

CURRICULUM VITAE

Richard S. Sutton

June 20, 2026

Address: 7-238 University Commons, University of Alberta, Edmonton, AB T6G 2E8, Canada

Web: <http://richsutton.com>

Degrees

- Doctor of Philosophy in Computer Science, University of Massachusetts at Amherst, 1984.
Thesis: “Temporal credit assignment in reinforcement learning.” Advisor: Andrew Barto
- Master of Science in Computer Science, University of Massachusetts at Amherst, 1980.
Thesis: “An adaptive network that constructs and uses an internal model of its world.”
Advisor: Andrew Barto
- Bachelor of Arts in Psychology, Stanford University, with honors and distinction, 1978.
Senior thesis: “A unified theory of expectation in classical and instrumental conditioning.”
- Honorary Doctor of Philosophy, University of Alberta, 2026

Current Employment and Professional Activities

- Quarter-time Professor, Department of Computing Science, University of Alberta, 2017–.
- Research Scientist and Co-founder, Oak Lab, 2026–.
- Advisor to ExperienceFlow.ai, Intrepid Growth Partners, Cisco, Equilibre, Nirvanic, and Artificial Agency.
- (pro bono) Chief Scientist, Openmind Global Research, including Director of Research of the Openmind Robot Interaction and Learning Research Lab in Beijing, a partnership with Tashan Embodied Intelligence, 2026—.

Past Employment and Professional Activities

- Research Scientist, Keen Technologies, 2023–2026.
- Co-founder, Openmind Research Institute, 2023.
- Distinguished Research Scientist, DeepMind Alberta, 2017–2023. Co-founded DeepMind’s first satellite research laboratory, in Edmonton Alberta.
- Co-founder, Knoggin Technologies Inc, 2017.

- Professor, Department of Computing Science, University of Alberta, 2003–2017
Founded and directed the Reinforcement Learning and Artificial Intelligence Laboratory, now comprising ten faculty PIs and more than 100 people.
- Principal Technical Staff Member, Artificial Intelligence Department, AT&T Shannon Laboratory, Florham Park, New Jersey, 1998–2002.
- Senior Research Scientist, Department of Computer Science, University of Massachusetts at Amherst, 1995–1998. Co-led (with A. Barto) the Adaptive Networks Laboratory.
- Principal Member of Technical Staff, Computer and Intelligent Systems Laboratory, GTE Laboratories, Waltham, Massachusetts, 1985–1994. From 1986 led the Connectionist Machine Learning Project, a group of 3-6 conducting basic and applied research.
- Senior Postdoctoral Research Associate, Department of Computer Science, University of Massachusetts at Amherst, 1984. Advisor: Andrew Barto.

Research

Pioneered and made repeated contributions to reinforcement learning, an approach to artificial and natural intelligence that emphasizes learning from first-person experience.

Most significant contributions (chronological):

- Temporal-difference models of animal learning (with Barto, 1981, 1990)
- Actor-critic and policy-gradient algorithms (1984, 2000)
- The theory of temporal-difference learning and the $TD(\lambda)$ algorithm (1988)
- The Dyna architecture for integrated learning, planning and reacting (1990, 2008)
- The IDBD algorithms for meta-gradient step-size optimization (1992)
- The standard textbook for reinforcement learning (with Barto, 1998, 2018)
- The options framework for temporal abstraction (with Precup, Singh, 1999)
- Gradient temporal-difference algorithms (with Maei, Szepesvari, 2008–10)
- The GVF framework for general efficient prediction (with others, 2011)
- Emphatic temporal-difference algorithms (with Mahmood, White, 2016)
- The bitter lesson of AI research (2019)
- The Oak/STOMP strategy for open-ended abstraction in state and time (with others, 2023–25)
- Continual backpropagation and loss of plasticity in deep learning (with others, 2024)

Selected Grants

- Amii RAP Grant, “Continual and Meta Supervised Learning,” May 2026, \$83,810.
- Amii RAP Grant, “Constructing State Representations for Continual Reinforcement Learning,” May 2025, \$83,875.
- CIFAR CCAI Chair Grant, 2021-2025, CAD\$300,000 over five years. Renewed at CAD\$500,000 2026-2030.
- DeepMind Research Grant, 2017-2022, funded at USD\$400,000/year.

- Google Research Grant, 2015, 2016, funded at USD\$200,000/year.
- iCORE Chair and Professorship Establishment Grant, “Reinforcement Learning and Artificial Intelligence,” September 1, 2003 – August 31, 2008, funded at CAD\$3,000,000. Principal investigator. Renewed through August 2013 at an additional CAD\$2,750,000. Renewed through August 2018 at an additional CAD\$3,000,000.
- Alberta Ingenuity Centre Grant, “Alberta Ingenuity Centre for Machine Learning,” April 2003 – March 2008, funded at CAD\$9,887,600. One of eight principal investigators. Renewed until March 2009 at CAD\$2,000,000. Renewed in 2009 for another five years at CAD\$10,000,000. Renewed in 2014 for an additional two years at CAD\$2,000,000/year.
- NSERC Discovery Grant, “Reinforcement Learning and Artificial Intelligence,” April, 2004 – March 2009, funded at CAD\$250,000. Principal investigator. Renewed in 2009 for five years at CAD\$190,000. Renewed in 2014 for five years at CAD\$310,000. Renewed in 2023 for five years at CAD\$205,000.
- CFI-LEF equipment grant, “Centre for Neural Interfaces and Rehabilitation Neuroscience,” 2013, funded at CAD\$3,076,491. One of 10 principal users.
- CIHR Collaborative Health Research Project (NSERC partnered) on “Robotic interface with Haptic Guidance and Artificial Intelligence for People with Disabilities,” 2014, funded at CAD\$111,000. One of eight investigators.
- NSERC Collaborative Research and Development Grant, with Nortel Networks and Bell Canada, “Learning and Prediction in High-dimensional Stochastic Domains,” September 2006 – August 2009, funded at CAD\$186,523. One of five principal investigators.
- Air Force Office of Scientific Research to the University of Massachusetts, “Stochastic Scheduling and Planning Using Reinforcement Learning,” 1996–2000, funded at USD\$446,570. Co-principal investigator with A. Barto.
- National Science Foundation to the University of Massachusetts, “Multiple Time Scale Reinforcement Learning,” 1995–1998, funded at USD\$157,261. Primary senior personnel (A. Barto, PI).

Honors

- *Honorary doctorate in Computing Science from the University of Alberta*, 2026
- *2026 IEEE Frank Rosenblatt Award* (with Andrew Barto), 2025
- *2024 ACM A. M. Turing Award* (with Andrew Barto), 2025
- *Fellow of the Royal Society of London for the Improvement of Natural Knowledge*, 2021
- *IJCAI Research Excellence Award*, 2021
- *Lifetime Achievement Award, Canadian Artificial Intelligence Association*, 2018
- *Fellow of the Royal Society of Canada*, 2016

- *Fellow of the Association for the Advancement of Artificial Intelligence*, 2001
- *President's Award, International Neural Network Society*, 2003
- *Outstanding Achievement in Research Award, University of Massachusetts Amherst, College of Computer Science*, 2013
- *Creative Destruction Lab's Ideas Award* (with Hinton and Bengio), 2016
- *Outstanding Leadership in Alberta Technology Award* from the Alberta Science and Technology Leadership Foundation, 2006 (one of eight principal investigators)
- *Classical Paper Award, Artificial Intelligence Journal* (with Precup and Singh), 2019
- *King Charles III Coronation Medal*, for contributions in Technology and Innovation, 2025
- *Warner Award*, GTE's highest award for technical achievement, 1994

Academic Outreach and Impact

Summer Schools

- Instructor at Reinforcement Learning Summer Schools, Montreal (2017), Toronto (2018), Edmonton (2019, 2025), Virtual (2021, 2022)
- Instructor at Machine Learning Summer Schools, Ile de Rey (2008), Austin (2015)
- Instructor at Cambridge University Neural Networks Summer School, 1993–1997
- Instructor at Cold Spring Harbor Summer School on Computational Neuroscience: Learning and Memory, July 1990

Tutorials on Reinforcement Learning at Conferences

- Conference on Neural Information Processing Systems, 2015
- Conference of the Association for the Advancement of Artificial Intelligence, 2010
- International Conference on Machine Learning, 1999
- Genetic Programming Conference, 1998
- Second Asia-Pacific Conference on Simulated Evolution and Learning, Australia, 1998
- National Conference on Artificial Intelligence, with L. Kaelbling, 1997
- National Conference on Computer Science – Mexico, 1997
- Conference on Neural Information Processing Systems, 1996, 2015

Meetings Organized

- Executive Committee, *Multidisciplinary Conference on Reinforcement Learning and Decision Making*, 2013, 2015 (area chair), 2017 (general chair), 2019, 2022, 2025
- Co-chair, *Multidisciplinary Symposium on Reinforcement Learning*, 2009
- Program co-chair, *National Conference on Artificial Intelligence*, 2002
- Chair, Workshop on *Reinforcement Learning*, Int'l. Conference on Machine Learning, 1993
- Co-chair, *Bellairs Workshop on Reinforcement Learning*, 2005—2020, 2023

- Co-chair, Workshop on *Reinforcement Learning–Benchmarks and Bakeoffs*, International Conference on Neural Information Processing Systems, 2004
- Co-chair, Workshop on *Predictive Representation of World Knowledge*, International Conference on Machine Learning, 2004
- Co-chair, NSF Workshop on *Neural Networks for Robotics and Control*, 1988

Editorial

- *Artificial Intelligence Journal*: associate editor, 2013–2020
- *Foundations and Trends in Machine Learning*: founding editorial board, 2007
- *Journal of Machine Learning Research*: editorial board, advisory board 1999–
- *J. of Artificial Intelligence Research*: action editor 1993–1996, advisory board 1996–2000
- *Machine Learning*: action editor 1989–94, editorial board 1994–2001
- *Connection Science*: action editor 1989–90, editorial board 1990–1995
- *Adaptive Behavior*: editorial board 1991–2000
- Guest Editor, special double issue of *Machine Learning* on Reinforcement Learning, 1992

