MIHIR AGARWAL

+1 (646) 235-7363 | ma4874@columbia.edu | New York, NY, USA | linkedin.com/in/mihir-agarwal/

EDUCATION

Columbia University **August 2025 - December 2026**
Master's, Data Science
- Unsupervised Learning, Applied Deep Learning, Probability and Statistics, Competitive Coding

Vellore Institute of Technology **July 2018 - July 2022**
Bachelor's, Computer Engineering *GPA: 8.97*
- Machine Learning, Artificial Intelligence, Databases, Operating Systems, Data Structures and Algorithms,

PROJECTS

NeuroLOB: Generative Market Intelligence - *Link to project* **New York, NY, USA**
Advisor: Prof. Nakul Verma | Columbia University *August 2025 - Present*
- Engineered an unsupervised hybrid generative architecture merging Neural Point Processes (NPP) with Denoising Diffusion Models and Causal Transformers to simulate asynchronous Limit Order Book (LOB) events.
- Benchmarked against 1D-Convolutional cGAN and NPP baselines, demonstrating superior fidelity in Kolmogorov-Smirnov (KS) scores and accurately reproducing market stylized facts like volatility clustering
- Deployed the model as a live Streamlit web application, enabling users to generate custom financial time-series data and perform real-time statistical evaluations of synthetic market scenarios.

PROFESSIONAL EXPERIENCE

Shell **Bengaluru, India**
Machine Learning Engineer *August 2022 - August 2025*
- Built an AI-powered HR Chatbot using a Retrieval-Augmented Generation (RAG) pipeline with LangChain and GPT, indexing Shell HR policy documents for quick, context-aware responses.
- Automated 68% of HR employee inquiries, streamlining internal communications, cutting response times by 50%, and saving $102k annualized savings based on cost-per-interaction and volume in H2 FY2025 in operational costs.
- Implemented document chunking, embedding generation, vector indexing, and semantic retrieval for both projects to ensure robust, scalable, and efficient intelligent automation systems.
- Developed a web application for Test Case Prioritization, enabling users to upload test cases, apply clustering and supervised learning techniques, and receive prioritized results, improving test suite accuracy and efficiency.
- Engineered machine learning models leveraging features like execution history, step complexity, and outcome status, leading to a 35% reduction in execution time and cost savings of $10k.

LISITE Lab Isep **Paris, France (remote)**
Researcher *January 2022 - August 2022*
- Co-authored and published a Springer journal paper on ML-based resource management in fog computing, proposing latency-reduction strategies that improved fog–cloud integration efficiency.
- Designed and evaluated scalable deployment models by estimating infrastructure requirements and performance metrics, ensuring optimized allocation of computing resources.

Vidrona **London, UK (remote)**
Machine Learning Intern *July 2020 - August 2022*
- Developed predictive maintenance solutions for insulator caps and guy adjusters by processing drone-captured images and videos with deep learning models (YOLOv3/5/8, Faster R-CNN) to detect electrical faults and automate inspections.
- Enhanced model accuracy through advanced image preprocessing techniques, including Histogram Equalization, noise reduction, and contrast adjustment.
- Optimized detection pipelines by fine-tuning thresholds and anchor boxes, saving 300+ hours of manual labor weekly.

SKILLS

Programming Languages:: Python, C/C++, Java, JavaScript, SQL, HTML/CSS, PHP, XML
Framework & Libraries: Transformers, TensorFlow, Pytorch, Scikit-Learn, Numpy, Pandas, Matplotlib, OpenCV
APIs & Backend Development: REST API, FastAPI, Flask, Node.js, .NET 8
ML Infrastructure: Model Evaluation, Optimization, Deployment, Monitoring, Debugging
ML & AI: Machine Learning, Deep Learning, Transformers, BERT, GPT-4, NLP, RAG, Large Language Models, Computer Vision, Feature Engineering, BLEU, ROUGE, Artificial Intelligence, Encoder Models

ACHIEVEMENTS AND EXTRA CURRICULAR

- Won Best Domain prize in IvyHacks : a joint hackathon hosted by 6 Ivy League universities
- Vice President, Shell Toastmaster Club: Completed Level 3 and conducted 45+ events to enhance public speaking and leadership skills

COURSES AND CERTIFICATES

PCAP Certified Python Programmer (Python Institute) | Cpp DS and Algorithms, Competitive Coding (Coding Ninja) | Neural Networks and Deep Learning (Coursera) | Deep Learning Nanodegree (Udacity)