SUNEEL BELKHALE

(408)705-3444 ◇ belkhale@stanford.edu ◇ suneel.belkhale.com
github.com/suneelbelkhale ◇ linkedin.com/in/suneel-b-032b1a101

EDUCATION

Stanford University *September 2020 - May 2025 (Expected)*
Computer Science, Ph.D. *GPA: 4.0/4.0*
Advisor: Prof. Dorsa Sadigh

University of California Berkeley *August 2019 - May 2020*
Electrical Engineering and Computer Science, M.S. *GPA: 3.95/4.0*
Selected Coursework: Adv. Robotics†, Convex Optimization†, Machine Learning Systems†, Deep Unsupervised Learning†, Nonlinear Control†

University of California Berkeley *August 2016 - August 2019*
Electrical Engineering and Computer Science, B.S. *GPA: 3.99/4.0*
Selected Coursework: Signals & Systems, Algorithms, Computer Architecture, Operating Systems, Adv. Analog Circuits, Control Theory, Artificial Intelligence, Machine Learning, Probability Theory, Computer Vision†, Deep Reinforcement Learning†.

RESEARCH EXPERIENCE

Stanford Intelligent and Interactive Autonomous Systems (ILIAD) *September 2020 - Present*
Researching practical, generalizable algorithms for real-world multi-task learning in robot manipulation. Two main lines of work: assistive feeding for people with mobility impairment, and general imitation learning methods for multi-task datasets (e.g., using play or other diverse human generated data). PI: Prof. Dorsa Sadigh.

RT-H: Robot Action Hierarchies using Language
Language provides a way to break down complex concepts into digestible pieces. Recent works in robot imitation learning use language-conditioned policies that predict actions given visual observations and the high-level task specified in language. These methods leverage the structure of natural language to share data between semantically similar tasks in multi-task datasets. However, as tasks become more semantically diverse, sharing data between tasks becomes harder, so learning to map high-level tasks to actions requires much more demonstration data. To bridge tasks and actions, we teach the robot the *language of actions*, describing low-level motions with more fine-grained phrases like "move arm forward". Predicting these language motions as an intermediate step between tasks and actions forces the policy to learn the shared structure of low-level motions across seemingly disparate tasks. Furthermore, a policy that is conditioned on language motions can easily be corrected during execution through human-specified language motions. This work was published at RSS 2024.

HYDRA: Hybrid Robot Actions for Imitation Learning
In imitation learning, policies often suffer from state distribution shift at test time, due to compounding errors in action prediction which lead to previously unseen states. In this work, we view the distribution shift problem through the lens of data "quality": a high quality dataset encourages the policy to stay in distribution at test time. We formalize quality as a set of properties, and guided by these properties, we introduce HYDRA, a method that leverages a hybrid action space with two levels of action abstractions: *sparse high-level waypoints* and *dense low-level actions*. HYDRA substantially outperforms baselines across challenging real world tasks like making coffee and toasting bread. This work was published at CoRL 2023.

PLATO: Learning Object Affordances from Play for Scalable Multi-Task Imitation Learning
In contrast to learning from demonstration data, human play data consists of unstructured, user-directed

† *graduate level course*

play with their environment. Play is cheaper to collect and obtains larger coverage of the state space for each task, and prior work showed that play results in better and more robust policy learning. Our work introduced an object-centric paradigm for learning from play which extracted object affordances and linked these to the robot behaviors that caused them, thus improving the policy performance across the board. This paper was published at CoRL 2022.

Assistive Feeding: Bite Transfer and Bite Acquisition
Millions of people in the US require caregivers to feed themselves, for example those with motor impairment. Robot-assisted feeding has the potential to improve the daily lives of these people, but represents a challenging manipulation and human modeling task. Feeding involves first acquiring deformable bites of various foods from a plate, and then transferring these bites to a user in a safe and comfortable manner. We have studied optimizing comfort for bite transfer, integrating haptics and vision for bite acquisition, and learning the bimanual task of scooping foods from a plate. These works have been published at ICRA 2022/23 and CoRL 2022.

Berkeley Robotics and Artificial Intelligence Lab (RAIL) *August 2017 - August 2020*
Developing sample efficient reinforcement learning, computer vision, and control methods for autonomous robotics. PI: Prof. Sergey Levine, Mentor: Gregory Kahn.

Online Adaptation to Changing Dynamics with Model-Based Reinforcement Learning
Robots deployed in the real world will need to adapt to changing conditions. We consider control for variable payloads that hang off a nano-quadcopter via a string. Changing the payload can break down existing control algorithms; our approach learns a dynamics model from data and learns a context variable to represent a range of dynamics. At test time we infer the context that best explains recent data using variational inference. With our solution, the quadcopter uses its experience to learn an effective control policy under different payloads and string lengths. In our final demo, we demonstrate a fully autonomous pick-up and drop-off sequence for a variable mass payload. This paper was published at ICRA 2021.

Generalization through Simulation: Integrating Simulated & Real Data in RL for Flight
Robotics experiments suffer from a lack of real world data because it is often very difficult to collect. We addressed this in the context of nano-quadcopters navigating indoor environments, where the fragility of the quadcopters made real world data collection very difficult. We discovered that the most sample efficient method learns visual diversity from a quadcopter simulator and learns physics from the real world. By mixing model-free learning in simulation and model-based learning in the real world, the quadcopter was able to navigate the full lengths of hallway environments with high consistency on just an hour of real data, while baselines failed at navigating even a few meters. Published in ICRA 2019.

