Shivaram Kalyanakrishnan: Curriculum Vitae

Department of Computer Science
The University of Texas at Austin
1616 Guadalupe St Suite 2.408
Austin TX 78701 USA

Office: ENS 32SE06
Phone: 1 512 925 0877
E-mail: shivaram@cs.utexas.edu
www: http://www.cs.utexas.edu/~shivaram

Research Summary

My research is motivated by the goal of creating intelligent systems, with an emphasis on learning, which facilitates adaptation in unknown environments. In particular I investigate the challenges behind matching the capabilities of theoretically-grounded algorithms with the practically relevant demands of applications. My Ph.D. thesis focuses on sequential decision making problems, and makes contributions to methods such as batch and model-based reinforcement learning, policy search methods, multi-armed bandits, and multiagent learning. Application domains include robot soccer, Tetris, and humanoid robotic control. In addition to my thesis, I have worked on the problem of fall prediction in humanoid robots; a relevant patent is pending. I actively participate and contribute to events such as RoboCup and the Reinforcement Learning Competition.

Education

University of Texas at Austin (Aug 2004 - present)

Ph.D., Computer Science
Anticipated Date of Graduation: Summer of 2011.
Advisor: Peter Stone
GPA: 3.79/4.00

Indian Institute of Technology Madras (July 2000 - July 2004)

B.Tech., Computer Science and Engineering
Project Supervisor: Deepak Khemani
CGPA: 9.22/10.00

Academic Experience

Graduate Research Assistantship

Summer 2005 - present: **Reinforcement learning**-related research under Peter Stone at UT Austin, summarized in publications.

Teaching Assistantship

Fall 2007: **"Reinforcement Learning: Theory and Practice"** taught by Peter Stone at UT Austin, involving interaction with graduate students, planning and teaching parts of the course.

Fall 2005: **"Elements of Databases"** taught by Glen Nuckolls at UT Austin, involving grading and monitoring Oracle programming laboratory.

Fall 2004 - Spring 2005: **"Contemporary Issues in Computer Science"** taught by John Messerly at UT Austin, involving grading student essays on subjects such as philosophy, ethics, and AI.

Spring 2003: **"Introduction to Computing"** taught by N.S.Narayanaswamy at IIT Madras, involving grading and monitoring C programming laboratory.

Professional Internship

July – October 2008: **Honda Research Institute**, Mountain View, CA, U.S.A.; developed a machine learning solution for humanoid fall prediction, in collaboration with Ambarish Goswami.

May – July 2003: **Tejas Networks**, Bangalore, India; examined solutions for "hot-swapping" storage devices in a network running under Linux.

Professional Activities

Invited Talks

CORAL Research Group, Computer Science Department, Carnegie Mellon University, Pittsburgh, PA, USA, 2011.

Department of Computer Science and Engineering, Indian Institute of Technology Bombay, Mumbai, India, 2010.

Germinait Solutions Private Limited, Mumbai, India, 2010.

Department of Computer Science and Engineering, Indian Institute of Technology Madras, Chennai, India, 2010.

Department of Computer Science and Engineering, Indian Institute of Technology Delhi, New Delhi, India, 2010.

ICML 2010 Workshop on Reinforcement Learning and Search in Very Large Spaces, Haifa, Israel, 2010.

Interactive Intelligence Lab, Department of Computer Science and Engineering, Indian Institute of Technology Madras, Chennai, India, 2009.

Programme Committee/Invited Reviewing

9^{th} European Workshop on Reinforcement Learning, 2011.

24^{th} Annual Conference on Learning Theory, 2011.

RoboCup International Symposium 2011, 2011.

2011 International Conference on Artificial Neural Networks, 2011.

Artificial Intelligence Journal, 2010.

2011 IEEE International Conference on Robotics and Automation, 2010.

Adaptive and Learning Agents Workshop at AAMAS 2010, 2010.

