The AAAI 2011 Robot Exhibition

Sonia Chernova, Zachary Dodds, Mike Stilman, David S. Touretzky, Andrea L. Thomaz

■ *In this article we report on the exhibits and challenges shown at the AAAI 2011 Robotics Program in San Francisco. The event included a broad demonstration of innovative research at the intersection of robotics and artificial intelligence. Through these multiyear challenge events, our goal has been to focus the research community's energy toward common platforms and common problems to work toward the greater goal of embodied AI.*

This year marked the 20th edition of the Robotics Program at AAAI, chaired by 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.

On the day before the exhibition the participants convened a workshop of 18 short talks. Each track's exhibitors presented a summary of their exhibit. In addition, four guest speakers provided a broader context for all of the exhibitors' efforts. The first guest speaker was the National Science Foundation's Sven Koenig, who highlighted several federal programs that support projects in embodied intelligence. Koenig also provided insights into some of these program's specific priorities, such as international collaborations and educational engagement. Guest speakers from Willow Garage and Bosch presented cutting-edge work with the PR2, Willow's mobile two-arm manipulator platform.[1] Bosch detailed its Remote Lab, which provides researchers anywhere with full access to the sensing and mobile manipulation capabilities of a PR2. Willow Garage featured some of its most recent work, in which point clouds (Anderson et al. 2011) are parsed not only to build generic three-dimensional scene models but also task-specific structures such as cabinet and drawer handles. Those structures, in turn, seed the automatic creation of task sequences for object retrieval in unconstrained human environments. Nataniel Dukan of Nao Robotics presented the workshop's final guest talk, a broad overview of humanoid robotics's current resources, along with a compelling vision for where those technologies will be in the next three to five years. Without providing specifics of Aldebaran's unannounced plans, Dukan hinted that the actuation and sensing needed for com-

 ISSN 0738-4602

mercial, full-size humanoid helpmates is within reach; the effective coordination of such complex resources, however, might present a longer-term challenge. The workshop concluded with a small-group brainstorming session on the software and hardware resources that they would like to see and will seek to create in the near future, starting from the work that was on exhibit the following days.

This year, all of the events in our program were continuing events, being run for at least the second time (Anderson et al. 2011). The small-scale manipulation challenge, chaired by David Touretzky and Mike Stilman, featured four teams competing in robotic chess. The Learning from Demonstration challenge, chaired by Sonia Chernova, featured five teams demonstrating advances in the field on a common problem domain and platform. The Robotics Education Track, chaired by Zachary Dodds, highlighted projects and platforms that have an impact on AI curricula and education.

In this article we provide a summary of each of these three events. In each case we provide some motivation about the particular AI problems that the challenge or event is designed to address. Then we provide a summary of the teams and their exhibits, and general lessons learned.

Small-Scale Manipulation Challenge: Robotic Chess

The small-scale manipulation challenge, chaired by David Touretzky from Carnegie Mellon and Mike Stilman from Georgia Tech, was once again chess, using the same rules as the 2010 competition. The event hosted three new teams and four new robots, each with distinct software and hardware designs. It culminated with an exciting robot-robot chess match on the same board.

The point of the challenge was to solve real-world perception and manipulation problems that do not require expensive, human-scale robots. Accurately moving pieces on a physical chess board, and recognizing opponent moves so as to play a legal but not necessarily inspiring game, was the order of the day. Points were awarded for making legal moves within the allotted 5 minutes per move, with extra points for more complex moves, that is, a capture or castle. Points were subtracted for errors such as knocking over pieces or failing to place a piece fully within its square. To keep the games short, each was limited to 10 moves per side. Contestants used open source chess engines for move generation.

