

PRIYANSH PATEL

London, UK · +44 7733877459 · priyansh2005p@gmail.com
linkedin.com/in/priyansh-patel-64ba742ba · github.com/PRIYANSH29-boop · portfolio-dun-two-93.vercel.app

PROFILE SUMMARY

First-Class Computer Science graduate and junior ML/AI engineer focused on machine learning for finance — fraud detection, quantitative signals, and the leakage-aware evaluation that makes a model trustworthy. Builds end to end: data pipelines, models, honest cost-aware backtests, and deployable apps. Biased toward honest metrics over leaderboard scores.

EDUCATION

University of East London — BSc (Hons) Computer Science | 2024–2026

- **First-Class Honours (85% average)**, awarded 2026. Coursework across software development, programming, data structures, algorithms, and web technologies.

TECHNICAL SKILLS

- **Languages:** Python, SQL, JavaScript / TypeScript
- **ML & Data:** scikit-learn, XGBoost, LightGBM (LambdaMART), SHAP, pandas, NumPy; walk-forward & purged/embargoed validation, learning-to-rank, imbalanced-data methods, feature engineering
- **LLM & AI:** Llama 3.3 (Groq), local LLMs (Ollama), retrieval-grounded explanation
- **Tools & Deployment:** Git / GitHub, GitHub Actions (CI), Docker, Streamlit, Vercel, Linux
- **Web:** React, Node.js, Express (full-stack background)

PROJECTS

RankAlpha — Cross-sectional equity ranking system (github.com/PRIYANSH29-boop/financeai)

- Built a LightGBM learning-to-rank (LambdaMART) model that ranks S&P 500 stocks into a market-neutral long/short book.
- Leakage control via walk-forward retraining + a 21-day embargo; beat a plain-momentum baseline after costs (Sharpe 1.14 vs 0.82).
- SHAP interpretability + a risk-managed, self-explaining Streamlit demo; 23-month forward paper-trading track.

FraudLens — Explainable credit-card fraud detection ([live demo](#) + [GitHub](#))

- XGBoost on a 0.17%-prevalence dataset; scored by the £ cost of errors instead of misleading accuracy.
- Diagnosed the data as the ceiling: an amount-ratio feature lifted PR-AUC 0.79→0.90 and high-value fraud recall from ~0% to 99.8%.

Anti-UAV Detection — Cross-domain evaluation (co-authored IEEE paper, University of East London)

- Built a YOLO11 → ResNet-18 → EfficientNet-B0 drone-detection cascade and honestly measured cross-domain generalisation.
- Showed near-perfect in-distribution scores partly collapse on unseen data, behind a perceptual-hash leakage audit; multi-seed and reproducible.

StockAI Terminal — ML + LLM stock-analysis prototype ([live demo](#))

- AdaBoost signals (walk-forward CV); Llama 3.3 (Groq) explanations grounded in the model's output — cites, never invents, gives no buy/sell advice.

EXPERIENCE

Intern — Red Bridge Library

- Introduced students to AI concepts and created educational content.

IT Volunteer — Compubits (Aug 2025)

- Web development (React, JavaScript); improved UI layout and responsiveness.