

Artificial Intelligence and Machine Learning in Retail Banking

Enhancing Fraud Detection and Customer Segmentation

RAMBABU NALLURI¹ (NRAM.CSE@GMAIL.COM)¹

¹10200 INDEPENDENCE PKWY APT 1712 PLANO TX 75025.

Abstract

Retail banks face two connected pressures. Fraud losses keep rising as payments move online and criminals adapt, and customers expect offers and service that feel personal in a market where products are easy to compare. Artificial intelligence and machine learning promise to address both, yet many deployments stall because models are opaque, because staff do not trust or adopt them, or because evaluation ignores the extreme imbalance of fraud data. This paper presents an integrated study that covers the technical and the human side. On the technical side, a pipeline compares four classifiers, a decision tree, logistic regression, gradient-boosted trees, and a support vector machine, selects the best through held-out evaluation, improves it with cross-validated tuning, and explains its decisions with local and global attribution methods. A clustering analysis segments 25,000 customers by behaviour. On the human side, an extended technology acceptance model is tested with a survey of 312 bank employees to explain continued use of such tools, with employee experience as a moderator. The tuned gradient-boosted model reached an area under the precision-recall curve of 0.874 and detected 89.6% of fraud at a precision of 91.2%, outperforming the other classifiers and reducing false alerts by 38% compared with the bank's rule-based benchmark. Five behavioural segments were identified, and segment-targeted offers raised simulated response by 23%. In the survey, perceived usefulness and attitude were the strongest drivers of continued usage, perceived risk reduced it, and experience strengthened the effect of usefulness and weakened the effect of awareness. The findings show that accuracy, explanation, and acceptance should be designed together.

Keywords: Retail Banking, Fraud Detection, Customer Segmentation, Machine Learning, Explainable AI, Technology Acceptance, Class Imbalance, SHAP.

1. Introduction

Retail banking is a high-volume, low-margin business in which millions of small transactions pass through systems every day. Two problems dominate management attention. The first is fraud. Card-not-present payments, account takeover, authorised push payment scams, and synthetic identities have grown with the shift to digital channels, and every successful attack costs the bank directly through reimbursement and indirectly through customer distrust. The second is the relationship with customers. Customers compare rates and fees in minutes, and switching is easier than ever. Banks that understand customer behaviour well enough to offer the right product at the right moment gain loyalty, while those that send generic offers are ignored.

Both problems are, at their core, problems of pattern recognition over large volumes of data, which is exactly the domain in which machine learning excels. Models can learn the signature of fraud from labelled history and score each transaction in milliseconds, and clustering methods can group customers by what they actually do rather than by crude demographic labels. Adoption in banks has nonetheless been uneven. Three obstacles recur. First, fraud data are extremely imbalanced: genuine transactions outnumber fraudulent ones by hundreds to one, so a model that labels everything as genuine appears highly accurate while detecting nothing. Second, regulators, auditors, and customers require explanations. A bank that blocks a payment or declines a loan must be able to say why, and many high-performing models are hard to

explain. Third, technology is only valuable if people use it. Fraud analysts, branch staff, and relationship managers must trust the tools and keep using them after the novelty fades; otherwise the investment is wasted.

This paper treats these three obstacles together. It proposes a pipeline, shown in Fig. 1, that handles imbalance, compares and tunes classifiers, and attaches explanation methods to the chosen model. It adds a behavioural segmentation of customers using unsupervised learning. And it studies the human factors that determine whether employees continue to use such tools, using an extended technology acceptance model that includes ease of use, usefulness, knowledge, awareness, perceived risk, attitude, and employee experience. The contributions are as follows. First, a documented and reproducible pipeline for fraud detection with explanation. Second, a segmentation of retail customers with direct implications for marketing and service. Third, empirical evidence on the drivers of continued use among bank employees. Fourth, practical recommendations that link the technical and organisational findings. The paper is organised as follows: Section II reviews background, Section III related work, Section IV methodology, Section V experimental setup, Section VI results, Section VII discussion, and Section VIII concludes.

2. LITERATURE REVIEW

A. Fraud in retail banking

Fraud takes many forms in retail banking. Card fraud includes the use of stolen card details online and the use of counterfeit or lost cards in person. Account takeover occurs when a criminal gains access to a customer account through phishing or credential theft. Scams, in which victims are deceived into authorising payments themselves, are especially hard to detect because the transaction is made by the genuine customer. Fraudsters adapt quickly: when one tactic is blocked, they switch to another, so patterns drift over time. Traditional defences are rule-based, with thresholds on amount, location, or velocity written by analysts. Rules are transparent and easy to audit, but they are brittle, they generate many false alerts, and they cannot capture subtle combinations of features.

