

# ZOEY HUANG

(+1)6508851969 ◊ huangzhuoyicn@gmail.com

<https://www.linkedin.com/in/zhuoyi-huang/>

## EDUCATION

---

### Stanford University

M.S. in Computer Science, Artificial Intelligence track (fully funded)

GPA: 4.1 / 4.0

Core Courses: CS 231A, Computer Vision: From 3D Reconstruction to Recognition, EE 277 & CS 234, Reinforcement Learning, CS 248 & CS 348, Interactive Computer Graphics, CS 224N & CS 224U, Natural Language Processing, CS 330, Multi-Task and Meta-Learning, CS 224B, Distributed Systems, etc.

Sep. 2021 - Jun. 2023

Palo Alto, USA

### Tongji University

B.S. in Computer Science, Information Security track

GPA: 94 / 100 (1/189)

Core Courses: Probability Theory & Statistics, Number theory, Discrete Mathematics, Database System, Operation System, Algorithm Analysis and Design, Pattern Recognition, Cryptography, Machine Learning, etc.

Sep. 2016 - Jul. 2020

Shanghai, China

## INDUSTRIAL EXPERIENCE

---

### xAI, RL Science & Coding RL recipe

Member of Technical Staff

Jan. 2026 - Now

Palo Alto, USA

- Work on latest large scale RL stability and sample efficiency recipe for coding models, achieved SOTA internal eval reward.
- Grok 4.2 stem task RL bigrun self-play RL recipe, surpassing RLVR recipe.

### Microsoft, Turing, MSAI

Senior Applied Scientist

May. 2025 - Jan. 2026

Mountain View & Redmond, USA

- Large scale RL reasoning model (OpenAI GPT-5.1 checkpoint) for complex long horizon tool use (code, canvas) and Computer Use Agent, with 1k nodes and 2k VMs for code execution, office operations and web browsing.

### Microsoft MAIDAP, Turing

Applied Scientist

July. 2023 - May. 2025

Mountain View & New England Research Center, USA

- Online policy distilled Phi-3.5 to enhance enterprise search relevance labeling, allowing automatic, scalable augmentation to noisy click data with high-quality relevance grades. Achieved comparable performance to OpenAI model with a 964% increase in speed (188 RPM to 2000 RPM) and a 99.9% reduction in cost, lowering expenses from \$10,000 to just \$12.50 for 1 million queries. ([Paper Fine-tuning Small Language Models as Efficient Search Relevance Labelers published internally Microsoft Machine Learning, AI & Data Science Conference 2025, Annual Best Paper](#))
- Multi-modal medical data generative models for electronic health records and time-series data via [conditional generative Latent Diffusion Models](#). Collaborated with (External Partner: MIT, Harvard, Duke, Mass General Brigham), and funded by the NIH. ([Learning ECG Representations via Poly-Window Contrastive Learning BHI 2025](#), [High-Fidelity Synthetic ECG Generation via Mel-Spectrogram Informed Diffusion Training NeurIPS GenAI4H 2025 Best Research Paper](#))
- Working with Office AI team on training generative models to generate personalized online hand-writing stroke sequences. ([TrInk: Ink Generation with Transformer Network EMNLP 2025](#))

### Meta, Creation ML Team

Software Engineering Intern

Jun. 2022 - Sept. 2022

Menlo Park, USA

- Developed end-to-end ML systems for Facebook Stories Video Generation, using instance segmentation to create dynamic photo-coloring video effects and sync with the soundtrack.
- Implemented [two variants](#) of the video effects. The OSPED (Original Exclusively Story Producer DAP) achieved a [0.032% gain over a 3-week](#) public test. The created videos were viewed by over [200k users](#), and the view-to-share rate reached approximately [3.5%](#), which was the [highest](#) among all current pre-creation effects on Facebook Stories at that time.

### Microsoft Research Asia, Visual Computing Group

Supervised by [Dr. Jingdong Wang](#)

Oct. 2020 - Feb. 2021

Beijing, China

- Refined the neural network structure of one-stage object detection algorithm by utilizing the backbone network HRNetV2 and referring to the prediction head of FCOS on mmdetection codebase.
- Implemented disentangled prediction head with Deformable Convolutions and Spatial Transformation Convolutions that focus on four boundary extremities features of detection boxes. The ultimate network structure got around 1.5 AP gain over traditional single head FCOS under same training schedules in ablation experiments.

