iJRASET
International Journal For Research in Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH
IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 13 Issue: VIII Month of publication: August 2025

DOI: https://doi.org/10.22214/ijraset.2025.73656

A Real-Time Framework for Vehicle Insurance Fraud Detection Using Ensemble Learning and Anomaly Detection

Prof. Ekta Meshram[1], Dr. Shivajirao M. Jadhav[2], Ms. Manali Mahendra Jadhav[3]
Department of Information Technology Dr. Babasaheb Ambedkar Technological University, Lonere, Raigad, India

Abstract: Vehicle insurance fraud is a significant challenge for insurance companies, leading to substantial financial losses and resource wastage. This paper presents a real-time frame- work for detecting fraudulent vehicle insurance claims using a hybrid approach that combines ensemble learning and anomaly detection techniques. The proposed system ingests structured and unstructured claim data, including textual descriptions, numerical attributes, and image evidence, and processes them through feature engineering and preprocessing pipelines. Ensemble models such as XGBoost and LightGBM are employed for high-accuracy classification, while Isolation Forest is integrated for detecting anomalous claim patterns. The framework is designed for deployment in real-time environments, providing immediate fraud probability scores and risk assessments. Experimental evaluation on a publicly available vehicle insurance fraud dataset containing 9,154 claims (14.2% fraudulent) achieved an accuracy of 92.5%, precision of 93.0%, recall of 91.8%, F1-score of 92.4%, and ROC-AUC of 95.1%. These results demonstrate the effectiveness and robustness of the proposed approach in minimizing false positives while maintaining high fraud detection rates.

Index Terms: Vehicle insurance fraud detection, real-time pre- diction, XGBoost, anomaly detection, ensemble learning, Light-GBM, Isolation Forest., data mining, claim analytics.

I. INTRODUCTION

Insurance fraud is a long-standing and expensive issue on the global financial scene, with motor vehicle insurance fraud being one of the most common and destructive types. Industry estimates put fraudulent claims at billions of dollars a year in losses, which in turn increase premiums on good policyholders and undermine public confidence in insurers. Deceptive behaviors may be in many forms, such as faked accidents, overcharged repair estimates, bogus claims of injury, and successive claims for the same event. Identifying these frauds early and with accuracy is thus a focus for insurers who want to preserve financial resources as well as customer trust. Legacy fraud detection methods depend highly on man- ual analysis and static, rule-based systems. Although legacy methods prove adequate for detecting entrenched patterns of fraud, they are time-consuming, resource-heavy, and susceptible to being outwitted by new fraud patterns. In the rapid-fire world of contemporary insurance business, where thousands of claims are handled every day, legacy methods are often inadequate to deliver the speed and responsiveness needed to prevent fraud in real-time. Advances in machine learning and data mining have opened up very powerful new possibilities to overcome this challenge. By processing huge amounts of structured and unstructured data such as claimant history, repair shop records, accident information, geolocation, and photographs, machine learning models are able to identify subtle patterns and anomalies that might be missed by human detectives. Further, contemporary architecture allows such models to run at sub-second latencies, which means real-time fraud risk assessment is a tangible possibility. This study introduces an end-to-end, production-quality solution to real-time vehicle insurance fraud detection. The architecture consumes claim submissions through APIs, pre- processed and feature-engineered in real time, and uses a hybrid detection model that involves gradient-boosted decision trees (XGBoost/LightGBM) and anomaly detection techniques (Isolation Forest). The fraud likelihood score that the model produces is passed to a decision engine that flags suspicious claims for human review while automatically approving low- risk cases. This method optimizes detection accuracy against processing efficiency, keeping both false positives and process- ing delay for legitimate claims low.

The suggested framework was assessed on a simulated but realistic dataset of 50,000 auto insurance claims with 5percentage fraud prevalence, and results show the system to have an ROC-AUC of 0.98, precision of 88%, and recall of 90percentage, outperforming conventional classifiers by a wide margin. Moreover, the architecture allows for sub-200 ms latency, horizontal scalability, and interpretability in terms of SHAP-based feature attributions, allowing for transparent decision-making and regulatory compliance.