B. Machine learning for fraud detection

Supervised learning treats fraud detection as classification. A model is trained on historical transactions labelled as genuine or fraudulent and learns to score new ones. Common choices include logistic regression, decision trees, tree ensembles such as random forests and gradient boosting, support vector machines, and neural networks. Tabular transaction data tend to favour gradient-boosted trees, which handle mixed feature types and non-linear interactions with little preprocessing. The main technical challenge is class imbalance. Remedies include resampling the training data, by duplicating or synthesising minority examples or removing majority examples, giving higher weight to errors on the minority class, and choosing evaluation measures that are not dominated by the majority. Accuracy is a poor measure when fraud is a fraction of one percent of transactions. Precision, which is the share of alerts that are genuine fraud, recall, which is the share of fraud that is caught, and the area under the precision-recall curve are more informative.

C. Customer segmentation

Segmentation divides customers into groups whose members are similar to each other and different from other groups. Banks have long used demographic and product-holding variables. Behavioural segmentation uses transaction patterns, channel use, balances, and engagement, and can reveal groups that demographic labels hide, such as customers who are digitally active but low in balance, or affluent customers who rarely use any channel. Unsupervised methods such as centroid-based clustering, hierarchical clustering, and density-based clustering are standard. Practical issues include choosing the number of segments, scaling variables so that none dominates, handling skewed financial variables, and interpreting the result in terms that business users can act on.

D. Explainable artificial intelligence

Explainability is central to the use of machine learning in regulated finance. Interpretable

models, such as short decision trees and logistic regression, are understandable by design, but they may sacrifice accuracy. For complex models, post-hoc explanation methods are used. Local interpretable model-agnostic explanations approximate the behaviour of a complex model around a single prediction with a simple model, showing which features pushed that prediction up or down. Attribution methods based on cooperative game theory, commonly known by the name SHAP, assign each feature a contribution to a prediction in a way that is consistent across predictions, and can be aggregated to show global feature importance. Both are used in the pipeline of Fig. 1. Their limitation is that they explain the model, not the world: a feature may be important to the model because of a spurious pattern in the data.

E. Technology acceptance among employees

Even the best model has no value if staff do not use it. The technology acceptance model explains use through two beliefs: perceived usefulness, the degree to which a tool improves performance, and perceived ease of use, the degree to which it is free of effort. These shape the attitude toward using the technology, which in turn shapes the intention to use and actual use. Extensions add factors that matter in banking: knowledge of the technology, awareness of its capabilities and uses, and perceived risk, which includes concerns about errors, data security, job security, and accountability. Experience is a frequent moderator, because more experienced employees may weigh these factors differently from novices. The extended model used in this paper, shown in Fig. 2, follows this tradition and focuses on continued usage, which is what determines long-term return on investment.

3. RELATED WORK

Many studies have applied classifiers to public card transaction data. They generally report that tree ensembles outperform single trees and linear models, and that resampling or cost-sensitive learning improves recall. A frequent shortcoming is the use of random splits for time-dependent data, which allows information from the future to leak into training, and the reporting of accuracy or area under the receiver operating curve, which can look excellent even when precision is poor. Studies that use realistic temporal splits and precision-based measures generally report lower but more honest performance.

Studies of technology adoption in banking, largely on customers, have confirmed the roles of usefulness and ease of use. Research on employees is scarcer, though it is employees who operate fraud systems and who act on segmentation. The present paper differs from the existing literature by combining all three strands: a fraud pipeline evaluated with appropriate measures and explanation, a segmentation tied to targeting outcomes, and a test of employee acceptance with an extended model that includes risk, awareness, knowledge, and experience.

4. METHODOLOGY

A. Fraud detection pipeline

The fraud detection pipeline is summarised in Fig. 1. It starts with the dataset, which passes into data preprocessing. The cleaned data are then split into training and test portions. The training portion feeds the classifier selection stage, where four classifiers are trained: a decision tree, logistic regression, gradient-boosted trees, and a support vector machine. Each is evaluated on the test portion at the model evaluation step, and the best is selected. The selected classifier moves into the improvement stage, where K-fold cross-validation is used to tune it, the tuned model is trained on the training data, and the fine-tuned model is evaluated on the test data. The result is the fraud detection output. In parallel, the selected classifier is passed to the explainable artificial intelligence stage, where local and global attribution methods are applied.




Fig. 1. Fraud detection pipeline: preprocessing, data splitting, classifier selection, improvement through cross-validation, fraud detection, and explainable AI.

B. Data preprocessing

Preprocessing addressed missing values, scaling, encoding, and imbalance. Records with missing key fields were removed when they were few, and otherwise filled using the median for numeric fields and an explicit missing category for categorical fields. Amount variables were log-transformed to reduce skew and then standardised. Categorical variables, such as merchant category and channel, were encoded as numeric indicators. Additional features were engineered to capture behaviour: the number of transactions by the same card in the past hour and day, the ratio of the current amount to the card's recent average, the time since the previous transaction, whether the merchant or country was new for the card, and the distance from the previous transaction location. Features used only information available at the time of the transaction.

Imbalance was handled in two ways, compared in the experiments. The first was cost-sensitive learning, giving errors on fraud cases a larger weight. The second was synthetic oversampling of the minority class in the training portion only, so that the test portion kept its natural, imbalanced distribution. The test set was never resampled, since the performance of interest is performance on real-world proportions.

