International Journal of Emerging Research in Engineering and Technology
Pearl Blue Research Group | Volume 7 Issue 3 PP 17-25, 2026
ISSN: 3050-922X | https://doi.org/10.63282/3050-922X.IJERET-V7I3P103

Original Article

An Explainable Generative AI Framework for Insurance Claims Intelligence and Risk Assessment

Kousik Mallick
MTech in AI, Department of Artificial Intelligence, Reva University Bangalore.

Received On: 18/05/2026 ***Revised On: 12/06/2026*** ***Accepted On: 23/06/2026*** ***Published On: 04/07/2026***

Abstract *- Insurance claims are becoming more complex, necessitating more sophisticated, clear, and automated risk evaluation processes. This research proposes a framework for the analysis of insurance claims and human decision support that combines the use of ML, explainable AI (XAI), and generative AI. The Insurance Claims and Policy Data dataset has 53,503 records, and these have been preprocessed using data cleaning, feature engineering, label encoding, and feature scaling of the data in a public repository. Several baseline models, namely Extra Trees Classifier (ETC), Gradient Boosting Classifier (GBC) and Multi-Layer Perceptron (MLP), are built and fused using a Hybrid Soft Voting Ensemble. To assess the proposed framework, the following metrics are used: Accuracy, Balanced Accuracy, Precision, Recall, Weighted F1-Score, Cohen's Kappa, MCC and ROC-AUC. The experimental findings show that compared to the individual models, the hybrid ensemble model has the best classification accuracy of 100%. Moreover, SHAP-based explainability helped to reveal the most critical features to capturing insurance risk prediction and Generative AI facilitated what-if risk scenario simulation, real-time claims intelligence, settlement estimation and automated decision support. The envisioned framework offers accurate, robust, and interpretable solution for smart insurance claims handling and risk assessment.*

Keywords *- Insurance Claims Intelligence, Insurance Risk Assessment, Explainable Artificial Intelligence (XAI), Generative AI and Fraud Detection.*

1. Introduction

An agreement between an insurer and an insured serve to financially safeguard the insured against a variety of hazards. Insurance is essential for the safety of both people and companies, and it also helps keep the economy stable [1]. An occurrence that leads the insured to submit a claim is called a claim event. In order to cover all current policies as well as any future claims from active policies, insurance firms must have enough money. Insurers pay for bodily harm or fatalities as well as other types of anticipated losses including pain, suffering, and mental distress under mandatory motor third-party liability insurance [2]. Motor liability insurance claims involving bodily harm need extensive modelling due to the fact that these claims often constitute the bulk of claims provisions and have a substantial impact on insurers' financial reserves [3][4]. By giving actuaries a thorough grasp of the claims portfolio and assisting in risk management, the creation of unique claims models enhances the accuracy of analysis. Improving profitability and maintaining stable rates in the insurance sector need early identification of the risk factors that result in expensive claims [5]. Predictive modelling using categorization approaches is useful for predicting claims and providing decision-making help in claims management since the severity of reported claims might vary substantially.

The ability to detect has improved a lot thanks to deep learning and other advanced AI techniques. Deep learning algorithms are trained on vast amounts of data to look for subtle, complex patterns that may be signs of fraud [6]. These models incorporate other factors like demographics, claim history, and behavioral patterns along with prior claims data to facilitate accuracy and improve operational efficiency[7]. These techniques can easily adjust to new data environments and fraud strategies to create a strong and proactive approach to fraud detection. This in turn can enhance decision making and financial loss. Deep learning models are especially suited to applications like insurance claim estimation and fraud detection, which involve complex and high-dimensional data.

The growing volume and complexity of claims are making insurance risk assessment and claim processing challenging and difficult. Current solutions tend to have limited prediction accuracy, lack transparency, and are not easily connected to real-time decision-making, which reduces their utility to identifying high-risk claims and enabling a fair settlement of claims. Thus, a smarter solution is required to bring together the right prediction, explainability, and generative AI to enhance insurance risk evaluation, claims intelligence, and decision-making for AI automation. The following are this work's main contributions:

- Developed a hybrid Soft Voting Ensemble by integrating Extra Trees, Gradient Boosting, and MLP classifiers for multi-class insurance risk prediction.
- Performed comprehensive data preprocessing and domain-specific feature engineering to improve predictive performance.
- Integrated SHAP-based Explainable AI for more comprehensible analysis of model predictions and feature importance.
- In-built Generative AI to simulate what if risk scenarios, estimate settlements and provide intelligent decision support.

- Implemented real-time insurance claim intelligence framework to route insurance claims, analyze fraud, and monitor operations.
- Performed an extensive evaluation with various classification performance metrics to prove the effectiveness and robustness of the proposed framework.

The proposed structure is justified because of the necessity for a smart insurance claims system to accurately recognize risks and deliver clear and actionable decision support. The traditional methodologies focus mainly on classification; this work incorporates a Hybrid Soft Voting Ensemble, SHAP explainability and Generative AI what-if scenario simulation within a single framework. The framework also provides real-time claims intelligence, including risk assessment, settlement cost estimation, fraud-aware analysis, and claim routing, providing comprehensive and comprehensible solutions to modern insurance claims management. The structure of the paper is as follows. The related work is presented in Section II, the suggested approach is explained in Section III, the experimental findings and analysis are covered in Section IV, and the study is concluded in Section V.

