

Michael Bowling

Curriculum Vitae

Department of Computing Science
University of Alberta
Edmonton, Alberta
Canada T6G 2E8

mbowling@ualberta.ca
<https://bowlingmh.github.io/>

Updated: August, 2026

RESEARCH INTERESTS	Artificial Intelligence. Specifically, learning in games, reinforcement learning, game theory, agent modelling, computer game environments, and mobile robots.	
ACADEMIC POSITIONS	Full Professor Department of Computing Science, University of Alberta	2013–Present
	Associate Professor Department of Computing Science, University of Alberta	2008–2013
	Adjunct Assistant Professor School of Computer Science, University of Waterloo	2006–2008
	Assistant Professor Department of Computing Science, University of Alberta	2003–2008
EDUCATION	Ph.D. in Computer Science, Carnegie Mellon University Advisor: Professor Manuela Veloso Award: Co-winner of the School of Computer Science Dissertation Award.	2003
	M.Sc., Computer Science Carnegie Mellon University	1999
	B.S., Math and Computer Science, Carnegie Mellon University University Honors	1996
AWARDS AND HONORS	Faculty of Science Service Award	2025
	Google DeepMind Chair	2024–2029
	AAAI Fellow	Elected 2021
	Canada CIFAR AI Chair	2021–2025
	Faculty of Science Research Fellowship	2017–2022
	Faculty of Science Research Award	2014
	Faculty of Science Innovation in Teaching Award Co-winner. This award recognizes creative contributions to the development of tools for the innovative presentation of coursework, new curriculum design, and the creation of an improved teaching and learning environment.	2014
	Honourable Mention for Alan Blizzard Award Past head instructor for the CMPUT250 teaching team that won this award. The Alan Blizzard award is a Canadian-wide award for collaborative teaching. Only 4 honourable	2011

mentions had been awarded in its first 12 years, in addition to the 14 winners and co-winners.

Department Research Award 2010

Teaching Unit Award 2009

Head instructor for the CMPUT250 teaching team that won the University of Alberta's Teaching Unit Award for "excellence that occurs as a result of the collaboration of instructors".

Second Man-Machine Poker Championship 2008

Led the Computer Poker Research Group in the development of Polaris, which had 3 wins, 2 losses, and 1 draw to become the first computer program to defeat professional poker players in a meaningful competition. Polaris also competed in the First Man-Machine Poker Championship in 2007 resulting in 1 win, 2 losses, and a draw.

AAAI Computer Poker Competition Champion 2006-2015

Led the University of Alberta team that has won a total of 23 of the 39 events over its first 10 years.

AI Video Awards Finalist 2008

Finalist for Best Educational Video.

PE Publishing Award 2007

Awarded the 2005 PE Publishing Award for the *Journal of Systems and Control Engineering* [17]. Co-author.

Best Student Paper Award 2007

At the 2007 IEEE International Symposium on Approximate Dynamic Programming and Reinforcement Learning [151]. Co-author (with Ph.D. student Tao Wang).

ASTech Award Winner 2006

A principal investigator in the Alberta Ingenuity Centre for Machine Learning, winning the Alberta Science and Technology Foundation award in the category *Outstanding Leadership in Alberta Technology*.

ASTech Finalist 2005

Alberta Science and Technology Foundation finalist in the category *Leaders of Tomorrow*.

CMU School of Computer Science Dissertation Award 2003

Co-winner. Nominated for the ACM Dissertation Award.

RoboCup Robot Soccer Champion Teams

RoboCup Small-Size League (CMU) 1998

American Open Small-Size League (CMU) 2003

IJCAI Excellent Paper 2001

Conference paper [135] invited as a fast-track submission to *Artificial Intelligence*.

Allen Newell Medal for Research Excellence 1998

The Allen Newell Medal for Research Excellence is awarded by the School of Com-

puter Science at Carnegie Mellon. It recognizes an outstanding body of work that epitomizes Allen Newell’s research style.

INDUSTRY RESEARCH **DeepMind**, Edmonton, AB July 2017–January 2024
Senior Staff Research Scientist and Site Lead.

Cogitai, Inc., Los Angeles, CA May 2016–April 2017
Visiting Researcher. Developing a platform for continual learning systems “that will let everyday things that sense and act get smarter with experience.”

Yahoo! Research, Santa Clara, CA February–June, 2010
Visiting Researcher. Worked on a complete redesign of the anti-spam system in order to achieve adversarial guarantees. Co-developed a new concept of regret based on locality that can be applied in domains where even off-line decision-making is computationally hard [81].

PUBLICATIONS I have focused my publications on highly refereed tier-one conferences, which in computer science are rigorously refereed and as valued as journal publications. The top conference venues in artificial intelligence often have citation impact factors higher than the top journals in the field with lower time to publication. **Boldfaced co-authors** note undergraduate students, graduate students, or post-doctoral fellows under my supervision at the time of the research.

Journal Articles

[1] Farzane Aminmansour, **Taher Jafferjee**, Ehsan Imani, Erin J. Talvitie, Michael Bowling, and Martha White. Multistep predecessor models and mitigating errors due to hallucinated value in dyna-style planning. *Journal of Artificial Intelligence Research*, 80:441–473, 2024.

[2] Zhe Wang, Petar Veličković, Daniel Hennes, Nenad Tomašev, Laurel Prince, Michael Kaisers, Yoram Bachrach, Romuald Elie, Li Kevin Wenliang, Federico Piccinini, William Spearman, Ian Graham, Jerome Connor, Yi Yang, Adrià Recasens, Mina Khan, Nathalie Beauguerlange, Pablo Sprechmann, Pol Moreno, Nicolas Heess, Michael Bowling, Demis Hassabis, and Karl Tuyls. TacticAI: an AI assistant for football tactics. *Nature Communications*, 15(1906), 2024.

[3] Martin Schmid, Matej Moravčík, Neil Burch, Rudolf Kadlec, Josh Davidson, Kevin Waugh, Nolan Bard, Finbarr Timbers, Marc Lanctot, G. Zacharias Holland, Elnaz Davoodi, Alden Christianson, and Michael Bowling. Student of games: A unified learning algorithm for both perfect and imperfect information games. *Science Advances*, 9(46):eadg3256, 2023.

[4] Marlos C. Machado, Andre Barreto, Doina Precup, and Michael Bowling. Temporal abstraction in reinforcement learning with the successor representation. *Journal of Machine Learning Research*, 24(80):1–69, 2023.

[5] Vojtěch Kovařík, **Martin Schmid**, Neil Burch, Michael Bowling, and Viliam Lisý. Rethinking formal models of partially observable multiagent decision making. *Artificial Intelligence*, 303:103645, 2022.

- [6] **Paniz Behboudian, Yash Satsangi**, Matthew E. Taylor, Anna Harutyunyan, and Michael Bowling. Policy invariant explicit shaping: an efficient alternative to reward shaping. *Neural Computing and Applications*, 34:1673–1686, 2022.
- [7] Cleyton R. Silva, Michael Bowling, and Levi H.S. Lelis. Teaching people by justifying tree search decisions: An empirical study in curling. *Journal of Artificial Intelligence Research*, 72:1083–1102, 2021.
- [8] Nolan Bard, Jakob N. Foerster, Sarath Chandar, Neil Burch, Marc Lanctot, H. Francis Song, Emilio Parisotto, Vincent Dumoulin, Subhodeep Moitra, Edward Hughes, Iain Dunning, Shihab Mourad, Hugo Larochelle, Marc G. Bellemare, and Michael Bowling. The Hanabi challenge: A new frontier for AI research. *Artificial Intelligence*, 280:103216, 2020.
- [9] **Marlos Machado**, Marc Bellemare, Erik Talvitie, Joel Veness, Matthew Hausknecht, and Michael Bowling. Revisiting the arcade learning environment: Evaluation protocols and open problems for general agents. *Journal of Artificial Intelligence Research*, 61:523–562, 2018.
- [10] Michael Bowling, **Neil Burch, Michael Johanson**, and Oskari Tammelin. Heads-up limit hold’em poker is solved. *Communications of the ACM*, 60(11):81–88, 2017.
- [11] Matej Moravčík, **Martin Schmid, Neil Burch, Viliam Lisý, Dustin Morrill, Nolan Bard, Trevor Davis, Kevin Waugh**, Michael Johanson, and Michael Bowling. Deepstack: Expert-level artificial intelligence in heads-up no-limit poker. *Science*, 356(6337):508–513, 2017.
- [12] Michael Bowling, **Neil Burch, Michael Johanson**, and Oskari Tammelin. Heads-up limit hold’em poker is solved. *Science*, 347(6218):145–149, 2015.
- [13] T.L. MacKay, **Nolan Bard**, Michael Bowling, and D.C. Hodgins. Do poker players know how good they are? Accuracy of poker skill estimation in online and offline players. *Computers in Human Behavior*, 31:419–424, 2014.
- [14] **Arash Afkanpour**, Csaba Szepesvári, and Michael Bowling. Alignment based kernel learning with a continuous set of base kernels. *Machine Learning*, 91(3):305–324, 2013.
- [15] **Marc G. Bellemare**, Yavar Naddaf, **Joel Veness**, and Michael Bowling. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47:253–279, 2013.
- [16] Daniel J. Lizotte, Michael Bowling, and Susan A. Murphy. Linear fitted-Q iteration with multiple reward functions. *Journal of Machine Learning Research*, 13:3253–3295, 2012.
- [17] Brett Browning, James Bruce, Michael Bowling, and Manuela Veloso. STP: Skills, tactics and plays for multi-robot control in adversarial environments. *Journal of Systems and Control Engineering*, 219(1):33–52, 2005.
- [18] Michael Bowling and Manuela Veloso. Existence of multiagent equilibria with limited agents. *Journal of Artificial Intelligence Research*, 22:353–384, 2004.
- [19] Michael Bowling and Manuela Veloso. Multiagent learning using a variable learn-

ing rate. *Artificial Intelligence*, 136:215–250, 2002.

