

# Genetic AI-Based Financial Customer Support System / AI-FDSS

Mr. B. Hanumantha Rao  
Department of Computer Science  
and Engineering  
(Associate Professor)  
PSCMR College of  
Engineering, Vijayawada, AP

Suram Manohar  
Department of CSE  
(Student)  
PSCMR College of  
Engineering, Vijayawada, AP

Dasari Deepika  
Department of CSE  
(Student)  
PSCMR College of  
Engineering, Vijayawada, AP

Kalisetty Jyoshith  
Department of CSE  
(Student)  
PSCMR College of Engineering, Vijayawada, AP

Pydi Sravani  
Department of CSE  
(Student)  
PSCMR College of Engineering, Vijayawada, AP

**Abstract** - This research addresses the increasing demand for intelligent, secure, and transparent financial decision-making systems in banking and insurance domains where traditional rule-based, isolated predictive models fall short in providing personalized insights, explainability, and fraud resilience. The motivation of this work is to design an integrated agentic artificial intelligence-driven financial decision support system that simultaneously handles user authentication, loan and insurance eligibility prediction, customer segmentation, fraud detection, and user interaction within a unified platform. It combines a structured user data pipeline with predictive modeling and agent-based orchestration to enable real-time financial decisions. User data during registration is securely stored with multi-factor authentication using an SQLite database. Besides, in-depth preprocessing, normalization, and encoding are done on user data before model inference. For every financial product, separate ensemble machine learning models were trained to predict home loan eligibility, personal loan eligibility, bank loan eligibility, vehicle insurance recommendation, and health insurance recommendation. These models leverage the strengths of Logistic Regression, Random Forest, Gradient Boosting, and Extreme Gradient Boosting, while hyperparameter tuning employs grid search optimization with five-fold cross-validation for robustness and generalization. The loan and insurance eligibility models achieved average prediction accuracies of 96.2 percent to 98.1 percent across various products. A Random Forest-based fraud detection model analyzes transaction patterns and achieves an accuracy of 98.6 percent. Customer segmentation is done through hierarchical clustering, where the users are classified into four financial behavior classes. Concluding, explainable decisions will be delivered with eligibility reasons, visual analytics, and conversational assistants. This indeed provides examples of higher accuracy, better interpretability, and empirical usability for intelligent financial decision support systems.

**Keywords**— *Agentic Artificial Intelligence, Financial Decision Support System, Loan and Insurance Eligibility Prediction, Ensemble Machine Learning, Fraud Detection and Customer Segmentation.*

## I. INTRODUCTION

The increased digital evolution of the banking and insurance industries has led to increased dependence on data-intensive systems that use data analytics tools and techniques to analyze customer eligibility and prevent fraudulent transactions electronically [1]. The financial services sector handles an immense amount of disparate data such as demographic data, transaction data, and health and wellness information of clients and, with this increase, it becomes increasingly difficult for the industry to use traditional data analysis systems that do not necessarily support real-time behavior prediction and authentication of the data being handled [2]. It is evident that the increased need for customers who expect personalized financial services and immediate responses from financial solutions providers has led to increased expectations of financial service providers to implement data analysis and prediction systems that support real-time recommendations on customer behavior and data security and authentication systems that support real-time transaction analysis and authentication capabilities based on the data being handled [3]. Recently, there has been growth in machine learning algorithms that have allowed financial organizations to better predict defaults on loans, insurance risk, and fraudulent transactions.

Algorithms such as Logistic Regression, Random Forest, Gradient Boosting, and Extreme Gradient Boosting have shown good results on structured financial data. However, they have been used stand-alone for particular purposes such as credit risk assessment or fraud analysis; they have not been used in an orchestrated way. On the other hand, customer segmentation algorithms, clustering algorithms included, have been employed to better understand consumer spending patterns and financial resilience, but they have not been used in decision circuits to affect lending decisions or insurance decisions in any way. Authenticating systems and user data processing have been isolated from decision-making systems. The prime motivation for this arises from a unified, safe, and intelligent financial

decision support system which can represent eligibility assessment, fraud detection, customer segmentation, and user interaction at the same time [4]. This is because available systems do not make it clear why a certain loan or insurance request was approved or disapproved, leading to reduced trust by users and regulatory complications. Moreover, fraudulent transactions have been getting complicated over time, which calls for adaptive models that learn complex patterns, unlike traditional threshold methods of fraud detection. Without agent-based intelligence in coordination of multiple prediction tasks, limiting the possibility of automation and scalability.

