




John D. Martin

✉ jdmartin86@gmail.com

🌐 jdmartin86.github.io

Education

- 2015 – 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  **M.Sc. Computer Science (Incomplete)** Columbia University.
- 2009 – 2012  **B.S. Physics & Aerospace Engineering** University of Maryland.

Publications

Conference Papers

- 1 **J. D. Martin**, M. Bowling, D. Abel, and W. Dabney, “Settling the Reward Hypothesis,” in *International Conference on Machine Learning*, PMLR, 2023.
- 2 R. Rafailov, K. B. Hatch, V. Koley, **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.
- 3 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.
- 4 **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.
- 5 **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.
- 6 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.
- 7 **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.
- 8 **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.








Articles

- 1 **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.
- 2 B. Burega, **J. D. Martin**, and M. Bowling, “Learning to prioritize planning updates in model-based reinforcement learning,” *NeurIPS Workshop on Meta Learning*, 2022.
- 3 **J. D. Martin**, “Time to take embodiment seriously,” *RLDM RL as Agency Workshop (Oral)*, 2022.
- 4 **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.
- 5 E. Saleh, **J. D. Martin**, A. Koop, A. Pourzarabi, and M. Bowling, “Should models be accurate?” *arXiv preprint arXiv:2205.10736*, 2022.
- 6 **J. D. Martin** and J. Modayil, “Adapting the function approximation architecture in online reinforcement learning,” *arXiv preprint arXiv:2106.09776*, 2021.
- 7 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.

Employment History





- 2023 –  **Adjunct Professor**, University of Alberta, Department of Computing Science.
- 2022 –  **Research Scientist**, Intel Labs.
Focus areas: reinforcement learning, code generation with LLMs, processor design automation
- 2021 – 2022  **Postdoctoral Fellow**, University of Alberta, Department of Computing Science.
Advisor: *Michael Bowling*
- Summer 2020  **Research Scientist Intern**, DeepMind.
Advisor: *Joseph Modayil*
- 2019 – 2020  **Student Researcher / Research Scientist Intern**, Google AI.
Advisor: *Marc G. Bellemare*
- 2017 – 2019  **Engineering 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.

Invited Talks



- 2024  **Reinforcement Learning and The Extended Mind Hypothesis.**
Cohere for AI virtual talk.
- 2023  **The Issaquah Plan.**
Seattle Minds and Machines Meetup, 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
- 2014  **Sikorsky R&D: Motion Planning for Autonomous Rotorcraft.**
Stevens Institute of Technology

Academic Service

Masters Thesis Advising

- 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

- 2024  **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.

Academic Service (continued)

Workflow Chair

2022  AAI.

Program Chair

2023  Barbados RL Workshop.

2021  NAAMII Winter AI School.

2020  ICML Reinforcement Learning Social.

Program Committee

2023  Nature Machine Intelligence.

 TMLR.

2021  ICLR.

2020–2022  NeurIPS.

2020–2024  ICML.

2019  AAI.

 CoRL.

2020  WAFR.

2019  RAL.

2018–2020  ICRA.

2017  IROS.

2020  JOE.

Mentor

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.

Skills

Languages  English, Nepalese.

Coding  Python, C, C++, R, \LaTeX , OCaml, ...



Libraries  JAX, Haiku, Tensorflow, Pandas, NumPy, ...

Miscellaneous Experience

Awards and Achievements

2019 – 2020  Robert Brooks Stanley Doctoral Fellow, Two-time recipient.

Miscellaneous Experience (continued)

- 2015  **Department of Homeland Security Doctoral Fellow.**
-  **Howard Hughes Award**, American Helicopter Society.

References

Available on Request