2. Literature Review

Recent studies have demonstrated the effectiveness of AI and ML techniques for insurance risk assessment and claims analysis. M. G. Rao et al. (2026) suggested approach incorporates supervised and unsupervised ML techniques for risk assessment and categorization. The most significant characteristics are found using SHAP. Additionally, experimental baselines are established to compare the performance of tree-based models with those of DL architectures like RNN and LSTM. The accuracy rate on the hybrid dataset reached by the proposed model was 91.27% [8]. Similarly, L. Suraiya et al. (2025) model is found to be applied effectively and at scale when integrating TensorFlow with Keras, allowing for the real-time processing of massive amounts of raw data. Experimental findings show a significant increase in prediction ability (97% accuracy) compared to more conventional methods; the key changes between the two are in precision, recall, and F1-score, which may account for the discrepancy in accuracy measures. Transparency and trust in decision-making are enhanced by combining SHAP with model interpretability [9].

L. P. P and G. Ramkumar (2025) suggested method uses Ridge Regression and Logistic Regression on historical claim datasets to successfully classify the claims' outcomes. At 96.88% vs 91.70, Ridge Regression is more accurate than Logistic Regression. The results show that insurance companies may benefit from machine learning's effect on risk management techniques, operational costs, and the accuracy of claim projections [10]. In another study, B. Al-Attar *et al*. (2025) presented an AI-powered system that can accurately identify, categorise, and evaluate the likelihood of insurance fraud claims. The system examines complicated fraud cases including health, property damage, automobile, and disability claims by combining supervised ML with sophisticated legal NLP methods. It does so by leveraging LegalBERT embeddings and structured claim information. The model has achieved classification accuracy of more than 94.6% on real-world annotated litigation data sets repeatedly, and F1 scores as high as 0.94 and AUC scores up to 0.94 [11].

Furthermore, A. Khan et al. (2024) explored predictive modeling, and many more results can be achieved with a variety of ML techniques. Use methods such as LR, RF, SVM and KNN to enhance accuracy of predictions. The study utilizes an insurance claims data set encompassing information on policyholders' characteristics, driving behavior, and prior claims. LR, RF, SVM, and KNN training and validation are supported by this dataset [12]. Finally, K. P. Rijanto et al. (2023) several ML models are examined using tenfold cross-validation, including the SVM, RF, Balanced Random Forest (BRF), and XGBoost classifier. Several classifiers are combined using the majority voting technique to improve prediction performance in the last phase. The SVM model, which uses the SMOTE oversampling approach, has an F1score of more than 98%. As a consequence, all models provide good results [13].

Although existing studies have demonstrated the effectiveness of ML, DL, ensemble learning, and explainable AI for insurance risk prediction, most focus primarily on classification accuracy or fraud detection. Limited research integrates hybrid ensemble learning, SHAP-based explainability, and Generative AI-driven what-if scenario simulation within a single framework. Moreover, existing approaches provide limited support for real-time claims intelligence, settlement estimation, and automated decision support. These shortcomings emphasise the need of an all-encompassing, interpretable AI system that can provide smart claims administration and accurate insurance risk assessment.

3. Methodology

The proposed methodology involves collecting and preprocessing the Insurance Claims and Policy Data dataset, followed by feature engineering, label encoding, and feature scaling using RobustScaler. The processed data is divided into testing and training sets, and the models Extra Trees, Gradient Boosting and MLP are trained and combined using the Hybrid Soft Voting Ensemble for insurance risk classification. The model is tested with well-established classification metrics and SHAP offers explainable predictions. Lastly, Generative AI applications include what-if risk scenario simulation, real-time claims intelligence, settlement estimation and decision support. The whole proposed methodology is illustrating in Fig. 1.

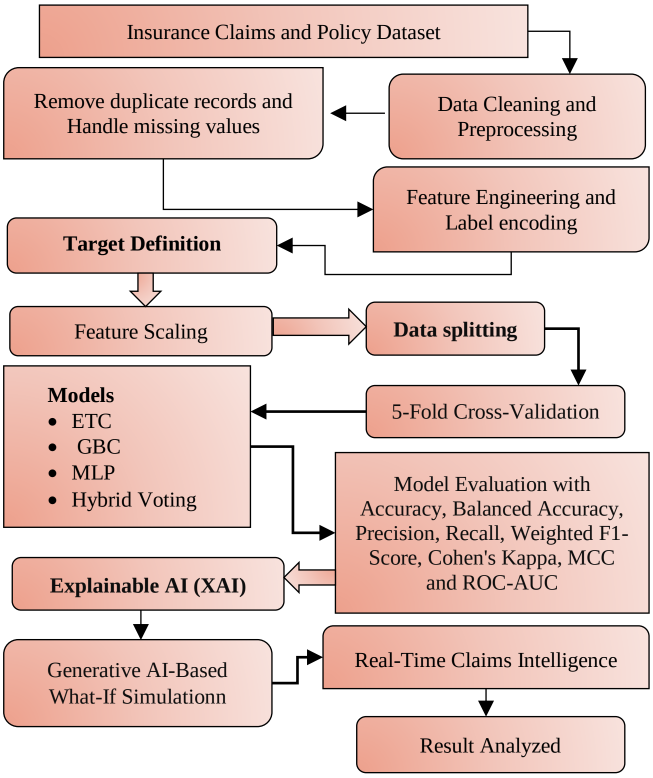

Fig 1: Flowchart of the Proposed Framework for Insurance Claims Intelligence and Risk Assessment

