

ALOSH DENNY

AI/ML Engineer · Researcher · Computer Science

✉ aloshdenny@gmail.com · [in alosh-denny](https://www.linkedin.com/in/alosh-denny)
🌐 github.com/aloshdenny · aoxo.pages.dev · 📍 Kerala, India

PROFESSIONAL SUMMARY

Comp Sci grad with extensive experience in AI/ML engineering, specializing in deep learning, LLM development, and AI safety research. Proven track record of deploying production-grade AI systems that drive significant user growth (50% MAU increase) and improve product ratings. Published researcher with 6 peer-reviewed publications in IEEE and Springer. Strong expertise in transformer architectures, multimodal AI, and AI acceleration technologies.

TECHNICAL SKILLS

Programming Languages: Python, C, C++, CUDA, Java, SQL, Verilog, SystemVerilog, Triton, LaTeX

AI/ML Frameworks: PyTorch, TensorFlow, JAX, Hugging Face, LangChain

Semiconductor + VLSI: LibreLane, OpenROAD, YoSys, GTKWave, Verilator, Magic, KLayout, Cocotb, Vivado **Cloud & Tools:** AWS, GCP, Docker, Modal, ChromaDB, Pinecone, Qdrant, ROS2, JMARS

Specializations: LLM Fine-tuning, RAG Systems, Computer Vision, NLP, Quantization, AI Safety

Languages: English (Fluent), Hindi (Fluent), Malayalam (Native), German (A1)

PROFESSIONAL EXPERIENCE

Fullstack AI Engineer

April 2025 – Feb 2026

Discern Security, Sunnyvale, CA (Remote)

- Architected and deployed agentic AI systems for autonomous threat detection, analysis, and cybersecurity response
- Developed **Agentic Studio**, an AI-driven platform for automated remediation, security reporting, and threat charting
- Implemented reinforcement learning algorithms for adaptive cybersecurity protocols and real-time threat mitigation
- Designed AI-aware Data Loss Prevention (DLP) and Data Security Posture Management (DSPM) systems
- Built **AI Query Editor** enabling natural language security investigations and insight generation
- Created internal AI automation workflows that accelerated product specification generation by 60%

Technical AI Safety Research Contributor

May 2025 – Oct 2025

Effective Thesis, Remote

- Investigated correlation between LLM context length and hallucination rates under retrieval-heavy conditions
- Designed controlled experiments varying context-window size to quantify evidence distortion and fabrication patterns
- Analyzed safety implications of context-conditioned hallucinations in goal-misaligned AI outputs
- Published research findings on detection mechanisms for deception-like behaviors in long-context models

AI Safety Researcher (Community)

May 2025 – Oct 2025

Effective Altruism, Remote

- Explored AI alignment risks from long-context model behavior focusing on coherent-but-false reasoning chains

- Connected empirical hallucination findings to deception threat models in extended reasoning scenarios
- Developed scalable evaluation methods for long-context reliability as foundation for anti-deception testing

Artificial Intelligence Engineer

Oct 2024 – Apr 2025

Eduport, Calicut, India

- Led development of multimodal LLMs specialized for Malayalam-language educational applications
- Architected and launched **Adapt AI**, resulting in 21% DAU increase, 23% WAU increase, and 50% MAU growth
- Automated mentor/student call QA, video search, and reasoning systems using RAG architecture on ChromaDB/Qdrant
- Deployed AI services on Modal's serverless infrastructure, reducing inference latency by 40%
- Improved app rating from 4.6 to 4.8, achieving #1 ranking for educational apps on Play Store in Kerala

Data Scientist Intern

Oct 2023 – Jun 2024

NATPAC (Kerala State Council for Science, Technology and Environment), Ernakulam

- Developed ensemble forecasting models for Cochin Port import/export traffic prediction through 2030
- Designed novel hybrid architecture combining LSTM, RNN, Prophet, XGBoost, CatBoost, and NGBoost
- Achieved 92% prediction accuracy using grid search hyperparameter optimization
- Delivered actionable insights for port infrastructure planning and resource allocation

Science and Geology Lead

Feb 2022 – Present

Team Horizon, CUSAT, Kielce, Poland

- Led autonomous Martian rover development for European Rover Challenge competition
- Conducted Martian geological research using JMARS (Java Mission-planning and Analysis for Remote Sensing)
- Built autonomous drone system using MATLAB, Raspberry Pi, and CARLA simulation environment

EDUCATION

Bachelor of Technology in Computer Science

2021 – 2025

School of Engineering, Cochin University of Science and Technology (CUSAT), Ernakulam, Kerala

PUBLICATIONS & RESEARCH

1. **A. Denny** et al., "Regularized Hybrid Deep Learning for DDoS Attack Prediction in Software-Defined IoT Networks," *IEEE ICPC2T*, 2024. DOI: 10.1109/ICPC2T60072.2024.10474620
2. **A. Denny** et al., "T2A: A Novel Text2ASMR Framework for Generating Autonomous Sensory Meridian Response Audio," *Preprint*, 2024. Available: ResearchGate
3. **A. Denny** et al., "IoT and ML for Information Management: A Smart Healthcare Perspective," *Springer Book Chapter 5*, 2024. DOI: 10.1007/978-981-97-5624-7_5
4. **A. Denny** et al., "Time Series Prediction of Container Freight Through Cochin Port Using Machine Learning Models," *CTSEM*, 2025. DOI: 10.1007/978-981-96-8444-1_7
5. **A. Denny** et al., "An Ensembled Deep Learning Model to Improve Diagnosis of Dementia Stages from Brain MRI Images," *IEEE ICSCC*, 2025. DOI: 10.1109/ICSCC66177.2025.11233305
6. **A. Denny** et al., "AI Research Methodologies in Cardiology," *Springer Book Chapter 6*, 2026 (In Press)

TOP PROJECTS

Superfloat · github.com/aloshdenny/superfloat

Implemented AI-first datatype and quantization pipelines optimizing neural networks for efficient hardware acceleration

Superfloat CUDA · github.com/aloshdenny/superfloat.cuda

A C, C++ and CUDA level implementation to support the Superfloat datatype and quantization framework

Superfloat GPU · github.com/aloshdenny/superfloat.gpu

RTL design + GDSII layout for a quad-core Q1.15 GPU built for high-throughput inference-only Superfloat

FLUX.1 [dev] Abliterated · huggingface.co/aoxo/flux.1dev-abliterated

Uncensored Flux Image gen model via surgical unlearning

Swaram · huggingface.co/aoxo/swaram

Advanced Malayalam text-to-speech model based on conditional VAE architecture for natural speech synthesis

Text2ASMR · [Demo Link](#)

End-to-end Diffusion-Transformer vocoder for prompt-guided ASMR audio generation