

Joshua Lum

jalum@usc.edu | www.linkedin.com/in/josh-lum | <https://github.com/Plum1234> | <https://joshlum.xyz>

EDUCATION

University of Southern California, Viterbi School of Engineering & Marshall School of Business May 2026
Bachelor of Science, Computer Science and Business Administration GPA: 3.72
Relevant Coursework: Artificial Intelligence: Principles and Foundations, Algorithms and Theory of Computing, Data Structures and Object-Oriented Design, Discrete Methods in Computer Science, Linear Algebra & Differential Equations
Leadership & Involvement: Head Peer Academic Leader, Chinese American Student Association (President)
Awards & Honors: 8VC Engineering Fellow (2026), Dean's List (All Semesters), Magna Cum Laude

EXPERIENCE

Campus – Series B (8VC/Founder's Fund portfolio) New York, NY
Software Engineering Intern (8VC Engineering Fellow) May 2026–August 2026

- Grew Sam (AI tutor) usage from 50 → 350 daily active users by shipping DM integration, file upload, and assignment context across a 4-service stack (Go backend, React frontend, Go WebSocket server, Python/LangChain)
- Built Sam's DM provisioning system (Postgres advisory locks, eligibility sweep cron) and extended it into a proactive outreach flow that confirms schedules via LLM tool-calling with auto-escalation to coaches
- Built a scalable Dagster pipeline to process 1,000+ student-coach transcripts into a compliant GCS bucket, applying LDA, BERTopic, and LLM-based extraction to surface recurring themes in student challenges

Uthana – Seed (A16Z portfolio) Los Angeles, CA
Product & Engineering Intern May 2025–August 2025

- Worked with product and engineering teams to iterate features and launch a live generative AI animation demo at SIGGRAPH 2025, viewed by 2,000+ game developers, ML/AI researchers, and industry stakeholders.
- Built an automated Python workflow that processed and ranked 10K+ user motion prompts using frequency (demand) and embedding-based cosine distance (novelty/coverage) to generate a monthly top-150 MoCap prompt batch and improve training data quality for Uthana's generative motion model

Tamagotchi Lab, USC Viterbi School of Engineering Los Angeles, CA
Undergraduate Student Researcher May 2024–May 2026

- Conducted ML/NLP interpretability research on how large language models (e.g., GPT-2) encode and use syntactic and semantic information, with applications to fairness-critical decision systems.
- Co-authored "Causal Interventions on Causal Paths: Mapping GPT-2's Reasoning From Syntax to Semantics," a paper studying attention mechanisms in transformer models, accepted to CaLM at NeurIPS 2024.
- Analyzed 30K+ sentence pairs across 10 syntactic and 5 semantic templates, using activation patching and ablation to identify 5+ GPT-2 layers driving causal reasoning across 144 attention heads.
- Co-authored "Evaluating Large Language Models for Fair and Reliable Organ Allocation," a paper analyzing fairness and bias in LLM-based kidney allocation decisions using real OPTN data

PROJECTS

[The Intern \(Proactive AI Coding Agent\)](#) | *Replit, Python, Claude Agent SDK, Slack API, Linear, GitHub API*

- Built an always-on AI agent that lives in Slack, triages Linear tickets, writes code on feature branches, opens draft GitHub PRs for human review, and demos on Replit using a multi-agent architecture on the Claude Agent SDK
- Set up continuous preview deployments on Replit via a GitHub Actions pipeline that auto-deploys every agent-opened PR to a live *.replit.app URL, then comments it back on the PR to give reviewers a one-click demo.
- Integrated Slack Socket Mode, Linear and GitHub MCP servers, GitHub App token minting, and Perseus semantic code search to give specialist subagents scoped, repo-limited tool access and guardrails for code safety.

[3D Interactive AI Companion](#) | *C#, OpenAI API, Microsoft Cognitive Services, Unity, Blender*

- Developed a 3D interactive AI companion in Unity, leveraging Blender for 3D modeling and animations, and integrating Microsoft Cognitive Services for speech recognition, synthesis, and return audio
- Used C# and OpenAI API to power dynamic conversational AI and syncing character animations and speech

SKILLS AND INTERESTS

Skills: Python, PyTorch, C++, C#, Go, Java, Javascript, React, Node.js, Typescript, SQL, Git
Interests: Music, Hiking, Pickleball, Spikeball, Snowboarding, Cycling, Gaming