Jia Hao Liang (Frank)

MIT | jhliang@mit.edu | (646) 553-9169 | github.com/flbjcy | flbjcy.github.io

Education

Massachusetts Institute of Technology, Cambridge, MA 2026–2029
Double major: Computer Science & Mathematics | Concentration: Linguistics

Columbia University, Columbia College, New York, NY 2025–2026
GPA: 4.16/4.00 (transferred to MIT). Science Research Fellow (CUSP). Dean's List, all terms.

Relevant Coursework

Current (MIT): Machine Learning, Deep Learning, TinyML, NLP, Introduction to Topology, Theory of Probability
Columbia: Honors Math A & B, Probability Theory, Quantum Computing, Modern Algebra I, Data Structures

Skills

Programming and tools: Python, PyTorch, pandas, NumPy, C++, Hugging Face, MATLAB, LaTeX
Theory: Machine learning, optimization, mathematical modeling, probability, statistics, game theory

Research & Experience

Computer Science and Artificial Intelligence Laboratory (CSAIL), MIT Sept 2026–Present
Student Researcher (UROP), advised by Prof. Vincent Sitzmann and Xingjian Bai

- Investigating agents as optimizers for neural network training through multiple approaches, including evolution strategies, per-layer optimizer code generation, scheduled reviews, and predictive monitoring.
- Tested agent-guided evolution strategies on Rosenbrock, Beale, Himmelblau, Ackley, and Rastrigin. Reduced loss by 65% versus fixed ES–Adam in a single-seed 2D pilot with 826 evaluations per run.
- Reduced CIFAR-10 FID-10K by 14.8%, from 54.85 to 46.71, using a predictive optimizer agent versus fixed-optimizer control at 4,000 updates in a single-seed trial with additional trial compute.

Software Systems Laboratory, Columbia University Jan 2025–Oct 2026
Student Researcher, advised by Prof. Junfeng Yang and Yun-Yun Tsai

- Developed RAFP for LLM lineage verification with fine-tuning stability bounds, reported in an arXiv preprint.
- Evaluated rare-region fingerprints across four base models and five fine-tuning datasets, achieving 93.33–100% average true positive rates. Separate black-box lineage tests against five unrelated models yielded 0% false positives.
- Investigated distributed fingerprinting with Shamir's secret sharing across Qwen2.5 agents, achieving ROC AUC 0.71 in detecting genuine and fine-tuned frameworks against a different-model control.

Transformer Language Model from Scratch, MIT Jun 2026–Aug 2026
Independent Machine Learning Project

- Built a 22.7M-parameter Transformer from scratch with a 10,000-token byte-level BPE tokenizer, rotary positional embeddings, causal self-attention, SwiGLU layers, and RMSNorm.
- Trained on TinyStories for 5,000 steps, reaching validation cross-entropy loss of about 1.80 with custom AdamW, cosine learning rate scheduling, and gradient clipping. Published model weights on Hugging Face.

Yau Mathcamp, Fudan University Jun 2024–Aug 2024
Group Leader, Best Project Award

- Studied the Inverse Galois Problem. Selected among 5 of 70 participants to present at the closing colloquium.

Selected Honors & Awards

Professor Van Amringe Mathematical Prize (sole freshman recipient) Apr 2026
Columbia Integration Bee (Grand Integrator, first freshman champion) Mar 2026
International Mathematical Olympiad (Honorable Mention), China MO (Bronze) Jul 2025
International Olympiad in Artificial Intelligence (Bronze Medal) Aug 2024

Leadership

Traders @ MIT, Executive Board Sept 2026–Present
Lead case creation and competition development for Traders@MIT as one of 8 selected board members.

Columbia Mathematics and Science Publications Sept 2025–May 2026
Board Member, Columbia Journal of Undergraduate Mathematics. Writer, Columbia Science Review. Editor, Columbia Undergraduate Science Journal. Published writing on knot theory and game theory.