

Shivaram Kalyanakrishnan: Curriculum Vitae

Department of Computer Science and Engineering
Indian Institute of Technology Bombay
Mumbai 400076 India

Phone: +91 22 2576 7704
E-mail: shivaram@cse.iitb.ac.in
www: <http://www.cse.iitb.ac.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 Technology Bombay.** *Associate Professor*, Department of Computer Science and Engineering. December 2018–present.
- **Indian Institute of Technology Bombay.** *Assistant Professor*, Department of Computer Science and Engineering. February 2015–December 2018.
- **Indian Institute of Science.** *INSPIRE Faculty Fellow*, Department of Computer Science and Automation. July 2014–January 2015.
- **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

- **Prof. Krithi Ramamritham Award for Creative Research** for the year 2016, awarded by IIT Bombay, November 2017.
- Among **AI's 10 to Watch**, list published by IEEE Intelligent Systems, January 2016.
- **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

- “**Advances in Intelligent and Learning Agents**”, graduate course, IIT Bombay, Spring 2021.
- “**Foundations of Intelligent and Learning Agents**”, graduate course, IIT Bombay, Autumn 2020.
- “**Certificate Program in Machine Learning & AI**”, CEP course, IIT Bombay, Summer 2020.
- “**Advances in Intelligent and Learning Agents**”, graduate course, IIT Bombay, Spring 2020.
- “**Foundations of Intelligent and Learning Agents**”, graduate course, IIT Bombay, Autumn 2019.
- “**Artificial Intelligence and Machine Learning**”, undergraduate course, IIT Bombay, Spring 2019.
- “**Foundations of Intelligent and Learning Agents**”, graduate course, IIT Bombay, Autumn 2018.
- “**Artificial Intelligence**”, undergraduate course, IIT Bombay, Spring 2018.
- “**Foundations of Intelligent and Learning Agents**”, graduate course, IIT Bombay, Autumn 2017.
- “**Theoretical Analysis of Policy Iteration**”, tutorial presented at the Twenty-sixth International Joint Conference on Artificial Intelligence (IJCAI 2017), Melbourne, Australia, August 2017.
- “**Artificial Intelligence**”, undergraduate course, IIT Bombay, Spring 2017.
- “**Foundations of Intelligent and Learning Agents**”, graduate course, IIT Bombay, Autumn 2016.
- “**Advances in Intelligent and Learning Agents**”, graduate course, IIT Bombay, Spring 2016.
- “**Foundations of Intelligent and Learning Agents**”, graduate course, IIT Bombay, Autumn 2015.

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.

Thesis Supervision

- Durgesh Samant, M.S., 2020, *Towards building agents that sense, explain and act optimally in RL settings.*
- Rohit Prasad, M.Tech., 2020, *Optimising a Real-time Scheduler for Railway Networks using Policy Search.*
- Suraj Singh, M.Tech., 2020, *Online Reinforcement Learning for Autonomous Driving.*
- Deep Karkhanis, B.Tech., 2020, *Tractable Policy Iteration in POMDPs.*
- Divyansh, B.Tech., 2019. *A Smartphone-based Document-scanning System.*
- Archit Gupta, B.Tech., 2019. *PAC-Optimal MDP Planning Algorithms.*
- Nilesh Gupta, B. Tech., 2019. *Learning complex behaviours and Keepaway in 3D Robocup Environment.*
- Rishabh Shah, B.Tech., 2019 (co-supervisor: Parag Chaudhuri). *Training Primitive Skills to Soccer Bots.*
- A. Srinath Naik, B.Tech., 2018 (co-supervisor: Siddhartha Chaudhuri). *Smart phone based book digitization system.*
- Vishal B. Bhavani, B.Tech., 2018 (co-supervisor: Siddhartha Chaudhuri). *Document Image De-warping using Deep Learning.*
- Samiran Roy, M.Tech., 2017. *A Bandit-based Framework for Self-evaluation in Autonomous Learning Agents.*
- A. Siddharth, M.Tech., 2017. *On the Effect of the Frequency of Decision Making in Temporal Difference Learning.*
- Mihir Kulkarni, B.Tech./M.Tech., 2017. *Barbicels: Towards a world-class Scrabble agent.*
- Anchit Gupta, B.Tech., 2017 (co-supervisor: Supratik Chakraborty). *Improved Strong Worst-case Upper Bounds for MDP Planning.* Joint winner, **Departmental Research Excellence Award**, B.Tech. Project category.
- Harshad Chavan, M.Tech., 2016 (primary supervisor: Pushpak Bhattacharyya). *Study of Reordering in Pivot Based SMT.* Nominal supervisor.
- Deepak Patil, M.Tech., 2016 (primary supervisor: Pushpak Bhattacharyya). *Study of Reordering in Pivot based Statistical Machine Translation.* Nominal supervisor.
- Shivam Garg, B.Tech., 2016 (primary supervisor: N. Hemachandra). *Adaptive Sliding-Window Based Approach for Multi-Armed Bandits.*