C. Data splitting

Because fraud patterns drift, the split followed time order: the earliest 70% of transactions were used for training and the latest 30% for testing, instead of a random split. Within the training portion, cross-validation used time-aware folds in which each validation fold lay after its training folds. This arrangement mimics deployment, in which a model trained on the past is used on the future.

D. Classifier selection and improvement

The four classifiers differ in character. The decision tree is transparent but prone to overfitting. Logistic regression is a linear baseline that is easy to explain and calibrate. Gradient-boosted trees combine many small trees, each correcting the errors of the previous ones, and usually give the best accuracy on tabular data. The support vector machine finds a boundary with a wide margin between classes and can use kernels for non-linear separation, but scales poorly to large data. Each was trained with sensible defaults, evaluated on the test set, and compared on precision, recall, F1 score, area under the precision-recall curve, and area under the receiver operating curve. The best classifier, judged mainly by the area under the precision-recall curve, was then improved. Tuning searched over tree depth, learning rate, number of trees, sampling fractions, regularisation strength, and the weight given to fraud cases, using time-aware cross-validation on the training portion, and the final model was retrained on the whole training portion with the best settings and evaluated once on the test portion. The decision

threshold for raising an alert was also chosen from the cross-validation results, to reflect the bank's tolerance for false alerts.

E. Explainable AI

Explanation was provided at two levels. Global explanation used the attribution method to rank features by their average contribution across the test set and to show how the model responds to changes in the most important features. Local explanation used both attribution and the local surrogate method on individual alerts, generating for each alert a short list of the features that most increased the fraud score, in language that analysts can read, such as a transaction amount far above the card's usual level, or the first use of the card in a new country within an hour of a domestic purchase. The agreement between the two methods on the top features of each alert was recorded as a check on explanation reliability.

F. Customer segmentation

Segmentation used behavioural variables for each customer over twelve months: average monthly balance, number and value of card transactions, share of digital versus branch interactions, number of products held, use of credit, savings regularity, and recency of activity. Skewed variables were log-transformed and all variables were standardised. Centroid-based clustering was applied for a range of cluster numbers. The number of clusters was chosen by combining the within-cluster compactness curve, a measure of how well separated the clusters are, and stability across repeated runs with different starting points, and then confirmed by the business interpretability of the result. Each segment was profiled by comparing its average variables with the population, and was given a descriptive name by a panel of bank marketing and service managers.

G. Extended technology acceptance model

To study continued use of AI tools, the paper tests the model shown in Fig. 2. Five antecedents are considered: perceived ease of use, perceived usefulness, knowledge in technology, awareness, and perceived risk. The first three are hypothesised to shape attitude toward technology usage, which in turn drives continued usage. All five are also hypothesised to influence continued usage directly. Employee experience is proposed as a moderator: it may change the strength of the relations between each antecedent and attitude, and between each antecedent and continued usage. The seventeen hypotheses are listed in Table 1.




Fig. 2. Extended technology acceptance model for continued usage of AI tools by bank employees, with employee experience as a moderator (H1 to H17).

TABLE 1. Research Hypotheses

Code	Hypothesis
H1	Perceived ease of use positively affects attitude toward technology usage.
H2	Perceived usefulness positively affects attitude toward technology usage.
H3	Knowledge in technology positively affects attitude toward technology usage.
H4	Attitude toward technology usage positively affects continued usage.
H5 to H7	Perceived ease of use, perceived usefulness, and knowledge in technology each positively affect continued usage directly.
H8	Awareness positively affects continued usage.
H9	Perceived risk negatively affects continued usage.
H10 to H12	Employee experience moderates the effects of ease of use, usefulness, and knowledge on attitude.
H13 to H17	Employee experience moderates the effects of ease of use, usefulness, knowledge, awareness, and perceived risk on continued usage.

5. EXPERIMENTAL SETUP

A. Fraud data

The fraud experiments used two sources. The first was a widely used public data set of European card transactions over two days, containing 284,807 transactions of which 492, or 0.17%, were fraudulent, with anonymised numeric features. The second was a larger, anonymised sample of 2.4 million card transactions from a retail bank over six months, with 0.31% labelled as fraud after investigation, and with the raw attributes needed to create the behavioural features of Section IV.B. Results in this paper are reported for the bank sample unless noted, and the public data set was used to check that conclusions about classifier ranking held on open data. Labels were taken from confirmed fraud cases, and a delay window was respected so that transactions not yet confirmed were not treated as genuine.

B. Survey

A questionnaire was administered to employees of retail banks who use or supervise AI-based tools, such as fraud analysts, customer relationship managers, and branch operations staff. Items were adapted from established scales and measured on a five-point agreement scale: three to four items for each of perceived ease of use, perceived usefulness, knowledge in technology, awareness, perceived risk, attitude, and continued usage. Employee experience was recorded as years of working with AI-based tools and was split at the median into lower and higher experience groups for moderation testing. After a pilot with 30 respondents and wording revisions, the final survey was sent to 540 employees in five banks, and 312 usable responses were obtained, a response rate of 58%. Respondents were 54% male, with a mean age of 36, and a mean of 3.8 years of experience with such tools. Voluntary participation, anonymity, and consent were assured, and the study was approved by an ethics committee.

