

Highlights

Proposes a hybrid demand-grounded dynamic pricing framework (CatBoost–PPO) that integrates gradient-boosted demand forecasting, adaptive reinforcement learning, vectorized price optimization, and price elasticity analysis within a single deployable application for non-technical e-commerce sellers.

- **Hybrid CatBoost–PPO ensemble pricing framework** eliminates the demand-grounding failure of standalone reinforcement learning, delivering 27.1% mean profit improvement ($p = 0.026$) across 80 products with zero products experiencing greater than 10% profit loss.
- **Vectorized O(1) batch inference strategy** reduces price optimization from $O(n)$ sequential model calls to a single batch inference call, enabling complete price recommendations in 115.92 ms and confirming real-time viability for interactive seller sessions.
- **Embedded price elasticity module** classifies product markets as elastic or inelastic using $\pm 10\%$ price perturbation, identifying Laptops & Computers as the only elastic category ($\epsilon = -1.325$) and supporting category-level pricing strategy guidance.
- **CatBoost demand model achieves $R^2 = 0.9877$** (MAE = 58.16 units) on a held-out test set of 2,304 records, substantially outperforming Linear Regression, Random Forest, and XGBoost across all evaluation metrics.
- **Cross-dataset generalisation confirmed on the UCI Online Retail II dataset** comprising 525,461 real UK retail transactions, achieving +11.8% mean profit improvement across 400 records in a fundamentally different retail context without retraining.