Professional Activities

SERVICE

- Consultant to E-committee, Supreme Court of India on National Judicial Data Grid, 2018.
- Member, Task force on Strategic Implementation of AI on National Security and Defence Needs, Department of Defence Production, Government of India, 2018.
- Member, 2015 Study Panel, One Hundred Year Study on Artificial Intelligence (AI100), 2016.

ORGANISATIONAL

- Chair, Demonstrations Program, 31st AAAI Conference on Artificial Intelligence, 2017.
- Organiser, Yahoo! Labs IISc Student Seminar, 2013.
- Technical Committee, 2nd Annual Reinforcement Learning Competition, 2008.
- Coordinator, UTCS Reinforcement Learning Reading Group, Spring 2006–Spring 2011.

SENIOR PROGRAMME COMMITTEE

- 25th International Joint Conference on Artificial Intelligence, 2016.
- 28th AAAI Conference, 2014.
- 23rd International Joint Conference on Artificial Intelligence, 2013.

PROGRAMME COMMITTEE/INVITED REVIEWING

- Neural Computation, 2019.
- 14th Inter-Research-Institute Student Seminar in Computer Science, IIT Gandhinagar, 2019.
- 32nd Annual Conference on Learning Theory, 2019.
- Sadhana, 2018.
- 5th Indian Control Conference, 2018.
- Artificial Intelligence Journal, 2018.
- 28th International Conference on Automated Planning and Scheduling, 2018.
- 31st Annual Conference on Neural Information Processing Systems, 2017.
- Journal of Artificial Intelligence Research, 2016.
- 30th Annual Conference on Neural Information Processing Systems, 2016.
- IEEE Transactions on Cognitive and Developmental Systems, 2016.
- 3rd ACM IKDD Conference on Data Science, 2016.
- Journal of Machine Learning Research, 2015.
- 32nd International Conference on Machine Learning, 2015.
- 2015 ACM SIGMOD/PODS Conference, 2015.
- AAAI 2015 Workshop on Multiagent Interaction without Prior Coordination, 2015.
- 23rd 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.