The following steps of proposed methodology are briefly discussing in below:

3.1. Data Collection

The proposed framework uses the Insurance Claims and Policy Data dataset [14] that is publicly available on Kaggle. Dataset contains 53,503 records and 30 attributes representing insurance claims (Customer Demographics, Policy, Claim, Credit Score, Driving Record, Coverage, and Risk Profile, which is the target class for prediction). The broadness of Dataset 2 allows for insurance risk assessment and explainable machine learning analysis.

The distribution of the four classes of insurance risk for the dataset is shown in Fig. 2. The Risk 3 class has the highest proportion with 31.1% (16,647 samples), followed by Risk 1 with 28.8% (15,393 samples), Risk 0 with 21.3% (11,405 samples), and Risk 2 with 18.8% (10,058 samples). A well-balanced class distribution is provided in the relatively well-balanced classes which helps to build a good multi-class insurance risk classification model and reduces the impact of class imbalance.

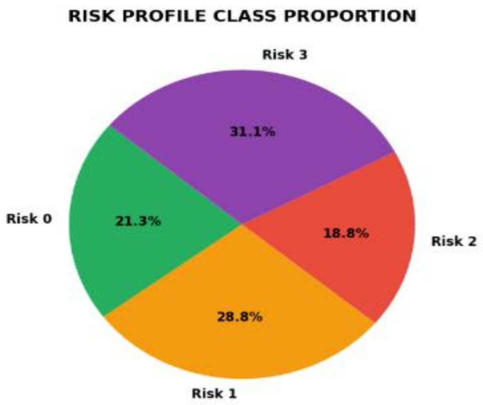

Fig 2: Risk Profile Class Distribution

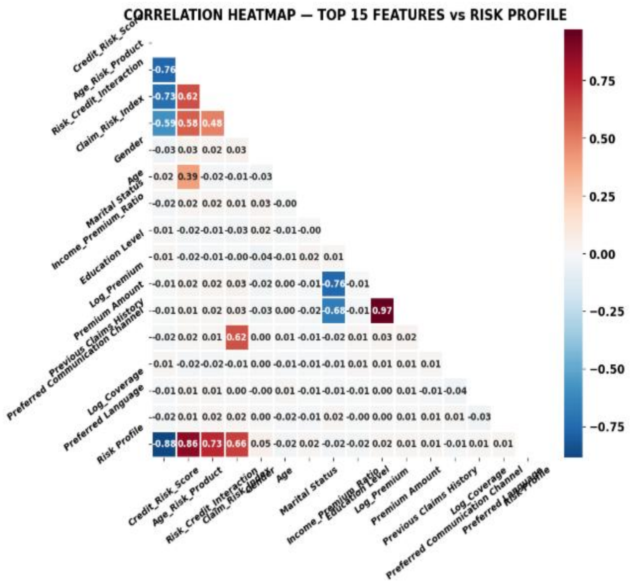

Fig 3: Correlation Heatmap of the Top Features for Insurance Risk Classification

The correlation heatmap of the top features vs. the Risk Profile target variable is shown in Fig. 3. Credit Risk Score (-0.88), Age_Risk_Product (0.86), Risk_Credit_Interaction (0.73), and Claim Risk Index (0.66) are the most correlated features with the target class, and thus, are most important for insurance risk prediction. Furthermore, Premium Amount and Premium_Coverage_Ratio are correlated positively (0.97) while Age_Risk_Product and Risk_Credit_Interaction are correlated negatively (-0.76).

3.2. Data Cleaning and Preprocessing

The collected datasets are preprocessed to improve data quality and ensure reliable model training. Records in both sets are duplicated and then deleted to avoid duplication. Numerical values that are missing are filled in with the median of the respective feature, and categorical values that are missing are filled in with the mode to ensure data consistency. The Dataset is preprocessed, with 53,503 records and 30 features respectively, in accordance with the subsequent feature engineering and model development.

3.3. Feature Engineering and Label Encoding

The main insurance data is used for feature engineering to improve predictive capabilities. Premium Coverage Ratio, Claim Risk Index, Credit Risk Score, Net Exposure, Claims per Premium, Deductible Ratio, Income Premium Ratio, Risk-Credit Interaction, Age-Risk Product, Log Income, Log Premium and Log Coverage are all removed and 12 domain specific features are generated. Label Encoding is used for the categorical attributes and only numerical features are considered for the model building process. This gave the processed dataset 40 numerical features and a total of 53,503 rows which provide a very complete set of features for predicting insurance risks.

