A Blockchain-Based Fraud Detection and Vehicle Damage Assessment System Using Machine Learning and Computer Vision

Wafa Ben Slama Souei[1], N'Gouari Gana Abdou Bachir[2] and Raoudha Ben Djemaa[1]
[1]*University of Sousse, ISITCOM, Sousse, Tunisia*
[2]*International multidisciplinary school, EPI Digital School, Sousse, Tunisia*

Keywords: Fraud Detection, Machine Learning, Mask R-CNN, Insurance Claims, Computer Vision, Vehicle Damage Assessment.

Abstract: Car insurance is a cornerstone of modern society, offering crucial financial protection in the event of accidents and vehicle-related damage. However, this intricate system is now grappling with a significant challenge: fraud manifesting in various forms, including staged accidents, fraudulent claims, and collusion between individuals, and it poses a serious threat to the integrity and long-term viability of car insurance. In this paper, we will propose an innovative approach to fraud detection in car insurance and reimbursement estimation. The proposed approach makes significant contributions to the field by introducing a new dataset of 5,483 images with corresponding labels. Fraud detection is performed using the XGBoost Classifier, which is known for its robustness in handling complex classification tasks. Damage detection is carried out using the Mask R-CNN model, enabling precise identification and segmentation of vehicle damages. The system integrates structured data fraud detection with image-based damage assessment, where Mask R-CNN results serve as an additional validation factor. This end-to-end approach enhances fraud detection accuracy by combining data-driven insights with visual evidence for more reliable claim verification. These advancements contribute to improving the accuracy and efficiency of automated fraud detection and reimbursement estimation systems.

1 INTRODUCTION

Automobile insurance is a cornerstone of modern society, providing essential financial protection in the event of accidents and vehicle-related damages. However, this complex system is increasingly challenged by a significant issue: fraud. Whether through staged accidents, falsified claims, or collusion between individuals, fraud poses a serious threat to the integrity and sustainability of the automobile insurance industry.

Insurance fraud takes many forms, including staging accidents, document falsification, and exaggerating damage claims. According to the International Association of Insurance Supervisors (IAIS), around 10% of auto insurance claims globally are fraudulent, leading to annual financial losses of approximately \$30 billion for insurance companies in the United States alone (PropertyCasualty360, 2009). This not only results in direct costs for insurers but also undermines policyholders' trust in the insurance system.

The consequences of insurance fraud are far-reaching. It not only inflicts substantial financial losses on insurance companies but also drives up premiums for honest policyholders. A study by the Coalition Against Insurance Fraud found that fraud costs American families an extra \$400 to \$700 annually in increased premiums, imposed to cover the financial burden of fraudulent activities (of Investigation, 2023). This creates an unfair situation for honest customers and places additional financial strain on households, which can, in turn, erode customer satisfaction and loyalty towards insurance providers.

Beyond financial costs, automobile insurance fraud also has broader societal impacts. It contributes to rising healthcare and emergency service expenses, as well as increased road congestion. For example, an analysis by the National Insurance Crime Bureau (NICB) found that staged accidents and fraudulent claims increase the demand for medical care and emergency services, which imposes additional costs on public infrastructure (National Insurance Crime Bureau, 2024). Moreover, such fraudulent activities can compromise road safety by encouraging reckless driving behaviors aimed at fabricating accidents. This, in turn, creates serious challenges for public

Souei, W. B. S., Bachir, N. G. A. and Ben Djemaa, R.
A Blockchain-Based Fraud Detection and Vehicle Damage Assessment System Using Machine Learning and Computer Vision.
DOI: 10.5220/0013500800003929

In *Proceedings of the 27th International Conference on Enterprise Information Systems (ICEIS 2025) - Volume 1*, pages 1123-1130
ISBN: 978-989-758-749-8; ISSN: 2184-4992

infrastructure management and healthcare services, which must respond to the growing demand for resources.

Addressing insurance fraud is a complex issue that demands a multifaceted approach. Insurance companies must not only detect fraud but also evaluate its severity and implement appropriate measures to combat it. Enhancing detection and assessment processes is critical to preserving the integrity and effectiveness of the insurance system. Artificial Intelligence (AI) has shown immense potential in transforming and enhancing computer systems across a wide range of industries. AI enables the automation of processes, the analysis of vast amounts of data with unparalleled accuracy, and the extraction of valuable insights to support decision-making.