Academic Service

- Schwartz Reisman Institute for Technology and Society Advisory Committee, 2021 —
- IEEE John von Neumann Medal Committee, 2023 —
- Awards/Fellows Committee of the Association for the Advancement of Artificial Intelligence, 2008–2011
- Executive Council, Association for the Advancement of Artificial Intelligence, 2007–2010
- Chief Scientific Advisor, Alberta Machine Intelligence Institute (Amii), 2018 —

Conference Program Committees

- International Conference on Machine Learning
- Neural Information Processing Systems
- Association for the Advancement of Artificial Intelligence
- International Joint Conference on Artificial Intelligence
- Multidisciplinary Conference on Reinforcement Learning and Decision Making
- International Conference on Learning Representations, 2019, 2020
- Conference on Uncertainty in Artificial Intelligence, 2005, 2009
- International Conference on the Simulation of Adaptive Behavior, 1990, 1992
- Cognitive Science Conference, 1986, 1987
- Autonomous Agents and Multi-Agent Systems, 2005
- Conference on Computational Learning Theory, 1996
- Canadian Conference on Artificial Intelligence, 2006
- Conference on Artificial General Intelligence, 2012

Courses

- *Reinforcement Learning II*, a second course on reinforcement learning for graduate students at the University of Alberta, Winter 2020–26

- *Intelligent Systems*, an introduction to artificial intelligence for undergraduates at the University of Alberta, Fall 2008–2017
- *Reinforcement Learning for Artificial Intelligence*, a graduate course at the University of Alberta, Fall 2003–2017
- *Non-procedural Programming Languages*, an undergraduate course at the University of Alberta, Fall 2006
- *Reinforcement Learning in Practice*, a graduate course at the University of Alberta, Winter 2005
- *Reinforcement Learning*, a graduate course at the University of Massachusetts, with A. Barto, Fall 1996 and 1997
- *Reinforcement Learning*, a special intensive course at the University of Uppsala, Sweden, April–May 1996
- *Reinforcement Learning*, a graduate seminar at the University of Massachusetts, with A. Barto, Fall 1995
- *Cybernetics of Adaptation and Learning*, assisted Prof. A. Barto, including some lecturing, with this graduate seminar at the University of Massachusetts, Fall 1981

Advanced Student Supervision and Outcomes

Currently Under Supervision at the University of Alberta

- Huizhen Yu, Research Associate
- Tian Tian, PhD
- Farzane Aminmasour, PhD
- Edan Meyer, PhD
- Henry Du, PhD
- Parash Rahman, PhD
- Andrew Freeman, MSc
- Patrick Spieler
- Erica Graves
- Andrew Freeman, MSc

Doctoral Students Supervised and Graduated

1. Doina Precup, “Temporal abstraction in reinforcement learning,” University of Massachusetts, 2000. Currently *associate professor at McGill University and research scientist at Google DeepMind, Montreal*
2. David Silver, “Reinforcement Learning and Simulation-Based Search in Computer Go,” University of Alberta, 2009. Currently *research scientist, Google Deepmind, London*
3. Hamid Maei, “Gradient Temporal-Difference Learning Algorithms,” University of Alberta, 2011. Currently at *Netflix*

4. Adam White, “Developing a Predictive Approach to Knowledge,” University of Alberta, 2015. Currently *associate professor, University of Alberta*
5. A. Rupam Mahmood, “Incremental Off-policy Reinforcement Learning Algorithms,” University of Alberta, 2017. Currently *associate professor, University of Alberta*
6. Sina Ghiassian, “Online Off-policy Prediction,” University of Alberta, 2022. Currently *research scientist, Spotify*
7. Katya Kudashkina, “Model-based Reinforcement Learning Methods for Developing Intelligent Assistants,” University of Guelph, 2022. Currently *Director of Engineering, Intelligent Search and Assistant, Dayforce*
8. Yi Wan, “Learning and Planning with the Average-Reward Formulation,” University of Alberta, 2023. Currently *Research Scientist at a startup*
9. Alexandra Kearney, “Letting the Agent Take the Wheel: Principles for Constructive and Predictive Knowledge,” University of Alberta, 2023. Currently *founder, Artificial Agency*
10. Banafsheh Rafiee, “State Construction in Reinforcement Learning,” University of Alberta, 2024. Currently *Research Scientist at a startup*
11. Kenny Young, “Leveraging Generic Problem Structure for Efficient Reinforcement Learning,” University of Alberta, 2024. Currently *Research Scientist at a startup*
12. Abhishek Naik, “Reinforcement Learning for Continuing Problems using Average Reward,” University of Alberta, 2024. Currently *Research Scientist at the Canadian Space Agency*
13. Kris De Asis, “Explorations in the Foundations of Value-based Reinforcement Learning,” University of Alberta, 2024. Currently *Research Fellow at Openmind Research Institute*
14. Khurram Javed, “Real-time Reinforcement Learning for Achieving Goals in Big Worlds,” University of Alberta, 2024. Currently *Research Scientist at Keen Technologies*
15. Shibhansh Dohare, “Learning Forever using Artificial Neural Networks,” University of Alberta, August 2025. Currently *Research Scientist, Keen Technologies*
16. Fernando (Juan) Hernandez-Garcia, “Selective Reinitialization Algorithms for Preventing Plasticity Loss in Artificial Neural Networks,” University of Alberta, December 2025. Currently *Research Scientist, Zyphra*

Postdoctoral Fellows and Research Associates “Graduated”

- Jae Young Lee, 2017. Currently *postdoctoral fellow, University of Waterloo*
- Harm van Seijen, 2016. Currently *research scientist, Sony AI*
- Joseph Modayil, 2015. Currently *research scientist, Keen Technologies*
- Patrick Pilarski, 2014. Currently *associate professor, Division of Physical Medicine and Rehabilitation, University of Alberta*
- Hado van Hasselt, 2014. Currently *research scientist, Google Deepmind, London*

- Thomas Degris, 2011. Currently *research scientist*, *Google Deepmind*, London
- Elliot Ludvig, 2009. Currently *associate professor of psychology*, *University of Warwick*
- Eric Wiewiora, 2009. Currently *research scientist*, *Startup finance company*, Chicago
- Shalabh Bhatnagar, 2009. Currently *professor of computer science and automation*, *Indian Institute of Science*, Bangalore
- Mohammad Ghavamzadeh, 2008. Currently *senior principal scientist*, *Amazon*,
- Mark Ring, 2006. *Co-founder Cogitai*

Masters Students Supervised and Graduated

- Amir Samani, University of Alberta, 2022, Senior AI Engineer, Intel
- Brendan Bennett, University of Alberta, 2020, rl.ai
- Cam Linke, University of Alberta, 2020, CEO Amii, Alberta
- Shibhansh Dohare, University of Alberta, 2020, PhD student at U. Alberta
- Parash Rahman, University of Alberta, 2020, Machine Learning Developer at Mojow
- Jingjiao Ni, University of Alberta, 2020
- Dylan Ashley, University of Alberta, 2020, PhD student at IDSIA, Switzerland
- Valliappa Chockalingam, University of Alberta, 2020, PhD student at U. Tokyo
- Chen Ma, University of Alberta, 2020, PhD student at U. Alberta
- Kenny Young, University of Alberta, 2018, PhD student at U. Alberta
- Fernando (Juan) Hernandez Garcia, University of Alberta, 2018, PhD student at U. Alberta
- Kris De Asis, University of Alberta, 2018, PhD student at U. Alberta
- Vivek Veeriah, University of Alberta, 2017, Research Scientist Google DeepMind
- Banafsheh Rafiee, University of Alberta, 2018, PhD student at U. Alberta
- Tian Tian, University of Alberta, 2018, PhD student at U. Alberta
- Shangdong Zhang, University of Alberta, 2018, Assistant Professor U. Virginia
- Kavosh Asadi, University of Alberta, 2015, AI Research Scientist Amazon
- Travis Dick, University of Alberta, 2014, Research Scientist Google
- Leah Hackman, University of Alberta, 2012, PhD student at U. Alberta
- Christian Denk, MSc, Technical University of Munich, 2012
- MahdiehSadat Miriam HosseinAbadi, University of Alberta, 2011, Microsoft
- Michael Delp, University of Alberta, 2010, Toyota Research
- Ashique Mahmood, University of Alberta, 2010, Assistant professor at U. Alberta
- Masoud Shahamiri, University of Alberta, 2008, Software Development Engineer at AWS
- Anna Koop, University of Alberta, 2007, Research Engineer, Google DeepMind
- Eddie Rafols, University of Alberta, 2006, Senior Software Developer at Shopify
- Adam White, University of Alberta, 2006, Assistant professor U. Alberta, founder RLcore
- Alborz Geramifard, University of Alberta, 2006, Research Scientist Director, Meta
- Cosmin Paduraru, University of Alberta, 2006, Research engineer, DeepMind, London
- Brian Tanner, University of Alberta, 2005, Founder, Artificial Agency
- Amy McGovern, University of Massachusetts, 1998, full professor at U. Oklahoma
- Doina Precup, University of Massachusetts, 1996, associate professor at McGill

