

Arjen Singh

714-801-6366 | arjensingh@outlook.com & asingh0504@g.ucla.edu | github.com/ArjenS-LA |
linkedin.com/in/arjensingh2005

EDUCATION

University of California, Los Angeles (UCLA)

Los Angeles, CA

B.S. Computer Science

June 2027

- Major GPA 3.7
- Focus Areas: Agentic Software Engineering, Computer Vision, NLP, Neural Networks
- Relevant Coursework: Machine Learning, Computer Vision, NLP, Algorithms, Databases, Dynamic Programming, Networks, Operating Systems, Computer Architecture, Robotics, Programming Languages
- Leadership: Finance Director – Punjabis at UCLA, Project Lead – IEEE

RELEVANT EXPERIENCE

LiDAR Engineering Intern @ The Aerospace Corporation

June 2026 – August 2026

- Characterized laser, detector, and supporting hardware behavior of a long-range LiDAR sensor; quantified performance limits and authored a customer-facing capability assessment for space deployment
- Built telescope-array simulation software modeling RSO detection in SEARCH mode and active/passive tracking; benchmarked MHT against PMDAF association algorithms across representative tracking scenarios
- Developed web and embedded tooling, including MCP servers, giving AI agents programmatic access to live production systems for automated testing, telemetry inspection, and iterative hardware/software validation
- Automated beam-profiling test pipelines integrating optics, lasers, control hardware, and characterization software
- Tech: C++, Python, Shell, Git, LLM agents, signal processing

SWE Intern (Manufacturing Operations) @ Astra Space

March 2026 - June 2026

- Deployed an ITAR-compliant LLM gateway (LiteLLM proxy), delivering model access to 150+ users enterprise-wide
- Built mBOM-backed job-queue tooling for manufacturing-floor scheduling and engineering support
- Surfaced real-time build and work-order status to Quality Control, replacing spreadsheet and whiteboard tracking with centralized production visibility and increasing defect detection to 98%
- Coordinated with Quality and Manufacturing Engineering to translate shop-floor requirements into software workflows
- Tech: Python, JavaScript, GraphQL, LiteLLM, LLM agents

SWE Intern (Operations) — Stealth Startup

November 2025 - January 2026

- Built RAG-based AI workflows for enterprise knowledge retrieval and operational task automation
- Developed automation pipelines that translated business processes into executable workflows across internal systems
- Integrated LLM-powered tooling with internal data sources & APIs to reduce manual operational work
- Partnered with Operations and Analytics to map enterprise workflows, identify automation opportunities, and translate business requirements into production AI systems
- Tech: Python, RAG, LLM APIs, REST APIs, Git, Docker

PROJECTS

Dynamic Programming for Belief-Aware Robot Navigation

Tech: Python, Dynamic Programming, POMDPs, Bayesian Filtering, Simulation

- Built a goal-conditioned dynamic programming planner for persistent robot navigation under partial observability, using precomputed state-value tables and online Bayesian belief updates
- Developed myopic and belief-aware policies that optimized expected navigation value while incorporating posterior entropy to balance goal progress with localization uncertainty
- Evaluated planners across symmetric, asymmetric, and bottleneck environments; belief-aware planning preserved 16/17 goal throughput in low-ambiguity settings and reduced MAP localization error from 1.44 to 0.29 in bottleneck scenarios

Context-Aware LLM Safety Supervisor

Tech: Python, PyTorch, HuggingFace Transformers, LLaMA, BERT, pandas

- Built a dual-model, inference-time moderation system that used a dedicated LLM supervisor to detect harmful escalation across multi-turn conversation history before chatbot response generation
- Developed a synthetic adversarial data pipeline generating 400 conversations / 7,494 turns with iterative similarity grading, bypass scoring, and revision against safety supervisors
- Implemented zero-shot, multi-shot, 8B LLM, and HateBERT supervisor variants with conversation-level F1, early-detection, bypass-rate, and latency evaluation
- Achieved 99.5% recall and 81% early detection with the 8B supervisor; characterized precision/recall and adversarial robustness tradeoffs across model architectures