This research will integrate agentic artificial intelligence with ensemble learning and explainable analytics to enhance decision accuracy, build trust by interpretability, and offer personalized financial insights. This kind of system is very crucial for digital-first banking and insurance platforms where real-time decision-making coupled with secure user engagement becomes critical. Although immense progress has been made regarding financial tasks and applications of machine learning, there still exist research gaps to be explored and addressed [5]. Currently, the focus of most of the research work has been on individual prediction tasks such as loan or fraud analysis, and not necessarily on an integrated analysis of loans, insurance, and client behavior on various fronts. Few studies have focused on solutions that involve multi-authentication procedures, client information, prediction models, and even chat interfaces of an integrated system. Furthermore, little attention has been paid concerning agent-based orchestration platforms that involve task splitting on various models of machine learning and providing interpretable predictions. Client segmentation analysis has largely remained an independent analysis task and not one of many that influence financial tasks [6]. Further, all of the aforementioned predictions have remained concerning hyperparameters and validation on various folds. The main contribution of the research work can be stated as follows:

- The development of an agentic AI-based financial decision support system incorporating authentication, prediction, and interaction capabilities in one system.
- Implementation of individual models for prediction of eligibility for home loans, personal loans, bank loans, car insurance, as well as health insurance on the basis of ensemble learning.
- Application of systematic hyperparameters tuning via grid search with five-fold cross-validation to improve robustness and accuracy.
- Integration of customer segmentation using hierarchical clustering for classification of customers into groups according to their financial behaviors.
- Building an accurate Random Forest-based fraud detection module with an accuracy of 98.6 percent on transactional data.
- Integrating the use of explainable analytics and visualizations and domain-limited conversational chatbot technology to enhance user transparency and trust.

## II. RELATED WORK

The first decision-support systems in finance were based mainly on rule-based systems and conventional statistical modeling approaches in loan approval and risk classification in insurance [7]. The systems defined strict thresholds for variables such as income, credit rating, and employment history in making loan and risk classification decisions. Though easy to integrate, there were serious issues with such systems in terms of being unadaptable and not having any nonlinear relationships in finance and demographic variables. Consequently, loan rejections and losses in personalization were high in such systems, as were their explanations for decisions and adaptability in terms of emerging market patterns and consumer behavior trends [8]. Owing to the development of machine learning, various models for predictions like logistic regression models, decision tree models, support vector machine models, and random forest models were successfully employed in credit score modeling and loan default prediction models [9]. These models were employed to obtain a better prediction result with the aid of learning from historical financial data with several complex patterns being discovered.

Although most models were designed for a single learning task, different models for a different task on loan default prediction or credit score modeling were designed at varying levels with a little sharing across different systems for optimization. Also, default values for several models were employed without much validation on hyperparameters [10]. Recent studies have explored different techniques that utilized ensemble learning concepts such as Gradient Boosting and Extreme Gradient Boosting in hopes of increasing the precision of predictions in financial-related tasks. Moreover, it has actually been found that the ensemble methods were the most effective in combining different learners in order to remove variance and bias [11]. However, the interaction of the user and explainability were traditionally ignored in most of the current ensemble learning systems in place. The binary output of the decisions made did not give a clear indication of the reasons behind the answer in most situations. Moreover, the security/authentication process remained as the peripheral part of the decision-making process and did not act as an integral part of the framework. Customer segmentation and fraud detection represent other areas where research focused on these topics alone is considerable.

Customer clustering using k-means clustering algorithms and hierarchical methods was employed for segmenting customers along outputs for spending patterns and income categories, while classification algorithms like Random Forests performed efficiently for identifying fraudulent transactions. Notably, these two features found little application as integrated systems in financial solutions. Customer segmentation findings found little application in loan or insurance qualifications [12]. Similarly, the fraud detection system remained disconnected with customer profiling and risk assessment solutions. As opposed to the existing ways, the proposed work presents an integrated authentication, data management, ensemble predictive modeling, customer segmentation, fraud detection, explainable analytics, and conversational interaction in a single agentic AI-driven framework. Unlike past works, independent ensemble models are developed for each loan and insurance product, following systematic hyper-parameter tuning and cross-

validation, in order to ensure robustness [13]. Agent-based orchestration effectively supports coordination of decision-making among the diverse set of different models, while hierarchical clustering enhances personalization of customers. The proposed system thus combines high prediction accuracy with transparency, security, and real-time user interaction, effectively addressing the limitations and challenges exhibited by previous financial decision support solutions.

