Aakash Madabhushi

+1 408-210-6122 | aakash.vardhan15@gmail.com | San Francisco Bay Area | Available Dec. 2026
linkedin.com/in/aakash-vardhan | github.com/aakashvardhan | aakashmadabhushi.com
Authorized to work in the U.S. on F-1 OPT, with STEM OPT eligibility — no sponsorship required until 2030

Education

San Jose State University – M.S. in Applied Data Science — Jan. 2025 – Dec. 2026 (expected)
Indian Institute of Science (IISc) – Graduate Certificate in AI & MLOps — Aug. 2024 – Jul. 2025
Illinois Institute of Technology – B.S. in Artificial Intelligence — Aug. 2018 – May 2023

Experience

Graduate Research Assistant — Mar. 2026 – Present
SJSU Embodied AI Lab — *San Jose, CA*

- Taught an SO-101 robot arm to pick and place on its own at **92% success** (46/50 real trials, 2.49 cm mean error) by training an Action Chunking Transformer, a policy mapping camera frames to motion.
- Kept a checkpoint with **67% higher action error** from shipping by picking checkpoints on held-out task accuracy, not validation loss; ruled a 99.9M-param vision-language-action model out of live control at 251 vs 18 ms/step.
- Restored an arm that broke after every power cycle by tracing the fault from camera to motor bus, down to a servo firmware bug that silently reassigned motor IDs; the same trace pinned all 4 failed trials on perception, not control.
- Published an open **49,633-frame** demonstration dataset with zero corrupted episodes by building a LeRobot recording pipeline that syncs arm joints with two 30 FPS cameras and rejects bad episodes before training.

Software Engineer — Jun. 2023 – Jun. 2024
Sapaad (restaurant-management SaaS) — *Dubai, UAE*

- Powered analytics and recommendation-model features for every customer on the platform at **5M+ rows/day**, by building production Python pipelines on Databricks (PySpark).
- Stopped 14 bad-data incidents (schema changes, stale feeds, failed loads) before they reached model training or serving, by adding automated Delta Lake data-quality checks in front of every pipeline.
- Cut load time on **75+ customer dashboards by 43%** (p95 12.4 s to 7.1 s) by profiling the query plans of 10+ Spark jobs, rewriting their slowest SQL, and caching reused results.

Computer Vision Engineer Intern — Sep. 2022 – Dec. 2022
Toothlens (dental AI) — *Remote, India*

- Delivered orthodontists a realistic after-treatment preview of a patient's own smile from one phone photo, as a working end-to-end service, by training a CycleGAN in PyTorch on unmatched before and after photos.
- Increased usable patient photos **20%** across 1,500+ real captures by assembling the unpaired before/after training sets and adding OpenCV preprocessing that screens out or corrects dark, off-angle, and partly blocked shots.

Projects

Low-Latency Super-Resolution on an Edge NPU | *PyTorch, ONNX, INT8, Mobilint NPU* — Apr. 2026

- Ran image super-resolution in real time on a low-power AI chip (**8.55 ms/image**, 1 dB+ PSNR gain) by designing a compact U-Net around its fixed 256×256 input and quantizing it to 8-bit integers via ONNX.

Multimodal Emotion Recognition | *PyTorch, CUDA, cross-attention, Grad-CAM* — Feb. – May 2026

- Matched models **30–50x larger** at 74.6% accuracy (4 emotions, actors unseen in training) with a 1.57M-param network that reads face and voice together, using cross-attention so each stream weighs the other's cues.

Explainable Network Threat Detection | *PyTorch Geometric, Airflow, Snowflake, Llama 3.3* — Feb. – May 2026

- Flagged suspicious machine pairs at **0.90 F1** (0.94 precision) across 27.1M network connections with a graph neural network fed by an Airflow, dbt and Snowflake pipeline; Llama 3.3 explains each alert in 0.87 s.

Technical Skills

Robot Learning: LeRobot, ACT, SmolVLA, VLA models, imitation learning, ROS 2, Isaac Sim, MuJoCo, sim-to-real
Machine Learning: PyTorch, CUDA, transformers, GNNs (PyG), GANs, ONNX, INT8 quantization, W&B
Data & Infra: Python, SQL, PySpark, Databricks, Airflow, dbt, Snowflake, Kafka, Docker, AWS