

Education	Massachusetts Institute of Technology Ph.D. in Electrical Engineering and Computer Science MIT Presidential fellowship	September 2019 - Present GPA: 5.00/5.00
	Massachusetts Institute of Technology S.M. in Electrical Engineering and Computer Science MIT Presidential fellowship	September 2017 - September 2019 GPA: 5.00/5.00
	University of California, Berkeley B.S. in Electrical Engineering and Computer Science Graduated with Highest Honors (top 3%)	August 2013 - May 2017 GPA: 3.95/4.00
Experience	Improbable AI Lab (Advisor: Pulkit Agarwal)	September 2019 - Present
	<ul style="list-style-type: none">Flexibly composing foundation models, each pretrained on different modality of internet data (such as vision-language model, video diffusion and action model), for decision-making in embodied agentsBuilding deep RL algorithms that leverages past offline interactions for faster online learning	
	Meta FAIR Research Internship (Advisor: Aravind Rajeswaran)	June 2023 - Present
	<ul style="list-style-type: none">Curated an evaluation dataset for Embodied Question Answering consisting of high-quality human generated questions representative of real-world use cases for Embodied AI agentsThe evaluation dataset is accompanied by strong baselines demonstrating the capabilities of publicly available VLMs and fine-tuned multi-frame VLMs, and highlighting the gap between state-of-the-art VLMs (GPT4V) and humans.	
	Google Brain Research Internship (Advisors: Ofir Nachum and Sergey Levine)	June 2020 - August 2020
<ul style="list-style-type: none">Worked on leveraging (unlabelled) offline data to extract skills which allows for improved offline reinforcement learning, online reinforcement learning and few-shot imitation learning.		
Learning and Intelligent Systems Group (Advisors: Leslie Kaelbling and Josh Tenenbaum)	September 2017 - September 2019	
<ul style="list-style-type: none">Worked on leveraging physics engine for dynamics model learning and controlWorked on curiosity driven exploration for reinforcement learning in environment with sparse rewards		
Madry Lab (with Aleksander Madry)	March 2019 - September 2019	
<ul style="list-style-type: none">Worked on robust style transfer for images and videos		
Berkeley Artificial Intelligence Research Lab (Advisors: Pieter Abbeel and Sergey Levine)	October 2014 - May 2017	
<ul style="list-style-type: none">Worked on deep learning for state estimationDeveloped model-based sample efficient reinforcement learning algorithm		
Publications	OpenEQA: Embodied Question Answering in the Era of Foundation Models Arjun Majumdar*, Anurag Ajay* , Xiaohan Zhang*, Pranav Putta, Sriram Yenamandra, Mikael Henaff, Sneha Silwal, Paul Mcvay, Oleksandr Maksymets, Sergio Arnaud, Karmesh Yadav, Qiyang Li, Ben Newman, Mohit Sharma, Vincent Berges, Shiqi Zhang, Pulkit Agrawal, Yonatan Bisk, Dhruv Batra, Mrinal Kalakrishnan, Franziska Meier, Chris Paxton, Sasha Sax, Aravind Rajeswaran. Preprint, 2023.	
	Compositional Foundation Models for Hierarchical Planning Anurag Ajay* , Seungwook Han*, Yilun Du*, Shuang Li, Abhi Gupta, Tommi Jaakkola, Joshua B. Tenenbaum, Leslie Kaelbling, Akash Srivastava, Pulkit Agrawal. Neural Information Processing Systems (NeurIPS), New Orleans, USA, December 2023.	

Statistical Learning under Heterogenous Distribution Shift

Max Simchowitz*, **Anurag Ajay***, Pulkit Agrawal, Akshay Krishnamurthy.
International Conference on Machine Learning (ICML), Hawaii, USA, July 2023.

Parallel-Q-Learning: Scaling Off-policy Reinforcement Learning

Zechu Li, Tao Chen, Zhang-Wei Hong, **Anurag Ajay**, Pulkit Agrawal.
International Conference on Machine Learning (ICML), Hawaii, USA, July 2023.