The four competitors differed greatly in their approaches. The winner, UAlbany Robotics, led by Michael Ferguson of the University at Albany, used a large mobile robot named Maxwell that stood by the table and reached out with an arm to move pieces (figure 1). Maxwell used a PrimeSense image-plus-depth sensor, the same technology used in the Microsoft Kinect, which has become popular for robotics experimentation due to its low cost and widespread availability. The sensor combines color images with point cloud (depth) information that can be helpful in segmenting a scene. Maxwell's sensor was mounted on a pan/tilt "head" and accompanied by a 5 degree-of-freedom arm, both constructed from Robotis Dynamixel servos. On-board computation was handled by a laptop plus an embedded real-time coprocessor for motion tasks. The software was written using ROS, the Robot Operating System from Willow Garage. Maxwell faced significant grasping challenges on the first day and Ferguson hand-crafted a new gripper overnight, returning to dominate the competition.

Hekateros, the entry from Road Narrows Robotics, LLC, was a fixed 5 degree-of-freedom arm built using a three-dimensional printer. The company used this opportunity to demonstrate and evaluate its new inexpensive robot arm in a real-world application. Although, like two of the other entries, motion was provided by Dynamixel servos, rather than the servos driving the joints directly, Hekateros used a belt and pulley system for driving its shoulder and elbow joints. Attached optical encoders provided improved position accuracy over what is available from the servos' internal, potentiometer-based position-sensing mechanism. The end-effector, which was specifically designed for manipulating chess pieces, also had a built-in camera. The software was built on Road Narrows' open source Bot Sense package. The Road Narrows team was led by Colin Horvat and included Kim Wheeler and Robin Knight.

The Griffins, undergraduate students Michael Lanighan and Jerod Sikrskyj from Canisius College, were advised by Professor Debra Burhans. The team built its robot from Lego NXT bricks, programmed in LeJos. The robot was tethered to a laptop that provided high-level control. The robot design took the form of a gantry that slid out over the chessboard and picked up pieces from directly overhead. Like the University of Alabama undergraduate team last year, the Griffins demonstrated that small-scale manipulation can be done effectively on a shoestring. The team came in third.

The fourth competitor, Carnegie Mellon's Tekkotsu Lab, advised by David Touretzky, returned this year with a new design: Calliope. Calliope consists of an iRobot Create mobile base, a 5 degree-of-freedom arm constructed from Dynamixel servos, a webcam on a pan and tilt mount, and a netbook running the Tekkotsu open source software framework.[2] Owen Watson, a summer intern from Florida A&M University, adopted Jonathan Coens' chess playing code for the Chiara robot to the Calliope (Coens 2010).[3] The Chiara