Doctoral Examination Committees

- Brett Daley, “Multistep Credit Assignment in Deep Reinforcement Learning,” University of Alberta, 2025. Currently *Research Scientist, Meta*
- Bo Liu, “From Diversity to Adaptivity: Effective Multitask Learning and Continual Learning Neural Architectures,” University of Texas at Austin, 2024
- Akhil Bagaria, “Skill Discovery for Exploration and Planning,” Brown University, 2024. Currently *Research Scientist, Amazon*
- Dongqi Han, “Toward a Cognitive Neurorobotic Agent That Can Abstract, Infer and Plan: Reinforcement Learning and Active Inference in Hierarchical and Partially Observable Tasks,” Okinawa Institute of Science and Technology, 2022
- Josiah Hanna, “Data Efficient Reinforcement Learning with Off-policy and Simulated Data,” University of Texas at Austin, 2019. Currently *Assistant Professor, University of Wisconsin*
- Marlos Machado, “Efficient Exploration in Reinforcement Learning through Time-Based Representations,” University of Alberta, 2018. Currently *Assistant Professor, University of Alberta*
- Ashley Dalrymple, “Machine Learning to Characterize Motor Patterns and Restore Walking after Neural Injury,” University of Alberta, 2018. Currently *Assistant Professor, University of British Columbia*
- Michael Thon, “Spectral Learning of Sequential Systems,” Jacobs University Bremen, 2017. Currently *computer vision engineer* at MOBIS Parts Europe N.V.
- Lakshmanan, K. “Online-Learning and Simulation based Algorithms for Stochastic Optimization,” Indian Institute of Science, 2012. Currently at *Associate Professor, CSE, Indian Institute of Science (BHU) Varanasi*
- Sun Yi, “On Generation of Representations for Reinforcement Learning,” Università della Svizzera Italiana, 2012. Became *Software Engineer, Google*
- Yasin Abbasi-Yadkori, “Online Learning for Linearly Parametrized Control Problems,” University of Alberta, 2012. Currently *Research Scientist, Sapient Intelligence*
- Erik Schuitema, “Reinforcement learning on autonomous humanoid robots,” Technical University of Delft, 2012. Currently *software and robotics engineer, Netherlands*
- Harm van Seijen, “Reinforcement Learning under Space and Time Constraints,” University of Amsterdam, 2011. Currently *Research Scientist, SonyAI*
- Amir-massoud Farahmand, “Regularization in Reinforcement Learning,” University of Alberta, 2011. Currently *Associate Professor, Polytechnique Montréal*
- Hado van Hasselt, “Insights in Reinforcement Learning,” SIKS, Dutch Research School for Information and Knowledge Systems, 2011. Currently *research scientist, Google Deepmind*
- Matthew Taylor, “Autonomous Inter-Task Transfer in Reinforcement Learning Domains,” University of Texas at Austin, 2008. Currently *Assistant Professor, University of Alberta*

- Tao Wang, “New Representations and Approximations for Sequential Decision Making under Uncertainty,” University of Alberta, 2007. Became *Machine Learning Specialist, San Francisco Bay area*
- Mohammad Al-Ansari, Northeastern University, 2001. Became *Director of Software Development, Oracle Health Sciences, Boston*
- Manfred Huber, “A hybrid architecture for adaptive robot control,” University of Massachusetts, 2000. Currently *Professor, Department of Computer Science and Engineering, University of Texas at Arlington*
- Gavin Rummery, “Problem solving with reinforcement learning,” Cambridge University, 1995. Became *Chief Technical Officer and Founder, Legendary Games, Nottingham, UK*
- Chen Tham, “Online function approximation for scaling up reinforcement learning,” Cambridge University, 1994. Became *Associate Professor, Department of Electrical and Computer Engineering, National University of Singapore*
- Jing Peng, “Efficient dynamic programming based learning for control,” Northeastern University, 1994. Became *Associate Professor, Department of Computer Science, Montclair State University*
- Satinder Singh, “Learning to solve Markovian decision processes,” University of Massachusetts, 1993. Currently *Senior Research Scientist, Google DeepMind, and Professor, Electrical Engineering and Computer Science Department, University of Michigan*
- Long-Ji Lin, “Reinforcement learning for robots using neural networks,” Carnegie-Mellon University, 1992. Became *Chief Scientist, Rocket Fuel, USA*
- Steve Whitehead, “Reinforcement learning for the adaptive control of perception and action,” University of Rochester, 1992
- Leslie Kaelbling, “Learning in embedded systems,” Stanford University, 1990. Currently *professor, Department of Electrical Engineering and Computer Science, MIT*

Selected Invited Presentations

- “The Future of AI: The Era of Experience and the Age of Design,” *83rd Tsinghua University Top-Talk Seminar*, Beijing, November 26, 2025
- “The OaK Architecture: A Vision of SuperIntelligence without Bitterness,” *Neural Information Processing Systems*, December 3, 2025
- “Shaping the Future of AI and Reinforcement Learning,” *National University of Singapore 120th Anniversary Distinguished Lecture Series*, June 6, 2025
- “The OaK Architecture: A Vision of SuperIntelligence from Experience,” *Reinforcement Learning Conference*, August 8, 2025
- “Decentralized Neural Networks,” *International Conference on Distributed Artificial Intelligence*, 2024
- “A Perspective on Intelligence,” *PTJC 20th Anniversary Celebration*, Sept 25, 2024

- “The Alberta Plan for Ambitious AI Research,” *European Workshop on Reinforcement Learning*, 2022
- “Maintaining Plasticity in Deep Continual Learning,” *Conference on Life-long Learning Agents*, 2022
- “Eyes on the Prize,” closing keynote, *Amii AI Week*, 2022
- “Gaps in the Foundations of Planning with Approximation,” *International Conference on Automated Planning and Scheduling*, 2021
- “The Increasing Role of Experience in AI,” award talk, *International Joint Conference on Artificial Intelligence*, 2021
- “Advice and Perspectives on RL Research Frontiers,” *Deep Learning and Reinforcement Learning Summer School*, 2019
- “Reinforcement Theories of Learning and Thinking,” *Conference of the Cognitive Science Society*, 2017
- “The Future of Artificial Intelligence Belongs to Search and Learning,” *University of Toronto*, October 27, 2017
- “Introduction to Reinforcement Learning with Function Approximation,” invited tutorial, *Conference on Neural Information Processing Systems*, Montreal, 2015
- “Creating Human-level AI: How and When?” *Future of AI: Opportunities and Challenges*, Puerto Rico, 2015
- “Temporal-difference Learning and the Coming of Artificial Intelligence,” *Distinguished Lecture Series*, *University of Massachusetts*, Amherst, Massachusetts, USA, 2014
- “Myths of Representation Learning,” *International Conference on Learning Representations*, Banff, Canada, 2014
- “Reinforcement Learning and Psychology: A Personal Story,” *47th Annual Meeting of the Society for Mathematical Psychology*, Quebec City, Canada, 2014
- “Scaling Life-long Learning,” *European Workshop on Reinforcement Learning*, Edinburgh, United Kingdom, 2012
- “Learning about Sensorimotor data,” *Conference on Neural Information Processing Systems*, Granada, Spain, 2011.
- “Learning about Sensorimotor data,” *21st Annual Conference of the Japanese Neural Network Society*, Okinawa, Japan, 2011
- “Beyond Reward: The Problem of Knowledge and Data,” *21st International Conference on Inductive Logic Programming*, Windsor Great Park, United Kingdom, 2011
- “Toward Learning Human-level Predictive Knowledge,” *Third Conference on Artificial General Intelligence*, Lugano, Switzerland, 2010
- “Deconstructing Reinforcement Learning,” *Multi-disciplinary Symposium on Reinforcement Learning*, Montreal, 2009
- “Mind and Time: A View of Constructivist Reinforcement Learning,” *European Workshop on Reinforcement Learning*, Lille, France, 2008
- “Reinforcement Learning’s Computational Theory of Mind,” *Workshop on Dopamine and its Role in Learning, Motivation, and Psychiatric Disorders*, McGill University, Montreal, 2005

- “Experience-Oriented Artificial Intelligence,” *Mathematics and Computer Science Department, University of Lethbridge, Canada*, 2005
- “Experience-Oriented Artificial Intelligence,” *Cognitive Science Department, University of California, San Diego, USA*, 2005
- “Open Theoretical Questions in Reinforcement Learning,” *European Conference on Computational Learning Theory, Nordkirchen, Germany*, 1999
- “Reinforcement Learning: How Far Can It Go?” *joint plenary talk at the International Conference on Machine Learning, Conference on Computational Learning Theory, and Conference on Uncertainty in Artificial Intelligence, Madison, Wisconsin, USA*, 1998
- “Reinforcement Learning: Past, Present, and Future,” *Second Asia-Pacific Conference on Simulated Evolution and Learning, Canberra, Australia*, 1998
- “Between MDPs and Semi-MDPs: Learning, Planning and Representing Knowledge at Multiple Temporal Scales,” *University Simon Bolivar, Caracas, Venezuela*, 1998
- “Computational Reinforcement Learning,” *Annual Meeting of the Society for the Quantitative Analysis of Behavior, Orlando, Florida*, 1998
- “Reinforcement Learning: Lessons for AI,” *International Joint Conference on Artificial Intelligence, Nagoya, Japan*, 1997
- “On the Significance of Markov Decision Processes,” *International Conference on Artificial Neural Networks, Switzerland*, 1997
- “Learning from Interaction,” *National Conference on Computer Science – Mexico*, 1997
- “The Value of Reinforcement Learning in Neurobiology,” *Gottingen Neurobiology Conference, Germany*, 1995
- “Temporal-Difference Learning,” *National Conference on Artificial Intelligence, Washington DC, USA*, 1993
- “Temporal-Difference Learning,” *IEEE International Conference on Neural Networks, San Francisco, USA*, 1993
- “Reinforcement Learning Architectures,” *International Symposium on Neural Information Processing, Kyushu, Japan*, 1992

Publications

These works have been cited about 180,000 times in total, according to Google scholar.

Books

1. Sutton, R. S., Barto, A. G., *Reinforcement Learning: An Introduction*, expanded second edition. MIT Press, 2018. Also translated into Chinese and Japanese.
2. Sutton, R. S., Barto, A. G., *Reinforcement Learning: An Introduction*. MIT Press, 1998. Also translated into Japanese and Russian.

