

ZIANG (ANDREW) TAN

Seattle, WA | ziantg2@uw.edu | 206-601-7855

[linkedin.com/in/ziang-tan](https://www.linkedin.com/in/ziang-tan) | github.com/ziangt2 | ziangtan.com | Availability: Jun.–Sep. 2027

EDUCATION

University of Washington | Seattle, WA

B.S. Electrical & Computer Engineering | B.S. Applied Mathematics: Data Science | Expected Dec 2027

Coursework: Data Structures & Algorithms, Machine Learning, Embedded Systems, Computer Architecture

TECHNICAL SKILLS

Programming: Python, SQL, C++, Java, JavaScript, C, Bash

ML / Retrieval: PyTorch, vision-language embeddings (SigLIP2), FAISS, BM25, RRF, recommendation systems, LLM tool use, grounded RAG, ranking evaluation, YOLO, OpenCV, NumPy, Pandas, SciPy, Weights & Biases, ONNX

Systems / Deployment: Docker, GCP Cloud Run, GitHub Actions, OIDC, Linux, Node.js, TensorRT, CUDA, NVIDIA Jetson Orin Nano, TCP/IP, FreeRTOS

RESEARCH EXPERIENCE

ARC Lab, University of Washington | Seattle, WA

Undergraduate Researcher — Tiny-Aircraft Detection for UAV Detect-and-Avoid | February 2026 – Present

- Tracked three-seed data-selection experiments in **Weights & Biases** using 300, 600, and 1,200-frame label budgets for synthetic-to-real adaptation.
- Prioritized missed detections and false-positive negatives during data selection, targeting the model's weakest failure modes throughout training.
- Lifted frozen-test recall@0.25 on sub-15-pixel aircraft from 0.92% to 1.94%, a 2.1x gain, at the 1,200-frame budget versus random sample selection.
- Designed a **YOLO26n** detector with P2 features, residual refinement, feature reinjection, and channel-spatial recalibration modules stacked together.
- Lifted sub-20-pixel recall@0.50 from 61.11% to 71.30%, a +10.19 percentage-point gain, while also cutting false positives captured per frame.
- Built reproducible training and evaluation across 72K synthetic and 1,799 real images, with tiled inference and frozen-test safeguards in place.

SELECTED PROJECTS

E-commerce Multimodal Search & Recommendation Platform | Python, PyTorch, SigLIP2, FAISS, BM25, LLM Agents / RAG, Docker, GCP, GitHub

- Built image and text retrieval using pretrained **SigLIP2** vision-language embeddings and **FAISS** exact search over 1,000 Amazon Berkeley Objects listings.
- Reached 89.29% image-to-product Recall@10 and a 0.7811 MRR score on 523 held-out image views, backed by image ID, byte, and pixel leakage audits.
- Added a **GPT-4.1 mini** tool-using agent with **BM25** and vector routing, metadata filters, and grounded **RAG** for handling every incoming user query.
- Achieved expected tool routing on 36 of 36 development queries and 28 of 36 nonempty grounded responses, with 6 abstentions and 2 errors logged.
- Deployed containerized retrieval API to **GCP Cloud Run**, reducing **Docker** image size from 4.62 GB to 1.80 GB, a 61% reduction, via CPU-only PyTorch.
- Implemented **GitHub Actions CI/CD** with **OIDC** authentication, measuring 708ms client p95 and 258ms server p95 across 200 sequential requests.
- Built collaborative, TF-IDF, and popularity retrieval with hybrid ranking on 330,270 synthetic events spanning 15,000 distinct registered users.
- Reached 0.2658 Recall@10 and 0.1374 NDCG@10 on a frozen 3,954-target temporal test set, backed by leakage checks and 20 automated test scenarios.

Edge ML Robotic Fruit Detection & Autonomous Sorting | Python, PyTorch, YOLO26n, TensorRT, OpenCV, Jetson Orin Nano, TCP/IP, mechArm

- Built end-to-end edge perception and robotic sorting on Jetson, sending vision-derived targets over **TCP/IP** to a Raspberry Pi-controlled mechArm.
- Audited 193 images and 346 bounding boxes drawn from independent video clips, confirming zero detected cross-split data source leakage anywhere.
- Exported **YOLO26n** to **ONNX** before compiling to FP16 **TensorRT**, lifting inference throughput from 20.83 to 27.57 FPS over PyTorch FP32 baseline.
- Validated 10 physical scenarios reaching 18/18 successful grasps, 15/18 vision-verified placements, and 2/2 successful robotic recovery attempts.
- Benchmarked **OpenCV/HSV** edge detection baseline separately at 52.39 FPS, providing a non-learned comparison point for the trained YOLO26n detector.

ADDITIONAL EXPERIENCE

University of Washington | Undergraduate Teaching Assistant, EE/CSE 474 | September 2025 – Present

- Supported students in **ESP32-S3** labs through code debugging and instruction on **FreeRTOS** scheduling, synchronization, and I2C hardware integration.

Shenzhen Volcos Technology Co., Ltd. | Engineering Intern | July 2025 – September 2025

- Developed sensor-integrated embedded prototypes and **Altium Designer** PCB layouts for motor control and environmental sensing hardware components.
- Supported board bring-up, hardware and software debugging, and functional validation across every embedded prototype iteration completed so far.