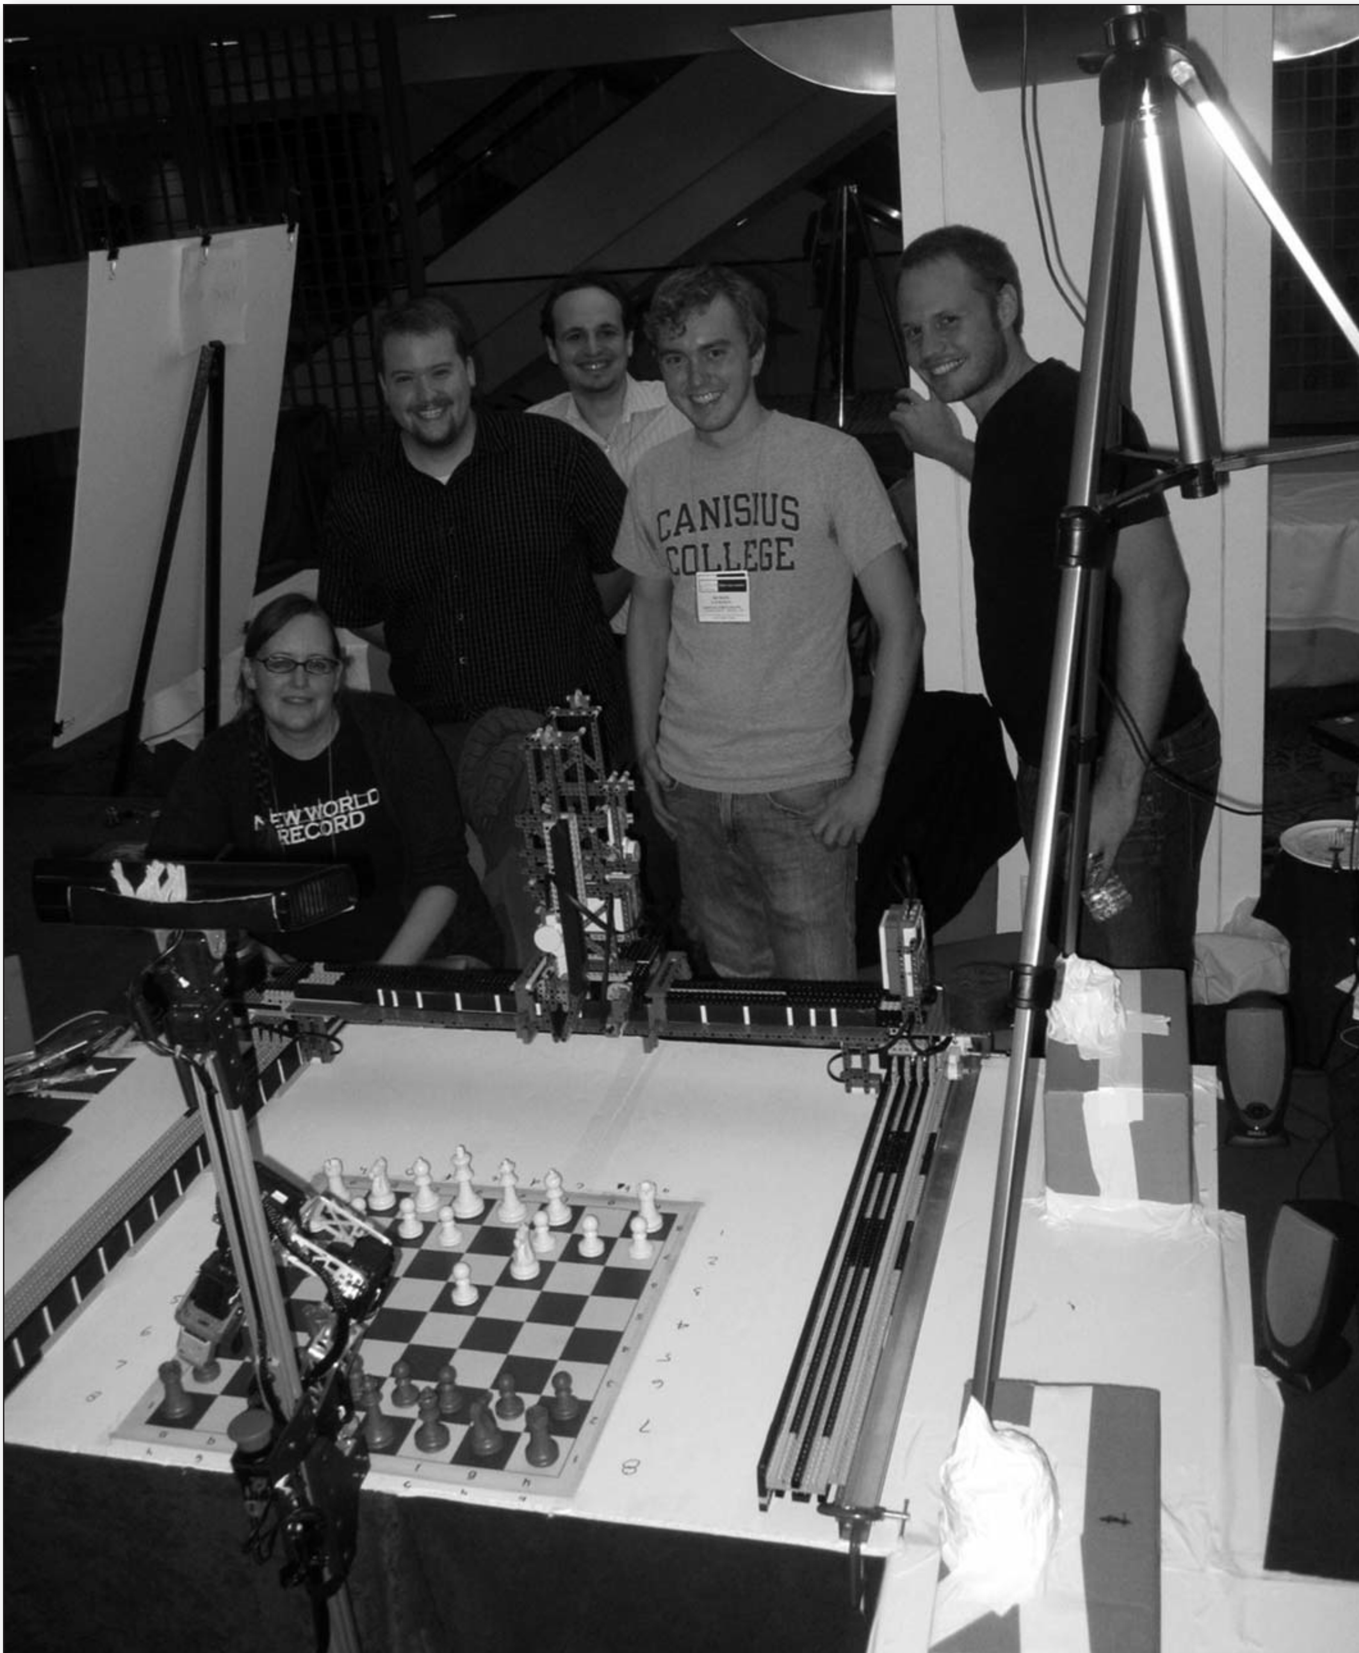

Figure 1. First Robot Versus Robot Chess Match, Griffins against Maxwell.

was a six-legged, one-armed robot developed by the Tekkotsu Lab that competed in 2010 using a gripper customized for chess pieces. Watson was assisted by Ethan Tira Thompson and Ashwin Iyengar. Like the Chiara, but unlike the other entries, Calliope's arm could not reach the entire board from one spot; the robot had to navigate around the board to reach the further ranks, performing true "mobile manipulation." Although the new arm performed well, glitches in the vision system prevented the robot from completing several games.

Vision remains a large challenge for robots. Three of the four competitors used nontraditional colors (blue and yellow pieces on a white and green checkerboard, or black and blue pieces on a red checkerboard) in an effort to simplify the segmentation problem. The winning team used traditional white and black pieces (on a blue and white checkerboard), but employed a depth sensor to help with segmentation. Progress in manipulator design continues. The Maxwell and Calliope arms were quite similar. Road Narrows' pulley-based design offered some performance advantages, at a cost of greater complexity and parts count. None of these arms included force feedback or could perform delicate manipulations, so there is much more to be done in the small-scale manipulation domain.

The robot chess challenge culminated with a unique after party, where the Griffins competed head-on against Albany's Maxwell on a single

Figure 2. Two Examples from the Learning from Demonstration Challenge, Learning Skills in the Cooking Domain.

Opposite Page: IIT's team presented a pizza-making robot that learned low-level motion trajectories required for rolling out dough. *Above:* The other four teams all used the PR2 platform; pictured here is Baris Akgun from Georgia Tech teaching the table-setting task.

board (see figure 1). Previously, each robot made moves on its own board, which were replicated by humans on the opponent's board. The game was not part of the tournament and the robots encountered a number of perceptual roadblocks. Yet, to our knowledge this was the first ever head-to-head competition between two chess-playing robots! Excitement about the match drew a crowd to watch the two vastly different robots, the Lego gantry of the Griffins and the lean and tall Maxwell, engage physically in the sport of chess.

