Bhushan Ladgaonkar

+91 9152963095 | bhushanladgaonkar@gmail.com | linkedin.com/in/bhushan-ladgaonkar | github.com/beeth73

Professional Summary

Security Researcher and Machine Learning Engineer specializing in **Adversarial AI** and **AI Security**. Published first-author researcher in adversarial robustness; proven experience engineering ML-based Network Intrusion Detection Systems for enterprise infrastructure and building hardware-optimized local LLM inference systems on Apple Silicon. Seeking to leverage expertise in adversarial robustness, simulation, and security architectures.

Education

Fr. Conceicao Rodrigues College of Engineering — Mumbai, India
Bachelor of Technology (B.Tech) in Computer Engineering — *Expected May 2028*

Experience

NetFortress Technologies Private Limited — Remote
Cybersecurity Research Intern — *June 2026 – Present*

- Assisting in product development and technical research for enterprise technology solutions, focusing heavily on cybersecurity and AI infrastructure.
- Supporting AI-related and cybersecurity initiatives through active participation in testing, execution activities, and the preparation of comprehensive technical documentation and reports.

Oil and Natural Gas Corporation (ONGC) — Mumbai, India
Summer Intern – Database and Security Group — *May 2025 – June 2025*

- Developed a prototype **Network Intrusion Detection System (NIDS)** on the UNSW-NB15 dataset (2M+ records) to identify zero-day threats, addressing limitations of traditional signature-based detection.
- Implemented **Random Forest** and **XGBoost** classifiers with **SMOTE**-based class imbalance handling and advanced feature engineering, reducing false positive rates while sustaining high detection sensitivity.
- Collaborated with the Database Group to document threat detection methodologies, translating raw data findings into actionable security intelligence.

Publications

Benchmarking Resilience: Asymmetric Latent Purification Inspired by Generative Diffusion Bottlenecks
Published: May 2026

First & Corresponding Author · IJRIAS · doi.org/10.51584/IJRIAS.2026.110400185

- Proposed the **Asymmetric Latent Purifier (ALP)**: a zero-shot structural defense leveraging a 64-channel bottleneck to shatter adversarial gradients prior to classification, eliminating the need for full generative models.
- Recovered CIFAR-10 accuracy to **32.06%** under iterative **PGD-7** attacks (**25.68 dB PSNR**) while optimizing Apple Silicon inference to **1.25 ms/sample** (**100×+** faster than iterative purifiers).

Technical Projects

Gemma-M4-Agent: Local LLM Inference System | *Python, MLX, FastAPI, ChromaDB, React* — Apr 2026 – Ongoing

- Built a full-stack local LLM agent around **Google Gemma** on Apple Silicon M4 via **MLX framework** – direct hardware-to-model ownership without wrappers such as Ollama or LM Studio; 8-bit quantized model footprint of **8.5 GB** on Unified Memory.
- Engineered a **Hybrid Memory Architecture** decoupling SSD-resident model weights from USB-persistent **ChromaDB** vector store (384-dim embeddings), enabling zero-copy GPU inference with persistent long-term memory across reboots.
- Implemented **dual-gate stop-token defense**: silicon-level halt via MLX kernel token IDs and software-level control-sequence buffer matching, achieving **100% stop-token accuracy** and preventing hallucinated output beyond semantic boundaries.

TISD-Edu-AI: Local-First RAG Tutor | *Python, MLX, ChromaDB, FastAPI, LLMs* — Apr 2026

- Engineered a privacy-focused on-device **RAG** pipeline using **Microsoft Phi-3 (3.8B)** and **ChromaDB**, achieving **41.42 tokens/sec** at 1.72 s latency with a 9.99 GB footprint on Apple Silicon via MLX.
- Ingested 4,600+ pages of academic PDFs into 6,400+ semantic chunks; validated with a custom evaluation suite achieving **81.1% semantic accuracy** across a 100-sample stress test.

Technical Skills

Languages: Python, C, SQL, Java
AI and ML: Machine Learning, Reinforcement Learning, RAG, LLMs, MCTS, PyTorch, MLX, ChromaDB, XGBoost
Security: Network Intrusion Detection (NIDS), Adversarial Robustness, OSINT, Reconnaissance, Linux (Kali, Ubuntu)
Developer Tools: Git, FastAPI, React, Node.js, Docker, VS Code, Jupyter Notebooks

Certifications

CS50 Cybersecurity – Harvard University
Comprehensive study of encryption, web security, network protocols, and defensive programming practices.

Offensive Security Operations – Cybrary
Hands-on training in penetration testing methodologies, exploitation techniques, and red team operations.

OSINT (Open-Source Intelligence) – Cybrary
Techniques for gathering, analyzing, and operationalizing publicly available intelligence for threat investigations.