3. Miller, W. T., Sutton, R. S., Werbos, P. J. (Eds.), *Neural Networks for Control*. MIT Press, 1991.
4. Sutton, R. S. (Ed.), *Reinforcement Learning*. Reprinting of a special issue of *Machine Learning Journal*. Kluwer Academic Press, 1992

Journal Articles

5. Dohare, S., Hernandez-Garcia, J.F., Lan, Q., Rahman, P., Sutton, R. S., Mahmood, A.R. (2024). Loss of plasticity in deep continual learning. *Nature* 632:768-774.
6. Ghiassian, S., Rafiee, B., Sutton, R. S. (2024). “Off-policy prediction learning: An empirical study of online algorithms.” *IEEE Transactions on Neural Networks and Learning Systems* 36(3):4477-4491.
7. Mathewson, K. W., Parker, A. S., Sherstan, C., Edwards, A. L., Sutton, R. S., Pilarski, P. M. “Communicative capital: A key resource for human-machine shared agency and collaborative capacity.” *Neural Computing and Applications* 35(23):16805-16819, 2023.
8. Javed, K., Shah, H., Sutton, R. S., White, M. “Scalable real-time recurrent learning using columnar-constructive networks. *Journal of Machine Learning Research* 24(256):1-34, 2023.
9. Sutton, R. S., Machado, M. C., Holland, G. Z., Timbers, D. S. F., Tanner, B., White, A. “Reward-respecting subtasks for model-based reinforcement learning.” *Artificial Intelligence*, 2023.
10. Rafiee, B., Abbas, Z., Ghiassian, S., Kumaraswamy, R., Sutton, R., Ludvig, E., White, A., “From eye-blinks to state construction: Diagnostic benchmarks for online representation learning.” *Adaptive Behavior*, 2022.
11. Silver, D., Singh, S., Precup, D., Sutton, R. S., “Reward is enough.” *Artificial Intelligence* 299:103535, 2021.
12. Barto, A. G., Sutton, R. S., Anderson, C. W., “Looking back on the actor–critic architecture.” *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 51(1), 2021.
13. Dalrymple, A., Roszko, D., Sutton, R., Mushahwar, V., “Pavlovian control of intraspinal microstimulation to produce over-ground walking.” *Journal of Neural Engineering* 17(3):036002, 2020.
14. Lee, J. Y., Sutton, R. S., “Policy iterations for reinforcement learning problems in continuous time and space — Fundamental theory and methods,” *Automatica* 126:109421, 2021.
15. Sutton, R. S., “John McCarthy’s definition of intelligence,” *Journal of Artificial General Intelligence* 11(2), 66-67, 2020.

16. Yu, H., Mahmood, A. R., Sutton, R. S., “On generalized Bellman equations and temporal-difference learning,” *Journal of Machine Learning Research* 19, 2018.
17. Travník, J. B., Mathewson, K. W., Sutton, R. S., Pilarski, P. M., “Reactive reinforcement learning in asynchronous environments,” *Frontiers in Robotics and AI*, June 16, 2018.
18. Ludvig, E. A., Mirian, M. S., Kehoe, E. J., Sutton, R. S., “Associative learning from replayed experience.” Submitted. <http://biorxiv.org/content/early/2017/01/16/100800>, 2017.
19. van Seijen, H., Mahmood, A. R., Pilarski, P. M., Machado, M. C., Sutton, R. S., “True online temporal-difference learning,” *Journal of Machine Learning Research* 17(145):1-40, 2016.
20. Sutton, R. S., Mahmood, A. R., White, M., “An emphatic approach to the problem of off-policy temporal-difference learning,” *Journal of Machine Learning Research* 17(73):1-29, 2016.
21. Edwards, A. L., Dawson, M. R., Hebert, J. S., Sherstan, C., Sutton, R. S., Chan, K. M., Pilarski, P. M., “Application of real-time machine learning to myoelectric prosthesis control: A case series in adaptive switching,” *Prosthetics and Orthotics International*, first published September 15, 2015.
22. Kehoe, E. J., Ludvig, E. A., Sutton, R. S., “Time course of the rabbit’s conditioned nictitating membrane movements during acquisition, extinction, and reacquisition,” *Learning and Memory* 21:585–590, Cold Spring Harbor Press, 2014.
23. Modayil, J., White, A., Sutton, R. S. “Multi-timescale nexting in a reinforcement learning robot,” *Adaptive Behavior* 22(2):146-160, Sage 2014.
24. Pilarski, P. M., Degris, T., Dawson, M. R., Chan, K. M., Hebert, J. S., Carey, J. P., Sutton, R. S., “Adaptive artificial limbs: A real-time approach to prediction and anticipation,” *IEEE Robotics and Automation Magazine* 20(1):58–64. March 2013.
25. Kehoe, E. J., Ludvig, E. A., Sutton, R. S., “Timing and cue competition in conditioning of the nictitating membrane response of the rabbit (*Oryctolagus cuniculus*),” *Learning and Memory* 20:97–102, Cold Spring Harbor Press 2013.
26. Ludvig, E. A., Sutton, R. S., Kehoe, E. J. “Evaluating the TD model of classical conditioning,” *Learning and Behavior* 40(3):305–319, Springer 2012.
27. Silver, D., Sutton, R. S., Müller, M., “Temporal-difference search in computer Go,” *Machine Learning* 87(2):183–219, Springer 2012.
28. Kehoe, E. J., Ludvig, E. A., Sutton, R. S., “Timing in trace conditioning of the nictitating membrane response of the rabbit (*Oryctolagus cuniculus*): Scalar, nonscalar, and adaptive features,” *Learning and Memory* 17:600-604, Cold Spring Harbor Press 2010.

29. Kehoe, E. J., Ludvig, E. A., Sutton, R. S., “Magnitude and timing of conditioned responses in delay and trace classical conditioning of the nictitating membrane response of the rabbit (*Oryctolagus cuniculus*),” *Behavioral Neuroscience* 123(5):1095–1101, 2009.
30. Kehoe, E. J., Olsen, K. N., Ludvig, E. A., Sutton, R. S., “Scalar timing varies with response magnitude in classical conditioning of the rabbit nictitating membrane response (*Oryctolagus cuniculus*),” *Behavioral Neuroscience* 123(1):212–217, 2009.
31. Bhatnagar, S., Sutton, R. S., Ghavamzadeh, M., Lee, M., “Natural actor–critic algorithms,” *Automatica* 45(11):2471–2482, 2009.
32. Ludvig, E. A., Sutton, R. S., Kehoe, E. J., “Stimulus representation and the timing of reward-prediction errors in models of the dopamine system,” *Neural Computation* 20(12):3034–3054, 2008.
33. Kehoe, E. J., Ludvig, E. A., Dudeney, J. E., Neufeld, J., Sutton, R. S., “Magnitude and timing of nictitating membrane movements during classical conditioning of the rabbit (*Oryctolagus cuniculus*),” *Behavioral Neuroscience* 122(2):471–476, 2008.
34. Stone, P., Sutton, R. S., Kuhlmann, G., “Reinforcement Learning for RoboCup-Soccer Keepaway,” *Adaptive Behavior* 13(3):165–188, 2005.
35. Sutton, R. S., Precup, D., Singh, S., “Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning,” *Artificial Intelligence*, 112, 1999. pp. 181–211. Won *Artificial Intelligence’s Classical Paper Award* in 2019.
36. Santamaria, J. C., Sutton, R. S., Ram, A., “Experiments with reinforcement learning in problems with continuous state and action spaces,” *Adaptive Behavior*, 6, No. 2, 1998, pp. 163–218.
37. Singh, S., Sutton, R. S., “Reinforcement learning with replacing eligibility traces,” *Machine Learning*, 22, 1996, pp. 123–158.
38. Sutton, R. S., “Learning to predict by the methods of temporal differences,” *Machine Learning*, 3, 1988, No. 1, pp. 9–44.
39. Moore, J., Desmond, J., Berthier, N., Blazis, D., Sutton, R. S., Barto, A. G., “Simulation of the classically conditioned nictitating membrane response by a neuron-like adaptive element: Response topography, neuronal firing, and interstimulus intervals,” *Behavioural Brain Research*, 21, 1986, pp. 143–154.
40. Barto, A. G., Sutton, R. S., Anderson, C., “Neuron-like adaptive elements that can solve difficult learning control problems,” *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-13, No. 5, 1983, pp. 834–846.
41. Barto, A. G., Sutton, R. S., “Simulation of anticipatory responses in classical conditioning by a neuron-like adaptive element,” *Behavioral Brain Research*, 4, 1982, pp. 221–235.

42. Barto, A. G., Anderson, C., Sutton, R. S., “Synthesis of nonlinear control surfaces by a layered associative network,” *Biological Cybernetics*, 43, 1982, pp. 175–185.
43. Sutton, R. S., Barto, A. G., “Toward a modern theory of adaptive networks: Expectation and prediction,” *Psychological Review*, 88, No. 2, 1981, pp. 135–170.
44. Sutton, R. S., Barto, A. G., “An adaptive network that constructs and uses an internal model of its world,” *Cognition and Brain Theory*, 4, 1981, pp. 217–246.
45. Barto, A. G., Sutton, R. S., Brouwer, P., “Associative search network: A reinforcement learning associative memory,” *Biological Cybernetics*, 40, 1981, pp. 201–211.
46. Barto, A., Sutton, R. S., “Landmark learning: An illustration of associative search,” *Biological Cybernetics*, 42, 1981, pp. 1–8.

