

W. BRADLEY KNOX

<http://www.bradknox.net>
bradknox@mit.edu

RESEARCH INTERESTS

I conduct research on intersections of (1) the control of robots and other agents, (2) machine learning (reinforcement learning in particular), (3) human-computer interaction, and (4) computational models of human behavior for cognitive science research. I'm primarily interested in human interaction with machine learning algorithms, especially when the human fills a teaching role.

EDUCATION AND RESEARCH POSITIONS

Postdoctoral Research, MIT Media Lab, October 2012 – July 2014

Developing robot learning companions for young children.

Advisor: **Cynthia Breazeal**

Ph.D. Computer Science, University of Texas at Austin, 2006 - 2012

Dissertation: **Learning from Human-Generated Reward**

Advisor: **Peter Stone**

B.S. Psychology, Texas A&M University at College Station, 1999 - 2003

Minor in Philosophy and completed pre-medical course requirements.

HONORS

Named to IEEE Artificial Intelligence's 10 To Watch, 2013

Bert Kay Dissertation Award, 2013

Award for best dissertation, Dept. of Computer Science at UT Austin

IFAAMAS-12 V. Lesser Distinguished Dissertation Award (runner-up), 2013

Chosen among 2012 PhD theses on autonomous agents or multiagent systems

ICSR Best Paper Award, 2013

Ro-Man CoTeSys Cognitive Robotics Best Paper Award (Finalist), 2012

AAMAS Pragnesh Jay Modi Best Student Paper Award, 2010

NSF Graduate Research Fellowship, 2008–2011

UT Austin's David Bruton, Jr. Graduate School Fellowship, 2008, 2009, 2010

Texas A&M University President's and Distinguished Scholar, 1999–2003

National Merit Scholar, 1999–2003

TEACHING EXPERIENCE

MIT, Fall 2013

Principal lecturer

Designed and led graduate course entitled Interactive Machine Learning

- Of the 21 Fall 2013 courses in the MIT Media Lab, this course *scored highest in student evaluations* on “[instructor] stimulated interest”, “subject’s learning objectives were met”, and “assignments contributed to my learning”. It tied for highest on “[instructor] helped me learn” and tied for the 3rd highest “overall rating of the subject”.

Website: <http://iml.media.mit.edu>

University of Texas at Austin, August 2006 – August 2012

Teaching assistant, August 2006 – May 2007.

- Autonomous Vehicles in Traffic
- Introduction to Programming in Java

- Overall student evaluation from two discussion sections: 4.7 and 4.3 (1–5 scale)

Texas A&M University Center for Academic Achievement, Fall 2001.
Supplemental Instruction leader for introductory psychology course

STUDENT MENTORING (graduate students in italics)

