

Kaushik Subramanian

Contact Information

1080 Mission Rd,
South San Francisco, CA - 94080,
USA

Mobile: +1 650 454 4701
E-mail: ksubrama@cc.gatech.edu
www: <http://www.cc.gatech.edu/~ksubrama>

Research Interests *Machine Learning, Deep Learning, Reinforcement Learning, Interactive Systems*

Education

Georgia Institute of Technology, USA August 2010 - Present
Doctor of Philosophy in Computer Science
Advised by Prof. Charles L. Isbell & Prof. Andrea L. Thomaz

Rutgers, The State University of New Jersey, USA August 2008 - May 2010
Master of Science in Electrical and Computer Engineering
Advised by Prof. Michael L. Littman

Anna University, India July 2004 - May 2008
Bachelor of Engineering in Electrical and Communication Engineering

Professional Experience

Intern at Honda Research Institute, Mountain View, USA June 2016 - August 2016
Worked as a member of the Systems & Control group on the design of decision-making algorithms and human-machine interaction for self-driving cars. Mentored by Kikuo Fujimura.

Intern at Microsoft Research, Redmond, USA June 2011 - August 2011
Completed a 3 month internship in the Computational User Experiences (CUE) group on activity recognition using machine learning for human muscle data. Mentored by Dan Morris and Desney Tan.

Intern at RWTH Aachen University, Germany June 2009 - August 2009
Completed a 3 month internship on robot learning from demonstration using Gaussian mixture models. Mentored by Prof. Gerhard Lakemeyer.

Intern at Texas Instruments, India May 2008 - July 2008
Completed a 3 month internship focusing on data compression algorithms in video encoding standards. Mentored by Ajit Gupte.

Research Experience

Research Assistant in the pFunk Machine Learning Lab, GATech August 2010 - Present
Worked with Monte Carlo methods, reinforcement learning and Bayesian methods assisted by human information to solve popular video games. Presently working on search space exploration strategies for sequential decision making problems.

Adjunct Member of the Real Life Reinforcement Learning Lab (RL3), Rutgers Jan 2009 - May 2010
Extended the theoretical properties of an existing RL algorithm and introduced a novel method for simultaneously learning the intentions of multiple people in uncontrolled environments.

Research Trainee at Waran Research FoundaTion (WARFT), India June 2006 - May 2008
Completed a 2 year research training program with specialization in signal processing with specific applications to computer vision.

Masters Dissertation

“HELP - Human assisted Efficient Learning Protocols” - an analysis into the effect of human interactions on learning algorithms, with focus on methods like learning by demonstration and apprenticeship learning. *Advisor - Prof. Michael Littman*

Publications

Journal
Luis C. Cobo, Kaushik Subramanian, Charles L. Isbell, Aaron D. Lanterman, Andrea L. Thomaz: *Abstraction from Demonstration for Efficient Reinforcement Learning in High-Dimensional Domains*. Appeared in Artificial Intelligence Journal (AIJ) 2014.

Conference

Kaushik Subramanian, Charles L. Isbell, Andrea L. Thomaz: *Exploration from Demonstration for Interactive Reinforcement Learning*. Appeared in AAMAS 2016, Singapore, May 2016.

(extended abstract) Himanshu Sahni, Brent Harrison, Kaushik Subramanian, Thomas Cederborg, Charles L. Isbell, Andrea L. Thomaz: *Policy Shaping in Domains with Multiple Optimal Policies*. Appeared in AAMAS 2016, Singapore, May 2016.

Shane Griffith, Kaushik Subramanian, Jonathan Scholz, Charles L. Isbell, Andrea L. Thomaz: *Policy Shaping: Integrating Human Feedback with Reinforcement Learning*. Appeared in NIPS 2013, Lake Tahoe, Nevada, USA, December 2013.

Monica Babes, Vukosi Marivate, Kaushik Subramanian, Michael L. Littman: *Apprenticeship Learning about Multiple Intentions*. Appeared in ICML 2011, Bellevue, Washington, USA, June 2011.

Thomas J. Walsh, Kaushik Subramanian, Michael L. Littman, Carlos Diuk: *Generalizing Apprenticeship Learning across Hypothesis Classes*. Appeared in ICML 2010, Haifa, Israel, June 2010.

Workshops

Kaushik Subramanian, Jonathan Scholz, Charles L. Isbell, Andrea L. Thomaz: *Efficient Exploration in Monte Carlo Tree Search using Human Action Abstractions*. Appeared in the FILM workshop at NIPS 2016, Barcelona, Spain, December 2016.

Baris Akgun, Kaushik Subramanian, Andrea Thomaz: *Novel Interaction Strategies for Learning from Teleoperation*. Appeared in the RLIHT Symposium at AAI 2012, Virginia, USA, November 2nd 2012.

