Feng Gao

fenggo@amazon.com • f.gao@ucla.edu • +1 (310) 569-4532 • https://fen9.github.io

EDUCATION

University of California, Los Angeles, Los Angeles, California, USA — Sep 2017 – Jun 2022

- Ph.D. in Statistics
  - Adviser: Prof. Mark S. Handcock, Prof. Ying Nian Wu
  - Ph.D. Thesis: Multi-Modal Robotic Learning, Reasoning and Planning
  - Research Areas: Artificial Intelligence, Computer Vision, Robotics
  - Focus Topics: Robotic Learning, Multi-Modal Visual Reasoning, Robot Planning
  - Funded by: DARPA SIMPLEX, DARPA XAI, ONR-MURI
  - Collaborated with Jet Propulsion Laboratory

University of Southern California, Los Angeles, California, USA — Aug 2015 – May 2017

- M.S. in Computer Science

University of Electronic Science and Technology of China (UESTC) — Sep 2011 – Jun 2015

- B.Eng. in Software Engineering
  - Graduated with various honors

RESEARCH EXPERIENCE

Amazon

- Applied Scientist (Rufus-Multi-Modal team) — Jul 2022 – Present
  - Core member of Rufus-MM. Full stack M-LLM including pre-training, IFT, alignment and evaluation.
  - Multiple paper accepted in CVPR, NeurIPS, ECCV, EMNLP.
  - Embodied AI and text to image generation (prior to Rufus in Alexa)
- Applied Scientist Intern — Jun 2021 – Nov 2021
  - One paper accepted by CVPR2022 in vision-language reasoning with external knowledge.

Center for Vision, Cognition, Learning and Autonomy, UCLA

- Graduate Student Researcher, Statistics Department — Sep 2017 – Dec 2020
  - Research Interest: Artificial Intelligence in Robotics
  - Supervisor: Prof.Song-Chun Zhu
  - Focus Areas: Robotics, Computer Vision
- Visiting Graduate Researcher, Statistics Department — Jul 2016 – Aug 2017
  - Project: Human Robot Collaboration

International Center for Artificial Intelligence and Robot Autonomy (CARA) — Sep 2017 – Sep 2019

- Robotics Research Engineer Intern
  - Research Projects: Visual Abstraction Reasoning, Human-Robot Collaboration
  - Supervisor: Dr. Yixin Zhu

PUBLICATIONS

JOURNALS

- A Tale of Two Explanations: Enhancing Human Trust by Explaining Robot Behavior
  M. Edmonds*, **F. Gao***, H. Liu*, X. Xie*, S. Qi, B. Rothrock, Y. Zhu, Y.N. Wu, H. Lu, S.-C. Zhu
  Science Robotics 18 Dec 2019: Vol. 4, Issue 37, eaay4663 (* Joint First Authors)
- Dark, Beyond Deep: A Paradigm Shift to Cognitive AI with Human-like Commonsense
  Y. Zhu, T. Gao, L. Fan, S. Huang, M. Edmonds, H. Liu, **F. Gao**, C. Zhang, S. Qi, Y.N. Wu, J.B. Tenenbaum, S.-C. Zhu
  Engineering, Special Issue on Artificial Intelligence, 2020

CONFERENCES

- Skews in the Phenomenon Space Hinder Generalization in Text-to-Image Generation
  Y. Chang, Y. Zhang, Z. Fang, Y.N. Wu, Y. Bisk, **F. Gao**
  The 18th European Conference on Computer Vision (ECCV 2024)
- VR-GS: A Physical Dynamics-Aware Interactive Gaussian Splatting System in Virtual Reality
  Y. Jiang, C. Yu, T. Xie, Y. Feng, H. Wang, M. Li, H. Lau, **F. Gao**, Y. Yang, C. Jiang
  ACM SIGGRAPH 2024

- Learning non-Markovian Decision-Making from State-only Sequences
  A. Qin, **F. Gao**, Q. Li, S.-C. Zhu, S. Xie
  37rd Conference on Neural Information Processing Systems (NeurIPS 2023)
- Masked Path Modeling for Vision-and-Language Navigation
  Z. Dou, **F Gao**, Nanyun Peng
  The 2023 Conference on Empirical Methods in Natural Language Processing 2023 (EMNLP 2023)
- GIVL: Improving Geographical Inclusivity of Vision-and-Language Models with Pre-Training Methods
  D. Yin, **F Gao**, G. Thattai, M. Johnston, K.W. Chang
  Conference on Computer Vision and Pattern Recognition 2023 (CVPR 2023)
- Transform-Retrieve-Generate: Natural Language-Centric Outside-Knowledge Visual Question Answering
  F. Gao, Qing Ping, Govind Thattai, Aishwarya Reganti, Ying Nian Wu, Prem Natarajan
  Conference on Computer Vision and Pattern Recognition 2022 (CVPR 2022)
- Learning Perceptual Inference by Contrasting
  C. Zhang, B. Jia, **F. Gao**, Y. Zhu, H. Lu, S.-C. Zhu
  33rd Conference on Neural Information Processing Systems (NeurIPS 2019, spotlight)
- VRGym: A Virtual Testbed for Physical and Interactive AI (Best Paper Award)
  X. Xie, H. Liu, Z. Zhang, Y. Qiu, **F. Gao**, S. Qi, Y. Zhu, S.-C Zhu
  Association for Computing Machinery Turing Celebration Conference (ACM TURC 2019)
- RAVEN: A Dataset for Relational and Analogical Visual rEasoNing
  C. Zhang*, **F. Gao***, B. Jia, Y. Zhu, S.-C. Zhu (* Joint First Authors)
  Conference on Computer Vision and Pattern Recognition 2019 (CVPR 2019)
- Unsupervised Learning of Hierarchical Models for Hand-Object Interactions using Tactile Glove
  X.Xie, H.Liu, M.Edmonds, **F. Gao**, S.Qi, Y.Zhu, B.Rothrock, S.-C. Zhu
  IEEE International Conference on Robotics and Automation 2018 (ICRA 2018)
- Feeling the Force: Integrating Force and Pose for Fluent Discovery through Imitation Learning to Open Medicine Bottles
  M. Edmonds*, **F. Gao***, X. Xie, H. Liu, S. Qi, Y. Zhu, B. Rothrock, S.-C. Zhu (* Joint First Authors)
  30th International Conference on Intelligent Robots and Systems (IROS 2017)
- A Glove-based System for Studying Hand-Object Manipulation via Pose and Force Sensing
  H. Liu, X. Xie, M. Millar, M. Edmonds, **F. Gao**, Y. Zhu, V. J. Santos, B. Rothrock, S.-C. Zhu
  30th International Conference on Intelligent Robots and Systems (IROS 2017)