RoboCup International Symposium 2009, 2009.

26^{th} International Conference on Machine Learning, 2009.

Journal of Artificial Intelligence Research, 2008.

10^{th} International Conference on Intelligent Autonomous Systems, 2008.

AAMAS 2006 Workshop on Adaptation and Learning in Autonomous Agents and Multiagent Systems, 2006.

Competitions

UT Austin Villa team: **1^{st} place**, RoboCup 3D Simulation Competition, Istanbul, Turkey, 2011.

UT Austin Villa team, RoboCup 3D Simulation Competition, Singapore, 2010.

UT Austin Villa team, RoboCup 3D Simulation Competition, Suzhou, China, 2008.

LARG team: **1^{st} place**, Tetris Event, First Annual Reinforcement Learning Competition, 2007.

UT Austin Villa team, RoboCup 3D Simulation Competition, Atlanta, GA, USA, 2007.

UT Austin Villa team: **2^{nd} place**, RoboCup Simulation Coach Competition, Bremen, Germany, 2006.

Organizational

Technical Committee: Second Annual Reinforcement Learning Competition, 2008.

Coordinator: UTCS Reinforcement Learning Reading Group, Spring 2006 - Spring 2011.

Publications

Journal articles

Characterizing Reinforcement Learning Methods through Parameterized Learning Problems, Shivaram Kalyanakrishnan and Peter Stone. *Machine Learning, 84(1–2): 205–247, July 2011.*

Learning to Predict Humanoid Fall, Shivaram Kalyanakrishnan and Ambarish Goswami. *International Journal of Humanoid Robotics, 8(2): 245–273, June 2011.*

Highly reviewed conferences

On Optimizing Interdependent Skills: A Case Study in Simulated 3D Humanoid Robot Soccer, Daniel Urieli, Patrick MacAlpine, Shivaram Kalyanakrishnan, Yinon Bentor and Peter Stone, *In Kagan Tumer, Pinar Yolum, Liz Sonenberg and Peter Stone, Proceedings of the Tenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2011), pp. 769–776, IFAAMAS, 2011.*

Efficient Selection of Multiple Bandit Arms: Theory and Practice, Shivaram Kalyanakrishnan and Peter Stone, *In Johannes Fürnkranz and Thorsten Joachims, editors, Proceedings of the Twenty-seventh International Conference on Machine Learning (ICML 2010), pp. 511–518, Omnipress, 2010.*

Predicting Falls of a Humanoid Robot through Machine Learning, Shivaram Kalyanakrishnan and Ambarish Goswami, *In Nestor Rychtyckyj and Daniel Shapiro, editors, Proceedings of the Twenty-second IAAI Conference on Artificial Intelligence (IAAI 2009), pp. 1793–1798, AAAI, 2010.*

An Empirical Analysis of Value Function-Based and Policy Search Reinforcement Learning, Shivaram Kalyanakrishnan and Peter Stone, *In Carles Sierra, Cristiano Castelfranchi, Keith S. Decker, and Jaime Simão Sichman, editors, Proceedings of the the Eighth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2009), pp. 749–756, IFAAMAS, 2009.*

Batch Reinforcement Learning in a Complex Domain, Shivaram Kalyanakrishnan and Peter Stone, *In Edmund H. Durfee, Makoto Yokoo, Michael N. Huhns, and Onn Shehory, editors, Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2007), pp. 650–657, IFAAMAS, 2007.*

Book Chapters

Three Humanoid Soccer Leagues: Comparison and Synthesis, Shivaram Kalyanakrishnan, Todd Hester, Michael Quinlan, Yinon Bentor and Peter Stone, *In Jacky Baltes, Michail G. Lagoudakis, Tadashi Naruse, and Saeed Shiry Ghidary, editors, RoboCup-2009: Robot Soccer World Cup XIII, pp. 140–152. Short paper.*

