

Zixian Ma

zixianma@cs.washington.edu | <https://zixianma.github.io>

EDUCATION

University of Washington

PhD in Computer Science and Engineering

9/2023 - now

Stanford University

BS with Honors and MS in Computer Science; Minor in Biology; GPA: 3.97

9/2018 – 6/2022

Courses: Artificial Intelligence: Principles and Techniques, Machine Learning, Natural Language Understanding, From Languages to Information, Convolutional Neural Network for Visual Recognition, Reinforcement Learning, Machine Learning with Graphs, Introduction to Human-Computer Interaction Design, Web Applications, Virtual People

PUBLICATIONS

Robin: Enhanced Visual Relationship Reasoning via Scene Graph Distillation

2024

Jae Sung Park, **Zixian Ma**, Linjie Li, Khyathi Chandu, Ximing Lu, Ali Farhadi, Yejin Choi, Ranjay Krishna
In submission

Task Me Anything

2024

Jieyu Zhang, Weikai Huang*, **Zixian Ma***, Oscar Michel, Dong He, Tanmay Gupta, Wei-Chiu Ma, Ali Farhadi, Aniruddha Kembhavi, Ranjay Krishna
The Thirty-Eighth Annual Conference on Neural Information Processing Systems (NeurIPS 2024)

NaturalBench: Evaluating Vision-Language Models on Natural Adversarial Samples

2024

Baiqi Li*, Zhiqiu Lin*, Wenxuan Peng*, Jean de Dieu Nyandwi*, Daniel Jiang, **Zixian Ma**, Simran Khanuja, Ranjay Krishna, Graham Neubig, Deva Ramanan
The Thirty-Eighth Annual Conference on Neural Information Processing Systems (NeurIPS 2024)

m&m's: A Benchmark to Evaluate Tool-Use for multi-step multi-modal Tasks

2024

Zixian Ma, Weikai Huang, Jieyu Zhang, Tanmay Gupta, Ranjay Krishna
The 18th European Conference on Computer Vision (ECCV 2024)

SugarCREPE: Fixing Hackable Benchmarks for Vision-Language Compositionality

2023

Cheng-Yu Hsieh*, Jieyu Zhang*, **Zixian Ma**, Aniruddha Kembhavi, Ranjay Krishna
The Thirty-Seventh Annual Conference on Neural Information Processing Systems (NeurIPS 2023)

CREPE: Can Foundation Vision-Language Models Reason Compositionally?

2023

Zixian Ma*, Jerry Hong*, Mustafa Omer Gul*, Mona Gandhi, Irena Gao, Ranjay Krishna
The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR 2023) (Highlight)

ELIGN: Expectation Alignment as a Multi-Agent Intrinsic Reward

2022

Zixian Ma, Rose E. Wang, Li Fei-Fei, Michael Bernstein, Ranjay Krishna
The Thirty-Sixth Annual Conference on Neural Information Processing Systems (NeurIPS 2022)

Model Sketching: Centering Concepts in Early-Stage Machine Learning Model Design

2022

Michelle Lam, **Zixian Ma**, Anne Li, Izequiel Freitas, Dakuo Wang, James Landay, Michael Bernstein
The ACM CHI Conference on Human Factors in Computing Systems (CHI 2023)

OpenAttack: An Open-source Textual Adversarial Attack Toolkit

2021

Guoyang Zeng, Fanchao Qi, Qianrui Zhou, Tingji Zhang, **Zixian Ma**, Bairu Hou, Yuan Zang, Zhiyuan Liu, Maosong Sun
The Joint Conference of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (ACL 2021): System Demonstrations

RESEARCH EXPERIENCE**Stanford Vision and Learning Lab**

4/2020 – 6/2022

Mentors: Prof. Li Fei-Fei, Dr. Ranjay Krishna (now Prof. Ranjay Krishna at University of Washington)

- Led an independent multi-agent collaboration project under the multi-agent reinforcement learning framework
- Formulated and implemented a novel multi-agent intrinsic reward ELIGN that incentivizes expectation alignment
- Conducted extensive experiments across 6 collaborative and competitive tasks in 2 multi-agent environments
- Wrote a full paper on the multi-agent intrinsic reward ELIGN, which got accepted to NeurIPS 2022

Stanford Human-Computer Interaction Group

4/2022 – 9/2022

Mentors: Prof. Michael Bernstein, Prof. James Landay

- Worked on the ModelSketching project advised by Profs. Michael Bernstein and James Landay (accepted to CHI 23)
- Wrote Python functions in the ModelSketchBook API for users to compare multimodal concepts with CLIP
- Evaluated the utility of model sketches in surfacing reviewers' biases on a manually collected food reviews dataset
- Conducted pilot and final user studies on the hateful memes detection task

WORK EXPERIENCE**Salesforce AI Research – Research Intern**

6/2024 – 9/2024

- Training multi-modal agents for vision-centric tasks

Meta – Software Engineer

12/2022 – 05/2023

- Improved the feed ranking algorithm for the Stories-in-Feed product

Google Research – Research Intern

8/2022 – 11/2022

- Evaluated LaMDA (128B) and PaLM (540B) on screen navigation tasks with various prompts
- Implemented prompt tuning and visual prefix tuning on top of transformers-based language models
- Finetuned LaMDA with prompt-tuning techniques and vision-language model CoCa on the MoTIF dataset

Facebook (currently Meta) – Software Engineering Intern

6/2021 – 8/2021

- Built and optimized multi-task multi-label models for stories ranking with multi-gate mixture of experts module
- Launched the models into production and reduced the company's multi-feed CPU usage by 0.6% (\$191,746)

AWARDS**Research***The Firestone Medal for Excellence in Undergraduate Research*

2022

- Awarded to the top 10 percent of Honors Theses across all disciplines and schools at Stanford each year

The Ben Wegbreit Prize for Undergraduate CS Research (The CS Department's Best Honors Thesis award)

2022

LEADERSHIP & EXTRACURRICULAR ACTIVITIES*Head Academic Lead @ BioX (a summer camp for high school students based in Shanghai, China)*

2019 – 2021

- TAed 24 high school students in Introduction to Bioinformatics and Computational Biology in 2019; designed and led the same course with three professors and three other student TAs in 2020; Co-organized BioX in 2021

Member @ She++

2018 - 2019

- Matched marginalized high school students, especially girls, to mentors in CS; organized events at the annual summit
-

SKILLS

Technical

Python, Pytorch, TensorFlow, C++, C, R, HTML&CSS, JavaScript, React, SQL

Language

English, Chinese (Mandarin and Cantonese)