Mohak Bhardwaj

mohak.bhardwaj@gmail.com mohakbhardwaj.github.io

Education

2019–2024 **University of Washington, School of Computer Science** Seattle, WA
Ph.D. in Computer Science
Advisor: Dr. Byron Boots

2018–2019 **Georgia Institute of Technology, College of Computing** Atlanta, GA
Ph.D. student in Robotics
Advisor: Dr. Byron Boots
Transferred to University of Washington in September 2019

2015–2016 **Carnegie Mellon University, School of Computer Science** Pittsburgh, PA
MS Robotic Systems Development
Advisors: Dr. Sebastian Scherer, Dr. John Dolan

2011–2015 **Indian Institute of Technology (BHU), Varanasi** India
B.Tech Mechanical Engineering

Research/Work Experience

Apr 2024–Present **Boston Dynamics**, *Staff Research Scientist* Waltham, MA
- Developing visuomotor policies for dexterous manipulation and agile whole-body control on the ATLAS robot using reinforcement and imitation learning.
- Led the technical development of a collaborative project with NVIDIA, culminating in a high-impact GTC demo and public video release.
- **Videos:** tinyurl.com/2775e283, tinyurl.com/462x6cpn

Sep 2019-Mar 2024 **University of Washington**, *Graduate Research Assistant* Seattle, WA
Advisor: Dr. Byron Boots
- Offline Reinforcement Learning: Developed an adversarial model-based offline RL framework with strong theoretical guarantees on robust performance improvement.
- Combining Model-based and Model-free Reinforcement Learning: Investigated hybrid methods combining model-predictive control and value function approximation for sample-efficient robot learning with provable reduction in model bias.
- Model-Predictive Control: Built a GPU-accelerated model-predictive control system for fast, reactive and collision free motion generation for robot manipulators using learned cost functions.
- Dynamic Non-prehensile Manipulation: Combined offline value function learning from demonstrations with MPC for real world non-prehensile object transport.

Jun 2022-Dec 2022 **Google DeepMind**, *Research Scientist Intern* London, UK
Mentors: Dr. Jonas Buchli, Dr. Markus Wulfmeier, Dr. Martin Riedmiller
- Developed a model-free RL framework for dynamic, airflow-based control of rigid bodies and deployed it on a real-world fluid dynamic control testbed.
- Explored offline RL to multi-task data re-use and efficient reward design.

Sep 2020-Dec 2020 **NVIDIA Research**, *Robotics Research Intern* Seattle, WA
Mentors: Dr. Dieter Fox, Dr. Byron Boots, Dr. Fabio Ramos, Dr. Balakumar Sundaralingam
- Led the development of STORM, a GPU-accelerated MPC framework enabling fast, reactive and collision free motion generation for robot manipulators with learned cost functions.
- Demonstrated robust performance on dynamic ball balancing, handling task constraints and obstacle avoidance with full joint-space control on a Franka Panda arm. **Link:** bit.ly/3y73HbW

May 2019-Oct 2019 **NVIDIA Research**, *Robotics Research Intern* Seattle, WA
Mentors: Dr. Ankur Handa, Dr. Dieter Fox, Dr. Byron Boots
Developed an algorithmic framework combining information theoretic MPC and entropy-regularized RL using learned *soft* Q-functions to mitigate short-horizon bias in path integral control.

Sep 2018-Aug 2019 **Georgia Institute of Technology**, *Graduate Research Assistant* Atlanta, GA
Advisor: Dr. Byron Boots
- Differentiable trajectory optimization: Developed a structured learning framework for learning factor graph parameters by representing Gaussian Process Motion Planning as a differentiable computation graph.
- Leveraging Experience in Lazy Search: Formulated lazy search as an MDP and developed an approach to learn edge evaluation policies via imitatiion of oracule selectors, acceleratin planning in complex, high-dimensional problems.

Dec 2017-July 2018 **Near Earth Autonomy**, *Robotics Engineer* Pittsburgh, PA
Designed an adaptive, sampling-based motion planning system for emergency landings under uncertainty for real-world UAVs.

Dec 2016-Nov 2017 **Carnegie Mellon University**, *Research Assistant* Pittsburgh, PA
Mentors: Dr. Sebastian Scherer, Dr. Sanjiban Choudhury
Formulated heuristic search as sequential decision-making and developed a self-supervised imitation learning framework for learning search policies. Demonstrated intelligent search behaviors in complex environments with up to 70× speedup over A* for real-world UAV motion planning.

May 2014-July 2014 **Indian Institute of Information Technology, Hyderabad**, *Research Intern* India
Mentor: Dr. Suril V. Shah
Developed an inverse kinematics based optimal control algorithm for visual servoing of dual-arm space manipulators with real-time singularity avoidance in a coupled arm-base dynamical system.

Publications

Journal Publications

[2] **Bhardwaj M.**, Choudhury S., Boots B., Srinivasa S., "Leveraging Experience in Lazy Search", Autonomous Robots (AuRo), 2021 **Link**: bit.ly/3Gpr6sO

[1] Choudhury S., **Bhardwaj M.**, Arora S., Kapoor A., Ranade G., Scherer S., Dey D., "Data-driven Planning via Imitation Learning", International Journal of Robotics Research (IJRR), 2018 **Link**: goo.gl/sgG7LJ **(Paper of the Year Finalist)**

Peer Reviewed Conference Publications

[12] Jawale N., Boots B., Sundraralingam B., **Bhardwaj M.***, "Dynamic Non-Prehensile Object Transport via Model-Predictive Reinforcement Learning (ICRA), 2025 **Link**: https://tinyurl.com/mu58xwve