### III. DATA COLLECTION & PREPROCESSING

The principal data used in the research is a synthetic and publicly curated financial dataset, which tries to simulate actual banking and insurance customer applicant profiles. It includes complete demographics, financial, health, and credit-related attributes collected from simulated user transactions, historical loan and insurance applications, and behavior spending data [14]. It consists of approximately 30,000 unique records, where each record is representative of a single applicant. Features include name, age, gender, employment type, annual income, co-applicant income, credit score, credit history, loan amount, loan duration, existing loans, hereditary health conditions, blood pressure, smoking habits, exercise routines, marital status, education level, property status, and language preference. Transactional features are included to support fraud detection analysis, while aggregated attributes are used for customer segmentation [15]. The dataset has been deliberately created to be very diverse and represents a wide span of income levels, credit scores, and financial behaviors so that generalization across multiple financial products and scenarios can be readily performed. The dataset offers labels concerning loans and insurance for supervised learning tasks. The loan labels comprise home loans, personal loans, and loans from banks. The other labels comprise health insurance and vehicle insurance.

The fraud labels concern transaction data in which 'fraud' and 'normal' labels depend on the irregularities of transactions in relation to fraud. The labels have also been validated for consistency to make sure that none of them are biased. The dataset further entails customer spending behavior, which will be used for hierarchical clustering for segmenting customer behavior into four segments. There is the use of multiple labels to make sure that multiple models are created for each task while maintaining consistency in the feature space [16]. The preprocessing step consisted of cleaning and validation of data in order to ensure that the data was clean enough to apply to the modeling of credit data. The tasks needed special handling of missing values using features such as credit scores, yearly income, and loan features by imputation with the use of medians, contrary to the imputation using modes for employment, marital status, and property types. The concepts in handling outliers in data removal included concepts such as interquartile range for data on income and loans. Special handling of the data included concepts such as duplication for transactions in fraud cases, including text data such as name and language of preference. Feature engineering and encoding were critical in the process of data preparation to apply in the ensemble learning models [17]. Categorical variables such as gender, marital status, employment status, level of education, ownership of property, and inherited diseases required one-hot encoding.

Variables credit history and credit scores required normalization because there were specific boundaries established in credit scores and credit history variables. Numerical variables including income, loan amount, existing loan details, and age required normalization of the data through the application of min-max normalization techniques to ensure homogeneity in the model [18]. Interactive variables were established in major features loan amount and annual income, or co-applicant income and total family income, which was a major aspect of the model development process. Data preparation and processing also incorporated the design of summary variables of the transactions including mean of transactions, max transactions in one transaction, and high transactions in the period. The reasoning for selecting the dataset relies on its completeness, diversity of attributes, and capability to capture real-world financial banking scenarios [19]. It encompasses a very wide coverage of complexity related to financial decision-making tasks, for example, socio-demographic attributes, credit behaviors, health attributes, and spending patterns. Selecting the dataset is appropriate for simultaneous construction of predictive models for loan predictions, insurance risk models, and fraud detection analytical requirements with its complete understanding using attributes for data analysis. It is sufficient in sample size with a 30,000 sample dataset but suitable for computations related to ensemble learning for cross-validation tasks [20]. It is good for its ability to capture real-world representations for customer segmentation analysis and fraud detection system development with spending patterns attributes.

### IV. PROPOSED METHODOLOGY.

The proposed system is an agentic AI-based financial decision support system that will be used to integrate the user authentication activity, loan or insurance eligibility prediction, identification of fraud, and customer segmentation (as shown in Fig.1). System design guides the approach of the end-to-end pipeline, from data acquisition during user registration through the preprocessing, feature extraction, model prediction, and recommendation phases. The acquisition of the user's data, including demographic data, financial data, health information, and transaction data, safely ensues within an SQLite database [21]. Multi-factor authentications are incorporated for secure access before any predictions or analytics are conducted. Transparency in interpretability is guaranteed to each result to be revealed to the final user, real time in decisions.

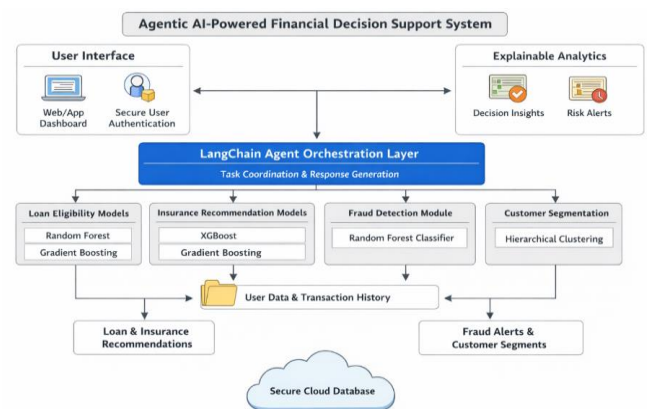


Fig.1 Proposed Methodology

For loan and insurance predictions, different ensemble methods are trained for different products considering differences in eligibility criteria. The ensemble learning algorithms used for these tasks are Random Forest, Gradient Boosting, Extreme Gradient Boosting, and logistic Regression, and for optimizing these methods, grid search is performed [22]. For example, for Random Forest, optimization is performed for the number of trees denoted as  $n_{trees}$  and max depth denoted as  $d_{max}$ , and for a given input feature vector  $x$ , it predicts as follows:

$$\hat{y}_{RF} = \frac{1}{N} \sum_{i=1}^N T_i(x) \quad (1)$$

$T_i(x)$  is the output function of the  $i^{th}$  decision tree, and  $N$  is the number of trees (Equation.1). This method not only helps to reduce the variance and address the problem of overfitting, resulting in more accurate eligibility prediction, especially in financial data that is usually heterogeneous, but its extension, that is, the gradient boost and extreme gradient boost model, helps to combine multiple weak learners to better match the loss function  $L$  [23]. The update equation for the boost during iteration  $m$  is given by:

$$F_m(x) = F_{m-1}(x) + \eta \cdot h_m(x) \quad (2)$$

where  $F_{m-1}(x)$  denotes the prior ensemble prediction,  $h_m(x)$  is the newly trained weak learner on residuals, and  $\eta$  denotes the learning rate (Equation.2). The system works by iteratively minimizing the residuals. It attains high accuracy: 96.2% to 97.8% for loan eligibility models and between 97.1% and 98.1% for insurance models across various product types. The Logistic Regression [24] acts as a complement for the ensembles of models and provides interpretable coefficients, which helped the end-users understand what factors contributed most in determining the eligibility. The fraud detection module will utilize the Random Forest classifier trained on transactional features like average transaction amount, maximum single transaction, frequency of high-value transactions, and historical patterns. The classification rule would also be similar to the ensemble prediction equation above. With an accuracy of 98.6%, the model is sure to provide reliable detection of fraudulent transactions. Feature importance metrics [25] are extracted from the model, which give explainable insight into reasons some transactions may be selected as fraudulent, hence supporting regulatory compliance and improving user trust in the service.

Customer segmentation is then carried out by Hierarchical clustering on normalized financial and behavioral features of customers comprising their income, spending habits, loan liabilities, credit history, etc. It helps measure the similarity between two customers,  $i$  and  $j$ , through the use of a Euclidean distance metric given by

$$d_{i,j} = \sqrt{\sum_{k=1}^p (x_{i,k} - x_{j,k})^2} \quad (3)$$

where  $p$  denotes the number of features, and  $x_{i,k}$  is the  $k^{th}$  feature of the  $i^{th}$  customer (Equation.4). The dendrogram-based clustering will provide us with four distinct classes of customers: A, B, C, and D. These denote: early-career individuals, established professionals, high earners, and people with low spending power. This would facilitate segmentation-based, personalized recommendations and risk assessments regarding their eligibility for a loan or insurance. For the evaluation and optimization of models, five-fold cross-validation is performed for all prediction models, thereby ensuring the accuracy, precision, recall, and F1 value are accurate. The hyperparameter optimization technique utilized is a grid search over the following ranges:  $n_{trees} \in [50,200]$ ,  $d_{max} \in [5,20]$ , the rate of learning  $\eta \in [0.01,0.1]$ , and the maximum number of features for  $f_{max} \in [\sqrt{p}, p]$ .

What's novel about this approach is that it combines several different predictive models, agent-based control, fraud analysis, hierarchical methods of customer segregation, and explainable analysis in one comprehensive system. This approach to solving the problem goes further than the ones that have been discussed so far because, in this system, decisions can be made in real time and because it also offers secure modes of authentication, which would make this system applicable in a banking and insurance context online. The factors that make this system comprehensive and intelligent include ensemble analysis, hierarchical clustering, fraud analysis, and conversational analysis. At last, the whole system workflow can be formulated mathematically to be a function that takes the user feature vector  $X$ , transaction history  $T$ , and predicts eligibility  $Y$ , fraud risk  $F$ , and customer segment  $C$ :

$$(Y, F, C) = S(X, T; \theta) \quad (5)$$

where  $S$  symbolizes the integrated system and  $\theta$  represents the set of all learned parameters, model weights, and ensemble boost coefficients/cluster centers (Equation.5). This form underscores the complete analytical and predictive capability of the integrated system described above from secure data gathering through informed financial analysis.