Book Chapters

47. Silver, D., Sutton, R. S. (in press). “Welcome to the Era of Experience.” To appear in *Designing an Intelligence*, G. Konidaris, Ed. MIT Press.
48. Tian, T., Sutton, R. S., “Extending sliding-step importance weighting from supervised learning to reinforcement learning.” *Best of IJCAI Workshops 2019*, Springer, 2020.
49. Stone, P., Sutton, R. S., “Keepaway soccer: A machine learning testbed.” In *Lecture Notes in Computer Science 2377*:207–237, Springer, 2002.
50. Stone, P., Sutton, R. S., Singh, S., “Reinforcement learning for 3 vs. 2 keepaway.” In *RoboCup-2000: Robot Soccer World Cup IV*, P. Stone, T. Balch, G. Kraetschmar, Eds. Springer Verlag, 2001.
51. Sutton, R. S. “Reinforcement learning.” In *The MIT Encyclopedia of the Cognitive Sciences*, R. Wilson, F. Keil, Eds., pp. 715–717, MIT Press, 1999.
52. Sutton, R. S., “On the significance of Markov decision processes.” In *Artificial Neural Networks – ICANN’97*, W. Gerstner, A. Germond, M. Hasler, J.-D. Nicoud, Eds., pp. 273–282. Springer, 1997.
53. Barto, A. G., Sutton, R. S., “Reinforcement learning in artificial intelligence.” In *Neural Network Models of Cognition*, Donahoe, J. W., Packard Dorsel, V., Eds., pp. 358–386, Elsevier, 1997.
54. Sutton, R. S., “First results with Dyna: An integrated architecture for learning, planning, and reacting.” In *Neural Networks for Control*, Miller, T., Sutton, R. S., Werbos, P., Eds., pp. 179–189, MIT Press, 1990.
55. Sutton, R. S., Barto, A. G., “Time-derivative models of Pavlovian reinforcement.” In *Learning and Computational Neuroscience*, M. Gabriel, J.W. Moore, Eds., pages 497–537, MIT Press, 1990.

56. Barto, A. G., Sutton, R. S., Watkins, C., “Learning and sequential decision making.” In *Learning and Computational Neuroscience*, M. Gabriel, J.W. Moore, Eds., pages 539–602, MIT Press, 1990.
57. Selfridge, O., Sutton, R. S., Anderson, C., “Selected bibliography on connectionism,” In: *Evolution, Learning, and Cognition*, Y.C. Lee, Ed., pp. 391–403, World Scientific Publishing, 1988.
58. Barto, A., Sutton, R. S., “Neural problem solving.” In: *Synaptic Modification, Neuron Selectivity, and Nervous System Organization*, W.B. Levy, J.A. Anderson, Eds., pp. 123–152, Lawrence Erlbaum, 1985.

Articles in Highly-refereed Conference Proceedings

59. Javed, K., Modayil, J. V., Kennickell, G., Sutton, R. S., Carmack, J. (2026). Physical Atari: A robust and accessible platform for real-time reinforcement learning on robots. Presented at the *Reinforcement Learning Conference*. To appear in *Reinforcement Learning Journal*. Accepted May 5, 2026.
60. Sharifnassab, A., Elsayed, E., De Asis, K., Mahmood, A. R., Sutton, R. S. (2026). Intentional updates for streaming reinforcement learning. In: *Proceedings of the International Conference on Machine Learning*. Also arXiv:2604.19033. Accepted April 30, 2026.
61. Sharifnassab, A., Salehkaleybar, S., Sutton, R. (2025). MetaOptimize: A framework for optimizing step sizes and other meta-parameters. In *42nd International Conference on Machine Learning*.
62. Javed, K., Sharifnassab, A., Sutton, R. S. (2024). SwiftTD: A fast and robust algorithm for temporal difference learning. In *Reinforcement Learning Conference*. Also *Reinforcement Learning Journal* 1-5:840-863.
63. Naik, A, Wan, Y., Tomar, M., Sutton, R. S. (2024). Reward centering. In *Reinforcement Learning Conference*. Also *Reinforcement Learning Journal* 1-5:1995-2016.
64. De Asis, K., Sutton, R. S. (2024). An idiosyncrasy of time-discretization in reinforcement learning. In *Reinforcement Learning Conference*. Also *Reinforcement Learning Journal* 1-5:1306-1316.
65. Sharifnassab, A., Sutton, R. S. “Toward efficient gradient-based value estimation.” In *International Conference on Machine Learning* (pp. 30827-30849). Also *Proceedings of Machine Learning Research* 202:30827-30849, 2023.
66. Tian, T., Young, K., Sutton, R. S., “Doubly-asynchronous value iteration: Making value iteration asynchronous in actions.” In *Advances in Neural Information Processing Systems* 35, 2022. Also arXiv:2207.01613.

67. Wan, Y., Naik, A., Sutton, R. “Average-reward learning and planning with options.” *Advances in Neural Information Processing Systems* 34, 22758-22769, 2021.
68. Zhang, S., Wan, Y., Naik, A., Sutton, R. S., Whiteson, S., “Average-reward off-policy policy evaluation with function approximation.” In *Proceedings of the International Conference on Machine Learning*, 2021.
69. Wan, Y., Naik, A., Sutton, R. S., “Learning and planning in average-reward Markov decision processes.” In *Proceedings of the International Conference on Machine Learning*, 2021.
70. De Asis, K., Chan, A., Pitis, S., Sutton, R., Graves, D., “Fixed-horizon temporal difference methods for stable reinforcement learning.” In *Proceedings of the AAAI Conference on Artificial Intelligence* 34(4):3741–3748, 2020.
71. Osband, I., Doron, Y., Hessel, M., Aslanides, J., Sezener, E., Saraiva, A., McKinney, K., Lattimore, T., Szepezsvari, C., Singh, S., Van Roy, B., Sutton, R., Silver, D., van Hasselt, H., “Behaviour suite for reinforcement learning.” In *Proceedings of the International Conference on Learning Representations*, 2020.
72. Wan, Y., Zaheer, M., White, A., White, M., Sutton, R.S., “Planning with expectation models.” In *Proceedings of the International Joint Conference on Artificial Intelligence*, 2019.
73. Rafiee, B., Ghiassian, S., White, A., Sutton, R.S., “Prediction in intelligence: An empirical comparison of off-policy algorithms on robots.” In *Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems*, 332-340, 2019.
74. Veeriah, V., van Seijen, H., Sutton, R. S., “Forward actor-critic for nonlinear function approximation in reinforcement learning.” In *Proceedings of the 16th Conference on Autonomous Agents and Multiagent Systems*, pp. 556-564, 2017.
75. De Asis, K., Hernandez-Garcia, J. F., Holland, G. Z., Sutton, R. S., “Multi-step reinforcement learning: A unifying algorithm.” *Proceedings of the Conference of the Association for the Advancement of Artificial Intelligence*, 2018.
76. De Asis, K., Sutton, R. S., “Per-decision multi-step temporal difference learning with control variates.” In *Proceedings of the 2018 Conference on Uncertainty in Artificial Intelligence*, 2018.
77. Sherstan, C., Ashley, D. R., Bennett, B., Young, K., White, A., White, M., Sutton, R. S., “Comparing direct and indirect temporal-difference methods for estimating the variance of the return.” In *Proceedings of the 2018 Conference on Uncertainty in Artificial Intelligence*, 2018.
78. Mahmood, A. R., Sutton, R. S., “Off-policy learning based on weighted importance sampling with linear computational complexity.” In *Proceedings of the 2015 Conference on Uncertainty in Artificial Intelligence*. Amsterdam, Netherlands, 2015.
79. van Seijen, H., Sutton, R.S., “A deeper look at planning as learning from replay.” In *Proceedings of the 32nd International Conference on Machine Learning*, Lille, France, 2015.

80. Mahmood, A. R., van Hasselt, H, Sutton, R. S., “Weighted importance sampling for off-policy learning with linear function approximation.” *Advances in Neural Information Processing Systems 27*, Montreal, Canada. 2014.
81. Yao, H., Modayil, J., Szepesvari, Cs., Bhatnagar, S., Sutton, R. S., “Universal Option Models.” *Advances in Neural Information Processing Systems 27*, Montreal, Canada, 2014.
82. van Hasselt, H., Mahmood, A. R., Sutton, R. S., “Off-policy TD(λ) with a true online equivalence.” In: *Proceedings of the 30th Conference on Uncertainty in Artificial Intelligence*, Quebec City, Canada, 2014.
83. Sutton, R. S., Mahmood, A. R., Precup, D., van Hasselt, H., “A new Q(λ) with interim forward view and Monte Carlo equivalence.” In *Proceedings of the 31st International Conference on Machine Learning*, Beijing, China, 2014.
84. van Seijen, H., Sutton, R. S., “True online TD(λ).” In *Proceedings of the 31st International Conference on Machine Learning*, Beijing, China, 2014.
85. van Seijen, H., Sutton, R. S., “Efficient planning in MDPs by small backups.” In *Proceedings of the 30th International Conference on Machine Learning*, pp. 361–369, Atlanta, USA, 2013.
86. Pilarski, P. M., Dick, T. B., Sutton, R. S., “Real-time prediction learning for the simultaneous actuation of multiple prosthetic joints.” In *Proc. of the 2013 IEEE International Conference on Rehabilitation Robotics*, Seattle, USA, 2013.
87. White, A., Modayil, J., Sutton, R. S., “Scaling life-long off-policy learning.” In *Proceedings of the IEEE Joint International Conference on Development and Learning—Epigenetics and Robotics*, San Diego, California, 2012.
88. Modayil, J., White, A., Pilarski, P. M., Sutton, R. S., “Acquiring a broad range of empirical knowledge in real time by temporal-difference learning.” In *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics*, Seoul, Korea, 2012.
89. Modayil, J., White, A., Sutton, R. S., “Multi-timescale nexting in a reinforcement learning robot.” In *Proceedings of the 2012 International Conference on Adaptive Behaviour, LNAI 7426*, pp. 299-309, T. Ziemke, C. Balkenius, and J. Hallam, Eds., Springer Heidelberg, 2012.
90. Degris, T., White, M., Sutton, R. S., “Off-policy actor-critic.” In *Proceedings of the 29th International Conference on Machine Learning*, Edinburgh, United Kingdom, 2012.
91. Pilarski, P. M., Dawson, M. R., Degris, T., Carey, J. P., Chan, K. M., Hebert, J. S., Sutton, R. S., “Towards prediction-based prosthetic control.” In *Proceedings of the 2012 Conference of the International Functional Electrical Stimulation Society*, September 9-12, Banff, Canada, 4 pages, 2012.
92. Pilarski, P. M., Dawson, M. R., Degris, T., Carey, J. P., Sutton, R. S., “Dynamic switching and real-time machine learning for improved human control of assistive biomedical robots.”

