly for advancing pieces and removing captured material. Also, competition rules required robots to withdraw from the board when not moving so as not to obstruct an opponent, so the Chiara backed away from the board after each move and walked up to it when it was its turn to move again.

University of Alabama

The University of Alabama featured the most unusual entry of the competition, first, because the team was composed entirely of undergraduates: Revarr Johnson, Daniel Thompson, Robert Phares, Daniel Gay, Tunch Tosoli, and Trey Davis, with a little help from graduate student Andrew McKenzie, all advised by Professor Monica D. Anderson. The second unusual feature was the shoestring budget: under $1000 total cost. (At the other extreme, Georgia Tech's Schunk arm and Intel/UW's custom arm both cost in the range of $100,000.)

Alabama built its robot from off-the-shelf components: an iRobot Create mobile base, a Lynxmotion AL5D hobbyist arm, and a wooden yardarm on which was mounted a Zii EGG handheld computer running the Android operating system. The Zii EGG, about the size of a smartphone, was positioned so it could look down on the board, but it was not just an eye: it was the brain of the robot.

The Lynxmotion arm was modified to have sufficient reach to cover the entire board, but the robot still needed the mobile base for two reasons. First, it had to remove the camera from the airspace above the board when it was the opponent's turn to move and bring it back into position when it was time to examine the board again. And second, it needed to back off from the board slightly in order for the arm to be able to reach the closest row of pieces without interference from the body. At the start of a match the robot was manually positioned relative to the board. It was then able to rely on the accuracy of the Create's motion to slide backward and forward for each move without relocalizing.

The Challenge

Each robot was able to make moves and play some chess. Everyone also suffered some difficulties: a failed auxiliary camera for Intel/UW, software issues for Georgia Tech, a failed gripper servo for CMU, and wireless problems for Alabama. The Intel/UW team still turned in a nearly flawless performance, earning the top score of 4980 points. Georgia Tech came in second at 4115 points. Carnegie Mellon took a big point hit when a gripper servo failed in the middle of the first match, putting it in third place overall at 3580 points. Alabama came in fourth, at a respectable 3270 points, and will no doubt inspire more undergraduates to enter the next competition.[2]

Vision remains a hard problem. Several teams were concerned about lighting conditions; we had to adjust the room lighting and reposition some robots in order for everyone to be able to play. It's unfortunate that roboticists still have such a hard time coping with glare; perhaps we should be talking more with our computer vision colleagues.

All of the teams recognized moves by detecting changes in occupancy of board squares, not by recognizing the shapes of pieces. Shape recognition from a video camera view of a chess board is much more difficult than detecting square occupancy due to the low resolution of the images and the subtlety of differences between some piece shapes. However, Intel/UW did attempt this problem, applying machine-learning techniques to high-dimensional feature vectors extracted from overhead images of pieces taken from the gripper camera.

We formulated the challenge as an open-ended event with room for teams to distinguish themselves by taking on harder tasks. One such task was entering a game in an arbitrary starting configuration. As this requires recognition of piece shapes, it was not attempted this year. Another task was to recognize when a piece has been knocked down and right it. This would require both enhanced perception and a manipulation algorithm capable of reorienting a piece from horizontal to vertical. A third task was for the robot to unpack and set up the chess game itself, which would combine the previous two piece-centered tasks with the problem of unrolling and positioning the plastic mat that depicts the board. A fourth, much harder task would be to play using nonstandard "artistic" pieces whose shapes would have to be learned at the start of the game.

The event generated a lot of enthusiasm among the competitors and those conference attendees who stopped by to observe. Since there is still much progress to be made in this area, the challenge will be repeated at AAAI 2011. Several modifications are being considered. One is to settle on standard board and piece colors so that robots can play on the same board. Another is to have separate divisions for large fixed arms vs. tabletop scale mobile robots, since the latter must solve localization and navigation problems in addition to board perception and piece manipulation.

Learning by Demonstration Challenge

A second event was the Learning by Demonstration Challenge. Learning by demonstration (LBD) research focuses on the development of algorithms that enable humans to teach robots new tasks by showing the robot what to do instead of by programming. Compared to other supervised learning methods being developed in the AI community, LBD assumes little or no robotics knowledge on behalf of the user and aims to provide a way for lay people to customize and develop robot behaviors.

