

Aayush Tiwari

aayushtiwari0110@gmail.com | 9336631824 | [linkedin.com/in/aayush-tiwari-4751a6253](https://www.linkedin.com/in/aayush-tiwari-4751a6253)
github.com/theeCzar | theeCzar.github.io

EDUCATION

Shri Ramswaroop Memorial University, Lucknow

2022 – 2026

Bachelor of Technology in Computer Science and Engineering

PROJECTS

Distributed URL Shortener + Analytics Platform

[Github.com/theeCzar/distributedURL](https://github.com/theeCzar/distributedURL)

- Built a URL shortener that uses Redis cache to handle repeat clicks without hitting the database, this makes redirects much faster
- Added a token bucket rate limiter using a Redis Lua script to handle many requests at once without breaking
- Created an analytics system with RabbitMQ that saves click data in batches, reducing the number of database calls

LLM Gateway: Eval, Guardrails & Security

[Github.com/theeCzar/LLM-Gateway](https://github.com/theeCzar/LLM-Gateway)

- Built a security-focused FastAPI gateway that sits in front of an open-source LLM served via Hugging Face's Inference API
- Added input guardrails that block prompt-injection attacks before they reach the model, this saves money on API calls and keeps the system safe
- Built an automated test harness that runs accuracy and safety checks on every code push, with a red-team suite that measures how well the guardrails catch attacks
- Added a tamper-evident audit log with SHA-256 hash chaining if anyone edits a log entry, the system immediately flags it
- Deployed the gateway to Hugging Face Spaces using Docker, with zero need for a local GPU

F1 Race Strategy Engine: Undercut / Overcut Prediction

[Github.com/theeCzar/F1](https://github.com/theeCzar/F1)

- Built a tool that predicts when a driver should pit, combining a tire-wear model with a survival model for pit timing validated on race-by-race splits to avoid data leakage
- Designed a Postgres database with 6 tables to store lap data, predictions, and test results; made sure re-running the same data never creates duplicate rows
- Set up two Airflow pipelines: one for daily data updates and another that only triggers model training when new data actually arrives

TECHNICAL SKILLS

Languages: Python, C, C++, JavaScript, HTML, SQL

ML / DL Frameworks: PyTorch, TensorFlow, Scikit-learn, Hugging Face Transformers, XGBoost, pandas, NumPy, Django, FastAPI, PySpark, LightGBM

Focus Areas: NLP (transformers, sentiment analysis), Generative AI, Computer Vision, Deep Learning, Applied Machine Learning

Generative AI / LLMs: Prompt Engineering, LLM Integration, RAG, Embeddings, MCP, Agentic AI

Tools & Platforms: AWS (EC2, Lambda, S3), Git, GitHub, Google Colab, Docker, CI/CD, GitHub Actions, Kafka, RabbitMQ, PostgreSQL, Vim, Hugging Face

Languages Spoken: Fluent in English and Hindi

INTERESTS

ML/AI research, football, formula 1, reading books, listening to music, making new food