

Shivaram Kalyanakrishnan: Curriculum Vitae

CSA Department
Indian Institute of Science
Bangalore 560012 India

Phone: +91 80 2293 2468 (Ext: 256)
E-mail: shivaram@csa.iisc.ernet.in
www: <http://www.csa.iisc.ernet.in/~shivaram>

Research Summary

I am a computer scientist with a specialisation in artificial intelligence. My research is motivated by the goal of creating intelligent agents, especially ones that can learn. In pursuit of this goal, I consider both theoretical and empirical questions, in areas such as sequential decision making, multiagent learning, multi-armed bandits, and humanoid robotics. I have applied my research in domains such as robot soccer, computer games, and on-line advertising.

Appointments

- **Indian Institute of Science.** *INSPIRE Faculty Fellow*, Department of Computer Science and Automation. July 2014–present.
- **Yahoo Labs Bangalore.** *Research Scientist*. February 2012–June 2014.
- **Honda Research Institute,** *Intern*. July–October 2008.
- **Tejas Networks.** *Intern*. May–July 2003.

Education

- **University of Texas at Austin.** August 2004–December 2011.
Ph.D., Computer Science
Advisor: Peter Stone
GPA: 3.7925/4.00
- **Indian Institute of Technology Madras.** July 2000–July 2004.
B.Tech., Computer Science and Engineering
Project Supervisor: Deepak Khemani
CGPA: 9.22/10.00

Honours and Awards

- **INSPIRE Faculty Fellowship**, awarded by the Government of India, December 2013.
- **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.

Teaching Assistantship

- “**Reinforcement Learning: Theory and Practice**”, graduate course taught by Peter Stone, UT Austin, Fall 2007.
- “**Elements of Databases**”, undergraduate course taught by Glen Nuckolls, UT Austin, Fall 2005.
- “**Contemporary Issues in Computer Science**”, undergraduate course taught by John Messerly, UT Austin, Fall 2004–Spring 2005.
- “**Introduction to Computing**”, undergraduate course taught by N. S. Narayanaswamy, IIT Madras, Spring 2003.

Professional Activities

INVITED TALKS

- “CS 729: Topics in Machine Learning”, upper-division undergraduate and graduate course, Department of Computer Science and Engineering, Indian Institute of Technology Bombay, Mumbai, India, September 2014. *An Introduction to Stochastic Multi-armed Bandits*.
- “CS 621: Artificial Intelligence”, upper-division undergraduate and graduate course, Department of Computer Science and Engineering, Indian Institute of Technology Bombay, Mumbai, India, September 2014. *RoboCup: A Grand Challenge for AI*.
- Department of Computer Science and Engineering, Indian Institute of Technology Delhi, New Delhi, India, April 2014. *PAC Subset Selection in Stochastic Multi-armed Bandits*.
- Research Promotion Workshop on Machine Learning and Social Networks, Department of Computer Science and Engineering, Indian Institute of Technology Kharagpur, Kharagpur, India, March 2014. *Panel Discussion: Research as a Career in Computer Science*.
- Research Promotion Workshop on Machine Learning and Social Networks, Department of Computer Science and Engineering, Indian Institute of Technology Kharagpur, Kharagpur, India, March 2014. *PAC Subset Selection in Stochastic Multi-armed Bandits*.
- “CS60073: Advanced Machine Learning”, graduate course, Department of Computer Science and Engineering, Indian Institute of Technology Kharagpur, Kharagpur, India, March 2014. *Tutorial: Exploration and Multi-armed Bandits*.
- Machine Learning Seminar Series, Department of Computer Science and Engineering, Indian Institute of Technology Kanpur, Kanpur, India, March 2014. *PAC Subset Selection in Stochastic Multi-armed Bandits*.
- Indo-US Lectures Week in Machine Learning, Game Theory and Optimization, Indo-US Joint Center for Advanced Research in Machine Learning, Game Theory and Optimization, Bengaluru, India, January 2014. *PAC Subset Selection in Stochastic Multi-armed Bandits*.
- School of Technology and Computer Science, Tata Institute of Fundamental Research, Mumbai, India, November 2013. *PAC Subset Selection in Stochastic Multi-armed Bandits*.
- NLP-AI Lecture, Department of Computer Science and Engineering, Indian Institute of Technology Bombay, Mumbai, India, October 2013. *Tutorial: Exploration and Multi-armed Bandits*.
- “CS 725: Foundations of Machine Learning”, upper-division undergraduate and graduate course, Department of Computer Science and Engineering, Indian Institute of Technology Bombay, Mumbai, India, October 2013. *Tutorial: An Introduction to Reinforcement Learning*.
- Department of Computer Science and Engineering, Indian Institute of Technology Bombay, Mumbai, India, October 2013. *PAC Subset Selection in Stochastic Multi-armed Bandits*.
- Department of Computer Science and Engineering, Indian Institute of Technology Madras, Chennai, India, October 2013. *PAC Subset Selection in Stochastic Multi-armed Bandits*.
- Intelligent Autonomous Systems Group, Informatics Institute, University of Amsterdam, Teleconference, August 2013. *PAC Subset Selection in Stochastic Multi-armed Bandits*.