Kaushik Subramanian, Andrea Thomaz, Charles Isbell: *Learning Options through Human Interaction*. Appeared in the ALIHT Workshop, IJCAI 2011, Barcelona, Spain, July 2011.

Kaushik Subramanian: *Task Space Behavior Learning for Humanoid Robots using Gaussian Mixture Regression*. Appeared in AAI 2010, Atlanta, USA, July 2010.

Technical Reports

Baris Akgun, Kaushik Subramanian, Jaeun Shim, Andrea Thomaz : *Learning Tasks and Skills Together From a Human Teacher*. Appeared in AAI 2011, San Francisco, USA, August 7th 2011.

Kaushik Subramanian and Michael Littman: *Efficient Apprenticeship Learning with Smart Humans*. Appeared in AAI 2010, Atlanta, USA, July 2010.

Professional Service

Organizing Committee Member for

- Adaptive and Learning Agents workshop at AAMAS 2017.
- Future of Interactive Learning Machines workshop at NIPS 2016.
- Interactive Machine Learning workshop at IJCAI 2016. Sponsored by the Artificial Intelligence Journal (AIJ, an Elsevier publication).

Program Committee Member for

- Adaptive Learning Agents workshop at AAMAS (2014 to 2016)
- Interactive Machine Learning workshop at IUI 2013
- Symposium on Robots Learning Interactively from Human Teachers at AAI 2012

Invited Reviewer for

- Autonomous Agents and Multi-Agent Systems (AAMAS) Journal 2017
- Adaptive Learning Agents 2014 Special Issue of Connection Science (Taylor & Francis Journal)
- Joint Conference on Development, Learning and Epigenetic Robotics, ICDL-EPIROB 2014
- International Conference on Intelligent Robots and Systems, IROS 2014
- Wiley Statistical Analysis and Data Mining Journal 2013

Volunteer for

- Autonomous Agents and Multiagent Systems (AAMAS) 2016 held in Singapore

- International Conference on Machine Learning (ICML) 2013 held in Atlanta, USA

Teaching Experience

Instructor for undergraduate Machine Learning (CS 4641) offered at GATech Summer 2015

Trusted Tester for Online MS in CS degree provided by GATech, Udacity and AT&T Spring 2014
Course - CS 7641 (Machine Learning) by Prof. Charles Isbell and Prof. Michael Littman

Teaching Assistant for

- CS 8803, Reinforcement Learning and Decision Making by Prof. Charles Isbell and Prof. Michael Littman (Online MS in CS degree course) Spring 2017
- CS 4001, Computing, Society and Professionalism by Prof. Charles E. Phillips Summer 2013
- CS 4641, Machine Learning by Prof. Charles Isbell Spring 2012

Honors and Awards

Travel Scholarship from the Graduate Student Council (GSC) in the CS department at Georgia Tech for NIPS 2016 in Barcelona, Spain.

Travel Scholarship for the Autonomous Agents and Multiagent Systems (AAMAS) 2016 conference held in Singapore.

Innovation Finalist at the Georgia Tech Research Innovation Conference (GTRIC) 2014 (1 of 9 finalists chosen from over 50 applicants).

Oral presentation slot given to one of four papers, out of 150 submissions, at the Reinforcement Learning and Decision Making Conference (RLDM) 2013.

IJCAI AI Video Award for Best Narration on an “Introduction to Reinforcement Learning” video made with Prof. Michael Littman in 2009.

Technical Skills

Programming - Python, Matlab
Softwares - Numpy/Scipy, Weka
SCMs - Git, Svn

Additional Training

Attended a two week summer school on “Deep Learning, Feature Learning” at UCLA, Los Angeles, USA in July 2012.

Worked on several personal projects using Arduino prototyping boards.

Coursework

Advanced Internet Computing	Prof. Ling Liu, GATech Spring 2015
Statistical Methods	Prof. Alexander Shapiro, GATech Spring 2014
Iterative Methods for Systems of Equations	Prof. Edmond Chow, GATech Fall 2013
Computability and Algorithms	Prof. Milena Mihail, GATech Spring 2013
Measurements in Human Integrated Systems	Prof. Karen Feigh, GATech Fall 2011
Human Robot Interaction	Prof. Andrea Thomaz, GATech Spring 2011
Robot Intelligent Planning	Prof. Mike Stilman, GATech Fall 2010
Machine Learning	Prof. Casimir Kulikowski, Rutgers Spring 2010
Artificial Intelligence	Prof. Casimir Kulikowski, Rutgers Fall 2008
Parallel and Distributed Computing	Prof. Manish Parashar, Rutgers Fall 2008