In the automobile insurance sector, AI is particularly effective at detecting fraudulent activities (Benedek et al., 2022). With advanced machine learning algorithms and predictive analytics, AI systems can identify patterns and anomalies in claims data that would typically escape traditional detection methods (Benedek and Nagy, 2023). This not only helps mitigate financial losses from fraud but also boosts operational efficiency and improves the overall experience for honest policyholders.

Given the growing challenge of fraud, innovative solutions for detecting and preventing it in automobile insurance are more critical than ever. Our project aims to address this need by developing an advanced AI model capable of identifying suspicious claims, assessing the level of fraud involved in each case, and determining the appropriate reimbursement percentage based on the findings.

The integration of blockchain and AI significantly enhances the performance and security of insurance and transport systems (Souki et al., 2024). Blockchain provides transparency (Souei et al., 2024), immutability, and decentralized validation (Souei et al., 2023), ensuring data integrity and reducing the risk of fraud in claims processing. AI, on the other hand, enables the system to analyze large datasets, such as vehicle damage reports, in real-time, ensuring accurate and efficient decision-making. Together, these technologies streamline the claims process, improve automation, and offer higher levels of trust, security, and reliability in the insurance sector. The main objective of our study is to design, develop, and evaluate an Artificial Intelligence (AI) model capable of effectively detecting fraud in automobile insurance claims. Specifically, our model aims to:

- Detect suspicious claims by analyzing data related to the accident, damages, and the policyholder's history.
- Assess the degree of fraud in each claim by evaluating the consistency of the information provided by the policyholder and cross-referencing it with known patterns of fraudulent behavior.
- Deploy the proposed solution using the blockchain technology to maintain a high level of security and accessibility.

Our model strives to be both accurate and fair, minimizing false positives (i.e., avoiding the wrongful classification of legitimate claims as fraudulent) while efficiently identifying proven cases of fraud.

2 RELATED WORK

Fraud detection in automobile insurance has been the focus of numerous recent studies, aiming to enhance the accuracy and efficiency of detection systems.

2.1 Fraude Detection

Quan (Quan, 2024) conducted an in-depth study on this subject by analyzing the landscape of automobile insurance fraud. His approach integrates machine learning models, such as logistic regression, decision trees, and discriminant analysis, to develop predictive models for fraud detection. The results showed that the logistic regression model provided the highest accuracy among the three models.

An innovative approach was presented by Yang et al. (Yang et al., 2023), who developed a multimodal learning framework for automobile insurance (AIML). This framework combines natural language processing and computer vision techniques with a knowledge-based algorithm to detect fraudulent behavior. AIML also incorporates a semi-automatic feature generation algorithm (SAFE) for processing automobile insurance data and a framework for handling visual data. The results demonstrated a significant improvement in the model's performance in detecting fraudulent behavior compared to models that used only structural data.

Kouach, el Attar, and El Hachloufi (Kouach et al., 2022) developed a novel automobile insurance fraud detection system using unsupervised learning. Their system employs Isolation Forest and Local Outlier Factor algorithms to identify fraudulent behavior by detecting anomalies. Isolation Forest isolates abnormal data efficiently, while Local Outlier Factor identifies data points that deviate from their neighbors. The integration of these algorithms allows their system to effectively detect fraudulent activities, offering a promising approach to reduce insurance companies' financial losses.

Table 1: Comparison of existing solutions based on various criteria.

Solution	Data Type	Models Used	Accuracy	Practical Implementation
Quan (2024)	Structured	Logistic Regression, Decision Tree, Discriminant Analysis	High for Logistic Regression	Limited
Yang et al. (2023)	Multimodal	NLP, Computer Vision, Knowledge-based Algorithms	Significant Improvement with AIML	Complex
Kouach et al. (2023)	Structured	Isolation Forest, Local Outlier Factor	Good but with False Positives	Calibration Required
Aziz et al. (2023)	Textual	Naïve Bayes	Good for Textual Data	Simple but Limited
Todevski (2024)	Structured	Logistic Regression, Gradient Boosting, Random Forest	High for Customer Retention	Customer Retention Specific
Zhang et al. (2020)	Images	Mask R-CNN, ResNet, FPN	Very High with Enhanced Precision	Complex but Effective
Widjojo et al. (2022)	Images	Mask R-CNN, EfficientNet, MobileNetV2	Very High with 91% F1 Score	Complex
Jayaseeli et al. (2021)	Images	Mask R-CNN	Very High with Improved Damage Detection	Complex
Our Solution	Structured and Unstructured	Logistic Regression, Random Forest, Gradient Boosting	Very High with Multimodal Combination	Adaptable and Scalable

Aziz, Fareedullah, & Mahmood (Aziz et al., 2022) proposed an automobile insurance fraud detection model using advanced machine learning techniques. Their approach primarily relies on the Naïve Bayes classifier, a simple yet powerful probabilistic classification algorithm. Naïve Bayes is known for its simplicity and effectiveness in data classification, especially with textual data. In the context of automobile insurance fraud detection, the model was adapted to identify potential fraud patterns by analyzing claim characteristics and comparing them to normal behavior patterns.

