

Shawn Liu

U.S. Permanent Resident, no sponsorship required

Research Engineer · neurosymbolic AI, secure ML inference, and LLM systems that ship
(510) 565-2523 · xl3704@columbia.edu · LinkedIn · github.com/ShawnnnLiu · OpenReview · shawnnnliu.github.io

EDUCATION

Columbia University · M.S. in Computer Science

Entered Fall 2026 · Expected January 2028

University of California, Irvine · B.S. in Computer Science, GPA 3.844 / 4.0 · Dean's Honor List, all quarters

September 2022 - June 2026

PUBLICATIONS

Cross-Modal Event Encoder: Bridging Image-Text Knowledge to Event Streams

WACV 2026 · arXiv:2412.03093

Co-author. Built the CLIP-event contrastive integration and the attention visualization pipeline.

Integrating Symbolic & Neural Mechanisms for Adversarially Robust HDC

ISLPED 2026 · DOI

Co-author. Ran the robustness evaluation and the symbolic-neural encoding experiments.

Geometric Priors for Generalizable World Models via Vector Symbolic Architecture

NeurIPS 2025 NeurReps Workshop · OpenReview

Co-author. Ran the VSA representation experiments and ablations on the world-model grid, then the MiniGrid follow-ups.

Optimal Hyperdimensional Representation for Learning & Cognitive Computation

Frontiers in AI 2026 · paper

Third author. Encoding-dimension analysis, decoding-accuracy heatmaps, and signal-to-noise distributions.

RESEARCH EXPERIENCE

BiasLab, UC Irvine · Undergraduate ML Researcher · Advisor: Prof. Mohsen Imani

June 2024 - June 2026

- Cut UAV carrier-landing failures by 65% against legacy optical markers in adverse-weather simulation with a PyTorch pose-estimation and symbolic-reasoning model for the U.S. Navy; the reasoning layer holds the landing whenever it sees a person on the deck.
- Aligned a trainable event-camera encoder to CLIP's frozen image-text space for zero-shot video anomaly detection: 57.7 / 64.1 / 54.4 AUC on XD-Violence / UCF-Crime / ShanghaiTech, +5.5 and +5.2 points over frame-based CLIP on the latter two. WACV 2026.
- Characterized CKKS-FHE noise in Microsoft SEAL under GNN operations and built a calibrated noise-injection simulation of encrypted graph inference; a noise-adaptive bootstrapping schedule cut projected overhead about 4x at over 90% accuracy in simulation.

BioIntelligence Lab, UC Irvine · Undergraduate ML Researcher · Advisor: Dr. Haleh Alimohamadi

December 2025 - July 2026

- Built a graph convolutional network over ESMFold-predicted peptide structures, one residue per node, fused with SVM geometric descriptors to predict antimicrobial MIC, with ablations isolating what 3D conformation adds.

Softcom Labs, Cal Poly Pomona · Student Researcher

December 2020 - April 2021

- Built a real-time audio recognition and alarm system, a lightweight CNN over Mel-spectrograms at 95% classification accuracy, deployed on a Raspberry Pi with a Flutter app on the Google Play Store.

ENGINEERING PROJECTS

Loop · LLM scheduling agent, live at loop-study.com Python, FastAPI, SQLite, React / TypeScript, Anthropic API

May 2026 - present

- Turns a career goal and weekly availability into a study plan on Google Calendar: five LLM nodes propose, a five-stage deterministic validator and a hash-checked human approval gate decide, with post-write verification and rollback. The LLM is never handed a calendar.
- A two-tier eval over committed LLM recordings, no model call in CI: tier one gates schema validity, post-repair invalid rate, and citation coverage at pinned floors; an LLM-judge tier scores prose and never gates. One fixture fails on purpose to prove the gate fires.
- Every prompt is pinned at two layers, SHA-256 of the system constant and of the full rendered frame; a prompt-byte change without a version bump fails make check. 4,826 backend and 326 frontend tests; 126 of 126 GitHub Actions runs green.
- Grounded plans in BM25 retrieval over a curated 242-document corpus, graded on 75 labeled queries at recall@5 0.7644, MRR 0.7144, nDCG@5 0.6910, floors pinned to the measured values; a dense-embedding hybrid scored +0.08 MRR and was not shipped.

