

Suraj Nair

surajnair.com, surajnair0220@gmail.com

- EDUCATION** **Stanford University**, Stanford, CA 2018-2023
Ph.D. in Computer Science
Thesis: [Scaling Deep Robotic Learning to Broad Real-World Data](#)
Advisors: Chelsea Finn, Silvio Savarese
- California Institute of Technology**, Pasadena, CA 2014-2018
Bachelor of Science in Computer Science GPA: 3.9/4.0
Advisor: Yisong Yue
- EXPERIENCE** **Physical Intelligence (Pi)**, Founding Team/Research Scientist 2024-Present
Toyota Research Institute (TRI), Research Scientist 2023-2024
Facebook AI Research, Research Intern/Student Researcher 2021-2022
Google Brain, Research Intern/Student Researcher 2018-2019
Stanford Vision and Learning Lab, Visiting Researcher 2017
- PUBLICATIONS & PREPRINTS** [26] Physical Intelligence. π_0 : A Vision-Language-Action Flow Model for General Robot Control. *Arxiv Preprint*. 2024
- [25] Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, **Suraj Nair**, Rafael Rafailov, Ethan Foster, Grace Lam, Pannag Sanketi, Quan Vuong, Thomas Kollar, Benjamin Burchfiel, Russ Tedrake, Dorsa Sadigh, Sergey Levine, Percy Liang, Chelsea Finn. OpenVLA: An Open-Source Vision-Language-Action Model. *Conference on Robot Learning (CoRL)*. 2024.
- [24] Alexander Khazatsky*, Karl Pertsch*, **Suraj Nair**, Ashwin Balakrishna, ..., Thomas Kollar, Sergey Levine, Chelsea Finn. DROID: A Large-Scale In-The-Wild Robot Manipulation Dataset. *Robotics: Science and Systems (RSS)*. 2024.
- [23] Siddharth Karamcheti, **Suraj Nair**, Ashwin Balakrishna, Percy Liang, Thomas Kollar, Dorsa Sadigh. Prismatic VLMs: Investigating the design space of visually-conditioned language models. *International Conference on Machine Learning (ICML)*. 2024.
- [22] Siddharth Karamcheti, **Suraj Nair**, Annie S. Chen, Thomas Kollar, Chelsea Finn, Dorsa Sadigh, Percy Liang. Language-Driven Representation Learning for Robotics. *Robotics Science and Systems (RSS)*. 2023.
- [21] Maximilian Du, **Suraj Nair**, Dorsa Sadigh, Chelsea Finn Behavior Retrieval: Few-Shot Imitation Learning by Querying Unlabeled Datasets. *Robotics Science and Systems (RSS)*. 2023.
- [20] **Suraj Nair**, Aravind Rajeswaran, Vikash Kumar, Chelsea Finn, Abhinav Gupta. R3M: A Universal Visual Representation for Robot Manipulation. *Conference on Robot Learning (CoRL)*. 2022.
- [19] Maximilian Du*, Olivia Y. Lee*, **Suraj Nair**, Chelsea Finn. Play it by Ear: Learning Skills amidst Occlusion through Audio-Visual Imitation Learning. *Robotics: Science and Systems (RSS)*. 2022
- [18] **Suraj Nair**, Eric Mitchell, Kevin Chen, Brian Ichter, Silvio Savarese, Chelsea Finn.

