

Aravind Rajeswaran

Research Scientist,
Fundamental AI Research (FAIR, Meta AI)
Menlo Park, California, USA

rajeswaran.aravind@gmail.com
[aravindr93.github.io](https://github.com/aravindr93)
[Google Scholar](#)
Phone: +1 (206)-393-8924

Education

- University of Washington Seattle** *Sep 2016 – June 2021*
Ph.D. in Computer Science & Engineering
Advisers: Profs. [Sham Kakade](#) and [Emo Todorov](#)
- Indian Institute of Technology Madras, BTech (Hons.)** *Aug 2011 – July 2015*
Advisers: Profs. [Balaraman Ravindran](#) and [Shankar Narasimhan](#)

Employment

- **FAIR Labs, Meta AI** (formerly Facebook) - Research Scientist
Mentors : [Dhruv Batra](#) (current) and [Abhinav Gupta](#) (previous)
Foundation models and large-scale pre-training for Embodied AI. *April 2021 - present*
- **Google Brain** - Research Internship
Mentors : [Vincent Vanhoucke](#) and [Vikash Kumar](#)
Algorithmic foundations for offline and model-based reinforcement learning. *June 2019 - May 2020*
- **OpenAI** - Research Internship
Mentors : [John Schulman](#)
Reinforcement learning for dexterous robot hand manipulation. *June 2017 - Sep 2017*

Academic awards

- Best paper award at the ICRA 2022 Scaling Robot Learning Workshop *2022*
- Best paper award finalist at the RSS 2022 Scaling Robot Learning Workshop *2022*
- J. P. Morgan PhD Fellowship in AI *2020*
- Facebook PhD fellowship finalist in ML *2020*
- Best paper award at IEEE SIMPAR *2018*
- University of Washington PhD fellowship *2016*
- Bhagyalakshmi and Krishna Ayengar award for best undergraduate thesis. *2015*

Publications

- [1] *MoDem: Accelerating Visual Model-Based Reinforcement Learning with Demonstrations*
N. Hansen, Y. Lin, H. Su, X. Wang, V. Kumar, A. Rajeswaran
International Conference on Learning Representations (**ICLR**) 2023.
- [2] *Real World Offline Reinforcement Learning with Realistic Data Source*
G. Zhou, L. Ke, S. Srinivasa, A. Gupta, A. Rajeswaran, V. Kumar
International Conference on Robotics and Automation (**ICRA**) 2023.
- [3] *R3M: A Universal Visual Representation for Robot Manipulation* (* equal contributions)
S. Nair, A. Rajeswaran, V. Kumar, C. Finn, A. Gupta
ICRA 2022 Scaling Robot Learning Workshop (**Best Paper Award**)
Conference on Robot Learning (**CoRL**), 2022.
- [4] *The (Un)Surprising Effectiveness of Pre-Trained Vision Models for Control*
A. Rajeswaran*, S. Parisi*, S. Purushwalkam, A. Gupta
International Conference on Machine Learning (**ICML**), 2022. (**Long Oral**)

