

Francis Fan

📞 484-802-6820 | ✉ francis.fan@yale.edu | 🏠 Philadelphia, PA | 🔗 LinkedIn | 🐙 Github | 📁 Portfolio

EDUCATION

Yale University

Aug 2023 – May 2027

Combined BS/MS in Computer Science

GPA: 3.96/4.00

- **Activities:** Yale Club Swim (Captain), Yale Computer Society, Yale Undergraduate Capital Partners (Associate)
- **Courses:** Operating Systems, Database Design, Data Structures, Systems Programming and Computer Organization, Algorithms, Parallel Programming, Machine Learning, Compilers and Interpreters, Computational Intelligence for Games, Quantum Computing, Discrete Mathematics, Software Engineering, Probability Theory

EXPERIENCE

Databricks

May 2026 - Aug 2026

Software Engineering Intern

Mountain View, CA

- Incoming intern on a data replication platform team building ELT pipelines that extract data from MySQL/TiDB sources and load analytical Delta Lake tables for internal analytics, monitoring, and reporting
- Developing an agentic operations platform for pipeline health monitoring, root cause analysis, oncall assistance, and lightweight mitigation across CDC, Kafka, Spark streaming, and Delta Lake infrastructure

Ramen Inc. (Series A Startup)

May 2025 - July 2025

Software Engineering Intern

Bala Cynwyd, PA

- Tested and engineered modular multi-step AI agents using MLflow, RAG pipelines, and ChromaDB to power user-adaptive UI components, boosting diagnostic engagement by 50% and reducing time-to-diagnosis by 35%
- Deployed AI agents with PostgreSQL, Kubernetes, and Docker, achieving 99.9% system uptime and enhancing output coherence through automated behavior validation against live database state and performance benchmarks

Yale Department of Computer Science

Sept 2023 – Dec 2025

Undergraduate Research Intern under Prof. Danny Rakita

New Haven, CT

- **Project 1: Diffusion Policy for Manipulation** - Implemented diffusion policy, reinforcement learning (RL), and behavioral cloning script on models trained by GANs using ROS, PyTorch, and Python on dual robotic arms, optimizing viewpoints and object manipulation to achieve a 90% success rate across various movements
- **Project 2: Robotics Generative AI Reinforcement Learning Pipeline** - Developed a high-throughput, autonomous reward-gen and task-analysis pipeline with PyTorch, ROS (publisher/subscriber nodes), and OpenAI Gymnasium, enabling continuous RL training in a virtual environment and reducing manual intervention by 25%

PROJECTS

Parallelizing Large Number Multiplication | C++, OpenMP, ParlayLib

Feb 2025 - May 2025

- Engineered high-performance implementations of Karatsuba and 3-way Toom–Cook multiplication algorithms for arbitrary-precision integers, leveraging cache-aware data structures and optimized vectorized digit arithmetic
- Designed novel parallelization of 3-way Toom–Cook using OpenMP and ParlayLib primitives, exploiting divide and conquer independence for fine-grained concurrency, achieved $535.25\times$ speedup, with asymptotic scaling analysis

Yale Clubs | TypeScript, React.js, Node.js, MongoDB, JWT

Sept 2024 - Present

- Built Yale's club management app with secure student authorization, reducing manual administrative workload
- Integrated AI-powered event verification and submission workflows to support scalable, secure club operations

Brain Matter Data Analysis | SHAP, Random Forests, SVM, Multiple Linear Regression

June 2022 - Aug 2023

- Utilized an interpretable model to quantify and confirm the cerebellum's importance in predicting risk tolerance and evaluated the performance of several machine learning techniques using 1000 runs of 10-fold cross-validation
- Published work using machine learning to assess the relationship between risk tolerance and brain gray matter volume as first author, and presented at the Organization for Human Brain Mapping (OHBM) Conference 2023

ADDITIONAL

Technical Skills: Python (PyTorch, Scikit-learn), C/C++, Java, Javascript, ReactJS, Typescript, HTML, XML, SQL, R, CSS, Racket, x86-64 assembly, Git, Github, PostgreSQL, Jupyter Notebook, L^AT_EX(Overleaf/R Markdown), VS Code, Robot Operating System, Linux, Windows, Bash, Excel, React, Node.js, Figma, Jira, Jenkins, Vercel, Render, WordPress

Preprints/Publications: X. Sun, F. Fan, Y. Chen, D. Rakita. "A Comparative Study on State-Action Spaces for Learning Viewpoint Selection and Manipulation with Diffusion Policy." arXiv preprint, 2024. arXiv:2409.14615; Xiatao Sun, Shuo Yang, Yinxing Chen, Francis Fan, Yiyan (Edgar) Liang, Daniel Rakita. "Dynamic Rank Adjustment in Diffusion Policies for Efficient and Flexible Training" RSS, 2025