

# Yunfan Jiang

🌐 WEBSITE: [yunfanj.ai](http://yunfanj.ai)

✉ EMAIL: [yunfanj@cs.stanford.edu](mailto:yunfanj@cs.stanford.edu)

📍 ADDRESS: Gates 220, 353 Jane Stanford Way, Stanford, CA 94305

**RESEARCH INTEREST** I am interested in the intersection of machine learning and robotics with a focus on **building generalist robots that perform diverse everyday activities.**

**EDUCATION** **Stanford University**, Stanford, CA Sept. 2023 – Now  
*Ph.D.* in Computer Science, advised by Prof. Fei-Fei Li

**Stanford University**, Stanford, CA June 2023  
*M.S.* in Electrical Engineering

**The University of Edinburgh**, Edinburgh, UK July 2020  
*B.Eng.* in Electronics & Electrical Engineering with First-Class Honors

**HONORS & AWARDS**

- **NVIDIA Graduate Fellowship** 2025
- IEEE ICRA 2024 Best Conference Paper Award 2024
- NeurIPS 2023 Scholar Award 2023
- **Stanford Engineering Exceptional Master's Student Award** 2023
- ICML Conference Travel Award 2023
- **NeurIPS 2022 Outstanding Datasets & Benchmarks Paper Award** 2022
- Ewart Farvis Prize (Outstanding Bachelor Thesis) 2020
- The University of Edinburgh School of Engineering Scholarship 2018, 2019

**EXPERIENCE** **Stanford Vision and Learning Lab** Stanford, CA  
Graduate Research Assistant Sept. 2023 – Now

- Advised by Prof. Fei-Fei Li.
- Built synergistic frameworks for learning whole-body manipulation for everyday activities [13, 14, 1, 2].
- Developed novel methods for robotic policy learning, with a focus on sim-to-real transfer by learning from human correction [11] and real-to-sim-to-real transfer for robust policy learning [12].

**NVIDIA Research** Santa Clara, CA  
Research Intern ◊ GEAR June 2025 – Now

- Hosted by Dr. Jim Fan and Prof. Yuke Zhu.
- Core member of the GR00T team, developing next-generation NVIDIA robot foundation models [18, S5].
- Conducted research on frontier robotics problems, including in-context learning [18], dexterous manipulation [15, 16], and real-to-sim-to-real transfer [17].

**Boston Dynamics AI Institute** (now RAI Institute)

Cambridge, MA

Research Intern

June 2023 – Aug. 2023

- Hosted by Dr. David Watkins and Dr. Jennifer Barry.
- Developed *VIMA+*, an extension of VIMA [7] to a real UR5e robot.
- Investigated video prediction models for robotic manipulation at scale.

**NVIDIA Research**

Santa Clara, CA

Research Intern ◊ AI Algorithm Team

June 2022 – Jan. 2023

- Hosted by Dr. Jim Fan, Prof. Yuke Zhu, and Prof. Anima Anandkumar.
- Developed embodied agents empowered by foundation models in various domains such as those related to robot learning [7] and open-ended video games [6, 9].
- Created a novel algorithm to enhance Transformer agents’ learning efficiency and generalization [8].

## REFEREED PUBLICATIONS

\* Equal contribution. † Equal advising.

