

MAX (LETIAN) FU

Berkeley, CA \diamond max-fu.github.io \diamond max.fu.letian@berkeley.edu

EDUCATION

University of California, Berkeley PhD. Electrical Engineering and Computer Sciences.	<i>August 2023 - May 2026 (Planned)</i>
University of California, Berkeley M.S. Electrical Engineering and Computer Sciences.	<i>August 2022 - May 2023</i>
University of California, Berkeley B.A. Computer Science and Applied Mathematics.	<i>August 2018 - May 2022</i>

SYSTEMS

- **GR00T-N1.6-3B**. Foundation Vision-Language-Action Model for Robotics. GEAR Team (alphabetical authors incl. **Letian Fu**). NVIDIA, 2025. Model Card.

PUBLICATIONS

(*denotes equal contribution)

- **Letian Fu***, Justin Yu*, Karim El-Refai*, Ethan Kou*, Haoru Xue*, Huang Huang, Wenli Xiao, Guanzhi Wang, Fei-Fei Li, Guanya Shi, Jiajun Wu, Shankar Sastry, Yuke Zhu, Ken Goldberg[†], Linxi “Jim” Fan[†]. CaP-X: A Framework for Benchmarking and Improving Coding Agents for Robot Manipulation. *In submission*, 2026.
- Ruijie Zheng, Dantong Niu, Yuqi Xie, Jing Wang, Mengda Xu, Yunfan Jiang, Fernando Castaeda, Fengyuan Hu, You Liang Tan, **Letian Fu**, Trevor Darrell, Furong Huang, Yuke Zhu, Danfei Xu, Linxi “Jim” Fan. EgoScale: Scaling Dexterous Manipulation with Diverse Egocentric Human Data. *In submission*, 2026.
- Zirui Wang*, Junyi Zhang*, Jiaxin Ge*, Long Lian, **Letian Fu**, Lisa Dunlap, Ken Goldberg, Xudong Wang, Ion Stoica, David M. Chan, Sewon Min, Joseph E. Gonzalez. VisGym: Diverse, Customizable, Scalable Environments for Multimodal Agents. *In submission*, 2026.
- Justin Yu*, **Letian Fu***, Huang Huang, Karim El-Refai, Rares Ambrus, Richard Cheng, Muhammad Zubair Irshad, Ken Goldberg. Real2Render2Real: Scaling Robotic Manipulation Data Without Dynamics Simulation or Robot Hardware. *Proceedings of 2025 Conference on Robot Learning (Oral)*. Seoul, Korea. Sep, 2025.
- Huang Huang*, Fangchen Liu*, **Letian Fu***, Tingfan Wu, Mustafa Mukadam, Jitendra Malik, Ken Goldberg, Pieter Abbeel. OTTER: A Vision-Language-Action Model with Text-Aware Visual Feature Extraction. *Proceedings of 2025 International Conference on Machine Learning*. Vancouver, Canada. Jul, 2025.
- **Letian Fu***, Huang Huang*, Gaurav Datta*, Lawrence Yunliang Chen, William Chung-Ho Panitch, Fangchen Liu, Hui Li, Ken Goldberg. In-Context Imitation Learning via Next-Token Prediction. *Proceedings of 2025 IEEE International Conference on Robotics and Automation*. Atlanta, US. May, 2025.
- Andrew Goldberg, Kavish Kondap, Tianshuang Qiu, Zehan Ma, **Letian Fu**, Justin Kerr, Huang Huang, Kaiyuan Chen, Kuan Fang, Ken Goldberg. Blox-Net: Generative Design-for-Robot-Assembly Using VLM Supervision, Physics Simulation, and a Robot with Reset. *Proceedings of 2025 IEEE International Conference on Robotics and Automation*. Atlanta, US. May, 2025.
- Kaiyuan Chen, **Letian Fu**, Siyuan Fu, Lawrence Yunliang Chen, Huang Huang, Kush Hari, Ashwin Balakrishna, Pannag R Sanketi, John Kubiatawicz, Ken Goldberg. Robo-DM: Efficient Robot Big Data Management. *Proceedings of 2025 IEEE International Conference on Robotics and Automation (Best Paper)*. Atlanta, US. May, 2025.
- **Letian Fu***, Long Lian*, Renhao Wang, Baifeng Shi, Xudong Wang, Adam Yala[†], Trevor Darrell[†], Alexei A Efros[†], Ken Goldberg[†]. Rethinking Patch Dependence for Masked Autoencoders. *2025 Transactions on Machine Learning Research*.

