

VEDANG AVAGHADE

+1(623)-275-1664 vavaghad@asu.edu LinkedIn Portfolio

EDUCATION

Arizona State University | Tempe, AZ

Aug 2025 - May 2027

Master of Science in Computer Science

- **Relevant Coursework:** Natural Language Processing, Agentic AI, Foundations of Algorithms, Statistical Machine Learning, Data Processing at Scale

MIT World Peace University | Pune, India

Jul 2019 - Jun 2023

Bachelor of Technology in Computer Science

- **Relevant Coursework:** Object Oriented Programming, Data Structures, Operating Systems, Database Management Systems, Artificial Intelligence, Machine Learning

WORK EXPERIENCE

Project Engineer | C-DAC, Pune, India

Nov 2024 - Jul 2025

- Directed full-stack engineering for civic portals (Flutter, .NET Core), deploying cross-platform apps to 150,000+ active users and reducing field reporting latency by 40%.
- Structured offline-first MySQL data pipelines and REST APIs with JWT authentication, enabling real-time GPS tracking and analytics for 98,500 km of rural road maintenance.

Software Engineer | Jombay, Pune, India

Oct 2023 - Apr 2024

- Built a mass mailing backend (Ruby on Rails, AWS, Docker) handling 100,000+ messages/month, automating HR workflows and cutting manual processing by 75%.
- Optimized mobile media pipelines by integrating FFmpeg video compression in Dart, reducing payload sizes by 50% and boosting low-bandwidth upload success rates by 40%.

Machine Learning Intern | C-DAC, Pune, India

Jun 2022 - Feb 2023

- Deployed a U-Net computer vision model on HPC GPUs for satellite road extraction, achieving 96% accuracy and accelerating urban planning mapping workflows by 70%.

PROJECTS

PEVAL: Multi-Agent Architecture for Tau-Bench | Python, LangGraph, vLLM | [GitHub](#)

- Developed the PEVAL (Plan-Execute-Validate-Analyze-Learn) LangGraph multi-agent framework (Planner, Executor, Critic) to resolve 6 core LLM failure modes in Tau-Bench.
- Constructed a Memory Kernel for stateful synchronization and a Global Learning Node writing to a persistent JSON store, breaking context deadlocks and enabling dynamic re-planning.
- Benchmarked the Qwen3 model family (4B, 8B, 14B, 32B) via vLLM on NVIDIA A100 GPUs with 4-bit quantization, elevating the 32B model's Pass-5 consistency metric from 0.0% to 30.0%.

VLM Multi-Agent for GUI Automation & Safety | Python, PyTorch, Qwen3-VL | [GitHub](#)

- Designed a Tri-Agent Vision-Language Model pipeline using Qwen3-VL on Intel Gaudi HPUs to execute desktop and web GUI automation (OmniACT) and multi-step workflows (AgentHARM).
- Eliminated geometric drift during UI navigation by engineering a parallel coordinate clustering algorithm, averaging 5 simultaneous inferences to boost OmniACT Action Scores from 40.2% to 51.5%.
- Implemented sequence-level risk-categorization auditing and a Native PyTorch Bridge, reducing the AgentHARM Attack Success Rate (ASR) from a 32.2% baseline to 2.9%.

Connector-Aware LLM Pretraining | Python, PyTorch, Llama 3.2 | [GitHub](#)

- Formulated a Connector Embedding Amplification mechanism for Llama 3.2 3B to counteract Gradient Starvation of structural tokens, targeting 150+ discourse connectors during pre-training.
- Integrated a 1.1x scalar amplification factor at the embedding layer, empirically validating signal persistence through Transformer residual connections despite RMSNorm attenuation.
- Fine-tuned on 64K scientific documents (bfloat16), demonstrating that the 1.1x gradient amplification accelerated logical token updates by 10% without triggering gradient explosion.

DeepDebris: Space Domain Awareness Platform | Python, PyTorch, RL, FastAPI, RAG | [GitHub](#)

- Programmed a satellite tracking platform combining SGP4 orbital physics with a Physics-Informed Neural Network (PINN) to reduce trajectory prediction error by 40%.
- Configured a Proximal Policy Optimization (PPO) reinforcement learning agent for fuel-optimal collision avoidance, deploying the inference engine via a FastAPI backend.
- Incorporated a RAG-driven LLM utilizing vector embeddings to interpret technical aerospace reports, enabling automated multi-agent negotiation for satellite conflict resolution.

TECHNICAL SKILLS

AI/ML & Alignment: Transformers, LLMs, VLMs, GNNs, Diffusion, Scaling, Word2Vec, RLHF, RLVR, PPO, DPO, TPO, GRPO, SFT

Agentic & Orchestration: MCP, Multi-Agent Systems, LangGraph, LangChain, Tool Calling, ReAct, Act, OpenClaw, Google ADK

RAG, Data & Benchmarks: VectorDBs, Graph DBs, Semantic Search, Knowledge Graphs, Tau-Bench, Tau-Trait

Languages & Compute: Python, C, C++, SQL, AWS, Docker, NVIDIA A100 GPUs, Intel Gaudi HPUs, TPUs