- Learning Language-Conditioned Robot Behavior from Offline Data and Crowd-Sourced Annotation. *Conference on Robot Learning (CoRL)*. 2021.
- [17] Bohan Wu, **Suraj Nair**, Li Fei-Fei*, Chelsea Finn*. Example-Driven Model-Based Reinforcement Learning for Solving Long-Horizon Visuomotor Tasks. *Conference on Robot Learning (CoRL)*. 2021.
- [16] Mohammad Babaeizadeh, Mohammad Taghi Saffar, **Suraj Nair**, Sergey Levine, Chelsea Finn, Dumitru Erhan. FitVid: Overfitting in Pixel-Level Video Prediction. *Arxiv Preprint*. 2021
- [15] Annie Chen, **Suraj Nair**, Chelsea Finn. Learning Generalizable Robotic Reward Functions from "In-The-Wild" Human Videos. *Robotics: Science and Systems (RSS)*. 2021
- [14] Bohan Wu, **Suraj Nair**, Roberto Martin-Martin, Li Fei-Fei*, Chelsea Finn*. Greedy Hierarchical Variational Autoencoders for Large-Scale Video Prediction, *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2021
- [13] Stephen Tian, **Suraj Nair**, Frederik Ebert, Sudeep Dasari, Benjamin Eysenbach, Chelsea Finn, Sergey Levine. Model-Based Visual Planning with Self-Supervised Functional Distances. *International Conference on Learning Representations (ICLR)*. 2021.
- [12] Annie Chen*, HyunJi Nam*, **Suraj Nair***, Chelsea Finn. Batch Exploration with Examples for Scalable Robotic Reinforcement Learning. *Robotics and Automation Letters (RA-L) and IEEE International Conference on Robotics and Automation (ICRA)*. 2021
- [11] Brijen Thananjeyan*, Ashwin Balakrishna*, **Suraj Nair**, Michael Luo, Krishnan Srinivasan, Minh Hwang, Joey E. Gonzalez, Chelsea Finn, Ken Goldberg. Recovery RL: Safe Reinforcement Learning with Learned Recovery Zones. *Robotics and Automation Letters (RA-L) and IEEE International Conference on Robotics and Automation (ICRA)*. 2021
- [10] **Suraj Nair**, Silvio Savarese, Chelsea Finn. Goal-Aware Prediction: Learning to Model What Matters. *International Conference on Machine Learning (ICML)*. 2020.
- [9] Henrik Marklund*, **Suraj Nair***, Chelsea Finn. Exact (Then Approximate) Dynamics Programming for Deep Reinforcement Learning *Workshop on Biases, Invariances, and Generalization in RL, International Conference on Machine Learning (ICML)*. 2020.
- [8] **Suraj Nair**, Chelsea Finn. Hierarchical Foresight: Self-Supervised Learning of Long-Horizon Tasks via Visual Subgoal Generation. *International Conference on Learning Representations (ICLR)*. 2020.
- [7] **Suraj Nair**, Mohammad Babaeizadeh, Chelsea Finn, Sergey Levine, Vikash Kumar. Time Reversal As Self-Supervision. *IEEE International Conference on Robotics and Automation (ICRA)*. 2020.
- [6] **Suraj Nair**, Yuke Zhu, Silvio Savarese, Li Fei-Fei. Causal Induction from Visual Observations for Goal Directed Tasks. *Workshop on Causal Machine Learning, Neural Information Processing Systems (NeurIPS)*. 2019.
- [5] Sudeep Dasari, Frederik Ebert, Stephen Tian, **Suraj Nair**, Bernadette Bucher, Karl Schmeckpeper, Siddharth Singh, Sergey Levine, Chelsea Finn. RoboNet: Large-Scale Multi-Robot Learning. *Conference on Robot Learning (CoRL)*. 2019.
- [4] De-An Huang*, **Suraj Nair***, Danfei Xu*, Yuke Zhu, Animesh Garg, Li Fei-Fei, Silvio Savarese, Juan Carlos Niebles. Neural Task Graphs: Generalizing to Unseen Tasks from a Single Video Demonstrations. *IEEE Conference on Computer Vision and Pattern*

Recognition (CVPR). 2019.

[3] Danfei Xu*, **Suraj Nair***, Yuke Zhu, Julian Gao, Animesh Garg, Li Fei-Fei, Silvio Savarese. Neural Task Programming: Learning to Generalize Across Hierarchical Tasks. *IEEE International Conference on Robotics and Automation (ICRA)*. 2018.

