

UTKARSH DUDEJA

utkarsh.dudeja@outlook.com | +91 8860349638 | India (Open to Remote) | [Portfolio](#)

SUMMARY

Software Engineer with 2.7+ years turning machine learning research into production-grade systems. I work across the full lifecycle—from prototyping and experimentation to optimizing and deploying models at scale—with deep hands-on experience in distributed data processing and GPU-accelerated computing. An IEEE-published researcher with strong mathematical foundations, I'm drawn to problems where rigorous research meets real-world deployment. Currently seeking applied research and AI/ML engineering roles.

SKILLS

ML & Deep Learning | PyTorch (DDP), TensorFlow, Keras, Scikit-learn, NumPy, LSTM, CNN, Transformers, Anomaly Detection, Time-Series Modeling

Advanced ML | Continual Learning, Model Compression (Quantization, Pruning)

NLP & GenAI | LangChain, LLMs, Hugging Face, RAG, Vector Databases, Sequence Tagging, Text Classification, LoRA, QLoRA

MLOps & Infrastructure | Docker, Kubernetes, Ray, Ollama, Redis, OpenSearch, MLflow, CI/CD Pipelines, Cassandra, Linux, Github

Data Engineering | PostgreSQL, Pyspark, Stream & Batch Processing Pipelines, FastAPI, GraphQL, Event-driven Architecture, Distributed Systems

Languages | Python (Advanced), SQL, Rust (Basic), Golang (Basic)

Networking & Protocols | TCP/IP, Modbus, MQTT, OPC-UA, Wireshark, Suricata, Zeek

PROFESSIONAL EXPERIENCE

WhizHack Technologies

Associate Software Development Engineer

Apr 2024 – Present

- Engineered a fully automated CVE intelligence pipeline processing **26 years** of vulnerability data with **100% ingestion reliability** to extract structured threat insights from unstructured vulnerability descriptions.
- Re-architected** the asset discovery system with a modular, event-driven design, cutting execution time from **~3 hours to ~5 minutes (98% faster)** and enabling **zero-downtime deployments** across **10+** latency-sensitive production environments.
- Architected a distributed, high-performance deep packet analysis system featuring adaptive packet capture, producer-consumer architecture with dynamic worker allocation, and parallel protocol parsing across **20+ industrial protocols**, delivering **~60% faster processing** in production.
- Optimized and maintained a real-time threat intelligence system with automated **multi-source ETL extraction** (MITRE, AbuseIPDB, TOR, Feodo, ThreatFox), GraphQL API layer, and concurrent bulk lookup — serving enriched threat data via **FastAPI with Redis caching**.
- Researched and developed** a novel multi-layer behavioral anomaly detection system that learns normal network behavior across topology, routine, and register-level physical values autonomously — enabling detection of **zero-day threats** including unauthorized writes, payload inflation and silent device substitutions under production workloads. **Solution endorsed by CTO; basis for a research paper and patent.**
- Built a causal-attribution and counterfactual analysis layer that **root-causes each detected incident** — ranking per-feature contributions to the deviation and generating **“what-if”** reasoning on which network controls would have prevented or caught it earlier — and correlates related anomalies across time and device pairs to surface **multi-stage activity**.
- Fine-tuned a **BERT-based binary classifier** to detect malicious **XSS payloads**, integrating **MLflow** for metrics logging and model tracking.

Software Development Intern

Jan 2024 – Mar 2024

- Contributed to the core R&D team** focused on developing security systems; gained hands-on experience with OpenSearch, Fluentd, and Kibana for large-scale data analysis and visualization.

PROJECTS

Distributed ML Recommendation System | *Kafka, Spark, PyTorch DDP, Ray Serve, Redis*

- Built an end-to-end distributed ML system with **real-time exactly-once streaming, multi-node training, and fault-tolerant scalable serving**; sustained **382 req/s at 190ms p50** with a **97.5% cache hit rate**, and validated recovery from broker, consumer, Spark, and replica failures via injected-fault tests with **zero data loss**.

ATLAS — Autonomous Multi-Source Threat Correlation & Response Agent | *Raspberry Pi 5, Ollama, Qwen 3.5, Langchain, CTI*

- Built an autonomous SOC triage agent using LangGraph and a locally hosted LLM that correlates Wazuh and Suricata telemetry, enriches indicators with threat intelligence, and drives automated response decisions behind prompt-injection guardrails.

PUBLICATIONS

Decoding Emotions: Emotion Classification from EEG Brain Signals using AI

— IEEE Xplore

Accurate Object Detection for Intelligent Self-Driving Car System for Mountainous Terrains

— Forthcoming

EDUCATION & CERTIFICATIONS

B.Tech in CSE (Hons. Specialization in AI & ML) — Amity University, Uttar Pradesh

2020 – 2024

Duolingo English Test - Advanced: CEFR C1 (135/160) [Link](#)