Is Conditional Generative Modeling all you need for Decision Making?

Anurag Ajay*, Yilun Du*, Abhi Gupta*, Josh Tenenbaum, Tommi Jaakkola, Pulkit Agrawal.
International Conference on Learning Representations (ICLR), Kigali, Rwanda, May 2023.

Oral Talk

Distributionally Adaptive Meta Reinforcement Learning

Anurag Ajay*, Abhishek Gupta*, Dibya Ghosh, Sergey Levine, Pulkit Agrawal.
Neural Information Processing Systems (NeurIPS), New Orleans, USA, December 2022.

Offline RL Policies should be trained to be Adaptive

Dibya Ghosh, **Anurag Ajay**, Pulkit Agrawal, Sergey Levine.
International Conference on Machine Learning (ICML), Baltimore, USA, July 2022.

Oral Talk

Overcoming The Spectral Bias of Neural Value Approximation

Ge Yang*, **Anurag Ajay***, Pulkit Agrawal.
International Conference on Learning Representations (ICLR), Virtual, May 2022.

Understanding the Generalization Gap in Visual Reinforcement Learning

Anurag Ajay*, Ge Yang*, Ofir Nachum, Pulkit Agrawal.
ICML Reinforcement Learning for Real Life Workshop, July 2021.

OPAL: Offline Primitive Discovery for Accelerating Offline Reinforcement Learning

Anurag Ajay, Aviral Kumar, Pulkit Agrawal, Sergey Levine, Ofir Nachum.
International Conference on Learning Representations (ICLR), Vienna, Austria, May 2021.

Learning Action Priors for Visuomotor transfer

Anurag Ajay, Pulkit Agrawal.
ICML Inductive biases, invariances and generalization in RL workshop, July 2020.

Long-Horizon Prediction and Uncertainty Propagation with Residual Point Contact Learners

Nima Fazeli, **Anurag Ajay**, Alberto Rodriguez.
IEEE International Conference on Robotics and Automation (ICRA), Paris, France, May 2020.

Learning Skill Hierarchies from Predicate Descriptions and Self-Supervision

Tom Silver*, Rohan Chitnis*, **Anurag Ajay**, Josh Tenenbaum, Leslie Pack Kaelbling.
AAAI Workshop on Generalization in Planning (GenPlan), 2020.

Learning to Navigate Endoscopic Capsule Robots

Mehmet Turan, Yasin Almalioglu, Hunter B Gilbert, Faisal Mahmood, Nicholas J Durr, Helder Araujo, Alp Eren Sar, **Anurag Ajay**, Metin Sitti.
Robotics and Automation Letters (RA-L), 2019.

Combining Physical Simulators and Object-Based Networks for Control

Anurag Ajay, Maria Bauza, Jiajun Wu, Nima Fazeli, Joshua B. Tenenbaum, Alberto Rodriguez, Leslie P. Kaelbling.
IEEE International Conference on Robotics and Automation (ICRA), Montreal, Canada, May 2019.

Augmenting Physical Simulators with Stochastic Neural Networks: Case Study of Planar Pushing and Bouncing

Anurag Ajay, Jiajun Wu, Nima Fazeli, Maria Bauza, Leslie P. Kaelbling, Joshua B. Tenenbaum, Alberto Rodriguez.
IEEE International Conference on Intelligent Robots and Systems (IROS), Madrid, Spain, October 2018.
Best Paper for Cognitive Robotics

Reset-Free Guided Policy Search: Efficient Deep Reinforcement Learning with Stochastic Initial States

Anurag Ajay*, William Montgomery*, Chelsea Finn, Pieter Abbeel, Sergey Levine. IEEE International Conference on Robotics and Automation (ICRA), Singapore, May 2017.

Backprop KF: Learning Discriminative Deterministic State Estimators

Tuomas Haarnoja, **Anurag Ajay**, Sergey Levine, Pieter Abbeel.

Neural Information Processing Systems (NIPS), Barcelona, Spain, December 2016.