INVITED TALKS

- Faculty Development Programme on Data Science and Data Analytics, Indian Institute of Technology Gandhinagar, Gandhinagar, India, February 2020. *Deep Reinforcement Learning*.
- Ubisoft India, Pune, India, January 2020. *Deep Reinforcement Learning*.
- Kishinchand Chellaram College, Mumbai, India, December 2019. *Artificial Intelligence: Opportunities and Challenges*.
- Indian Institute of Management Kozhikode, Kozhikode, India, September 2019. *Deep Reinforcement Learning*.
- Cotton Corporation of India, Navi Mumbai, India, July 2019. *Relevance of AI/ML to the Cotton Corporation of India*.
- Department of Electronics Engineering, Rizvi College of Engineering, Mumbai, India, July 2019. *Artificial Intelligence: A Natural Pursuit*.
- Forum for Medical Ethics Society (FMES), Mumbai, India, July 2019. Panelist, *Emerging AI technology in health care in India, health equity and justice: Critical reflections and charting out way forward*.
- IMCs India Calling Conference, IMC (Chamber of Commerce and Industry), Mumbai, India, March 2019. Panelist, *AI the Game Changer—An Overview on Opportunities and Changing the Way of Doing Business*.
- Symposium on Artificial Intelligence/Machine Learning: Science and Society, Indian Academy of Sciences, Bengaluru, India, February 2019. *Towards a Better Theoretical Understanding of AI: The Case of Policy Iteration*.
- Computational Optimization Laboratory, Stanford University, Stanford, USA, October 2018. *Improved Strong Upper Bounds for Policy Iteration*.
- Stanford Intelligent Systems Laboratory, Stanford University, Stanford, USA, October 2018. *Reinforcement Learning with Imperfect Representations*.
- School of Technology and Computer Science, Tata Institute of Fundamental Research, Mumbai, India, August 2018. *Improved Strong Upper Bounds for Policy Iteration*.
- Mazagon Dock Shipbuilders Limited, Mumbai, India, May 2018. *Artificial Intelligence: A Natural Pursuit*.
- BARC-BRNS Workshop on Machine Learning and Artificial Intelligence, Bhabha Atomic Research Centre, Mumbai, May 2018. *Reinforcement Learning*.
- BARC-BRNS Workshop on Machine Learning and Artificial Intelligence, Bhabha Atomic Research Centre, Mumbai, May 2018. *Exploration and On-line Learning*.
- International Conference on Knowledge Organization, Library and Information Management, Chennai, India, October 2017. *The Explore-Exploit Tradeoff in Recommender Systems*.
- RoboVR, Mumbai, India, November 2017. Panelist, *Emerging Robotics and Robot Sports in India & World*.
- In-Service Course for PGT Computer Science, Kendriya Vidyalaya IIT Powai, Mumbai, India, May 2017. *Artificial Intelligence: A Natural Pursuit*.
- THINK Research Club, Vidyalankar Institute of Technology, Mumbai, India, March 2017. *Deep Reinforcement Learning*.
- Mumbai Artificial Intelligence Meetup, Mumbai, India, February 2017. *Deep Reinforcement Learning*.
- Carnegie India Global Technology Summit, Bengaluru, India, December 2016. *Artificial Intelligence: A Natural Pursuit*.

- Microsoft Machine Learning & Data Sciences Conference, Bengaluru, India, August 2016. *PAC Subset Selection in Stochastic Multi-Armed Bandits*.
- 2nd Indian Workshop on Machine Learning, Kanpur, India, July 2016. *Deep Reinforcement Learning*.
- Sixth Indo-American Frontiers of Science Symposium (IAFOS-2015), Irvine, USA, August 2015. *Reinforcement Learning*.
- Workshop on Non-convex Optimization for Machine Learning, Indian Institute of Technology Bombay, Mumbai, India, June 2015. *Improved Expected Running Time for MDP Planning*.
- Recent Advances in Reinforcement Learning Workshop 2015, Indian Institute of Technology Madras, Chennai, India, March 2015. *Learning with Imperfect Representations*.
- Recent Advances in Reinforcement Learning Workshop 2015, Indian Institute of Technology Madras, Chennai, India, March 2015. *An Improved Bound for MDP Planning*.
- Tutorial and Workshop on Learning and Related Probabilistic Applications, Tata Institute of Fundamental Research, Mumbai, India, February 2015. *Improved Expected Running Time for MDP Planning*.
- “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*.

COMPETITIONS

- IIT Bombay team: **1st place**, 2048 Controller Competition, GECCO 2015, Madrid, Spain, 2015.
- 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.

Publications

JOURNALS

- **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.