Learning from Demonstration Challenge

Algorithms for robot learning from demonstration seek to enable human users to expand the capabilities of robots through interactive teaching instead of explicit programming. The aim of the Learning from Demonstration Challenge is to promote technological innovation and comparative assessment in this area, through an organized challenge. In previous years, challenge participants demonstrated their research techniques on a set of unrelated tasks. To better enable comparisons and cooperation, the 2011 Learning from Demonstration Challenge focused on the central theme of food preparation. The event consisted of two components, a table-setting challenge featuring a standard domain and robot platform, the Willow Garage PR2 robot, and an open demo component for any cooking-related tasks (see figure 2).

The table-setting challenge involved setting the table for breakfast for two people by selecting objects from one table and placing them into the appropriate location on a second table. The task requires a combination of both low-level skills and high-level reasoning, allowing a variety of learning from demonstration techniques to be showcased. The event was sponsored by the Bosch Research and Technology Center, which provided the PR2 robot, a custom software framework, and remote access to the robot through the web for all participants. Four teams participated in the table-setting challenge.

The team from Ecole Polytechnique Federale de Lausanne (EPFL), advised by Aude Billard, focused on learning low-level motion primitives, such as reaching for a glass or pouring milk, from demonstration. The learning method, called Stable Estimator of Dynamical Systems (SEDS) (Khansari-Zadeh and Billard 2011), models motions as a nonlinear autonomous dynamical system and defines sufficient conditions to ensure global asymptotic stability at the target.

The team from the Massachusetts Institute of Technology, advised by Brian Williams, also focused on the learning of low-level motion primitives from demonstration. Their approach employed probabilistic flow tubes to infer the desired state region at each time step from the data provided by the demonstrations (Dong and Williams 2010).

The team from the Georgia Institute of Technology, advised by Andrea Thomaz, used kinesthetic teaching to teach motion sequences, such as grasp-carry-place. Kinesthetic teaching enables the human teacher to physically guide the robot in performing the skill. The learning algorithm represented the trajectory using a sparse set of consecutive keyframes that achieve the skill when connected together (Akgun et al 2012).

The team from Worcester Polytechnic Institute, advised by Sonia Chernova, focused on learning the task as a high level plan consisting of subgoals. Using the Behavior Networks algorithm (Nicolescu and Mataric 2002), the robot created a history of the subgoals reached during the demonstration process, and reproduced the task by planning over these subgoals.

Finally, the team from the Italian Institute of Technology, advised by Sylvain Calinon, participated in the open demonstration component of the Learning from Demonstration Challenge. The team presented a pizza-making robot, with learning focusing on the acquisition of low-level motion trajectories required for rolling out pizza dough. The robot used both force and motion information from a kinesthetic demonstration, combined with visual feedback, to learn a parametric statistical representation of the force and position or orientation in Cartesian space. The teacher demonstrated how to flatten pizza dough by associating each movement with parameters extracted from vision such as center of mass or eigencomponents of the dough. Depending on the current shape of the dough, the rolling pin in the robot's grasp was used at different places and angles. The robot also learned when to stop the task of flattening the dough in order to move to the next step of the pizza preparation.

The design of the Learning from Demonstration Challenge around a central theme has been a great success, resulting in increased collaboration between participants on future projects, and a submission of a coauthored paper. Critically, the event was not organized as a competition, and no winner was selected among the participants. Instead, the goal of the event was to highlight the diversity of the learning from demonstration field and to initiate a discussion of how existing techniques could be combined in the future to develop more robust learning algorithms. Although none of the teams in the table-setting challenge completed the full task independently, the event has created opportunities for future advancements in this area through the establishment of a shared standard API that leverages the Bosch PR2 Remote Lab. Since the event, members across several teams have began exploring how low-level motion trajectory learning techniques can be integrated with high-level reasoning to solve the full table-setting task.