The Learning by Demonstration Challenge, held in conjunction with the AAAI Robotics Exhibition, is the premier forum for the presentation of advanced demonstration learning systems. The challenge is designed as an annual, multiday event that enables participants to showcase their latest findings through presentations and live interactive demos.

LBD algorithms incorporate a broad range of learning techniques, including autonomous learning through exploration, statistical inference, Bayesian learning, regression, and planning. Demonstration techniques also vary greatly based on the method of interaction between the human and the robot and on the content of the information exchanged. Example techniques include teleoperation, verbal commands, human demonstrations recorded through motion capture, and kinesthetic teaching. Differences in the physical embodiment of different robotic platforms further add to the diversity of approaches currently being explored in this research field.

The broad aim of the LBD Challenge is to promote technological innovation in this research area by facilitating discussion of the latest developments, providing a venue for showcasing cutting-edge research, and encouraging comparative assessment through a series of organized challenges. The LBD exhibit has taken place in two AAAI/IJCAI events prior to this one and has consisted of an open demonstration for researchers to showcase the latest research of their choice. Toward the goal of comparative assessment, this year's event consisted of two components: an open demonstration and an optional challenge event. The 2010 LBD Challenge included participants from 6 institutions showcasing 10 different projects.

Open Demonstration Portion

Three teams participated in the open-demonstration portion of the LBD exhibit: BBN Technologies, Carnegie Mellon University, and Rutgers University.

Jacob Beal, Alice Leung, and Robert Laddaga from BBN Technologies presented a demonstration of spectrum curricula learning, which enables a human teacher to advance the capabilities of engineered systems through a series of lessons in an agent-based instructional framework. The researchers presented seven different curricula that allowed visitors to train a different set of soccer skills in a three versus two version of RoboCup soccer keepaway.

Cetin Mericli and Manuela Veloso from Carnegie Mellon University demonstrated learning a biped walk behavior from real-time human feedback using the Aldebaran Nao humanoid robot. The researchers used a Nintendo Wii Remote (Wiimote) wireless game controller to modify the joint angles of the robot in real time, resulting in significant improvements in robot stability.

Kaushik Subramanian and Michael Littman from Rutgers University presented a generalized apprenticeship learning system for conjunctive learning tasks using a Lego Mindstorms robot. The researchers' demo allowed visitors to teach the robot to locate, pick up, and deliver an object from one randomly selected location to another.

The Challenge Portion

Given the diversity of approaches and problems that people are tackling in the LBD field, it is difficult to pose a challenge task that is relevant to a broad section of the community. This year, as our first attempt to do so, we provided a common set of objects (various sizes of colorful foam blocks) for the task learning domain. These objects were known ahead of time, and the blocks were the only aspect of standardization in the challenge. In coming years the challenge will move toward a standard task that all teams focus on learning. Teams from Georgia Institute of Technology, Brown University, and Ecole Polytechnique Fédérale de Lausanne (EPFL) chose to participate in this LBD challenge portion of the event.

The Georgia Tech team was headed by professor Andrea L. Thomaz and included graduate students Crystal Chao, Michael Geilniak, and Jaewook Yoo. They presented an interactive learning demo using their humanoid robot, Simon (figure 3). The hardware is a 24 degree-of-freedom upper torso by Meka Robotics with a 13-DOF custom head designed by Carla Diana and built by Meka. Simon perceived the color of the foam blocks with its eye cameras, and categorized blocks as large or small by the finger position needed for grasping. It used two microphones for ambient sound localization, and a microphone for speech input from the human teacher. Simon's software is developed in the group's Creatures codebase (C6), which originates from Bruce Blumberg's cognitive architecture for software learning agents.

The Georgia Tech team focused on goal learning in an active learning setting. A person issues spo-

Figure 3. Simon, the Georgia Tech Robot, Sorting Colored Blocks for the Learning by Demonstration Challenge.

ken commands to give Simon examples of what kind of objects go where. They physically hand Simon an object and then tell the robot where to put it. Additionally, Simon could look at the objects on the table, select one that would maximally improve its model of the task, and ask the human, "What do we do with this one?" The Simon robot learned several sorting tasks from humans throughout the week, learning to put all the red blocks in small bins, or that green blocks could go anywhere, or that small yellow blocks go in the big blue bin.

