

# Benjamin Liu

bliu8@seas.upenn.edu | linkedin.com/in/benliu8 | github.com/bliu8

## EDUCATION

---

### University of Pennsylvania

*B.S.E., Computer Science*

May 2028

GPA: 3.94/4.00

– **Coursework:** Big Data Analytics, Linear Algebra, AI Lab, Data Structures & Algorithms, LLM Agents

## PROJECTS & RESEARCH

---

### High-Performance Vector Search Engine & RAG System

Nov 2025 – Present

- Built a high-performance vector search engine in C++ with a custom **HNSW** index.
- Achieved **40k+ QPS @ 97% recall** on 1M vectors (**76×** faster than 12-thread brute-force, M4 Pro).
- Improved distance computation **2.3×** via ARM NEON SIMD, cache-aligned layouts, and prefetching.
- Implemented hybrid dense+BM25 retrieval with RRF fusion; Python bindings via **pybind11** for **RAG**.
- Parallelized index build to **22k inserts/s** (1M in **45s**) using OpenMP + fine-grained locks.

### ML Developer — Penn Institute for Immunology and Immune Health

Nov 2025 – Present

- Co-developed a **ResNet-based** model to detect spectral unmixing errors in high-parameter flow cytometry.
- Engineered an **end-to-end pipeline** (preprocess → inference → QC), enabling automation at scale.
- Achieved **99% accuracy** on unmixing error detection and reduced latency from **hours to seconds**.
- Deployed department-wide via Docker and Pennsieve; **sole undergraduate** contributor.
- **First author** abstract submitted to CYTO 2026 (*manuscript in preparation*).

### AI Research Assistant — Shen Lab, Perelman School of Medicine at Penn

Oct 2025 – Present

- Fine-tuning **transformer models** to automate cytometry gating; following up on a *Nature Communications* study.
- Revised and validated **~4,000** gated populations across **135** cytometry files using OMIQ.

## PROFESSIONAL EXPERIENCE

---

### Incoming AI/ML Engineering Intern — SailPoint

Jun 2026 – Aug 2026

### Founding AI Engineer (Part-Time) — Aurbit

Nov 2025 – Present

- Built injury-prevention agent using LLM tool-calling and Supabase to analyze Garmin health data.
- Designed **PostgreSQL** lookup tools to calculate personalized risk scores and recovery metrics.
- Engineered multi-step logic processing heart rate variability and workload ratios via **Deno** edge functions.
- Shipped production features to a **200+ user** beta, refining the **React** frontend and decision engine.

### Research Intern — Stanford University School of Medicine

Jun 2024 – Aug 2025

- Co-designed and prototyped thermoelectric cooling glove to prevent chemotherapy-induced peripheral neuropathy.
- Filed a **provisional patent** and presented to **Stanford Cancer Institute**.
- Conducted physiology experiments using laser Doppler flowmetry and heat-flux sensors.
- Pitched at NASA as **Conrad Challenge Global Finalist (top 5 globally)**.

## TECHNICAL SKILLS

---

**Languages:** C++, Python, Java, TypeScript/JavaScript, Swift, SQL

**ML/AI:** PyTorch, TensorFlow, OpenCV, LangGraph, LLM APIs (OpenAI, Anthropic, Together)

**Systems:** Multithreading, OpenMP, ARM NEON SIMD, cache/memory optimization

**Backend/Infra:** PostgreSQL, Redis, Docker

**Web/Deployment:** React, Next.js, Vercel, Supabase, Firebase, Upstash

**Tools:** Git/GitHub, CMake, pybind11