Zhuokai Zhao

zhuokai@uchicago.edu | zhuokai-zhao.com | https://github.com/zhuokaizhao | https://www.linkedin.com/in/zhuokai-zhao-a9385169/

Education

University of Chicago — *Chicago, IL*

Ph.D. in Computer Science. Advisor: Prof. Yuxin Chen. Committee: Prof. Michael Maire and Prof. Bo Li — *2018 - 2024 (expected)*

- Concentrations: Data-efficient, Trustworthy, and Robust Multimodal Deep Learning.
- Award: Doctoral Fellowship, Center for Data and Computing (CDAC), Data Science Institute (DSI)

Johns Hopkins University — *Baltimore, MD*

M.S. in Robotics. Advisors: Prof. Nassir Navab and Prof. Russell Taylor. — *2016 - 2018*

University of Illinois at Urbana-Champaign — *Champaign, IL*

B.S. in Electrical Engineering. Advisor: Prof. Seth Hutchinson. — *2013 - 2016*

Work Experience

Twitch, a subsidiary of Amazon — *New York, NY*

Applied Scientist Intern — *Jun, 2023 - Aug, 2023*

- Proposed a multimodal model fusing features from both streaming (vision) and metadata (text) to conduct real-time violent contents moderation.
- Constructed a custom dataset supporting the training and testing of the proposed multimodal approach.

Meta AI — *Menlo Park, CA*

Machine Learning Intern and Part-time Student Researcher — *Jun, 2022 - Nov, 2022*

- Implemented User-Centric Ranking, a new formulation in recommender system that trains better-converged models on substantially larger datasets.

Kitware Inc. — *Clifton Park, NY*

R&D Intern — *Jun, 2019 - Sep, 2019*

- Contributed to open-source libraries including the Visualization Toolkit (VTK), Open Chemistry, and Tomviz.

JD.COM Silicon Valley Research Center — *Mountain View, CA*

Research Intern — *May, 2018 - Sep, 2018*

- Developed body-garment registration (virtual try-on) using both rigid and non-rigid (collision-detection based optimization) point cloud alignments.
- Developed in-house simulation engine based on ARCSim to generate dynamic body-garment animations.

Siemens Corporate Research — *Princeton, NJ*

Research Intern — *May, 2017 - Aug, 2017*

- Developed an algorithm that reconstructs patients in real-time from 2D stream with limited field of view. Integrated AR for interactive visualizations.

Selected Publications (* indicates co-first authorship.)

(1) **Zhuokai Zhao***, Yiming Zhang*, Zhaorun Chen, Zhili Feng, Zenghui Ding, and Yining Sun. RankCLIP: Ranking-Consistent Language-Image Pre-training. In *submission*, 2024

(2) **Zhuokai Zhao***, Zhaorun Chen*, Wenjie Qu, Zichen Wen, Zhiguang Han, Zhihong Zhu, Jiaheng Zhang, and Huaxiu Yao. PANDORA: Detailed LLM Jailbreaking via Collaborated Phishing Agents with Decomposed Reasoning. In *In submission. Preliminary version appeared in ICLR Workshop on Secure and Trustworthy Large Language Models*, 2024

(3) Zhaorun Chen, **Zhuokai Zhao**, Dandan Liang, Zhihong Zhu, Liang Gong, Chengliang Liu, Chenglin Miao, and Lu Su. EscIRL: Evolving Self-Contrastive IRL for Trajectory Prediction in Autonomous Driving. In *submission*, 2024

(4) Zhaorun Chen*, **Zhuokai Zhao***, Hongyin Luo, Huaxiu Yao, Bo Li, and Jiawei Zhou. HALC: Object Hallucination Reduction via Adaptive Focal-Contrast Decoding. In *submission. Preliminary version appeared in ICLR Workshop on Reliable and Responsible Foundation Models*, 2024

(5) **Zhuokai Zhao**, Yibo Jiang, and Yuxin Chen. Direct Acquisition Optimization for Low-Budget Active Learning. In *submission*, 2024

(6) Zhaorun Chen*, **Zhuokai Zhao***, Zhihong Zhu*, Ruiqi Zhang, Xiang Li, Bhiksha Raj, and Huaxiu Yao. AutoPRM: Automating Procedural Supervision for Multi-Step Reasoning via Controllable Question Decomposition. In *Annual Conference of the NAACL*, 2024

(7) Zhaorun Chen, **Zhuokai Zhao**, Tairan He, Binhao Chen, Xuhao Zhao, Liang Gong, and Chengliang Liu. Safe Reinforcement Learning via Hierarchical Adaptive Chance-Constraint Safeguards. In *submission*, 2024

(8) **Zhuokai Zhao**, Harish Palani, Tianyi Liu, Lena Evans, and Ruth Toner. Multi-Modality Guidance Network for Missing Modality Inference in Content Moderation. In *submission*, 2023

(9) **Zhuokai Zhao**, Yang Yang, Wenjie Hu, and Shuang Yang. Breaking the Curse of Quality Saturation with User-Centric Ranking. In *29th Conference on Knowledge Discovery and Data Mining (KDD)*, 2023

(10) **Zhuokai Zhao**, Takumi Matsuzawa, William Irvine, Michael Maire, and Gordon L Kindlmann. Evaluating Machine Learning Models with NERO: Non-Equivariance Revealed on Orbits. 2023

Skills

Programming Python (PyTorch, TensorFlow, NumPy, PySide, and etc.), C/C++, SQL (Redshift, Presto), JavaScript (Svelte, React).

Miscellaneous Linux, Git, LaTeX, Shell (Bash/Zsh), CMAKE.