

AADITYA RAJ BANIYA

itsadityabaniya@gmail.com | 737-315-0786 | [linkedin.com/in/aadityabaniya](https://www.linkedin.com/in/aadityabaniya) | github.com/aditya-baniya-ai | adityabaniya.netlify.app

PROFESSIONAL SUMMARY

Computer Vision and ML Engineer with production-grade experience optimizing and deploying inference pipelines on edge hardware, building full-stack AI applications, and leading end-to-end research. Recipient of \$13,000 in research grants and competition prizes. Seeking ML/CV engineering roles at companies building at the intersection of deep learning and real-world systems.

SKILLS

ML and Edge: PyTorch, YOLOv8, TensorRT, ONNX Runtime, OpenCV, MediaPipe, Transformers, Jetson AGX Orin, Isaac Sim

Languages: Python, C++, Java, TypeScript, SQL

Infrastructure: Git, Docker, Linux, AWS, CI/CD

Backend and Frontend: FastAPI, Flask, REST APIs, MySQL, RAG, OpenAI and Llama APIs, React.js, Next.js, WebRTC, LiveKit

EDUCATION

Texas State University

Bachelor of Science in Computer Science, Honors Scholar. President's List. GPA: 3.8

San Marcos, TX

Expected May 2027

- Relevant Coursework:** Machine Learning, Data Structures & Algorithms, Object-Oriented Programming, Applied Linear Algebra, Software Engineering
- Organizations:** .EXE, ACM AI, SIAM Math Club, IEEE, CODE PATH ORG

EXPERIENCE

Undergraduate Research Assistant, Human Digital Twin and Computer Vision

Sep 2024 – Present

Ingram School of Engineering, Texas State University — San Marcos, TX

- Voice Alert System for Visually Impaired Workers** - Built a real-time MediaPipe pose estimation pipeline detecting hazardous factory actions and triggering voice alerts, deployed at a San Antonio facility and reducing 85% of potential workplace injury risks
- Makerspace Sentinel (PPE Compliance)** - Optimized a CBAM-enhanced YOLOv8 model for PPE compliance detection, validated across SH17 and R2PPE benchmarks, and deployed to NVIDIA AGX Orin via TensorRT at 25–30ms inference
- Human Action Recognition (HAR)** - Improved human action recognition accuracy by 10% on industrial footage by developing a dual-stream pipeline combining YOLOv8 object detection, MediaPipe pose estimation, and a TimeSformer trained on Assembly101

Tech Lead, MOSTify (Computer Vision Startup)

Sep 2024 – Present

BobCatalyst Innovation Accelerator Program, Texas State University — San Marcos, TX

- Lead technical vision for a CV startup applying human action recognition and computer vision to reduce 3–5 days of manual industrial time-study analysis to 1–2 hours of semi-automated work
- Conducted 40+ customer discovery interviews at industry events and built the CV inference pipeline end-to-end, from dataset collection to model training and optimization, backed by a \$5,000 BIAP II grant

Software Development Intern

April 2023 – Dec 2023

Orgware Construct Pvt. Ltd. — Nepal

- Built a Flask and React app connecting rural businesses to IT teams, cutting average issue resolution time to under 20 minutes
- Redesigned the MySQL schema and rewrote slow queries to support 100 concurrent users without performance degradation

PROJECTS

Makerspace Sentinel | YOLOv8-CBAM, MediaPipe, TensorRT, Nvidia Orin

Jan 2025 – Present

Real-Time PPE Compliance Detection

- Designed end-to-end model pipeline from dataset selection to a CBAM-enhanced YOLOv8 fine-tuned on three PPE classes
- Validated on 18,000+ unseen images and deployed to NVIDIA AGX Orin via TensorRT at 25–30ms; running live in Ingram, backed by a \$3,000 Makerspace Digital Twin grant

ReelStudio | Python, FastAPI, GPT-4o Vision, FAL, Kling, FFmpeg, Next.js | [Link](#)

Mar 2026 – Mar 2026

AI Cinematic Reel Generator

- Built a 10-layer agentic pipeline converting real estate photos into beat-synced cinematic reels, with three agents classifying rooms, planning shot order, and generating Kling clips cut on BPM timing via FFmpeg
- Won 1st prize of \$1,000 for Best Use of FAL API

NexHacks Hackathon | CMU | Python, LiveKit, WebRTC, OpenAI API, OpenCV | [Link](#)

Jan 2026 – Jan 2026

Autonomous Surveillance Robot

- Built a LiveKit robot on Raspberry Pi with a Next.js web console for real-time home surveillance over WebRTC
- Added OpenCV ArUco marker detection for autonomous indoor navigation and the OpenAI API for two-way voice

AWARDS & HONORS

- \$13,000 in grants and prizes: \$5,000 BIAP II (MOSTify) | \$3,500 Quanta Fellowship (URA) | \$3,000 Makerspace Digital Twin | \$1,000 Best Use of FAL API (ReelStudio) | \$500 Bobcat Innovation Challenge
- President, Data Analytics Association | Secretary, Google Developer Group TXST | President's List and Dean's List, 6 semesters