Twissa Modi

Machine Learning Engineer

Singapore +65-87402729 twissamodiofficial@gmail.com linkedin.com/in/twissamodi
github.com/twissamodiofficial

SUMMARY

Machine Learning Engineer with a focus on applied machine learning, deep learning, and model development, with experience across time-series forecasting, computer vision, NLP, and large language models. Experienced in developing, fine-tuning, evaluating, and deploying ML models using PyTorch, Hugging Face, and Scikit-learn, with 3+ years of production software engineering experience at Microsoft. Currently pursuing an MS in Artificial Intelligence at Nanyang Technological University (NTU), Singapore.

PROFESSIONAL EXPERIENCE

Software Engineer Intern – AI Agents 07/2026 – 12/2026 | Singapore
Robert Bosch (SEA) Pte Ltd
- Designing an NL-to-SQL analysis system via Copilot Studio custom connectors to on-prem services, grounding query generation in database metadata to enable natural-language cost-centre spend analysis for business users
- Building an Agent-to-Agent (A2A) bridge between Copilot Studio and Gemini Enterprise agents, including secure Azure AD authentication, to enable cross-platform data and workflow interoperability

Software Engineer 08/2022 – 07/2025 | Hyderabad, India
Microsoft
- Led end-to-end telemetry implementation for Windows Share Sheet, defining product usage metrics and partnering with the Data Science team to deliver actionable analytics dashboards.
- Re-architected backend services for Microsoft Family Safety using C#/.NET, reducing compute costs by 16%, and developed the App Category feature for more granular application controls.
- Co-invented and filed a U.S. patent, “Drop Tray Engine(s) for Providing a User-Based Drag-and-Drop Action Experience” (Docket 503629-US01; filed Feb 2025), for a Windows feature reaching millions of users.

Software Engineer Intern 02/2022 – 07/2022 | Bangalore, India
McKinsey Digital Labs
Built interactive data visualization dashboards using React.js for client-facing applications, translating complex datasets into actionable insights for business stakeholders.

EDUCATION

Master of Science in Artificial Intelligence 08/2025 – 12/2026 | Singapore
Nanyang Technological University
GPA: 4.7/5.0 Relevant Coursework: Machine Learning, Deep Learning, Natural Language Processing, Large Language Models, Computer Vision

Bachelor of Engineering in Computer Engineering 07/2018 – 06/2022 | India
Thapar Institute of Engineering & Technology
CGPA: 9.54/10 Scholarship holder

TECHNICAL SKILLS

Programming & Data
Python, C++, C#, SQL, Pandas, NumPy, Matplotlib, Seaborn

Machine Learning
Scikit-learn, Feature Engineering, Model Evaluation, Cross-Validation, Hyperparameter Optimization

Deep Learning & Computer Vision
PyTorch, CNNs, Transfer Learning, Model Fine-tuning, ResNet, GradCAM

LLM & NLP
Hugging Face Transformers, BERT, GPT, LLaMA, SFT, LoRA/PEFT

Infrastructure
FastAPI, Docker, Hugging Face Spaces, Azure, SLURM/GPU Clusters, Git, GitHub

Agentic AI & RAG
LangGraph, LangChain, RAG, Vector Databases, Tool Calling

PROJECTS

Day-Ahead Electricity Load Forecasting
Time Series Forecasting | *Probabilistic Machine Learning*
- Built a day-ahead electricity load forecasting pipeline using ENTSO-E grid data, achieving 2.65% MAPE with calibrated probabilistic forecasts using Conformalized Quantile Regression (CQR) and 89.9% prediction interval coverage.
- Deployed an automated forecasting pipeline with daily data ingestion, model inference, and monitoring using GitHub Actions and Supabase.

Rossmann Sales Forecasting
Time Series Forecasting | *Machine Learning*
- Built an end-to-end LightGBM sales forecasting pipeline for the Rossmann Kaggle competition, using Optuna hyperparameter optimization, log-transformed targets, and leakage-safe temporal features to achieve 0.1199 RMSPE on the held-out test set.
- Engineered lag and rolling features (49/56/63 days) alongside promotion and competition-period features to capture temporal, seasonal, and business-driven demand patterns.

Serengeti Wildlife Classifier
Deep Learning | *Computer Vision*
- Fine-tuned ResNet18 for 11-class wildlife image classification using transfer learning, data augmentation, and class-balanced sampling, achieving 84.0% test accuracy on a held-out split grouped by camera site.
- Applied GradCAM for model interpretability and deployed the inference service using FastAPI and Docker for reproducible serving.

Banking Intent Classifier
Efficient LLM Fine-tuning
- Fine-tuned Qwen2.5-0.5B on BANKING77 using LoRA and supervised fine-tuning (SFT), achieving 84.6% accuracy while training only 0.12% of model parameters.

CERTIFICATIONS & ACTIVITIES

Machine Learning Specialisation
DeepLearning.AI (Coursera)

Foundations of AI
Microsoft

Volunteer
Women Devs SG

Member
Women in Tech – Microsoft