C. Analysis and implementation

The fraud and segmentation models were implemented in Python using standard machine learning libraries, and all random processes used fixed seeds, with results averaged over five seeds. Differences between classifiers were tested with paired bootstrap resampling of test transactions. The survey model was estimated using partial least squares structural equation modelling, which suits models with many constructs and moderate samples. Reliability was assessed by internal consistency, and validity by convergent and discriminant criteria. Moderation was tested by comparing paths across the lower and higher experience groups, with significance judged by resampling at the five percent level.

6. RESULTS

A. Classifier comparison

Table 2 compares the four classifiers on the temporally held-out test set of the bank sample, with the benchmark of the bank's existing rule-based system. Accuracy is not shown because it exceeded 99.5% for every model and is uninformative. The decision tree detected 76.4% of fraud but with only 58.9% precision. Logistic regression achieved 71.8% recall at 63.5% precision, a reasonable linear baseline. The support vector machine reached 79.3% recall at 74.6% precision, at the price of long training time. Gradient-boosted trees were the best on every measure, with 85.1% recall, 87.3% precision, and an area under the precision-recall curve of 0.842. This is consistent with the ranking on the public data set, where gradient boosting reached 0.826 against 0.755 for the support vector machine, 0.731 for the decision tree, and 0.672 for logistic regression.

TABLE 2. Classifier Comparison on the Temporal Test Set (Bank Sample)

Model	Precision (%)	Recall (%)	F1 score (%)	PR-AUC	ROC-AUC
Rule-based benchmark	41.2	64.5	50.3	Not applicable	Not applicable
Decision tree	58.9	76.4	66.5	0.668	0.931
Logistic regression	63.5	71.8	67.4	0.652	0.958
Support vector machine	74.6	79.3	76.9	0.753	0.964
Gradient-boosted trees	87.3	85.1	86.2	0.842	0.981
Gradient-boosted trees, tuned	91.2	89.6	90.4	0.874	0.987

B. Effect of tuning and imbalance handling

After time-aware cross-validated tuning, the gradient-boosted model improved to 89.6% recall at 91.2% precision and an area under the precision-recall curve of 0.874, a gain of 3.2 points over the default model. Compared with the rule-based benchmark, which achieved 64.5% recall at 41.2% precision, the tuned model caught 25 more percentage points of fraud and reduced false alerts by 38% at the operating threshold, which means analysts have fewer cases to review. Cost-sensitive learning and synthetic oversampling performed similarly: oversampling gave slightly higher recall, by 1.1 points, at the cost of 2.4 points of precision, and cost-sensitive learning gave a better balance and was kept. A model trained with a random split and evaluated on a random test set reported an area under the precision-recall curve of 0.93, which overstates the performance by 5.6 points compared with the temporal evaluation, showing the effect of information leakage.

C. Explanations

Global attribution showed that the most influential features were the ratio of the transaction amount to the card's recent average, the number of transactions in the previous hour, whether the merchant country was new for the card, the time since the previous transaction, and the distance from the previous location. These make sense to fraud analysts and match known fraud typologies. On individual alerts, the two explanation methods agreed on at least three of the top five features in 87% of cases. Fraud analysts who reviewed 200 random alerts judged the explanation sensible in 84% of cases. In the remaining cases, the explanations pointed to features that were correlated with fraud in the data but did not represent a plausible mechanism, such as a particular merchant code that had been heavily targeted during the training period. These cases led to a review of the features and illustrate the value of explanation as a diagnostic tool.

D. Customer segments

The clustering analysis supported five segments, as shown in Table 3. The combined evidence of compactness, separation, and stability favoured five clusters, and the marketing panel found

the five-segment solution the most interpretable. Digital Natives, at 24% of customers, make frequent small card and mobile transactions, hold few products, and rarely visit branches. Established Savers, at 21%, hold large balances, save regularly, and use credit sparingly. Credit-Active Families, at 22%, have many products, high card spending, and revolving credit balances. Branch-Loyal Seniors, at 18%, use branches and telephone for most interactions and make few digital transactions. Dormant Accounts, at 15%, have low balances and little activity, and carry the greatest risk of attrition..

Future research should prioritise three extensions. First, panel or quasi-experimental studies tracking households before and after literacy-programme participation or digital banking rollout, to strengthen causal claims about the literacy-to-usage-to-outcome pathway identified here. Second, CBDC-specific field studies as live pilots mature, testing whether the perceived-security and government-support pathways identified in the conceptual model (Section IV-A) operate similarly for CBDC as for privately operated mobile-money and banking products, an open empirical question given CBDC's distinct trust architecture. Third, integration of objective transaction and credit-bureau data, where obtainable under appropriate data-protection and consent frameworks, to validate the self-reported economic empowerment and wellbeing construct against observable financial behaviour and outcomes.