AAAI 2011 Robot Exhibition: Education Track

For many undergraduate and precollege students, robotics serves to motivate the study of artificial intelligence. The AAAI robotics exhibition's education track leverages this motivation: it provides students a venue for their research projects and offers educators a forum to share innovative curricula. During the workshop, attendees brainstormed both frustrations and opportunities present in today's teaching of AI through physically embodied agents. Those discussions made it clear that educators continue to seek hardware, software, and curricular resources that will excite their students about the topics featured at AAAI. The 2011 education track exhibits approached this mandate in three distinct ways: first, by bridging undergraduate courses and programs with AI topics of current research interest; second, by improving support for the curricular use of robots not yet typical outside of research labs, and, third, by showcasing novel educational robot platforms and the programming suites that support them.

Teams from the City University of New York focused on engaging students with two very active areas of AI research: multirobot exploration and localization and learning spatial tasks by demonstration. Six participants from CUNY's Research

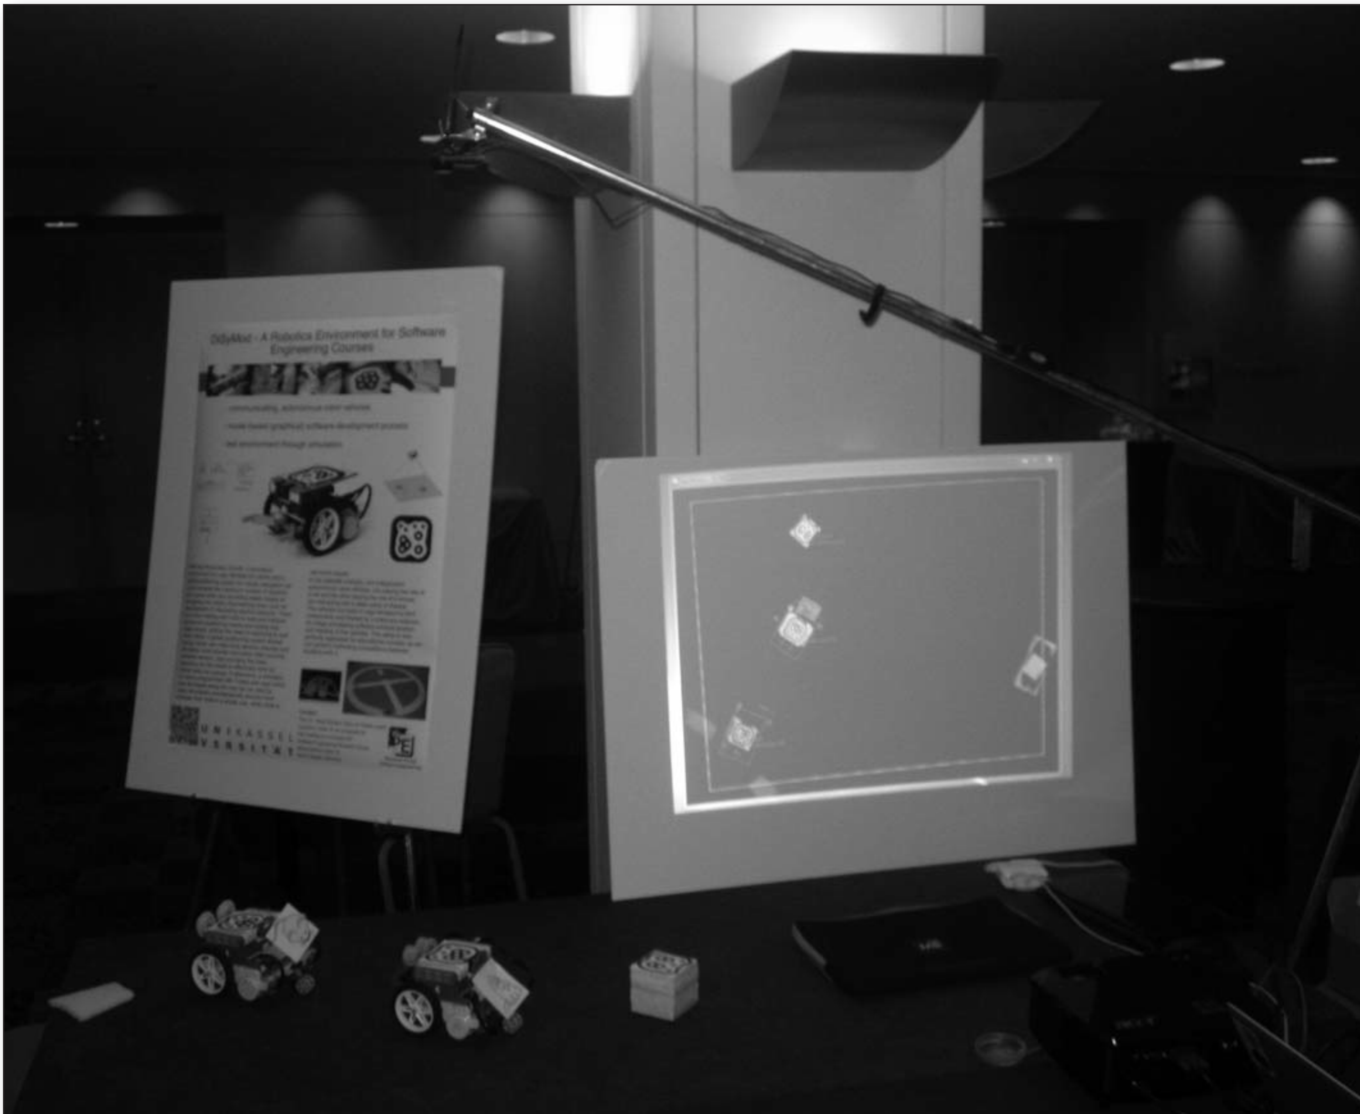

Figure 3. University of Kassel's Exhibit.

In this exhibit of curriculum, students developed software engineering experience by specifying the behaviors of distinct agents. At the exhibition, observers could specify roles for physical agents. In this example, a Lego cat would chase a mouse, which alternately hid away in safety and sneaked out to grab cheese from the table.

Experiences for Undergraduates (REU) program presented their work in coordinated robot navigation. Using the iRobot Create, the Scribbler, and the surveyor, the CUNY team demonstrated autonomous map-building that was simultaneously accurate and accessibly cost-effective. The Scribbler and Surveyor are small, inexpensive robots