Yasi Zhang

+1 424 977 9466 | yasminzhang@ucla.edu | LinkedIn | Google Scholar | Website | Los Angeles, USA

Education

University of California, Los Angeles — Los Angeles, USA
PhD student in Statistics and Data Science — *Sep 2022 – Sep 2026 (expected)*

- **PhD Advisors:** Prof. Ying Nian Wu and Prof. Oscar Leong
- **Research Interests:** Generative AI, Multimodality, Representation Learning, Reinforcement Learning
- **Tools and Programming:** Python, PyTorch, Jax, Flax, Git, LaTeX, Linux

Fudan University — Shanghai, China
B.S. Data Science and Big Data Technology; First-Class Honor — *Sep 2018 – Jun 2022*

Personal Statement

I enjoy research from both industry and academic perspectives, having gained multiple experiences by collaborating and interning at a range of leading tech companies. My current research goal is to build a unified foundation for powerful multimodal generative models capable of understanding, reasoning, and generating high-dimensional data across diverse modalities.

Selected Industry Experiences

ByteDance Seed, Student Researcher — *Mar 2026 – Present*
Manager: Dr. Peng Wang

- Worked on the scaling study of multimodal representation learning in the pre-training stage of Omni models.

Google Research, Student Researcher — *Mar 2025 – Sep 2025*
Manager: Dr. Junfeng He, Co-mentor: Dr. Michal Lukasik

- Worked on post-training and reward modeling for text-to-image/video/design/editing generation.
- First-authored one paper on text-to-image generation currently under internal embargo-period strategic review.
- Contributed to Gemini 3 Flash in image understanding involving visual common sense.
- First-authored one paper on regression-aware reinforcement learning currently under review.

Amazon AWS AI Labs, Applied Scientist Intern — *Jun 2024 – Sep 2024*
Mentor: Dr. Jiarui Cai

- Worked on LLM-enhanced synthetic data generation to boost object detection with Grounding-DINO.

Selected Research Experiences

Reinforcement Learning for Multi-Turn Image Editing — UCLA×Apple AIML
Advisor: Dr. Ying Nian Wu [Under Review] — *Dec 2025 – Mar 2026*

- Developed MT-EditFlow, a flow-matching reinforcement learning framework that optimizes reward signals specifically for multi-turn image editing.
- Mitigated compounding errors and exposure bias inherent in iterative editing processes and boosted the FLUX.1-Kontext-dev model's turn-3 performance by 6.85 points, and FLUX.2-Klein's performance by 2.89 points.

Regression-Aware Reinforcement Learning for LLM-as-a-Judge — UCLA×Google Research
Advisors: Dr. Michal Lukasik [Under Review] — *Sep 2025 – Jan 2026*

- Proposed REAL, a regression-aware RL framework for LLM-as-a-Judge, employing the generalized policy gradient to maximize regression rewards by jointly optimizing Chain-of-Thought exploration and numeric score prediction.
- Achieved SOTA correlation performance across multiple evaluation benchmarks by training Mistral2-7B, Qwen3-8B, and Qwen3-32B with Verl, demonstrating strong out-of-domain generalization.

EdiVal-Agent for Multi-Turn Image Editing — UCLA×UT-Austin×Microsoft AI
Advisors: Dr. Lijuan Wang [ICLR 2026] [Chinese Blog] — *May 2025 – Sep 2025*

- Developed EdiVal-Agent, an object-centric evaluation framework for multi-turn image editing, using expert tools to assess instruction following, content consistency, and visual quality.
- Built EdiVal-Bench, covering 9 instruction types and 11+ SOTA editing models, outperforming zero-shot VLMs in alignment with human judgment and identifying key failure modes like exposure bias.

Restoration Score Distillation for One-Step High-Quality Generation — UCLA×Microsoft AI
Advisors: Dr. Mingyuan Zhou, Dr. Ying Nian Wu, and Dr. Oscar Leong [ICLR 2026] — *Dec 2024 – May 2025*

- Introduced Restoration Score Distillation, a two-phase framework that trains one-step generators directly from corrupted data by theoretically capturing the clean distribution's eigenspace.
- Outperformed multi-step diffusion models on denoising, super-resolution, deblurring, inpainting, and multi-coil MRI in sample quality, achieving up to 30× faster sampling without clean or few-shot data.

Flow Matching Priors for Linear Inverse Image Problems UCLA×Amazon

Advisors: Dr. Oscar Leong and Dr. Ying Nian Wu [NeurIPS 2024] *Dec 2023 – May 2024*

- Developed an iterative algorithm with theoretical justification to efficiently approximate the MAP estimator for linear inverse problems using flow matching priors.
- Validated the approach through experiments on super-resolution, deblurring, inpainting, and compressed sensing, outperforming existing flow-based methods.