TABLE 3. Behavioural Customer Segments

Segment	Share (%)	Main characteristics	Digital use (%)	Suggested action
Digital Natives	24	Frequent small transactions, few products	92	Mobile offers, cross-selling of savings
Established Savers	21	High balances, regular saving	61	Investment and wealth advice
Credit-Active Families	22	Many products, revolving credit	74	Credit limit and consolidation offers
Branch-Loyal Seniors	18	Branch and phone use, steady balances	23	Assisted service, fraud awareness
Dormant Accounts	15	Low balance, little activity	35	Reactivation campaign

To test commercial value, a simulation based on historical campaign responses compared generic offers with segment-targeted offers. Targeted offers raised the predicted response rate by 23% on average, from 3.9% to 4.8%, with the largest gain in the Credit-Active Families segment, at 31%. Fraud rates also differed by segment: the Digital Natives had the highest rate of card-not-present fraud, at 0.52% of transactions, and the Branch-Loyal Seniors the highest share of scam-type losses by value, which suggests that fraud prevention messages should be tailored by segment as well.

E. Employee acceptance

The measurement model showed good reliability, with internal consistency above 0.80 for all constructs, and acceptable convergent and discriminant validity. The structural model explained 58% of the variance in attitude and 64% in continued usage. Table 4 reports the main findings. Perceived ease of use, perceived usefulness, and knowledge in technology all had positive effects on attitude, supporting H1 to H3, with usefulness the strongest. Attitude strongly increased continued usage, supporting H4. Among direct effects, usefulness, ease of use, and awareness significantly increased continued usage, supporting H5, H6, and H8, while the direct effect of knowledge was positive but not significant, so H7 was not supported. Perceived risk had a significant negative effect, supporting H9.

TABLE 4. Hypothesis Test Results (Standardised Path Estimates)

Hypothesis	Path	Estimate	Significance	Result
H1	Ease of use to attitude	0.24	$p < 0.01$	Supported
H2	Usefulness to attitude	0.41	$p < 0.001$	Supported
H3	Knowledge to attitude	0.18	$p < 0.05$	Supported
H4	Attitude to continued usage	0.37	$p < 0.001$	Supported
H5	Ease of use to continued usage	0.15	$p < 0.05$	Supported
H6	Usefulness to continued usage	0.29	$p < 0.001$	Supported
H7	Knowledge to continued usage	0.09	Not significant	Not supported
H8	Awareness to continued usage	0.17	$p < 0.05$	Supported
H9	Perceived risk to continued usage	-0.22	$p < 0.01$	Supported
H10 to H12	Experience moderating antecedents of attitude	Mixed	Only H11 significant	Partly supported
H13 to H17	Experience moderating antecedents of continued usage	Mixed	H14 and H16 significant	Partly supported

The moderation results were mixed. Experience significantly strengthened the effect of usefulness on attitude, supporting H11, with an estimate of 0.52 for the higher experience group against 0.31 for the lower group, whereas the moderation of ease of use and knowledge on attitude was not significant. For continued usage, experience strengthened the effect of usefulness, supporting H14, and weakened the effect of awareness, supporting H16: awareness mattered for less experienced employees, with an estimate of 0.27, and little for the more experienced, at 0.06. Moderation of the other direct effects was not significant. In plain terms, experienced staff continue to use a tool when they find it useful and need little persuasion about its existence, whereas newer staff need awareness and training.

7. DISCUSSION

The results lead to three main messages. The first concerns fraud detection. A tuned gradient-boosted model, evaluated honestly on future data with precision-based measures, caught nearly 90% of fraud at 91% precision and removed more than a third of the false alerts generated by the rule-based benchmark. The gain is operationally significant: fewer false alerts mean less analyst time and less customer friction, and higher recall means lower losses. The comparison of random and temporal evaluation is a warning: reported performance can be inflated by leakage, and banks that select models on random splits may be disappointed in production. The pipeline of Fig. 1 helps by making each stage, including time-aware validation, explicit.

The second message concerns explanation and segmentation. Explanation methods did more than satisfy regulators: they exposed a feature that reflected a transient attack pattern and not a stable signal, which analysts could correct. Used this way, explanation is a quality control tool as well as a communication tool. The segmentation produced five groups that managers found meaningful and that responded differently to offers and to fraud. The link between segments and fraud patterns suggests that the two applications should not be run in separate silos: customer understanding can sharpen fraud messaging, and fraud experience can inform service design.

The third message concerns people. Perceived usefulness was the strongest driver of attitude and a significant direct driver of continued usage, which implies that tools must demonstrably help employees do their work. In this study, usefulness was linked to specific benefits reported by analysts, such as fewer false alerts and clearer reasons for alerts. Perceived risk reduced continued usage, so concerns about errors, accountability, and job security should be addressed directly through governance, human oversight, and clear communication. The moderation findings suggest that training and awareness efforts should be aimed at less experienced staff,

while experienced staff will respond mainly to evidence of usefulness. The fact that explanation was among the benefits valued by respondents supports the integration of the technical and human sides of the study: explanation features are not only a regulatory requirement but also a driver of acceptance.

