

Xin Yu

814-769-6357 | xmy5152@psu.edu | [Linkedin](#) | [Google Scholar](#)

EDUCATION

The Pennsylvania State University

PhD candidate in Statistics (GPA 4.0/4.0)

State College, PA

Aug. 2023 - May. 2027(Expected)

Nankai University

Bachelor's degree in Statistics (GPA 3.90/4.0)

Tianjin, China

Aug. 2019 - May. 2023

RESEARCH INTERESTS

- *Reinforcement Learning in Foundation Models:* Developing RL-based approaches to align and adapt multimodal foundation models to human preferences, enabling efficient and responsive AI agents capable of interactive reasoning, planning, and vision–language understanding in real-world environments.

TECHNICAL SKILL

- Coding: Python(Pytorch), PEFT in Hugging Face, R, MATLAB, SQL.
- Math and Probability: Stochastic differential equation, advanced linear algebra, and optimization theory.

EXPERIENCE

Ongoing project

Aug. 2025 – Now

SLDM Member

Penn State, USA

- Introduced a decoding-time user-level preference alignment method for LLMs, effectively mitigating reward over-optimization and improving agent-level personalization stability..
- Applied LLMs to medical text understanding (EHRs), enhancing interpretability in diagnosis prediction — demonstrating domain-specific generalization of foundation models.

Bytedance Seed

May. 2025 – Now

Research Scientist Intern

Seattle, USA

- Developed vision–language reinforcement learning frameworks for multimodal reasoning tasks such as math–vision question answering and image segmentation.
- Designed a parameter-efficient reinforcement fine-tuning pipeline for `volcengine/ver1`, improving convergence efficiency in large-scale distributed settings.
- Formulated a list-wise human preference optimization algorithm for diffusion models, enabling personalized text-to-image generation without explicit reward modeling.

Microsoft Research Asia

Jan. 2023 – May. 2023

Intelligent Multimedia Group Intern

Beijing, China

- Developed light-weight convolutional architectures for visual recognition under data imbalance, improving model efficiency and generalization.

PUBLICATIONS

1. Understanding the Statistical Accuracy-Communication Trade-off in Personalized Federated Learning

Xin Yu, Zelin He, Ying Sun, Lingzhou Xue, Runze Li.

Published at The Forty-Second International Conference on Machine Learning, Vancouver

- Established the first minimax optimal statistical guarantee for the objective in FedProx under data heterogeneity with mild assumptions and proposed a guideline for the hyperparameter strategy
- Designed an algorithm, provably showing that stronger personalization reduces communication complexity without increasing the computation cost overhead in the algorithmic convergence
- Conducted simulations in various datasets and models to verify our findings indicating the potential benefit of the class of regularization based personalized federated learning

2. AltLoRA: Towards Better Gradient Approximation in Low-Rank Adaptation with Alternating Projections

Xin Yu, Yujia Wang, Jinghui Chen, Lingzhou Xue

Published at The Thirty-Ninth Annual Conference on Neural Information Processing Systems, San Diego

- Proposed AltLoRA, a novel parameter-efficient fine-tuning method that improves gradient approximation via alternating projections, addressing limitations of existing LoRA variants in LLM post-training.
- Demonstrated that AltLoRA enables stable momentum integration and transformation-invariant learning without additional memory cost.
- Theoretical analysis and extensive experiments show that AltLoRA consistently outperforms LoRA and LoRA-Pro across multiple tasks while maintaining memory efficiency.

PREPRINTS

A New Inexact Manifold Proximal Linear Algorithm with Adaptive Stopping Criteria

Submitted to JMLR

- Proposed an adaptive algorithm to solve nonsmooth composite problem efficiently, which would be applied for Sparse PAC, Sparse Spectral Clustering and etc
- Established the total complexity for our algorithm and proved that the update convergences to the stationary point

Each Rank Could be an Expert: Single-Ranked Mixture of Experts LoRA for Multi-task Learning

Submitted to KDD 2026

- Unified single-LoRA and multi-LoRA MoE into a novel framework by treating each rank as an independent expert with dynamic rank-wise activation.
- Proved that finer-grained expert partitioning expands parameter diversity and improves error bounds, validated by consistent performance gains in multi-task LLM adaptation.
- Theoretical analysis and extensive experiments show that AltLoRA consistently outperforms LoRA and LoRA-Pro across multiple tasks while maintaining memory efficiency.

PrunedLoRA: Efficient and Robust Gradient-Based Structural Pruning for Low-rank Adaptation

Submitted to ICLR 2026

- Proposed PrunedLoRA, a gradient-based structural pruning framework for low-rank adaptation. It achieves provable robustness, and on Llama-3-8B matches or surpasses full fine-tuning performance on GSM8K and HumanEval, while remaining parameter-efficient.

Towards Better Optimization For Listwise Preference in Diffusion Models

Submitted to ICLR 2026

- Proposed the first listwise preference optimization framework for diffusion models, eliminating the need for reward models and leveraging richer human ranking signals.
- Provided theoretical guarantees showing listwise optimization is more stable and no worse than pairwise DPO, with broader applicability to Diffusion-DPO methods.
- Demonstrated superior performance on Stable Diffusion 1.5 and SDXL across text-to-image generation, image editing, and personalized preference alignment tasks.

AWARDS

Best Paper in C.R Conference 2025 Jack and Eleanor Pettit Scholarship in Science 2024-2025
University Graduate Fellowship 2023-2024

Second Prize, National Undergraduate Academic Competition 2023

Second Prize, National High School Mathematics Competition 2018

Second Prize, National High School Physics Competition 2018

SERVICE

Review: ICML 2025, Neurips 2025, ICLR 2026

Statistical Consulting: 2024-2025