




Adit Whorra

✉ aditwhorra@gmail.com 📞 +91-98992-45540  LinkedIn  GitHub  Medium

Machine Learning Scientist with 4+ years of experience developing and deploying scalable NLP and computer vision systems. Published researcher (IJCAI) with a strong record of translating deep learning research into production-ready solutions. Skilled in building end-to-end ML pipelines using Python and PyTorch, and experienced with advanced architectures including Transformers, GNNs, and RNNs. Driven to bridge cutting-edge AI research with real-world business impact through scalable, enterprise-grade solutions.


EDUCATION

- Master of Science (M.S.) in Computer Science**, *Technical University of Delft, The Netherlands* 2023 – 2025
Specialization: Artificial Intelligence
Grade: 8.8/10 (Cum Laude)
Thesis: "Modeling Multivariate Time Series with Spatio-Temporal Simplicial Complexes"
Relevant Courses: Natural Language Processing, Deep Reinforcement Learning, ML for Graph Data, Deep Learning
- Bachelor of Science (B.S.) in Computer Science**, *Ashoka University, India* 2015 – 2019
Specialization: Computer Science (Major), Mathematics (Minor)
Grade: 3.79/4.00 (Magna Cum Laude)
Thesis: "Generative Question Answering System for Legal Contracts"
Relevant Courses: Machine Learning, Unstructured Information Processing, Information Retrieval

WORK EXPERIENCE

- Associate Machine Learning Scientist - II**, *Wadhvani AI* Dec 2025 – Present
Building ASR models on Indic languages at scale in India.
- Teaching Assistant**, *Technical University of Delft* 2023 – 2025
- Led weekly labs for MSc courses like ML for Graph Data and Deep Learning, teaching core ML along with GNNs, CNNs, and transformers to 100+ graduate students.
 - Developed PyTorch tutorials and coding exercises, bridging ML theory with practical implementation.
 - Mentored student research projects, guiding experimental design and technical implementation.
- Founding Machine Learning Engineer**, *A'Alda (Japan-India)* 2022 – 2023
- Led a team of 5 engineers to develop a first-of-its-kind computer vision-based biometric recognition system for animals, achieving 90+% accuracy and establishing technical feasibility for automated veterinary identification systems.
 - Designed and deployed a Flask-based REST API application for multi-site data collection across 5 veterinary hospitals (India & Japan). Integrated an Airflow-driven ETL pipeline with OpenCV, CVAT, and AWS S3 to automate data ingestion, preprocessing, and annotation, enabling scalable dataset versioning for 3,000+ unique animal samples.
 - Developed and optimized deep metric learning models in PyTorch for feature extraction, experimenting with triplet and contrastive losses, ResNet backbones, augmentations, and hard example mining. Built structured benchmarking workflows using MLFlow to streamline model iteration and evaluation.
 - Deployed recognition system using TorchServe and Docker with Django, PostgreSQL, Redis, and Celery, enabling scalable and reliable model inference. Implemented a human-in-the-loop evaluation workflow for continuous online model assessment and improvement.
 - Mentored junior engineers and established best practices, including code reviews and comprehensive documentation that enabled seamless knowledge transfer and system maintainability.
- Machine Learning Engineer**, *SpotDraft* 2019 – 2022
- Built and deployed scalable NLP systems using RNNs and transformer-based architectures (BERT, RoBERTa & DistilBERT) to automate contract drafting, compliance checking, and metadata extraction on AWS infrastructure with PyTorch, Flask, and Kubernetes.
 - Pre-trained domain-specific transformer models on 100K+ legal documents, improving classification F1 by 20% over non-pretrained baselines. Applied attention-head pruning and quantization to reduce model size by 35% and double inference speed for real-time contract analysis.
 - Developed few-shot classification systems using entailment-based prompt engineering, reducing new client onboarding time from weeks to days and creating a reusable framework for rapid adaptation to new legal domains.
 - Built dense retrieval systems using transformer embeddings and FAISS for semantic contract search. Led construction of a 200K+ entity knowledge graph enabling downstream GNN models for clause-level reasoning and compliance risk detection.
 - Owned the end-to-end ML development lifecycle, including model packaging, RESTful API integration, CI/CD automation, and scalable model serving on AWS using Docker, TorchServe, and Kubernetes.
 - Conducted applied R&D by reproducing and adapting recent NLP research papers, integrating promising techniques into production systems, and guiding internal model improvements.
- Research Intern**, *New York University (Stern School of Business)* 2018 – 2019
- Built interactive dashboards analyzing 900+ economic variables across 150+ countries for policy research at NYU.
 - Developed end-to-end data pipelines in R and created visualizations using Tableau for cross-country analysis.

