

John D. Martin

✉ jdmartin86@gmail.com

🌐 [jdmartin86.github.io](https://github.com/jdmartin86)

in [johndmartinjr](https://www.linkedin.com/in/johndmartinjr)

🐦 [@ejdmartin86](https://twitter.com/ejdmartin86)

My research focuses on the science of artificial intelligence, framed around the concepts of reinforcement learning. I want to uncover the general laws that give rise to artificial minds, and use these insights to design competent decision making systems for complex problems.

Education

- 2021 **Ph.D. Mechanical Engineering**, Stevens Institute of Technology.
Advisor: *Brendan Englot*
Thesis: *Reinforcement Learning Algorithms for Representing and Managing Uncertainty in Robotics.*
- 2013 – 2015 **Graduate studies in Computer Science**, Columbia University.
- 2012 **B.S. Physics, B.S. Aerospace Engineering**, University of Maryland.

Employment History

- 2024 – ⋯ **Research Fellow**, Openmind Research Institute.
I conduct fundamental research in artificial intelligence with a focus on reinforcement learning. Specifically, I focus on how an agent can exploit its environment as an effective form of memory.
- 2023 – ⋯ **Adjunct Professor**, University of Alberta, Department of Computing Science.
I supervise graduate students studying artificial intelligence and reinforcement learning. My students study algorithms for model-based RL in the continual setting.
- 2022 – 2024 **Research Scientist**, Intel Labs.
I performed applied and fundamental research. My applied work focused on the detection of CPU errors with code-specialized LLMs, fine-tuned for particular failure modalities. My fundamental work focused on sparsity in neural networks, model-based RL, and the Reward Hypothesis.
- 2021 – 2022 **Postdoctoral Fellow**, University of Alberta, Department of Computing Science.
Advisor: *Michael Bowling*
I studied the limits of reward as a mechanism to express goals, meta-learning algorithms for reinforcement learning, and algorithms for sample-efficient planning.
- Summer 2020 **Research Scientist Intern**, DeepMind.
Advisor: *Joseph Modayil*
I studied RL algorithms for constructing sparse neural networks from unorganized observations.
- 2019 – 2020 **Student Researcher / Research Scientist Intern**, Google Brain.
Advisor: *Marc G. Bellemare*
I studied algorithms for detecting catastrophic interference in deep neural networks.
- 2017 – 2019 **Technical Consultant**, Piasecki Aircraft.
Focus areas: Conceptual design of experimental aircraft, proposal writing.
- 2012 – 2015 **Robotics and Flight Controls Engineer**, Sikorsky Aircraft.
Focus areas: Design of motion planning and control algorithms, automation of full-scale S-76.



Publications

Articles

- 1 **J. D. Martin**, F. Mince, E. Saleh, and A. Pajak, “Artifacts as Memory Beyond the Agent Boundary,” *arXiv preprint arXiv:2604.08756 (Under Review at RLC)*, 2026.
- 2 E. S. Lupu, P. Spieler, K. Javed, *et al.*, “The Open Ant: A Robot Platform for Reinforcement Learning Research,” *Under Review at RLC*, 2026.
- 3 F. Dave-Louis, **J. D. Martin**, and M. Bowling, “On the Interplay Between Sparsity and Training in Deep Reinforcement Learning,” *arXiv preprint arXiv:2501.16729*, 2025.
- 4 **J. D. Martin**, “The Methodological Tangle of AI Research,” *RLDM (Oral)*, 2025.

- 5 **J. D. Martin**, B. Burega, L. Kapeluck, and M. Bowling, “Meta-Gradient Search Control: A Method for Improving the Efficiency of Dyna-style Planning,” *arXiv preprint arXiv:2406.19561*, 2024.
- 6 R. Rafailov, K. B. Hatch, V. Kolev, **J. D. Martin**, M. Phielipp, and C. Finn, “MOTO: Offline Pre-training to Online Fine-tuning for Model-based Robot Learning,” in *7th Annual Conference on Robot Learning*, 2023.
- 7 **J. D. Martin**, M. Bowling, D. Abel, and W. Dabney, “Settling the Reward Hypothesis,” in *International Conference on Machine Learning*, PMLR, 2023.
- 8 B. Burega, **J. D. Martin**, and M. Bowling, “Learning to Prioritize Planning Updates in Model-based Reinforcement Learning,” *NeurIPS Workshop on Meta Learning*, 2022.
- 9 E. Saleh, **J. D. Martin**, A. Koop, A. Pourzarabi, and M. Bowling, “Should Models Be Accurate?” *arXiv preprint arXiv:2205.10736*, 2022.
- 10 **J. D. Martin**, P. Szenher, X. Lin, and B. Englot, “The Stochastic Road Network Environment for Robust Reinforcement Learning,” *ICRA Workshop on Releasing Robots into the Wild*, 2022.
- 11 **J. D. Martin**, “Time to Take Embodiment Seriously,” *RLDM RL as Agency Workshop (Oral)*, 2022.
- 12 **J. D. Martin** and J. Modayil, “Adapting the Function Approximation Architecture in Online Reinforcement Learning,” *arXiv preprint arXiv:2106.09776*, 2021.
- 13 **J. D. Martin**, “Reinforcement Learning Algorithms for Representing and Managing Uncertainty in Robotics,” Ph.D. dissertation, Stevens Institute of Technology, Hoboken, New Jersey, USA, 2021.
- 14 F. Chen, **J. D. Martin**, Y. Huang, J. Wang, and B. Englot, “Autonomous exploration under uncertainty via deep reinforcement learning on graphs,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2020, pp. 6140–6147.
- 15 J. McConnell, **J. D. Martin**, and B. Englot, “Fusing concurrent orthogonal wide-aperture sonar images for dense underwater 3D reconstruction,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2020, pp. 1653–1660.
- 16 W. Fedus, D. Ghosh, **J. D. Martin**, M. G. Bellemare, Y. Bengio, and H. Larochelle, “On Catastrophic Interference in Atari 2600 Games,” *arXiv preprint arXiv:2002.12499*, 2020.
- 17 **J. D. Martin**, M. Lyskawinski, X. Li, and B. Englot, “Stochastically Dominant Distributional Reinforcement Learning,” in *International Conference on Machine Learning*, PMLR, 2020, pp. 6745–6754.
- 18 **J. D. Martin**, K. Doherty, C. Cyr, B. Englot, and J. Leonard, “Variational filtering with copula models for SLAM,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2020, pp. 5066–5073.
- 19 **J. D. Martin**, J. Wang, and B. Englot, “Sparse Gaussian Process Temporal Difference Learning for Marine Robot Navigation,” in *Conference on Robot Learning*, PMLR, 2018, pp. 179–189.
- 20 **J. D. Martin** and B. Englot, “Extending model-based policy gradients for robots in heteroscedastic environments,” in *Conference on Robot Learning*, PMLR, 2017, pp. 438–447.