- **Letian Fu**, Gaurav Datta*, Huang Huang*, William Chung-Ho Panitch*, Jaimyn Drake*, Joseph Ortiz, Mustafa Mukadam, Mike Lambeta, Roberto Calandra, Ken Goldberg. A Touch, Vision, and Language Dataset for Multimodal Alignment. *Proceedings of 2024 International Conference on Machine Learning (Oral)*. Vienna, Austria. July 2024.
- Adam Rashid, Chung Min Kim, Justin Kerr, **Letian Fu**, Kush Hari, Ayah Ahmad, Kaiyuan Chen, Huang Huang, Marcus Gualtieri, Michael Wang, Christian Juette, Nan Tian, Liu Ren, Ken Goldberg. Lifelong LERF: Local 3D Semantic Inventory Monitoring Using FogROS2. *Proceedings of 2024 IEEE International Conference on Robotics and Automation*. Yokohama, Japan. May 2024.
- Ilija Radosavovic, Baifeng Shi, **Letian Fu**, Ken Goldberg, Trevor Darrell, Jitendra Malik. Robot Learning with Sensorimotor Pre-training. *Proceedings of 2023 Conference on Robot Learning (Oral)*. Atlanta, US. Nov, 2023.
- **Letian Fu**, Huang Huang, Lars Berscheid, Hui Li, Ken Goldberg, Sachin Chitta. Safely Learning Visuo-Tactile Feedback Policies in Real For Industrial Insertion. *Accepted by 2023 IEEE International Conference on Robotics and Automation*. London, UK. May, 2023.
- Justin Kerr, **Letian Fu**, Huang Huang, Jeffrey Ichnowski, Matthew Tancik, Yahav Avigal, Angjoo Kanazawa, Ken Goldberg. EvoNeRF: Evolving NeRF for Sequential Robot Grasping. *Proceedings of 2022 Conference on Robot Learning (Oral)*. Auckland, NZ. Dec. 2022.
- Huang Huang*, **Letian Fu***, Michael Danielczuk, Chung Min Kim, Zachary Tam, Jeffrey Ichnowski, Anelia Angelova, Brian Ichter, Ken Goldberg. Mechanical Search on Shelves with Efficient Stacking and Destacking of Objects. *Proceedings of 2022 International Symposium on Robotics Research*. Geneva, Switzerland. Sep, 2022.
- **Letian Fu**, Michael Danielczuk, Ashwin Balakrishna, Daniel S. Brown, Jeffrey Ichnowski, Eugen Solowjow, Ken Goldberg. LEGS: Learning Efficient Grasp Sets for Exploratory Grasping. *Proceedings of 2022 IEEE International Conference on Robotics and Automation*. Philadelphia, PA. May, 2022.
- Huang Huang, Michael Danielczuk, Chung Min Kim, **Letian Fu**, Zachary Tam, Jeffrey Ichnowski, Anelia Angelova, Brian Ichter, Ken Goldberg. Mechanical Search on Shelves using a Novel Bluction Tool. *Proceedings of 2022 IEEE International Conference on Robotics and Automation*. Philadelphia, PA. May, 2022.
- Yeshwant Reddy Chillakuru, Kyle Kranen, Vishnu Doppalapudi, Zhangyuan Xiong, **Letian Fu**, Aarash Heydari, Aditya Sheth, Youngho Seo, Thienkhai Vu, Jae Ho Sohn. High precision localization of pulmonary nodules on chest CT utilizing axial slice number labels. *BMC Med Imaging* 21, 66 (2021).

RESEARCH AND WORK

Graduate Student Researcher

December 2020 - Present

Berkeley Artificial Intelligence Research, AUTOLAB

Advised by Prof. Ken Goldberg. Currently working on vision, multimodal, and sensorimotor pre-training so that robots can quickly adapt and learn new skills; works accepted to ICML, ICRA, ISRR, and CoRL.

Research Intern

January 2025 - Present

NVIDIA

Advised by Linxi Fan and Prof. Yuke Zhu under the Generalist Embodied Agent Research (GEAR) team. Research on creating more general and scalable vision language action models for robotics (i.e. GR00T N1.6).

Robotics Research Intern

March 2022 - September 2022

Autodesk

Advised by Sachin Chitta and Hui Li; researched on enabling robots to safely perform industrial insertion tasks with vision, tactile, and force-torque feedback; research conducted on a Franka Emika robot; the resulting publication is submitted to ICRA 2023.

Machine Learning Intern

June 2020 - September 2020

Apple

Advised by Daniel Ulbricht and Mohammad Haris Baig; applied computer vision and deep learning to internal development; researched, designed and implemented real-time semantic segmentation algorithms; improved semantic segmentation performance by leveraging geometrical priors; designed new metrics and benchmarked the developed algorithms; developed model evaluation and visualization pipelines.

TEACHING

EECS C106B/206B (Spring 2026): Robotic Manipulation and Interaction
EECS 127 (Fall 2021): Optimization Models in Engineering
CS 182/282A (Spring 2021): Deep Neural Network Architecture
CS 170 (Fall 2020): Efficient Algorithms and Intractable Problems
Math 1B (Summer 2019): Calculus II

MISC.

Reviewer: ICLR, ICML, CVPR, NeurIPS, ICRA, IROS, ISRR, CASE, CoRL, RAL, L4DC

INVITED TALKS AND PANELS

In-Context Imitation Learning as Next-Token Prediction. Amazon Frontier AI & Robotics Org. February 2025.
Aligning Touch, Vision, and Language for Multimodal Perception. NeurIPS: 2nd Workshop on Touch Processing. Dec 2024.
In-Context Imitation Learning as Next-Token Prediction. Google DeepMind Robotics. September 2024.
Aligning Touch, Vision, and Language for Multimodal Perception. Multimodal AI webinar at Twelve Labs. June 2024.
Aligning Touch, Vision, and Language for Multimodal Perception. Embodied AI Seminar at the Fundamental AI Research (FAIR) team at Meta. Mar 2024.