[2] Men-Andrin Meier, Zachary E Ross, Anshul Ramachandran, Ashwin Balakrishna, **Suraj Nair**, Peter Kundzicz, Zefeng Li, Jennifer Andrews, Egill Hauksson, Yisong Yue. Reliable RealTime Seismic Signal/Noise Discrimination With Machine Learning. *Journal of Geophysical Research: Solid Earth*. 2019.

[1] **Suraj Nair**, Anshul Ramachandran, Peter Kundzicz. Annotated Reconstruction of 3D Spaces Using Drones. *IEEE MIT URTC*. 2017. **Best Paper Presentation**.

INVITED TALKS

Efficiently Scaling Data for Robot Learning
6th Robot Learning Workshop @ NeurIPS 2023 . December 2023

Supervising Robot Learning with Human Video Data
3rd International Ego4D Workshop @ CVPR 2023 Keynote . June 2023

Scaling Deep Robotic Learning to Broad Real-World Data
Boston Dynamics AI Institute. December 2022
Tesla. December 2022
Toyota Research Institute (TRI). November 2022
Allen Institute for AI (AI2). November 2022
Dyson Robotics. October 2022

Supervising Robot Learning with Language and Video from the Web
Columbia Artificial Intelligence and Robotics Lab March 2023
Covariant.ai. October 2022
MILA. September 2022
Toyota Research Institute (TRI). June 2022
Nuro.ai. April 2022
University of Cambridge Language Technology Lab Seminar. November 2021

Time Reversal as Self-Supervision
Berkeley Robotic Artificial Intelligence and Learning Lab. October 2018

TEACHING

Head Teaching Assistant, Stanford University
CS 330: Deep Multi-Task and Meta Learning 2022
Teaching Assistant, Stanford University
CS 330: Deep Multi-Task and Meta Learning 2019, 2020
Teaching Assistant, California Institute of Technology
CS/EE 155: Machine Learning/Data Mining 2017
CS 121: Introduction to Relational Databases 2016

AWARDS & HONORS

Best Paper Award Finalist (Robotics Science and Systems) 2023
For "Language-Driven Representation Learning for Robotics."

Best Paper Award (ICRA Scaling Robot Learning Workshop) 2022
For "R3M: A Universal Visual Representation for Robot Manipulation"

Robotics: Science and Systems (RSS) Pioneer 2022
Selected as one of 30 top early career researchers in robotics

ICLR Highlighted Reviewer Award 2021, 2022
Awarded to top 10% of reviewers

Stanford Nominee for Apple ML/AI PhD Fellowship 2020
Selected as one of 5 university nominees

National Science Foundation Graduate Research Fellowship 2018-2021

Best Paper Presentation - IEEE MIT URTC 2017

Caltech Summer Undergraduate Research Fellowship Recipient 2017
1st Place GE Digital Intern Hackathon 2016

PROFESSIONAL *Workshop Organization:*

ACTIVITIES Co-Organizer: Deep Reinforcement Learning Workshop at NeurIPS 2022
Co-Organizer: Workshop on Learning from Diverse, Offline Data at RSS 2022
Co-Organizer: Workshop on Learning from Diverse, Offline Data at ICRA 2023

Paper Reviewing:

Neural Information Processing Systems (NeurIPS) 2020-2022
International Conference on Machine Learning (ICML) 2020-2022
International Conference on Learning Representations (ICLR) 2019-2022
IEEE International Conference on Robotics and Automation (ICRA) 2019-2021
Conference on Robot Learning (CoRL) 2021, 2022
Robotics Science and Systems (RSS) 2023
IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2019, 2020
International Conference on Computer Vision (ICCV) 2021

**ADVISING &
MENTORSHIP**

Niveditha Iyer B.S., Stanford
Patricia Strutz B.S., Stanford
Olivia Lee B.S., Stanford
Maximilian Du B.S., Stanford Next: Ph.D. CS, Stanford
HyunJi Nam B.S. Stanford, Next: Software engineer at ScaleAI
Annie Chen B.S. Stanford, Next: Ph.D. CS, Stanford