Lauren Lee

Brooklyn, NY | (917) 922-9915 | lauren.lee.ll2243@yale.edu | linkedin.com/in/laurenwylee | github.com/laurenwylee

Education

Yale University — New Haven, CT

B.S./M.S. in Computer Science, GPA: 3.93/4.00 — *Expected May 2028*

Coursework: Operating Systems, Compilers & Interpreters, Database Design, Parallel Programming, Systems Programming, Algorithms, Machine Learning, Computer Vision, Data Structures, Probability Theory, Linear Algebra

Experience

Software Development Intern — Jun 2026 – Aug. 2026

Dow Jones & Company — *New York, NY*

- Accelerated the Unified-to-Indexed Design migration across the Wall Street Journal and multiple Dow Jones brands by building a shared React/TypeScript component library, eliminating redundant per-property reimplementation and consolidating UI updates into a single maintained codebase.
- Built an AI-powered E2E test generation platform that constructs a knowledge graph of dependencies, routes, components, and pages across every engineering repo, then uses a LangChain agent to synthesize Playwright test suites automatically filed and tracked through Jira workflows.

Machine Learning Research Intern — May 2025 – Aug. 2025

JC STEM Lab of Data Science Foundations, HKUST — *Hong Kong, CN*

- Improved multi-hop question-answering accuracy by +1.0 pp EM and +0.53 pp F1 on 2WikiMultihopQA by building a hybrid retrieval pipeline that combined dense vector search with keyword-based retrieval over millions of indexed passages, tuning the balance between both methods to find the optimal configuration.
- Cut hallucination rate by $\sim$0.8% (1.364 $\rightarrow$ 1.353 per generation pass) by designing an adaptive retrieval trigger that detected model uncertainty mid-generation and fetched additional context on demand, with filtering to suppress false triggers on low-information tokens.
- Optimized LLM inference throughput on A100 GPUs by running systematic ablation studies on Llama 2 7B, identifying a non-obvious hyperparameter sweet spot where a specific retrieval mix improved both accuracy metrics simultaneously while all adjacent configurations degraded below baseline.

Projects

DeepLearn | *Developer, HackMIT* — Sept. 2024

- Won first-place sponsor prize for best Llama integration at HackMIT (100+ teams) by building a multimodal deep learning platform that jointly reasoned over computer vision and speech representations using a Llama-powered Q&A layer.
- Used NVIDIA SegFormer for real-time scene segmentation, Whisper for speech transcription, and perceptual image hashing for key-frame extraction, fusing outputs into a shared multimodal context for LLM inference.

CertiKOS-based OS Kernel | *Operating Systems, Yale University* — Jan 2026 – May 2026

- Built a full x86 OS kernel in C demonstrating hardware-software co-design across 6 labs: bootloader (real-mode to protected-mode transition), buddy allocator for physical memory, two-level page table virtual memory, ELF process loading with ring-3 user context setup, trap and syscall dispatch, preemptive thread scheduling via assembly context switch, and a buffer-cached filesystem with multiprocessor support (LAPIC/IOAPIC/AP boot).

Yale Clubs | *Developer, Yale Computer Society* — Oct. 2024 – Present

- 12,000+ students served by a production full-stack platform built with end-to-end system design: Next.js + Python backend, REST APIs, indexed MongoDB document relations, and optimized query patterns managing subscriptions across 500+ organizations.
- Cut search complexity from $O(n \cdot m)$ naive substring matching to $O(m)$ prefix traversal by implementing a trie-based search engine with multi-token prefix matching and tag-category filtering, enabling real-time search under concurrent user load.

Skills

Languages: Python, C++, C, Java, TypeScript, JavaScript, SQL, OCaml
Frameworks & Libraries: PyTorch, TensorFlow, NumPy, pandas, React, Next.js, Node.js, Tailwind CSS
Tools & Infrastructure: Git, Docker, Linux, CUDA, AWS, Elasticsearch, MongoDB
Developer Tools: Claude Code, GitHub Copilot, Codex