Yijun Pan

yijun.pan@yale.edu | +1 734-882-7842 | charles-pyj.github.io

EDUCATIONAL BACKGROUND

Yale University **New Haven, CT | Aug 2025 - May 2027 (expected)**
Master of Science in Computer Science (Selective 2-year Research Track)

University of Michigan **Ann Arbor, MI | Aug 2023 - May 2025**
Bachelor of Science in Data Science, **summa cum laude** GPA: 4.00 / 4.00

Shanghai Jiao Tong University **Shanghai, China | Sep 2021 - Jun 2025**
Bachelor of Science in Electrical and Computer Engineering

TECHNICAL SKILLS

- **Programming:** Python, C/C++, SQL, JavaScript, Git, Linux, Slurm
- **ML frameworks:** PyTorch, scikit-learn, Hugging Face Transformers, Weights & Biases

MACHINE LEARNING & RESEARCH EXPERIENCE

Research Intern, Alibaba Inc. (AliStar Top Talent Program) **May 2026 - Present**
Business Arena: Benchmarking LLM Agents in a Realistic Marketplace | **business-arena.site.accio.ai**

- Built Business Arena, a long-horizon environment for autonomous business agents, with action attribution linking financial outcomes to tool calls for **fine-grained credit assignment** in post-training.
- Engineered a **white-box agent harness** with exact save/fork/load across context, workspace, and market state; trace-search test-time scaling improved final net worth by **16.7%** over the strongest independent rollout.
- Built a **long-horizon post-training pipeline** that samples consequential decision states, constructs outcome-linked SFT supervision, and applies **Single-Rollout Asynchronous Optimization (SAO)** to sparse, delayed rewards.

Student Researcher, Snap Inc. & Yale University **Sep 2025 - Feb 2026**
FlexRec: Adapting LLM-based Recommenders for Flexible Needs via Reinforcement Learning **(Submitted to NeurIPS)**

- Proposed **FlexRec**, an RL post-training framework that aligns LLM recommenders with dynamic user needs, achieving up to **59% higher NDCG@5** and **109.4% higher Recall@5** across multi-need ranking tasks.
- Introduced a **swap-based item-level reward** to address coarse sequence-level credit assignment, providing dense supervision for individual item placements in autoregressive ranking.
- Engineered an **uncertainty-aware RL training pipeline** for sparse and noisy interaction feedback, improving robustness by down-weighting unreliable reward signals.

Machine Learning Research Intern (Remote), CMU **Jan 2025 - May 2025**
DATE-LM: Benchmarking Data Attribution Evaluation for LLMs **(First-author paper at NeurIPS 2025)**

- Defined and implemented end-to-end data-attribution pipelines for LLMs, translating goals such as safety filtering and factual tracing into metrics, dataset splits, and PyTorch workflows; guided a **3-person team** through sprint planning and code reviews.
- Collaborated with external research teams to standardize a common API for model training, evaluation, and attribution, enabling seamless integration of new submissions.
- Developed an automated CI/CD pipeline and Hugging Face leaderboard, defining submission schemas and evaluation protocols and publishing rankings and metrics.

Machine Learning Research Intern (Remote), UIUC & USC **Jun 2024 - Dec 2024**
Detecting Unsafe Training Data via Data Attribution Methods **(First-author paper at ICML 2026)**

- Built an unsafe-alignment pipeline that injects **<1% toxic data** into benign corpora, increasing model attack success by up to **50%** and exposing critical safety gaps.
- Developed an unsafe training-data detection pipeline for Llama-3-8B, Gemma-29B, and Phi-3.5, improving detection AUC by up to **7.5%** over SOTA classifiers in jailbreak settings and **44.1%** in gender-bias settings.
- Designed a token-wise gradient-selection algorithm to isolate high-contributing features under noisy model responses, reducing retraining attack success rate by up to **45%** versus baselines.