Mohak Bhardwaj

✉ *mohakb@cs.washington.edu*
mohakbhardwaj.github.io

Research Statement

My research aims at enabling flexible robot decision-making systems that continually improve a robot's performance with experience. I approach this by combining classical planning & control algorithms with data-driven techniques. Specifically, my research topics include integrating model-predictive control with reinforcement learning, imitation learning for accelerated motion planning, differentiable trajectory optimization and offline reinforcement learning.

Education

2019– **University of Washington, School of Computer Science**,
Ph.D. Computer Science, *Advisor: Dr. Byron Boots*, *GPA:3.95/4.0*.

2018–2019 **Georgia Institute of Technology, College of Computing**,
Ph.D. student in Robotics, *Advisor: Dr. Byron Boots*, *GPA:4.0/4.0*.
Transferred to University of Washington in September 2019

2015–2016 **Carnegie Mellon University, School of Computer Science**,
Master of Science in Robotic Systems Development, *Advisors: Dr. Sebastian Scherer and Dr. John M. Dolan*, *QPA:3.83/4.0*.

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

Research Experience

Sep 2021-June 2022 **Learning Adversarial Models for Offline Reinforcement Learning**, *University of Washington*.

- Developed a framework for adversarial model-based offline RL that can improve over arbitrary reference policies by optimizing for worst case performance difference.
- Demonstrated theoretically and empirically that the framework exhibits robust policy improvement over the reference for any admissible hyperparamter regardless of data coverage.

Sep 2020-Aug2021 **STORM: An Integrated Framework for Fast Joint-Space Model-Predictive Control for Reactive Manipulation**, *NVIDIA/University of Washington*.

- Developed a GPU-accelerated model-predictive control (MPC) framework for reactive real-world manipulation that incorporates learned components and enables smooth, reactive motions at a 125Hz control frequency.
- Demonstrated efficacy on dynamic ball balancing, handling task constraints and obstacle avoidance with full joint-space control on a Franka Panda arm. **Link:** bit.ly/3y73HbW

Jan 2020-Sep 2020 **Accelerating Reinforcement Learning by Blending Model-Predictive Control and Value Function Approximation**, *University of Washington*.

- Proposed a framework for sample efficient reinforcement learning (RL) that leverages approximate models by blending MPC with a value function learned from real system data to mitigate the effects of model bias.
- Provided a theoretical analysis that shows how errors from inaccurate models in MPC and value function estimation in RL can be balanced, and demonstrated performance comparable to MPC with access to true dynamics on high-dimensional control tasks.

May 2019-Oct 2019 **Information Theoretic Model Predictive Q-Learning**, *NVIDIA/University of Washington*.
Developed a principled framework for combining information theoretic MPC and entropy regularized RL that shows how a learned *soft* Q-function can overcome short horizon bias in path integral control.

Sep 2018-Aug 2019 **Differentiable Continuous Time Trajectory Optimization**, *Georgia Institute of Technology*.
Developed a structured learning framework for learning factor graph parameters by representing Gaussian Process Motion Planning as a differentiable computation graph.

Sep 2018-Jan 2019 **Leveraging Experience in Lazy Search for Accelerated Motion Planning**, *Georgia Institute of Technology*.
Formulated lazy search as a Markov Decision Process and developed an approach for learning effective edge evaluation policies by imitating oracular selectors. Learned edge selectors accelerate planning in obstacle rich, high-dimensional problems compared to conventional hand-designed heuristics.

Dec 2016-July 2017 **Learning Heuristic Search via Imitation**, *Carnegie Mellon University*.
Proposed a formulation of heuristic search as sequential decision making and developed an algorithmic framework to learn heuristic policies via self-supervised imitation learning. The approach learns intelligent search policies on complex environment distributions and provides 70x speedup over classic A* search for real world motion planning on UAVs.

May 2014-July 2014 **Visual Servoing and Singularity Avoidance for Dual Arm Space Robot**, *IIIT-Hyderabad*.
Developed inverse kinematics based optimal control algorithms for visual servoing of space manipulators with real-time singularity avoidance in a coupled arm-base dynamic 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)**

Conference Publications

[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

Work Experience

Jun- Nov 2022 **DeepMind Inc., London, UK**, *Research Scientist Intern*, Manager: Martin Riedmiller, Mentors: Jonas Buchli, Arun Byranvan, Markus Wulfmeier, Abbas Abdolmaleki, Thomas Lampe.
Research on reinforcement learning for dynamic real-world control systems. (Project details under NDA)

Dec 2020- Jun 2022
Sep 2019-Sep 2020 **University of Washington, Robot Learning Lab**, *Research Assistant*, Advisor: Dr. Byron Boots.
Research on model-predictive control and reinforcement learning for dynamic control tasks.

Sep 2020-Dec 2020
May 2019-Aug 2019 **NVIDIA Seattle Robotics Lab**, *Research Intern*, Mentors: Dr. B. Sundaralingam, Dr. F. Ramos, Dr. A. Handa, Dr. B. Boots, Dr. D. Fox.
Research on GPU accelerated model-predictive control for dynamic manipulation and improving path-integral MPC with entropy-regularized reinforcement learning.

Aug 2018-May 2019 **Georgia Institute of Technology, Robot Learning Lab**, *Research Assistant*, Advisor: Dr. Byron Boots.
Imitation and self-supervised learning for search-based planning & trajectory optimization.

Dec 2017-July 2018 **Near Earth Autonomy**, *Robotics Engineer*.
Adaptive motion planning under uncertainty for real-world UAVs.

Mar 2017-Dec 2017 **Carnegie Mellon University, AirLab**, *Extern*, Advisor: Dr. Sebastian Scherer.
Reinforcement and imitation learning applied to search based planning; Planning under uncertainty.

May 2016-Aug 2016 **Qualcomm R&D**, *Intern, Autonomous Driving*, Manager: Sebastian Mounier.
S.L.A.M and multi-sensor calibration for autonomous cars.

May 2014-Aug 2014 **IIIT-Hyderabad, Robotics Research Institute**, *Research Intern*, Advisor: Dr. Suril V. Shah.
Research on optimal control algorithms for space manipulators.

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, November 2020.
- Workshop on Machine Learning for Planning & Control, ICRA 2020.

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

Teaching Experience

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

Honors

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

2015 **Institute Color Award from IIT,Varanasi**.
Outstanding extra-curricular achievements.

Open-Source Code

S.T.O.R.M: A GPU accelerated toolkit for model-predictive control for robots. Link: bit.ly/3y73HbW

Search as Imitation Learning: Tensorflow pipeline for learning heuristic policies for search based motion planning. Link: goo.gl/YXkQAC

Python Motion Planning: Easy-to-use motion planning library geared towards learning for planning research Link: goo.gl/88shhJ

Deep RL with OpenAI Gym: Modular pipeline for developing and testing RL agents with OpenAI gym environments. Link: goo.gl/8tkFC4

Technical Skills

Languages C++, Python, MATLAB

Software Pytorch, TensorFlow, OpenAI Gym, MuJoCo, ROS, OMPL, OpenCV