- Forum for Artificial Intelligence, The University of Texas at Austin, Austin, TX, USA, August 2013. *PAC Subset Selection in Stochastic Multi-armed Bandits*.
- “E0 270: Machine Learning”, graduate course, Department of Computer Science and Automation, Indian Institute of Science, Bengaluru, India, March 2013. *An Introduction to Reinforcement Learning*.
- Machine Learning Special Interest Group, Department of Computer Science and Automation, Indian Institute of Science, Bengaluru, India, October 2012. *PAC Subset Selection in Stochastic Multi-armed Bandits*.
- CORAL Research Group, Computer Science Department, Carnegie Mellon University, Pittsburgh, PA, USA, April 2011. *Learning Methods for Sequential Decision Making in Practice*.
- Department of Computer Science and Engineering, Indian Institute of Technology Bombay, Mumbai, India, December 2010. *Learning Methods for Sequential Decision Making in Practice*.
- Germinait Solutions Private Limited, Mumbai, India, December 2010. *Learning Methods for Sequential Decision Making in Practice*.
- Department of Computer Science and Engineering, Indian Institute of Technology Madras, Chennai, India, December 2010. *Learning Methods for Sequential Decision Making in Practice*.
- Department of Computer Science and Engineering, Indian Institute of Technology Delhi, New Delhi, India, December 2010. *Learning Methods for Sequential Decision Making in Practice*.
- ICML 2010 Workshop on Reinforcement Learning and Search in Very Large Spaces, Haifa, Israel, June 2010. *Learning Methods for Sequential Decision Making in Practice*.
- Interactive Intelligence Lab, Department of Computer Science and Engineering, Indian Institute of Technology Madras, Chennai, India, February 2009. *An Empirical Analysis of Value Function-Based and Policy Search Reinforcement Learning*.

SENIOR PROGRAMME COMMITTEE

- 28th AAAI Conference, 2014.
- 23rd International Joint Conference on Artificial Intelligence, 2013.

PROGRAMME COMMITTEE/INVITED REVIEWING

- 2014 ACM International Conference on Information and Knowledge Management, 2014.
- AAAI 2014 Workshop on Multiagent Interaction without Prior Coordination, 2014.
- 2014 Conference on Learning Theory, 2014.
- Journal of Machine Learning Research, 2012.
- Journal of Autonomous Agents and Multi-Agent Systems, 2012.
- International Journal of Adaptive Control and Signal Processing, 2012.
- 10th European Workshop on Reinforcement Learning, 2012.
- Neural Networks, 2012.
- 29th International Conference on Machine Learning, 2012.
- Adaptive and Learning Agents Workshop at AAMAS 2012, 2012.
- 9th European Workshop on Reinforcement Learning, 2011.
- 24th 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.
- 26th International Conference on Machine Learning, 2009.
- Journal of Artificial Intelligence Research, 2008.
- 10th 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: **1st 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: **1st 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: **2nd place**, RoboCup Simulation Coach Competition, Bremen, Germany, 2006.

ORGANISATIONAL

- Organiser: Yahoo! Labs IISc Student Seminar, 2013.
- Technical Committee: Second Annual Reinforcement Learning Competition, 2008.
- Coordinator: UTCS Reinforcement Learning Reading Group, Spring 2006–Spring 2011.

Publications

JOURNAL ARTICLES

- **Direction-Changing Fall Control of Humanoid Robots: Theory and Experiments**, Ambarish Goswami, Seung-kook Yun, Umashankar Nagarajan, Sung-Hee Lee, KangKang Yin, and Shivaram Kalyanakrishnan, *Autonomous Robots*, 36(3): 199–223, March 2014.
- **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 Building Decision Trees from Large-scale Data in Applications of On-line Advertising**, Shivaram Kalyanakrishnan, Deepthi Singh, and Ravi Kant, *In Proceedings of the 2014 ACM International Conference on Information and Knowledge Management (CIKM 2014)*. To appear.
- **GEV-Canonical Regression for Accurate Binary Class Probability Estimation when One Class is Rare**, Arpit Agarwal, Harikrishna Narasimhan, Shivaram Kalyanakrishnan, and Shivani Agarwal, *JMLR Workshop and Conference Proceedings (International Conference on Machine Learning, 2014)*, 32(1): 1989–1997, 2014.
- **Information Complexity in Bandit Subset Selection**, Emilie Kaufmann and Shivaram Kalyanakrishnan, *JMLR Workshop and Conference Proceedings (Conference on Learning Theory, 2013)*, 30: 228–251, 2013.

- **PAC Subset Selection in Stochastic Multi-armed Bandits**, Shivaram Kalyanakrishnan, Ambuj Tewari, Peter Auer, and Peter Stone, *In John Langford and Joelle Pineau, Editors, Proceedings of the Twenty-ninth International Conference on Machine Learning (ICML 2012)*, pp. 655–662, Omnipress, 2012.
- **UT Austin Villa 2011: A Champion Agent in the RoboCup 3D Soccer Simulation Competition**, Patrick MacAlpine, Daniel Urieli, Samuel Barrett, Shivaram Kalyanakrishnan, Francisco Barrera, Adrian Lopez-Mobilia, Nicolae Ştiurcă, Victor Vu, and Peter Stone, *In Vincent Conitzer, Michael Winikoff, Wiebe van der Hoek, and Lin Padgham, Editors, Proceedings of the Eleventh International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, pp. 129–136, IFAAMAS, 2012.
- **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, Editors, 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 Rychtycky 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

- **UT Austin Villa 2011 3D Simulation Team Report**, Patrick MacAlpine, Daniel Urieli, Samuel Barrett, Shivaram Kalyanakrishnan, Francisco Barrera, Adrian Lopez-Mobilia, Nicolae Știurcă, Victor Vu, and Peter Stone, *Technical Report AI11-10, The University of Texas at Austin, Department of Computer Science, AI Laboratory, 2011*.
- **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

- **Machine Learning Approach for Predicting Humanoid Robot Fall**, Ambarish Goswami and Shivaram Kalyanakrishnan. *US Patent 8,554,370, issued October 8, 2013*.

Biographical

- **Date of Birth:** January 3, 1983.
- **Citizenship:** India.
- **Languages:** Tamil, English, Hindi.