The 17th Annual AAAI Robot Exhibition and Manipulation and Mobility Workshop

Monica Anderson, Odest Chadwicke Jenkins, and Paul Oh

■ *The AAAI 2008 Workshop on Mobility and Manipulation (held during the 23rd AAAI Conference on Artificial Intelligence) showcased advances in mobility and manipulation through a half-day workshop and an exhibition. The workshop focused on possible solutions to both technical and organizational challenges to mobility and manipulation research. This article presents the highlights of that discussion along with the content of the accompanying exhibits.*

Robots can often be found in homes vacuuming or on battlefields diffusing bombs. Fortunately, these applications can be successful through simple repetitive behaviors or remote human operation. However, useful autonomy needed for operation in general situations requires advanced mobility and manipulation. Opening doors, retrieving specific items, and maneuvering in cluttered environments are required for useful deployment in anything but the most controlled environment. The mobile manipulation skills necessary to perform tasks in arbitrary environments may not result from current approaches to robotics and AI. Moving toward true robot autonomy may require new paradigms, hardware, and ways of thinking.

The goal of the AAAI 2008 Workshop on Mobility and Manipulation was not only to demonstrate current research successes to the AAAI community but also to road-map future mobility and manipulation challenges that create synergies between artificial intelligence and robotics. The half- day workshop included both a session on the exhibits and a panel discussion. The panel consisted of five prominent researchers who led a discussion of future directions for mobility and manipulation research.

The half-day workshop was followed by a two-day combined exhibition of robot demonstrations from both the Mobility and Manipu-

 ISSN 0738-4602

Figure 1. The Snake Robot.

Students from Adler Planetarium summer robotics camp for middle school girls watch the snake robot from Carnegie Mellon University as it provides a live video feed.

Figure 2. Stanford University Tested the STAIR Platform.

STAIR was tested in two new buildings on 20 different types of doors.

lation and the Robots and Creativity workshops (held separately). The exhibits were collectively judged, with second place awarded to David Touretzky of Carnegie Mellon University for "The Chiara: A New Platform for Robotic Manipulation" and third place awarded to Oliver Brock from the University of Massachusetts Amherst for "Interactive Perception for Manipulation in Unstructured Environments." First place was awarded to a participant from the Robots and Creativity workshop.

We estimate that there were more than 200 visitors from the conference and the general public. Through special invitation, 20 students from the Adler Planetarium summer robotics camp for middle school girls attended the exhibition as a field trip (see figure 1). Student interaction with the exhibits ranged from sitting in robotic vehicles to watching robots climb legs. Each exhibit was chosen for how it defined a different approach to the problem as well as demonstration of results with real robots. Within this article, we explore mobility and manipulation in terms of the contribution of each exhibit. Approaches and challenges as discussed at the workshop are detailed in future directions.

Approaches

The primary challenge in mobility and manipulation is the complexity of unstructured environments in which a priori models are not available. One method for managing novel experiences is learning. Three exhibitions advocated decidedly different learning approaches for application in cluttered, dynamic, unknown environments.

Andrew Ng of Stanford University (along with students Ashutosh Saxena and Ellen Klingbeil) focuses on opening arbitrary doors through learning a few visual keypoints, such as the location and type of door handle. Such keypoints are enough to plan an action for opening doors. Similar function implies similar form, and hence similar visual features. A vision-based supervised learning algorithm helps to achieve generalization and therefore applies to new doors. The STAIR (Stanford AI robot) (figure 2) platform succeeded in opening doors 91.2 percent in a total of 34 trials.

Oliver Brock of University of Massachusetts Amherst (along with students Dov Katz and Jacqueline Kenney) integrates perception, action, and learning to create a foundation for intelligent behavior in novel, cluttered environments. Perception capabilities acquire relevant task information without making assumptions about the environment. Learning-manipulation capabilities allow the generalization and transfer of knowledge to new objects. The exhibit showcased the UMan (University of Massachusetts mobile manipulator) (figure 3) actively learning manipulation parameters of novel objects and transferring knowledge to decrease the learning time of new objects.

Jianna Zhang of Western Washington University applies machine learning to robots targeted to home health care. A reinforcement-learning algorithm uses sensor inputs to determine the safety of obeying the user's commands. Commands that appear, in specific environments, to endanger the robot are gradually disobeyed. The exhibition prototype used the standard Lego Mindstorms kit as a way to construct a low-cost, high-functioning learning robot platform (figure 4).

