OMKAR THORVE

+91-8999424867 ◇ Pune, Maharashtra, India
omkar.thorve2001@gmail.com ◇ Linkedin ◇ Github

OVERVIEW

Machine Learning Engineer with hands-on experience in computer vision and LLM-based systems, including Retrieval-Augmented Generation (RAG) and agentic workflows. Currently pursuing M.Tech in Artificial Intelligence and Machine Learning with strong experience in Python, deep learning, and model development.

EDUCATION

M.Tech in Artificial Intelligence and Machine Learning, Symbiosis Institute of Technology 2024 - 2026
Relevant Coursework: Supervised and Unsupervised ML, Deep Learning, Machine Vision, NLP, Reinforcement Learning, Graph Neural Networks, GANs, LLMs, Agentic AI

B.Tech in Mechanical Engineering, Deogiri Institute of Engineering and Management Studies, Aurangabad 2020 - 2023
Relevant Coursework: Image Processing, Linear Algebra, Calculus, Numerical Analysis, Statistics, Control Systems, Python, Linux

SKILLS

Languages	Python, C, SQL
Libraries	Pytorch, Tensorflow, Numpy, Scipy, Flask, NLTK, OpenCV, LangGraph, LangSmith, GCP Gymnasium, Langchain, Hugging Face Transformers
Technical Skills	OOPs, NLP, LLMs, CNN, RNN, GNN, AutoEncoders, Transformers, GANs, MCP, RAG, Image Processing and Computer Vision, Git, Docker, Streamlit, Exploratory Data Analysis, Power BI, Agentic AI

EXPERIENCE

John Deere India Pvt. Ltd. - AI/ML Intern Jun 2025 – Dec 2025

- Applied **object detection and segmentation models** (YOLO, OpenCV, SAM2, DETR, ViT, Grounding DINO) to automate manual processes, reducing task time by **4+ hours** and improving efficiency.
- Fine-tuned **Large Language Models** on custom image datasets, achieving a **50% improvement in accuracy** and significantly reducing hallucinations in outputs.
- Built a **speech-to-text pipeline** using Whisper, GPT-4o Transcribe, GPT-4o Mini, with a JavaScript-based frontend and backend, cutting manual typing effort by **2 hours** and boosting transcription accuracy by **10%**.
- Designed a **multi-agent workflow** leveraging **Graph Neural Networks (Knowledge Graphs)** and deployed on **AWS (EC2, S3, DynamoDB)**, enabling authors to save **4+ hours** in document workflows.
- Implemented **cross-encoder reranker algorithms** for semantic retrieval, outperforming similarity search approaches. Used **Databricks** for large-scale data extraction (EDL) and **OpenAI embeddings + NLP** for evaluation.
- Conducted experiments on **hallucination reduction techniques** for deployed LLMs, improving reliability and robustness of outputs.

K.S.J Recruitment Pvt. Ltd. - Data Analyst Intern Feb 2023 – Feb 2024

- Created and maintained **Power BI dashboards** for recruitment metrics, providing actionable insights to optimize hiring strategies.
- Performed **data cleaning and analysis using Excel and Python**, ensuring accurate and efficient reporting.
- Automated data workflows, improving efficiency and reducing manual effort by 20%.
- Collaborated with cross-functional teams to streamline recruitment data pipelines and enhance decision-making processes.

PROJECTS

Domain-Specific Retrieval-Augmented Generation (RAG) System for HVAC Technical Documentation (In Development): Building a production-grade domain-specific RAG system enabling natural language querying over HVAC technical documentation using LangChain, HuggingFace Embeddings, and transformer-based LLMs. Implementing an end-to-end pipeline covering document ingestion, cleaning, semantic chunking, embedding generation, and vector-based retrieval using Pinecone, grounded in the EnergyPlus engineering documentation corpus. The system aims to deliver context-aware, domain-grounded answers by combining dense vector retrieval with LLM generation, supporting accurate technical Q&A over complex HVAC engineering references.