3.4. Target Definition

The Risk Profile attribute is chosen as target variable for the insurance risk classification task with multiple classes, with the Customer ID column not taken into account in the model training. A feature matrix is created using a total of 38 input features. The distribution of the target variable (insurance risk level) is complete: 11,405 samples (21.3%) in Risk Level 0, 15,393 samples (28.8%) in Risk Level 1, 10,058 samples (18.8%) in Risk Level 2, and 16,647 samples (31.1%) in Risk Level 3, which provided a comprehensive distribution for insurance risk assessment.

3.5. Data Scaling and Train–Test Split

The input features are normalized using the RobustScaler to minimize the influence of outliers and improve model performance. The processed data is then split into 80% training (42,802 samples) and 20% testing (10,701 samples) datasets with StratifiedSampling to maintain the class distribution. In addition, 5-fold Stratified Cross-Validation has been used to perform robust and unbiased model training evaluation.

3.6. Machine Learning Models

To evaluate the proposed insurance risk assessment framework, three classification models Extra Trees Classifier (ETC), GBC, and MLP are implemented as baseline models, while a Hybrid Soft Voting Ensemble integrating all three classifiers is developed as the proposed model to enhance prediction accuracy, robustness, and generalization performance.

3.6.1. Extra Tree Classifier

The first baseline model used for insurance risk classification is the Extra Trees Classifier (ETC). ETC is an EnsembleLearning algorithm based on multiple DT generated by randomizing the data, and by combining the results of those trees by MajorityVoting, it achieves better classification accuracy and also significantly reduces the overfitting. The model is configured with 10 decision trees (n_estimators = 10), no restriction on tree depth (max_depth = None), minimum samples split of 3, minimum samples leaf of 1, square-root feature selection (max_features = 'sqrt'), bootstrap sampling, and balanced class weights. Moreover, Out-of-Bag (OOB) scoring is turned on, and random_state = 42 to ensure the results are reproducible.

3.6.2. Gradient Boosting Classifier

The second baseline model is Gradient Boosting Classifier (GBC) which is utilized to improve the prediction performance by decreasing the errors made by the previous trees in sequence. Classifier: 10 boosting stages (n_estimators = 10), learning rate = 0.1, maximum tree depth = 3, subsampling = 0.8, minimum samples split = 6, minimum samples leaf = 3 and square-root feature selection. Random_state = 42 for reproducibility, and the model parameters have been optimized for better generalization and avoiding overfitting.

3.6.3. MLP Neural Network Classifier

This decision is made because the Multi-Layer Perceptron (MLP) classifier is considered the most representative deep learning model for handling complex nonlinear relationships between insurance features. One hidden layer with 32 neurons (hidden_layer_sizes = (32,)) and tanh activation function are used, along with the Stochastic Gradient Descent (SGD) optimizer. L2 regularization (alpha = 10.0), constant learning rate (0.05), early stopping and 25% validation split are used to improve the generalization. The network is trained for a maximum of 10 iterations, and random_state = 42 is used for consistent results during training.

3.6.4. Hybrid Model- Soft Voting Ensemble

The proposed framework uses a Soft Voting Ensemble strategy to incorporate the advantages of Extra Trees Classifier, Gradient Boosting Classifier, and MLP Neural Network. The ensemble predicts the final insurance risk class by averaging the class probability outputs of the three base learners and selecting the class with the highest combined probability. For the ensemble, the Extra Trees model had 10 trees, bootstrap sampling and minimum leaf size of 2 with balanced class weights. The Gradient Boosting model used 10 estimators, learning rate 0.1, max depth 1 and subsampling 80%. The MLP is made up of three hidden layers (256, 128, and 64 neurons), ReLU activation, L2 regularization with alpha = 0.001, 100 training iterations and early stopping with 10% validation split. The soft voting strategy successfully fused the complementary advantage of various classifiers and is able to build a strong and highly accurate insurance risk prediction model.

3.7. Evaluation Metrics

Test Accuracy, Balanced Accuracy, Precision, Recall, Cohen's Kappa, Matthews Correlation Coefficient (MCC), and ROC-AUC are used to assess the effectiveness of the suggested insurance risk categorization framework. Test Accuracy measures the overall classification performance, while Balanced Accuracy accounts for class-wise prediction performance. Precision, Recall and the Weighted F1-Score measure the accuracy and completeness of the predictions. Cohen's Kappa and MCC are metrics used to measure agreement between and reliability of the classification results, especially in multi-class problems. The following matrix are calculated in Equations (1) to (6).

$$Accuracy = \frac{TP+TN}{TP+Fp+TN+FN} \quad (1)$$

$$Precision = \frac{TP}{TP+FP} \quad (2)$$

$$Recall = \frac{\mathrm{TP}}{TP+FN} \quad (3)$$

$$F1-score = 2 \times \frac{Precision \times Recall}{Precision + Recall} \quad (4)$$

$$\text{Balanced Accuracy} = \frac{1}{C}\sum_{i=1}^{C} Recall_i \quad (5)$$

$$\text{MCC} = \frac{TP \times TN - FP \times FN}{\sqrt{(TP+FP)(TP+FN)(TN+FP)(TN+FN)}} \quad (6)$$

The **ROC-AUC** indicates the quality of the model's separation of different risk classes. The suggested model's accuracy, robustness, and dependability are evaluated by these measures taken as a whole.