Complex environments sometimes require a new approach to platforms and mechanisms. Howie Choset of Carnegie Mellon University (along with students Eric Rippey, Donnie Cober, and Frederick Layton) demonstrated a snake robot. This platform's novel locomotion provides enhanced access possibilities over traditional wheeled and bipedal platforms. However, such mobility requires different gait- and path-planning approaches. Through teleoperation, the robot can be directed to climb a cylindrical object like a leg or pole (figure 5).

Jerry Weinberg of Southern Illinois University (along with student Jeff Croxell) incorporates kinesthetic intelligence by providing an automated process for a robot to discover the kinematic equations that govern the movement of its limbs. Motivated by the difficulty of setting up or changing robot arms, researchers use accelerometers and encoder data to measure link lengths.

Two exhibits focused on advanced computing power on less-customized platforms. David Touretzky of Carnegie Mellon University (along with students Ethan Tira-Thompson and Charleston Manning) demonstrated the Chiara, a hexapod with manipulation capability (see figure 6). This platform meets both research and educational needs by being relatively low cost. The open-source design and incorporation of commodity parts allows the robot to be maintained in house. An included 1 GHz processor with 1 GB of RAM and an 80 GB hard drive enables high-level vision algorithms, posture control, and manipulation primitives.

Jerry Weinberg and Kim Wheeler of RoadNarrows Robotics showcased the SkewlZone brain pack. Augmented with a sensor network, it provides real-time computing power and control for off-the-shelf legged robots. Integration with the Kondo KHR2 (see figure 7) results in high-level motion primitives accessible by both on- and off-board controller programs. Both the Chiara and SkewlZone platforms utilize open-source software such as Linux and Tekkotsu to encourage sharing of high-level components between research groups, reducing software development time and increasing opportunities for vertical integration.

Some entries were geared toward algorithms rather than hardware. David A. Gustafson of Kansas State University (along with students Andrew King, Michael Marlen, Aaron Chavez, Aaron Chleborad, Alex Richardson, Eric Vanneval, and Jon Hoag) demonstrated a semantic vision algorithm (figure 8). The objective is to study the effectiveness of building just-in-time set-superset ontologies from the web and using these to eliminate irrelevant images and to establish the context of an object.

Jerry Weinberg and student Ross Mead (with assistance from George Engel and Ryan Krauss) presented a distributed control algorithm for robots

Figure 3. Students Demonstrate the UMan Platform.

UMan focuses on manipulation of arbitrary objects for which there is no a priori model. This exhibit won third place in the overall competition.

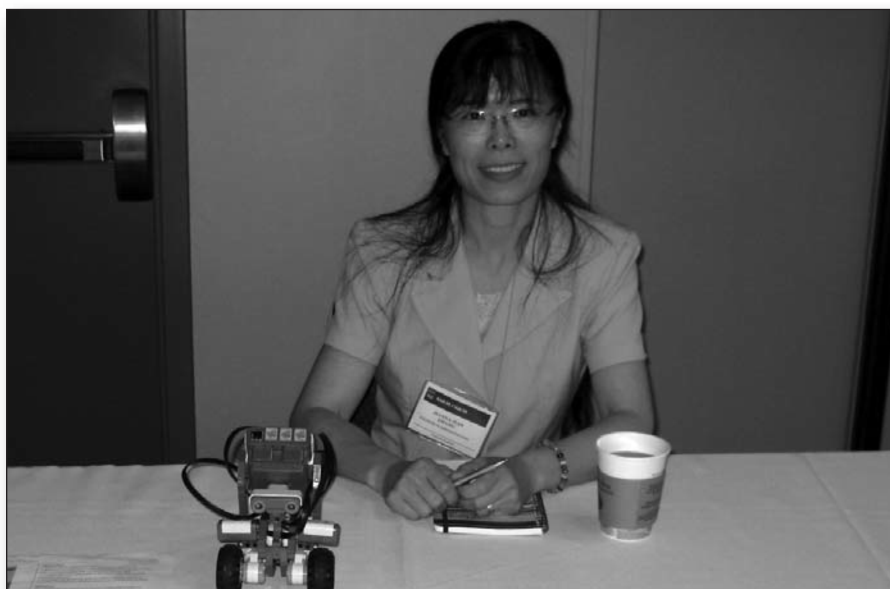

Figure 4. Jianna Zhang (Western Washington University) Demonstrates Reinforcement Learning on a Lego Mindstorms Kit.