In *Proceedings of the Fourth IEEE International Conference on Biomedical Robotics and Biomechatronics (BioRob)*, June 24-27, Roma, Italy, 7 pages, 2012.

93. Degris, T., Pilarski, P. M., Sutton, R. S., “Model-free reinforcement learning with continuous action in practice.” In *Proceedings of the American Control Conference*, Montreal, Canada, 2012.
94. Mahmood, A. R., Sutton, R. S., Degris, T., Pilarski, P. M., “Tuning-free step-size adaptation.” In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, Kyoto, Japan, 2012.
95. Pilarski, P. M., Dawson, M. R., Degris, T., Fahimi, F., Carey, J. P., Sutton, R. S., “Online human training of a myoelectric prosthesis control via actor-critic reinforcement learning.” In *Proceedings of the 2011 IEEE International Conference on Rehabilitation Robotics*, Zurich, Switzerland, pp. 134-140, 2011.
96. Sutton, R. S., Modayil, J., Delp, M., Degris, T., Pilarski, P. M., White, A., Precup, D., “Horde: A scalable real-time architecture for learning knowledge from unsupervised sensorimotor interaction.” In *Proceedings of the Tenth International Conference on Autonomous Agents and Multiagent Systems*, pp. 761–768, Taipei, Taiwan, 2011.
97. Maei, H. R., Szepesvari, Cs., Bhatnagar, S., Sutton, R. S., “Toward off-policy learning control with function approximation.” In *Proceedings of the 27th International Conference on Machine Learning*, Haifa, Israel. Omnipress, 2010.
98. Maei, H. R., Szepesvari, Cs., Bhatnagar, S., Precup, D., Silver, D., Sutton, R. S., “Convergent temporal-difference learning with arbitrary smooth function approximation.” *Advances in Neural Information Processing Systems 22*. MIT Press, 2009.
99. Sutton, R. S., Maei, H. R., Precup, D., Bhatnagar, S., Silver, D., Szepesvari, Cs., Wiewiora, E., “Fast gradient-descent methods for temporal-difference learning with linear function approximation.” *Proceedings of the 26th International Conference on Machine Learning*, 2009.
100. Sutton, R. S., Szepesvari, Cs., Maei, H. R., “A convergent $O(n)$ algorithm for off-policy temporal-difference learning with linear function approximation,” *Advances in Neural Information Processing Systems 21*. MIT Press, 2008.
101. Ludvig, E., Sutton, R. S., Verbeek, E., Kehoe, E. J., “A computational model of hippocampal function in trace conditioning,” *Advances in Neural Information Processing Systems 21*, pp. 993–1000. MIT Press, 2008.
102. Cutumisu, M., Szafron, D., Bowling, M., Sutton, R. S., “Agent learning using action-dependent learning rates in computer role-playing games,” *Proceedings of the 4th Artificial Intelligence and Interactive Digital Entertainment Conference*, pp. 22-29, 2008.
103. Sutton, R. S., Szepesvari, Cs., Geramifard, A., Bowling, M., “Dyna-style planning with linear function approximation and prioritized sweeping,” *Proceedings of the 24th Conference on Uncertainty in Artificial Intelligence*, pp. 528–536, 2008.

104. Silver, D., Sutton, R. S., Müller, M., “Sample-based learning and search with permanent and transient memories,” *Proceedings of the 25th International Conference on Machine Learning*, 2008.
105. Bhatnagar, S., Sutton, R. S., Ghavamzadeh, M., Lee, M., “Incremental natural actor-critic algorithms,” *Advances in Neural Information Processing Systems 20*, 2007.
106. Sutton, R. S., Koop, A., Silver, D., “On the role of tracking in stationary environments,” *Proceedings of the 24th International Conference on Machine Learning*, 2007.
107. Silver, D., Sutton, R. S., Müller, M., “Reinforcement learning of local shape in the game of Go,” *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, 2007.
108. Geramifard, A., Bowling, M., Zinkevich, M., Sutton, R. S., “iLSTD: Eligibility traces and convergence analysis,” *Advances in Neural Information Processing Systems 19*, 2006.
109. Geramifard, A., Bowling, M., Sutton, R. S., “Incremental least-squares temporal difference learning,” *Proceedings of the Twenty-First National Conference on Artificial Intelligence*, pp. 356-361, 2006.
110. Precup, D., Sutton, R. S., Paduraru, C., Koop, A., Singh, S., “Off-policy learning with options and recognizers,” *Advances in Neural Information Processing Systems 18*, 2006.
111. Sutton, R. S., Rafols, E. J., Koop, A., “Temporal abstraction in temporal-difference networks,” *Advances in Neural Information Processing Systems 18*, 2006.
112. Tanner, B., Sutton, R. S., “TD(λ) networks: Temporal-difference networks with eligibility traces,” *Proceedings of the 22nd International Conference on Machine Learning (ICML)*, Bonn, Germany, 2005.
113. Rafols, E. J., Ring, M. B., Sutton, R. S., Tanner, B., “Using predictive representations to improve generalization in reinforcement learning,” *Proceedings of the 19th International Joint Conference on Artificial Intelligence*, Edinburgh, Scotland, 2005.
114. Tanner, B., Sutton, R. S., “Temporal-difference networks with history,” *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence (IJCAI)*, Edinburgh, Scotland, 2005.
115. Sutton, R. S., Tanner, B., “Temporal-difference networks,” *Advances in Neural Information Processing Systems 17*. MIT Press, 2005.
116. Littman, M. L., Sutton, R. S., Singh, S., “Predictive representations of state,” *Advances in Neural Information Processing Systems 14*. MIT Press, 2002.
117. Stone, P., Sutton, R. S., “Scaling reinforcement learning toward RoboCup soccer,” *Proceedings of the 18th International Conference on Machine Learning*, pp. 537–544. Morgan Kaufmann, 2001.

118. Precup, D., Sutton, R. S., Dasgupta, S., “Off-policy temporal-difference learning with function approximation,” *Proceedings of the 18th International Conference on Machine Learning*, pp. 417–424. Morgan Kaufmann, 2001.
119. Sutton, R. S., McAllester, D., Singh, S., Mansour, Y., “Policy gradient methods for reinforcement learning with function approximation,” *Advances in Neural Information Processing Systems 12* (Proceedings of the 1999 conference), pp. 1057–1063. MIT Press, 2000.
120. Precup, D., Sutton, R. S., Singh, S., “Eligibility traces for off-policy policy evaluation,” *Proceedings of the 17th International Conference on Machine Learning*, pp. 759–766. Morgan Kaufmann, 2000.
121. Sutton, R. S., Singh, S., Precup, D., Ravindran, B., “Improved switching among temporally abstract actions,” *Advances in Neural Information Processing Systems 11* (Proceedings of the 1998 conference), pp. 1066–1072. MIT Press, 1999.
122. Moll, R., Barto, A. G., Perkins, T. J., Sutton, R. S., “Learning instance-independent value functions to enhance local search,” *Advances in Neural Information Processing Systems 11* (Proceedings of the 1998 conference), pp. 1017–1023. MIT Press, 1999.
123. Sutton, R. S., Precup, D., Singh, S., “Intra-option learning about temporally abstract actions,” *Proceedings of the 15th International Conference on Machine Learning*, pp. 556–564. Morgan Kaufmann, 1998.
124. Precup, D., Sutton, R. S., Singh, S., “Theoretical results on reinforcement learning with temporally abstract options,” *Machine Learning: ECML-98, Proceedings of the 10th European Conference on Machine Learning*, pp. 382–393. Springer Verlag, 1998.
125. Precup, D., Sutton, R. S., “Multi-time models for temporally abstract planning,” *Advances in Neural Information Processing Systems 10* (Proceedings of the 1997 conference), pp. 1050–1056. MIT Press, 1998.
126. Precup, D., Sutton, R. S., “Exponentiated gradient methods for reinforcement learning,” *Proceedings of the 14th International Conference on Machine Learning*, pp. 272–277. Morgan Kaufmann, 1997.
127. Sutton, R. S., “Generalization in reinforcement learning: Successful examples using sparse coarse coding,” *Advances in Neural Information Processing Systems 8* (Proceedings of the 1995 conference), pp. 1038–1044. MIT Press, 1996.
128. Sutton, R. S. “TD models: Modeling the world at a mixture of time scales,” *Proceedings of the Twelfth International Conference on Machine Learning*, pp. 531–539. Morgan Kaufmann, 1995.
129. Sutton, R. S., Whitehead, S. “Online learning with random representations,” *Proceedings of the Tenth International Conference on Machine Learning*, pp. 314–321. Morgan Kaufmann, 1993.