[20] Manuela Veloso, Michael Bowling, Sorin Achim, Kwun Han, and Peter Stone. The CMUnited-98 champion small robot team. *Advanced Robotics*, 2000. An earlier version appeared in *RoboCup-98: Robot Soccer World Cup II, Asada and Kitano (Eds.)*, Springer, 1999, pages 77–92. A shorter version appeared in the *AI Magazine*, 21:29–36.

Refereed Conferences

[21] **Montaser Mohammedalamen** and Michael Bowling. Generalization in monitored Markov decision processes (Mon-MDPs). In *Proceedings of the Third Reinforcement Learning Conference (RLC)*, 2026. Accept: 34%.

[22] **Adrian Orenstein, Jiajing Chen, Gwyneth Anne Delos Santos, Bayley Sapara**, and Michael Bowling. Toward agents that reason about their computation. In *Proceedings of the Third Reinforcement Learning Conference (RLC)*, 2026. Accept: 34%.

[23] David Sychrovský, Christopher Solinas, Revan MacQueen, Kevin A. Wang, James R. Wright, Nathan R. Sturtevant, and Michael Bowling. Approximating Nash equilibria in general-sum games via meta-learning. In *Proceedings of the Twenty-Fifth International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2026. Accept: 25%.

[24] Jiabin Fan, Guoqing Luo, Michael Bowling, and Lili Mou. KETCHUP: K-step return estimation for sequential knowledge distillation. In *Findings of the Association for Computational Linguistics: EACL 2026*, 2026. Accept: 36%.

[25] David Abel, Michael Bowling, André Barreto, Will Dabney, Shi Dong, Steven Stenberg Hansen, Anna Harutyunyan, Khimya Khetarpal, Clare Lyle, Razvan Pascanu, Georgios Piliouras, Doina Precup, Jonathan Richens, Mark Rowland, Tom Schaul, and Satinder Singh. Plasticity as the mirror of empowerment. In *Advances in Neural Information Processing Systems 38 (NeurIPS)*, 2025. Spotlight (688 of 5290 accepted), Accept: 25%.

[26] **Anna Koop**, Michael Bowling, and Michael Bradley Johanson. Using partition-tree weighting and MAML for continual and online learning. In *Proceedings of the Fourth Conference on Lifelong Learning Agents (CoLLAs)*, 2025.

[27] Esraa Elelimy, David Szepesvari, Martha White, and Michael Bowling. Rethinking the foundations of continual reinforcement learning. In *Proceedings of the Second Reinforcement Learning Conference (RLC)*, 2025. Accept: 39%.

[28] **Alireza Kazemipour**, Matthew E. Taylor, and Michael Bowling. Model-based exploration in monitored Markov decision processes. In *Proceedings of the Forty-Second International Conference on Machine Learning (ICML)*, 2025. To Appear, Accept: 27%.

[29] Esraa Elelimy, Adam White, Michael Bowling, and Martha White. Real-time recurrent learning using trace units in reinforcement learning. In *Advances in Neural Information Processing Systems 37 (NeurIPS)*, 2024. Accept: 25.8%.

- [30] **Simone Parisi, Alireza Kazemipour**, and Michael Bowling. Beyond optimism: Exploration with partially observable rewards. In *Advances in Neural Information Processing Systems 37 (NeurIPS)*, 2024. Accept: 25.8%.
- [31] **Jacob Adkins**, Michael Bowling, and Adam White. A method for evaluating hyperparameter sensitivity in reinforcement learning. In *Advances in Neural Information Processing Systems 37 (NeurIPS)*, 2024. Accept: 25.8%.
- [32] David Sychrovsky, Michal Sustr, Elnaz Davoodi, Michael Bowling, Marc Lanctot, and Martin Schmid. Learning not to regret. In *Proceedings of the Thirty-Eighth Conference on Artificial Intelligence (AAAI)*, pages 15202–15210, 2024. Accept: 24%.
- [33] **Simone Parisi, Montaser Mohammedalamen, Alireza Kazemipour**, Matthew E. Taylor, and Michael Bowling. Monitored Markov decision processes. In *Proceedings of the Twenty-Third International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1549–1557, 2024. Accept: 25%.
- [34] **Diego Gomez**, Michael Bowling, and Marlos C. Machado. Proper Laplacian representation learning. In *Proceedings of the Eleventh International Conference on Learning Representations (ICLR)*, 2024. To Appear, Accept: 31%.
- [35] Michael Bowling, **John D. Martin**, David Abel, and Will Dabney. Settling the reward hypothesis. In *Proceedings of the Fortieth International Conference on Machine Learning (ICML)*, pages 3003–3020, 2023. Accept: 28%.
- [36] **Alex Trudeau** and Michael Bowling. Targeted search control in alphazero for effective policy improvement. In *Proceedings of the Twenty-Second International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 842–840, 2023. Accept: 23%.
- [37] Finbarr Timbers, Nolan Bard, Edward Lockhart, Marc Lanctot, Martin Schmid, Neil Burch, Julian Schrittwieser, Thomas Hubert, and Michael Bowling. Approximate exploitability: Learning a best response. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence (IJCAI)*, pages 3487–3493, 2022. Accept: 15%.
- [38] Levi Lelis, João Nova, Eugene Chen, Nathan Sturtevant, Carrie Demmans Epp, and Michael Bowling. Learning curricula for humans: An empirical study with puzzles from The Witness. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence (IJCAI)*, pages 3877–3883, 2022. Accept: 15%.
- [39] **Dustin Morrill**, Ryan D’Orazio, Marc Lanctot, Reza Sarfati, James Wright, Amy Greenwald, and Michael Bowling. Efficient deviation types and learning for hindsight rationality in extensive-form games. In *Proceedings of the Thirty-Eighth International Conference on Machine Learning (ICML)*, pages 7818–7828, 2021. Accept: 21.5%.
- [40] Michal Šustr, **Martin Schmid**, Matej Moravčík, Neil Burch, Marc Lanctot, and Michael Bowling. Sound search algorithms in imperfect information games. In *Proceedings of the Twentieth International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2021. Extended Abstract, Accept: 40.2%.
- [41] Samuel Sokota, Edward Lockhart, Finbarr Timbers, Elnaz Davoodi, Ryan D’Orazio, Neil Burch, **Martin Schmid**, Michael Bowling, and Marc Lanctot. Solv-

ing common-payoff games with approximate policy iteration. In *Proceedings of the Thirty-Fifth Conference on Artificial Intelligence (AAAI)*, 2021. Accept: 21.3%.

[42] **Dustin Morrill**, Ryan D’Orazio, Reza Sarfati, Marc Lanctot, James R. Wright, Amy Greenwald, and Michael Bowling. Hindsight and sequential rationality of correlated play. In *Proceedings of the Thirty-Fifth Conference on Artificial Intelligence (AAAI)*, 2021. Accept: 21.3%.

[43] **Zaheen Farraz Ahmad**, Levi H. S. Lelis, and Michael Bowling. Marginal utility for planning in continuous or large discrete action spaces. In *Advances in Neural Information Processing Systems 33 (NeurIPS)*, 2020. Accept: 20.1%.

[44] **Trevor Davis**, **Martin Schmid**, and Michael Bowling. Low-variance and zero-variance baselines for extensive-form games. In *Proceedings of the Thirty-Seventh International Conference on Machine Learning (ICML)*, 2020. Accept: 21.8%.

