

# Irena Gao

i-gao.github.io · irena@cs.stanford.edu

## Education

**Ph.D. in Computer Science**, *Stanford University*. 2023–present

Rotation advisors: Carlos Guestrin

**M.S. in Statistics**, *Stanford University*. GPA: 4.03 2021–2023

**B.S. in Computer Science with Honors and Distinction**, *Stanford University*. GPA: 4.06 2018–2023

Thesis advisor: Percy Liang

## Work Experience

Research Intern, **Microsoft Research** 2022

Supervisor: Marco Tulio Ribeiro

## Publications

\* denotes equal contribution.

- [1] Anas Awadalla\*, Irena Gao\*, Joshua Gardner, Jack Hessel, Yusuf Hanafy, Wanrong Zhu, Kalyani Marathe, Yonatan Bitton, Samir Gadre, Shiori Sagawa, Jenia Jitsev, Simon Kornblith, Pang Wei Koh, Gabriel Ilharco, Mitchell Wortsman, and Ludwig Schmidt. Openflamingo: An open-source framework for training large autoregressive vision-language models, 2023.
- [2] Irena Gao, Gabriel Ilharco, Scott Lundberg, and Marco Tulio Ribeiro. Adaptive testing of computer vision models. In *International Conference on Computer Vision*, 2023. **Oral presentation (top 5.9% of accepted papers)**.
- [3] Irena Gao\*, Shiori Sagawa\*, Pang Wei Koh, Tatsunori Hashimoto, and Percy Liang. Out-of-domain robustness via targeted augmentations. In *International Conference on Machine Learning*, 2023
- [4] Nicholas Carlini, Milad Nasr, Christopher A Choquette-Choo, Matthew Jagielski, Irena Gao, Anas Awadalla, Pang Wei Koh, Daphne Ippolito, Katherine Lee, Florian Tramèr, and Ludwig Schmidt. Are aligned neural networks adversarially aligned? In *Neural Information Processing Systems*, 2023.
- [5] Zixian Ma\*, Jerry Hong\*, Mustafa Omer Gul\*, Mona Gandhi, Irena Gao, and Ranjay Krishna. Crepe: Can vision-language foundation models reason compositionally? In *Conference on Computer Vision and Pattern Recognition*, 2023. **Highlight (top 10% of accepted papers)**.
- [6] Shiori Sagawa\*, Pang Wei Koh\*, Tony Lee\*, Irena Gao\*, Sang Michael Xie, Kendrick Shen, Ananya Kumar, Weihua Hu, Michihiro Yasunaga, Henrik Marklund, Sara Beery, Etienne David, Ian Stavness, Wei Guo, Jure Leskovec, Kate Saenko, Tatsunori Hashimoto, Sergey Levine, Chelsea Finn, and Percy Liang. Extending the wilds benchmark for unsupervised adaptation. In *International Conference on Learning Representations*, 2022. **Oral presentation**.
- [7] Pang Wei Koh\*, Shiori Sagawa\*, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanus Phillips, Irena Gao, Tony Lee, Etienne David, Ian Stavness, Wei Guo, Berton Earnshaw, Imran Haque, Sara Beery, Jure Leskovec, Anshul Kundaje, Emma Pierson, Sergey Levine, Chelsea Finn, and Percy Liang. Wilds: A benchmark of in-the-wild distribution shifts. In *International Conference on Machine Learning*, pp. 5637–5664. PMLR, 2021. **Oral presentation**.
- [8] Mark E Whiting, Irena Gao\*, Michelle Xing\*, N’godjigui Junior Diarrassouba\*, Tonya Nguyen\*, and Michael S Bernstein. Parallel worlds: Repeated initializations of the same team to improve team viability. *proceedings of the ACM on Human-Computer Interaction*, 4(CSCW1):1–22, 2020.

## Service

Conferences and workshops

Reviewer, NeurIPS Workshop on Distribution Shifts 2023

Reviewer, NeurIPS 2023

Reviewer, CVPR 2023

Reviewer, ICML Principles of Distribution Shifts Workshop 2022

Reviewer, NeurIPS Workshop on Distribution Shifts 2022

Reviewer, NeurIPS Workshop on Distribution Shifts 2021

Community

Mentor, Stanford CS Mentoring Program 2023

Mentor, Women’s Community Center STEM Mentorship Program 2023

## Teaching

Course Assistant, CS221: <i>Artificial Intelligence: Principles and Techniques</i>	2023
Course Assistant, CS224N: <i>Natural Language Processing with Deep Learning</i>	2023
Course Assistant, CS221: <i>Artificial Intelligence: Principles and Techniques</i>	2022
Section Leader, CS106A/B: <i>Programming Methods &amp; Abstractions</i>	2019–2020