Attention Map Alignment with Text-to-Image Diffusion Models UCLA

Advisor: Dr. Ying Nian Wu [ECCV 2024] *Jul 2023 – Nov 2023*

- Proposed an object-conditioned energy-based attention alignment method with intensity regularization to mitigate object neglect and attribute mismatch in text-to-image diffusion models.
- Achieved top performance in multi-object and attribute alignment on both automated metrics and human evaluation, enhancing cross-attention editing with more semantically meaningful attention maps.

Seleted Publications

1. Jiahui Huang*, **Yasi Zhang***, Tianyu Chen, Shu Wang, Jianwen Xie, Oscar Leong, Mingyuan Zhou, Nanzhu Wang, Ying Nian Wu *MT-EditFlow: Reinforcement Learning for Multi-Turn Image Editing with Flow Matching.* Under Submission.
2. **Yasi Zhang***, Tianyu Chen*, Mingyuan Zhou, Oscar Leong, Ying Nian Wu, Michal Lukasik. *REAL: Regression-Aware Reinforcement Learning for LLM-as-a-Judge.* Under Submission. [Paper]
3. Tianyu Chen*, **Yasi Zhang***, Zhi Zhang, Peiyu Yu, Shu Wang, Zhendong Wang, Kevin Lin, Xiaofei Wang, Zhengyuan Yang, Linjie Li, Cheng-Han Chiang, Jianwen Xie, Oscar Leong†, Lijuan Wang†, Ying Nian Wu†, Mingyuan Zhou†. *EdiVal-Agent: An Object-Centric Framework for Automated, Fine-Grained Evaluation of Multi-Turn Editing.* **International Conference on Learning Representations (ICLR) 2026**. [Paper] [Website]
4. **Yasi Zhang***, Tianyu Chen*, Zhendong Wang, Ying Nian Wu, Mingyuan Zhou†, Oscar Leong†. *Score Distillation Beyond Acceleration: Generative Modeling from Corrupted Data.* **International Conference on Learning Representations (ICLR) 2026**. [Paper]
5. Peiyu Yu, Dinghuai Zhang, Hengzhi He, Xiaojian Ma, Ruiyao Miao, Yifan Lu, **Yasi Zhang**, Deqian Kong, Ruiqi Gao, Jianwen Xie, Guang Cheng, Ying Nian Wu. *Latent Energy-Based Odyssey: Black-Box Optimization via Expanded Exploration in the Energy-Based Latent Space.* **International Conference on Learning Representations (ICLR) 2026**. [Paper]
6. Johannes Hertrich, Hok Shing Wong, Alexander Denker, Stanislas Ducotterd, Zhenghan Fang, Markus Haltmeier, Željko Kereta, Erich Kobler, Oscar Leong, Mohammad Sadegh Salehi, Carola-Bibiane Schönlieb, Johannes Schwab, Zakhar Shumaylov, Jeremias Sulam, German Shâma Wache, Martin Zach, **Yasi Zhang**, Matthias J. Ehrhardt, Sebastian Neumayer. *Learning Regularization Functionals: A Comparative Study.* Book chapter under review by Springer. [Book Chapter]
7. **Yasi Zhang**, Oscar Leong. *Learning Difference-of-Convex Regularizers for Inverse Problems: A Flexible Framework with Theoretical Guarantees.* Under submission. [Paper]
8. Zhi Zhang, Chris Chow, **Yasi Zhang**, Yanchao Sun, Haochen Zhang, Eric Hanchen Jiang, Han Liu, Furong Huang, Yuchen Cui, Oscar Hernan Madrid Padilla. *Statistical Guarantees for Lifelong Reinforcement Learning using PAC-Bayesian Theory.* **International Conference on Artificial Intelligence and Statistics (AISTATS) 2025**. [Paper]
9. **Yasi Zhang**, Peiyu Yu, Yaxuan Zhu, Yingshan Chang, Feng Gao, Ying Nian Wu, Oscar Leong. *Flow Priors for Linear Inverse Problems via Iterative Corrupted Trajectory Matching.* **Conference on Neural Information Processing Systems (NeurIPS) 2024**. [Paper] [Code]
10. **Yasi Zhang**, Peiyu Yu, Ying Nian Wu. *Object-Conditioned Energy-Based Attention Map Alignment in Text-to-Image Diffusion Models.* **European Conference on Computer Vision (ECCV) 2024**. [Paper] [Code]
11. Yingshan Chang, **Yasi Zhang**, Jacob Zhiyuan Fang, Ying Nian Wu, Yonatan Bisk, Feng Gao. *Skews in the Phenomenon Space Hinder Generalization in Text-to-Image Generation.* **European Conference on Computer Vision (ECCV) 2024**. [Paper] [Code]