- [5] *CIC: Contrastive Intrinsic Control for Unsupervised Skill Discovery*
M. Laskin, H. Liu, X.B. Peng, D. Yarats, A. Rajeswaran, P. Abbeel
Advances in Neural Information Processing Systems (**NeurIPS**) 2022.
- [6] *Can Foundation Models Perform Zero-Shot Task Specification For Robot Manipulation?*
Y. Cui, S. Niekum, A. Gupta, V. Kumar, A. Rajeswaran
RSS 2022 Scaling Robot Learning Workshop. (**Best Paper Award Finalist**)
Learning for Dynamics and Control (**L4DC**), 2022.
- [7] *Decision Transformer: Reinforcement Learning via Sequence Modeling*
L. Chen*, K. Lu*, A. Rajeswaran, K. Lee, A. Grover, M. Laskin,
P. Abbeel, A. Srinivas, I. Mordatch
Advances in Neural Information Processing Systems (**NeurIPS**), 2021.
- [8] *Visual Adversarial Imitation Learning using Variational Models*
R. Rafailov, T. Yu, A. Rajeswaran, C. Finn
Advances in Neural Information Processing Systems (**NeurIPS**), 2021.
- [9] *COMBO: Conservative Offline Model-Based Policy Optimization*
T. Yu*, A. Kumar*, R. Rafailov, A. Rajeswaran, S. Levine, C. Finn
Advances in Neural Information Processing Systems (**NeurIPS**), 2021.
- [10] *Reinforcement Learning with Latent Flow*
W. Shang*, X. Wang*, A. Srinivas, A. Rajeswaran, Y. Gao, P. Abbeel, M. Laskin
Advances in Neural Information Processing Systems (**NeurIPS**), 2021.
- [11] *Behavioral Priors & Dynamics Models: Improving Performance and Domain Transfer in Offline RL*
C. Cang, A. Rajeswaran, P. Abbeel, M. Laskin
Pre-print 2021, *arXiv: 2106.09119*
- [12] *Offline Reinforcement Learning from Images with Latent Space Models*
R. Rafailov*, T. Yu*, A. Rajeswaran, C. Finn
Learning for Dynamics and Control (**L4DC**), 2021.
- [13] *MOReL: Model-Based Offline Reinforcement Learning*
R. Kidambi*, A. Rajeswaran*, P. Netrapalli, T. Joachims
Advances in Neural Information Processing Systems (**NeurIPS**), 2020.
- [14] *A Game Theoretic Framework for Model Based Reinforcement Learning*
A. Rajeswaran, I. Mordatch, V. Kumar
International Conference on Machine Learning (**ICML**), 2020.
- [15] *Lyceum: An efficient and scalable ecosystem for robot learning.*
C. Summers, K. Lowrey, A. Rajeswaran, S. Srinivasa, E. Todorov
Learning for Dynamics and Control (**L4DC**), 2020.
- [16] *Meta-Learning with Implicit Gradients.*
A. Rajeswaran*, C. Finn*, S. Kakade, S. Levine
Advances in Neural Information Processing Systems (**NeurIPS**), 2019.
- [17] *Online Meta-Learning.*
C. Finn*, A. Rajeswaran*, S. Kakade, S. Levine
International Conference on Machine Learning (**ICML**), 2019.
- [18] *Plan Online, Learn Offline: Efficient Learning and Exploration via Model-Based Control.*
K. Lowrey*, A. Rajeswaran*, S. Kakade, E. Todorov, I. Mordatch
International Conference on Learning Representations (**ICLR**), 2019.
- [19] *Dexterous Manipulation with Deep Reinforcement Learning: Efficient, General, and Low Cost.*
H. Zhu, A. Gupta, A. Rajeswaran, S. Levine, V. Kumar
International Conference on Robotics and Automation (**ICRA**), 2019.

- [20] *Reinforcement learning for non-prehensile manipulation: Transfer from simulation to physical system.*
K. Lowrey, S. Kolev, J. Dao, A. Rajeswaran, E. Todorov,
IEEE SIMPAR, 2018 (**Best Paper Award**)
- [21] *Variance Reduction for Policy Gradient Using Action-Dependent Factorized Baselines.*
C. Wu, A. Rajeswaran, Y. Duan, V. Kumar, A. Bayen, S. Kakade, I. Mordatch, P. Abbeel
International Conference on Learning Representations (ICLR), 2018. (**Full Oral**)
- [22] *Divide-and-Conquer Reinforcement Learning.*
D. Ghosh, A. Singh, A. Rajeswaran, V. Kumar, S. Levine
International Conference on Learning Representations (ICLR), 2018.
- [23] *Learning complex dexterous manipulation with deep reinforcement learning and demonstrations.*
A. Rajeswaran*, V. Kumar*, A. Gupta, G. Vezzani, J. Schulman, E. Todorov, S. Levine
Proceedings of Robotics: Science and Systems (RSS), 2018.
- [24] *Towards generalization and simplicity in continuous control.*
A. Rajeswaran, K. Lowrey, E. Todorov, S. Kakade
Advances in Neural Information Processing Systems (NIPS), 2017.
- [25] *EPOpt: Learning robust neural network policies using model ensembles.*
A. Rajeswaran, S. Ghotra, B. Ravindran, S. Levine
International Conference on Learning Representations (ICLR), 2017.
- [26] *Identifying Topology of Power Distribution Networks Based on Smart Meter Data.*
S. Jayadev, N. Bhatt, R. Pasumarthy, A. Rajeswaran
IEEE Transactions on Smart Grid, 2017.
- [27] *A Graph Partitioning Approach for Leak Detection in Water Distribution Networks.*
A. Rajeswaran, S. Narasimhan, S. Narasimhan
Computers & Chemical Engineering, 2017.