## RESEARCH & PUBLICATIONS

---

### Long-Horizon Causal Inference Through Multi-Modal Simulation

Supervised by Prof. Jiajun Wu and Prof. Tobias Gerstenberg

Jan. 2023 - Now

Stanford University, Stanford Vision & Learning Lab, Stanford Causality in Cognition Lab

Stanford, USA

- **MARPLE: A Benchmark for Long-Horizon Inference.** Zhuoyi Huang, Emily Jin, Jan-Philipp Fränken, Weiyu Liu, Hannah Cha, Erik Brockbank, Sarah Wu, Ruohan Zhang, Jiajun Wu, Tobias Gerstenberg. *NeurIPS D&B track 2024*
- **Whodunnit? Inferring what happened from multimodal evidence.** Sarah A Wu, Erik Brockbank, Hannah Cha, Jan-Philipp Fränken, Emily Jin, Zhuoyi Huang, Weiyu Liu, Ruohan Zhang, Jiajun Wu, Tobias Gerstenberg. *COGSCI 2024*
- **Mini-BEHAVIOR: A Procedurally Generated Benchmark for Long-horizon Decision-Making in Embodied AI.** Emily Jin, Jiaheng Hu, Zhuoyi Huang, Ruohan Zhang, Jiajun Wu, Li Fei-Fei, and Roberto Martín-Martín. *Agent Learning in Open-Endedness 2023.*

### Few-Shot Activity Parsing on MOMA-LARGE

Supervised by Prof. Fei-Fei Li and Prof. Ehsan Adeli

Sep. 2021 - Apr. 2022

Stanford University, Stanford Partnership in AI-Assisted Care

Stanford, USA

- **MOMA-LRG: Language-Refined Graphs for Multi-Object Multi-Actor Activity Parsing.** Zelin Luo, Zane Durante, Linden Li, Wanze Xie, Ruochen Liu, Emily Jin, Zhuoyi Huang, Lun Yu Li, Jiajun Wu, Juan Carlos Niebles, Ehsan Adeli, Li Fei-Fei. *NeurIPS D&B track 2022*

### Bayesian LASSO Thompson Sampling Algorithm in Multi-armed Bandit Problem

Supervised by Prof. Zhengyuan Zhou

Jun. 2019 - Sep. 2019

Stanford University, Information System Laboratory

Stanford, USA

- Improved maximized cumulative rewards in contextual bandit problem in the case of high-dimensional feature vectors and sparse regression coefficients.
- Implemented Bootstrap Lasso LinUCB algorithm which outperforms traditional LinUCB algorithms in the high-dimensional feature settings.
- Proposed Bayesian LASSO Thompson Sampling algorithm to solve contextual multi-bandits problem from a Bayesian statistical view.

## SERVICE EXPERIENCE

---

**Teaching Assistant** at Stanford Computer Science Department for CS231N, Deep learning for Computer Vision; CS131, Computr Vision: Foundations and Applications; CS224W, Machine Learning with Graphs; CS246, Mining Massive Data Sets.

**Peer Reviewer** for MIDL 2024, ICML 2024, ICLR 2025, AAI 2025 workshop GenPlan, ICLR 2026.

## PATENTS

---

**Online Ink Generation Using A Transformer Network** Z. Huang, et al. MS# 505283-US01

**Generating Modified Prompts Based On Feedback** Z. Huang, et al. MS# 502518-US01

**An Emergency Wireless Charging System Based on Tethered UAV.** H. Deng, Z. Huang, Q. Liu, M. Xiong, et al. *National Invention Patent 2019, 201910294599.7.*

**A Targeting and Positioning System based on Segregate Laser Cavity.** H. Deng, Q. Liu, M. Liu, W. Fang, Z. Huang, M. Xiong et al. *National Invention Patent 2019, 201910294606.3.*

## HONORS & AWARDS

---

|  |                                 |
|--|---------------------------------|
| NeurIPS GenAI4H Oral   | 2025                            |
| Microsoft Global Hackathon Hack for Industry-Honorable Mention (Accept ratio: 0.98%) | 2024                            |
| Fusion Fund Venture Fellowship (Accept ratio: 6.25%)                                 | 2023                            |
| Excellent Graduates in Shanghai (Award ratio 0.2%)                                   | 2020                            |
| The National Scholarship for full 3 academic years (Award ratio 0.2%)                | 2016-2017, 2017-2018, 2018-2019 |

Excellent Student, Pacemaker to Excellent Student (Award ratio 0.15%)  
Cyrus Tang Scholarship for Moral Education (Award ratio 0.18%)

*2016-2017, 2017-2018*  
*2017-2018*