- **Sam Spaulding**, September 2013 – current
During Sam’s first year as a master’s student in the MIT Media Lab, I directly oversaw his contributions to our robotic learning companion project. Our work has been submitted to RSS.
- **Max Molina**, Summer 2013
I mentored Max Molina, an MIT undergraduate researcher, contributing to an ongoing project of mine that involves programming a laser-and-camera-based pet toy by user demonstrations of laser control. Max fine-tuned camera-based sensing of a pet’s location and orientation and later mapped this pet information to the coordinate system of the laser waypoints given by control.
- **Kyle Liberti**, Summer 2013
I mentored Kyle, an undergraduate researcher from Boston University, as he learned the fundamentals of reinforcement learning (partially through a reading group with Adam, Robert, and Luis from this list) and later translated over 100 reward functions from free-form description to Java code for an ongoing project with expected publication next year.
- **Luis Sanmiguel**, Summer 2013
Luis, also an MIT undergraduate researcher, conducted summer research on incorporating characteristics of animal learning into reinforcement learning algorithms, making progress on an ongoing project.
- **Patrick McCabe**, Spring 2013
I mentored Patrick, an MIT undergraduate researcher, as he drove early progress on the aforementioned pet-game project, integrating the system’s hardware components and initiating camera-based sensing of the pet.
- **Adam Yala**, January 2013 – December 2013
I mentored MIT undergraduate researcher Adam Yala, who with my guidance built a module for detecting vocal pitch and intensity in real time. This *prosody* module has been used in various human-robot interaction projects within our research group at MIT. Adam, like all of the undergraduates on this page, began his research with me as a *freshman*.
- **Robert Radway**, January 2013 – December 2013
I instructed Robert Radway, also an MIT undergraduate researcher, as he designed and conducted experiments on using past training sessions with a human trainer to bias reinforcement learning in a new session.
- **Jin Joo Lee**, October 2012 – December 2013
I co-supervised Jin Joo Lee’s PhD research (with Professor Cynthia Breazeal) at MIT, which focused on modeling interpersonal trust from nonverbal behavior. Her research under my supervision resulted in a journal article in *Frontiers in Psychology* [4].
- **Adam Setapen**, January 2010 – July 2010
I mentored Adam while he was completing a master’s degree at UT Austin, where we collaborated to investigate learning from human feedback, resulting in a paper at a *AAAI Spring 2011 Symposium*. Adam subsequently earned a second master’s degree from the MIT Media Lab and now works for the personal robotics startup Romotive, where he leads artificial intelligence development.

PUBLICATIONS

Journal articles

- [1] W. Bradley Knox and Peter Stone. **Framing Reinforcement Learning from Human Reward: Reward Positivity, Temporal Discounting, Episodicity, and Performance.** Under review.
- [2] Saleema Amershi, Maya Cakmak, W. Bradley Knox, and Todd Kulesza. **Power to the People: The Role of Humans in Interactive Machine Learning.** *AI Magazine*. 35 (4): pp. 105-120. January 2015.
- [3] A. Ross Otto, W. Bradley Knox, Art Markman, and Bradley C. Love. **Behavioral and Physiological Signatures of Reflective Exploratory Choice.** *Cognitive, Affective, & Behavioral Neuroscience*. 1-17. March 2014.
- [4] Jin Joo Lee, W. Bradley Knox, Jolie Baumann, Cynthia Breazeal, and David DeSteno. **Computationally Modeling Interpersonal Trust.** *Frontiers in Psychology*, 4, 893. 2013.
- [5] W. Bradley Knox, Brian Glass, Bradley C. Love, W. Todd Maddox, and Peter Stone. **How Humans Teach Agents: A New Experimental Perspective.** *International Journal of Social Robotics. Special Issue on Robot Learning from Demonstration*. 2012.
- [6] W. Bradley Knox, A. Ross Otto, Peter Stone, and Bradley C. Love. **The Nature of Belief-Directed Exploratory Choice by Human Decision-makers.** *Frontiers of Psychology*, 2, 398. *Special Issue in Cognitive Science specialty on the Dynamics of Decision-Making*. 2012.
- [7] Juhyun Lee, W. Bradley Knox, and Peter Stone. **Inter-Classifier Feedback for Human-Robot Interaction in a Domestic Setting.** *Journal of Physical Agents*, 2(2):41–50, July 2008. Special Issue on Human Interaction with Domestic Robots.