The team from Brown University, headed by professor Odest Chadwicke Jenkins, included Sarah Osentoski, Graylin Jay, Christopher Crick, Jesse Butterfield, Gal Peleg, and Sam Pucci. The Brown entry focused on creating a community resource for shared robotics experimentation and large-scale learning from demonstration (figure 4). Motivated to create interfaces that allow nonroboticists and even nonprogrammers to shape robot behavior, the team recognizes that current robot research has been limited by the prevalence of "one-off" solutions focused on specialized and proprietary software and often exotic hardware. Consequently, there has been a lack of critical scientific independent verification through reproduction and comparison of results. Instead of building on previous work, the tendency is to reinvent existing functionality and advocate research ideas using proof-of-concept demonstrations. A community resource for shared robotic experimentation will allow researchers to compare to and build off of previous research.

To this end, their demonstration showcased a selection of freely available tools, provided by the Brown robotics lab (the brown-ros-pkg). Their goal is to create in an environment in which users can teleoperate robot platforms through video-game-style web-based interfaces and run their own custom controllers. The tools are built upon ROS, Willow Garage's robot middleware system. They have provided a lightweight JavaScript binding for ROS called rosjs, which allows robot applications to be built within the web browser, providing a way to create functionality quickly on a multitude of platforms. Visualization tools are provided through the gscam node, as well as ar recog, which wraps the ARToolkit. The use of such augmented reality tags has been helpful for localization in navigation, teleoperation of humanoid robots, and people following. All of these tools can be used with a robot running ROS. The brown-ros-pkg provides drivers for two off-the-shelf platforms: the iRobot Create and the Aldebaran Nao.

The team from EPFL, headed by professor Aude Billard, included Brenna Argall, Florent D'Halluin, Dan Grollman, S. Mohammad Khansari Zadeh, Basilio Noris, and Eric Sauser. Their demonstration consisted of multiple humanoid robot platforms performing a variety of tasks. The first platform is the iCub robot, a child-sized 57 DOF humanoid, and the second is the 28 DOF Hoap3 robot from Fujitsu. Since their robots are located in Switzerland, they participated through video demonstrations.

As a demonstration of research with stable dynamic systems, EPFL showed the iCub robot performing a directed object sorting task. Specifically, the robot uses its hand to tap a ball such that its motion is directed toward a goal position and enters a basket located at the goal. The stability of the system is shown by the ability to achieve the goal from multiple starting positions. To demonstrate their work on combining multiple reference frames within dynamical systems, they showed the Hoap3 robot mimicking feeding an infant, by spooning food into a doll's mouth. The two reference frames in this case are the bowl containing the food and the mouth of the doll. Additionally, their work with policy adaptation through tactile corrections was shown on the iCub robot performing grasp positioning tasks. The demonstrations indicate how to position its end effector to grasp a single object at a particular location. Tactile corrections are then used to refine this behavior, as well as to adapt the behavior to accomplish other tasks, in particular, tasks that grasp the object at a different location and that grasp a novel object, which requires a different hand orientation.

Moving forward, one of the key long-term goals of the LBD Challenge event is to facilitate a more direct comparison between approaches through the introduction of a standard domain. However, the diversity of constraints imposed by the different robotic platforms makes it impossible to define a single task that all teams are able to perform. To address this challenge, we are currently working toward developing a remote open-access robotics facility that will enable 2011 challenge participants to interactively train a remote robot in a standard domain, and the details of this exciting 2011 LBD Challenge can be found at www.lfd-challenge.org.