130. Gluck, M., Glauthier, P., Sutton, R. S. "Adaptation of cue-specific learning rates in network models of human category learning," *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, pp. 540–545. Erlbaum, 1992.
131. Sutton, R. S. "Adapting bias by gradient descent: An incremental version of Delta-Bar-Delta," *Proceedings of the Tenth National Conference on Artificial Intelligence*, pp. 171–176. MIT/AAAI Press, 1992.
132. Sanger, T., Sutton, R. S., Matheus, C., "Iterative construction of sparse polynomial approximations," *Advances in Neural Information Processing Systems 4*, pp. 1064–1071. Morgan Kaufmann, 1992.
133. Sutton, R. S., "Planning by incremental dynamic programming," *Proceedings of the Eighth International Workshop on Machine Learning*, pp. 353–357. Morgan Kaufmann, 1991.
134. Sutton, R. S., Matheus, C., "Learning polynomial functions by feature construction," *Proceedings of the Eighth International Workshop on Machine Learning*, pp. 208–212. Morgan Kaufmann, 1991.
135. Sutton, R. S., Barto, A. G., Williams, R., "Reinforcement learning is direct adaptive optimal control," *Proceedings of the American Control Conference*, pp. 2143–2146, 1991. Also published in *IEEE Control Systems Magazine* 12, No. 2, 19–22, 1992.
136. Sutton, R. S., "Integrated modeling and control based on reinforcement learning and dynamic programming," *Advances in Neural Information Processing Systems 3*, pp. 471–478. MIT Press, 1991.
137. Sutton, R. S., "Reinforcement learning architectures for animats," *Proceedings of the First International Conference on Simulation of Adaptive Behavior: From Animals to Animats*, pp. 288–296, 1991.
138. Sutton, R. S., "Integrated architectures for learning, planning, and reacting based on approximating dynamic programming," *Proceedings of the Seventh International Conference on Machine Learning*, pp. 216–224, 1990. Reprinted in *Artificial Neural Networks: Concepts and Theory*, edited by P. Mehra and B. Wah, IEEE Computer Society Press, 1992.
139. Barto, A. G., Sutton, R. S., Watkins, C., "Sequential decision problems and neural networks," *Advances in Neural Information Processing Systems 2*, pp. 686–693. MIT Press, 1990.
140. Sutton, R. S., Barto, A. G., "A temporal-difference model of classical conditioning," *Proceedings of the Ninth Annual Conference of the Cognitive Science Society*, pp. 355–378, 1987.
141. Sutton, R. S., "Two problems with backpropagation and other steepest-descent learning procedures for networks," *Proceedings of the Eighth Annual Conference of the Cognitive Science Society*, pp. 823–831, 1986. Reprinted in *Artificial Neural Networks: Concepts and Theory*, edited by P. Mehra and B. Wah, IEEE Computer Society Press, 1992.

142. Selfridge, O., Sutton, R. S., Barto, A. G., “Training and tracking in robotics,” *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, pp. 670–672, 1985.
143. Sutton, R. S., Pinette, B., “The learning of world models by connectionist networks,” *Proceedings of the Seventh Annual Conference of the Cognitive Science Society*, pp. 54–64, 1985.
144. Moore, J., Desmond, J., Berthier, N., Blazis, D., Sutton, R. S., Barto, A. G., “Connectionist learning in real time: Sutton-Barto adaptive element and classical conditioning of the nictitating membrane response,” *Proceedings of the Seventh Annual Conference of the Cognitive Science Society*, pp. 318–322, 1985.

Other Articles in Proceedings

145. Javed, K., Sutton, R. (2025). Swift-Sarsa: Fast and Robust Linear Control. In: *Multi-disciplinary Conference on Reinforcement Learning And Decision-Making (RLDM)*. Also arXiv:2507.19539.
146. Hernandez-Garcia, J.F., Dohare, S., Luo, J., Sutton, R.S. (2025). “Reinitializing weights vs units for maintaining plasticity in neural networks.” In *Proceedings of the Fourth Conference on Lifelong Learning Agents*, 2025. Also ArXiv:2508.00212.
147. Rafiee, B., Ghiassian, S., Jin, J., Sutton, R., Luo, J., White, A. “Auxiliary task discovery through generate-and-test.” In *Proceedings of the Second Conference on Lifelong Learning Agents*, 2023. Also arXiv:2210.14361.
148. De Asis, K., Graves, E., Sutton, R. S. (2023). “Value-aware importance weighting for off-policy reinforcement learning.” In *Proceedings of the Second Conference on Lifelong Learning Agents*, 2023. Also arXiv:2306.15625.
149. Wan, Y., Sutton, R. S. (2022). On convergence of average-reward off-policy control algorithms in weakly-communicating MDPs. *14th International OPT Workshop on Optimization for Machine Learning*.
150. Sutton, R. S., “The quest for a common model of the intelligent decision maker.” In *Multi-disciplinary Conference on Reinforcement Learning and Decision Making*, 2022.
151. Javed, K., Shah, H., Sutton, R. S., “Scalable online state construction using recurrent networks.” In *Multi-disciplinary Conference on Reinforcement Learning and Decision Making*, 2022.
152. Wan, Y., Sutton, R. S., “Discovering options by minimizing the number of composed options to solve multiple tasks.” In *Multi-disciplinary Conference on Reinforcement Learning and Decision Making*, 2022.
153. De Asis, K. Chan, A., Sutton, R. S., “Inverse policy evaluation for value-based decision making.” In *Multi-disciplinary Conference on Reinforcement Learning and Decision Making*, 2022.

154. Dohare, S., Mahmood, R., Sutton, R. S., “Continual backprop: Stochastic gradient descent with persistent randomness.” In *Multi-disciplinary Conference on Reinforcement Learning and Decision Making*, 2022.
155. Sutton, R. S., Machado, M. C., Holland, G. Z., Szepesvari, D., Timbers, F., Tanner, B., White, A. “Reward-respecting subtasks for model-based reinforcement learning.” In *Multi-disciplinary Conference on Reinforcement Learning and Decision Making*, 2022.
156. Naik, A., Sutton, R. S., “Multi-step average-reward prediction via differential TD(λ).” In *Multi-disciplinary Conference on Reinforcement Learning and Decision Making*, 2022.
157. Wan, Y., Zaheer, M., White, M., Sutton, R. S., “Model-based reinforcement learning with non-linear expectation models and stochastic environments.” *Workshop of the Federated Artificial Intelligence Meeting*, Stockholm, Sweden, 2018.
158. Yu, H., Mahmood, A. R., Sutton, R. S., “On generalized Bellman equations and temporal-difference learning.” *Proceedings of the Canadian Artificial Intelligence Conference*, 2017.
159. Barto, A. G., Thomas, P. S., Sutton, R. S., “Some recent applications of reinforcement learning. In *Proceedings of the Eighteenth Yale Workshop on Adaptive and Learning Systems*, 2017.
160. Zhang, S., Sutton, R. S., “A deeper look at experience replay. *Symposium on Deep Reinforcement Learning at the 31st Conference on Neural Information Processing Systems*. ArXiv:1712.01275. 2017.
161. Veeriah, V., Pilarski, P. M., Sutton, R. S. “Face valuing: Training user interfaces with facial expressions and reinforcement learning.” *Int. J. Conference on Artificial Intelligence, Workshop on Interactive Machine Learning*, New York. Also arXiv 1606.02807. 2016.
162. van Seijen, H., Mahmood, A. R., Pilarski, P. M., Sutton, R. S. “An empirical evaluation of true online TD(λ).” In *Proceedings of the 2015 European Workshop on Reinforcement Learning*. Lille, France. 2015.
163. Mahmood, A. R., Yu, H., White, M., Sutton, R. S. “Emphatic temporal-difference learning.” In *Proceedings of the 2015 European Workshop on Reinforcement Learning*. Lille, France. 2015.
164. Edwards, A. L., Dawson, M. R., Hebert, J. S., Sutton, R. S., Chan, K. M., Pilarski, P.M. “Adaptive switching in practice: Improving myoelectric prosthesis performance through reinforcement learning.” In: *Proceedings of MEC14: Myoelectric Controls Symposium*, Fredericton, New Brunswick, 2014.
165. Edwards, A. L., Kearney, A., Dawson, M. R., Sutton, R. S., Pilarski, P. M., “Temporal-difference learning to assist human decision making during the control of an artificial limb.” In: *Proceedings of the First Interdisciplinary Conference on Reinforcement Learning and Decision Making*, 2013. <http://arxiv.org/abs/1309.4714>.

166. Silver, D., Sutton, R. S., Müller, M. “Temporal-difference search in computer go.” In: *Proceedings of the International Conference on Automated Planning and Scheduling*, 2013. Two-page extended abstract.
167. Mahmood, A. R., Sutton, R. S., “Representation search through generate and test.” In: *Proceedings of the AAAI Workshop on Learning Rich Representations from Low-Level Sensors*, Bellevue, Washington, USA, 2013.
168. Pilarski, P. M., Sutton, R. S., “Between instruction and reward: Human-prompted switching.” In: *Proceedings of the AAAI Fall Symposium on Robots Learning Interactively from Human Teachers*, 2012. AAAI Technical Report FS-12-07, pp. 46–52.
169. Sutton, R. S., “Beyond reward: The problem of knowledge and data.” Extended abstract in the *Proceedings of the 21st International Conference on Inductive Logic Programming*, Muggleton, S., Watanabe, H., Tamaddoni-Nezhad, A., and Chen, J. Eds. Lecture Notes in Computer Science, Springer, 2012
170. Modayil, J., White, A., Pilarski, P. M., Sutton, R. S., “Acquiring diverse predictive knowledge in real time by temporal-difference learning.” In *International Workshop on Evolutionary and Reinforcement Learning for Autonomous Robot Systems*, Montpellier, France, 2012.
171. Degris, T, Pilarski, P. M., Sutton, R. S., “Apprentissage par renforcement sans modèle et avec action continue.” *Journées Francophones sur la Planification, la Décision, et l'Apprentissage pour la conduite de systèmes*, Nancy, France, 2012.
172. Maei, H. R., Sutton, R. S., “GQ(λ): A general gradient algorithm for temporal-difference prediction learning with eligibility traces.” In *Proceedings of the Third Conference on Artificial General Intelligence*, Lugano, Switzerland, 2010.
173. Modayil, J., Pilarski, P. M., White, A., Degris, T., Sutton, R. S., “Off-policy knowledge maintenance for robots.” *Robotics Science and Systems Workshop (Towards Closing the Loop: Active Learning for Robotics)*. Extended abstract and poster, 2010.
174. Sutton, R. S., “The grand challenge of predictive empirical abstract knowledge.” *Working Notes of the IJCAI-09 Workshop on Grand Challenges for Reasoning from Experiences*, 2009.
175. Sutton, R. S., “Open theoretical questions in reinforcement learning,” *Proceedings of the Fourth European Conference on Computational Learning Theory*, pp. 11–17. Springer-Verlag, 1999.
176. McGovern, A., Precup, D., Ravindran, B., Singh, S., Sutton, R. S., “Hierarchical Optimal Control of MDPs,” *Proceedings of the Tenth Yale Workshop on Adaptive and Learning Systems*, pp. 186–191. Yale University, New Haven, CT, 1998.
177. Precup, D., Sutton, R. S., “Multi-time models for reinforcement learning,” *Proceedings of the ICML'97 Workshop on Modeling in Reinforcement Learning*, 1997.