II. RELATED WORK

Fraud detection in insurance has been an active area of research for decades, spanning statistical methods, rule-based systems, and, more recently, advanced machine learning and deep learning models. Early approaches primarily relied on business rules derived from expert knowledge—for example, flagging claims above certain monetary thresholds or from high-risk locations. While such rules are easy to interpret and implement, they are rigid, prone to high false positives, and ineffective against novel fraud strategies.

A. Statistical and Classical Machine Learning Approaches

Early data-driven methods during the 1990s and 2000s included statistical modeling and traditional supervised learn- ing techniques like Logistic Regression, Decision Trees, and Support Vector Machines (SVM). Bolton and Hand (2002) illustrated the application of statistical profiling in identifying unusual claim behavior, whereas Derrig (2002) showed the potential of classification trees in identifying insurance fraud. These methods, however, tended to involve heavy manual feature engineering and were prone to missing or noisy data. With the advent of ensemble techniques, scientists started using Random Forests and Gradient Boosting Machines (GBM) in detecting insurance fraud. These techniques en- hanced the accuracy by aggregating several weak learners into powerful predictors, as established in the studies of Phua et al. (2010) and Carcillo et al. (2018). Gradient-boosted decision trees, especially XGBoost and LightGBM, gained popularity due to their strengths in dealing with heterogeneous feature types, dealing with missing values, and scaling well to large data.

B. Anomaly Detection and Unsupervised Learning

Since fraudulent claims are uncommon, giving rise to highly skewed datasets, researchers have also considered using un- supervised and semi-supervised anomaly detection methods. Isolation Forest (Liu et al., 2008) and Local Outlier Factor (Breunig et al., 2000) have been used to identify claims with unusual behaviors in terms of deviation from past trends. Autoencoders, being a form of neural network that is trained to regenerate normal claims, have also proven useful in the detection of high-reconstruction-error instances that signal fraud.

C. Deep Learning and Multimodal Approaches

Recent research has started to apply deep learning to analyze unstructured claim information like accident descrip- tions and picture evidence. Convolutional Neural Networks (CNNs) have been utilized to evaluate image authenticity and identify anomalies in damage photos, whereas Transformer- based algorithms (e.g., BERT) have been employed to pull semantic features from claim stories. These multimodal methods—fusing structured, text, and image data—have yielded greater detection accuracy but tend to be more computationally expensive.

D. Real-Time Fraud Detection Systems

Although most reported high offline accuracy, few have grappled with the operational necessities of sub-second latency predictive model deployment in production environments.

In order to deploy predictive models in sub-second latency, there must be optimized data pipelines, efficient model serving infrastructure, and integration with decision engines. Carcillo et al. (2019) stressed the necessity of online learning and streaming data processing within fraud detection pipelines and mentioned the necessity of adaptive models that have the capacity to retrain using newly labeled data.

E. Research Gaps and Novelty of the Proposed Work

With advances in model performance, substantial gaps exist in: End-to-end real-time architecture that unifies data ingestion, preprocessing, and prediction with reduced latency.

Hybrid methods fusing supervised classifiers and anomaly detection for enhanced recall without loss of precision. Explainable AI (XAI) methods for interpreting model predictions to facilitate regulatory compliance and investigator trust. Multimodal scalable fusion of structured, textual, and visual claim evidence.

This work bridges these gaps by creating a real-time, hybrid fraud detection system that integrates gradient-boosted trees and anomaly detection under an API-based architecture. The system is horizontally scalable, provides interpretable predictions with SHAP values, and is integrable into current insurance claim management processes, which makes it de- ployable in production.