designed for hands-on classroom investigation at the high school and undergraduate levels. A subset of CUNY's exhibitors also presented a project overlapping with the learning from demonstration research track. In contrast to the teams using the PR2, however, CUNY's work featured an AIBO robot accomplishing maze-navigation and obstacle-avoidance tasks after a human teacher presents a small number of exemplar choices through a graphical interface.

As research topics find their way into the undergraduate experience, a challenge arises: new curricula are needed to support them. A group from the University of Kassel exhibited an innovative curriculum in which students developed software engineering experience specifying the behaviors of distinct agents (figure 3). Those behaviors were then tested both through simulation and embodied in Lego robots. The physical realization relied on an overhead webcamera and fiducial markers: the lower-level details of control and localization were abstracted away, allowing the students to reason about the behaviors without the overhead of programming all of them from scratch. At the exhibition, observers could specify a variety of roles for the physical agents. In one memorable example, a Lego cat would chase a mouse, which alternately hid away in safety and sneaked out to grab cheese from the table.

A team of Research Experience for Undergraduates (REU) students from Harvey Mudd College also showed a live localization demo, in their case implemented using an ARDrone quadrotor helicopter. That platform and a Kinect-equipped Create showed the possibilities of Willow Garage's open TurtleBot platform and ROS software in service of new curricula. With high-resolution ranging and aerial autonomy now accessibly inexpensive, the use of those resources can now be considered across a wider range of the CS and engineering curriculum than even one or two years ago. For instance, these coordinated ground and aerial robot teams now form the basis for a CS2-level programming course at Harvey Mudd College.

