

Xi Xiao

+1-659-232-2378 | xxiao@uab.edu | [Homepage](#)
 [Google Scholar](#) |  [LinkedIn](#) |  [GitHub](#)
Birmingham, AL 35243, United States

RESEARCH INTERESTS

- My research primarily focuses on the **visual and multimodal reasoning / understanding, post-training with reinforcement learning, and post-training for personalized intelligence.**

EDUCATION

- **University of Alabama at Birmingham** Jan 2024 – Dec 2027 (Expected)
Ph.D. Student in Computer Science Birmingham, United States
- **Sichuan University Jincheng College** Aug 2019 – Jun 2023
B.Eng. in Artificial Intelligence Chengdu, China

PROFESSIONAL EXPERIENCE

- **Oak Ridge National Laboratory**  Aug 2026 – Present
Research Scientist Intern (Advisor: [Dr. Xiao Wang](#)) Oak Ridge, United States
 - Co-Lead the **AI4Science Foundation Model Team**, overseeing **foundation model training, architecture design, and parallel training optimization.**
 - Leading the design of a scalable **Vision Foundation Model architecture** and the pretraining on the Frontier exascale supercomputer.
- **Amazon AGI Foundations**  May 2026 – Aug 2026
Applied Scientist Intern (Advisor: [Dr. Eric Zhao](#) and [Dr. Youngeun Kim](#)) Seattle, United States
 - Developed **reinforcement learning-based post-training methods for Amazon Nova multimodal foundation models**, targeting stronger **MLLM and agentic reasoning** capabilities.
 - Built scalable **training and evaluation pipelines** for multimodal reasoning and long-horizon decision-making; conducted systematic ablations and failure analyses to improve reasoning quality, robustness, and reliability.
 - **Outcome:** One manuscript submitted to ICLR 2027.
- **Oak Ridge National Laboratory**  May 2025 – May 2026
Research Scientist Intern (Advisor: [Dr. Xiao Wang](#)) Oak Ridge, United States
 - Conducted **pretraining of a tens-of-billions parameter foundation model (ORBIT-2)** on high-dimensional nonlinear climate datasets using the Frontier exascale supercomputer.
 - Optimized large-scale distributed training with **FSDP, DDP, and TILLING**, achieving significant improvements in training scalability and efficiency.
 - Applied **parameter-efficient fine-tuning** for downstream adaptation, enabling efficient adaptation of the foundation model to climate impact tasks.
 - **Outcomes:** One ECCV 2026, One ICME 2026, and One SC 2025. ([Best Paper Award, ACM Gordon Bell Prize Finalist](#)).

SELECTED PUBLICATIONS

PLEASE SEE MY GOOGLE SCHOLAR FOR FULL PUBLICATIONS

- **[ECCV 2026] Staying VIGILant: Mitigating Visual Laziness via Counterfactual Visual Alignment in MLLMs**
[MLLM Post-Training](#) [Model Alignment](#) [MLLM Reasoning](#)
[Xi Xiao](#), Chen Liu, Chih-Ting Liao, Yunbei Zhang, Qizhen Lan, Yuxiang Wei, Lin Zhao, Janet Wang, Jianyang Gu, Muchao Ye, Tianyang Wang, Hao Xu.
Project: [VIGIL](#)
- **[ECCV 2026] Layer-Specific Prompt Fusion Discovery via Differentiable Search in Vision Foundation Models**
[Vision Foundation Models](#) [Foundation Model Adaptation](#) [Differentiable Search](#)
[Xi Xiao](#), Xingjian Li, Yunbei Zhang, Cheng Han, Tianming Liu, Tianyang Wang, Runmin Jiang, Jihun Hamm, Xiao Wang, Min Xu.
Project: [Prompt-Fusion-Discovery](#)
- **[ACL 2026] Not All Directions Matter: Towards Structured and Task-Aware Low-Rank Model Adaptation**
[LLM Adaptation](#) [LoRA & PEFT](#) [Structured Rank Allocation](#)
[Xi Xiao](#), Chenrui Ma, Yunbei Zhang, Chen Liu, Zhuxuanzi Wang, Yanshu Li, Lin Zhao, Guosheng Hu, Tianyang Wang, Hao Xu.
Project: [StructLoRA](#) Paper: [arXiv](#) Code: [GitHub](#)