Highly refereed conference proceedings

- [8] Guangliang Li, Hayley Hung, Shimon Whiteson, and W. Bradley Knox. **Learning from Human Reward Benefits from Socio-competitive Feedback.** In Proceedings of the IEEE International Conference on Development and Learning and on Epigenetic Robots (ICDL-EPIROB).
- [9] W. Bradley Knox, Cynthia Breazeal, and Peter Stone. **Training a Robot via Human Feedback: A Case Study.** In Proceedings of the International Conference on Social Robotics (ICSR), October 2013. 58% acceptance rate. **Best Paper** (93 submissions; 1 award).
- [10] Guangliang Li, Hayley Hung, Shimon Whiteson, and W. Bradley Knox. **Using Informative Behavior to Increase Engagement in the TAMER Framework.** In Proceedings of the Twelfth International Conference on Autonomous Agents and Multiagent Systems (AAMAS), May 2013. 23% acceptance rate.
- [11] W. Bradley Knox and Peter Stone. **Learning Non-Myopically from Human-Generated Reward.** In *Proceedings of the International Conference on Intelligent User Interfaces (IUI)*, March 2013. ~20% acceptance rate.
- [12] W. Bradley Knox and Peter Stone. **Reinforcement Learning with Human and MDP Reward.** In *Proceedings of the Eleventh International Conference on Autonomous Agents and Multiagent Systems (AAMAS-12)*, June 2012. 20% acceptance rate.

- [13] W. Bradley Knox and Peter Stone. **Combining Manual Feedback with Subsequent MDP Reward Signals for Reinforcement Learning.** In *Proceedings of the Ninth International Conference on Autonomous Agents and Multiagent Systems (AAMAS-10)*, pages 5-12, May 2010. 24% acceptance rate. **Best Student Paper** (684 submissions; 1 student award).
- [14] W. Bradley Knox and Peter Stone. **Interactively Shaping Agents via Human Reinforcement: The TAMER Framework.** In *Proceedings of The Fifth International Conference on Knowledge Capture (K-CAP 2009)*, pages 9-16, September 2009. 26% acceptance rate.
- [15] W. Bradley Knox and Peter Stone. **TAMER: Training of an Agent Manually via Evaluative Reinforcement.** In *IEEE 7th International Conference on Development and Learning (ICDL-08)*, pages 292-297, August 2008. Poster presentation. 33% acceptance rate, additional 35% for posters.
- [16] Gregory Kuhlmann, William B. Knox, and Peter Stone. **Know Thine Enemy: A Champion RoboCup Coach Agent.** In *Proceedings of the Twenty-First National Conference on Artificial Intelligence (AAAI-06)*, pages 1463-1468, July 2006. 30% acceptance rate.

Highly refereed symposium proceedings

- [17] W. Bradley Knox and Peter Stone. **Reinforcement Learning from Human Reward: Discounting in Episodic Tasks.** In *Proceedings of the 21st IEEE International Symposium on Robot and Human Interactive Communication (Ro-Man)*. September 2012. 70% acceptance rate. **Finalist for CoTeSys Cognitive Robotics Best Paper Award** (223 submissions; 4 nominees for award; 3 different Best Paper Awards).

Book chapter

- [18] W. Bradley Knox, Juhyun Lee, and Peter Stone. **Domestic Interaction on a Segway Base.** In L. Iocchi, H. Matsubara, A. Weitzenfeld, and C. Zhou, editors, *Robocup 2008: Robot Soccer World Cup XII*, pages 519-531. Springer-Verlag, Berlin, 2009. 62% acceptance rate.

Lightly refereed conference and symposium proceedings

- [19] Jin Joo Lee, W. Bradley Knox, and Cynthia Breazeal. **Modeling the Dynamics of Nonverbal Behavior on Interpersonal Trust for Human-Robot Interactions.** In *AAAI Spring 2013 Symposium on Trust and Autonomous Systems*, March 2013.
- [20] W. Bradley Knox, Cynthia Breazeal, and Peter Stone. **Learning from feedback on actions past and intended.** In *Proceedings of 7th ACM/IEEE International Conference on Human-Robot Interaction, Late-Breaking Reports Session*. March 2012.
- [21] W. Bradley Knox, Adam Setapen, and Peter Stone. **Reinforcement Learning with Human Feedback in Mountain Car.** In *AAAI Spring 2011 Symposium entitled Help Me Help You: Bridging the Gaps in Human-Agent Collaboration*, March 2011.
- [22] W. Bradley Knox, Ian Fasel and Peter Stone. **Design Principles for Creating Human-Shapable Agents.** In *AAAI Spring 2009 Symposium on Agents that Learn from Human Teachers*, March 2009.

