Ye Marn Aung (Jared Aung)

Daly City, CA 94015| (253)-345-2360 | jaredaungfr@gmail.com | https://github.com/JaredAung | https://www.linkedin.com/in/ye-marn-aung/

EDUCATION

San Francisco State University — **San Francisco, CA**
Bachelor of Science, Computer Science (Minor in Mathematics) — Expected December 2026

PROFESSIONAL EXPERIENCE

AI Engineering Extern - Outamation — March 2026 - Present

- Build an end-to-end document intelligence pipeline that classifies and extracts structured data from mortgages, invoices, and contracts Tesseract, PaddleOCR, and EasyOCR for RAG workflows
- Benchmarked layout-aware parsing technique across document formats, identifying OCR-LLM combination that improved extraction efficiency and accuracy on multi-column and table-heavy layouts

PROJECT EXPERIENCE

HowToKeepPlantsAlive: AI Plant Recommendation System & Care Platform — **Github**
LangGraph, PyTorch, Python, MongoDB, Gemini LLM, Voyage AI, Cohere Reranker, DVC, Prefect, MLFlow, Pinecone

- Engineered a two-tower recommendation system using Feast-backed features and holdout-based offline evaluation, outperforming a cosine-similarity baseline by 6.6× on Recall@10, 3.8× on NDCG@10, and 4× on Hit@10 while reducing inference latency from 63.7ms to 2.2ms
- Developed a LangGraph multi-agent system capable of plant Q&A, side-by-side comparison, exploring recommendations, live web search via Tavily, and supported by a hybrid RAG pipeline with cosine similarity, BM25 reranking and recursive/sliding window chunking
- Designed a synthetic data generation framework that simulated user-plant interactions for training and evaluating on a data-scarce domain, successfully capturing 8 out of 10 target behavioral patterns
- Containerized the FastAPI backend with Docker and deployed on Railway and Vercel, serving a catalogue of 1000+ plants

GaitorGate: Full-Stack Search Engine for AI Tools — **Github**
Team Lead | LAMP Stack, Flask, AWS EC2, Google Gemini

- Led a team of 6 to develop a search platform for 100+ AI tools using keyword or NLP-based search, winning 1st place out of 12 teams in a software engineering competition
- Designed and optimized a relational database schema and search pipeline using SQL, enabling natural language query with ~0.4s average latency
- Collaborated cross-functionally with designers, frontend and backend engineers, and QA in an Agile environment to improve team productivity and workflow efficiency

Board2Board: Chess Utility Tool for Over-the-Board (OTB) Game Recognition — **Github**
Keras, Python, OpenCV, ResNet50 Model, Scikit-Learn, Matplotlib, Scipy, Scikit-Image, TensorFlow

- Designed a computer vision pipeline with OpenCV to detect, warp and segment real-life chessboard images into 64 cropped square for piece classification
- Fine-tuned a ResNet50 model on a custom dataset of 2,000+ images, achieving 94% classification accuracy across all classes
- Developed an adaptive thresholding system using linear regression to handle variable lighting and image conditions, improving successful board detection rate by 4× across 150 diverse real-world unseen board images

Additional SKILLS

Programming Languages: Python, Java, JavaScript, TypeScript, Rust

Databases: MySQL, MongoDB, PostgreSQL, Pinecone**,** Supabase

Cloud/DevOps: AWS, Azure, Docker, Railway

ML/AI: Keras, TensorFlow, Pytorch, LangGraph, LangChain, OpenCV, Hugging Face