- [18] **Yunfan Jiang**, Yevgen Chebotar, Ruijie Zheng, Fengyuan Hu, Yunhao Ge, Jimmy Wu, Tianyuan Dai, Scott Reed, Li Fei-Fei<sup>†</sup>, Yuke Zhu<sup>†</sup>, and Linxi Fan<sup>†</sup>. “RoboTTT: Context Scaling for Robot Policies”. In: *10th Annual Conference on Robot Learning (CoRL)*. 2026. **Oral Presentation**.
- [17] Nadun Ranawaka\*, Josiah Wong\*, Wei-Lin Pai, Wei-Teng Chu, Tianyuan Dai, Masoud Moghani, Hang Yin, **Yunfan Jiang**, Wesley Durbano\*, Brandon Huynh\*, Yu Fang, Danfei Xu, Ruohan Zhang, Li Fei-Fei, Linxi Fan, Bowen Wen, Ajay Mandlekar<sup>†</sup>, and Yuke Zhu<sup>†</sup>. “SimFoundry: Modular and Automated Scene Generation for Policy Learning and Evaluation”. In: *10th Annual Conference on Robot Learning (CoRL)*. 2026.
- [16] Dantong Niu\*, Zhuoyang Liu\*, Zekai Wang\*, Boning Shao, Zhao-Heng Yin, Anirudh Pai, Yuvan Sharma, Stefano Saravalle, Ruijie Zheng, Jing Wang, Ryan Punamiya, Mengda Xu, Yuqi Xie, **Yunfan Jiang**, Letian Fu, Konstantinos Kallidromitis, Matteo Gioia, Junyi Zhang, Jiaxin Ge, Haiwen Feng, Fabio Galasso, Wei Zhan, David M. Chan, Yutong Bai, Roei Herzig, Jiahui Lei, Li Fei-Fei, Ken Goldberg, Jitendra Malik, Pieter Abbeel, Yuke Zhu, Danfei Xu, Linxi Fan, and Trevor Darrell. “T-Rex: Tactile-Reactive Dexterous Manipulation”. In: *10th Annual Conference on Robot Learning (CoRL)*. 2026.
- [15] Ruijie Zheng\*, Dantong Niu\*, Yuqi Xie\*, Jing Wang, Mengda Xu, **Yunfan Jiang**, Fernando Castañeda, Fengyuan Hu, You Liang Tan, Letian Fu, Trevor Darrell, Furong Huang, Yuke Zhu<sup>†</sup>, Danfei Xu<sup>†</sup>, and Linxi Fan<sup>†</sup>. “EgoScale: Scaling Dexterous Manipulation with Diverse Egocentric Human Data”. In: *10th Annual Conference on Robot Learning (CoRL)*. 2026. **Oral Presentation**.
- [14] Chengshu Li\*, Mengdi Xu\*, Arpit Bahety\*, Hang Yin\*, **Yunfan Jiang**, Huang Huang, Josiah Wong, Sujay Garlanka, Cem Gokmen, Ruohan Zhang, Weiyu Liu, Jiajun Wu, Roberto Martín-Martín, and Li Fei-Fei. “MoMaGen: Generating Demonstrations under Soft and Hard Constraints for Multi-Step Bimanual Mobile Manipulation”. In: *The Fourteenth International Conference on Learning Representations (ICLR)*. 2026.
- [13] **Yunfan Jiang**, Ruohan Zhang, Josiah Wong, Chen Wang, Yanjie Ze, Hang Yin, Cem Gokmen, Shuran Song, Jiajun Wu, and Li Fei-Fei. “BEHAVIOR Robot Suite: Streamlining Real-World Whole-Body Manipulation for Everyday Household Activities”. In: *9th Annual Conference on Robot Learning (CoRL)*. 2025. Also **Oral Presentation** at RSS 2025 Whole-body Control and Bimanual Manipulation Workshop.

- [12] Tianyuan Dai\*, Josiah Wong\*, **Yunfan Jiang**, Chen Wang, Cem Gokmen, Ruohan Zhang, Jiajun Wu, and Li Fei-Fei. “Automated Creation of Digital Cousins for Robust Policy Learning”. In: *8th Annual Conference on Robot Learning (CoRL)*. 2024.
- [11] **Yunfan Jiang**, Chen Wang, Ruohan Zhang, Jiajun Wu, and Li Fei-Fei. “TRANSIC: Sim-to-Real Policy Transfer by Learning from Online Correction”. In: *8th Annual Conference on Robot Learning (CoRL)*. 2024.
- [10] Open X-Embodiment Collaboration (including **Yunfan Jiang**). “Open X-Embodiment: Robotic Learning Datasets and RT-X Models”. In: *IEEE International Conference on Robotics and Automation (ICRA)*. 2024, pp. 6892–6903. DOI: 10.1109/ICRA57147.2024.10611477. **Best Conference Paper Award**.
- [9] Guanzhi Wang, Yuqi Xie, **Yunfan Jiang**\*, Ajay Mandlekar\*, Chaowei Xiao, Yuke Zhu, Linxi Fan<sup>†</sup>, and Anima Anandkumar<sup>†</sup>. “Voyager: An Open-Ended Embodied Agent with Large Language Models”. In: *Transactions on Machine Learning Research* (2024). Also **Oral Presentation** at NeurIPS 2023 Agent Learning in Open-Endedness Workshop and Intrinsically Motivated Open-Ended Learning Workshop.
- [8] Lucy Xiaoyang Shi\*, **Yunfan Jiang**\*, Jake Grigsby, Linxi Fan<sup>†</sup>, and Yuke Zhu<sup>†</sup>. “Cross-Episodic Curriculum for Transformer Agents”. In: *Conference on Neural Information Processing Systems (NeurIPS)*. 2023.
- [7] **Yunfan Jiang**, Agrim Gupta\*, Zichen Zhang\*, Guanzhi Wang\*, Yongqiang Dou, Yanjun Chen, Li Fei-Fei, Anima Anandkumar, Yuke Zhu<sup>†</sup>, and Linxi Fan<sup>†</sup>. “VIMA: General Robot Manipulation with Multimodal Prompts”. In: *International Conference on Machine Learning (ICML)*. 2023. Also **Oral Presentation** at NeurIPS 2022 Foundation Models for Decision Making Workshop.
- [6] Linxi Fan, Guanzhi Wang\*, **Yunfan Jiang**\*, Ajay Mandlekar, Yuncong Yang, Haoyi Zhu, Andrew Tang, De-An Huang, Yuke Zhu<sup>†</sup>, and Anima Anandkumar<sup>†</sup>. “MineDojo: Building Open-Ended Embodied Agents with Internet-Scale Knowledge”. In: *Conference on Neural Information Processing Systems (NeurIPS), Datasets and Benchmarks Track*. 2022. **Outstanding Paper Award, Featured Paper Presentation**.
- [5] Yueyi Jiang, **Yunfan Jiang**, Liu Leqi, and Piotr Winkielman. “Many Ways to Be Lonely: Fine-Grained Characterization of Loneliness and Its Potential Changes in COVID-19”. In: *Proceedings of the International AAAI Conference on Web and Social Media (ICWSM)* 16.1 (2022), pp. 405–416.
- [4] **Yunfan Jiang**, Jingjing Si, Rui Zhang, Godwin Enemali, Bin Zhou, Hugh McCann, and Chang Liu. “CSTNet: A Dual-Branch Convolutional Neural Network for Imaging of Reactive Flows Using Chemical Species Tomography”. In: *IEEE Transactions on Neural Networks and Learning Systems* 34.11 (2023), pp. 9248–9258. DOI: 10.1109/TNNLS.2022.3157689. Submitted in 2020.