Learning Complementary Multiagent Behaviors: A Case Study, Shivaram Kalyanakrishnan and Peter Stone, *In Jacky Baltes, Michail G. Lagoudakis, Tadashi Naruse, and Saeed Shiry Ghidary, editors, RoboCup-2009: Robot Soccer World Cup XIII, pp. 153–165.*

Model-based Reinforcement Learning in a Complex Domain, Shivaram Kalyanakrishnan, Peter Stone and Yaxin Liu, *In Ubbo Visser, Fernando Ribeiro, Takeshi Ohashi, and Frank Dellaert, editors, RoboCup-2007: Robot Soccer World Cup XI, pp. 171–183, Springer Verlag, Berlin, 2008.*

Half Field Offense in RoboCup Soccer: A Multiagent Reinforcement Learning Case Study, Shivaram Kalyanakrishnan, Yaxin Liu and Peter Stone, *In Gerhard Lakemeyer, Elizabeth Sklar, Domenico Sorrenti, and Tomoichi Takahashi, editors, RoboCup-2006: Robot Soccer World Cup X, pp. 72–85, Springer Verlag, Berlin, 2007.*

Workshops and Symposia

On Learning with Imperfect Representations, Shivaram Kalyanakrishnan and Peter Stone, *In Proceedings of the 2011 IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning (ADPRL 2011), pp. 17–24, IEEE, 2011.*

Integrating Value Function-Based and Policy Search Methods for Sequential Decision Making, Shivaram Kalyanakrishnan and Peter Stone, *Multidisciplinary Symposium on Reinforcement Learning (MSRL 2009). Extended abstract.*

Technical Reports

The UT Austin Villa 3D Simulation Soccer Team 2008, Shivaram Kalyanakrishnan, Yinon Bentor and Peter Stone, *Technical Report AI09-01, The University of Texas at Austin, Department of Computer Science, AI Laboratory, 2009.*

The UT Austin Villa 3D Simulation Soccer Team 2007, Shivaram Kalyanakrishnan and Peter Stone, *Technical Report AI07-348, The University of Texas at Austin, Department of Computer Science, AI Laboratory, 2007.*

Patents

A Machine Learning Approach for Predicting Humanoid Robot Fall, Ambarish Goswami and Shivaram Kalyanakrishnan. *Patent pending.*

Honours and Awards

Best Student Paper Award, RoboCup International Symposium 2009, Graz, Austria. Paper title: *Learning Complementary Multiagent Behaviors: A Case Study.*

Nominee for Best Student Paper Award, AAMAS 2007, Honolulu, Hawai'i, USA. Paper title: *Batch Reinforcement Learning in a Complex Domain.*

Best Student Paper Award, RoboCup International Symposium 2006, Bremen, Germany. Paper title: *Half Field Offense in RoboCup Soccer: A Multiagent Reinforcement Learning Case Study.*

All-India Rank 75, IIT Joint Entrance Examination 2000 (out of approximately 200,000 students).

First Rank in high school class of 2000 (of approximately 60), Vidya Mandir Adyar, with a score of 483/500 (Central Board of Secondary Education).

Recipient of the **National Talent Search scholarship** awarded by the Government of India (to the top 0.01% of students), 1999.

Graduate Course Work

Distributed Computing I (Jayadev Misra, Fall 2006)
Programming Languages (William Cook, Spring 2006)
Computer Graphics (Bill Mark, Fall 2005)
Computational Learning Theory (Adam Klivans, Fall 2005)
Numerical Analysis: Int., App., Quad., Diff. Eq. (Alan Cline, Spring 2005)
Knowledge Representation ad Reasoning (Bruce Porter, Spring 2005)
Introduction to Mathematical Logic (Vladimir Lifschitz, Fall 2004)
Machine Learning (Raymond J. Mooney, Fall 2004)
Supervised Teaching in Computer Science (Roger Priebe, Fall 2004)

Citizenship

India