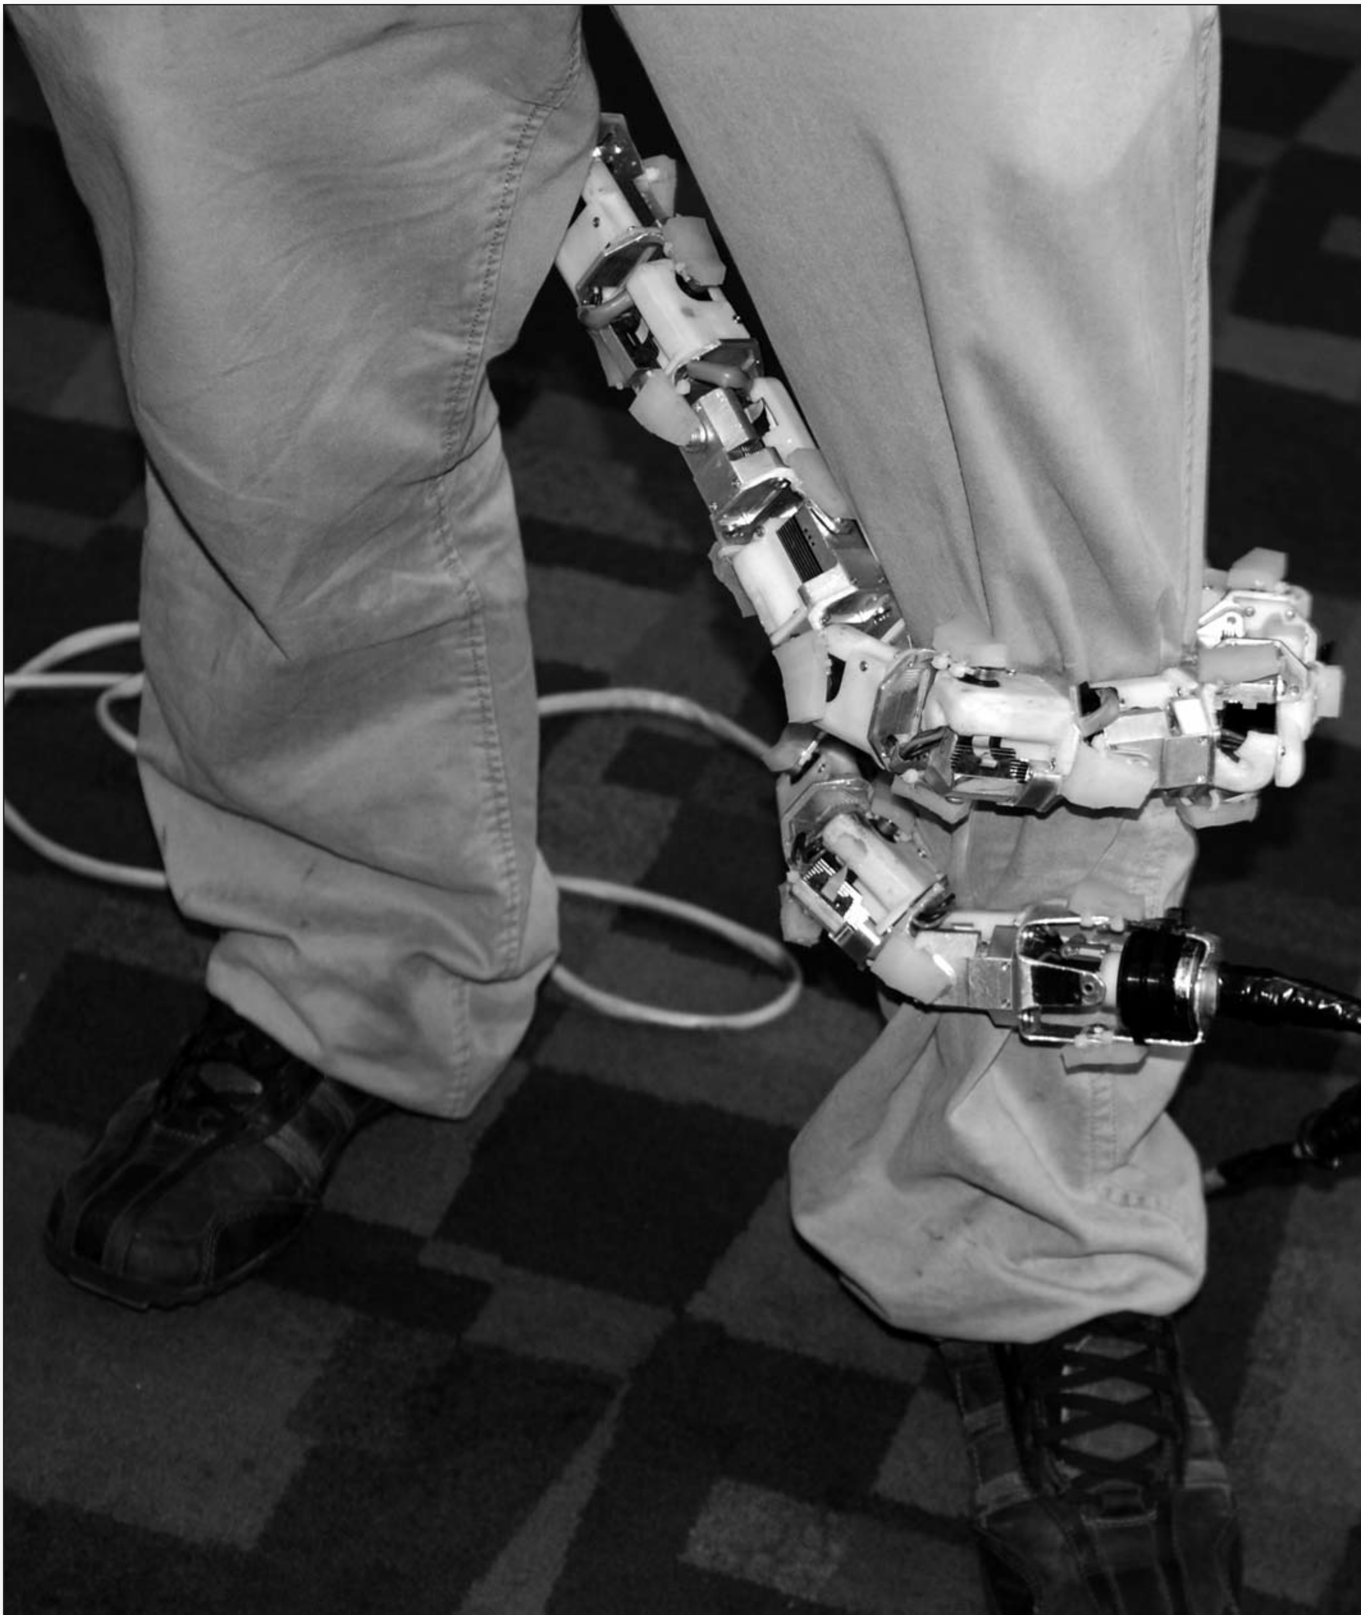

Figure 5. Howie Choset's Students Demonstrate a 16 DOF Snake Robot.

The robot is based on hobby servos covered with dragon skin for enhanced gripping capabilities. This exhibit won a blue ribbon for robot hardware.

in grid formations. Self-organization results from treating each robot as a cell in cellular automation, where the robot reacts to changes in its neighborhood. Robots can move in formation and change orientation based on a single command.

The final exhibit presented by Jerry Weinberg and student Ross Mead was Road Runner, a single-passenger autonomous electric vehicle (see figure 9). This vehicle enables human-robot interface (HRI) research through interactions with the passenger. Basic services include braking, steering, and stopping. An OpenCV-based vision algorithm provides path following.

Future Directions

Five prominent figures in robotics research participated in a panel discussion. Each panelist discussed ways of motivating future research through challenges at the intersection of artificial intelligence and robotics. Panelists suggested several challenges within the military and home health-care domains. Searching a vehicle entails opening the doors, trunk, and glove compartment, as well as moving occluding items out of the way, and uses manipulation of arbitrary objects and requires grasping handles. By contrast, bimanual manipulation is needed to unzip a bag and selectively remove its contents. Moving a large object, like a casualty, shows footstep planning that takes the awkwardness of a nonrigid object into consideration. Suggested challenges in the home health-care domain include pouring a drink, cleaning a home, or loading the dishwasher. However, most discussion centered on the ability of the current approaches, methods, and infrastructure to solve current manipulations problems. Panelists and attendees focused on issues that posed impediments to advancing mobility and shared their vision on how true autonomy and useful human interaction may be accomplished.