WORK EXPERIENCE

NVIDIA — Santa Clara *Summer 2019*
Autonomous Vehicles Perception Intern

· Researching ML-based radar object detection and classification methods and long term tracking / filtering. Developed probabilistic filtering method for altitude estimation of radar objects using prior radar detections.

NVIDIA — Santa Clara *Summer 2018*
Autonomous Vehicles Validation Intern

· Realtime validation of the simulated autonomous vehicle (AV) perception DNN output by comparing with ground truth data. Developed C++ code directly within the AV Drivestack for injecting ground truth data into the planning pipeline. Created AV Scenario Editor tool for traffic scenario design.

NASA Ames Research Center — Mountain View *Summer 2017*
Robotics Intern

· Implemented Stereo & RGB-D SLAM in indoor environments on a custom quadcopter. Designed and machined parts of quadcopter chassis. Presented at annual Poster Symposium.

PUBLICATIONS

Suneel Belkhale, Tianli Ding, Ted Xiao, Pierre Sermanet, Quan Vuong, Jonathan Tompson, Yevgen Chebotar*, Debidatta Dwibedi*, Dorsa Sadigh*. RT-H: Robot Action Hierarchies using Language. RSS 2024.

Many authors. DROID: A Large-Scale In-The-Wild Robot Manipulation Dataset. RSS 2024

Open X-Embodiment Collaboration. Open X-Embodiment: Robotic Learning Datasets And RT-X Models.

Priya Sundaresan, **Suneel Belkhale**, Dorsa Sadigh, Jeannette Bohg. KITE: Keypoint-Conditioned Policies for Semantic Manipulation. CoRL 2023.

Suneel Belkhale, Yuchen Cui, Dorsa Sadigh. HYDRA: Hybrid Robot Actions for Imitation Learning. CoRL 2023.

Lorenzo Shaikewitz*, Yilin Wu*, **Suneel Belkhale***, Jennifer Grannen, Priya Sundaresan & Dorsa Sadigh. In-Mouth Robotic Bite Transfer with Visual and Haptic Sensing. *In submission, ICRA 2023.*

Suneel Belkhale, Dorsa Sadigh. PLATO: Predicting Latent Affordances Through Object-Centric Play. CoRL 2022.

Priya Sundaresan, **Suneel Belkhale**, & Dorsa Sadigh. Learning Visuo-Haptic Skewering Strategies for Robot-Assisted Feeding. CoRL 2022.

Jennifer Grannen, Yilin Wu, **Suneel Belkhale**, & Dorsa Sadigh. Learning Bimanual Scooping Policies for Food Acquisiton. CoRL 2022.

Suneel Belkhale, Ethan Gordon, Tracy Chen, Siddhartha Srinivasa, Tapomayukh Bhattacharjee, & Dorsa Sadigh. Balancing Efficiency and Comfort in Robot-Assisted Bite-Transfer. ICRA 2022.

Suneel Belkhale, Rachel Li, Gregory Kahn, Rowan McAllister, Roberto Calandra, & Sergey Levine. Model-based meta-reinforcement learning for flight with suspended payloads. ICRA 2021.

Katie Kang*, **Suneel Belkhale***, Gregory Kahn*, Pieter Abbeel, & Sergey Levine.
Generalization through simulation: Integrating simulated and real data into deep reinforcement learning for vision-based autonomous flight. ICRA 2019.

AWARDS

Regent and Chancellor's Scholarship *UC Berkeley*
CalHacks Fellowship Award *UC Berkeley*

ACADEMIC CLUBS

Machine Learning @ Berkeley *August 2016 - May 2020*
Officer / Project Manager for industry driven projects and application projects using ML.

Unmanned Aerial Vehicles @ Berkeley *August 2016 - May 2020*
Officer / Project Manager for unmanned aerial systems design, control, and autonomy.

PROJECTS

AutoQuad — UAVs @ Berkeley & ML @ Berkeley *2018*
Created realistic simulator for quadrotors using Unity 5. Utilized simulator to develop and train RL algorithms for autonomous navigation in outdoor environments.

Unity Reinforcement Learning — ML @ Berkeley *2017*
Implemented RL methods (e.g. Hierarchical RL) for Shadow Tactics game (sparse reward environment).

Battery Hotswap — UAVs @ Berkeley *2016 - 2017*

Developed custom hexacopter to allow smaller drones to land on and replace batteries while in air. Designed parts, electronics, and autonomous docking algorithm.

Tessellate — CalHacks 4.0 *2017 - 2018*
Built a 3D Scanner to reconstruct STL meshes from real world objects. Won CalHacks Fellowship Award to create final product through Skydeck partnership.

Structured VAE Learning — CS 280 Computer Vision *2019*
Researching methods to structure latent spaces for VAE. Studied how this affects learning, reconstruction, and generalization.

Temporally Linked Latent Spaces — CS 294-158 Deep Unsupervised Learning *2020*
Training generative models to support time interpolation in latent space (e.g. video frame interpolation given a start and end frame).