MURTAZA DALAL

mdalal@andrew.cmu.edu | mihdalal.github.io | linkedin.com/in/murtaza-dalal/

I am a Robotics PhD interested in building generalist agents that are capable of solving any task of interest. To that end, my research work focuses on unlocking generalization in robot learning.

Education

Carnegie Mellon University: Robotics PhD, Expected Graduation: December 2024 **GPA:** 3.95
Advisor: Ruslan Salakhutdinov
National Science Foundation Fellowship, Apple Funded Scholar

University of California Berkeley: B.S. Electrical Engineering and Computer Science, May 15, 2020 **GPA:** 3.972
Highest Honors in Electrical Engineering and Computer Sciences (top 3% of class)
UC Berkeley Regent's and Chancellor's Scholar (top 2% of class), Accel Scholars (Network led by Accel Partners), Edward F. Kraft Award (4.0 during freshman year), HKN (EECS Honor Society)

Publications

- **Dalal*, M.,** Yang*, J., Mendonca, R., Khaky, Y., Salakhutdinov, R., Pathak, D. (2024). Neural MP: A Policy Learning Framework for Generalist Neural Motion Planners. (In submission)
- **Dalal* M.,** Liu*, M., Talbott, W., Zhang, J., Pathak, D., Salakhutdinov, R. (2024). Local Policies enable Zero-shot Long-Horizon Manipulation. (In submission)
- **Dalal, M.**, Chiruvolu, T., Chaplot, D., Salakhutdinov, R. (2024). Plan-Seq-Learn: Language Model Guided RL for Solving Long Horizon Robotics Tasks. International Conference on Learning Representations (**ICLR 2024**)
- **Dalal, M.**, Mandlekar, A., Garrett, C., Handa, A., Salakhutdinov, R., Fox, D. (2023). Imitating Task and Motion Planning using Visuomotor Transformers. Conference on Robot Learning. **(CoRL 2023)**
- **Dalal, M.,** Pathak, D., Salakhutdinov, R. (2021). Accelerating Robotic Reinforcement Learning via Parameterized Action Primitives. Neural Information Processing Systems. (**NeurIPS 2021)**
- Chaplot, D, **Dalal, M.**, Gupta, S., Malik, J., Salakhutdinov, R. (2021). SEAL: Self-supervised Embodied Active Learning using Exploration and 3D Consistency. Neural Information Processing Systems. (**NeurIPS 2021)**
- Nair*, A., Gupta*, A., **Dalal, M.,** and Levine, S. (2020). Accelerating Online Reinforcement Learning with Offline Datasets.
- Singh, A., Jang, E., Irpan A., Kappler, D., **Dalal, M.**, Levine, S., Khansari, M., Finn, C. Scalable Multi-Task Imitation Learning with Autonomous Improvement. International Conference on Robotics and Automation (**ICRA 2020**).
- Pong*, V., **Dalal*, M.**, Lin*, S., Nair, A., Bahl, S., and Levine, S. (2020). Skew-Fit: State-Covering Self-Supervised Reinforcement Learning. International Conference on Machine Learning (**ICML 2020**).
- Nair*, A., Pong*, V., **Dalal, M.**, Bahl, S., Lin, S., and Levine, S. (2018). Visual Reinforcement Learning with Imagined Goals. Neural Information Processing Systems (**NeurIPS 2018 spotlight**).
- Pong*, V., Gu*, S., **Dalal, M.**, and Levine, S. (2017). Temporal Difference Models: Model-Free Deep RL for Model-Based Control. International Conference on Learning Representations (**ICLR 2018**).
- Haarnoja*, T., Pong, V., Zhou, A., **Dalal, M.**, Abbeel, P., and Levine, S. (2017). Composable Deep Reinforcement Learning for Robotic Manipulation. International Conference on Robotics and Automation (**ICRA 2018**).
- **Dalal*, M.**, Li*, A., Taori*, R. (2019). Auto-regressive Models: What are They Good For? NeurIPS 2019 Workshop on Information Theory and Machine Learning (**NeurIPS ITML 2019**).