Foothold • transfer-credit navigator, live at foothold-transfer.com solo • Stellar Pathfinders 2026 entry

July 31 - August 21, 2026

- Built the ASSIST articulation pipeline: 34,975 requests at one per second, zero 429s, and per-agreement fault isolation that turns any parse failure into a typed exclusion code; 31,236 agreements and 352,024 articulation rows for 115 colleges and 15 UC/CSU campuses.
- Wrote a deterministic evaluator over AllOf / AnyOf / course-leaf expression trees with nine typed finding codes; no clock, no PRNG, stable ordering, and a CI fixture asserts the TypeScript triage board equals the Python one.
- Counted N-from requirement groups by courses, not sections: the section-counting reading overstated 23,016 groups and turned pick-one-of-N into complete-all-N. Left ASSIST template overrides unmodelled because ASSIST's own renderer never applies them.
- One Claude node drafts appeal letters; every course code and agreement key is validated against the findings bundle, and two failed repairs fall back to a deterministic template. 870 backend and 89 frontend tests, 298 invalid-input fixtures.

Crash Anticipation • dashcam collision advisories PyTorch, YOLO11n • github.com/ShawnnnLiu/Crash-Anticipation

November 2025 - July 2026

- Distilled a 21.88M-parameter VideoMAE-S teacher into a 2.95M-parameter MobileNetV2 + GRU student whose causal recurrence makes each new frame one CNN forward: 2.45 ms per step, 408 fps on an RTX 5070 Ti.
- On 300 CCD crash and 301 DAD no-crash videos: video AP 0.9998, a mean 3.05 s of warning before impact, 3 false alarms.
- Threw out the first benchmark after it scored AP 1.0 on a model that output $p = 1.0$ on every input, because both manifests contained only positives; rebuilt it with 1,130 real no-crash videos and time-to-event labels with a guard band that contributes no gradient.
- A symbolic layer binds each advisory to a ByteTrack ID: looming as the least-squares slope of log box size over a 6-frame history, a projected-miss test in object widths, and time-to-collision only once the collision-course test holds two consecutive frames.

Arrhythmia Classifier • MIT-BIH leakage study github.com/ShawnnnLiu/Arrhythmia_Classifier • CS 184A, team of 3

November 2025 - August 2026

- Showed that beat-wise splits leak patient identity on MIT-BIH: the same 64,966-parameter CNN scores 98.36% beat-wise and 83.88% patient-wise, and a 443,527-parameter LSTM autoencoder drops from 98.95% to 82.06%. 48 records, 109,445 beats, six classes.
- Wrote a patient-wise split search: 50,000 sampled assignments out of 2,498,640,144, every class present in the test set, one two-record subject never straddling a split, scored by KL plus L1 distance to the overall label distribution; the chosen split differs from it by 3.53%.
- A leakage guard raises before training if any same-subject group spans two splits; macro F1 is reported over present classes alongside all-class macro, since absent classes score zero and conflate test composition with model quality.

PROFESSIONAL EXPERIENCE

AdamsFoods • Full-Stack Developer • adamsfoodswholesale.com

Summers 2024 and 2025

- Built and shipped a live wholesale e-commerce platform in React and Node / Express with S3 signed-URL media and JWT admin routes.
- Built the internal inventory system tracking 300+ SKUs; it replaced handwritten logs and saves staff 5-7 hours a week.

Commit the Change, UC Irvine • Full-Stack Developer

November 2023 - June 2024

- Built pro-bono donation-tracking and authentication software for the nonprofit Feeding Pets of the Homeless in React, Firebase, and SQL.

TEACHING

Certified Learning Assistant, ICS 33: Intermediate Python • Prof. Alex Thornton • UC Irvine

April 2025 - June 2025

- Certified peer educator (UNI STU 176) supporting students in lecture and lab; weekly planning with faculty and TAs.

SKILLS

Languages: Python, TypeScript / JavaScript, SQL, C / C++

ML / AI: PyTorch, scikit-learn, CLIP / ViT / VideoMAE, knowledge distillation, graph neural networks, hyperdimensional computing

Secure ML: CKKS fully homomorphic encryption (Microsoft SEAL), noise characterization, encrypted-inference simulation

LLM systems: Anthropic API, structured outputs, bounded repair loops, two-layer prompt hash-pinning

Evaluation and retrieval: recordings-based eval harnesses with CI gates, BM25 / FTS5 retrieval, retrieval eval (recall@k, MRR, nDCG)

Systems: FastAPI, Pydantic v2, SQLite / FTS5, PostgreSQL, React, Node / Express, Google OAuth / JWT

Infra: Docker, Fly.io, AWS S3, GitHub Actions CI/CD, Linux

Spoken: English (fluent), Mandarin (fluent)