Brian Yang

brianhyang@cmu.edu
byang.org

Research Interests

My research focuses on developing learning-based control algorithms at the intersection of generative modeling, model-based planning, and reinforcement learning. In my PhD, I have primarily focused on applying these ideas to autonomous driving. During my undergrad, I was fortunate enough to have the opportunity to work on a variety of real-world robot learning problems ranging from low-cost manipulation to microrobot locomotion.

Education

Carnegie Mellon University, PhD in Robotics 2020 – Present
Advised by Katerina Fragkiadaki and Jeff Schneider

University of California, Berkeley, BA in Computer Science 2016 – 2020

Research Experience

Carnegie Mellon University 2020 – Present
PI: Katerina Fragkiadaki and Jeff Schneider

- Designing scalable, controllable agents for autonomous driving using generative models and model-based RL.
- Proposed *Diffusion-ES*, which uses gradient-free search to guide diffusion sampling with black-box rewards.
    - This achieves state-of-the-art performance in nuPlan, beating all other learning-based methods.
    - Can use few-shot LLM prompting to enable instruction-following with zero paired training data.
- Helped close the loop between prediction and planning on AV systems by proposing *P2DBM*, which can retrofit pretrained motion forecasting models to use them downstream for fully reactive planning in CARLA.
- Used offline counterfactual reactive simulation to train cost-aware driving policies from offline logs.

Meta AI 2019
PI: Roberto Calandra

- Developed a system for dexterous in-hand manipulation using model-based RL and tactile sensors (*DIGIT*).
- Trained video prediction models over tactile feature maps and used them for MPC on a real hardware platform.

University of California, Berkeley 2017 – 2020
PI: Sergey Levine and Dinesh Jayaraman

- Researched morphology-agnostic visual robotic control using self-recognition and visual servoing (*MAVRIC*).
- Engineered an open-source low-cost robotics platform (*REPLAB*) for benchmarking robot learning algorithms.
- Assisted in building a fully autonomous real-world mobile navigation and grasping system (*ReLMM*).

PI: Kristofer Pister

- Used Bayesian optimization to perform millimeter-scale microrobot locomotion from a handful of trials.
- Created a data-efficient algorithm for jointly optimizing microrobot morphologies and gaits (*HPC-BBO*).

Publications

A. Villaflor, **B. Yang**, K. Fragkiadaki, J. Dolan, J. Schneider. “Learning Driving Policies with Offline Counterfactual Reactive Simulation.” *In submission*.

B. Yang, H. Su, N. Gkanatsios, T. Ke, A. Jain, J. Schneider, K. Fragkiadaki. “Diffusion-ES: Gradient-free Planning with Diffusion for Autonomous Driving and Zero-Shot Instruction Following.” In *CVPR*, 2024.

A. Villaflor, **B. Yang**, H. Su, K. Fragkiadaki, J. Dolan, J. Schneider. “Tractable Joint Prediction and Planning Over Discrete Behavior Modes for Urban Driving.” In *ICRA*, 2024.

C. Sun, J. Orbik, C. Devin, **B. Yang**, A. Gupta, G. Berseth, S. Levine. “Fully autonomous real-world reinforcement learning with applications to mobile manipulation.” In *CoRL*, 2022.

B. Yang, D. Jayaraman, G. Berseth, A. Efros, S. Levine. “Morphology-Agnostic Visual Robotic Control.” In *ICRA* and *RA-L*, 2020.

M. Lambeta, P. Chou, S. Tian, **B. Yang**, B. Maloon, V. Most, D. Stroud, R. Santos, A. Byagowi, G. Kammerer, D. Jayaraman, R. Calandra. “Digit: A novel design for a low-cost compact high-resolution tactile sensor with application to in-hand manipulation.” In *ICRA* and *RA-L*, 2020.

B. Yang, J. Zhang, V. Pong, S. Levine, D. Jayaraman. “REPLAB: A reproducible low-cost arm benchmark platform for robotic learning.” In *ICRA,* 2019.

T. Liao, G. Wang, **B. Yang**, R. Lee, K. Pister, S. Levine, R. Calandra. “Data-efficient learning of morphology and controller for a microrobot.” In *ICRA*, 2019.

B. Yang, G. Wang, R. Calandra, D. Contreras, S. Levine, K. Pister. “Learning flexible and reusable locomotion primitives for a microrobot.” In *ICRA* and *RA-L,* 2018.

Teaching

- Spring 2024, 16-811: Introduction to Robot Learning. Teaching assistant.
- Fall 2022, 10-703: Deep Reinforcement Learning. Teaching assistant.
- Spring 2017, CS 61A: Structure and Interpretation of Computer Programs. Academic intern.