## TECHNICAL REPORTS

- [3] Chenyuan Zhou\*, Yining Hong\*, Huang Huang, **Yunfan Jiang**, Wenlong Huang, David Held, Vitor Campagnolo Guizilini, Yue Wang, Li Fei-Fei, Jiajun Wu<sup>†</sup>, Leonidas Guibas<sup>†</sup>, and Yejin Choi<sup>†</sup>. *T2-z: Test-Time Optimization of Environment Latents for Robotic World Models*. 2026. arXiv: 2610.22364 [cs.RO].
- [2] Wensi Ai\*, Christopher Luey\*, Karthik Dharmarajan\*, Evans Han, **Yunfan Jiang**, Stef Ren, Jianwen Xie, Jiajun Wu, Li Fei-Fei<sup>†</sup>, and Ruohan Zhang<sup>†</sup>. *I3L: Iterative Interactive Imitation Learning for Whole-Body Mobile Manipulation*. 2026. arXiv: 2610.17748 [cs.RO].

- [1] Evans Han, **Yunfan Jiang**, Yingke Wang\*, Haoyue Xiao\*, Huang Huang, Jianwen Xie, Jiajun Wu, Li Fei-Fei, and Ruohan Zhang. *StereoPolicy: Improving Robotic Manipulation Policies via Stereo Perception*. 2026. arXiv: 2605.09989 [cs.RO].

## SOFTWARE

- [S5] *NVIDIA Isaac GR00T N1.6*. <https://github.com/NVIDIA/Isaac-GR00T>. **8.1K+ GitHub Stars**.
- [S4] *BEHAVIOR-Robot-Suite*. <https://github.com/behavior-robot-suite/brs-algo> and <https://github.com/behavior-robot-suite/brs-ctrl>. **300+ GitHub Stars**.
- [S3] *TRANSIC*. <https://github.com/transic-robot/transic> and <https://github.com/transic-robot/transic-envs>. **150+ GitHub Stars**.
- [S2] *VIMA*. <https://github.com/vimalabs/VIMA> and <https://github.com/vimalabs/VIMABench>. **1.1K+ GitHub Stars**.
- [S1] *MineDojo*. <https://github.com/MineDojo/MineDojo> and <https://github.com/MineDojo/MineCLIP>. **2.5K+ GitHub Stars**.