4. Result Analysis and Discussion

This experimental setup made use of a 64 GB RAM "Intel Core" i7 CPU, "Windows 10" OS, "Jupyter Notebook" platform for development, and the "Python" programming language. The performance comparison of the above ML models Extra Trees, Gradient Boosting, MLP Neural Network and the proposed Hybrid Voting model is shown in Table I with different classification metrics. The evaluation comprises of Test Accuracy, Balanced Accuracy, Weighted F1-Score, Precision, Recall, Cohen'sKappa, MCC, ROC-AUC, cross validation accuracy and training accuracy. When compared with other models, the proposed Hybrid Voting model records the highest performance in all the evaluation metrics with a cross validation accuracy of 100% and standard deviation of 0.0, which shows good robustness and generalization ability. The Extra Trees classifer also performs well with a test accuracy of 99.87%, while the MLP Neural Network and Gradient Boosting models give an accuracy of 97.37% and 93.99% respectively. The results shown that in every instance, the suggested hybrid ensemble model outperforms the individual classifiers when it comes to insurance risk categorization.

Table 1: Performance Comparison of Machine Learning Models for Insurance Risk Classification

Model	Test Accuracy	Balanced Accuracy	Weighted F1-Score	Precision	Recall	Cohen's Kappa	MCC	ROC-AUC	CV Accuracy (Mean)	CV Accuracy (Std)	CV F1-Score (Mean)	Training Accuracy
Extra Trees	0.9987	0.9984	0.9987	0.9987	0.9987	0.9982	0.9982	1.0000	0.9995	0.0003	0.9995	1.0000
Gradient Boosting	0.9399	0.9206	0.9363	0.9465	0.9399	0.9180	0.9218	0.9992	0.9448	0.0065	0.9417	0.9450
MLP Neural Network	0.9737	0.9651	0.9732	0.9756	0.9737	0.9643	0.9652	0.9999	0.9790	0.0060	0.9786	0.9790
Hybrid Voting (Proposed)	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	0.0000	1.0000	1.0000

Fig. 4 presents the confusion matrix of the proposed Hybrid Voting model for insurance risk classification. The model correctly classified 2,281 Risk 0, 3,079 Risk 1, 2,012 Risk 2, and 3,329 Risk 3 samples, with no misclassifications in any class. The optimal diagonal values achieve the accuracy of 100% in classification and this shows the effectiveness and reliability of the proposed model.

The findings of the proposed Hybrid Voting model for explainability study are presented in Fig. 5. The feature importance plot shows that the biggest two features contributing to insurance risk prediction are Credit Risk Score (0.0997) and Age_Risk_Product (0.0989), followed by Claim Risk Index (0.0724) and Risk_Credit_Interaction (0.0641). SHAP beeswarm plots also show the impact of the individual features' values on the prediction of Risk 0 (Low) or Risk 3 (High) classes. High SHAP values are associated with higher probabilities of a specific risk class while low values lead to reduced contributions. In general, the SHAP method offers clear insights into the weighting process of the insurance-claim classification model, highlighting the most important factors determining the classification.

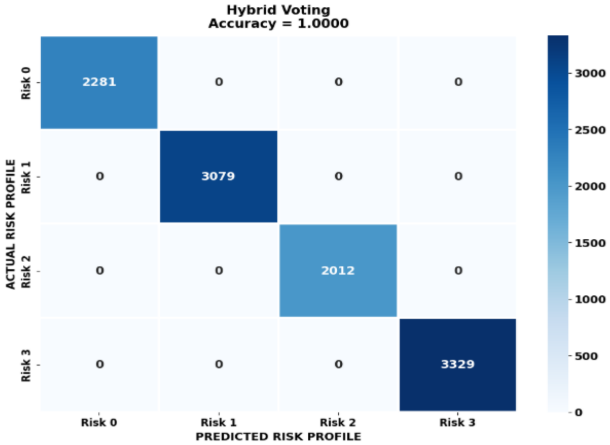

Fig 4: Confusion Matrix of the Proposed Hybrid Voting Model for Insurance Risk Classification