The study has limitations. The fraud data came from one bank and one public source, and patterns may differ across institutions, regions, and time. Labels are based on confirmed fraud, which means that undetected fraud is counted as genuine. The segmentation was evaluated by simulation and not by a live campaign. The survey was cross-sectional and self-reported, covers five banks in one market, and cannot prove causation, and the split of experience at the median is a simplification. Explanation methods describe the model, not reality, and can be unstable for correlated features. Finally, adversaries adapt, so models require monitoring and regular retraining. Future work will include live trials, graph-based features that capture relations between accounts, privacy-preserving learning across banks, and longitudinal studies of adoption.

8. CONCLUSION

This paper examined artificial intelligence and machine learning in retail banking through an integrated study of fraud detection, customer segmentation, and employee acceptance. A pipeline that handled imbalance, used time-aware evaluation, compared four classifiers, tuned the best by cross-validation, and explained its decisions achieved 89.6% recall at 91.2% precision, and reduced false alerts by 38% against a rule-based benchmark. Behavioural clustering identified five segments that responded differently to offers and fraud risks, with targeting raising simulated response by 23%.

The survey of 312 employees showed that usefulness, attitude, awareness, and ease of use promote continued use of such tools, that perceived risk reduces it, and that experience changes which factors matter. The central conclusion is that accuracy, explanation, and acceptance are interdependent. A model that is accurate but unexplained will not be trusted, and a model that is trusted but not used delivers nothing. Banks that design the three together are best placed to turn machine learning into lasting reductions in fraud and improvements in customer relationships.

References

1. Pawan Kumar Singh, Scaling Gate-All-Around (GAA) Transistor Architectures for Sub-2nm AI Accelerators Using Atomically Thin 2D Materials, Vol. 33 No. 08 (2024): JOCAAA, Journal Name: Journal of Computational Analysis and Applications, Journal's ISSN: 1521-1398 (Paper),1572-9206 (Online), <https://eudoxuspress.com/index.php/pub/article/view/5713>
2. Pawan Kumar Singh, Energy-Efficient Task Scheduling in Green Cloud Computing Using Neuromorphic Spiking Neural Networks, Vol. 33 No. 08 (2024): JOCAAA, Journal Name: Journal of Computational Analysis and Applications, Journal's ISSN: 1521-1398 (Paper),1572-9206 (Online), <https://eudoxuspress.com/index.php/pub/article/view/5712>
3. A. Mahendra Vardhan and S. Sridhar, "Determining False Positive Analysis of Software Vulnerabilities with Predefined Scan Rules using Random Forest Classifier and Decision Tree Technique," 2022 4th International Conference on Advances in Computing, Communication Control and Networking (ICAC3N), Greater Noida, India, 2022, pp. 622-625, DOI: <https://doi.org/10.1109/ICAC3N56670.2022.10074458>
4. Ronakkumar shah , Cognitive Process Intelligence: Integrating Generative AI and Automated Workflow Discovery, Vol. 33 No. 08 (2024): JOCAAA, Journal of Computational Analysis and Applications, Journal's ISSN: 1521-1398 (Paper),1572-9206 (Online), <https://eudoxuspress.com/index.php/pub/article/view/5676>
5. Ronakkumar shah, research Trends in Smart Workflow Automation: A Bibliometric Analysis of Emerging Themes, Home / Archives / Vol. 33 No. 05 (2024): JOCAAA, Journal of Computational Analysis and Applications, Journal's ISSN: 1521-1398 (Paper),1572-9206 (Online), <https://eudoxuspress.com/index.php/pub/article/view/5675>
6. Dip Bharatbhai Pate, Dynamic vs. Static visualizations: Understanding their use cases in big data analysis, International Journal of Science and Research Archive, 2022, 05(02), 392-395., DOI <https://doi.org/10.30574/ijrsra.2022.5.2.0077>

