ABHAV VOHRA

(347) 416-7797 voh.abhav@gmail.com github.com/abhavvohra linkedin.com/in/abhavvohra abhavvohra.github.io
Brooklyn, NY

EDUCATION

M.S. in Computer Engineering, New York University — Sept 2022 – Dec 2024
Relevant Coursework: Data Structures, Deep Learning, Reinforcement Learning, Advanced ML, Large Language Models, Natural Language Processing, Probability & Stochastic Processes — GPA: 3.67/4.0

B.Tech in Electronics and Communication Engineering, GGSIPU Delhi, India — Aug 2017 – June 2021

WORK EXPERIENCE

AI Systems Engineer, Easley Dunn Productions, Inc. (Remote, US) — June 2025 - Present

- Developed multi-agent deep reinforcement learning system for Monster Gridiron, a Unity-based American football game, implementing PPO (Proximal Policy Optimization) networks with Unity ML-Agents to enable intelligent offensive team control.
- Implemented self-play training architecture enabling agents to learn coordinated offensive strategies through iterative competition against defensive opponents.

AI Engineer (Intern), Treevah LLC (Remote, US) — Feb 2025 - May 2025

- Developed an intelligent file organization system leveraging OpenAI GPT-4o and LangGraph to create an automated workflow that analyzes uploaded text documents, extracts semantic content, and categorizes files based on content similarity.
- Implemented human-in-the-loop validation with user confirmation prompts before file transfers to ensure accuracy and user control.

ML Engineer, Vitalth Forgers Pvt. Ltd. (New Delhi, India) — May 2022 - Aug 2022

- Developed influenza detection system using fine-tuned transformer models on vital sign time-series data (temperature, Pulse, Respiratory Rate, SpO2), achieving 88% accuracy.
- Implemented an LSTM network for influenza classification, processing sequential vital sign data to distinguish flu from other respiratory conditions with 18% reduction in false positives.

Software Engineer, Antriksh Labs Pvt. Ltd. (Remote, India) — June 2020 - Dec 2021

- Developed end-to-end anomaly detection pipeline using scikit-learn, and unsupervised learning to analyze sensor data from robotics equipment, predicting failures with 85% accuracy and optimizing maintenance schedules.
- Built time-series forecasting models with PyTorch for multi-sensor IoT data streams, creating real-time dashboards that achieved 30% reduction in unplanned downtime across robotics automation systems.

PROJECTS

Arxiv Research Agent

- Developed an automated academic research system using LangGraph's multi-agent architecture, integrating ArXiv API with StateGraph orchestration to streamline scholarly paper retrieval, analysis, and synthesis for researchers and academics.
- Built comprehensive workflow that leverages reflection agents for iterative improvement, vector database integration for semantic search, and structured report generation with proper citations to enhance research productivity and accuracy.

Multi-Agent Technical Documentation RAG System

- Built a technical documentation assistant using AI agents (API, code generation, troubleshooting), processing open-source codebases with query routing and parallel execution to provide comprehensive developer solutions.
- Deployed FastAPI service with ChromaDB vector database, semantic chunking strategies, and confidence scoring systems for multi-domain knowledge synthesis.

TECHNICAL SKILLS

Programming Languages:	Python , SQL , C, C++, C#
ML/AI Frameworks:	PyTorch , Scikit-learn, Hugging Face Transformers, SpaCy, XGBoost, OpenCV
AI/LLM Engineering:	LangChain, LangGraph, LlamaIndex, CrewAI, Neo4j, RAG Systems, Graph RAG
MLOps & Development:	MLflow, Weights & Biases, Gradio, FastAPI, Docker, Kubernetes, Jupyter, PySpark, AWS (SageMaker, EC2), Git, CI/CD (Jenkins/GitHub Actions), Airflow
Deep Learning & Architecture:	Transformers, Attention Mechanisms, Natural Language Processing (NLP), CNNs, GANs, VAEs, Diffusion Models, Graph Neural Networks, Multi-modal Models, Large Language Models (LLMs), RAGs, Prompt Engineering