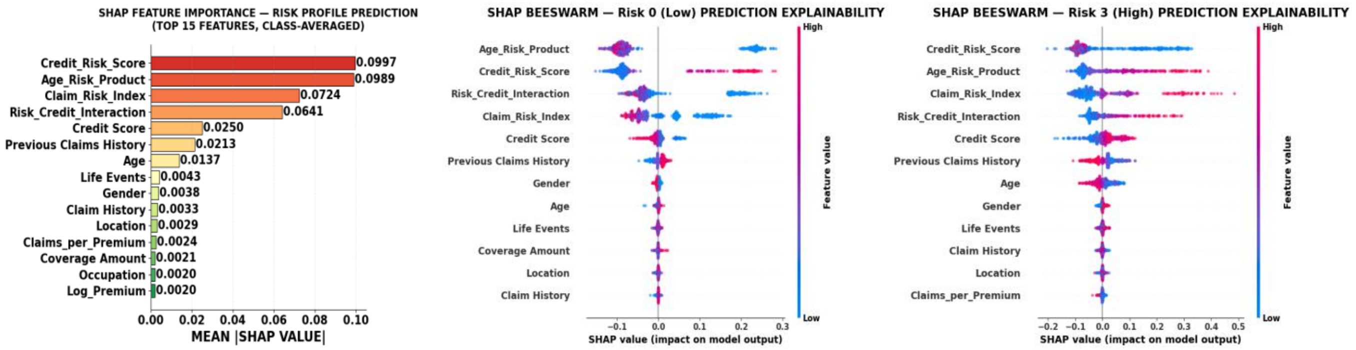

Fig 5: SHAP-Based Feature Importance and Explainability of the Proposed Hybrid Voting Model

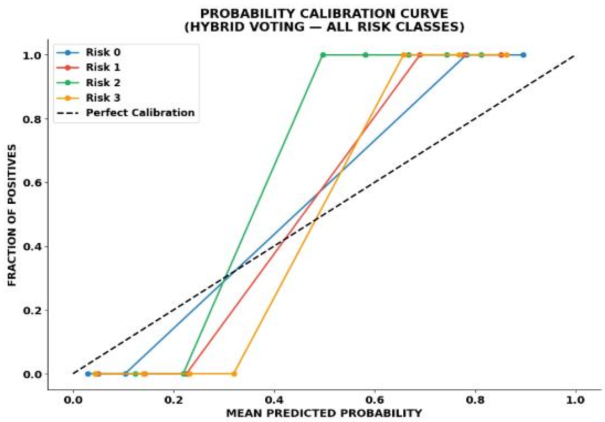

Fig 6: Probability Calibration Curve of the Proposed Hybrid Voting Model

Fig. 6 shows the Probability Calibration curves for all four categories of insurance risk (Risk 0–Risk 3) of the proposed Hybrid Voting model. The predicted probabilities are closely aligned to the perfect calibration line, which suggests that the model generates well calibrated and reliable prediction of probabilities. The R-squared values for all plots are close to 1.0 for positive predictions, indicating strong confidence and consistency in all risk classes. The outcomes demonstrate that the suggested approach may provide precise probability forecasts in addition to superior categorization results.

Table II presents the results of the Generative AI-based What-If Risk Scenario Simulation for different insurance claim scenarios. The baseline case is classified as Moderate Risk (Risk 1) with an estimated settlement of $17,200. Improving the customer's credit score from 640 to 760 reduces the predicted risk to Low (Risk 0) and decreases the estimated settlement to $5,200, resulting in an Auto-Approve recommendation. In contrast, the grantor considered the scenario of a Young High-Risk Driver with prior claims of 1 to 4 and the settlement goes up to $71,200 (Very High Risk – Risk 3) and requires Escalate + Audit actions. A Life Event – Job Loss scenario is considered to be High Risk (Risk 2) with a settlement estimate of $41,200, and a Senior Review recommendation. The findings illustrate the potential of the proposed framework to replicate scenarios of changes in risk and aid intelligent insurance decision making.

Table 2: Generative AI-Based What-If Risk Scenario Simulation Results

Scenario	Predicted Risk Level	Credit Score	Prior Claims	Estimated Settlement ($)	Recommended Action
Baseline	Moderate (Risk 1)	640	1	17,200	Standard Review
Credit Score Improved	Low (Risk 0)	760	1	5,200	Auto-Approve
Prior Claims Spike	Very High (Risk 3)	640	4	71,200	Escalate + Audit
Life Event – Job Loss	High (Risk 2)	640	1	41,200	Senior Review
Young High-Risk Driver	Very High (Risk 3)	510	1	71,200	Escalate + Audit

The proposed framework has generated insurance claims intelligence in real time as displayed in Table III. The five claims compared are CLM-01014, identified as High Risk, with a risk probability of 0.5593 and estimated settlement of $72,891, which is subject to Senior Review. CLM-01011 and CLM-01012 would be categorized as Moderate Risk and would be reviewed with a Standard Review; CLM-01013 and CLM-01015 are categorized as Low Risk with a low risk probability (0.1836 and 0.1035) and are automatically approved. The results show that the framework can effectively evaluate the risks, predict the settlement amount and offer decision support for efficient insurance claim management, which is helpful for the companies.