SELECTED PUBLICATIONS & RESEARCH

- Exploring Equity of Climate Policies using Multi-Agent Multi-Objective Reinforcement Learning** , 2025
Published at IJCAI 2025
- Published first-author paper at IJCAI 2025, presenting an AI framework that helps countries negotiate climate policies by automatically finding strategies that balance conflicting objectives like economic growth and emission control.
 - Developed the framework using Multi-Objective Multi-Agent Reinforcement Learning, implemented using MOMALand and PyTorch, where agents learn emission control and savings rate policies optimizing multiple conflicting objectives simultaneously while operating in a climate-economy simulation environment.
 - Released an open-source implementation enabling stakeholders to explore these trade-offs in real-world applications, supporting international climate negotiations and policy decision-making.
- Modeling Multivariate Time Series with Spatio-Temporal Simplicial Complexes (Master's Thesis)**, 2025
Advisor: Dr. Elvin Isufi, TU Delft
- Pioneered a novel spatio-temporal GNN architecture that extends traditional pairwise graph representations to model higher-order multi-node relationships that better capture real-world data dependencies in multivariate time series data (awarded Cum Laude).
 - Conducted comprehensive evaluation against 10+ baseline models across forecasting and imputation tasks in three real-world domains (traffic, energy, air pollution), implementing hyperparameter optimization using Optuna, experiment tracking with Weights & Biases, and large-scale experimentation on HPC clusters with GPU acceleration.
 - Achieved comparable performance to state-of-the-art methods with significantly improved parameter and data efficiency and stronger theoretical foundations through spectral analysis.

PROFESSIONAL COURSES

- AI Agents Fundamentals**, *HuggingFace* 2025
Training in modern LLM agentic frameworks like LangChain, LlamaIndex, and HuggingFace smolagents.
- AI Alignment**, *BlueDot Impact* 2024
Program on AI alignment and safety principles for LLMs and advanced AI systems.

SKILLS

Programming Languages

Python · C++ · R

Development & Deployment

Docker · AWS · Kubernetes · FastAPI · Flask · Django · TorchServe
· SQL · PostgreSQL · Redis · Celery · AirFlow · CI/CD · Git

Optimization & Experimentation

Optuna · Weights & Biases · MLFlow · TensorBoard

ML Frameworks & Libraries

PyTorch · PyTorch Geometric · TensorFlow · Hugging Face ·
LangChain · LlamaIndex · OpenCV · FAISS · scikit-learn

Data Processing & Visualisation

NumPy · Pandas · Tableau

AWARDS

- IJCAI Conference Grant**, *International Joint Conference on Artificial Intelligence* 2025
Awarded a competitive travel and conference grant to present my research and serve as a volunteer at IJCAI in Montreal, Canada.
- Academic Honours** 2019 – 2015
Graduated Cum Laude from TU Delft (M.S., 2025) and Magna Cum Laude from Ashoka University (B.S., 2019).
- University Organizer Fellowship Award**, *Open Philanthropy* 2024
Awarded personal grants (>\$8000) to run an AI Safety community and pursue independent upskilling in AI Safety.
- Research Excellence Award**, *Ashoka University* 2019
Awarded to the graduating student with the best thesis project in Computer Science at Ashoka University.