Robert Mandelbaum (DARPA) witnesses the opportunity for robots in military tasks (figure 10). He envisions that robots should go where people go and do what people do. Future advances rely upon four urgent technol-

Figure 6. Ethan Tira-Thompson of Carnegie Mellon Demonstrates the Chiara.

Chiara is an open-source platform constructed from commodity parts. It combines low cost with high computational ability, providing support for vision and dexterous manipulation through the Tekkotsu software framework. This exhibit won second place in the overall competition.

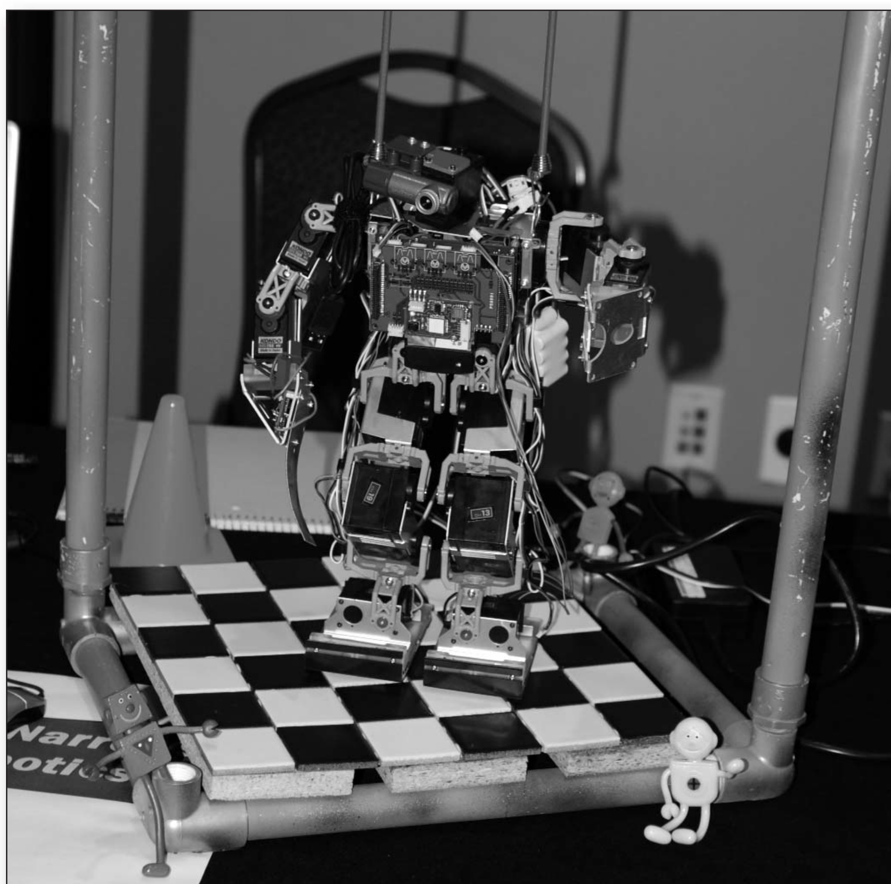

Figure 7. The RoadNarrows SkewlZone Brain Pack Integrated with the Kondo KHR2.

Figure 8. Semantic Vision System from Kansas State University Recognizes a Book Using the Sematic Net.

ogy needs. First, better operator situational awareness; second, fewer physical limitations; third, better autonomous control; and fourth, coordinated mobility and manipulation. Situational awareness should include hapic or three-dimensional displays that model the workspace in real time using low bandwidth. New displays should include extra information such as the range to an object. New robots should overcome current limitations by providing a better strength-to-weight ratio and improved dexterity. Force sensing should be included in all fingers and joints, not just the wrists. Vector touch-sensitive skin would provide sensitivity greater than the amount of force needed to pick up an object to enable the ability to determine when an object is slipping. Coordination of mobility and manipulation can enable canonical tasks such as opening a door and then moving through it.