[45] **Marlos Machado**, Marc Bellemare, and Michael Bowling. Count-based exploration with the successor representation. In *Proceedings of the Thirty-Fourth Conference on Artificial Intelligence (AAAI)*, 2020. Accept: 21%.

[46] **Fushan Li** and Michael Bowling. Easy-to-teach language emergence from multi-agent communication games. In *Advances in Neural Information Processing Systems 32 (NeurIPS)*, 2019. Accept: 21%.

[47] Jakob N. Foerster, H. Francis Song, Edward Hughes, Neil Burch, Iain Dunning, Shimon Whiteson, Matthew Botvinick, and Michael Bowling. Bayesian action decoder for deep multi-agent reinforcement learning. In *Proceedings of the Thirty-Sixth International Conference on Machine Learning (ICML)*, pages 1942–1951, 2019. Accept: 22.6%.

[48] **Trevor Davis**, Kevin Waugh, and Michael Bowling. Solving large extensive-form games with strategy constraints. In *Proceedings of the Thirty-Third Conference on Artificial Intelligence (AAAI)*, pages 1861–1868, 2019. Accept: 16.2%.

[49] Srinivasan Sriram, Marc Lanctot, Vinicius Zambaldi, Julien Perolat, Karl Tuyls, Remi Munos, and Michael Bowling. Actor-critic policy optimization in partially observable multiagent environments. In *Advances in Neural Information Processing Systems 31 (NeurIPS)*, pages 3426–3439, 2018. Accept: 20.8%.

[50] **Martin Schmid**, Neil Burch, Marc Lanctot, **Matej Moravčík**, Rudolf Kadlec, and Michael Bowling. Variance reduction in monte carlo regret minimization for extensive games using baselines. In *Proceedings of the Thirty-Second Conference on Artificial Intelligence (AAAI)*, pages 2157–2164, 2018. Accept: 16.2%.

[51] **Neil Burch**, **Martin Schmid**, **Matej Moravčík**, **Dustin Morrill**, and Michael Bowling. AIVAT: A new variance reduction technique for agent evaluation in imperfect information games. In *Proceedings of the Thirty-Second Conference on Artificial Intelligence (AAAI)*, pages 949–956, 2018. Accept: 24.6%.

[52] **Marlos Machado**, Marc Gendron Bellemare, and Michael Bowling. A Laplacian framework for option discovery in reinforcement learning. In *Proceedings of the Thirty-Fourth International Conference on Machine Learning (ICML)*, pages 2295–2304, 2017. Accept: 25%.

- [53] Kieran Milan, Joel Veness, James Kirkpatrick, Michael Bowling, **Anna Koop**, and Demis Hassabis. The forget-me-not process. In *Advances in Neural Information Processing Systems 29 (NIPS)*, 2016. Accept: 23%.
- [54] Zaheen Farraz Ahmad, Robert C. Holte, and Michael Bowling. Action selection for hammer shots in curling. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 561–567, 2016. Accept: 24%.
- [55] Timothy Yee, Viliam Lisý, and Michael Bowling. Monte Carlo tree search in continuous action spaces with execution uncertainty. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 690–696, 2016. Accept: 24%.
- [56] Yitao Liang, Marlos Machado, Erik Talvitie, and Michael Bowling. State of the art control of atari games using shallow reinforcement learning. In *Proceedings of the Fifteenth International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 485–493, 2016. Accept: 25%.
- [57] **Viliam Lisý**, **Trevor Davis**, and Michael Bowling. Counterfactual regret minimization in sequential security games. In *Proceedings of the Thirtieth Conference on Artificial Intelligence (AAAI)*, pages 544–550, 2016. Accept: 26%.
- [58] Oskari Tammelin, Neil Burch, Michael Johanson, and Michael Bowling. Solving heads-up limit texas hold'em. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 645–652, 2015. Accept: 29%.
- [59] Viliam Lisý, Marc Lanctot, and Michael Bowling. Online Monte Carlo counterfactual regret minimization for search in imperfect information games. In *Proceedings of the Fourteenth International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 27–36, 2015. Accept: 25%.
- [60] **Nolan Bard**, Deon Nicholas, Csaba Szepesvari, and Michael Bowling. Decision-theoretic clustering of strategies. In *Proceedings of the Fourteenth International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 17–25, 2015. Accept: 25%.
- [61] **James Neufeld**, Michael Bowling, and Dale Schuurmanns. Variance reduction via antithetic Markov chains. In *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 708–716, 2015. Accept: 29%.
- [62] Kevin Waugh, **Dustin Morrill**, J. Andrew Bagnell, and Michael Bowling. Solving games with functional regret estimation. In *Proceedings of the Twenty-Ninth Conference on Artificial Intelligence (AAAI)*, pages 2138–2145, 2015. Accept: 27%.
- [63] **Martha White**, Junfeng Wen, Michael Bowling, and Dale Schuurmanns. Optimal estimation of multivariate ARMA models. In *Proceedings of the Twenty-Ninth Conference on Artificial Intelligence (AAAI)*, pages 3080–3086, 2015. Accept: 27%.
- [64] Pascal Poupart, Aarti Malhotra, Pei Pei, Kee-Eung Kim, Bongseok Goh, and Michael Bowling. Approximate linear programming for constrained partially observable Markov decision processes. In *Proceedings of the Twenty-Ninth Conference on Artificial Intelligence (AAAI)*, pages 3080–3086, 2015. Accept: 27%.

- [65] **Sriram Srinivasan**, Erik Talvitie, and Michael Bowling. Improving exploration in UCT using local manifolds. In *Proceedings of the Twenty-Ninth Conference on Artificial Intelligence (AAAI)*, pages 3386–3392, 2015. Accept: 27%.
- [66] **Ujjwal Das Gupta**, Erik Talvitie, and Michael Bowling. Policy tree: Adaptive representation for policy gradient. In *Proceedings of the Twenty-Ninth Conference on Artificial Intelligence (AAAI)*, pages 2547–2553, 2015. Accept: 27%.
- [67] **Trevor Davis**, **Neil Burch**, and Michael Bowling. Using response functions to measure strategy strength. In *Proceedings of the Twenty-Eighth Conference on Artificial Intelligence (AAAI)*, pages 630–636, 2014. Accept: 28%.
- [68] **Neil Burch**, **Michael Johanson**, and Michael Bowling. Solving imperfect information games using decomposition. In *Proceedings of the Twenty-Eighth Conference on Artificial Intelligence (AAAI)*, pages 602–608, 2014. Accept: 28%.
- [69] **Nolan Bard**, **Michael Johanson**, and Michael Bowling. Asymmetric abstractions for adversarial settings. In *Proceedings of the Thirteenth International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 501–508, 2014. Accept: 24%.
- [70] **Parisa Mazrooei**, **Chris Archibald**, and Michael Bowling. Automating collusion detection in sequential games. In *Proceedings of the Twenty-Seventh Conference on Artificial Intelligence (AAAI)*, pages 675–682, 2013. Accept: 29%.
- [71] **Chris Rayner**, Nathan Sturtevant, and Michael Bowling. Subset selection of search heuristics. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence (IJCAI)*, pages 637–643, 2013. Accept: 28%.
- [72] **Marc Bellemare**, **Joel Veness**, and Michael Bowling. Bayesian learning of recursively factored environments. In *Proceedings of the Thirtieth International Conference on Machine Learning (ICML)*, pages 1211–1219, 2013. Accept: 24%.
- [73] **Arash Afkanpour**, András György, Csaba Szepesvári, and Michael Bowling. A randomized mirror descent algorithm for large scale multiple kernel learning. In *Proceedings of the Thirtieth International Conference on Machine Learning (ICML)*, pages 374–382, 2013. Accept: 24%.
- [74] **Joel Veness**, **Martha White**, Michael Bowling, and András György. Partition tree weighting. In *Proceedings of the Data Compression Conference (DCC)*, pages 321–330, 2013.
- [75] **Michael Johanson**, **Neil Burch**, Richard Valenzano, and Michael Bowling. Evaluating state-space abstractions in extensive-form games. In *Proceedings of the Twelfth International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 271–278, 2013. Accept: 23%.
- [76] **Nolan Bard**, **Michael Johanson**, **Neil Burch**, and Michael Bowling. Online implicit agent modelling. In *Proceedings of the Twelfth International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 255–262, 2013. Accept: 23%.
- [77] **Joshua Davidson**, **Christopher Archibald**, and Michael Bowling. Baseline:

Practical control variates for agent evaluation in zero-sum domains. In *Proceedings of the Twelfth International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1005–1012, 2013. Accept: 23%.

[78] **Katherine Chen** and Michael Bowling. Tractable objectives for robust policy optimization. In *Advances in Neural Information Processing Systems 25 (NIPS)*, pages 2078–2086, 2012. Accept: 25%.