SwishFit: AI-Powered Basketball Training Platform: Built a full-stack basketball training and workout management web application using React.js (Vite), Node.js, Express.js, and MongoDB Atlas, with Google Gemini 2.5 Flash API for AI-generated personalized workout plans. Implemented role-based features for players (progress tracking, leaderboard, interactive charts), coaches (workout creation, player management, AI assistant), and admins (platform analytics), with JWT authentication, bcrypt password hashing, and Helmet.js security headers. Deployed frontend on Vercel and backend on Render with MongoDB Atlas, achieving 70% bundle size reduction, 90% faster database queries via indexing, and response caching. (Demo Link)

Bosch CNC Machining Anomaly Detector: Developed a CNC Anomaly Detector integrating deep learning and mechanical engineering insights for real-time fault detection, with 4096-point vibration signal analysis and feature extraction strategies aligned with CNC dynamics. Achieved 86% anomaly detection accuracy with a 1.2% false positive rate, 84% precision, and 98% specificity. Deployed an interactive monitoring dashboard using Streamlit for real-time anomaly visualization. (Demo Link)

Agentic Industrial Diagnostics: A Multi-Agent LLM Framework for Explainable Root Cause Analysis: Designed an exploratory multi-agent diagnostic framework integrating LSTM-based autoencoder anomaly detection, ontology-driven knowledge graph reasoning, and LLM-based agents for explainable and adaptive root cause analysis (RCA). The architecture comprises four specialized agents - Diagnostic, Reasoning, Planning, and Learning - enabling hybrid symbolic–subsymbolic reasoning via OWL/RDF knowledge graphs, SWRL rules, and TransE/ComplEx embeddings. Achieved 84.6% RCA accuracy on the AI4I dataset and 72% zero-shot cross-domain accuracy on MetroPT, with dynamic knowledge updates through operator feedback and meta-learning. (Demo Link)

Mild Steel Degradation Analysis Using Microscopic Imaging and Deep Learning: Developed a deep learning-based corrosion detection framework integrating CNNs, color segmentation, and edge detection for real-time mild steel degradation analysis using microscopic imagery. Applied advanced image preprocessing including noise reduction, contrast enhancement, and edge sharpening to improve model robustness and detection reliability. Achieved a 15% increase in F1-score and a 10% improvement in precision over baseline models. Built a web-based corrosion detection application enabling real-time image analysis with corrosion zone highlighting, 90%+ user accuracy in severity classification, and automated corrosion report generation. (Demo Link)

Traffic Signal Control System Building a solution using Yolov5 and Reinforcement Learning. For building an adaptive traffic control system used Actor-Critic method and Yolov5 and vision transformers to monitoring purposes. (Project Docs)

PUBLICATIONS

- **GangaFlow: A Multi-Model Deep Learning Framework for Real-Time River Pollution Detection and Analysis Using Drone Imagery** (IEEE - 2025 ICoICC) (Link)
- **Agentic Industrial Diagnostics: A Multi-Agent LLM Framework for Explainable Root Cause Analysis** (Under Review - Springer)

ACHIEVEMENTS

- **Secured the first runner-up position at The Great Ninja Hack 2K25 (DYPCET Kolhapur)** for the Mild Steel Degradation Analysis Using Microscopic Imaging and Deep Learning Project.
- **Secured the second Runner-up position in the prestigious "Best Manager" competition (NICMAR University)**.
- **Achieved 1st Position in Smart India Hackathon (Internal at SIT)** for the Traffic Signal Control Project.
- **Served as Zonal Basketball Team Captain**, led college team to a zonal championship, and competed in State and National tournaments.

CERTIFICATIONS

- **Generative AI Engineering with LLMs** by IBM.
- **Data Analysis and Visualization with Power BI** by Microsoft.
- **AWS Cloud Training** by Unnati Development Pvt Ltd.