Lightly refereed but non-archival symposium and workshop papers

- [23] W. Bradley Knox, Samuel Spaulding, and Cynthia Breazeal. **Learning Social Interaction from the Wizard: A Proposal.** In *3rd Workshop on Machine Learning for Interactive Systems: Bridging the Gap Between Perception, Action and Communication (MLIS-2014)*. July 2014.

- [24] W. Bradley Knox and Peter Stone. **Learning Objectives for Numeric Human Feedback.** Poster presented at the 1st Multidisciplinary Conference on Reinforcement Learning and Decision Making (RLDM), October 2013.
- [25] W. Bradley Knox and Peter Stone. **Augmenting Reinforcement Learning with Human Feedback.** Poster presented at the *ICML 2011 Workshop on New Developments In Imitation Learning*, June 2011.
- [26] W. Bradley Knox, Matthew Taylor, and Peter Stone. **Understanding Human Teaching Modalities in Reinforcement Learning Environments: A Preliminary Report.** In *2011 IJCAI Workshop on Agents Learning Interactively from Human Teachers (ALIHT)*, July 2011.
- [27] W. Bradley Knox and Peter Stone. **Interactively Shaping Agents via Human Reinforcement: The TAMER Framework.** In *Workshop on Analysis and Design of Algorithms for Interactive Machine Learning at NIPS 2009 (ADA-IML'09)*, poster presentation, July 2009.
- [28] W. Bradley Knox and Peter Stone. **Interactive Shaping of a Tetris Agent Using the TAMER Framework.** Technical report in *Proceedings of 2009 AAI Robot Workshop at IJCAI-09*, July 2009.
- [29] W. Bradley Knox and Ole Mengshoel. **Diagnosis and Reconfiguration using Bayesian Networks: An Electrical Power System Case Study.** In *IJCAI 2009 Workshop on Self-* and Autonomous Systems*, July 2009.

Other peer-reviewed publications

- [30] A. Ross Otto, W. Bradley Knox, Tyler Davis, Art B. Markman, and Bradley C. Love. **Optimal Belief-Based Exploration by Human Decision-Makers and its Behavioral and Physiological Signatures.** Poster presented at the *52nd Annual Meeting of the Psychonomic Society*. Seattle, WA, November 2011.
- [31] A. Ross Otto, W. Bradley Knox, Bradley C. Love, Sam Gershman, Yael Niv, Darrell A. Worthy, W. Todd Maddox, J. M. Hotaling, J. R. Busemeyer, and R. M. Shiffrin. **Symposium: Computational, Neuroscientific, and Lifespan Perspectives on the Exploration-Exploitation Dilemma.** Presented at the *33rd Annual Conference of the Cognitive Science Society*. Boston, MA, July 2011.
- [32] A. Ross Otto, W. Bradley Knox, Tyler Davis, Arthur B. Markman, Bradley C. Love. **Anticipation- and Outcome-Related Skin Conductance Responses in an Exploration-Exploitation Task.** Poster presented at the *Eighteenth Annual Meeting of the Cognitive Neuroscience Society*, April 2011.
- [33] W. Bradley Knox and Peter Stone. **Training a Tetris Agent via Interactive Shaping: A Demonstration of the TAMER Framework.** In proceedings of *2010 AAMAS Demonstration Session*, May 2010. [archived]
- [34] W. Bradley Knox. **Shaping Agents via Human Reinforcement.** In *Proceedings of The AAI/SIGGART Doctoral Consortium at IJCAI-09*, July 2009.
- [35] W. Bradley Knox, Juhyun Lee, and Peter Stone. **Person Recognition on a Segway Robot: A Video of UT Austin Villa Robocup@Home 2007 Finals Demonstration.** *2008 IEEE International Conference on Robotics and Automation*, May 2008.