[79] **Marc G. Bellemare, Joel Veness**, and Michael Bowling. Sketch-based linear value function approximation. In *Advances in Neural Information Processing Systems 25 (NIPS)*, pages 2222–2230, 2012. Accept: 25%.

[80] **Marc Lanctot**, Richard Gibson, **Neil Burch**, and Michael Bowling. No-regret learning in extensive-form games with imperfect recall. In *Proceedings of the Twenty-Ninth International Conference on Machine Learning (ICML)*, pages 65–72, 2012. Accept: 27%.

[81] Michael Bowling and Martin Zinkevich. On local regret. In *Proceedings of the Twenty-Ninth International Conference on Machine Learning (ICML)*, pages 1631–1638, 2012. Accept: 27%.

[82] **Michael Johanson, Nolan Bard, Neil Burch**, and Michael Bowling. Finding optimal abstract strategies in extensive form games. In *Proceedings of the Twenty-Sixth Conference on Artificial Intelligence (AAAI)*, pages 1371–1379, 2012. Accept: 26%.

[83] **Marc G. Bellemare, Joel Veness**, and Michael Bowling. Investigating contingency awareness using Atari 2600 games. In *Proceedings of the Twenty-Sixth Conference on Artificial Intelligence (AAAI)*, pages 864–871, 2012. Accept: 26%.

[84] Richard Gibson, **Marc Lanctot, Neil Burch**, Duane Szafron, and Michael Bowling. Generalized sampling and variance in counterfactual regret minimization. In *Proceedings of the Twenty-Sixth Conference on Artificial Intelligence (AAAI)*, pages 1355–1361, 2012. A longer version is available as a University of Alberta Technical Report, TR12-02., Accept: 26%.

[85] **Joel Veness**, Kee Siong Ng, Marcus Hutter, and Michael Bowling. Context tree switching. In *Proceedings of the Data Compression Conference (DCC)*, pages 327–336, 2012.

[86] **Michael Johanson, Nolan Bard, Marc Lanctot**, Richard Gibson, and Michael Bowling. Efficient Nash equilibrium approximation through Monte Carlo counterfactual regret minimization. In *Proceedings of the Eleventh International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 837–844, 2012. Accept: 20%.

[87] **Joel Veness, Marc Lanctot**, and Michael Bowling. Variance reduction in Monte Carlo tree search. In *Advances in Neural Information Processing Systems 24 (NIPS)*, pages 1836–1844, 2011. Accept: 22%.

[88] **Chris Rayner**, Michael Bowling, and Nathan Sturtevant. Euclidean heuristic optimization. In *Proceedings of the Twenty-Fifth Conference on Artificial Intelligence (AAAI)*, pages 81–86, 2011. Accept: 25%.

- [89] **Michael Johanson**, Michael Bowling, **Kevin Waugh**, and Martin Zinkevich. Accelerating best response calculation in large extensive games. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI)*, pages 258–265, 2011. Accept: 30%.
- [90] Daniel J. Lizotte, Michael Bowling, and Susan A. Murphy. Efficient reinforcement learning with multiple reward functions for randomized controlled trial analysis. In *Proceedings of the Twenty-Seventh International Conference on Machine Learning (ICML)*, pages 695–702, 2010. Accept: 26%.
- [91] **Marc Lanctot**, **Kevin Waugh**, Martin Zinkevich, and Michael Bowling. Monte Carlo sampling for regret minimization in extensive games. In *Advances in Neural Information Processing Systems 22 (NIPS)*, pages 1078–1086, 2009. A longer version is available as a University of Alberta Technical Report, TR09-15. An earlier version appeared in the *COLT Workshop on On-Line Learning with Limited Feedback* (2009), Accept: 24%.
- [92] **Kevin Waugh**, **Nolan Bard**, and Michael Bowling. Strategy grafting in extensive games. In *Advances in Neural Information Processing Systems 22 (NIPS)*, pages 2026–2034, 2009. Accept: 24%.
- [93] **Martha White** and Michael Bowling. Learning a value analysis tool for agent evaluation. In *Proceedings of the Twenty-First International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1976–1981, 2009. Accept: 26%.
- [94] **David Schnizlein**, Michael Bowling, and Duane Szafron. Probabilistic state translation in extensive games with large action sets. In *Proceedings of the Twenty-First International Joint Conference on Artificial Intelligence (IJCAI)*, pages 276–284, 2009. Accept: 26%.
- [95] **Kevin Waugh**, **Dave Schnizlein**, Michael Bowling, and Duane Szafron. Abstraction pathologies in extensive games. In *Proceedings of the Eighth International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 781–788, 2009. Accept: 22%.
- [96] **Michael Johanson** and Michael Bowling. Data biased robust counter strategies. In *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 264–271, 2009. Accept: 30%.
- [97] Maria Cutumisu, Duane Szafron, Michael Bowling, and Richard S. Sutton. Agent learning using action-dependent learning rates in computer role-playing games. In *Proceedings of the Fourth Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE)*, 2008.
- [98] Richard Sutton, Csaba Szepesvari, **Alborz Geramifard**, and Michael Bowling. Dyna-style planning with linear function approximation and prioritized sweeping. In *Proceedings of the Twenty-Fourth Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 528–536, 2008. Accept: 28%.
- [99] Umar Syed, Robert Schapire, and Michael Bowling. Apprenticeship learning using linear programming. In *Proceedings of the Twenty-Fifth International Conference on Machine Learning (ICML)*, pages 1032–1039, 2008. Accept: 27%.

- [100] Michael Bowling, **Michael Johanson**, Neil Burch, and Duane Szafron. Strategy evaluation in extensive games with importance sampling. In *Proceedings of the Twenty-Fifth International Conference on Machine Learning (ICML)*, pages 72–79, 2008. Accept: 27%.
- [101] Michael Bowling, **Alborz Geramifard**, and David Wingate. Sigma point policy iteration. In *Proceedings of the Seventh International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 379–386, 2008. Accept: 22%.
- [102] **James Neufeld**, Jason Roberts, Stephen Walsh, Michael Sokolsky, Adam Milstein, and Michael Bowling. Autonomous geocaching: Navigation and goal finding in outdoor domains. In *Proceedings of the Seventh International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 47–54, 2008. Accept: 22%.
- [103] Nathan Sturtevant, H. James Hoover, Jonathan Schaeffer, Sean Gouglas, Michael Bowling, Finnegan Southey, Matthew Bouchard, and Ghassan Zabaneh. Multidisciplinary students and instructors: A second-year games course. In *Proceedings of the Thirty-Ninth ACM Technical Symposium on Computer Science Education (SIGCSE)*, pages 383–387, 2008. Accept: 31%.
- [104] **Tao Wang**, Daniel Lizotte, Michael Bowling, and Dale Schuurmans. Stable dual dynamic programming. In *Advances in Neural Information Processing Systems 20 (NIPS)*, pages 713–720, 2008. Accept: 22%.
- [105] **Martin Zinkevich**, **Michael Johanson**, Michael Bowling, and Carmelo Piccione. Regret minimization in games with incomplete information. In *Advances in Neural Information Processing Systems 20 (NIPS)*, pages 905–912, 2008. Accept: 22%.
- [106] **Michael Johanson**, **Martin Zinkevich**, and Michael Bowling. Computing robust counter-strategies. In *Advances in Neural Information Processing Systems 20 (NIPS)*, pages 1128–1135, 2008. A longer version is available as a University of Alberta Technical Report, TR07-15, Accept: 22%.
- [107] **Martin Zinkevich**, Michael Bowling, and **Neil Burch**. A new algorithm for generating equilibria in massive zero-sum games. In *Proceedings of the Twenty-Second Conference on Artificial Intelligence (AAAI)*, pages 788–793, 2007. Accept: 27%.
- [108] **Nolan Bard** and Michael Bowling. Particle filtering for dynamic agent modelling in simplified poker. In *Proceedings of the Twenty-Second Conference on Artificial Intelligence (AAAI)*, pages 515–521, 2007. Accept: 27%.
- [109] Daniel Lizotte, **Tao Wang**, Michael Bowling, and Dale Schuurmans. Automatic gait optimization with gaussian process regression. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 944–949, 2007. Accept: 35%.
- [110] **Alborz Geramifard**, Michael Bowling, **Martin Zinkevich**, and Richard S. Sutton. iLSTD: Eligibility traces and convergence analysis. In *Advances in Neural Information Processing Systems 19 (NIPS)*, pages 441–448, 2007. Accept: 24%.
- [111] **Martin Zinkevich**, Michael Bowling, **Nolan Bard**, **Morgan Kan**, and **Darse Billings**. Optimal unbiased estimators for evaluating agent performance. In *Proceed-*

ings of the Twenty-First National Conference on Artificial Intelligence (AAAI), pages 573–578, 2006. Accept: 30%.

