

Vaddarapu Karthik Yadav

📞 +91-8143094363 ✉ vaddarapukarthikyadav@gmail.com 🔗 LinkedIn 🌐 GitHub

EDUCATION

Vellore Institute of Technology, Vellore
B.Tech in Computer Science and Engineering

Sep 2022 – Jul 2026
CGPA: 8.03 / 10

RESEARCH EXPERIENCE

Hawkes Process + Transformer for LOB Mid-Price Prediction | *PyTorch, C++, Numba* Jan 2026 – Present

- Developed a deep learning system using cross-attention Transformers fusing limit order book (LOB) state representations with Hawkes process intensity features, targeting real-world, large-scale high-frequency trading prediction — a core machine learning research problem in financial microstructure.
- Ran controlled ablation experiments (feature ablation, regime robustness, latency sensitivity) over 5-seed training runs on RTX A4500/A4000 GPU clusters; achieved +6.13pp accuracy over LOB-only Transformer baseline ($p < 0.0001$, Cohen's $d = 7.95$).
- Applied quantization-aware training (QAT/INT8 via ONNX) targeting deployment in large-scale, low-latency intelligent systems; engineered Numba JIT-compiled Hawkes feature modules ($\tau = 50\text{ms}$, $\tau = 500\text{ms}$) for real-time inference.
- Paper accepted: AIAI 2026* (Springer LNCS / IFIP AICT) — contributing to foundational research in machine learning for financial microstructure; demonstrates experience publishing in major international conferences.

Energy-Efficient WSN Protocol | *Python, Network Simulation* | doi 2024 – 2025

- Co-authored peer-reviewed research on energy-efficient networking protocols, contributing to scientific advancement in wireless sensor systems; designed a multi-hop communication protocol improving network coverage by 30%. *Published: Procedia Computer Science, Elsevier (ICECMSN 2025)*.

PROJECTS

High-Performance Quantitative Trading Engine | *C++, STL, Multi-threading, Python* 2025

- Built a high-concurrency trading engine using lock-free atomic operations for nanosecond-level order processing; eliminated race conditions and memory leaks via Valgrind/GDB profiling, addressing real-world, large-scale concurrency and reliability challenges in production software engineering.
- Scripted Python automation tools to benchmark order book latency and analyze system throughput, ensuring the engine meets low-latency requirements.

AI Interview Simulator | *Flask, PostgreSQL, RAG, Google Gemini API* Jan 2025 – Apr 2025

- Built a RAG-powered interview simulator using Google Gemini LLM to generate tailored interview questions; reduced operational costs by ~91% vs. prompt-only approaches by routing to smaller models via retrieval architecture, demonstrating applied machine learning and data science in a production system.
- Paper accepted: IEEE OTCON 5.0* (5th OPJU International Technology Conference, June 2026) — peer-reviewed publication on AI-powered interview platforms.

Data Embedding and Encryption via Video Steganography | *C++, OpenCV, OpenSSL, FFmpeg, CMake* 2025

- Designed a secure file-hiding tool using AES-256 encryption and LSB steganography, ensuring dual-layer confidentiality in video streams.
- Implemented a lossless OpenCV + FFmpeg pipeline with PBKDF2 key derivation, achieving 100% data recovery across tests.

TECHNICAL SKILLS

ML & AI (Deep Learning / Machine Learning): PyTorch, Transformers (cross-attention, multi-head), Hawkes Processes, QAT/INT8, ONNX, NumPy, Pandas, scikit-learn, Numba JIT, LLMs, RAG

Systems & Software Engineering: C/C++, Linux, Multi-threading, Lock-free Data Structures, Memory Management, GDB, Valgrind

Languages: Python, C++, C, SQL, Bash, JavaScript

Infrastructure: Docker, AWS, Git/GitHub, Kafka, PostgreSQL, Redis

AWARDS & ACHIEVEMENTS

- AIAI 2026 (Springer LNCS)** — First-author paper accepted at international peer-reviewed venue; experience publishing in major conferences aligning with scientific advancement goals.
- IEEE OTCON 5.0** — Paper accepted on AI-powered interview platforms, OPJU International Technology Conference 2026.
- Procedia / Elsevier** — Co-author publication, ICECMSN 2025; peer-reviewed contribution to research communities.
- Amazon HackOn Season 5** — Top 75 out of 52,000 participants (link).
- DevSoc'25 Best Project** — Media & Entertainment track, India's largest student-run hackathon (140+ teams).