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 (current)*
Computer Science, Ph.D.
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*
I work on Vision Language Action Models (VLAs) for robotics, specifically leveraging language as a reasoning modality to enable large VLAs to reason about robot actions. Language-based action reasoning enables robot policies to scale better to diverse robot demonstration and cross-embodiment datasets, generalize better to unseen tasks, and to respond better to human feedback. PI: Prof. Dorsa Sadigh.

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.

INDUSTRY EXPERIENCE

Google Brain — Mountain View *June 2023 - Feb 2024*
Student Researcher

Investigating language-based reasoning in large Vision Language Action models. Developed RT-Hierarchy (RT-H), a VLA that reasons through a hierarchy of action representations: specifically, we predict *language motions* (e.g., "move forward", "open the gripper") before predicting robot actions. We found this action hierarchy enabled (1) easily ingesting diverse datasets (2) human-interpretable policies through reasoning, and (3) intervention in language space to correct errors and learn from these corrections efficiently.

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

[16] **Suneel Belkhale**, Dorsa Sadigh. MiniVLA: A Better VLA with a Smaller Footprint. https://ai.stanford.edu/blog/minivla/

[15] Suvir Mirchandani, **Suneel Belkhale**, Joey Hejna, Evelyn Choi, Md Sazzad Islam, Dorsa Sadigh. So You Think You Can Scale Up Autonomous Robot Data Collection? *CoRL 2024*.

[14] **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*.

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

[12] *Open X-Embodiment Collaboration*. Open X-Embodiment: Robotic Learning Datasets And RT-X Models. *ICRA 2024*

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

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

[9] **Suneel Belkhale**, Yuchen Cui, Dorsa Sadigh. Data Quality in Imitation Learning. *NeurIPS 2023*.

[8] Andy Shih, **Suneel Belkhale**, Dorsa Sadigh, Nima Anari. Parallel Sampling of Diffusion Models. *NeurIPS 2023*.

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

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

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

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

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

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

[1] 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*