CONFERENCES

- **Regret Minimisation in Multi-Armed Bandits Using Bounded Arm Memory**, Arghya Roy Chaudhuri and Shivaram Kalyanakrishnan, *In Proceedings of the Thirty-fourth AAAI Conference on Artificial Intelligence (AAAI 2020)*. To appear. Acceptance rate: 20.6%
- **A Tighter Analysis of Randomised Policy Iteration**, Meet Taraviya and Shivaram Kalyanakrishnan, *In Proceedings of the 2019 Conference on Uncertainty in Artificial Intelligence (UAI 2019)*, ID 174, AUAI Press, 2019. Acceptance rate: 26.3%
- **PAC Identification of Many Good Arms in Stochastic Multi-Armed Bandits**, Arghya Roy Chaudhuri and Shivaram Kalyanakrishnan, *In Proceedings of the Thirty-sixth International Conference on Machine Learning (ICML 2019)*, pp. 991–1000, PMLR, 2019. Acceptance rate: 22.6%
- **Quantile-Regret Minimisation in Infinitely Many-Armed Bandits**, Arghya Roy Chaudhuri and Shivaram Kalyanakrishnan, *In Proceedings of the 2018 Conference on Uncertainty in Artificial Intelligence (UAI 2018)*, pp. 425–434, AUAI Press, 2018. Acceptance rate: 30.9%
- **Opportunities and Challenges for Artificial Intelligence in India**, Shivaram Kalyanakrishnan, Rahul Alex Panicker, Sarayu Natarajan, and Shreya Rao, *In Proceedings of the First AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society (AIES 2018)*, pp. 164–170, ACM, 2018. Acceptance rate: 37.0%
- **Improved Strong Worst-case Upper Bounds for MDP Planning**, Anchit Gupta and Shivaram Kalyanakrishnan, *In Proceedings of the Twenty-sixth International Joint Conference on Artificial Intelligence (IJCAI 2017)*, pp. 1788–1794, IJCAI, 2017. Acceptance rate: 26.0%
- **PAC Identification of a Bandit Arm Relative to a Reward Quantile**, Arghya Roy Chaudhuri and Shivaram Kalyanakrishnan, *In Proceedings of the Thirty-first AAAI Conference on Artificial Intelligence (AAAI 2017)*, pp. 1777–1783, AAAI Press, 2017. Acceptance rate: 24.7%
- **Batch-Switching Policy Iteration**, Shivaram Kalyanakrishnan, Utkarsh Mall, and Ritish Goyal, *In Proceedings of the Twenty-fifth International Joint Conference on Artificial Intelligence (IJCAI 2016)*, pp. 3147–3153, AAAI Press, 2016. Acceptance rate: 25%
- **Randomised Procedures for Initialising and Switching Actions in Policy Iteration**, Shivaram Kalyanakrishnan, Neeldhara Misra, and Aditya Gopalan, *In Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI 2016)*, pp. 3145–3151, AAAI Press, 2016. Acceptance rate: 26%
- **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 Twenty-third ACM International Conference on Information and Knowledge Management (CIKM 2014)*, pp. 669–678, ACM, 2014. Acceptance rate: 21%
- **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. Acceptance rate: 25%
- **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. Acceptance rate: N/A
- **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. Acceptance rate: 27.3%

- **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 Știurcă, 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. Acceptance rate: 20.4%
- **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. Acceptance rate: 22.1%
- **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. Acceptance rate: 25.5%
- **Predicting Falls of a Humanoid Robot through Machine Learning**, Shivaram Kalyanakrishnan and Ambarish Goswami, *In Nestor Rychtickyj and Daniel Shapiro, Editors, Proceedings of the Twenty-second IAAI Conference on Artificial Intelligence (IAAI 2009)*, pp. 1793–1798, AAAI, 2010. Acceptance rate: N/A
- **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 Eighth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2009)*, pp. 749–756, IFAAMAS, 2009. Acceptance rate: 22.3%
- **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. Acceptance rate: 22.4%

WORKSHOPS AND SYMPOSIA

- **Half Field Offense: An Environment for Multiagent Learning and Ad Hoc Teamwork**, Matthew Hausknecht, Prannoy Mupparaju, Sandeep Subramanian, Shivaram Kalyanakrishnan, and Peter Stone, *In Proceedings of the Adaptive and Learning Agents Workshop 2016, 2016*. Available at http://ala2016.csc.liv.ac.uk/ALA2016_Proceedings.pdf.
- **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.
- **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, Springer, 2010. 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, Springer, 2010.
- **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. Available at <http://msrl09.rl-community.org/abstracts>.
- **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.*

TECHNICAL REPORTS

- **Artificial Intelligence and Life in 2030**, Peter Stone, Rodney Brooks, Erik Brynjolfs-son, Ryan Calo, Oren Etzioni, Greg Hager, Julia Hirschberg, Shivaram Kalyanakrishnan, Ece Kamar, Sarit Kraus, Kevin Leyton-Brown, David Parkes, William Press, AnnaLee Saxenian, Julie Shah, Milind Tambe, and Astro Teller, *One Hundred Year Study on Artificial Intelligence: Report of the 2015-2016 Study Panel, Stanford University, Stanford, CA, September 2016.*
- **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.*
- **Learning Methods for Sequential Decision Making with Imperfect Representations**, Shivaram Kalyanakrishnan, *Ph.D. dissertation, published as UT Austin Computer Science Technical Report TR-11-41, 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.

(Updated April 27, 2021.)