Jessica Ho

(224) 200-4230 | jho32690@usc.edu | LinkedIn | GitHub

EDUCATION

University of Southern California **Aug 2023 – Dec 2026**

B.S. Computer Science and Business Administration *Los Angeles, CA*

- **Relevant Courses:** Data Structures and Object Oriented Design, Software Engineering in C++, Algorithms and Theory of Computing, Deep Learning for AI, Computer Systems and Architecture
- **Involvement:** AthenaHacks, HackSC, Business Technology Club, Society of Women Engineers, Women in Computing

EXPERIENCE

Toyota Financial Services **May 2026 – Aug 2026**

Software Engineering Intern *Plano, Texas*

- Built and deployed a document processing pipeline using **AWS Bedrock** to extract and map key fields from unstructured insurance documents into structured data, automating manual data lookup and reducing turnaround time per document by about **80%** across over **6,000** documents annually
- Designed a document ingestion and storage pipeline using **S3 presigned URLs** for secure direct uploads, with raw files stored in **S3** and **Bedrock** extracted structured data persisted in **DynamoDB** for downstream workflows
- Provisioned serverless infrastructure with **AWS CDK**, deploying **API Gateway** endpoints and **event-driven Lambda** functions with **CloudWatch monitoring**, automated end-to-end tests validating document ingestion, extraction, and **DynamoDB** persistence through **CI/CD**

Chosn AI **Jan 2025 – May 2025**

Software Development Intern *Los Angeles, CA*

- Developed cross-platform journaling features in **Flutter** for voice capture and video recording, using **Firebase** authentication and **Cloud Firestore** to persist user profiles, journal entries, and personalization data, contributing to over **50 active users** within the first month
- Built a **FastAPI** service using HumeAI's emotion recognition API to transform voice and text inputs into emotion-aware signals for personalized responses, increasing average session duration by **40%**
- Reduced backend API latency from **400ms to 95ms** by implementing **Redis** as a distributed caching layer for high-demand video and voice metadata, minimizing **Cloud Firestore** read loads during traffic

PROJECTS

ProductBee | Next.js, TailwindCSS, Auth0, Supabase, Gemini API | Live Demo

- Built a product planning dashboard using **Gemini API** that converts product ideas into PRDs and feature roadmaps, generating structured requirements and development timelines for cross-functional teams
- Implemented feature prioritization and Kanban visual workflows in **Next.js**, using **Supabase** to store projects, feature metadata, priorities, and task status to support shared project planning and progress tracking
- Integrated **Auth0-based** authentication and role-based access controls **(RBAC)** to separate permissions across product, engineering, and stakeholder users and protect project-specific workflows

TwitchBits | NumPy, Pandas, PostgreSQL, Tableau

- Built a streamer analytics pipeline that ingests livestream data from **Twitch API** to analyze viewer behavior, game performance, and channel growth over time
- Automated **Python** data collection at scheduled intervals and designed a **PostgreSQL schema** to persist stream sessions, viewer counts, VOD metadata, and game activity for historical analysis
- Wrote **NumPy** and **Pandas** analysis workflows to aggregate viewer trends, top-performing games, and stream performance, cleaning data and exporting results to **Tableau** dashboards for visual exploration

SKILLS

Programming Languages: C, C++, HTML/CSS, Java, JavaScript, Python, R, SQL

Frameworks: FastAPI, Flutter, Next.js, PyTorch, React.js, RESTful APIs, TensorFlow

Cloud & DevOps: AWS Services, Docker, GitHub Actions, Jenkins

Tools: Agile, Confluence, Firebase, Git, Redis, Jira, MySQL, NoSQL, PostgreSQL