The study's results demonstrated that the Naïve Bayes model achieved significant accuracy in fraud detection, making it an effective tool for insurance companies to combat fraud and reduce financial losses.

And finally, another example is Todevski (Todevski et al., 2021), who developed an artificial intelligence model aimed at increasing the customer base for insurance companies. Their approach utilizes three AI models: logistic regression, gradient boosting, and random forest. These models were used to predict whether a potential customer would remain with the company, with a prediction probability of 81%.

2.2 Car Damage Detection

These works significantly contribute to automobile insurance claims management by providing an automated and accurate method for vehicle damage assessment.

Zhang et al. (Zhang et al., 2020) developed an enhanced Mask R-CNN algorithm for vehicle damage detection. They improved detection accuracy by using an optimized ResNet with Feature Pyramid Network (FPN) and fine-tuned the Region Proposal Network (RPN) for better region proposals. The model, trained with COCO dataset weights, achieved significant performance gains, providing more accurate damage identification and reducing manual evaluation costs.

Widjojo et al. (Widjojo et al., 2022) developed a deep learning system for vehicle damage detection and classification using transfer learning. Their approach includes damage segmentation with Mask R-CNN, damage detection with EfficientNet, and damage classification with MobileNetV2. The system combines segmentation outputs with feature extraction for improved classification accuracy, achieving an F1 score of 91%. Their integrated model offers better accuracy and resource management compared to other CNN models, enhancing damage assessment and insurance claims processing.

Jayaseeli et al. (Jayaseeli et al., 2021) utilized Mask R-CNN for vehicle damage detection and cost assessment. Their model, trained with COCO dataset weights and fine-tuned on damaged vehicle images, uses precise annotation and a ”color splash” visualization technique to highlight damage. This approach improves detection accuracy and cost assessment transparency, reducing insurance claims processing costs and fraud risks while enhancing repair estimate precision and claims evaluation efficiency.

2.3 Synthesis

After a thorough analysis of Table 1, it appears that the previously presented work has certain shortcomings. Firstly, many studies primarily focus on the development and validation of fraud detection models in a laboratory setting. However, there is a gap in research concerning the implementation of these models in an operational environment within insurance companies. The challenges related to integrating existing systems, managing organizational change, and user acceptance are often underestimated. A thorough understanding of the practical obstacles to adopting these technologies and strategies to overcome them is crucial to ensure a successful transition from theoretical models to practical applications.

Secondly, most current research focuses on leveraging structured data, such as tabular data from insurance company databases. Nevertheless, a considerable amount of relevant information is contained in unstructured data, such as images of damaged vehicles, accident videos, and text from claims reports. Integrating this unstructured data into fraud detection models could enhance their accuracy and robustness. Advanced techniques such as natural language processing (NLP) and image recognition need to be explored to leverage these diverse data sources.

Finally, there is a notable lack of research on the actual impact of fraud detection models on reducing fraud and improving the profitability of insurance companies. Current studies often focus on model performance metrics such as precision, recall, and F1-score, without examining their effectiveness in a real-world context. Longitudinal studies are essential to assess the long-term impact of these models on fraud prevention, customer satisfaction, and financial gains. This also includes analyzing the costs associated with the implementation and maintenance of these systems.

3 PROPOSED SOLUTION

This research work focused on the development and evaluation of artificial intelligence models for fraud detection in car insurance and reimbursement estimation. The proposed approach illustrated in Figure 1

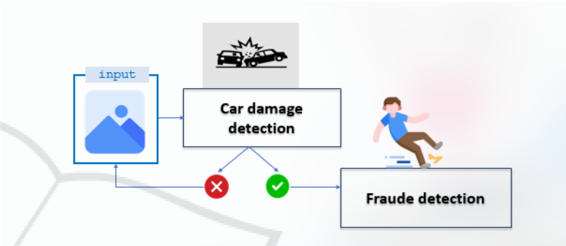

Figure 1: Architecture of the Proposed Solution.

makes significant contributions to the field of fraud detection in car insurance and reimbursement estimation;

- Creating a new data set of 5,483 images and labels.
- The detection of fraud is performed based on the XGBoost Classifier.
- The detection of damage is performed based on the Mask R-CNN model.