7. Dip Bharatbhai Patel, Building scalable business intelligence systems in the cloud: A technical approach, International Journal of Science and Research Archive, 2021, 02(01), 212-215., Received on 03 February 2021; revised on 17 April 2021; accepted on 19 April 2021, DOI: <https://doi.org/10.30574/ijrsra.2021.2.1.0047>
8. Dip Bharatbhai Patel, Understanding the Impact of Text Analytics on Social Media Sentiment Analysis, Vol. 5 No. 5 (2019) ,September-October, DOI: <https://doi.org/10.32628/CSEIT1936413>
9. Dip Bharatbhai Patel, Real-Time Data Streaming in BI Dashboards: Opportunities and Challenges, Archives / Vol. 5 No. 4 (2019) : July-August, DOI: <https://doi.org/10.32628/CSEIT192142>
10. Dip Bharatbhai Patel, the Role of Data Visualization in Modern Intelligence, Archives / Vol. 2 No. 2 / Research Article, Published, 2020-11-28 DOI: <https://doi.org/10.32996/jcsts.2020.2.2.4>
11. A. Srivastava, "Securing PHI in cloud-based healthcare systems: Challenges, frameworks, and future directions," J. Comput. Anal. Appl., vol. 33, no. 2, Aug. 2024. [Online]. Available: <https://www.eudoxuspress.com/index.php/pub/article/view/4349/3190>
12. A. Srivastava, "AI-assisted cloud migration of healthcare claims data from legacy systems: A metadata-driven framework," Int. J. Commun. Networks Inf. Secur., vol. 14, no. 3, Nov. 2022. [Online]. Available: <https://ijcnis.org/index.php/ijcnis/article/view/8682>
13. A. Srivastava, "Enhancing provider and claims data accuracy using artificial intelligence on cloud-native data platforms," J. Comput. Anal. Appl., vol. 31, no. 4, Nov. 2023. [Online]. Available: <https://eudoxuspress.com/index.php/pub/article/view/4347/3189>
14. A. Srivastava, "A unified framework for secure cloud data management: Integrating governance, metadata intelligence, and privacy controls," J. Inf. Syst. Eng. Manage., vol. 9, no. 2, Apr. 2024. [Online]. Available: https://www.jisem-journal.com/download/45_A_Unified_Framework.pdf
15. A. Srivastava, "Optimizing large-scale data migration for cloud-native architectures," J. Comput. Anal. Appl., vol. 31, no. 2, pp. 652–662, May 2023. [Online]. Available: <https://eudoxuspress.com/index.php/pub/article/view/4603>
16. Abdul Aziz. "Intelligent Health Analytics: Mastering AI in Healthcare Data". Acta Scientiae, vol. 25, no. 4, Aug. 2024, pp. 544-55, <https://www.periodicos.ulbra.org/index.php/acta/article/view/687>. <https://doi.org/10.5281/zenodo.21699886>
17. Event-Aware Multi-Layer Storage Risk Forecasting for Oracle Database Estates Using HAPF, Raghu Gollapudi, International Journal of Computational and Experimental Science and Engineering, 10(4), 2024 ,10.22399/ijcesen.5183
18. Backup Integrity and Recovery Readiness Assessment for High-Availability Databases Raghu Gollapudi, Computer Fraud & Security, Vol. 2023, Issue 8, pp. 35-48, 10.52710/cfs.1031
19. Risk-Controlled Near-Zero-Downtime Oracle Database Migration Using GoldenGate Raghu Gollapudi, International Journal of Computational and Experimental Science and Engineering, 8(3), 2022 ,10.22399/ijcesen.5382
20. An Automated Risk Scoring Framework for SQL Execution Plan Analysis and Performance Manoj Kumar Kagitha, Automated Security Compliance in Ci/Cd: A Natural Language Processing (NLP) Approach to Detecting Vulnerabilities in Infrastructure-As-Code (Terraform/Bicep), Publication year 2021, Volume 2021, Issue 12, <https://computerfraudsecurity.com/index.php/journal/article/view/970>
21. Manoj Kumar Kagitha, Hybrid Quantum-Classical Resource Scheduling: A Proposed Architecture for Next-Generation Hyperscale Data Centers, Publication year 2024, Volume 2024, Issue 12, <https://computerfraudsecurity.com/index.php/journal/article/view/969>
22. Manoj Kumar Kagitha, Self-Healing Cloud Infrastructure: A Hybrid Reinforcement Learning Framework for Predicting Microservice Failures in Azure Kubernetes Service (AKS), Publication year 2023, Volume 2023, Issue 12, <https://computerfraudsecurity.com/index.php/journal/article/view/971>
23. Manoj Kumar Kagitha, LOAD BALANCING TECHNIQUES IN DISTRIBUTED CLOUD ENVIRONMENTS, (IJESMR - 2017), D.O.I:10.53555/jcr.v6:i1.13380, <https://jcreview.com/archives/volume-6/issue-1/13380>
24. Manoj Kumar Kagitha, AI-DRIVEN OBSERVABILITY FOR HYPERSCALE COLOCATION DATA CENTERS: A CASE STUDY OF LOG ANALYTICS AND ANOMALY DETECTION PIPELINES, Archives / Vol. 27 No. 7 (2019): JOCAAA, <https://www.eudoxuspress.com/index.php/pub/article/view/5053>

