

Swarnim Tripathi

+91 7607260185 | swarnim.tr@gmail.com | linkedin.com/in/swarnim-trip | github.com/tripathiji1312 | [Portfolio](#)
Lucknow, Uttar Pradesh, India

EDUCATION

Vellore Institute of Technology

Bachelor of Technology in Computer Science and Engineering

Chennai, India

July 2024 – May 2028

TECHNICAL SKILLS

Programming Languages: Python, C++, JavaScript, SQL, Kotlin, Rust, HTML, CSS

Machine Learning and AI: PyTorch, TensorFlow, scikit-learn, Transformers, Fine-tuning, LoRA, Neural Networks, NLP, Reinforcement Learning, Stable-Baselines3, LSTM, DQN, Sentence-Transformers

Frameworks and Libraries: FastAPI, Flask, Django, Pandas, NumPy, Matplotlib, XGBoost, Docker, Docker Compose

Databases and Storage: PostgreSQL, MongoDB, SQL, NoSQL, TimescaleDB

DevOps and Cloud: Google Colab T4 GPU, Git, Linux, CI/CD, Vercel, Render, Netlify, Model Deployment

Data Science: Feature Engineering, Data Preprocessing, Statistical Analysis, Web Scraping (Stack Exchange API), A/B Testing

PROJECTS

Aura - AI-Powered Diabetes Management | *Python, Flask, LSTM, DQN, PostgreSQL, NLP* [GitHub](#) | 2025

- Engineered 3-layer LSTM (50 neurons/layer) on OhioT1DM medical dataset, forecasting glucose 1-2 hours ahead with 85% accuracy (± 20 mg/dL); predicted 12 points ahead at 10-min intervals
- Trained Deep Q-Network in simglucose simulator with custom reward function penalizing hypoglycemia 10x; state space: [glucose, trend, active_insulin]; achieved 15% Time-in-Range (70-180 mg/dL) improvement
- Implemented NLP engine using regex extraction to parse dietary/activity logs; Orchestrated full-stack Flask app with real-time Chart.js dashboards and automated FPDF medical reports; targets 422M+ diabetes patients
- Containerized with Docker Compose, released on Vercel (frontend) and Render (backend + PostgreSQL/TimescaleDB); optimized time-series queries for 1000+ glucose readings/user

Promptimus - AI Prompt Optimizer | *Python, PyTorch, T5, LoRA, RAG*

[GitHub](#) | 2025

- Fine-tuned 220M-parameter T5 via LoRA on 200K custom prompt-optimization pairs; implemented RAG pipeline with constrained beam search guided by Sentence-Transformers semantic similarity; ROUGE-L 0.72 (28% improvement)
- Curated 324 high-quality Q&A pairs from Stack Exchange API; fine-tuned TinyLlama-1.1B via QLoRA (4-bit quantization) on Colab T4 GPU in 2 minutes; achieved 89% command generation accuracy
- Engineered safety validation layer blocking 50+ dangerous command patterns; implemented dry-run preview system; identified parser limitation (failed 3/7 test cases with \$ prefix), proposed regex-based solution
- Production-ready Docker container with sub-2s inference; logged 150+ interactions to JSONL traces; open-sourced with 95% safety compliance score

LearnBuddy - Adaptive Learning Platform | *Python, Flask, PostgreSQL, RL*

[GitHub](#) | 2025

- Architected epsilon-greedy Multi-Armed Bandit with heuristic overlay, analyzing sliding window of 12 user attempts; balanced exploration-exploitation for personalized difficulty adaptation
- Set confidence threshold at 75% for difficulty advancement; heuristic triggers: 2 consecutive correct (80% success) → level up; ≥ 30% success → intervention; maintained 92% user satisfaction
- Integrated Sentence-Transformers for semantic answer grading; built WCAG-compliant accessibility (keyboard nav, high-contrast themes); deployed via Docker Compose to Netlify with 15+ beta testers

CERTIFICATIONS

Machine Learning Specialization

DeepLearning.AI and Stanford University

June 2025

Coursera

Data Analysis with Python

freeCodeCamp

May 2025

freeCodeCamp