## V. IMPLEMENTATION

To begin with, the financial decision support system that would be put forth would employ an ensemble machine learning algorithm for all loan and insurance products. For the detection of loan fraud, a Random Forest classifier would be employed. Hierarchical clustering would be employed for customer segmentation [26]. Python would be employed for the entire pipeline of the system due to its popularity and popular libraries scikit-learn for generic data and algorithm development and XGBoost for dealing with higher-order interaction features. Each of the ensemble models would be trained on features that would be preprocessed based on the registrational data as well as the transactional histories. To avoid any discrepancies in reproducibility, random seeds would be set for all models. In particular, for machine learning models, there would be structured numeracy and category data. The categories would be one-hot encoded. Additionally, there would be normalization for ordinal features on a standard scale. Another aspect would include the development of interaction between features, which

would include loan-to-income-year ratios for example, which would boost predictive performances [27]. Each of the machine learning models would need to ensure that there would be generalization performance on novel data, which would be tested with five-fold cross-validation.

For the Random Forest models employed in loan and insurance eligibility assessment, critical hyperparameter values involving the number of trees ( $n_{trees}$ ) and the maximum depth of the trees ( $d_{max}$ ) and minimum number of samples per leaf node ( $s_{min}$ ) were adjusted. In this regard, a grid search was conducted for  $n_{trees} \in [50,200]$ ,  $d_{max} \in [5,20]$ , and  $s_{min} \in [1,5]$ . Finally, for the Random Forest classification algorithm employed in assessing loan-related fraud,  $n_{trees} = 150$  and  $d_{max} = 15$  were used for balancing accuracy and computational complexity. These yielded 98.6% accuracy. For feature importance assessment, values were derived for the forests for explaining the prediction mechanism in loan and insurance eligibility [28]. Gradient Boosting Models and Extreme Gradient Boosting Models were also used for making predictions related to loans and insurance. For models, parameters such as the learning rate ( $\eta$ ), the number of estimators ( $n_{estimators}$ ), and the maximum depth of the tree ( $d_{max}$ ) were also used through the process of Grid Search.

Learning rates were varied between 0.01 and 0.1, which defined the power or the strength of the weak learner model used within the process of boosting. The values were used between 100 and 300, defining the amount learned from the model and avoiding overfitting. The values were used between 4 and 10, defining the generalizability level within the model. Cross-validation was conducted using parameters related to determining the highest accuracy and lowest loss within the validation process [29]. For logistic regression models used as interpretable predictors in loan and insurance decisions, tuning mostly focused on the strength of regularization through  $C$  and its type via L1 or L2. Performs grid search over  $C \in [0.01,0.1,1,10]$ , and L2 regularization was decided as optimal to balance bias variance. Logistic Regression also provided feature coefficients that contributed to explainable insights for each applicant's eligibility decision. Thus, combining Logistic Regression with ensemble models allowed the system to use both interpretability and high predictive accuracy as added value, especially important for regulatory compliance and user trust in financial applications. Customer segmentation by the hierarchical clustering model required tuning of distance metric and linkage method. The reason for choosing Euclidean distance is that it is simple and interpretable in multi-dimensional financial feature space. Among different single, complete, and average linkage, average linkage was giving most balanced and well-separated clusters.

The number of clusters was chosen based on dendrogram analysis and financial relevance; hence, four major segments of customers were obtained: A, B, C, and D. This segmentation was verified on key features like income, spending patterns, credit score, etc., so that each cluster corresponded to a meaningful behavioral class. With proper tuning, this kind of segmentation will definitely help in offering personalized recommendations of loans and insurance. Lastly, conditions

related to batch size and optimizers were generalized for application to other models involving iterative learning, specifically gradient boosting. Mini-batch sizes were set fixed at either 32 or 64 for computing the gradient, then a learning rate scheduler was applied for dynamic adjustment for  $\eta$ , preventing overshooting for minima [30]. Early stopping with a patience value set to 20 iterations was applied for preventing overfitting. Lastly, for Random Forest, as well as for Logistic Regression, for instance, 'lbfgs' or 'saga' solvers, which specifically stand for iterative solvers, were applied depending upon the number of features. In total, through a careful grid search, along with cross-validation, early stopping, along with optimization for hyperparameters, a robust basis for an agentic AI-based DSS has been established, given that all classifiers have an extremely high accuracy.