Emerging platforms generate excitement among AI researchers and educators alike, and the education track featured two robots and software suites that fundamentally change the questions undergraduates can investigate in an AI course. Ethan Tira-Thompson presented "What's New in Tekkotsu," a mature software scaffold that supports several legged and wheeled robot platforms with manipulation capabilities. Already in use at many schools, the Tekkotsu exhibit showed off some its latest capabilities by guiding a Create-based robot with a mounted servo arm in the chess-playing challenge. Manipulation and legged locomotion are also the differentiating capabilities of the Nao humanoid, which Aldebaran's Nataniel Dukan demonstrated to onlookers at the exhibition. The standard small-humanoid of RoboCup, the Nao and its software provide a single point of entry for investigating anthropomorphic tasks, just as Tekkotsu's Chiara and Calliope platforms offer accessible starting points for students and educators to experiment with sensing and manipulation atop a hexapedal or wheeled robot.

The AAAI education track hopes to inspire AI practitioners to ask how physical platforms might enhance the classes and projects their students pursue. Curricular innovations are as important as research insights in making AI a compelling and accessible field to as broad an audience as possible. The 2011 education track suggests that the study of AI and robotics in 10 years will involve very different topics, tools, software, and hardware than it does today. In particular, there are powerful sensing and actuation capabilities now considered the province of relatively few research labs, for example, dense, realtime, three-dimensional reconstruction, legged or aerial locomotion, and mobile manipulation. As the workshop showed, these capabilities are quickly become mature and accessible enough to become resources for teaching AI, not only objects of study within the field. We look forward both to the curricular opportunities these changes will foster and to the ideas and energy a wide cohort of students will bring to AI Robotics in the future.

Conclusion

The AAAI 2011 Robotics Program in San Francisco included a broad demonstration of innovative research at the intersection of robotics and artificial intelligence. Through these multiyear challenge events, our goal has been to focus the research community's energy toward common platforms and common problems to work toward the greater goal of embodied AI.

Acknowledgments

We would like to thank the National Science Foundation for its generous sponsorship of the AAAI 2011 Robotics Program in the form of student travel grants.

Notes

1. Throughout this article we mention many specific brand names of platforms and products; this is not meant as an advertisement for these companies or products. This is meant to provide the interested reader with details about what is in standard use in robotics research and education.

2. Available at Tekkotsu.org.

3. See the Chiara and Calliope website, chiara-robot.org.

References

Akgun, B.; Cakmak, M.;Yoo, J. W.; Thomaz, A. L. 2012. Trajectories and Keyframes for Kinesthetic Teaching: A Human-Robot Interaction Perspective. In *Proceedings of the International Conference on Human-Robot Interaction* (HRI). New York: Association for Computing Machinery.

Anderson, M.; Chernova, S.; Dodds, Z.; Thomaz, A. L.; Touretsky, D. 2011. Report on the AAAI 2010 Robot Exhibition. *AI Magazine* 32(3): 109–118.

Coens, J. 2010. Taking Tekkotsu out of the Plane. Masters thesis Technical report CMU-CS-10-139, Computer Science Department, Carnegie Mellon University, Pittsburgh, PA.

Dong, S., and Williams, B. 2010. Motion Learning in Variable Environments Using Probabilistic Flow Tubes. In *Proceedings of the 2011 IEEE International Conference on Robotics and Automation* (ICRA), 1976–1981. Piscataway, NJ: Institute of Electrical and Electronics Engineers.

Khansari-Zadeh, S. M., and Billard, A. 2011. Learning Stable Nonlinear Dynamical Systems with Gaussian Mixture Models. *IEEE Transactions on Robotics* 27(5): 943–957.

Nicolescu, M., Mataric, M. J. 2002. A Hierarchical Architecture for Behavior-Based Robots. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems* (AAMAS), 227–233. New York: Association for Computing Machinery.

Sonia Chernova is an assistant professor in the Computer Science Department and 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 CS education more generally.

Mike Stilman is an assistant professor in the School of Interactive Computing at the Georgia Institute of Technology.

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.

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.