[11] Fishman A., Walsman A., **Bhardwaj M.***, Yuan W., Sundaralingam B., Boots B., Fox D., "Avoid Everything: Model-Free Collision Avoidance with Expert-Guided Fine-Tuning", Conference on Robot Learning (CORL), 2024 **Link**: shorturl.at/W3vtF

[10] **Bhardwaj M.***, Lampe T., Neunert M, Romano F., Abdolmaleki A., Byravan A., Wulfmeier M., Riedmiller M., Buchli J., "Real-World Fluid Directed Rigid Body Control via Deep Reinforcement Learning ", Learning for Dynamics & Control Conference (L4DC), 2024 **Link**: shorturl.at/CuaJf

[9] Nakamura S., Higuera Arias C., **Bhardwaj M.***, Boots B., "Robotic System Performing Dynamic Interaction in Human-Robot Cooperative Work for Assembly Operation ", IEEE/SICE International Symposium on System Integration (SII), 2024 **Link**: shorturl.at/Zf7si

[8] **Bhardwaj M.***, Xie T *, Boots B., Jiang N., Cheng C., "Adversarial Model for Offline Reinforcement Learning", Conference on Neural Information Processing Systems (NeuRIPS), 2023 **Link**: shorturl.at/fP128

[7] **Bhardwaj M.**, Sundaralingam B., Mousavian A., Ratliff N., Fox D., Ramos F., Boots B., "STORM: An Integrated Framework for Fast Joint-Space Model-Predictive Control for Reactive Manipulation", Conference on Robot Learning (CoRL), 2021 **Link**: bit.ly/3ePBWNK **(Among top 6% selected for oral presentation)**

[6] **Bhardwaj M.**, Choudhury S., Boots B., "Blending MPC & Value Function Approximation for Efficient Reinforcement Learning", International Conference on Learning Representations (ICLR), 2021 **Link**: bit.ly/3i9VxtN

[5] **Bhardwaj M.**, Handa A., Fox D., Boots B., "Information Theoretic Model Predictive Q-Learning", Learning for Dynamics and Control (L4DC), 2020 **Link**: bit.ly/2TyrhPT

[4] **Bhardwaj, M.**, Boots B., Mukadam M., "Differentiable Gaussian Process Motion Planning", International Conference on Robotics and Automation (ICRA), 2020 **Link**: bit.ly/3x2AcXu

[3] **Bhardwaj, M.**, Choudhury S., Boots B., Srinivasa S., "Leveraging Experience in Lazy Search", Robotics: Science and Systems (RSS), 2019 **Link**: bit.ly/2T13MKt

[2] **Bhardwaj M.**, Choudhury S., Scherer S., "Learning Heuristic Search via Imitation", Conference on Robotic Learning (CoRL), 2017 **Link**: goo.gl/cPo2yQ

[1] Mithun, P., Anurag, V. V., **Bhardwaj, M.**, Shah, S. V., "Real-Time Dynamic Singularity Avoidance while Visual Servoing of a Dual-Arm Space Robot", Advances in Robotics (AIR), 2015 **Link**: goo.gl/j1uVLg

Workshop Publications

[2] Xie, T*, **Bhardwaj M.**, Jian N., Cheng C., Boots B., "ARMOR: A Model-based Framework for Improving Arbitrary Baseline Policies with Offline Data.", Workshop on Offline Reinforcement Learning, NeurIPS, 2022 **Link**: tinyurl.com/mtkvxpxy

[1] **Bhardwaj M.**, Handa A., Fox D., Boots B., "Information Theoretic Model Predictive Q-Learning", Workshop on Machine Learning for Planning and Control, ICRA, 2020 **Link**: bit.ly/3x7C95c

Professional Activities

Reviewer
- **Journals:** IEEE Robotics and Automation Letters (RA-L), IEEE Transactions on Robotics (T-RO)).
- **Conferences**: IEEE International Conference on Robotics and Automation (ICRA), Conference on Robot Learning (CoRL), IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Learning for Dynamics and Control (L4DC).

Volunteer Reader for PhD applications at Paul G. Allen School of Computer Science & Engineering, University of Washington

Invited Talks and Posters

Improving Model-Predictive Control with Reinforcement Learning
- University of Washington Robotics Colloquium, Nov 2020.
- ICRA Workshop on Machine Learning for Planning & Control, 2020.

Adversarial Model for Offline Reinforcement Learning
- NeuRIPS Workshop on Offline Reinforcement Learning, 2022.

Teaching

Fall 2023 **Graduate Teaching Assistant, University of Washington**, *CSE-446: Machine Learning (Instructor: Prof. Kevin Jamieson)*

Spring 2020 **Graduate Teaching Assistant, University of Washington**, *CSE-599W: Reinforcement Learning (Instructor: Prof. Byron Boots)*

Honors

2025 **Peak Power Award, Boston Dynamics**
For leading technical effort on vision-based dexterous manipulation for the eAtlas GTC video.

2018 **Finalist, IJRR Paper of the Year**
"Data Driven Planning via Imitation Learning."

2015 **Institute Color Award, IIT Varanasi**
Awarded for outstanding extra-curricular achievements.

Open-Source Code

S.T.O.R.M: A GPU accelerated MPC toolkit for robot control. *[bit.ly/3y73HbW]*

Search as Imitation Learning: Tensorflow pipeline for learning heuristic policies for motion planning. *[goo.gl/YXkQAC]*

Python Motion Planning: Lightweight library for motion planning research *[goo.gl/88shhJ]*

Technical Skills

Languages	C++, Python
Frameworks & Tools	Pytorch, JAX, MuJoCo, IsaacLab, ROS, OpenCV