Table 3: Real-Time Insurance Claims Intelligence Stream Analysis

Claim ID	Policy Type	Region	Risk Level	Risk Probability	Credit Score	Prior Claims	Claim Amount ($)	Estimated Settlement ($)	Recommended Action
CLM-01011	Liability	South	Moderate	0.4443	474	3	59,416	43,062	Standard Review
CLM-01012	Auto	Central	Moderate	0.3207	429	0	25,790	18,593	Standard Review
CLM-01013	Life	North	Low	0.1836	812	0	57,080	51,013	Auto-Approve
CLM-01014	Liability	South	High	0.5593	624	5	135,256	72,891	Senior Review
CLM-01015	Property	North	Low	0.1035	743	0	5,224	4,056	Auto-Approve

Table 4: Claims-AI Session Summary Report

Category	Metric	Value
Session Information	Stream Window	08:00 – 14:00
	Total Claims Processed	30
Adjuster Routing	Auto-Approved	2
	Standard Review	23
	Senior Review	3
	Escalated + Audit	2
Risk Intelligence	High-Risk Claims	5 (16.7%)
	Critical Risk Claims	0
	Fraud-Flagged Claims	2 (6.7%)
Financial Exposure	Total Claims Filed	$43,545
	Total Estimated Payout	$1,447,817
	Total Premium Pool	$77,443
	Loss Ratio	1869.5% (Critical)

Table IV shows the performance of the Claims-AI framework on the real-time insurance claims processing session. There are a total of 30 claims: 23 are sent to Standard Review, 3 to Senior Review, 2 are automatically approved, and 2 are sent to audit. The framework detected 5 high-risk claims (16.7%), 2 fraud-flagged claims (6.7%) and no critical-risk claims. A total estimated payout of $1,447,817 and a premium pool of $77,443, with a loss ratio of 1869.5%, the system's ability to monitor operational risk and financial exposure in real time can be seen in financial analysis.

4.1. Comparative Analysis

Table V displays the results of comparing the suggested Hybrid Voting model to other ML methods. This model have highest accuracy (1.0000) and F1-Score (1.0000) among all current and baseline techniques. The MLP Neural Network has an F1 score of 0.9786 and an accuracy of 0.9790, whereas Extra Trees is the best performing individual model with an F1 score of 0.9995 and an accuracy of 0.9990. The proposed hybrid ensemble shows better classification performance and more stable prediction performance than previous methods such as RF, XGBoost, LR, LASSO, and DT in insurance risk assessment.

Table 5: Comparative Performance Analysis of the Proposed Hybrid Voting Model and Existing Classification Methods

Models	Accuracy	F1-Score
Extra Trees	0.9995	0.9995
Gradient Boosting	0.9448	0.9417
MLP Neural Network	0.9790	0.9786
Hybrid Voting	1.0000	1.0000
Random Forest[15]	0.975	0.975
XGB[16]	0.7761	0.6856
LR[17]	0.871	0.931
LASSO[17]	0.977	0.986
DT[18]	0.9540	-

4.2. Limitations and Future Work

The framework proposed has good performance in insurance risk assessment, which contains some defects. The study is assessed based on a single dataset, available in the public, which is not likely to represent all the insurance environments in reality. Furthermore, the framework is geared towards structured claim information and doesn't include unstructured data like claim documents, images, or customer communications. Future directions include validation of the framework with multi-source data and real-world datasets, integrating multimodal data and large language models (LLMs), enhancing fraud detection using graph-based learning, and applying the framework in cloud-based and real-time insurance systems for large-scale and adaptive insurance decision-making support.

5. Conclusion and Future Work

Insurance claim systems that can be intelligent, accurate and transparent are increasingly required in the rapidly changing world of digital insurance. The present research suggested Explainable Generative AI Framework for Insurance claims Intelligence and Risk Assessment with Hybrid Soft Voting Ensemble model based on Extra Trees, Gradient Boosting and MLP model. The proposed framework achieved 100% accuracy in classification of all the data which was better than the models themselves. The enhancements of SHAP-based explainability, in addition to model transparency, focused on the most significant risk factors, and Generative AI brought what-if risk scenario simulation, claims intelligence in real time, settlement estimation and decision support. The results of the experiments back the proposed model of intelligent and efficient risk assessment and automated claims processing in the insurance industry which could offer a genuine, reliable and understandable solution.

References

[1] A. R. Aitha, “Explainable Agentic AI Framework for Automated Insurance Fraud Detection and Predictive Risk Intelligence,” *Int. J. Adv. Res. Comput. Sci. Technol.*, vol. 9, no. 03, May 2026, doi: 10.15662/IJARCST.2026.0903009.

[2] E. Brati, A. Braimllari, and A. Gjeçi, “Machine Learning Applications for Predicting High-Cost Claims Using Insurance Data,” *Data*, vol. 10, no. 6, Jun. 2025, doi: 10.3390/data10060090.

[3] K. Amistapuram and P. Pandey, “Generative AI for Claims Exceptions and Investigations: Enhancing Resolution Efficiency in Complex Insurance Processes,” *Lex localis - J. Local Self-Government*, vol. 23, no. S6, pp. 7320–7337, Oct. 2025, doi: 10.52152/y5j9v602.

[4] K. McDonnell, B. Sheehan, and F. Murphy, “Bridging transparency in insurance claims prediction: A comparative study of explainable AI and traditional linear models using vehicle telematics data,” *Technol. Forecast. Soc. Change*, vol. 223, p. 124418, Feb. 2026, doi: 10.1016/j.techfore.2025.124418.

