

Mahmoud El-Bahnasawi – ML Engineer

6 October City, Egypt · +20 109 261 5675 · m.elbahnasawi.ai@gmail.com

[LinkedIn](#) | [GitHub](#) | [W&B](#)

Summary

ML Engineer with **hands-on experience** building and deploying efficient **NLP** systems. **Fine-tuned** BERTweet on 530k+ tweets to achieve **84.5% macro-F1**, delivering **sub-second inference** via a FastAPI endpoint. I am eager to apply and deepen my skills in **NLP** and **MLOps** to scale **production-grade** intelligent systems.

Experience

AI Research Intern [GitHub](#) [Streamlit Demo](#)

Sep 2024 – Oct 2024 | Zewail City of Science and Technology (*Supervised by Dr. Doaa Shawky*)

- **Architected** a **graph-based extractive summarization** framework combining **SBERT** embeddings with **PageRank**, reaching **ROUGE L2-Norm ≈ 0.41** on CNN/DailyMail at the best config.
- **Engineered** a **multi-embedder pipeline** (SBERT, TF-IDF, BoW) with entity-aware graph construction and a reproducible **ROUGE** evaluation suite.
- **Ran** an **ablation study** on graph connectivity types (**sentence–sentence**, **sentence–entity**, **entity–entity**), showing sentence–sentence edges as the primary driver
- **Deployed** an interactive **Streamlit** app for real-time summarization and graph visualization; optimized the pipeline to **~ 0.04 s/article** through efficient similarity computation and sparse graph building.

Publications

Efficient Hate Speech Detection: A Three-Layer LoRA-Tuned BERTweet Framework

November 2025 · arXiv:2511.06051 [cs.CL] | [PDF](#)

- **Introduced** a parameter-efficient framework that achieves 94% of the performance of a 14B parameter **LLM** (SafePhi) using a model 100x smaller (**134M vs. 14B parameters**).
- **Demonstrated** superior efficiency via **LoRA fine-tuning**, requiring **only 1.87M trainable** parameters (1.37% of BERTweet's total) while maintaining competitive performance.
- **Outperformed** traditional BERT-based approaches through **strategic** dataset unification and **optimized** fine-tuning methodology.

Projects

AI-Powered Hate Speech Moderation [GitHub](#) · [Streamlit Demo](#) · [W&B Logs](#)

June 2025 · Bachelor Graduation Project | BERTweet, LoRA, Hugging Face, W&B, FastAPI, Supabase

- **Engineered** a production-ready three-tier moderation stack—rule-based slur filter → LoRA-tuned BERTweet (530K tweets) → Supabase feedback loop for continuous learning from real-world deployment.
- **Optimized** training efficiency using W&B experiment tracking and free Kaggle T4 GPU, completing 3 epochs in ≈ 2 hours with 84.5% Macro-F1 / 0.68 MCC, beating MetaHate by +4.5 Macro-F1.
- **Deployed** end-to-end system with sub-second inference via FastAPI, complemented by a Streamlit demo and live browser extension filtering Twitter & YouTube in real-time.

- **Research** published as [arXiv:2511.06051](#), demonstrating 94% performance of 14B parameter models with 100x smaller architecture.

Analyzing U.S. Crime Data · [GitHub](#)

Mar 2024 | Python, Scikit-learn, XGBoost, Plotly, UMAP, RFE, VIF

- **Analyzed** 1991-2021 U.S. crime data: Identified larceny-theft dominance, Alaska's 40% violent crime rate, and post-2020 homicide surge via Plotly dashboards. Tested firearm laws' impact (t-test, $p=0.383$), finding **no significant** crime reduction link (2010-2020).
- Built 98-100% **accurate** recidivism predictor: RFE/VIF + XGBoost **outperformed** PCA/UMAP models in **efficiency** for 3-year re-offense risk, following KNN imputation and feature engineering.

Walmart Sales Forecasting Dashboard · [Live Demo](#) | [GitHub](#)

Feb 2024 | Python, Plotly Dash, Scikit-learn, KMeans, GMM, Docker, Render

- Built a **Docker-packaged, multi-tab Plotly Dash dashboard** and deployed it to Render, enabling analysts to explore three years of Walmart sales across 45 stores and 98 departments (time-series, holiday vs non-holiday, CPI & fuel overlays).
- **Clustered stores (4 tiers, \$88 M → \$25 M) and departments (3 tiers: 22 top, 55 mid, 4 laggards) using K-Means & GMM**, guiding merchandising strategy and targeted marketing spend.

Education

- B.Sc. Communications & Information Engineering – Zewail City of Science & Technology, June 2025
- **Relevant courses: ML/DL; CV; Data/Stats; SW Eng; Cloud.**

Courses

Applied Data Science Lab [[Badge](#)]

Sep 2023 | WorldQuant University

- Completed 8 end-to-end projects encompassing data extraction, cleaning, exploratory analysis, SQL querying, visualization, and full machine-learning pipelines with Pandas and Scikit-learn.

Machine Learning Specialization [[Certificate](#)]

Nov 2023 | Coursera (Andrew Ng)

- supervised/unsupervised ML, neural nets, recommenders, RL fundamentals

Skills

- **Core Stack:** SQL · Pandas, PyTorch · scikit-learn · HF Transformers
- **MLOps & Deployment:** Docker, FastAPI · Weights & Biases · HF Spaces · Render · Supabase
- **Familiar & Learning:** RAG, Azure
- **Soft skills:** Systems thinking, collaborative, technical communication, pragmatic execution.