RESEARCH COMPETITIONS

- **Robocup 2007 @Home Competition**, July 2007
2nd Place
Home assistant robots from 11 teams representing 5 continents performed basic tasks relevant to assisting in a realistic home environment.
- **Robocup 2006 Coach Competition**, July 2006
2nd Place
Opponent modeling in the domain of simulated soccer.
- **Robocup 2005 Coach Competition**, July 2005
World Champion
Resulted in AAAI-2006 publication.

EXTENDED VISITS AND INTERNSHIPS

MIT, May 2011 – August 2011.

Visiting Student at the Media Lab

Implemented dissertation research on the social robot Nexi.

Data Mining Lab at Apple Inc., Summer 2010.

Data Mining Intern

Improved effectiveness of model for detecting fraudulent orders from the online Apple Store by (1) constructing relational features over other orders that are connected by matching attributes such as order zip code and credit card number used and (2) flagging suspicious orders by the expected financial loss rather than the probability of fraud. On real historical data, my improvements catch an additional \$3.5 million of attempted fraud money annually, accounting for almost half of fraud dollars that the previous model lost.

NASA Ames (contracted by Mission Critical Technologies, Inc.) Summer 2008.

Research intern

Developed service agent that diagnoses faults within an electric power system and takes actions to supply power as demanded in the face of the faults. Agent models system with dynamic Bayesian networks.

PRESS

- “Everything Is Not a Damn Robot”. Drake Bennett. *Bloomberg Businessweek*. October 2013

INVITED TALKS

- “Learning from the Wizard: Programming Social Interaction through Teleoperated Demonstration” – **Google**, September 2014
- “Robot Learning from Human Teachers” – **Rethink Robotics**, July 2014
- “Learning Control from Human-Generated Reward”
 - Texas A&M**, Computer Science and Engineering Seminar, December 2013
 - Univ. of Texas at Austin**, Research Colloquium at School of Information, November 2013
 - Brown University**, Computer Science, November 2013
 - Texas A&M**, Computer Engineering and Systems Group Seminar, February 2013
 - Vecna Medical**, February 2013
 - University of Washington**, July 2012
 - MIT**, CSAIL, April 2012
- “Robot Learning from Human Input” - guest lecture for **MIT course on Human-Robot**

- Interaction**, December 2012
- Plenary talk at **AAAI 2012 Fall Symposium Series**, representing the symposium Robots Learning Interactively from Human Teachers (RLIHT), 2012
- “Discounting Human Rewards” and “RL Applet”, **7th Barbados Workshop on Reinforcement Learning**, April 2012
- “Interactively Shaping Agents via Human Reinforcement”
 - Georgia Tech**, Charles Isbell’s and Andrea Thomaz’s research group, May 2011
 - 5th Barbados Workshop on Reinforcement Learning**, April 2010
 - University of Texas at Austin**, guest lecturer for Peter Stone’s course Artificial Intelligence, March 2010
 - Univ. of Southern California**, Milind Tambe’s TEAMCORE research group, July 2009

ACADEMIC ORGANIZATION

- **Robotics Co-chair**, AAAI Conference, 2015.
- **Organizing Committee**, Workshop on Interactive Machine Learning at IUI, 2013
- **Organizing Committee**, Symposium on Robots Learning Interactively from Human Teachers (RLIHT), part of AAAI 2012 Fall Symposium Series, 2012
- **Co-chair**, Workshop on Agents Learning Interactively from Human Teachers (ALIHT) at IJCAI, 2011 (led this 2 day workshop that had 14 paper presentations and 46 attendees)
- **Organizing Committee**, Workshop on Agents Learning Interactively from Human Teachers at AAMAS, 2010
- **Organizer**, UT Reading group for Agents that Learn from Humans, 2009 – 2010