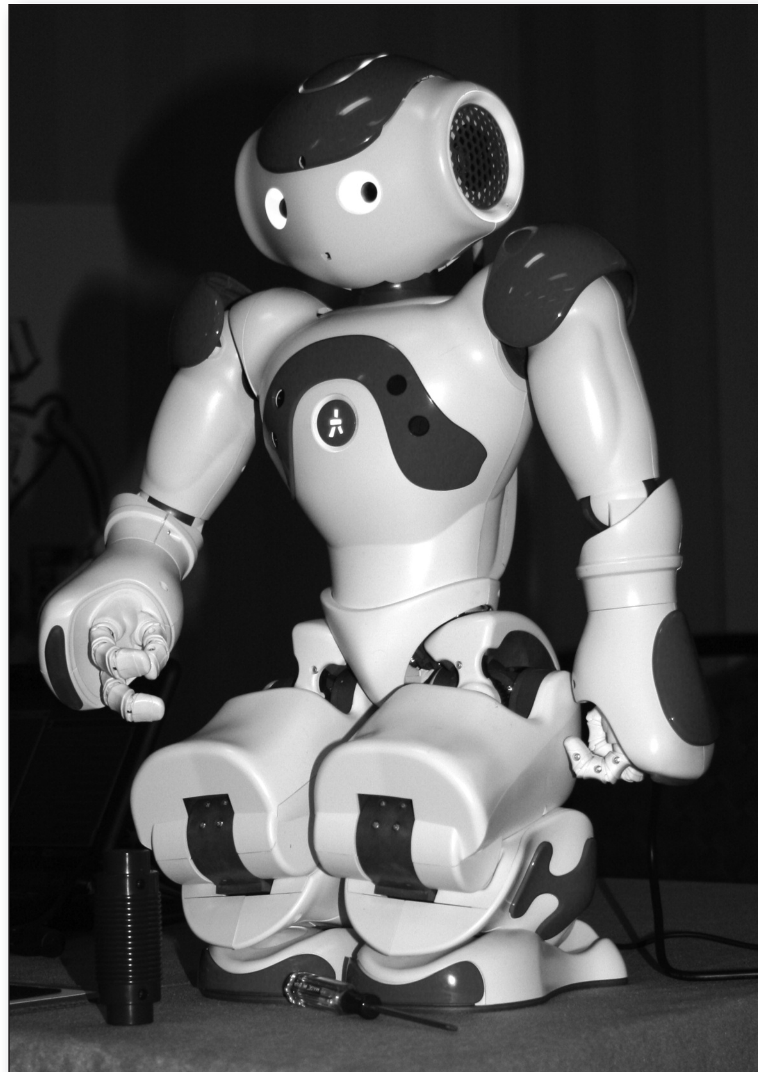

Figure 4. The Brown University Team's Demonstration of their Robotics Tools ROS Package on the Aldebaran Nao Platform.

AAAI 2010 Robot Exhibition: Education Track

Robots have long been a crucial foundation for AI research, but in recent years they have also become increasingly important resources for undergraduate and precollege education. For several years now, the education track of the AAAI Robotics Program has highlighted student- and educator-led robotic projects that integrate AI research and AI education, especially undergraduate education; broaden participation in AI and AAAI; and showcase exemplary hardware, curricular, and software resources for teaching AI and AI robotics. Participation in the 2010 education track more than doubled from 2009: a total of nine exhibitions from 13 schools participated in Atlanta.

Creative Use of Robotics in the Classroom

Educators teaching with physical platforms face a fundamental trade-off between the power and the accessibility of the software interfaces they provide. We had several exhibits highlighting the creative use of robotics in the classroom, from the University of Alabama, Rose-Hulman Institute of Technology, Duke University, and the City University, New York.

The first two examples featured novel resources that make students' interaction "as simple as possible, but not simpler." For example, PREOP, or Providing Robotic Experiences through Object-Based Programming, was exhibited by the University of Alabama's Monica D. Anderson and her students. PREOP builds atop the graphical story-telling environment of Alice, which has become a popular interface for introducing students to programming. By replacing the stock avatars with robots such as iRobot Creates or Lego NXT constructions, PREOP invites students to explore physically embodied computation at the same time—or even before—they learn to program.

Matt Boutell and student Andrew Hettlinger (Rose-Hulman Institute of Technology) exhibited a second interface that promotes student learning of robotics and computer science. Their software acts as a drop-in replacement for an iRobot Create; a visualizer then provides a beautifully and realistically rendered three-dimensional model of the simulation's ongoing behavior. As a result, students can prototype—and run initial tests—of Python-based robot programs even when away from the platforms themselves. In contrast to many robot simulators, Rose Hulman's is standalone: installing and running it require no support libraries and less than a minute of time. Leveraging Creates and this simulator, Rose Hulman's CSSE 120 course teaches software development to first-year students through robotics tasks.