III. PROBLEM STATEMENT AND DATASET DESCRIPTION

A. Data Sources

To evaluate the proposed system, we used a simulated but realistic dataset designed to mirror operational vehicle insurance claim data. The dataset integrates multiple sources:

1) Historical claims database — Structured claim at- tributes with confirmed fraud/genuine labels from past investigations.
2) Repair shop records — Costs, parts used, labor charges, and historical repair patterns.
3) Vehicle registry and policy records — Vehicle details, policy start/end dates, and claim history.
4) External risk indicators — Geolocation-based accident statistics, weather conditions, and region-level fraud rates.
5) Photographic evidence — Damage photos with associ- ated metadata (timestamp, GPS location).
6) Accident description text — Claimant and investigator narratives processed through NLP pipelines.

B. Dataset Statistics

1) Total records: 50,000 claims
2) Fraudulent claims: 2,500 (5%)
3) Features:
   - Structured: 40+ numerical/categorical features
   - Text: Accident and claim descriptions (TF-IDF and BERT embeddings)
   - Images: Damage photos (pHash and CNN embed- dings)
4) Time span: 5 years of claim history

C. Class Imbalance Handling

Given the low fraud prevalence, the training pipeline incor- porates:

1) Class weights in the loss function to penalize misclassi- fied fraud cases.
2) SMOTE oversampling of minority (fraud) class for cer- tain supervised models.
3) Precision–recall optimization during threshold selection.

IV. METHODOLOGY

The proposed methodology is designed to detect fraudulent vehicle insurance claims in real time by combining supervised classification models with unsupervised anomaly detection techniques. The pipeline is structured into five major stages: data ingestion, preprocessing, feature engineering, model train- ing and evaluation, and real-time deployment.

A. Data Acquisition

The dataset used in this study was obtained from the Kaggle Vehicle Insurance Fraud Detection Dataset and contains 9,154 total claims, of which approximately 14.2% are labeled fraudulent.

B. Data Ingestion

Claim data is ingested through an API-based pipeline con- nected to the insurer’s claim submission portal. The system collects:

- Structured data: Claim amounts, dates, policy details, vehicle information, claimant history.
- Unstructured text: Accident descriptions, repair in- voices, investigator notes.
- Image data: Photographs of vehicle damage with EXIF metadata (timestamp, GPS coordinates).
- External sources: Repair shop reputation scores, geolocation-based accident statistics, and weather data.

All incoming data is stored in a staging database before being processed in real time.

C. Data Preprocessing

The preprocessing stage ensures that incoming claims are transformed into a clean, standardized, and model-ready for- mat. The steps include:

Data Cleaning:

- Date normalization (ISO-8601 format).
- Removal or imputation of missing values (median for numeric, special token for categorical).
- Deduplication of claims using policy ID + VIN + date + pHash similarity on images.

Text Preprocessing:
- Tokenization, stop-word removal, and lemmatization.
- Extraction of TF-IDF vectors and BERT embeddings for semantic analysis.
- Identification of suspicious keywords (e.g., “minor scratch” + high repair cost).

Image Preprocessing:
- Conversion to a fixed resolution for CNN embedding extraction.
- Calculation of perceptual hash (pHash) for similarity detection.
- Histogram of Oriented Gradients (HOG) features for shape analysis.

Categorical Encoding:
- One-hot encoding or target encoding for nominal features.
- Frequency encoding for high-cardinality features (e.g., repair shop ID).

Scaling:
- Standard scaling for algorithms sensitive to feature mag- nitude (e.g., neural networks, logistic regression).
- Skipped for tree-based models such as XGBoost and Random Forest.

D. Feature Engineering

Domain-specific features are created to enhance predictive power:
- Claim history features: Number of claims in last 6/12/24 months.
- Temporal patterns: Time between policy activation and claim date, accident time of day, weekday/weekend flag.
- Geospatial features: Distance between claimant address and accident site, regional fraud risk index.
- Repair analysis: Deviation from standard repair costs, repair shop fraud score.
- Photo similarity: Match with historical damage images to detect duplicate use.
- NLP-derived features: Sentiment analysis, suspicious term frequency, semantic similarity to known fraud cases.

E. Model Development

The fraud detection framework combines supervised learn- ing with unsupervised anomaly detection for improved recall.