[112] **Armita Kaboli**, Michael Bowling, and Petr Musilek. Bayesian calibration for Monte Carlo localization. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence (AAAI)*, pages 964–969, 2006. Accept: 30%.

[113] **Achim Rettinger**, **Martin Zinkevich**, and Michael Bowling. Boosting expert ensembles for rapid concept recall. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence (AAAI)*, pages 464–469, 2006. Accept: 30%.

[114] **Alborz Geramifard**, Michael Bowling, and Richard S. Sutton. Incremental least-squares temporal difference learning. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence (AAAI)*, pages 356–361, 2006. Accept: 30%.

[115] **Tao Wang**, Michael Bowling, Pascal Poupart, and Dale Schuurmans. Compact, convex upper bound iteration for approximate POMDP planning. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence (AAAI)*, pages 1245–1251, 2006. Accept: 30%.

[116] Nathan Sturtevant, **Martin Zinkevich**, and Michael Bowling. ProbMaxn: Opponent modeling in n-player games. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence (AAAI)*, pages 1057–1063, 2006. Accept: 30%.

[117] Michael Bowling, **Dana Wilkinson**, and Ali Ghodsi. Subjective mapping. In *New Scientific and Technical Advances in Research (NECTAR) of the Proceedings of the Twenty-First National Conference on Artificial Intelligence (AAAI)*, pages 1569–1572, 2006. Accept: 57%.

[118] Michael Bowling, **Peter McCracken**, Michael James, **James Neufeld**, and **Dana Wilkinson**. Learning predictive state representations using non-blind policies. In *Proceedings of the Twenty-Third International Conference on Machine Learning (ICML)*, pages 129–136, 2006. Accept: 20%.

[119] Nathan Sturtevant and Michael Bowling. Robust game play against unknown opponents. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 713–719, 2006. Accept: 23%.

[120] **Peter McCracken** and Michael Bowling. Online discovery and learning of predictive state representations. In *Advances in Neural Information Processing Systems 18 (NIPS)*, pages 875–882, 2006. Accept: 27%.

[121] Michael Bowling, **Dana Wilkinson**, Ali Ghodsi, and Adam Milstein. Subjective localization with action respecting embedding. In *Proceedings of the International Symposium of Robotics Research (ISRR)*, 2005. Accept: 30%.

[122] Finnegan Southey, Michael Bowling, Bryce Larson, Carmelo Piccione, Neil Burch, Darse Billings, and **Chris Rayner**. Bayes’ bluff: Opponent modelling in poker. In *Proceedings of the Twenty-First Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 550–558, 2005. Accept: 34%.

[123] **Tao Wang**, Daniel Lizotte, Michael Bowling, and Dale Schuurmans. Bayesian sparse sampling for on-line reward optimization. In *Proceedings of the Twenty-Second*

International Conference on Machine Learning (ICML), pages 961–968, 2005. Accept: 27%.

[124] Michael Bowling, Ali Ghodsi, and **Dana Wilkinson**. Action respecting embedding. In *Proceedings of the Twenty-Second International Conference on Machine Learning (ICML)*, pages 65–72, 2005. Accept: 27%.

[125] Michael Bowling and **Peter McCracken**. Coordination and adaptation in impromptu teams. In *Proceedings of the Twentieth National Conference on Artificial Intelligence (AAAI)*, pages 53–58, 2005. Accept: 18%.

[126] **Dana Wilkinson**, Michael Bowling, and Ali Ghodsi. Learning subjective representations for planning. In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 889–894, 2005. Accept: 18%.

[127] Michael Bowling. Convergence and no-regret in multiagent learning. In *Advances in Neural Information Processing Systems 17 (NIPS)*, pages 209–216, 2005. Accept: 25%.

[128] Darse Billings, Aaron Davidson, Terence Schauenberg, Neil Burch, Michael Bowling, Robert Holte, Jonathan Schaeffer, and Duane Szafron. Game tree search with adaptation in stochastic imperfect information games. In *Computers and Games (CG)*, 2004. Accept: 58%.

[129] Michael Bowling, Brett Browning, and Manuela Veloso. Plays as effective multiagent plans enabling opponent-adaptive play selection. In *Proceedings of the Fourteenth International Conference on Automated Planning and Scheduling (ICAPS)*, pages 376–383, 2004. Accept: 31%.

[130] Michael Bowling and Manuela Veloso. Simultaneous adversarial multi-robot learning. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 699–704, 2003. Accept: 21%.

[131] James Bruce, Michael Bowling, Brett Browning, and Manuela Veloso. Multi-robot team response to a multi-robot opponent team. In *Proceedings of the 2003 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2281–2286, 2003. An earlier version appeared in the *IROS Workshop on Collaborative Robotics* (2002).

[132] Gal A. Kaminka and Michael Bowling. Towards robust teams with many agents. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 729–736, 2002. Accept: 26%.

[133] Brett Browning, Michael Bowling, and Manuela Veloso. Improbability filtering for rejecting false positives. In *Proceedings of the 2002 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3038–3043, 2002. Accept: 59%.

[134] Rune Jensen, Manuela Veloso, and Michael Bowling. OBDD-based optimistic and strong cyclic adversarial planning. In *Proceedings of the Sixth European Conference on Planning (ECP)*, pages 265–276, 2001. Accept: 50%.

[135] Michael Bowling and Manuela Veloso. Rational and convergent learning in stochastic games. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1021–1026, 2001. Accept: 24%.

- [136] Michael Bowling and Manuela Veloso. Convergence of gradient dynamics with a variable learning rate. In *Proceedings of the Eighteenth International Conference on Machine Learning (ICML)*, pages 27–34, 2001. Accept: 34%.
- [137] Michael Bowling. Convergence problems of general-sum multiagent reinforcement learning. In *Proceedings of the Seventeenth International Conference on Machine Learning (ICML)*, pages 89–94, 2000.
- [138] Michael Bowling and Manuela Veloso. Bounding the suboptimality of reusing subproblems. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1340–1345, 1999. An earlier version appeared in the *Proceedings of the NIPS Workshop on Abstraction in Reinforcement Learning* (1998), Accept: 26%.
- [139] Manuela Veloso, Michael Bowling, Sorin Achim, Kwun Han, and Peter Stone. CMUnited-98: a team of robotic soccer agents. In *Proceedings of the Eleventh Innovative Applications of Artificial Intelligence (IAAI)*, pages 891–896, 1999.

Additional Refereed Workshops, Symposia, and Letters

- [140] Michael Bowling and Esraa Elelimy. Rethinking the foundations of continual reinforcement learning. In *The Sixth Multi-disciplinary Conference on Reinforcement Learning and Decision Making*, 2025.
- [141] David Abel, André Barreto, Michael Bowling, Will Dabney, Shi Dong, Steven Hansen, Anna Harutyunyan, Khimya Khetarpal, Clare Lyle, Razvan Pascanu, Georgios Piliouras, Doina Precup, Jonathan Richens, Mark Rowland, Tom Schaul, and Satinder Singh. Agency is frame-dependent. In *The Sixth Multi-disciplinary Conference on Reinforcement Learning and Decision Making*, 2025.
- [142] **Esra’a Saleh, John D. Martin, Anna Koop, Arash Pourzarabi**, and Michael Bowling. Should models be accurate? In *The Fifth Multi-disciplinary Conference on Reinforcement Learning and Decision Making*, 2022.
- [143] **John D. Martin**, Joseph Modayil, **Fatima Davelouis Gallardo**, and Michael Bowling. Adapting the function approximation architecture in online reinforcement learning. In *The Fifth Multi-disciplinary Conference on Reinforcement Learning and Decision Making*, 2022.
- [144] **Dustin Morrill**, Amy R. Greenwald, and Michael Bowling. The partially observable history process. In *The AAI Workshop on Reinforcement Learning in Games*, 2022.
- [145] **Zach Holland**, Erik Talvitie, and Michael Bowling. The effect of planning shape on dyna-style planning in high-dimensional state spaces. In *ICML Workshop on Prediction and Generative Modeling in Reinforcement Learning*, 2018.
- [146] **Viliam Lisý** and Michael Bowling. Equilibrium approximation quality of current no-limit poker bots. In *AAAI Workshop on Computer Poker and Imperfect Information Games*, volume WS-17, 2017.
- [147] Marlos C. Machado and Michael Bowling. Learning purposeful behaviour in the absence of rewards. In *ICML Workshop on Abstraction in Reinforcement Learning*,

2016.

[148] Martin Zinkevich, Michael Bowling, and Michael Wunder. The lemonade stand game competition: Solving unsolvable games. *ACM SIGecom Exchanges*, 10(1):35–38, 2011.