Younger students were the target of another innovative curriculum at the exhibition. Duke graduate students Mac Mason and Gavin Taylor presented an Intensive Introductory Robotics Course Without Prerequisites, an ambitious, three-week AI course for seventh graders in Duke's Talent Identification Program (TIP). Mason and Taylor teach fundamental AI robotics topics including localization, image processing, and multiagent coordination through web cameras and laptops atop iRobot Creates. Their fast-paced approach has succeeded for three years running. What's more, Mason and Taylor have honed the materials to the point where others have successfully adapted and deployed their work, a key test of curricular maturity.

Similar AI topics lay at the heart of the exhibit from the City University of New York. Elizabeth Sklar, Simon Parsons, and students from CUNY's Research Experiences for Undergraduates showed a unified hardware and software framework for exploring shared-autonomy systems. Their scaffolding leverages Player/Stage software and four distinct hardware platforms: the Sony AIBO, Lego NXT, Surveyor SRV-1, and a Parallax Scribbler. The resulting system enables these robots to handle the low-level details of navigation and object-identification tasks in teams. Human operators collaborate with the robot team by specifying high-level goals, such as an object to be explored or a task to be performed.

Undergraduate Research Results

Several other teams joined CUNY in exhibiting undergraduate research results: over 20 undergraduates attended AAAI as part of the robotics workshop and exhibition. Exhibitors included Florida A&M University, Carnegie Mellon University, Bard College, and Harvey Mudd College.

Florida A&M University's Owen Watson demonstrated his work on Calliope, a new robot being jointly developed by Carnegie Mellon, RoPro Design, and FAMU. Calliope is a mobile manipulation platform that augments an iRobot Create with a netbook, a pan/tilt camera, and an arm with gripper. It's the latest in a family of platforms running the Tekkotsu software framework created by CMU's David Touretzky and Ethan Tira-Thompson. Developed specifically for teaching AI and the computer science side of robotics, Tekkotsu-based robots are open source, but also available commercially from RoPro Design for researchers and educators who don't want to build their own.

Like Calliope, the IMP was a Create-based robot. Short for Intelligent Mobile Projector, the IMP was exhibited by Bard College's Keith O'Hara and students Anis Zaman and Aaron Strauss. The IMP's portable projector and netbook's web camera form a feedback loop that human users may join and guide. In one application, the IMP projected only a part of a large mural onto one of the room's walls. Observers then used laser pointers either to add virtual graffiti to the mural or to indicate which direction IMP should move in order to extend the work. They also demonstrated a mixed-reality interactive game that used the floor in front of the IMP as its playing surface.

A team of students from Harvey Mudd College exhibited another Create-based platform, named PixelLaser. This robot sought to measure the range to obstacles by distinguishing the texture of the traversable area from that of walls, legs, and other barriers. As if in confirmation of the exhibition's pedagogical value, the team managed to speed up its nearest-neighbors classification by an order of magnitude during the event's first evening. On the second day, the robot wandered the hall safely and autonomously using only the netbook's web camera for input.

Although the exhibits noted so far make it clear how popular the iRobot Create has become as a foundation for AI and robotics education, there were several schools demonstrating platforms very different from that ubiquitous standard. ShakeTime (figure 5) embodied the most unusual form factor: it is a fruit-adorned rotating platter with an anthropomorphic bottle opener at its center. Carnegie Mellon Univerity graduate students Marynel Vazquez, Alexander May, and Wei-Hsuan Chen exhibited ShakeTime, whose purpose is to host a reaction-time game. Throughout the exhibition, visitors would join one of ShakeTime's four stations and grab an input buzzer. The robot would start a game by subtly shaking one of its plastic fruits. Then it proceeded to shake fruits in random order, challenging participants to be the first to buzz in when the initial fruit was shaken a second time. Winners would enjoy a grape or other treat; losers had to settle for carrots or broccoli.