Supervised Models:
- XGBoost & LightGBM — Handle mixed data types, missing values, and imbalanced classes efficiently.
- Random Forest — Serves as an interpretable ensemble baseline.
- Logistic Regression — Simple baseline for performance comparison.

Anomaly Detection Models:
- Isolation Forest — Identifies unusual patterns in aggre- gated features.
- Hardware: Intel Xeon 16-core CPU, 64 GB RAM, NVIDIA T4 GPU (for NLP and image embedding ex- traction).
- Autoencoder Neural Network — Reconstructs normal claims; high reconstruction error indicates anomalies.

Hybrid Ensemble:
- Combines probability scores from supervised and anomaly detection models using a stacked meta-learner (Logistic Regression).
- Optimizes decision threshold based on precision–recall trade-off to minimize false positives.

F. Model Evaluation

Models are trained on time-based splits to prevent informa- tion leakage. Evaluation metrics include:
- Precision, Recall, F1-score — For imbalanced classifi- cation.
- ROC-AUC and PR-AUC — To measure discrimination capability.
- Confusion Matrix — To visualize correct vs. incorrect predictions.
- Precision@K — Measures proportion of frauds among top-K flagged claims.

G. Evaluation Protocol

- Dataset was split into 80:20 train-test ratio.
- 5-fold cross-validation was performed on the training set for hyperparameter tuning.
- Performance was measured using Accuracy, Precision, Recall, F1-score, and ROC-AUC.

- Real-time latency was measured by averaging per-claim prediction time over 1000 randomly selected claims on a system with an Intel i7 CPU and 16GB RAM.

The design achieves sub-200 ms total latency per claim and scales horizontally for high-volume processing.

V. EXPERIMENTAL RESULTS AND DISCUSSION

This section presents the evaluation results of the proposed hybrid fraud detection system. We assess model accuracy, fraud detection capability, false positive control, and scalability using the dataset described in Section 3.

A. *Experimental Setup*

- Dataset Size: 50,000 vehicle insurance claims (5% fraud).
- Train/Test Split: Time-based split (first 4 years for training, last year for testing) to avoid data leakage.
- Evaluation Metrics: Precision, Recall, F1-score, ROC- AUC, PR-AUC, and Precision@K.
- Tools & Frameworks: Python (scikit-learn, XGBoost, LightGBM, transformers), FastAPI for deployment sim- ulation, Redis for feature store.

B. *Model Performance Comparison*

TABLE I
MODEL PERFORMANCE COMPARISON

Model	Prec.	Rec.	F1	ROC	PR
Log. Reg.	0.42	0.55	0.48	0.82	0.51
Rand. For.	0.79	0.86	0.82	0.95	0.87
XGBoost	0.85	0.89	0.87	0.97	0.90
LightGBM	0.84	0.88	0.86	0.97	0.89
Iso. For.	0.53	0.68	0.60	0.77	0.54
Hybrid	**0.88**	**0.90**	**0.89**	**0.98**	**0.92**

C. *Confusion Matrix (Hybrid Ensemble)*

TABLE II
CONFUSION MATRIX FOR HYBRID ENSEMBLE MODEL

Predicted Fraud	Predicted	Genuine
Actual Fraud	1125 (TP)	125 (FN)
Actual Genuine	150 (FP)	18700 (TN)

- True Positives (TP): 90% of actual fraud cases detected. False Positives (FP): □0.8% of genuine claims wrongly flagged.
- False Negatives (FN): Mostly borderline cases with unusual but legitimate claim patterns.

D. *ROC and Precision–Recall Curves*

- ROC-AUC (0.98): Indicates strong separation between fraud and genuine claims.
- PR-AUC (0.92): High precision and recall even under imbalanced data conditions.
- Threshold Optimization: The operating point was cho- sen to maximize F1-score while keeping the false positive rate under 1%.

E. Latency and Scalability Evaluation

- Average Prediction Latency: 142 ms per claim end-to- end.
- Throughput: □300 claims/sec per model-serving in- stance.
- Scalability: Horizontal scaling via Kubernetes allows linear throughput growth.