The research of Charlie Kemp (Georgia Institute of Technology) centers on coordinated mobility and manipulation. Kemp advocates empirical experimentation with robots operating in the real world. Through this methodology, it has been shown that some problems that have been considered hard succumb to straightforward solutions. For example, recent results have shown that many objects can be grasped using controller-based approaches without detailed models or model-based planning. Kemp also believes the field can promote meaningful progress by evaluating robotic systems in the context of real-world problems, such as assisting the motor impaired.

Oliver Brock (University of Massachusetts Amherst) (figure 11) perceives the need for a paradigm shift. The current paradigm focuses on developing well-defined capabilities, specified by a researcher and believed to fully address a particular task. If the task to be performed by the robot is well-defined, the existing paradigm can lead to impressive results, as demonstrated by the DARPA Grand Challenge. However, the existing paradigm does not enable robustness in situations that differ qualitatively from those anticipated at design time. To achieve autonomy that is not confined to those anticipated circumstances, robots have to learn in such a way that skills transfer across tasks, across domains, and across levels of complexity. If we learn to put a peg into a hole, it should teach us something about putting a key into a keyhole.

Robotics needs a paradigm shift to meet technology targets. Robots must be able to generalize and find new ways of solving a problem by themselves. Better sensors alone cannot solve the problem of autonomous manipulation. In fact, high-resolution sensors provide more information that

Figure 9. The Road Runner.

The Road Runner is a single-passenger autonomous electric vehicle developed by Southern Illinois University. The current sensor suite (sonar, GPS, shaft encoders, a color web camera) provides path following and waypoint navigation. This exhibit won a blue ribbon for integration.

Figure 10. Robert Mandelbaum Discusses Opportunities for Artificial Intelligence in Robotics.

Figure 11. Oliver Brock Discusses the Future Integration of AAAI and Robotics.

Figure 12. Stewart Tansley Discusses Approaches to Organizational Challenges.

further increases the complexity of the problem. To be able to operate successfully in unstructured, open environments and across a variety of tasks, robots must be able to identify within the abundance of sensor data the information salient to the current task.

In contrast, Andrew Ng promotes a standard platform for robotics to enable software sharing in the same way that IBM-compatible PCs enabled sharing of spreadsheets. Without a standard platform, each research group must create its own hardware. Software sharing is difficult since everyone has a different hardware platform.

Stewart Tansley (Microsoft Research) (figure 12) discussed current organizational challenges. He sees robotics as an ongoing inspiration for research within the AI community and the larger technology research community. Many successes in robotics are artificial intelligence successes. Since robotics is very accessible to a broad population, showcasing robotics at AAAI can increase the general prestige and interest in the conference overall. Although the hardware aspect of robotics belongs in other communities in his view, AI algorithms applied to robotics exemplifies contemporary AI techniques in demanding real-world applications. Robotics is where the "AI rubber hits the road."

Bridging the two communities can be accomplished by focusing on areas of combined interest. Human-robot interfaces can include natural language interfaces as well as learning and cognition, all of which are central to AI research. AI-based robotics should become part of the main conference through integration into existing tracks, not only as a separate track. The application of AI techniques in robotic contexts provides new knowledge that can help to delineate limitations and constraints. This helps to shape future research directions. Joint road-mapping exercises can link challenges to research directions within the AAAI community.

The robotics community can do more to support artificial intelligence research too. Practical tutorials on acquiring and using current robot technology would help AI labs move toward more robotic applications.

Finally, there is the systems perspective of robotics. This provides a rare opportunity for students to learn systems engineering for complex AI systems. Students must use a systems design approach, which gives them valuable experience for a wide variety of careers in research and industry.

Conclusion

The goal for future AAAI robotics challenges is to demonstrate artificial intelligence techniques on physical platforms. Robotic challenges will be coordinated with similar efforts such as challenges at ICRA. Proposed robotic challenges for the 21st International Joint Conference on Artificial Intelligence (IJCAI-09) include learning from demonstration, mobile manipulation, multirobot teaming, and an undergraduate challenge.

Monica Anderson received her Ph.D. in computer science from the University of Minnesota in 2007. She currently is an assistant professor at the University of Alabama.

Odest Chadwicke Jenkins is an assistant professor at Brown University.

Paul Oh is an associate professor at Drexel University.