ShakeTime's twist was that it didn't always play fair. In order to maintain interest among all participants, the robot occasionally rigged the outcome, allowing late buzzers to win if the system considered them likely to lose interest otherwise. And ShakeTime did generate interest. Many bowls of grapes and carrots were consumed during two days of ShakeTime! Its creators use the game to study human-robot interaction and, specifically, how people's sense of a task's fairness and their self-efficacy affect their willingness to continue.

Figure 5. ShakeTime, an Entry in the Robotics Education Track from a Group of CMU Graduate Students, Explores Human-Robot Interaction through a Robot Initiated Reaction-Time Game.

Exhibitors from the Humanoid Obstacle Run Challenge

The final set of education-track exhibitors comprised several teams within the Humanoid Obstacle Run challenge. This event involved five schools (Bryn Mawr College, Colby College, University of Pennsylvania, Virginia Tech, and Drexel University) whose collaboration is making humanoid robots more broadly accessible as a research and education resource.

Drexel's Paul Oh organized both the Humanoid Obstacle Run and the international collaboration that has brought the Hubo robots (figure 6) to so many U.S. institutions.

Doug Blank and students from Bryn Mawr College showed off their custom minihumanoid, for which they had created software interfaces with which introductory computer science students could program the robot.

Bruce Maxwell and his students from Colby College invited visitors to play "Simon Says" against a simulated version of the human-scale Jaemi Hubo, displayed in the opposite corner of the hall. The Colby team's simulation software serves as a foundation that allows any group to specify tasks to Jaemi Hubo, whether they have access to the physical platform or not.

Daniel Lee of the University of Pennsylvania, Virginia Tech's Dennis Hong, and their students exhibited several mini-Hubos and custom humanoids that could walk through a raised obstacle course of boxes to be avoided and ridges to be stepped over. During several scheduled runs, their robots showed the ability to navigate through the environment nimbly and safely.

That these humanoid challenges were so well attended by AAAI participants attests to the progress that AI and robotics education has enjoyed over the past several years. By the same token, the exhibitors, students, and researchers attending the 2010 robotics exhibition show that today's educational opportunities in physically instantiated AI are more accessible, inviting, and challenging than ever.

Conclusion

As we look back and reflect on the successful exhibits of the 2010 event in Atlanta, a common theme across the three exhibits is excitement about what lies ahead in the AAAI 2011 Robotics