PRE-PRINTS

- Atlas3D: Physically Constrained Self-Supporting Text-to-3D for Simulation and Fabrication
  Y. Chen, T. Xie, Z. Zong, X. Li, **F. Gao**, Y. Yang, Y.N. Wu, C. Jiang
  arxiv:2405.18515
- GarmentDreamer: 3DGS Guided Garment Synthesis with Diverse Geometry and Texture Details
  B. Li*, X. Li*, Y. Jiang*, T. Xie, **F. Gao**, H. Wang, Y. Yang, C. Jiang
  arxiv:2405.12420
- Planning as In-Painting:
  A Diffusion-Based Embodied Task Planning Framework for Environments under Uncertainty
  C. Yang, T. Wu, X. Gao, K.W. Chang, **F. Gao**
  arxiv:2312.01097
- TPA-Net: Generate A Dataset for Text to Physics-based Animation
  Y. Qiu, **F. Gao**, M. Li, G. Thattai, Y. Yang, C. Jiang
  arXiv:2211.13887

WORKSHOPS

- Towards Reasoning-Aware Explainable VQA
  R. Vaideeswaran, **F. Gao**, A. Mathur, G. Thattai
  36rd Conference on Neural Information Processing Systems, TSRML workshop (NeurIPS 2022 TSRML)

PRESENTATIONS **Oral Presentations**

- Feeling the Force: Integrating Force and Pose for Fluent Discovery through Imitation Learning to OpenMedicine Bottles
  IROS 2017, Vancouver, Canada

Poster Presentations

- GIVL: Improving Geographical Inclusivity of Vision-and-Language Models with Pre-Training Methods
  CVPR 2023, Virtual Presentation, Vancouver, Canada
- Transform-Retrieve-Generate: Natural Language-Centric Outside-Knowledge
  CVPR 2022, New Orleans, USA
- RAVEN: A Dataset for Relational and Analogical Visual rEasoNing
  MURI 2019, Edinburgh, UK
  CVPR 2019, Long Beach, USA

AWARDS & SCHOLARSHIPS

- Doctoral Fellowship, UCLA — 2020 – 2022
- Outstanding Reviewer, CVPR — 2019
  For reviewers contributed at least two reviews noted as excellent by area chairs
- Doctoral Student Travel Grants, UCLA — 2017 – 2022
- First Class People's Scholarships, UESTC — 2012 – 2014
  For Top 5% students in their major
- Honor Award of Graduation, UESTC — Jun 2015
  For student who got top graduate school offers

PROFESSIONAL APPOINTMENTS & SERVICES

Conference Reviewer

- Reviewer, CVPR — 2019-2021, 2023-2024
- Reviewer, ICLR — 2022
- Reviewer, NeurIPS Dataset Track — 2021
- Reviewer, NeurIPS — 2020-2022
- Reviewer, ECCV — 2020
- Reviewer, AAAI — 2020, 2021
- Reviewer, ICCV — 2019
- Reviewer, ICRA — 2018

Conference Organization

- Student Organizer, MURI Annual Review Meeting, UCLA Aug 2017

Program Reviewer

- Reviewer, Fall 2018,2019,2020 UCLA Computer Science Master's Program, UCLA
- Reviewer, 2018 Cross-disciplinary Scholars in Science and Technology program (CSST), UCLA

Teaching Service

- Teaching Associate, Stats425, UCLA — 2022 Winter
- Teaching Assistant, Stats20, Stats100B, Stats102C, UCLA — 2020 Fall, 2021 Winter, 2021 Spring
- Reader, Statistics Department, UCLA — 2017 Fall, 2021 Fall, 2022 Winter
- Teaching Assistant, C Programming, UESTC — 2014 Fall

PROFESSIONAL SKILLS

Programming Language
Python, C++, Matlab

Frameworks & Softwares
PyTorch, Tensorflow, ROS, Matlab