Experience

Senior Robotics Engineer at Apple (Contract) *May 2024-Current*

- Leading research project on scaling AI models for robotics
- Built pipeline to train large-scale models for manipulation using reinforcement learning and imitation learning

Research Intern at NVIDIA Seattle Robotics Lab *May 2022 – May 2023*

- Led research project aiming to leverage task and motion planner supervision for imitation learning
- Developed an imitation system that can solve **over 300** long-horizon manipulation tasks involving up to **8** and **72** different objects, achieving success rates of **over 70%**
- Wrote imitation pipeline that sped up training process by more than **10x**

MURTAZA DALAL

mdalal@andrew.cmu.edu | mihdalal.github.io | linkedin.com/in/murtaza-dalal/

AI Resident at Google X, The Moonshot Factory *May 2019 – September 2019*

- Led research project at the intersection of meta-learning, reinforcement learning (RL) and robotics
- Trained RL agents to solve sparse-reward manipulation tasks sample efficiently by leveraging demonstrations
- Analyzed critical problem with using pre-trained actors in actor-critic RL algorithms

Undergraduate Robotics and AI Researcher under Professor Sergey Levine *June 2017 – May 2020*

- Achieved state-of-the-art performance on unsupervised vision-based robot reaching, pushing and grasping
- Trained robot to learn unsupervised vision-based door opening in 5 hours
- Published 4 papers to peer reviewed research conferences with under 30% acceptance rate
- Constructed a complete software package for using Rethink Sawyer robots with reinforcement learning algorithms and integrated the Microsoft Kinect to enable vision-based tasks, used by RAIL Lab at UC Berkeley
    - https://github.com/mihdalal/sawyer_control
- Designed a variety of robotic environments for testing RL algorithms in simulation
    - https://github.com/vitchyr/multiworld.git

CS 188 Teaching Assistant (Introduction to Artificial Intelligence) *January 2019 - May 2019*

- Taught weekly discussions and held office hours, tested and gave feedback for exams

CS 189 Tutor (Introduction to Machine Learning) *June 2018 - December 2018*

- Provided individualized tutoring services of up to 10 hours per week
- Helped students understand the linear algebra and probability as well as intuition behind ML methods

NASA Ames Research Center Embedded Systems Intern *June 2015 - August 2015*

- Designed a payload for a 3DR quad-copter to hover over a target and adjust for lateral positional differences
- Wrote object tracking software in Visual Studio utilizing C/C++ and OpenCV
- Operated in a team of eight interns to construct a cube satellite network using Arduino and bread-board sensors

Technical Skills:
Languages: Python (experienced), C++, C, Java, Go
Operating Systems: Linux (Ubuntu), Windows, MacOS
Experience: Reinforcement Learning, Deep Learning, Data Science, Robotic Operating System (ROS), Robotics, PyTorch, Tensorflow, OpenAI Gym, NumPy, Computer Vision (OpenCV), Slurm, Git, UAVs, MapReduce, MuJoCo, Habitat, Isaac Gym

Relevant Course Work (CMU):
Optimal Control and Reinforcement Learning (A), Learning for Action and Perception (A), Convex Optimization (A), Probabilistic Graphical Models (A), Kinematics Dynamics Control (A-), Computer Vision (A), Math for Robotics (A)

Relevant Course Work (UC Berkeley):
Operating Systems (A), Advanced Robotics (A), Computer Security (A), Deep Unsupervised Learning (A), Real Analysis (A), Signals and Systems (A), Deep Reinforcement Learning (A), Machine Learning (A), Sequential Decision Making (A), Machine Structures (A+), Probability and Random Processes (A), Convex Optimization (A+), Artificial Intelligence (A), Algorithms (A-), Circuits and Control (A-), Linear Algebra (A+), Discrete Math (A), Data Structures (A)

MURTAZA DALAL

mdalal@andrew.cmu.edu | mihdalal.github.io | linkedin.com/in/murtaza-dalal/

Projects

SC2AI, Leveraging Deep RL to Play StarCraft 2: *Machine Learning @ Berkeley* *August 2018 - November 2018*

- As Project Manager, led a team to implement improvements on Google DeepMind's baselines for StarCraft 2
- Developed project scope and direction, designed software interface, and implemented base RL algorithm

AutoQuad, Reinforcement Learning for UAVs: Machine Learning @ Berkeley *February 2018 - May 2018*

- Implemented behavioral cloning to solve UAV location reaching tasks using state information to within .01m
- Implemented and adapted the DQN algorithm to solve UAV location reaching tasks using vision to within .15m

Delighting Real World Images to Repurpose for Video Games: Machine Learning @ Berkeley *February 2017 - May 2017*

- Constructed a machine learning pipeline that used convolutional neural networks along with state-of-the-art image processing algorithms to break down images into delighted meshes

UAV Midair Battery Swap: UAVs @ Berkeley *September 2016 - May 2017*

- Used OpenCV to design a highly accurate distance approximation algorithm for a single camera
- Developed Arduino control system to facilitate a safe and secure battery swap atop the hex-copter

$200 Prosthetic Elbow: Range of Motion Project Robotics Competition 1st Place Winner *February 2017 - May 2017*

- Wrote software to utilize neural input to accurately apply motor torques and constructed software interface

Automated Snack Delivery Robot: Dorm Ex Machina Robotics Competition 3rd place Winner *September 2016 - November 2016*

- Designed software for handling user input, motor control, line sensing, and computer vision