Invited Talks

- 2026  **Artifacts as Memory Beyond the Agent Boundary.**
 University of Massachusetts, Amherst
 Institute of Philosophy, School of Advanced Studies, University of London
 Noah’s Ark Lab, Toronto
 Alberta Machine Intelligence Institute, Edmonton
 UTAR, Kuala Lumpur, Malaysia
- 2025  **AI for the Mind: The Quest Ever Thus.**
 California Institute of Technology
 NASA Jet Propulsion Laboratory

Invited Talks (continued)

- 2024 ■ **The Methodological Tangle of AI Research.**
Trinity College (RLDM), Dublin
University of Alberta, Edmonton
Google DeepMind, Seattle
- 2024 ■ **Reinforcement Learning and The Extended Mind Hypothesis.**
Cohere for AI virtual talk
- 2023 ■ **The Issaquah Plan.**
Google DeepMind, Seattle
- 2022 ■ **Learning to Prioritize Planning Updates in Model-based Reinforcement Learning.**
University of Massachusetts, Amherst
- 2021 ■ **Adapting the Function Approximation Architecture in Online Reinforcement Learning.**
Google AI, Sparsity Reading Group
- 2020 ■ **Uncertainty, Perception, and Their Lessons for Creating General-purpose Robots.**
University of California, Berkeley
- 2019 ■ **From Tasks to Timescales: A path to generalization in reinforcement learning.**
Massachusetts Institute of Technology
DeepMind, Edmonton
Google Robotics, New York

Academic Service

Masters Thesis Advising

- 2025 ■ **Haruto Tanaka**, University of Alberta, thesis committee member.
- 2024 – ■ **Deepak Ranganatha Sastry Mamillapalli**, University of Alberta, co-advised with Matt Taylor.
- 2023 – ■ **Luke Kapeluck**, University of Alberta, co-advised with Michael Bowling.
- 2022 – 2023 ■ **Bradley Burega**, University of Alberta, co-advised with Michael Bowling.
- 2021 – 2024 ■ **Fatima Davelouis**, University of Alberta, co-advised with Michael Bowling.

Organizer

- 2026 ■ **Reinforcement Learning and Philosophy Workshop:**
- 2025 ■ **Saving the Phenomena of Minds:** An RLDM workshop for examining RL as an computational theory of mental phenomena.
- 2024 – 2026 ■ **Finding the Frame Workshop:** An RLC workshop for examining conceptual frameworks in RL.
- 2023 – ■ **Seattle Minds and Machines Meetup:** a seminar series for Reinforcement Learning in Computer Science and Computational Neuroscience researchers in the Seattle-area.

Workflow Chair








- 2022 ■ AAI.

Program Chair

- 2026 ■ **Openmind Winter School on AI and RL.**
- 2023 ■ **Barbados RL Workshop.**
- 2021 ■ **NAAMII Winter AI School.**
- 2020 ■ **ICML RL Social.**

Academic Service (continued)

Reviewer

- 2026  ICML, RLC, Workshops.
- 2025  ICML, RLC, RLDM, TMLR, CogSci, AAI, Workshops.
- 2024  ICML, TMLR, Workshops.
- 2023  ICML, TMLR, Nature Machine Intelligence.
- 2022  ICML, NeurIPS, RLDM.
- 2021  ICML, NeurIPS, ICLR.
- 2020  ICML, NeurIPS, ICRA, WAFR, Workshops.

Mentor

- 2023  Cohere Labs Community.
- 2022  Neuromatch Academy.
- 2020  NeurIPS New in ML Workshop.

Teaching Experience



Primary Instructor

- Winter 2021  RL Lecture Series, Nepal Applied Mathematics and Informatics Institute.



Guest Lecturer

- 2017, 2020, 2021  Advanced Robotics, Stevens Institute of Technology.

Technical Skills

- Coding  Python, \LaTeX , C, C++, R, OCaml, ...
- Libraries  JAX, Haiku, NumPy, Tensorflow, Pandas, Docker, Kubernetes, ROS, ...

Honors and Awards

- 2019 – 2020  Robert Brooks Stanley Doctoral Fellow, Two-time recipient.
- 2015  Department of Homeland Security Doctoral Fellow.
-  Howard Hughes Award, American Helicopter Society.

References

Available on Request