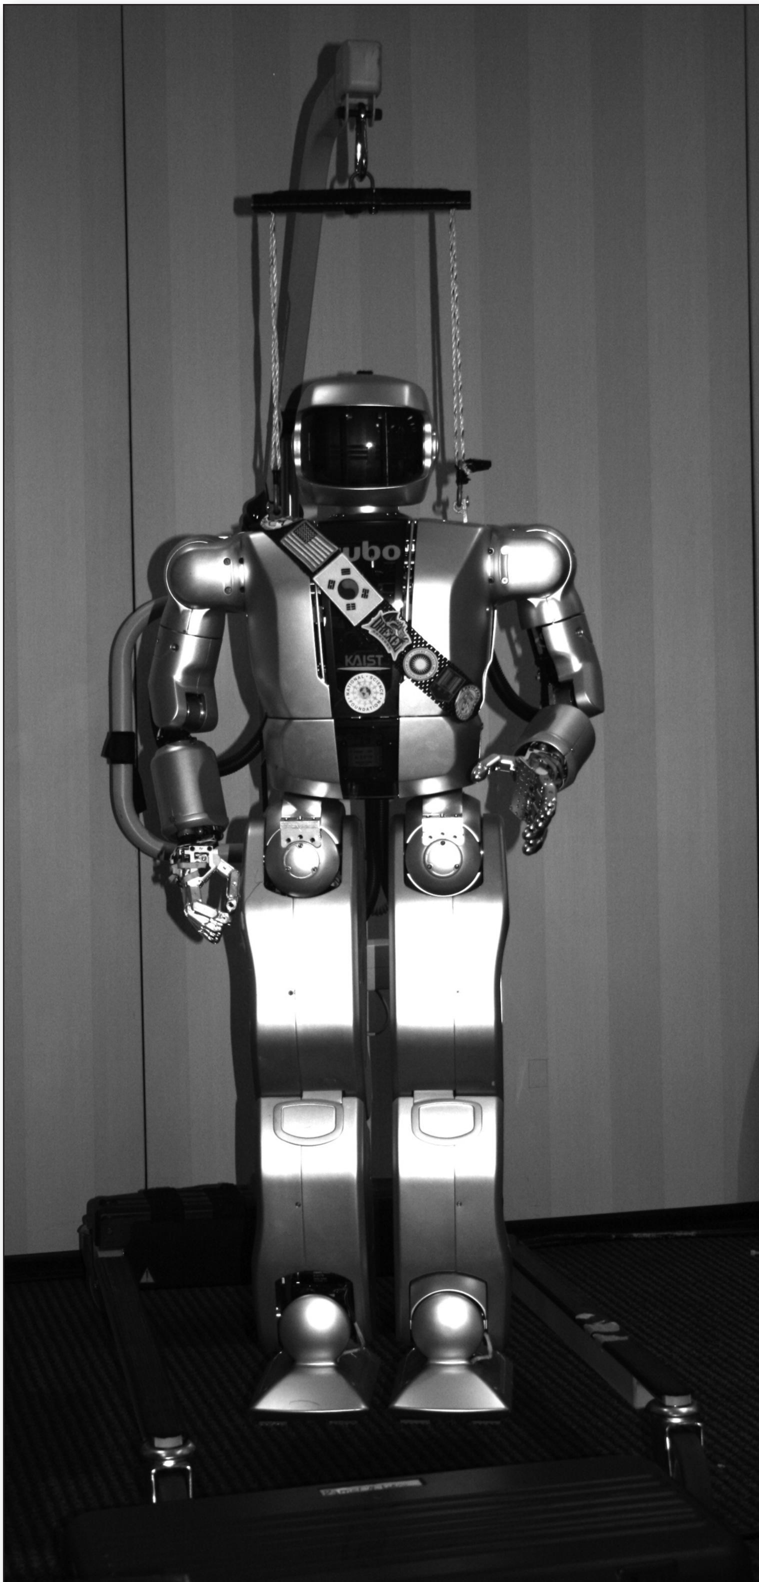

Figure 6. The Hubo Robot, Used by Five Teams that Participated in the Humanoid Obstacle Run Challenge.

Program in San Francisco. This year has been one step on our multiyear vision of the Robotics Program. We had multiple challenges, demonstrating state of the art research at the intersection of AI and robotics. Next year the Learning by Demonstration and Small-Scale Manipulation challenge events will run for their second year. This embodies our vision for the AAAI robotics program, a community of researchers focusing on interesting problems, over multiple years, fostering a continuous and scientific dialog around the field of physically grounded AI.

Acknowledgments

We would like to thank the generous sponsors of the AAAI 2010 Robotics Program: The National Science Foundation, DARPA, iRobot, and Microsoft.

Note

1. A video montage of the four teams can be seen at bit.ly/bqGDws. Additional footage is available at Chiara-Robot.org/Challenge.

2. See the slides prepared by C. Jenkins, "The Future of AAAI Robotics: Challenges and Directions for Building Upward" (www.informatik.uni-freiburg.de/helmert/events/ijcai09-ws/talks/talk-jenkins.pdf).

Reference

Coens, J. A. 2010. Taking Tekkotsu out of the Plane. Masters thesis. Technical report CMU-CS-10-139, Computer Science Department, Carnegie Mellon University. (reports-archive.adm.cs.cmu.edu/anon/2010/CMU-CS-10-139.pdf).

Monica D. Anderson is an assistant professor in the Computer Science Department at The University of Alabama. Her research focuses on multirobot systems and robotics software architectures.

Sonia Chernova is an assistant professor in the Robotics Engineering Program at Worcester Polytechnic Institute. Her research focuses on interactive robot learning, adjustable autonomy, and human-robot interaction.

Zachary Dodds is a professor of computer science at Harvey Mudd College. His interests include vision-based robotics, AI education, and computer-science education more generally.

Andrea L. Thomaz is an assistant professor in the School of Interactive Computing at the Georgia Institute of Technology. She conducts research in the domain of human-robot interaction and interactive machine learning.

David S. Touretzky is a research professor of computer science and codirector of the graduate training program of the Center for the Neural Basis of Cognition at Carnegie Mellon University.