[149] **Kevin Waugh**, **Martin Zinkevich**, **Michael Johanson**, Morgan Kan, **David Schnizlein**, and Michael Bowling. A practical use of imperfect recall. In *Proceedings of the Eighth Symposium on Abstraction, Reformulation and Approximation (SARA)*, pages 175–182, 2009. Accept: 62%.

[150] Michael Biggs, Ali Ghodsi, **Dana Wilkinson**, and Michael Bowling. Scalable action respecting embedding. In *Proceedings of the Tenth International Symposium on Artificial Intelligence and Mathematics (ISAIM)*, 2008.

[151] **Tao Wang**, Michael Bowling, and Dale Schuurmans. Dual representations for dynamic programming and reinforcement learning. In *Proceedings of the 2007 IEEE International Symposium on Approximate Dynamic Programming and Reinforcement Learning*, pages 44–51, 2007. Accept: 75%.

[152] **Peter McCracken** and Michael Bowling. Safe strategies for agent modelling in games. In *AAAI Fall Symposium on Artificial Multi-agent Learning*, 2004.

[153] Michael Bowling, Brett Browning, **Allen Chang**, and Manuela Veloso. Plays as team plans for cooperation and adaptation. In *IJCAI Workshop on Issues in Designing Physical Agents for Dynamic Real-Time Environments: World Modelling, Planning, Learning, and Communicating*, 2003.

[154] Michael Bowling and Manuela Veloso. Scalable learning in stochastic games. In *AAAI Workshop on Game Theoretic and Decision Theoretic Agents*, 2002.

[155] Michael Bowling, Rune Jensen, and Manuela Veloso. A formalization of equilibria for multiagent planning. In *AAAI Workshop on Planning with and for Multiagent Systems*, 2002.

[156] Michael Bowling and Manuela Veloso. Motion control in dynamic multi-robot environments. In *Proceedings of the 1999 IEEE International Symposium on Computational Intelligence in Robotics and Automation (CIRA)*, pages 168–173, 1999.

[157] Manuela Veloso, Peter Stone, and Michael Bowling. Anticipation as a key for collaboration in a team of agents: A case study in robotic soccer. In *Proceedings of SPIE Sensor Fusion and Decentralized Control in Robotic Systems II*, volume 3839, 1999.

[158] Michael Bowling and Manuela Veloso. Reusing learned policies between similar problems. In *Proceedings of the AI*IA-98 Workshop on New Trends in Robotics*, 1998.

[159] Michael Bowling, Peter Stone, and Manuela Veloso. Predictive memory for an inaccessible environment. In *Working Notes of the IROS-96 Workshop on RoboCup*, 1996.

Unrefereed Book Chapters, Technical Reports, and Theses

- [160] Richard S. Sutton, Michael Bowling, and Patrick M. Pilarski. The Alberta plan for AI research, 2023.
- [161] Jesse Farebrother, Marlos C. Machado, and Michael Bowling. Generalization and regularization in DQN, 2020.
- [162] Michael Bowling, Rune Jensen, and Manuela Veloso. Multiagent planning in the presence of multiple goals. In Rene Jorna, Wout van Wezel, and Alex Meystel, editors, *Planning in Intelligent Systems: Aspects, Motivations, and Methods*, Intelligent Series, pages 345–371. Wiley, 2005.
- [163] Michael Bowling. *Multiagent Learning in the Presence of Agents with Limitations*. PhD thesis, Computer Science Department, Carnegie Mellon University, Pittsburgh, PA, 2003. Available as technical report CMU-CS-03-118.
- [164] Michael Bowling and Manuela Veloso. An analysis of stochastic game theory for multiagent reinforcement learning. Technical report CMU-CS-00-165, Computer Science Department, Carnegie Mellon University, 2000.
- [165] Michael Bowling. Robocup rescue: Agent development kit, version 0.4, 2000. Manual.
- [166] Jr. David M. McKeown, Michael Bowling, G. Edward Bulwinkle, Steven Douglas Cochran, Stephen J. Ford, Wilson A. Harvey, Dirk Kalp, Chris McGlone, Jeff McMahill, Michael F. Polis, Jefferey A. Shufelt, and Daniel Yocum. Research in image understanding and automated cartography: 1995-1996. In *Proceedings of the DARPA Image Understanding Workshop*, pages 779–812, 1997.

ACADEMIC
SERVICE**Journal Action Editor***Transactions on Machine Learning Research*

2022–Present

Journal Associate Editor*Artificial Intelligence*

2021–2024

Associate Program Chair*Association for the Advancement of Artificial Intelligence*

2023

Journal Associate Editor*Journal of Machine Learning Research*

2015–2017

Journal Associate Editor*Journal of Artificial Intelligence Research*

2009–2015

Senior Program Committee Member

International Conference on Machine Learning (ICML) 2006,2008,2012,2020

International Joint Conference on Artificial Intelligence (IJCAI) 2009,2011

Conference on Artificial Intelligence (AAAI) 2007,2013,2014,2016,2018,2019

Conference on Autonomous Agents and Multiagent Systems (AAMAS) 2006,2014

Neural Information Processing Systems (NeurIPS) 2008,2009,2014,2019

Artificial Intelligence and Statistics (AISTATS) 2010

Program Committee Member

Conference on Artificial Intelligence (AAAI)	2004–2006,2008,2010,2015
Neural Information Processing Systems (NIPS)	2006,2007,2011–2013
International Conference on Machine Learning (ICML)	2002–2005,2007,2008,2025
International Joint Conference on Artificial Intelligence (IJCAI)	2005–2007
International Conference on Autonomous Agents and Multiagent Systems (AAMAS)	2002,2004,2008,2012
ACM Electronic Commerce (EC)	2012
Multi-disciplinary Conference on Reinforcement Learning and Decision Making (RLDM)	2013,2014
Conference on AI and Interactive Digital Entertainment (AIIDE)	2010–2012
Robotics: Science and Systems (RSS)	2006,2007
International Conference on Artificial Intelligence and Statistics (AISTATS)	2008,2009
Symposium on Abstraction, Reformulation and Approximation (SARA)	2009
International Symposium on Artificial Intelligence and Mathematics (ISAIM)	2008
Canadian Conference on Artificial Intelligence (AI)	2004
RoboCup Symposium	2000,2001–2003

Journal Editorial Board

<i>Machine Learning</i>	2006–2020
-------------------------	-----------

AAAI Feigenbaum Prize Committee	2020
--	------

AAAI Games Night Co-Czar	2016,2017,2018
---------------------------------	----------------

AI Index Steering Committee	2016
One-Hundred Year Study on Artificial Intelligence (AI-100), Stanford University	

Doctoral Consortium Mentor	2011–2012
For the AAAI Doctoral Consortium.	

Steering Committee	2009–2017
Annual Computer Poker Competition.	

Special Issue Guest Co-Editor	2006
For <i>Machine Learning</i> on “Machine Learning and Games”. Co-editors with Johannes Fuernkranz, Thore Graepel, and Ron Musick.	

Co-Chair of the AI Video Awards	2009
At the 2009 International Joint Conference on Artificial Intelligence (IJCAI).	

Intelligent Systems Demonstrations Co-Chair	2008
For the AAAI Conference on Artificial Intelligence.	

Volunteer Chair	2006
For the International Conference on Machine Learning (ICML).	

Chair of the Educational Track of the AI Video Awards	2008
As part of the AAAI Conference on Artificial Intelligence.	

Exhibitions Co-Chair	2005
For the International Conference on Intelligent Robots and Systems (IROS). Co-chaired	

with Bruce Digney.

Workshop Organization

AAAI Workshop on “General Competency in Video Games”	2015
AAAI Workshop on “Applied Adversarial Reasoning and Risk Modelling”	2011
NIPS Workshop on “Learning and Planning from Batch Time Series Data”	2010
AAAI Workshop on “Multiagent Learning”	2005
AAAI Fall Symposium on “Artificial Multi-Agent Learning”	2005
CG Workshop on “Agents in Computer Games” (Co-chaired with Regis Vincent)	2002

RoboCup Organization

2000–2004

Chair of the RoboCup small-size league technical committee for 2004, which is responsible for competition rules and guiding the league’s research agenda. Chair of the small-size league competition for the RoboCup American Open in 2002 (co-chair), 2003, and 2004. I am also a member of the organizing committee for the RoboCup special interest group on multiagent learning started in 2001. I was head referee for the RoboCup small-size league in 2000 and 2001.