- **[TMLR 2026] Adapting Vision Foundation Models with Cascaded Semantics**
Vision Foundation Models Foundation Model Adaptation Semantic Adaptation
Xi Xiao, Xingjian Li, Cheng Han, Tianyang Wang, Guosheng Hu, Yunbei Zhang, Lin Zhao, Runmin Jiang, Xi Li, Xiao Wang, Min Xu.
Project: [Cascaded Semantics](#) Paper: [OpenReview](#)
- **[TMLR 2026] Prompt-based Adaptation in Large-scale Vision Models: A Survey**
Vision Foundation Models Foundation Model Adaptation Survey
Xi Xiao, Yunbei Zhang, Lin Zhao, Yiyang Liu, Xiaoying Liao, Zheda Mai, Xingjian Li, Xiao Wang, Hao Xu, Jihun Hamm, Xue Lin, Min Xu, Qifan Wang, Tianyang Wang, Cheng Han.
Paper: [arXiv](#)
- **[ACM MM 2026] Rethinking Layer-Wise Information Allocation for Vision Foundation Model Adaptation**
Vision Foundation Models Foundation Model Adaptation Information Bottleneck
Xi Xiao*, Yuqi Li*, Yunbei Zhang, Lin Zhao, Yu Li, Aiden Zhao, Tianyang Wang, Hao Xu, Yingli Tian. (* Co-first author)
Paper: [arXiv](#) Code: [GitHub](#)
- **[ACM MM 2025] Visual Instance-Aware Prompt Tuning**
Vision Foundation Models Foundation Model Adaptation Instance-Aware Adaptation
Xi Xiao, Yunbei Zhang, Xingjian Li, Tianyang Wang, Yuxiang Wei, Xiao Wang, Jihun Hamm, Min Xu.
Paper: [ACM DL](#) Code: [GitHub](#)
- **[ICML 2026] Dispersion Loss Counteracts Embedding Condensation and Improves Generalization in Small Language Models**
Language Models Representation Learning Embedding Regularization
Xi Xiao*, Chen Liu*, Xingzhi Sun*, Alexandre Van Tassel*, Ke Xu, Kristof Reimann, Danqi Liao, Mark Gerstein, Tianyang Wang, Xiao Wang, Smita Krishnaswamy. (* Equal contribution / Co-first authors)
Project: [LM-Dispersion](#) Paper: [arXiv](#)
- **[ICCV 2025] MagicID: Hybrid Preference Optimization for ID-Consistent and Dynamic-Preserved Video Customization**
Video Diffusion Models Preference Optimization Identity Preservation
Xi Xiao*, Hengjia Li*, Lifan Jiang*, Tianyang Wang, Hongwei Yi, Boxi Wu, Deng Cai. (* Co-first author)
Project: [MagicID](#) Paper: [CVF](#) Code: [GitHub](#)
- **[NeurIPS 2026] Drift Flow Matching**
Generative Models Flow Matching Test-Time Scaling
Chenrui Ma, Xi Xiao, Lin Zhao, Tianyang Wang, Ferdinando Fioretto, Yanning Shen.
Paper: [arXiv](#)
- **[NeurIPS 2026] Stochastic Interpolants via Conditional Dependent Coupling**
Generative Models Stochastic Interpolants End-to-End Generation
Chenrui Ma, Xi Xiao, Tianyang Wang, Xiao Wang, Yanning Shen.
Paper: [arXiv](#)
- **[ECCV 2026] SpaMEM: Benchmarking Dynamic Spatial Reasoning via Perception-Memory Integration in Embodied Environments**
Embodied AI Spatial Reasoning Perception-Memory Integration
Chih-Ting Liao, Xi Xiao, Chunlei Meng, Zhangquan Chen, Yitong Qiao, Weilin Zhou, Tianyang Wang, Xu Zheng, Xin Cao.
Paper: [arXiv](#)
- **[CVPR 2026] Learning Straight Flows: Variational Flow Matching for Efficient Generation**
Generative Models Flow Matching Variational Inference
Chenrui Ma, Xi Xiao, Tianyang Wang, Xiao Wang, Yanning Shen.
Paper: [arXiv](#)
- **[AAAI 2026] CAD-VAE: Leveraging Correlation-Aware Latents for Comprehensive Fair Disentanglement**
Variational Autoencoders Disentangled Representation Algorithmic Fairness
Chenrui Ma, Xi Xiao, Tianyang Wang, Xiao Wang, Hao Zhang, Yanning Shen.
Paper: [arXiv](#) Code: [GitHub](#)
- **[NeurIPS 2026] Leveraging Latent Visual Reasoning in Silence**
MLLM Post-Training Reinforcement Learning Latent Visual Reasoning
Dongyao Zhu, Zhen Wang, Xi Xiao, Han Jiang, Saeed Vahidian, Wei-Lun Chao, Tanya Berger-Wolf, Yu Su, Raju Vatsavai, Jianyang Gu.
Paper: [arXiv](#)
- **[NeurIPS 2026] UtoMe: Observation-Uncertainty-Guided Token Merging for Weather Foundation Models**
AI for Science Uncertainty-Guided Token Merging Weather Foundation Models
Yunbei Zhang, Janet Wang, Xi Xiao, Jihun Hamm, Xiao Wang.
Conference page: [NeurIPS](#)