[5] V. Rengaraj, Y. Lnu, Y. Patel, and S. Vinod, “Decoding-Enhanced Bidirectional Encoder Representations with Attention-Based Pooling Mechanism for Insurance Fraud Detections,” in *2026 3rd International Conference on Research Methodologies in Knowledge Management, Artificial Intelligence and Telecommunication Engineering (RMKMATE)*, 2026, pp. 1–9. doi: 10.1109/RMKMATE69073.2026.11518574.

[6] P. A. Kumar and S. Sountharrajan, “Insurance claims estimation and fraud detection with optimized deep learning techniques,” *Sci. Rep.*, vol. 15, no. 1, p. 27296, Jul. 2025, doi: 10.1038/s41598-025-12848-0.

[7] J. Mishra, S. Rai, B. Moharana, V. K. Singh, S. Dey, and T. Sarkar, “Revolutionizing Financial Fraud Detection in Healthcare Insurance Through Blockchain & AI-Driven Predictive Analysis,” in *2025 International Conference on Technology Enabled Economic Changes (InTech)*, IEEE, Feb. 2025, pp. 1516–1522. doi: 10.1109/InTech64186.2025.11198190.

[8] M. G. Rao, V. B. S, P. Rao H, M. Mascarenhas, A. P. Kumar, and S. Shanbhag, “A Data Fusion–Driven Framework for Predicting Insurance Risk Using Explainable Ensemble and Deep Learning Approaches,” in *2026 International Conference on Artificial Intelligence and Data Engineering (AIDE)*, IEEE, Feb. 2026, pp. 851–856. doi: 10.1109/AIDE69088.2026.11544448.

[9] L. Suraiya, B. P. C, F. R. Sumarna, J. R. Machireddy, A. A, and M. C. Janaki, “Machine Learning-Based Risk Prediction Framework for Insurance Fraud Detection and Claim Optimization,” in *2025 International Conference on Recent Innovation in Science Engineering and Technology (ICRISET)*, IEEE, Aug. 2025, pp. 1–6. doi: 10.1109/ICRISET64803.2025.11254689.

[10] L. P. P and G. Ramkumar, “Insurance Claim Prediction and Risk Assessment Using Machine Learning,” in *2025 1st International Conference on Emerging Innovation and Digital Technology (ICEIDT)*, IEEE, Nov. 2025, pp. 288–292. doi: 10.1109/ICEIDT66693.2025.11473722.

[11] B. Al-Attar *et al.*, “Supervised Machine Learning Framework for Detection, Classification, and Risk Assessment of Fraudulent Legal Claims in Insurance Litigation,” in *2025 3rd International Conference on Cyber Resilience (ICCR)*, IEEE, Jul. 2025, pp. 1–7. doi: 10.1109/ICCR67387.2025.11291936.

[12] A. Khan, K. Dhungana, S. Gyawali, and I. Sharma, “Forecasting Automobile Insurance Claims: Examining the Interplay of Policyholder Traits for Enhanced Predictive Insights,” in *2024 2nd International Conference on Intelligent Data Communication Technologies and Internet of Things (IDCIoT)*, IEEE, Jan. 2024, pp. 881–886. doi: 10.1109/IDCIoT59759.2024.10467373.

[13] K. P. Rijanto, D. A. Putri, P. Calvin, I. S. Edbert, and D. Suhartono, “Machine Learning Algorithms with Oversampling Methods for Automobile Insurance Fraud Detection,” in *2023 IEEE 21st Student Conference on Research and Development (SCOReD)*, IEEE, Dec. 2023,

pp. 138–143. doi: 10.1109/SCOReD60679.2023.10563582.

[14] S. Raval, “Insurance Claims and Policy Data,” Kaggle. [Online]. Available: https://www.kaggle.com/datasets/ravalsmit/insurance-claims-and-policy-data

[15] A. Polamarasetti *et al.*, “Enhancing Risk Assessment in Auto Insurance with Data-Driven Insights using Machine Learning,” *J. Contemp. Educ. Theory Artif. Intell.*, vol. 3, 2023, doi: 10.71010/2023/JCETAI-e104-1.

[16] S. K. Majhi, S. Bhatachharya, R. Pradhan, and S. Biswal, “Fuzzy clustering using salp swarm algorithm for automobile insurance fraud detection,” *J. Intell. Fuzzy Syst.*, vol. 36, no. 3, pp. 2333–2344, 2019.

[17] H. Moon, Y. Pu, C. Ceglia, and others, “A predictive modeling for detecting fraudulent automobile insurance claims,” *Theor. Econ. Lett.*, vol. 9, no. 06, p. 1886, 2019.

[18] A. S. Elbhrawy, M. A. Belal, and M. S. Hassanein, “AIRA-ML: Auto Insurance Risk Assessment-Machine Learning Model using Resampling Methods,” *Int. J. Adv. Comput. Sci. Appl.*, vol. 14, no. 9, pp. 633–641, 2023, doi: 10.14569/IJACSA.2023.0140966.