UNIVERSITY SERVICE

Alberta Machine Intelligence Institute

Chair of the Fellows Membership Committee	2023–Present
Chair of the Fellows Advisory Committee	2019–Present
Chair of Resource Allocation Panel	2018–2022
Member of the Pan-Canadian AI Strategy National Program Committee	2018–2021

University

ESIO Trailhead Program Liasion	2024
Chair of the AI Cohort Hiring Committee	2024–2025
Member representative on Compute Canada	2014–2016
Co-Director of the Game Development Certificate	2012–2022
TLAT Sub-Committee on eTextbooks	2012

Faculty of Science

Faculty of Science Awards Selection Committee	2014–2016
Faculty of Science representative on Faculty of Arts Council	2013–2016
Faculty Evaluation Committee	2010–2011

Departmental

Faculty Recruiting Committee	2024 (2 Committees)
Scholarship Review Committee	2023
Faculty Recruiting Committee	2015–2016, 2018–2019, 2023
Graduate Admissions Committee	2018–Present
Graduate Program Committee	2014–2017
Undergraduate Curriculum Committee	2011–2013
Faculty Recruiting Committee	2010
Undergraduate Program Committee	2008–2009
Computing Infrastructure Committee	2007–2008
Graduate Program Committee	2005–2006
Distinguished Lecture Co-Organizer	2006–2007

RESEARCH GRANTS	Canada CIFAR AI Chair	2026–2030
	\$150,000 per year.	
	NSERC Alliance — Alberta Innovates Advance: Stream I	2026–2028
	\$20,000 per year.	
	Google DeepMind Chair	2024–2029
	\$50,000 per year.	
	NSERC Discovery Grant	2024–2029
	\$55,000 per year. “Useful Even If Inaccurate: More Effective and Efficient Use of World Models in Reinforcement Learning”.	
	Canada CIFAR AI Chair	2021–2025
	\$120,000 per year.	
	NSERC Discovery Grant	2018–2024
	\$74,000 per year. “Human-Scale Decision Making in Uncertain Environments”.	
	Alberta Economic Development and Trade (EDT)	2017–2018
	\$2,000,000. “Alberta Machine Intelligence Institute (Amii)”. Joint with R. Goebel, R. Greiner, R. Holte, P. Pilarski, D. Schuurmans, R. Sutton, C. Szepesvari, M. White, Y. Yasui, O. Zaiane.	
	IBM Faculty Award	2016
	\$51,500. “Deep and Decomposable Game Theoretic Learning Algorithms”.	
	Alberta Innovates Research Centre	2014-2017
	\$2,000,000 per year (on average). “Alberta Innovates Centre for Machine Learning (AICML)”. Joint with R. Goebel, R. Greiner, R. Holte, D. Schuurmans, R. Sutton, C. Szepesvari, Y. Yasui, O. Zaiane. Third round of funding for AICML.	
	Alberta Gaming Research Institute Grant	2012–2013
	\$10,000 total. “Tools for the Analysis of Multi-Player Poker Games”. Joint with C. Archibald.	
	NSERC Discovery Grant	2012–2017
	\$42,000 per year. “Human-Scale Game Theory in Imperfect Information Settings”.	
	Kule Institute for Advanced Study	2011–2012
	\$14,720 total. “Creating a Gaming Community Platform Driven by Students in Real World Classrooms”. Joint with S. Gouglas, G. Rockwell, P. Boechler, M. Carbonaro.	
	Alberta Gaming Research Institute Grant	2009–2010
	\$10,000 total. “Calling Their Bluff: Investigating the Accuracy of Skill Assessment Among Poker Players”. Joint with T. L. MacKay, D. Hodgins.	
	Alberta Innovates Research Centre	2009–2014
	\$2,000,000 per year (on average). “Alberta Innovates Centre for Machine Learning (AICML)”. Joint with R. Goebel, R. Greiner, R. Holte, R. Mitchell, D. Schuurmans, R. Sutton, C. Szepesvari, Y. Yasui, O. Zaiane. Second round of funding for AICML.	
	Alberta Health Services	2009–2011

\$1,000,000 total. “Ingenuity Research Grant”. Joint with R. Goebel, R. Greiner, R. Holte, R. Mitchell, D. Schuurmans, R. Sutton, C. Szepesvari, Y. Yasui, O. Zaiane.

NSERC Discovery Grant 2007–2012
\$28,000 per year. “Data-Driven Statistical Agent Modelling”.

Alberta Ingenuity New Faculty Grant 2006–2008
\$100,000 per year. “Subjective Models for Autonomous Robots”.

NSERC Strategic Project 2005–2007
\$195,000 per year. “Intelligent Agents for Interactive Entertainment”. Joint with J. Schaeffer, M. Buro, M. Mueller, D. Schuurmans, and D. Szafron.

Alberta Ingenuity Research Centre 2003–2008
\$1,800,000 per year (on average). “Alberta Ingenuity Centre for Machine Learning (AICML)”. Joint with R. Greiner, R. Goebel, R. Holte, J. Schaeffer, D. Schuurmans, R. Sutton, C. Szepesvari. The centre was started in 2002. I joined in 2003.

NSERC Discovery Grant 2004–2007
\$26,034 per year. “Learning in Impromptu Multiagent Systems”.

University of Alberta Startup Funds 2003
\$80,000.

SUPERVISION

Post-Doctoral Fellows

S. Parisi (11/2022–01/2025).

J. Martin (5/2021–10/2022).

Y. Satsangi (3/2019–3/2020).

V. Lisy (8/2015–2/2017).

K. Waugh (9/2015–8/2016).

C. Archibald (9/2011–9/2013).

J. Veness (1/2011–10/2012).

M. Zinkevich (9/2005–8/2007).

Y. Engel (3/2005–8/2007).

Ph.D. Students

P. M. Panahi (1/2026–Present; with A. White).

A. Kazemipour (9/2025–Present).

J. Adkins (1/2025–Present; with A. White).

D. Gomez (1/2025–Present; with M. Machado).

A. Orenstein (8/2023–Present).

A. Koop (4/2012–Present).

M. Mohammedalamen (1/2022–Present). Thesis title: “Reinforcement Learning with Partially Observable Rewards”

- Z. Ahmad (1/2017–9/2025); with L. Lelis). Thesis title: “Strengthening Integration Between Search and Learning for Planning”
- T. Davis (6/2015–8/2024). Withdrew.
- D. Morrill (2/2016–6/2022). Thesis title: “Hindsight Rational Learning for Sequential Decision-Making: Foundations and Experimental Applications”
- M. Schmid (2/2016–8/2021; with M. Hladik). Thesis title: “Search in Imperfect Information Games”
- M. Machado (9/2013–2/2019). Thesis title: “Efficient Exploration in Reinforcement Learning through Time-Based Representations”
- N. Burch (9/2011–9/2017). Thesis title: “Time and Space: Why Imperfect Information Games are Hard”
- K. Chen (9/2009–9/2015). Withdrew to co-found Alieo Games.
- J. Neufeld (4/2008–3/2016). Thesis title: “Adaptive Monte Carlo Integration”.
- N. Bard (3/2008–3/2016). Thesis title: “Online Agent Modelling in Human-Scale Problems”.
- M. Johanson (9/2008–1/2016). Thesis title: “Robust Strategies and Counter-Strategies: From Superhuman to Optimal Play”.
- M. White (12/2009–12/2014). Thesis title: “Regularized Factor Models”
- C. Rayner (9/2008–11/2014). Thesis title: “Optimization for Heuristic Search”.
- M. Gendron-Bellemare (4/2011–7/2013; with J. Veness). Thesis title: “Fast, Scalable Algorithms for Reinforcement Learning in High Dimensional Domains”.
- A. Afkanpour (3/2007–3/2013; with C. Szepesvári). Thesis title: “Multiple Kernel Learning with Many Kernels”.
- M. Lanctot (9/2008–11/2012). Thesis title: “Monte Carlo Sampling and Regret Minimization for Equilibrium Computation and Decision-Making in Large Extensive Form Games”.
- D. Wilkinson (4/2006–11/2007; with D. Schuurmans). Thesis title: “Subjective Mapping”.
- T. Wang (10/2003–8/2007; with D. Schuurmans). Thesis title: “New Representations and Approximations for Sequential Decision Making under Uncertainty”.