- **[Findings of EMNLP 2026] Dynamic Hub-and-Spoke Memory for Streaming Video Understanding**
[Streaming Video Understanding](#) [Long-Range Memory](#) [Training-Free VLM](#)
 Xinru Jiang, Lin Zhao, [Xi Xiao](#), Yunbei Zhang, Janet Wang, Chenrui Ma, Haolin Li, Yanzhi Wang, Yifan Gong, Octavia Camps.
 Paper: [arXiv](#)
- **[CVPR 2026] fMRI-LM: Towards a Universal Foundation Model for Language-Aligned fMRI Understanding**
[Multimodal Foundation Models](#) [fMRI Representation Learning](#) [Language Alignment](#)
 Yuxiang Wei, Yanteng Zhang, [Xi Xiao](#), Tianyang Wang, Chenxuan Qian, Vince Calhoun.
 Paper: [arXiv](#) Code: [GitHub](#)
- **[NeurIPS 2025] MoRE-Brain: Routed Mixture of Experts for Interpretable and Generalizable Cross-Subject fMRI Visual Decoding**
[Mixture of Experts](#) [fMRI Decoding](#) [Cross-Subject Generalization](#)
 Yuxiang Wei, Yanteng Zhang, [Xi Xiao](#), Tianyang Wang, Xiao Wang, Vince Calhoun.
 Paper: [arXiv](#) Code: [GitHub](#)
- **[CVPR 2026] HierAmp: Coarse-to-Fine Autoregressive Amplification for Generative Dataset Distillation**
[Dataset Distillation](#) [Autoregressive Modeling](#) [Coarse-to-Fine Generation](#)
 Lin Zhao, Xinru Jiang, [Xi Xiao](#), Qihui Fan, Lei Lu, Yanzhi Wang, Xue Lin, Octavia Camps, Pu Zhao, Jianyang Gu.
 Paper: [arXiv](#)
- **[Findings of ACL 2026] HyperAdaLoRA: Accelerating LoRA Rank Allocation During Training via Hypernetworks without Sacrificing Performance**
[LLM Adaptation](#) [LoRA & PEFT](#) [Hypernetwork Rank Allocation](#)
 Hao Zhang, Zhenjia Li, Runfeng Bao, [Xi Xiao](#), Yifan Gao, Heng Zhang, Shuyang Zhang, Bo Huang, Yuhang Wu, Tianyang Wang, Hao Xu.
 Paper: [arXiv](#)
- **[ECCV 2026] Less Tokens, Better Forecasts: Sparse Residual Routing for Efficient Weather Prediction**
[AI for Science](#) [Sparse Residual Routing](#) [Weather Forecasting](#)
 Janet Wang, Yunbei Zhang, Lin Zhao, [Xi Xiao](#), Jihun Hamm, Xiao Wang.
 Paper: [OpenReview](#)
- **[ECCV 2026] Personalize Your Large Vision-Language Models With In-Context Prompt Tuning**
[Vision-Language Models](#) [Foundation Model Adaptation](#) [Personalization](#)
 Yanshu Li, Jiaqian Li, Kuai Yu, [Xi Xiao](#), Dongfang Liu, Tianyang Wang, Ruixiang Tang.
 Paper: [OpenReview](#)
- **[Findings of EMNLP 2025] Sensitivity-LoRA: Low-Load Sensitivity-Based Fine-Tuning for Large Language Models**
[LLM Adaptation](#) [LoRA & PEFT](#) [Sensitivity-Based Tuning](#)
 Hao Zhang, Bo Huang, Zhenjia Li, [Xi Xiao](#), Hui Yi Leong, Zumeng Zhang, Xinwei Long, Tianyang Wang, Hao Xu.
 Paper: [ACL Anthology](#)
- **[SC 2025] ORBIT-2: Scaling Exascale Vision Foundation Models for Weather and Climate Downscaling**
[AI for Science](#) [Exascale Pretraining](#) [Climate Foundation Models](#)
 Xiao Wang, Jong-Youl Choi, Takuya Kurihaya, Isaac Lyngaas, Hong-Jun Yoon, [Xi Xiao](#), Ming Fan, Nasik Muhammad Nafis, Aristeidis Tsaris, Ashwin M. Aji, Maliha Hossain, Mohamed Wahib, Dali Wang, Peter Thornton, Prasanna Balaprakash, Moetasim Ashfaq, Dan Lu.
 Paper: [ACM DL](#) Story: [AMD](#)
- **[COLM 2025] M²IV: Towards Efficient and Fine-grained Multimodal In-Context Learning in Large Vision-Language Models**
[Vision-Language Models](#) [Multimodal In-Context Learning](#) [Representation Engineering](#)
 Yanshu Li, Hongyang He, Yi Cao, Qisen Cheng, Xiang Fu, [Xi Xiao](#), Tianyang Wang, Ruixiang Tang.
 Paper: [arXiv](#)

HONORS AND AWARDS

- | | |
|--|------|
| • ACM Gordon Bell Prize Finalist | 2025 |
| • SC 2025 Best Paper Award | 2025 |
| • Tinker Research Grant (\$5,000 USD), Thinking Machines Lab | 2025 |