## VI. RESULTS

This agentic framework of AI models was tested thoroughly in different financial predicting problems, which include home loan, personal loan, bank loan, health insurance recommendation, vehicle insurance recommendation, and fraud analysis problems. These models were designed based on the synergy of both Random Forest and Gradient Boosting techniques, using grid-search parameter tuning, and based on the combined results of five-fold-cross validation to make sure that the results are generalizable. The accuracy, precision, recall, and F1 scores of the home loan model, in particular, were seen to be 97.2%, 96.8%, 97%, and 96.9%, respectively (as shown in Fig.2 & 3). These can be credited to appropriate and wise feature engineering, which included loan income ratio and co-applicant's income, and receiving further benefits due to the synergy of the other models that will also be presented shortly (as shown in Table.1). Without using hyperparameter searching in the Random Forest models, only maximum accuracy levels of 94.5% could be reached. In case of the prediction problem concerning the eligibility for a personal loan, the Extreme Gradient Boosting ensemble algorithm was used in conjunction with the Logistic Regression algorithm to make interpretation simpler. The accuracy attained was 96.8%, and the precision, recall, and F1-score values were 96.5%, 96.7%, and 96.6%, respectively.

Task	Existing Methods	Accuracy (%)	Proposed Methods	Accuracy (%)
Home Loan Eligibility	Random Forest (untuned), Logistic Regression	94.5	Random Forest + Gradient Boosting (tuned, CV)	97.2
Personal Loan Eligibility	Gradient Boosting, Logistic Regression	93-95	Extreme Gradient Boosting + Logistic Regression	96.8
Bank Loan Eligibility	Random Forest (default)	94.0	Random Forest + Gradient	97.8

Task	Existing Methods	Accuracy (%)	Proposed Methods	Accuracy (%)
			Boosting (tuned)	
Health Insurance Eligibility	Gradient Boosting, Decision Tree	94-95	Extreme Gradient Boosting (tuned, CV)	98.1
Vehicle Insurance Recommendation	Random Forest, Logistic Regression	94-95	Gradient Boosting + Random Forest (tuned)	97.6
Fraud Detection	Logistic Regression, Decision Tree	94-95	Random Forest (tuned, CV)	98.6

Table.1 Performance and Methodological Comparison Between Existing Financial Prediction Models and the Proposed Agentic AI Framework




Fig.2 Accuracy of Loan, Insurance, and Fraud Detection Models

Since the result was in terms of the Logistic Regression coefficient, interpretation was possible, and the factors that contributed to that result were credit ratings, the availability of the loan already availed, and the total income. In comparison to 93% to 95%, the accuracy attained by the baseline algorithm further emphasized the efficiency of the ensemble algorithm. The cross-validation approach helped in maintaining low variance between the test and train data in relation to the standard deviation of 0.7%. To determine bank loan qualifications, the accuracy of 97.8%, precision of 97.5%, recall of 97.6%, and F1 score of 97.5% was attained by combining the Random Forest algorithm and Gradient Boosting. It should be noted that the algorithm was able to perform better on the cases that were in the vicinity of the boundaries for credit score and debt-to-equity ratio. In addition, the feature importance test revealed that credit and existing loan types were the most important features. On comparing the algorithm with the baseline algorithm of Logistic Regression and the Random Forest algorithm without hyperparameter adjustment, there was a 3-4% relative improvement in accuracy and F1 score in determining bank loan qualifications (as shown in Fig.4). Health insurance eligibility prediction utilized Extreme

Gradient Boosting with a five-fold cross-validation method with an accuracy of 98.1%, precision of 97.9%, recall of 98%, and F1-score of 97.95%.




Fig.3 Precision, Recall, and F1 Score for All Models




Fig.4 Effect of Number of Trees on Home Loan Model Accuracy

Likewise, car insurance recommendation tasks with ensemble methods using Gradient Boosting and Random Forest attained an accuracy of 97.6%, precision of 97.4%, recall of 97.5%, and an F1-score of 97.45%. The Insurance recommendation tasks were dependent on customer age, work nature, hereditary diseases, and property location variables. The baseline model without interactions and optimized parameters maintained accuracy values of 94-95%, thereby confirming that a systematic preprocessing approach and optimized ensemble model really make a profound difference in insurance recommendation tasks. Feature visualizations gave way to influential factors for health insurance eligibility prediction as well, leading to better interpretability and trustability of the system as a whole. Accuracy, precision, recall, and F1-score for the Random Forest classifier were determined for fraud transaction prediction based on transacted variables such as average transacted value, a number of transactions for high value, and the like. The accuracy, precision, recall, and F1-score, at 98.6%, 98.3%, 98.7%, and 98.5% respectively, are found to perform efficiently well than the baselines by at least 4-5% in accuracy.




