

Pranav Achar

Dublin, CA | achar.pranav@gmail.com | github.com/PranavAchar01 | linkedin.com/in/pranavachar | pranavachar.com

Robot learning, computer vision and AI agents. Seeking Summer 2027 software and machine learning internships.

EDUCATION

Los Angeles Mission College, A.S. Computer Science & Mathematics (transfer track) **Expected transfer: Fall 2027**
4.0 GPA at LA Mission College. Also coursework at Diablo Valley College and Coastline College. 2026 to present
Coursework: Calculus I and II, Programming in C++; *in progress:* Calculus III, Advanced C/C++; *Spring 2027:* Linear Algebra & Differential Equations, Discrete Mathematics, Data Structures.

PROJECTS AND AWARDS

Etch [\[code\]](#) [\[live\]](#) | **1st Place**, Recursive Agents Hackathon (\$14,000 prize) Oct 2026
• Nemotron agent turns a spoken request into a routed KiCad board of real parts, checked by ERC, DRC and ngspice.

Understudy [\[code\]](#) [\[live\]](#) | **1st Place**, Runway API Hackathon (\$10,000 prize) Sep 2026
• Turns one sentence into Runway-generated, physics-checked robot demonstrations and trains a vision-language-action model (SmolVLA) on a rented GPU. In simulation, the policy scored 43/50 vs 44/50 for perfect scripted demos.

Player Two [\[code\]](#) [\[live\]](#) | **1st Place**, Executable World AI Hackathon; **2nd Place**, Battle of the Personal Brains Sep 2026
• Built an agent that turns a plain-language task into robot training data: it finds Creative Commons video, tracks pose with MediaPipe, and retargets it onto Unitree G1, SO-101 and Franka Panda robots via inverse kinematics in MuJoCo.
• One unattended run screened 68 videos in under 17 minutes and exported the 22 that passed deterministic checks as a LeRobot v3.0 dataset. Rebuilt on AWS Strands with Cognee and Docker.

Foundry [\[code\]](#) [\[live\]](#) | **1st Place**, Superserve Track, terac One-Person-Company Hackathon: agent that builds and runs a storefront for each person in a real audience. Aug 2026

Finalist (top 6 of 115), Abridge x Anthropic x Lightspeed Hackathon; **1st Place**, Apify Track, Build Club Jul 2026

SixthSense [\[code\]](#) | **2nd Place**, ExecuTorch Hackathon (Qualcomm x Meta) Jun 2026
• On-device navigation copilot for blind users, [featured on the PyTorch Foundation blog](#). Built the ExecuTorch vision pipeline and LLM voice agent on a Galaxy S25 that runs Depth-Anything-V2, YOLOv11, Qwen2.5 and OCR fully offline.

HealthFlow [\[code\]](#) | **1st Place**, Scalekit x Apify Hackathon: 9-agent voice-to-EHR pipeline. May 2026

RESEARCH

Independent Researcher, vision-language models (VLMs) for camera-based event monitoring Jul 2026 to present
• Researching vision-language models that run on the camera device, notice and flag observable events in video (a fall, a person down) and save reviewable clips. VLM work in progress.
• Pose-based prototype fall model on a public held-out benchmark it never trained on (UR Fall): flagged all 14 scored falls, a median 2.6 s after onset. 100+ GPU hours of training.

AI Research Intern, UC Santa Cruz AIEA Lab: multi-step logical reasoning in LLMs May 2026 to present
• Measured how a wrong intermediate step carries into the final answer in chain-of-thought reasoning, for three small open LLMs (Qwen2.5 1.5B and 3B, Llama 3.2 3B) run locally on ProofWriter, GSM8K and FOLIO. Wrote checkers that test every stated claim and calculation against ground truth.
• Found that 63 to 72% of wrong GSM8K answers had correct arithmetic on the wrong plan, and that a verifier which removed every checkable step error still did not beat plain chain-of-thought for either Qwen model.

EXPERIENCE

Engineering Intern, Opterra Energy Solutions Jun to Aug 2026
• Designed the system and built the financial model for a Heat-Recovery-as-a-Service proposal to sell AI data center waste heat as district heat (about \$3.5M/yr per 10 MW site, 2 to 4 year payback), and wrote the proposal.

Freelance Web Developer 2024 to 2026
• Built client websites: rebuilt the Peninsula Tamil School site in Next.js (load time 4.2s to 1.8s, Lighthouse 68 to 98), a 150+ product catalog site for lab-materials distributor Anushma LLC, and a site for the cybersecurity nonprofit CPFA.

OPEN SOURCE

18 merged pull requests to statsmodels (14), Biome (2), LanceDB (1) and Lightning CSS (1): 11 statsmodels bug fixes (32M PyPI downloads a month; two issues open since 2018), GLM docs, a LanceDB client fix, and CSS-parser fixes.

SKILLS

Languages: Python, TypeScript, JavaScript, C++, Kotlin, Java, SQL, Rust

ML and robotics: PyTorch, ExecuTorch, pose estimation (MediaPipe, RTMPose), YOLO, MuJoCo, LeRobot, NumPy

Agents and systems: Anthropic SDK (Claude), LangGraph, LangChain, Next.js, React, Node.js, Postgres, Docker