W. Bradley Knox

bradknox@mit.edu

EDUCATION

- **University of Texas at Austin**
  Doctorate of Philosophy, Computer Science, September 2006 - August 2012
  Dissertation: **Learning from Human-Generated Reward**
  Advisor: **Peter Stone**

  Master of Science, Computer Science, September 2006 - May 2010

  Post-Baccalaureate, Computer Science, January 2005 – May 2006
  - 3.92 GPA (39 hours of computer science and math)

- **Texas A&M University at College Station**
  Bachelor of Science, Psychology, September 1999 - December 2003
  - Minor in Philosophy and completed Pre-Medical course requirements.
  - 3.94 Overall GPA (Summa Cum Laude), 3.88 Science GPA
  - University Honors

HONORS

- **Nominated for UT-Austin's Outstanding Dissertation Award**, 2013
  Selected by UT-Austin Dept. of Computer Science; one nominee per department
- **Finalist: Ro-Man CoTeSys Cognitive Robotics Best Paper Award,** 2012
- **Nominated by UT-Austin: Dolores Zohrab Liebmann Fellowship,** 2011
  University limited to 3 nominees from 30 departmentally nominated applications.
- **AAMAS Pragnesh Jay Modi Best Student Paper Award,** 2010
- **NSF Graduate Research Fellowship,** 2008–2011
- **UT-Austin: 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

PROFESSIONAL EXPERIENCE

- **MIT**, October 2012 – current.
  Postdoctoral fellow at the Media Lab
  Working on Socially Assistive Robots, an NSF Expeditions project that aims to create socially interactive, physically embodied intelligent tutors for young children.
- **University of Texas at Austin**, August 2006 – August 2012
  - *Research assistant,* September 2011– August 2012, May 2007 – May 2008,

and *NSF Graduate Research Fellow,* September 2008 – August 2011
Conducted original research in **reinforcement learning, robotics, human-robot interaction, and cognitive science**.

– *Teaching assistant, August 2006 – May 2007.*
• Autonomous Vehicles in Traffic – For an undergraduate research class that created code for our robot car entry into the DARPA Grand Challenge, I wrote assignments, created a simulation environment of upcoming DARPA site visit location, and facilitated meetings on software architecture and various brainstorming.
• Introduction to Programming in Java – led two discussion sections, organized exam reviews, wrote quizzes, graded assignments and exams
- Overall student evaluation from discussion sections: 4.7/5.0 and 4.3/5.0

• **MIT**, May 2011 – August 2011.
Visiting Student at the Media Lab
Joined Prof. Cynthia Breazeal's Personal Robots Group. Implemented TAMER on the social robot Nexi, and further work on the project is ongoing.

• **Data Mining Lab at Apple Inc.,** May 2010 – August 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.)** June 2008 – August 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.

• **Texas A&M University Center for Academic Achievement,** September 2001 – December 2001.
Supplemental Instruction leader
Prepared and led three one-hour tutoring sessions each week for an Introduction to Psychology class with more than 100 students at peak attendance.

PUBLICATIONS

Journal articles

[3] 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.

[2] 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.

[1] 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.

Book chapter

[1] 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.

Archival conference proceedings

[8] 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.

[7] 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.

[6] 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.

[5] 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.

[4] 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).**

[3] 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.

[2] 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.

[1] 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.

Archival symposium proceedings

[4] 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.

[3] 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).**

[2] 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.

[1] 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.

Refereed (but non-archival) symposium and workshop papers

[5] 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.

[4] 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.

[3] 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.

[2] W. Bradley Knox and Peter Stone. **Interactive Shaping of a Tetris Agent Using the TAMER Framework.** Technical report in *Proceedings of 2009 AAAI Robot Workshop at IJCAI-09*, July 2009.

[1] 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

[6] 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.

[5] 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.

[4] 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.

[3] 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]

[2] W. Bradley Knox. **Shaping Agents via Human Reinforcement.** In *Proceedings of The AAAI/SIGGART Doctoral Consortium at IJCAI-09*, July 2009.

[1] 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.

INVITED TALKS

- Plenary talk at AAAI 2012 Fall Symposium Series, representing the symposium Robots Learning Interactively from Human Teachers (RLIHT), 2012
- "Learning Control from Human-Generated Reward"
  - **–University of Washington**, July 2012
  - **–MIT**, April 2012
  - **–Texas A&M**, February 2013
  - **–Vecna Medical**, February 2013
- "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
  - **–University of Southern California,** Milind Tambe's TEAMCORE research group, July 2009
-

ACADEMIC ORGANIZATION

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

SERVICE

- **Program Committee**, International Conference on Machine Learning (ICML), 2012, 2013
- **Program Committee**, Association for the Advancement of Artificial Intelligence (AAAI), 2013
- **Program Committee**, International Joint Conferences on Artificial Intelligence (IJCAI), 2013
- **Reviewer**, Machine Learning Journal (MLJ), 2013, 2012, 2010 Special Issue on Empirical Evaluations in Reinforcement Learning (two reviews)
- **Reviewer**, International Journal of Social Robotics (IJSR), 2012, 2013
- **Reviewer**, Artificial Intelligence Journal (AIJ), 2012
- **Program Committee**, European Workshop on Reinforcement Learning (EWRL), 2012
- **Reviewer**, IEEE International Conference on Robotics and Automation (ICRA), 2012, 2011, 2010
- **Reviewer**, ACM/IEEE International Conference on Human-Robot Interaction (HRI), 2012, 2009
- **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 starting in April 2012.

COMMUNITY SERVICE

- **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.