Mentoring

Interns & Residents

- Philipp Wu (PhD at UC Berkeley)
- Shikhar Bahl (PhD at CMU)
- Nicklas Hansen (PhD at UCSD)
- Mandi Zhao (PhD at Columbia)
- Suraj Nair (PhD at Stanford)
- Allan Zhou (PhD at Stanford)
- Liyiming Ke (PhD at UW Seattle)
- Yuchen Cui (PhD at UT Austin)

University Students

- Aryan Jain (UC Berkeley BS/MS)
- Ethao Guo (UC Berkeley BS/MS)
- Rafael Rafailov (Stanford MS → Stanford PhD)
- Kevin Lu (UC Berkeley BS → Stanford PhD)
- Catherine Cang (UC Berkeley BS → Plaid)
- Ben Evans (UW BS/MS → NYU PhD)
- Divye Jain (UW BS/MS → Google SWE)
- Sarvjeet Ghotra (IIT-M → MILA PhD)

Invited Talks

- Model-Based Offline Reinforcement Learning. IISc DeepRL workshop (Golden Jubilee event). *2021*
- Model-Based Offline Reinforcement Learning. TWIML podcast. *2020*
- Recent advances in model-based RL. CILVR Lab, NYU. *2020*
- Data-driven models for efficient Reinforcement Learning. MIT. *2020*
- Data-driven models for efficient Reinforcement Learning. Google Brain. *2020*
- Data-driven models for efficient Reinforcement Learning. DeepMind. *2020*
- Data-driven models for efficient Reinforcement Learning. Microsoft Research. *2020*

- Data-driven models for efficient Reinforcement Learning. Facebook AI Research. *2020*
- Data-driven models for efficient Reinforcement Learning. UC Berkeley. *2020*
- Data-driven models for efficient Reinforcement Learning. SAIL Lab, Stanford University. *2020*
- POLO: A new framework for model-based control and learning. [Informs](#) annual meeting. *2019*
- Towards embodied artificial intelligence. CMU and FAIR Pittsburgh. *2019*
- Accelerating robot learning. UW CSE affiliates day. *2018*
- Towards generalization and simplicity in continuous control. OpenAI. *2017*

Professional Service and Teaching

Course Instructor and TA

- Fully designed and **taught** a special topics course at UW on deep RL for robotics. [[course website](#)]
- Teaching assistant for advanced graduate level machine learning courses at UW.

Workshops Organized

- Pretraining for Robot Learning ([website](#)), CoRL 2022.
- 3rd Offline RL workshop: Offline RL as a “Launchpad” ([website](#)), NeurIPS 2022.
- Object Representations for Learning and Reasoning ([website](#)), NeurIPS 2020.
- Generative Modeling and Model-Based Reasoning for Robotics and AI ([website](#)), ICML 2019.

Reviewing and Program Committee

- NeurIPS (2018, 2019, 2020, 2021, 2022)
- ICML (2018, 2019, 2020, 2021, 2023)
- ICLR (2019, 2020, 2021)
- CoRL (2019, 2020, 2021)

References

- [Dr. Sham Kakade](#), Professor (CSE & Statistics), Harvard University.
- [Dr. Emo Todorov](#), Affiliate Professor, University of Washington. Lead developer of [MuJoCo](#).
- [Dr. Pieter Abbeel](#), Professor (EECS), UC Berkeley. Co-Founder and Chief Scientist, [Covariant](#).
- [Dr. Dhruv Batra](#), Research Director, Meta AI. Associate Professor, Georgia Tech.
- [Dr. Abhinav Gupta](#), Associate Professor, CMU.
- [Dr. Sergey Levine](#), Assistant Professor (EECS), UC Berkeley.