SERVICE

- **Program Committee**, Assoc. for the Advancement of AI (AAAI), 2015, 2014, 2013
- **Reviewer**, ACM/IEEE Inter. Conf. on Human-Robot Interaction (HRI), 2015, 2014, 2012, 2009
- **Program Committee**, Reinforcement Learning and Decision-Making Conf. (RLDM), 2015
- **Program Committee**, Inter. Conf. on Autonomous Agents and Multiagent Sys. (AAMAS), 2014
- **Program Committee**, Inter. Conf. on Intelligent User Interfaces (IUI), 2014
- **Reviewer**, IEEE Transactions on Autonomous Mental Development, 2013
- **Program Committee**, Inter. Conf. on Machine Learning (ICML), 2013, 2012
- **Program Committee**, Inter. Joint Confs. on AI (IJCAI), 2013
- **Reviewer**, Machine Learning Journal (MLJ), 2013, 2012, 2010
- **Reviewer**, International Journal of Social Robotics (IJSR), 2013, 2012
- **Reviewer**, Artificial Intelligence Journal (AIJ), 2012
- **Program Committee**, European Workshop on Reinforcement Learning (EWRL), 2012
- **Reviewer**, IEEE Inter. Conference on Robotics and Automation (ICRA), 2012, 2011, 2010
- **Program Committee**, AAMAS Workshop on Adaptive Learning Agents, 2011, 2010
- **Reviewer**, American Control Conference (ACC), 2012
- **Reviewer**, Advances in Complex Systems (ACS) Journal, 2011
- **Co-reviewer**, Neural Information Processing Systems (NIPS) Conference, 2009

RELEVANT ACTIVITIES

- **Toastmasters International**, May 2009 – May 2011
Member

Gave numerous prepared speeches, speech evaluations, and impromptu speeches to develop my public speaking skills.

- **Coldtowne Theater**, August 2010 – October 2012
Improv comedy performer and student
Organized and performed on multiple troupes, including the winner of a two-month audience-choice competition. Performed shows every other week April–October 2012.

COMMUNITY SERVICE

- **Personal Robots Group in the MIT Media Lab**, November 2012 – current
Presenter for Non-Academic Visitors
Shared robotics research with at least 20 different audiences, ranging from young children to corporate and military representatives.
- **SXSW Interactive Festival**, March 2011
Presenter at Plutopia
Gave interactive demonstration of my research at this family-friendly event that exhibits technology to the public.
- **Explore UT**, March 2007, 2008, 2009, and 2010
Lab guide
Shared research with Central Texas community at university-wide “open house”.
- **Citizen Schools**, February 2009 – May 2009
Apprenticeship Instructor
Taught 10-week, after-school course on robotics at a low-income middle school.
- **UTCS Gradfest**, January 2007 – April 2007
Committee Member
Planned and ran department orientation for potential PhD students by serving on a committee of faculty, staff, and another graduate student. Managed transportation and lodging, social events for a weekend, and any requests by the potential students.
- **Graduate Representative Association of Computer Science**, September 2006 – September 2007
Executive Committee Member
Represented interests of graduate students to department and university. Achievements included revamping both monthly and weekly departmental social events, resulting in more than doubled attendance.
- **Mountain Mover’s Mission International in Honduras**, June 2004 – July 2004
Premedical volunteer
Set up and ran pharmacy and took vitals during medical brigades into rural, impoverished areas of Honduras.

Interests: unplanned exploration, performing improvised plays, and bicycle-based social outings

Citizenship: U.S.

REFERENCES

Cynthia Breazeal

Associate Professor

Massachusetts Institute of Technology

Media Lab

cynthiab@media.mit.edu

Please carbon copy Cynthia's administrative assistant, Polly Guggenheim, at pobs@media.mit.edu.

Michael Littman

Professor

Brown University

Department of Computer Science

mlittman@cs.brown.edu

Peter Stone

Associate Professor

University of Texas at Austin

Department of Computer Science

pstone@cs.utexas.edu

Andrea Thomaz

Associate Professor

Georgia Institute of Technology

School of Interactive Computing

athomaz@cc.gatech.edu