

# SURIYA PRASATH PARMESWARAN

[Linked in](#)

| +44 7887299164

| [suriyavictor6@gmail.com](mailto:suriyavictor6@gmail.com)

| [Portfolio](#)

| [Github](#)

---

## Professional summary

Data Science graduate building ML systems around real business constraints. If analysts can only review 500 transactions a day, which ones? If you have a £10k retention budget, who do you call? I've fine-tuned domain-specific transformers, built MILP optimisers, and implemented drift monitoring — three end-to-end projects, real datasets, real numbers.

---

## Education

University of Leeds, Leeds, UK. ( *MSc in Data Science and Analytics* )

(Sep 2024 – Sep 2025)

Graduate research in machine learning, NLP, and risk modelling; strong academic performance in Machine Learning, Statistical Learning, and Natural Language Processing

---

## Projects

### CyberSec Severity Intelligence Engine ( *MSc Dissertation* )

- Rebuilt v1 (LDA + Random Forest, 64% acc) into a dual-model NLP pipeline — TF-IDF bigrams + fine-tuned SecBERT on 29k real NVD CVEs
- Achieved 74% accuracy and 0.81 F1 on CRITICAL class; matched GPT-4o zero-shot accuracy at ~100× lower inference cost
- Added SHAP token-level explainability, stratified 3-fold CV, and SMOTE — replacing a single train/test split with reproducible evaluation
- Deployed end-to-end: Streamlit live app + MLflow experiment tracking + GitHub Actions CI/CD on Python 3.10/3.11

### Fraud Triage Engine

- Cut fraud review costs 70% (£37,500 → £11,390) across 284k transactions via cost-matrix threshold optimisation
- Achieved 99× lift at K=500 with LightGBM + XGBoost ensemble; chronological splits exposed overfitting hidden by random splits
- Added SHAP explainability and drift monitoring — detected 31% PR-AUC drop (0.849 → 0.582) with auto retraining alerts .

### Churn-to-Profit Decision Engine

- Delivered 8.15× ROI and £81,534 profit at £10k budget — outperformed churn-first targeting by 37.7%
  - Built T-learner uplift model + MILP optimiser (PuLP); exact optimisation beat greedy by £13,823
  - Validated with 2,000-run Monte Carlo — profitable across every simulation, £56,719 downside floor
- 

## Professional Experience

### Shree Kay Vee Automation – *IoT & Analytics Intern*

(Jan 2023 – Apr 2023)

- Worked with a team analyzing IoT sensor data from industrial equipment to explore patterns related to machine performance and potential failures.
  - Cleaned and preprocessed IoT sensor datasets using Python and SQL to support equipment performance analysis.
  - Contributed to building basic dashboards and visualisation to track operational metrics and equipment performance.
  - Supported the development of data ingestion and processing scripts used in internal analytics workflows.
- 

## Skills

**Machine Learning & NLP:** Python, SQL, PyTorch, HuggingFace Transformers, scikit-learn, LightGBM, XGBoost, SHAP, Optuna, SMOTE, Isotonic Calibration, Transfer Learning

**Time Series & Forecasting:** SARIMA, ARIMA, Ridge Regression, Temporal Validation, Concept Drift Monitoring

**Optimisation & Simulation:** MILP (PuLP), Monte Carlo Simulation, Hyperparameter Search (Optuna), Uplift Modelling

**Data Engineering:** pandas, NumPy, DuckDB, SQL, Feature Engineering, ETL Pipelines

**MLOps & Deployment:** MLflow, GitHub Actions CI/CD, Docker, Streamlit, FastAPI, pytest