## SPEECHES

- “*RoboTTT: Context Scaling for Robot Policies*”
- Invited Talk at Skild AI Aug. 2026
  - Invited Talk at Enigma Journal Club, Tolias Lab, Stanford Aug. 2026
  - Invited Talk at REAL @ Stanford, Host: Prof. Shuran Song Aug. 2026
- “*BEHAVIOR Robot Suite: Streamlining Real-World Whole-Body Manipulation for Everyday Household Activities*”
- Oral Presentation at RSS 2025 Whole-body Control and Bimanual Manipulation Workshop, Los Angeles, CA June 2025
  - Invited Talk at Moonshot AI Mar. 2025
- “*VIMA: General Robot Manipulation with Multimodal Prompts*” [pdf]
- Invited Talk at Boston Dynamics AI Institute Feb. 2023
  - Oral Presentation at NeurIPS 2022 Foundation Models for Decision Making Workshop, New Orleans, LA Dec. 2022
  - Invited Talk at Prof. Anima Anandkumar’s Caltech Group Nov. 2022
  - Invited Talk at Inspir.ai Oct. 2022
- “*MineDojo: Building Open-Ended Embodied Agents with Internet-Scale Knowledge*” [pdf]
- Lecture at Stanford CS422 Interactive and Embodied Learning Feb. 2023
  - Invited Talk at Inspir.ai Oct. 2022
  - Co-Presentation at Prof. Anima Anandkumar’s Caltech Group Aug. 2022

**PROFESSIONAL SERVICES**   **Workshop & Challenge Organizer**

- *Organizer*, 1st BEHAVIOR Challenge, Conference on Neural Information Processing Systems (NeurIPS), 2025 [link]
- *Organizer*, Mobile Manipulation: Emerging Opportunities & Contemporary Challenges, Robotics: Science and Systems (RSS), 2025 [link]

**Workshop Program Committee**

- 2nd Workshop on Foundation Models for Decision Making, Conference on Neural Information Processing Systems (NeurIPS), 2023 [link]

**Journal Reviewer**

- IEEE Robotics and Automation Letters (RA-L)

**Conference Reviewer**

- Conference on Robot Learning (CoRL)
- International Conference on Machine Learning (ICML)
- International Conference on Learning Representations (ICLR)
- Conference on Neural Information Processing Systems (NeurIPS)
- Conference on Computer Vision and Pattern Recognition (CVPR)
- International Conference on Robotics and Automation (ICRA)
- International Conference on Intelligent Robots and Systems (IROS)
- AAAI Conference on Artificial Intelligence (AAAI)

**Workshop Proposal Reviewer**

- International Conference on Learning Representations (ICLR)

**TEACHING**

**Course Assistant**, Stanford University

- CS221 Artificial Intelligence: Principles and Techniques, Autumn 2026, Instructor: Prof. Percy Liang
- CS231N Deep Learning for Computer Vision, Spring 2025, Instructors: Prof. Fei-Fei Li and Prof. Ehsan Adeli

**Course Grader**, Stanford University

- ENGR76 Information Science and Engineering, Spring 2023, Instructor: Prof. Ayfer Özgür
- EE364A Convex Optimization I, Winter 2023, Instructor: Prof. Stephen Boyd
- EE277 Reinforcement Learning: Behaviors and Applications, Autumn 2021, Instructor: Prof. Benjamin Van Roy
- EE236A Modern Optics, Autumn 2021, Instructor: Dr. Mohammad Zaman

**SELECTED MEDIA COVERAGE**

- [M6] “State of AI Report 2026,” by Nathan Benaich, *Air Street Capital*. Oct. 8, 2026. URL: <https://www.stateof.ai/2026>.
- [M5] “NVIDIA Awards up to \$60,000 Research Fellowships to PhD Students,” by Sylvia Chanak, *NVIDIA Blog*. Dec. 4, 2025. URL: <https://blogs.nvidia.com/blog/graduate-fellowship-recipients-2026-2027/>.

- [M4] “BEHAVIOR Challenge Charts the Way Forward for Domestic Robotics,” by Andrew Myers, *Stanford HAI News*. Sept. 22, 2025. URL: <https://hai.stanford.edu/news/behavior-challenge-charts-the-way-forward-for-domestic-robotics>.
- [M3] Stanford Electrical Engineering 2023 Commencement Ceremony and Awards. June 20, 2023. URL: <https://ee.stanford.edu/2023-commencement-ceremony-and-awards>.
- [M2] NVIDIA GTC Jensen Huang Keynote. Mar. 21, 2023. URL: <https://www.nvidia.com/en-us/on-demand/session/gtcspring23-s52226/>.
- [M1] “Building Generally Capable AI Agents with MineDojo,” by Nathan Horrocks, *NVIDIA Technical Blog*. July 1, 2022. URL: <https://developer.nvidia.com/blog/building-generally-capable-ai-agents-with-minedojo/>.