

Raymond Dankwa

Selected Analytics and Strategy Project Summaries

This public evidence pack presents concise, sanitised project summaries. University names, module codes, student identifiers, personal contact details, assessment declarations, team-member identifiers and unpublished assessment files have been excluded.

Professional focus

Business analytics, predictive modelling, strategy simulation, HR and compliance reporting, digital innovation and evidence-led decision support.

Predicting AI Job Salaries with Machine Learning

A comparative modelling study using a global AI and data-jobs salary dataset containing approximately 90,000 records and 35 variables.

Approach

Data cleaning, median imputation, categorical encoding, correlation analysis, an 80/20 train-test split and comparison of Linear Regression, Random Forest and XGBoost. Model performance was evaluated using MAE, RMSE and R-squared.

Selected findings

XGBoost achieved the strongest reported performance (R^2 0.94), followed by Random Forest (R^2 0.91) and Linear Regression (R^2 0.77). Experience, bonus, country, satisfaction and job role emerged as influential predictors.

Skills demonstrated

Python/R-style predictive workflow, model comparison, diagnostics, feature importance, labour-market interpretation and evidence-based HR implications.

Publication note

The full assessed report is not publicly reproduced. This summary retains only portfolio-relevant evidence and excludes assessment metadata and institutional identifiers.

Retail Sales Forecasting and Profit Optimisation

A proof-of-concept retail analytics project using an open Global Superstore dataset as a proxy for a multinational furniture retailer.

Approach

Python and Jupyter-based exploratory analysis, feature engineering, seasonal forecasting, regression, classification, customer segmentation and market-basket analysis. Techniques included SARIMA, ARIMA, ETS, Decision Trees, Logistic Regression and K-Means.

Selected findings

SARIMA captured seasonal demand more effectively than non-seasonal ARIMA, with a reported MAPE of 18.97%. The analysis also highlighted discount-driven margin erosion, distinct customer clusters and the value of transparent classification rules for profit risk.

Skills demonstrated

Time-series forecasting, classification, clustering, profitability analysis, ethical data practice, teamwork and business recommendation design.

Publication note

This was collaborative work. The public summary does not identify other contributors or reproduce the original team submission.

Voltex Motors Business Simulation

A four-round competitive simulation evaluating the financial, operational and strategic development of a simulated automotive manufacturer.

Approach

Round-by-round analysis of production, pricing, profitability, finance, workforce decisions, market position and SMART objectives. Strategic interpretation was linked to operating outcomes and team decision-making.

Selected findings

The simulated company moved from a £63.36m post-tax loss in Round 1 to £549.88m post-tax profit in Round 4, achieved 100% sell-through in every round and increased shareholders' funds from £446.64m to £998.79m. Raymond's reflection focused on translating financial data into strategic narrative.

Skills demonstrated

Business simulation, strategic analysis, financial interpretation, SMART objectives, operations, teamwork reflection and recommendations.

Publication note

Only a sanitised summary is published; the original assessed document remains private.

Telecommunications Customer Churn Modelling

A customer-churn modelling notebook designed to prepare telecom customer data and compare classification approaches.

Approach

Data cleaning, conversion of total charges, categorical encoding, stratified train-test splitting, feature scaling and comparison of Logistic Regression, Decision Tree and Random Forest models. The workflow includes precision, recall, F1, confusion matrix, ROC/AUC, cross-validation and feature importance.

Selected findings

The uploaded notebook establishes a reproducible end-to-end modelling structure. Because the supplied notebook does not include preserved final outputs, this portfolio does not claim specific performance scores.

Skills demonstrated

Python, pandas, scikit-learn, model evaluation, class balancing, cross-validation, feature importance and responsible reporting.

Publication note

The code notebook is available as technical evidence. The accompanying customer dataset is not redistributed through this portfolio.