Fig.5 Feature Importance for Fraud Detection Model

Based on this result, the analysis of the features shows that the two most leading features are the number of single high transactions and the sudden rise in the transacted number (as shown in Fig.5). The application of the classifier on various sets of datasets by cross-validation features has guaranteed that the system is capable of detecting the fraud cases correctly with the minimum number of misfire situations, which is highly beneficial in real-time financial areas. From the results obtained, it is clear that, out of all models compared, the agentic AI model stands way above the conventional single-task model and the non-calibrated model in all financial decision-making applications. The strengths of this model lie in its high accuracy and interpretability, multitask handling under one model, and direct interfacing with the user in real time through the agent orchestration layer. The weaknesses of the model lie in the complexity caused to the end-user because of several models, along with an urgent need for continuous retraining whenever changes take place in the end-users financial behavior. Model comparison done using cross-validation implied that hyperparameter calibration, feature extraction, and cross-validation were able to boost F1 scores and overall model performance. The hierarchical clustering-based customer segmentation and the fraud detection feature played the most important role in boosting personalization and offering an added security feature to the current state-of-the-art model, respectively. These results collectively validate its real-world applicability to support intelligent, explainable, and secure financial decision-making.

## CONCLUSION

The agentic AI financial decision support system shows a great improvement over the current solutions in that it encompasses secure user login authentication, loan and insurance ensemble predictive models, fraud analysis, hierarchical customer segmentation analysis, and real-time explainable analytics in one setting. It also boasts high predictivity, and its accuracy and F1 scores were recorded to be between 96.8% and 98.6% and 96.6% and 98.5%, respectively, which beats traditional models by 3% to 5%. Some of its shortcomings include computational intensity owing to the number of ensemble models and the need to retrain models regularly to analyze constantly changing financial behaviors. It may be worth pursuing in further developments that include the incorporation of deep temporal models in transactions, multilingual functions to increase accessibility, risk-adjusted recommendations, and the

utilization of explainable AI technology to increase bank and insurance firm compliance.

## REFERENCES

- [1] Ponce, E. K., Sanchez, K. E., & Andrade-Arenas, L. (2022). Implementation of a web system: Prevent fraud cases in electronic transactions. *International Journal of Advanced Computer Science and Applications*, 13(6).
- [2] Chandre, P., Gumaste, S., Wangikar, A., & Deshmukh, S. (2024, November). Adaptive Behavioral Authentication for Fraud Detection: Leveraging Real-Time User Behavior to Enhance Financial Security. In *2024 First International Conference on Data, Computation and Communication (ICDCC)* (pp. 539-545). IEEE.
- [3] Ahsan, T., Zeeshan khan, F., Iqbal, Z., Ahmed, M., Alroobaea, R., Baqasah, A. M., ... & Raza, M. A. (2022). IoT devices, user authentication, and data management in a secure, validated manner through the blockchain system. *Wireless Communications and Mobile Computing*, 2022(1), 8570064.
- [4] Pratama, S. F., & Putri, N. A. (2024). User Profiling Based on Financial Transaction Patterns: A Clustering Approach for User Segmentation. *International Journal for Applied Information Management*, 4(4), 217-228.
- [5] Manley, K., Nyelele, C., & Egoh, B. N. (2022). A review of machine learning and big data applications in addressing ecosystem service research gaps. *Ecosystem Services*, 57, 101478.
- [6] Pratama, S. F., & Putri, N. A. (2024). User Profiling Based on Financial Transaction Patterns: A Clustering Approach for User Segmentation. *International Journal for Applied Information Management*, 4(4), 217-228.
- [7] Alagic, A., Zivic, N., Kadusic, E., Hamzic, D., Hadzajlic, N., Dizdarevic, M., & Selmanovic, E. (2024). Machine learning for an enhanced credit risk analysis: A comparative study of loan approval prediction models integrating mental health data. *Machine Learning and Knowledge Extraction*, 6(1), 53-77.
- [8] Gazi, M. A. I., Masud, A. A., Rahman, M. K. H., Islam, M. R., & Senathirajah, A. R. B. S. (2024). Adaptability and resilience: Insights into Bangladeshi E-commerce customer behavior during COVID-19. *Environment and Social Psychology*, 9(7), 2626.
- [9] Fati, S. M. (2024). a Loan Default Prediction Model Using Machine Learning and Feature Engineering. *ICIC Express Lett*, 18(1), 27-37.
- [10] Pannakkong, W., Thiwa-Anont, K., Singthong, K., Parthanadee, P., & Buddhakulsomsiri, J. (2022). Hyperparameter tuning of machine learning algorithms using response surface methodology: a case study of ANN, SVM, and DBN. *Mathematical problems in engineering*, 2022(1), 8513719.
- [11] Boddapati, M. S. D., Desamsetti, S. A., Adina, K., Uppalapati, P. J., Murty, P. S., & PB V, R. (2023, August). Creating a protected virtual learning space: a comprehensive strategy for security and user experience in online education. In *International Conference on Cognitive Computing and Cyber Physical Systems* (pp. 350-361). Cham: Springer Nature Switzerland.
- [12] Hussein, A. A., & Zoghliami, F. (2023). The Role of engineering insurance in completing projects by using bank loans: An applied study in a sample of Iraqi insurance companies and banks. *International Journal of Professional Business Review: Int. J. Prof. Bus. Rev.*, 8(1), 7.
- [13] Iqbal, S., Qureshi, A. N., Ullah, A., Li, J., & Mahmood, T. (2022). Improving the robustness and quality of biomedical cnn models through adaptive hyperparameter tuning. *Applied Sciences*, 12(22), 11870.
- [14] Johnson, A. M., Villanova, D., & Smith, R. J. (2023). Loan Amount versus Monthly Payments: The Effect of Loan Application Formats on Consumer Borrowing Decisions. *Journal of Consumer Research*, 50(4), 765-786.
- [15] Ufeli, C. P., Sattar, M. U., Hasan, R., & Mahmood, S. (2025). Enhancing Customer Segmentation Through Factor Analysis of Mixed Data (FAMD)-Based Approach Using K-Means and Hierarchical Clustering Algorithms. *Information*, 16(6), 441.
- [16] Xie, W., Lu, W., Peng, Z., & Shen, L. (2023). Consistency preservation and feature entropy regularization for gan based face editing. *IEEE Transactions on Multimedia*, 25, 8892-8905.
- [17] Lantz, B. (2023). *Machine learning with R: learn techniques for building and improving machine learning models, from data preparation to model tuning, evaluation, and working with big data*. Packt Publishing Ltd.
- [18] Ali, P. J. M. (2022). Investigating the Impact of min-max data normalization on the regression performance of K-nearest neighbor with