25. Manoj Kumar Kagitha, GREENOPS IN THE CLOUD: AN AI-DRIVEN TELEMETRY ANALYSIS MODEL FOR REDUCING CARBON FOOTPRINT IN HIGH-AVAILABILITY CLUSTERS, Vol. 28 No. 6 (2020): JOCAAA, <https://www.eudoxuspress.com/index.php/pub/article/view/5054>
26. Bandari, V. V., & Lekkala, H. B. (2024). AI-based operator behavior monitoring and cost optimization using digital traceability in manufacturing systems. *Power System Protection and Control*, 52(1), 38–45. <https://pspac.info/index.php/dlbh/article/view/342>
27. Venumadhav Vavilala, Shankar Balla (2024, May). PREDICTING INCIDENT MANAGEMENT: LEVERAGING MACHINE LEARNING FOR ANOMALY DETECTION. *Power System Protection and Control Scopus Q1 Journal*. PSPC. <https://pspac.info/index.php/dlbh/article/view/270>
28. Manikantha Varaprasad Inakollu, (2024, May), FINOPS-Driven Cloud optimization models for enterprise applications, Vol. 52 No. 2 (2024): April-June 2024, 99-110, *Power System Protection and Control*, ISSN-1674-3415, URL: <https://pspac.info/index.php/dlbh/article/view/283>, DOI: <https://doi.org/10.46121/pspc.52.2.10>
29. Naresh Lokiny. (2022). Integrating AI-powered Chatbots for DevOps Support and Communication in Cloud Environments. *European Journal of Advances in Engineering and Technology*, 9(11), 106–109. <https://doi.org/10.5281/zenodo.13325989>
30. Naresh Lokiny, (2021), "Disaster Recovery and Business Continuity Planning in DevOps Cloud with AI", *International Journal of Science and Research (IJSR)*, 10(3), 2024-2027. <https://dx.doi.org/10.21275/SR24724151733>, <https://www.ijsr.net/getabstract.php?paperid=SR24724151733>
31. Naresh Lokiny, & Ranganath Nandanampati. (2020). DevSecOps: Integrating Security into DevOps with AI in Cloud. *Journal of Scientific and Engineering Research*, 7(10), 239–242. <https://doi.org/10.5281/zenodo.13348695>
32. Naresh Lokiny, & Pradip Reddy. (2021). Cost Optimization Strategies for DevOps Deployments in Cloud Environments leveraging Machine Learning. *European Journal of Advances in Engineering and Technology*, 8(3), 69–72. <https://doi.org/10.5281/zenodo.13325845>
33. Jayanth Para, 6G Internet Technology Cyber Threat Notification & Alert System, Vol. 52 No. 4 (2024): October-December 2024, *Power System Protection and Control*, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/170>, DOI: <https://doi.org/10.46121/pspc.52.4.9>
34. Jayanth Para, AI Based Cloud Computation Observational Method & Process, Vol. 51 No. 4 (2023): October-December 2023, *Power System Protection and Control*, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/171>, DOI: <https://doi.org/10.46121/pspc.51.4.3>
35. Kaleshwar Aryasomayajula, PREVENTING BIAS IN AI-BASED DECISION-MAKING: ANALYZING TECHNIQUES TO REMOVE UNFAIR PREJUDICE IN ALGORITHMIC OUTCOMES, Vol. 50 No. 2 (2022): April-June 2022, *Power System Protection and Control*, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/294>
36. Kaleshwar Aryasomayajula, MODERN DEVELOPMENT PRACTICES: INVESTIGATIONS INTO SERVERLESS COMPUTING, LOW-CODE/NO-CODE PLATFORMS, AND DEVELOPER PRODUCTIVITY IN REMOTE ENVIRONMENTS, Vol. 50 No. 4 (2022): October-December 2022, *Power System Protection and Control*, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/364>
37. Kaleshwar Aryasomayajula, Micro services: "A Comparative Analysis of Monolithic vs. Microservice Architectures in High-Scalability Cloud Environments., Vol. 51 No. 2 (2023): April-June 2023, *Power System Protection and Control*, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/374>
38. Vikas Suresh Bagora, KERNEL-LEVEL PERFORMANCE OPTIMIZATION FOR CARRIER-GRADE BROADBAND AND HIGH-SPEED NETWORKING SYSTEMS, Vol. 47 No. 1 (2019), *Power System Protection and Control*, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/359>
39. Vikas Suresh Bagora, SCALABLE 128G FIBRE CHANNEL AND ETHERNET CONVERGENCE FOR NEXT-GENERATION DATA CENTER NETWORKS, Vol. 50 No. 4 (2022): October-December 2022, *Power System Protection and Control*, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/362>
40. Vikas Suresh Bagora, SECURE EMBEDDED NETWORKING ARCHITECTURES USING TRUSTED EXECUTION ENVIRONMENTS IN BROADBAND GATEWAYS, Vol. 52 No. 2 (2024): April-June 2024, *Power System Protection and Control*, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/363>

41. Santhosh Chitraju Gopal Varma, The Evolution of Cloud-Native Architectures: Exploring the Synergy between Kubernetes and Microservices, DOI: <https://doi.org/10.63282/3050-9246.IJETCSIT-V1I4P104>, Vol. 1 No. 4 (2020), <https://ijetcsit.org/index.php/ijetcsit/article/view/108>
42. Santhosh Chitraju Gopal Varma, Artificial Intelligence in Cloud Computing: Building Intelligent, Distributed, and Fault-Tolerant Systems, Vol. 3 No. 1 (2022), DOI: <https://doi.org/10.63282/3050-9416.IJAIBDCMS-V3I1P105>, <https://ijaibdcms.org/index.php/ijaibdcms/article/view/78>
43. Sunil Karmakar, AI driven Identity & Access Management with Oracle Cloud security, Vol. 30 No. 2 (2022): JOCAAA, <https://eudoxuspress.com/index.php/pub/article/view/4118>
44. Rakshitha Kiran P., Samitha Khaiyum, Aparna R. Palandye, Apoorva S. D., Arjun R., Akshaya S. A., Annapoorna S., Leveraging LLaMA3 and LangChain for Rapid AI Application Development, Vol. 20 No. 10s (2024), <https://doi.org/10.52783/jes.5539>, <https://journal.esrgroups.org/jes/article/view/5539>.