

Xintian Pan

☎ +1 475-227-5113 | ✉ xintian.pan@yale.edu | 🌐 Personal Website | 🐙 GitHub | 🔗 LinkedIn | 📍 New Haven, CT

Professional Experience

Snowflake Bellevue, WA
Research Intern (AI/ML) May 2026 – Aug 2026

Education

Yale University New Haven, CT
M.S. in Statistics and Data Science Aug 2025 – Dec 2026 (Expected)

Nanjing University Nanjing, China
B.S. in Computer Science Sept 2021 – Jul 2025
GPA: 4.51/5.0 (Rank: 3/103)

University of California, Berkeley Berkeley, CA
Exchange Student, Berkeley International Study Program Jan 2024 – May 2024
GPA: 4.0/4.0

Publications

- 2025 Jianliang He, **Xintian Pan**, Siyu Chen, Zhuoran Yang. "In-Context Linear Regression Demystified: Training Dynamics and Mechanistic Interpretability of Multi-Head Softmax Attention." arXiv preprint arXiv:2503.12734, 2025.
- 2025 Ed Li, Junyu Ren, **Xintian Pan**, Cat Yan, Chuanhao Li, Dirk Bergemann, Zhuoran Yang. "Build Your Personalized Research Group: A Multiagent Framework for Continual and Interactive Science Automation." arXiv preprint arXiv:2510.15624, 2025.

Research Experience

Snowflake Bellevue, WA
Research Intern (AI/ML) May 2026 – Aug 2026

Manager: Zhewei Yao, Mentor: Boyi Liu, Snowflake AI Research

- Research direction: LLM post-training, agentic RL.
- Developed **SWARM**, multi-agent RL training for deep research, continued jointly with Yale (*ongoing*).
- Explored RL algorithms for multi-agent training, including GRPO, DPPO, CISPO, and SAO.

Yale University New Haven, CT
Research Assistant Aug 2025 – Present

Supervisor: Prof. Zhuoran Yang, Department of Statistics and Data Science

- Research direction: LLM pre-training and post-training, agentic RL, and evaluation of LLM agents.
- Explored Nanotron for LLM pre-training and conducted LLM post-training with RL frameworks such as VERL and SLIME.
- Contributed to **freephdlabor**, an open-source multiagent framework that automates the full research lifecycle, from hypothesis generation through experimentation to publication-ready manuscripts.
- Developing **ForensicsBench**, a benchmark evaluating LLM agents on end-to-end digital forensics investigations, and **AD-Bench**, an Attack-and-Defense CTF-style benchmark evaluating LLM agents' offensive and defensive cybersecurity capabilities (*ongoing*).
- Continuing **SWARM**, multi-agent RL training for deep research, jointly with Snowflake (*ongoing*).

Yale University New Haven, CT
Student Intern Jul 2024 – Jan 2025

Supervisor: Prof. Zhuoran Yang, Department of Statistics and Data Science

- Research direction: Mechanistic Interpretability
- Identified the kernel estimator learned by 1-Layer Transformers for linear regression task and discovered the training dynamics and learned patterns.

Projects

SWARM

May 2026 – Present

- Designed an RL training recipe for multi-agent deep research systems, in which an orchestrator delegates sub-tasks to one or more executors, collects their findings, and issues new tasks until confident in the answer.
- Designed joint training of the orchestrator and executors under a shared policy, sampling executor trajectories from each orchestrator rollout based on its outcome (successful executors from correct rollouts, failed executors from incorrect ones).
- Prevented model collapse from invalid tool calls by ending orchestrator rollouts with zero reward and truncating executor trajectories at their first invalid call.
- Extended PRIME-RL to schedule executor requests alongside orchestrator requests, balancing load across inference GPUs.
- Trained Qwen3.5-4B on BrowseComp-Plus with LLM-as-a-Judge rewards, and observed the orchestrator learning parallel workflows.

AD-Bench

May 2026 – Present

- Built a game-like benchmark modeled on Attack-and-Defense CTF to evaluate LLM agents' offensive and defensive cybersecurity capabilities, pairing one model as a red (attacker) agent against another as a blue (defender) agent over an isolated Docker match network.
- Curated 300+ intentionally vulnerable services from open-source Docker environments, and agents play blackbox without knowing the underlying vulnerability.
- Designed the game infrastructure on the open-source A&D CTF framework ForcAD, running the service and each agent in separate Docker containers with role-based access to the service, and isolating each service match so that matches run in parallel.
- Built a global checker that periodically plants flags, verifies service availability, and scores both sides independently of agent steps, plus automatic recovery that restores a service snapshot when a faulty patch crashes the service.
- Developed ReAct-style red and blue agents that play asynchronously.

ForensicsBench

Jan 2026 – Present

- Built a benchmark of 14 labs (237 questions) curated from open-source forensic challenges to evaluate LLM agents on end-to-end digital forensics investigations, spanning network, memory, and disk forensics, log analysis, and malware analysis.
- Developed a ReAct-style LLM agent that investigates evidence inside per-lab Docker environments, supporting multiple LLM APIs via LiteLLM.
- Designed hybrid grading combining string matching for factual questions with rubric-based LLM-as-a-Judge for analytical ones, plus post-hoc trajectory analysis.

Freephdlabor | github.com/ltjed/freephdlabor

Aug 2025 – Oct 2025

- Developed ReviewerAgent for freephdlabor multiagent system for evaluation of generated papers.
- Developed resume mechanism which supports the multiagent system to resume memory from pervious run and a non-blocking mechanism which allows real-time human intervention.

Awards & Honors

First Class People's Scholarship, Nanjing University (**Top 5%**)
Second Prize Scholarship for Basic Subjects, Nanjing University

Skills

Languages	Python, C/C++ , Java, TypeScript, SQL, ~~TeX~~
Training & RL	PyTorch, Megatron, Nanotron, PRIME-RL, VERL, SLIME
Inference & Agents	vLLM, SGLang, LiteLLM
Infrastructure	Linux, Docker, Kubernetes, Git, AWS, Weights & Biases