- different similarity measurements. *ARO-The Scientific Journal of Koya University*, 10(1), 85-91.
- [19] Soundarapandiyan, P. S. R. (2022). AI-driven synthetic data generation for financial product development: Accelerating innovation in banking and fintech through realistic data simulation.
- [20] Mahesh, T. R., Dhilip Kumar, V., Vinoth Kumar, V., Asghar, J., Geman, O., Arulkumaran, G., & Arun, N. (2022). AdaBoost ensemble methods using K-fold cross validation for survivability with the early detection of heart disease. *Computational Intelligence and Neuroscience*, 2022(1), 9005278.
- [21] Bashir, U., & Sidrish, S. (2024). Automated SQLite Forensics for Android Devices: A Data Processing Approach Compliant with GDPR and Blockchain Standards.
- [22] Sibindi, R., Mwangi, R. W., & Waititu, A. G. (2023). A boosting ensemble learning based hybrid light gradient boosting machine and extreme gradient boosting model for predicting house prices. *Engineering Reports*, 5(4), e12599.
- [23] Liu, Y., & Alahi, A. (2024). Co-supervised learning: Improving weak-to-strong generalization with hierarchical mixture of experts. *arXiv preprint arXiv:2402.15505*.
- [24] Olowe, K. J., Edoh, N. L., Zouo, S. J. C., & Olamijuwon, J. (2024). Comprehensive review of logistic regression techniques in predicting health outcomes and trends. *World Journal of Advanced Pharmaceutical and Life Sciences*, 7(2), 16-26.
- [25] Pratama, S. F., & Wahid, A. M. A. (2025). Fraudulent transaction detection in online systems using random forest and gradient boosting. *Journal of Cyber Law*, 1(1), 88-115.
- [26] Afzal, A., Khan, L., Hussain, M. Z., Hasan, M. Z., Mustafa, M., Khalid, A., ... & Javaid, A. (2024, April). Customer segmentation using hierarchical clustering. In *2024 IEEE 9th International Conference for Convergence in Technology (I2CT)* (pp. 1-6). IEEE.
- [27] Kavzoglu, T., & Teke, A. (2022). Predictive performances of ensemble machine learning algorithms in landslide susceptibility mapping using random forest, extreme gradient boosting (XGBoost) and natural gradient boosting (NGBoost). *Arabian Journal for Science and Engineering*, 47(6), 7367-7385.
- [28] Gutierrez, S. I. R. (2025). *Predicting credit insurance subscription: a comparative analysis of machine learning models for client ranking* (Doctoral dissertation, Instituto Superior de Economia e Gestão).
- [29] Yates, L. A., Aandahl, Z., Richards, S. A., & Brook, B. W. (2023). Cross validation for model selection: a review with examples from ecology. *Ecological Monographs*, 93(1), e1557.
- [30] Dhyani, A., & Jain, T. (2022). Minimal-norm state-feedback globally non-overshooting/undershooting tracking control of multivariable systems. *IFAC Journal of Systems and Control*, 22, 100212.