178. McGovern, A., Sutton, R. S., Fagg, A. H., "Roles of macro-actions in accelerating reinforcement learning," *Proceedings of the 1997 Grace Hopper Celebration of Women in Computing*, pp. 13–17, 1997.
179. Precup, D., Sutton, R. S., Singh, S. P., "Planning with closed-loop macro actions," *Working notes of the 1997 AAAI Fall Symposium on Model-directed Autonomous Systems*, pp. 70–76, 1997.
180. Kuvayev, L., Sutton, R. S., "Model-based reinforcement learning with an approximate, learned model," *Proceedings of the Ninth Yale Workshop on Adaptive and Learning Systems*, pp. 101–105. Yale University, New Haven, CT, 1996.
181. Mehra, R. K., Ravichandran, B., Sutton, R. S., "Adaptive intelligent scheduling for ATM networks," *Proceedings of the Ninth Yale Workshop on Adaptive and Learning Systems*, pp. 106–111. Yale University, New Haven, CT, 1996.
182. Sutton, R. S., "On the virtues of linear learning and trajectory distributions." Abstract in *Proceedings of the Workshop on Value Function Approximation, Conference on Neural Information Processing Systems*, 1995.
183. Sutton, R. S., Singh, S. "On step size and bias in temporal-difference learning," *Proceedings of the Eighth Yale Workshop on Adaptive and Learning Systems*, pp. 91–96. Yale University, New Haven, CT, 1994.
184. Sutton, R. S., "Gain adaptation beats least squares?," *Proceedings of the Seventh Yale Workshop on Adaptive and Learning Systems*, pp. 161–166. Yale University, New Haven, CT, 1992.
185. Sutton, R. S., "Reinforcement learning architectures," *Proceedings ISKIT'92 International Symposium on Neural Information Processing*, pp. 211–216. Fukuoka, Japan, 1992.
186. Sutton, R. S., "Dyna, an integrated architecture for learning, planning and reacting," *Working Notes of the 1991 AAAI Spring Symposium on Integrated Intelligent Architectures and SIGART Bulletin 2*, pp. 160–163, 1991.
187. Whitehead, S., Sutton, R. S., Ballard, D., "Advances in reinforcement learning and their implications for intelligent control," *Proceedings of the Fifth IEEE International Symposium on Intelligent Control*, pp. 1289–1297, 1990.
188. Anderson, C., Franklin, J., Sutton, R. S., "Learning a nonlinear model of a manufacturing process using multilayer connectionist networks," *Proceedings of the Fifth IEEE International Symposium on Intelligent Control*, pp. 404–409, 1990.
189. Sutton, R. S., "Artificial intelligence by dynamic programming," *Proceedings of the Sixth Yale Workshop on Adaptive and Learning Systems*, pp. 89–95. Yale University, New Haven, CT, 1990.

190. Franklin, J., Sutton, R. S., Anderson, C., Selfridge, O., Schwartz, D., “Connectionist learning control at GTE laboratories,” *Intelligent Control and Adaptive Systems*, G. Rodriguez, Ed., Proc. SPIE 1196, pp. 242–253, 1990.
191. Sutton, R. S., “Artificial intelligence as a control problem: Comments on the relationship between machine learning and intelligent control,” *Proceedings of the IEEE International Symposium on Intelligent Control 1988*, pp. 500–507, 1989.
192. Franklin, J., Sutton, R. S., Anderson, C. “Application of connectionist learning methods to manufacturing process monitoring,” *Proceedings of the IEEE International Symposium on Intelligent Control 1988*, pp. 709–712, 1989.
193. Sutton, R. S., “Convergence theory for a new kind of prediction learning,” *Proceedings of the 1988 Workshop on Computational Learning Theory*, pp. 421–422 1988.
194. Sutton, R. S., “Learning distributed, searchable, internal models,” *Proceedings of the Distributed Artificial Intelligence Workshop*, pp. 287–289, 1985.
195. Barto, A. G., Sutton, R. S., Anderson, C., “Spatial learning simulation systems,” *Proceedings of the 10th IMACS World Congress on Systems Simulation and Scientific Computation*, pp. 204–206, 1982.

Other Publications

196. Sharifnassab, A., Salehkaleybar, S., Sutton, R. “Metaoptimize: A framework for optimizing step sizes and other meta-parameters.” ArXiv:2402.02342, 2024
197. Javed, K., Sutton, R.S. “The big world hypothesis and its ramifications for artificial intelligence.” Presented at the Reinforcement Learning Conference. On Openreview.net. 2024.
198. Degris, T., Javed, K., Sharifnassab, A., Liu, Y., Sutton, R.S. “Step-size optimization for continual learning.” ArXiv 2401.17401, 2024.
199. Dohare, S., Hernandez-Garcia, J.F., Rahman, P., Sutton, R. S., Mahmood, A.R. “Loss of plasticity in deep continual learning.” ArXiv:2210.14361, 2023.
200. Sutton, R. S., Bowling, M., Pilarski, P. M. “The Alberta plan for AI research.” ArXiv:2208.11173, 2023.
201. Sutton, R. S. “A history of meta-gradient: Gradient Methods for meta-learning.” ArXiv:2202.09701, 2022
202. Sutton, R., “The bitter lesson.” Influential blog entry, 2019.
203. Mahmood, A. R., Yu, H., Sutton, R. S., “Multi-step off-policy learning without importance sampling ratios.” ArXiv:1702.03006, 2017.

204. Young, K. J., Sutton, R. S., Yang, S., “Integrating episodic memory into a reinforcement learning agent using reservoir sampling. ArXiv:1806.00540, 2018.
205. Pilarski, P. M., Sutton, R. S., Mathewson, K. W., Sherstan, C., Parker, A. S., Edwards, A. L., “Communicative capital for prosthetic agents.” ArXiv:1711.03676, 2017.
206. Kearney, A., Veeriah, V., Travník, J. B., Sutton, R. S., Pilarski, P. M., “TIDBD: Adapting temporal-difference step-sizes through stochastic meta-descent.” ArXiv:1804.03334, 2018.
207. Ghiassian, S., Yu, H., Rafiee, B., Sutton, R. S., “Two geometric input transformation methods for fast online reinforcement learning with neural nets.” ArXiv:1805.07476, 2018.
208. De Asis, K., Bennett, B., Sutton, R. S., “Predicting periodicity with temporal difference learning,” ArXiv:1809.07435, 2018.
209. van Hasselt, H., Sutton, R. S., “Learning to predict independent of span,” ArXiv:1508.04582, 2015.
210. “Interview with Richard S. Sutton,” by Verena Heidrich-Meisner. *Kuenstliche Intelligenz*, pp. 41–43. 2009.
211. Sutton, R. S., “Reinforcement learning: Past, present, and future.” Extended abstract in *Simulated Evolution and Learning*, McKay, B., Yao, X., Newton, C. S., Kim, J.-H., Furuhashi, T., Eds. Published as *Lecture Notes in Computer Science 1585*, pp. 195–197, Springer, 1999.
212. McGovern, A., Sutton, R. S., “Macro-actions in reinforcement learning: An empirical analysis.” Technical Report 98-70, Department of Computer Science, University of Massachusetts, Amherst, MA 01003, 1998.
213. Precup, D., Sutton, R. S., “Empirical comparison of gradient descent and exponentiated gradient descent in supervised and reinforcement learning.” Technical Report UM-CS-1996-070, Department of Computer Science, University of Massachusetts, Amherst, MA 01003, 1996.
214. Sutton, R. S., “Machines that learn and mimic the brain.” In *ACCESS, GTE's Journal of Science and Technology*, 1992. Reprinted in *Stethoscope Quarterly*, Spring 1993.
215. Sutton, R. S., “The challenge of reinforcement learning.” Introduction to a special issue on reinforcement learning, *Machine Learning* 8, No 3/4, pp. 225–227, 1992.
216. Sutton, R. S., “Implementation details of the TD(λ) procedure for the case of vector predictions and backpropagation.” GTE Laboratories Technical Report TR87-509.1, GTE Laboratories, 40 Sylvan Road, Waltham, MA 02254, 1989.
217. Sutton, R. S., “NADALINE: A normalized adaptive linear element that learns efficiently.” GTE Laboratories Technical Report TR88-509.4, GTE Laboratories, 40 Sylvan Road, Waltham, MA 02254, 1988.

218. Sutton, R. S., “Temporal credit assignment in reinforcement learning,” Ph.D. dissertation, University of Massachusetts, Amherst, MA. Published as COINS Technical Report 84-2, 1984.
219. Barto, A. G., Sutton, R. S., “Goal seeking components for adaptive intelligence: An initial assessment.” Air Force Wright Aeronautical Laboratories Technical Report, AFWAL-TR-81-1070, Wright-Patterson Air Force Base, Ohio, 1981.
220. Sutton, R. S., “Single channel theory: A neuronal theory of learning,” *Brain Theory Newsletter* 3, No. 3/4, pp. 72–75, 1978. (earliest publication)

Patents

221. Sutton, R. S., “Apparatus for Machine Learning.” United States Patent No. 5,946,675, August 31, 1999. Assigned to GTE Laboratories Incorporated, Waltham, Mass.
222. Sutton, R. S., “Method and Apparatus for Machine Learning.” United States Patent No. 6,249,781, June 19, 2001. Assigned to Verizon Laboratories Incorporated, Waltham, Mass.