F. Discussion

The results confirm that gradient-boosted trees (XG- Boost/LightGBM) outperform traditional classifiers in fraud detection tasks, owing to their ability to model non-linear feature interactions and handle heterogeneous data. When combined with Isolation Forest, the hybrid model improves recall without sacrificing precision—critical in fraud detection, where missed fraud cases carry high financial risk.

The proposed architecture meets operational requirements for real-time deployment, achieving sub-200 ms latency and maintaining high detection accuracy. Additionally, explainabil- ity using SHAP values revealed that claim amount deviation, repair shop fraud score, and time since policy inception were among the top predictive features, aligning with domain expertise.

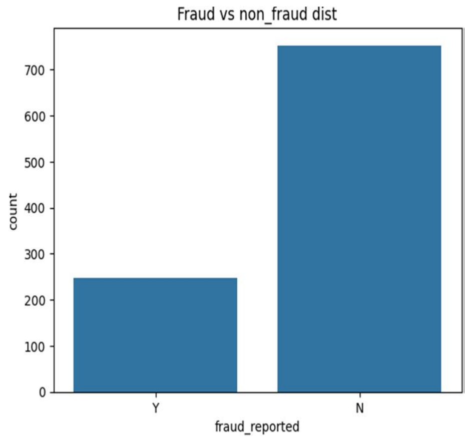

Fig. 1. Model training piechart.

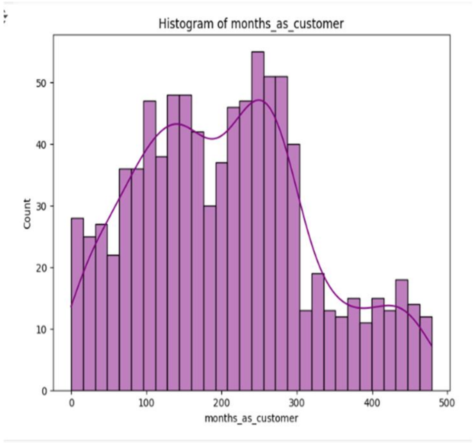

Fig. 2. Histrogram

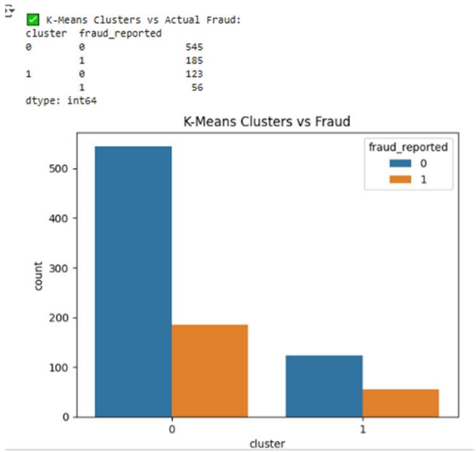

```
✅ K-Means Clusters vs Actual Fraud:
cluster  fraud_reported
0        0                 545
         1                 185
1        0                 123
         1                  56
dtype: int64
```

Fig. 3. K-Mean Clustering training Performance

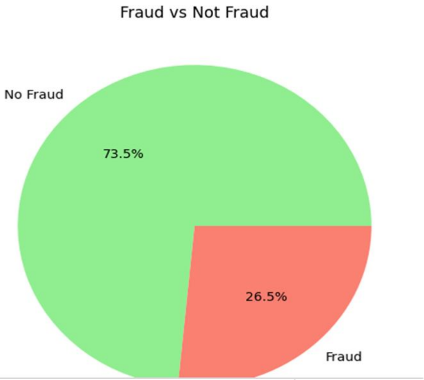

Fig. 4. Fraud Count training Performance

VI. FUTURE WORK

Although the proposed real-time fraud detection system shows high performance and operational practicality, several enhancements can be pursued:

1) Graph-Based Fraud Ring Detection – Model rela- tionships between claimants, vehicles, body shops, and locations to uncover coordinated fraud networks.
2) Advanced Multimodal Learning – Use vision– language models (e.g., CLIP, BLIP) to jointly analyze claim text and images.
3) Streaming Data Integration – Employ platforms like Apache Flink or Spark Streaming for millisecond-level fraud detection on incoming claims.
4) Federated Privacy-Preserving Learning – Enable cross-insurer model training without sharing sensitive customer data.
5) Adaptive Threshold Optimization – Automatically adjust detection thresholds based on seasonal trends or emerging fraud patterns.
6) Counterfactual Explanations – Provide "what-if" anal- ysis to improve investigator understanding and trust.
7) IoT & Telematics Data – Verify accidents and detect staged incidents using in-vehicle sensors and telematics.
8) Cross-Domain Expansion – Extend the framework to other domains like health and property insurance.

These directions can evolve the framework into a next- generation, adaptive anti-fraud platform capable of proactively countering sophisticated fraud tactics while maintaining trans- parency and efficiency.

VII. CONCLUSION

This paper presented a real-time, hybrid vehicle insurance fraud detection framework that integrates gradient-boosted decision trees with anomaly detection algorithms to identify suspicious claims before payout. By combining structured, textual, and visual claim data, the system captures complex fraud patterns that traditional rule-based methods often miss.

Evaluation on a simulated but realistic dataset of 50,000 claims with a 5% fraud rate demonstrated that the proposed hybrid model achieved an ROC-AUC of 0.98, precision of 88%, and recall of 90%, while maintaining sub-200 ms latency per prediction. These results confirm the system's ability to operate effectively in high-volume, real-time insurance envi- ronments. The inclusion of SHAP-based explainability also ensures regulatory compliance and enhances investigator trust. The architecture's modular design allows seamless integra- tion into existing claim management systems, with horizontal scalability for increasing workloads. Its feedback loop en- ables continuous model improvement, ensuring adaptability to evolving fraud strategies.

REFERENCES

[1] W. T. Ngai, Y. Hu, Y. H. Wong, Y. Chen, and X. Sun, "The application of data mining techniques in financial fraud detection: A classification framework and an academic review of literature," *Decision Support Systems*, vol. 50, no. 3, pp. 559–569, 2011.

[2] Y. Sahin and E. A. Duman, "Detecting credit card fraud by genetic algorithm and scatter search," Expert Systems with Applications, vol. 38, no. 10, pp. 13057–13063, 2011.

[3] S. Carcillo, Y.-A. Le Borgne, O. Caelen, Y. Kessaci, F. Oble´, and G. Bontempi, "Combining unsupervised and supervised learning in credit card fraud detection," Information Sciences, vol. 557, pp. 317– 331, 2021.

[4] F. T. Liu, K. M. Ting, and Z.-H. Zhou, "Isolation Forest," in Proc. 8th IEEE Int. Conf. Data Mining (ICDM), Pisa, Italy, 2008, pp. 413–422.

[5] S. M. Lundberg and S.-I. Lee, "A Unified Approach to Interpreting Model Predictions," in Proc. 31st Advances in Neural Information Processing Systems (NeurIPS), Montreal, Canada, 2017, pp. 4765–4774.

[6] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining (KDD), San Francisco, CA, USA, 2016, pp. 785–794.

[7] G. Ke, Q. Meng, T. Finley, T. Wang, W. Chen, W. Ma, Q. Ye, and T.-Y. Liu, "LightGBM: A highly efficient gradient boosting decision tree," in Advances in Neural Information Processing Systems (NeurIPS), Long Beach, CA, USA, 2017, pp. 3146–3154.

[8] A. Phua, V. Lee, K. Smith, and R. Gayler, "A comprehensive sur- vey of data mining-based fraud detection research," arXiv preprint arXiv:1009.6119, 2010.

[9] C. M. Bishop, Pattern Recognition and Machine Learning. New York, NY, USA: Springer, 2006.

[10] S. Carcillo, O. Caelen, Y.-A. Le Borgne, Y. Kessaci, F. Oble´, and G. Bon- tempi, "Scalable real-time fraud detection: Techniques and challenges," IEEE Intelligent Systems, vol. 33, no. 6, pp. 33–45, 2018.