M.Sc. Students

- J. Chen (9/2026–Present).
- G. A. Delos Santos (1/2026–Present; with B. Kacsmar).
- S. Akins (3/2025–Present).
- M. Horgen (9/2024–Present).
- L. Kapeluck (3/2023–Present; with J. Martin).
- A. Kazemipour (4/2024–8/2025; with M. Taylor). Thesis title: “Model-Based Exploration in Monitored Markov Decision Processes”
- M. R. Hanzaki (2/2024–8/2025). Thesis title: “Using Discrete Overlapping Partitions for Count-based Exploration”

- J. Adkins (3/2023–1/2025; with A. White). Thesis title: “A Method for Evaluating Hyperparameter Sensitivity in Reinforcement Learning”
- F. Gallardo Davelouis (1/2019–1/2024; with J. Martin). Thesis title: “On the benefits of sparsity in value function approximators for Reinforcement Learning”
- D. Gomez (1/2023–12/2023; with M. Machado). Thesis title: “Proper Laplacian Representation Learning”
- B. Burega (4/2020–12/2023; with J. Martin). Thesis title: “Adaptive Search Control through Meta-Gradient Reinforcement Learning”
- E. Saleh (1/2021–9/2023). Thesis title: “Should Models Be Accurate?”
- A. Trudeau (10/2018–12/2022). Thesis title: “Targeted Search in AlphaZero for Effective Policy Improvement”
- P. Behboudian (4/2018–2/2020; with M. Taylor). Thesis title: “Useful Policy Invariant Shaping from Arbitrary Advice”
- A. Pourzarabi (4/2018–1/2020). Rest in peace.
- T. Jafferjee (1/2018–9/2019; with M. White). Thesis title: “Chasing Hallucinated Value: A Pitfall of Dyna Style Algorithms with Imperfect Environment Models”
- F. Li (4/2016–8/2019; with D. Schuurmans). Thesis title: “Ease-of-teaching and Language Structure in Emergent Communication”
- R. Post (4/2015–12/2018). Withdrew.
- Z. Holland (4/2016–8/2018; with E. Talvitie). Thesis title: “The Effect of Planning Shape on Dyna-style Planning in High-dimensional State Spaces”
- T. Yee (4/2014–12/2017). Withdrew.
- D. Morrill (1/2013–2/2016). Thesis title: “Using Regret Estimation to Solve Games Compactly”.
- T. Davis (9/2012–6/2015). Thesis title: “Using Response Functions for Strategy Training and Evaluation”.
- U. das Gupta (4/2013–11/2014; with E. Talvitie). Thesis title: “Adaptive Representation for Policy Gradient”.
- S. Srinivasan (4/2013–11/2014; with E. Talvitie). Thesis title: “State Generalization in UCT”.
- J. Davidson (4/2011–1/2014; with C. Archibald). Thesis title: “The Baseline Approach to Agent Evaluation”.
- M. Vafadoost (4/2011–8/2013; with J. Veness). Thesis title: “Temporal Abstraction in Monte Carlo Tree Search”.
- P. Mazrooei (4/2011–9/2012; with C. Archibald). Thesis title: “Collusion Detection in Sequential Games”.
- M. Joya (3/2010–6/2012; with D. Schuurmans). Thesis title: “Methods for Automatic Heart Sound Identification”.
- K. Davison (3/2006–10/2010). Withdrew.
- Y. Naddaf (3/2008–3/2010). Thesis title: “Learning to Play Generic Atari 2600 Console Games”.
- M. White (9/2008–12/2009; with D. Schuurmans). Thesis title: “A General Framework

for Reducing Variance in Agent Evaluation”.

D. Schnizlein (3/2008–7/2009; with D. Szafron). Thesis title: “State Translation in No-Limit Poker”.

K. Waugh (3/2008–7/2009; with D. Schuurmans). Thesis title: “Abstraction in Large Extensive Games”.

J. Neufeld (3/2006–4/2008). Thesis title: “Autonomous Outdoor Navigation and Goal Finding”.

N. Bard (3/2005–3/2008). Thesis title: “Using State Estimation for Dynamic Opponent Modelling”.

M. Johanson (3/2005–10/2007). Thesis title: “Robust Strategies and Counter-Strategies: Building a Champion Level Computer Poker Player”.

A. Kaboli (1/2005–9/2007; with P. Musilek). Thesis title: “Bayesian Calibration for Monte Carlo Localization”.

A. Geramifard (3/2005–12/2006; with R. Sutton). Thesis title: “Incremental Least-Squares Temporal Difference Learning”.

P. McCracken (3/2004–9/2005). Thesis title: “An Online Algorithm for Learning Predictive State Representations”.

Undergraduate Summer Researchers

C. Rayner (2005), T. Pittman (2005), J. Ho (2005), M. Lednicky (2006–2008), G. Ward (2012–2013), B. Kuzma (2013–2014), J. Chu (2013), J. Emery (2014–2015), N. Westbury (2015), N. Yasui (2015), J. Ryan (2017; HS student), J. Farebrother (2018–2020), J. Chen (2023–2024), G. A. Delos Santos (2024), B. Sapara (2024), Y. A. Kumar (2025), R. Dovganyuk (2026)

Visiting Students / Internships

A. Rettinger (3/2005–6/2005), M. Mohammedalamen (2/2020–12/2020), D. Sychrovsky (1/2024–6/2024), M. Azpeitia (2/2026–5/2026).

THESIS EXAMINATION

Ph.D. Examiner: C. Kharyal (5/2026; Alberta, CS), S. Poo Hernandez (3/2025; Alberta, CS), K. De Asis (7/2024; Alberta, CS), K. Young (3/2024; Alberta, CS), M. Schlegl (7/2023; Alberta, CS), C.K. Ling (5/2023; CMU, SCS), J. Mei (8/2021; Alberta, CS), C. Sherstan (8/2020; Alberta, CS), P. Joulani (12/2018; Alberta, CS), R. Mahmood (4/2018; Alberta, CS), F. Mirzazadeh (11/2016; Alberta, CS), A. White (5/2015; Alberta, CS), S. Ganzfried (4/2015; CMU, CS), R. Zhao (3/2015; Alberta, CS), J. Hawkin (8/2014; Alberta, CS), R. Gibson (12/2013; Alberta, CS), L. Lelis (7/2013; Alberta, CS), D. Chakraborty (8/2013; UTexas-Austin, CS), G. Bartok (6/2012; Alberta, CS), A. m. Farahmand (9/2011; Alberta, CS), J. Long (8/2011; Alberta, CS), D. Lizotte (7/2008; Alberta, CS), A. Milstein (3/2008; Waterloo, CS),

M.Sc. Examiner: R. MacQueen (8/2023; Alberta, CS), H. Farrahi (4/2021; Alberta, CS), V. Chockalingam (8/2020; Alberta, CS), C. Ma (1/2020; Alberta, CS), W. Chung (8/2019; Alberta, CS), T. Sajed (5/2019; Alberta, CS), B. Rafiee (8/2018; Alberta, CS), Z. Ahmad (1/2016; Alberta, CS), Y. Wu (7/2016; Alberta, CS), A. Edwards (4/2016;

Alberta, Rehab), K. Atui (9/2015; Alberta, CS), S. Hernandez (9/2015; Alberta, CS), P. Jin (12/2014; Alberta, CS), M. Bastani (11/2013; Alberta, CS), A. Sanabria (9/2013; Alberta, CS), J. Wen (7/2013; Alberta, CS), A. Karamaoun (6/2013; Alberta, Physics), M. Burden (8/2012; Alberta, HuCo), D. Huntley (10/2011; Alberta, CS), A. Sharifi (9/2010; Alberta, CS), S. Hladky (6/2009; Alberta, CS), D. Rayner (6/2008; Alberta, CS), A. Visatemongkolchai (9/2007; Alberta, CS), P. Pytlak (8/2007; Alberta, ECE), C. Onuczko (5/2007; Alberta, CS), G. Xaio (2/2007; Alberta CS), M. Kan (1/2007; Alberta CS), F. Ahmad (1/2007; Alberta, ECE), E. Rafols (9/2006; Alberta, CS), C. Paduraru (12/2006; Alberta, CS), B. Hoehn (12/2005; Alberta, CS), R. Guanlao (9/2004; Alberta, ECE), S. Verret (5/2004; Alberta, CS)

TEACHING

Undergraduate Courses

INTD350: Game Design Principles and Practice (W2013,F2013,W2015,W2016,W2018–W2026)

CMPUT450: Computers and Games (W2024,W2026)

CMPUT274/275: Honors Intro to Computing I/II (F2011–W2014)

CMPUT101: Introduction to Computing (W2011)

CMPUT229: Computer Organization and Architecture I (F2008)

CMPUT250: Computers and Games (F2005–W2009,F2015,W2016)

CMPUT412: Experimental Mobile Robotics (W2005,W2006)

CMPUT204: Algorithms I (W2004,F2004)

Graduate Courses

CMPUT655: Reinforcement Learning I (F2024)

CMPUT658: Applied Game Theory (F2010)

CMPUT608: Robotics Challenge (F2005,F2007)

CMPUT608: Learning in Multiagent Systems (W2004)

15-889: Multiagent Systems, Theory and Practice (2001)

Co-taught at CMU with Prof. Manuela Veloso and Dr. Gal Kaminka.

PERSONAL
INFORMATION

Citizenship: United States, Canada

Family: Married with three boys.

Hobbies: Board game connoisseur, curling, disc golf.