4 RESEARCH METHODOLOGY AND STEPS

Previous research has underscored that existing deep learning-based vehicle damage recognition techniques often overlook the damage volume necessary for auto insurance claims. In addition, accurately identifying various types of vehicle damage and assessing their severity is critical for practical applications. To address these shortcomings, this paper introduces a prototype system called CES (Cost Estimation System), designed to detect damaged vehicle parts, evaluate the extent of the damage, and estimate the total claim cost. The CRISP-DM methodology was employed to guide the development of this prototype, ensuring its accuracy and alignment with business objectives (Chapman et al., 1999).

4.1 DataSet

The data collection for this project was carried out using two main sources: Kaggle and data.gov. The data from Kaggle includes rich and well-documented datasets on car insurance claims, while data.gov provides additional details on accidents, such as speed, engine temperature, and fuel consumption. These reliable sources allowed for the creation of a solid foundation for analysis, with comprehensive information on drivers, insured vehicles, and claims.

Additionally, a dataset of 1,400 annotated images, including segmentation masks and bounding boxes, was used to train a Mask R-CNN model. It consists of multiple images of vehicles categorized into three levels of damage as illustrated in Figure 2: soft, medium, and hard. These images were divided into three groups: 900 for training, 300 for validation, and 200 for testing, all accompanied by JSON files containing detailed metadata. The images were preprocessed to ensure quality, resized, normalized, and enhanced using augmentation techniques such as rotation and exposure adjustments. The total dataset includes 5,483 images and an equal number of labels, ensuring a robust foundation for training and evaluating the Mask R-CNN model. Data preprocessing involved resizing, augmenting, and splitting the dataset to optimize model performance while maintaining a balance between subsets. This preparation ensures the model's robustness and its ability to generalize effectively in detecting and classifying vehicle damage.

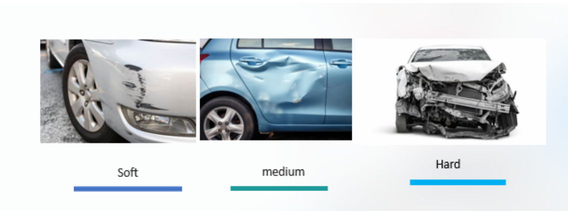

Figure 2: Car Damage Types.

4.2 Evaluation Metrics

In our work, we evaluated the performance of the developed models using various metrics (Naidu et al., 2023): precision and recall to measure the accuracy of positive predictions and the proportion of actual positives detected, respectively. We also used the loss function to quantify the model's prediction errors during training and the F1-score to balance precision and recall, especially for imbalanced datasets. Finally, the mean Average Precision (mAP) was employed to assess overall detection accuracy across all classes in object detection tasks. These metrics are essential to ensure the robustness and reliability of our recommendation system.

4.3 Models Training and Testing

In this stage, the focus is on selecting the appropriate model for the task at hand. Different modeling techniques are considered based on the nature of the data and the problem being solved. The goal is to choose a model that aligns with the desired outcomes and can effectively learn from the data to make accurate predictions.

4.3.1 Damage Detection

The detection of vehicle damage is performed using the Mask R-CNN model, a state-of-the-art deep learning framework for instance segmentation. This model was trained on a dataset of annotated vehicle images, where damages were categorized into minor, moderate, and severe levels. Each image was labeled with segmentation masks and bounding boxes to enable precise localization of damaged areas. Preprocessing techniques such as image resizing, normalization, and data augmentation (rotation, brightness adjustments, and noise reduction) were applied to enhance the model's robustness. The performance of the Mask R-CNN model was evaluated using key metrics such as Intersection over Union (IoU), mean Average Precision (mAP), and detection accuracy. The results showed in Figure 3 that the model achieved an IoU score of 0.85 and an mAP of 0.80, demonstrating its ability to accurately segment and classify vehicle damages. The precision-recall analysis indicated a high confidence level in detecting damaged regions while minimizing false positives. These findings confirm that the Mask R-CNN model is highly effective in automating vehicle damage assessment, providing insurers with a reliable tool for faster and more objective claim evaluations.

4.3.2 Fraude Detection

The detection of fraud in automobile insurance claims is carried out using the XGBoost Classifier, a powerful gradient boosting algorithm known for its efficiency and high predictive accuracy. The model was trained on a dataset containing key features related to claims, including policyholder information, claim history, and vehicle attributes. Preprocessing steps such as feature selection, categorical encoding, and

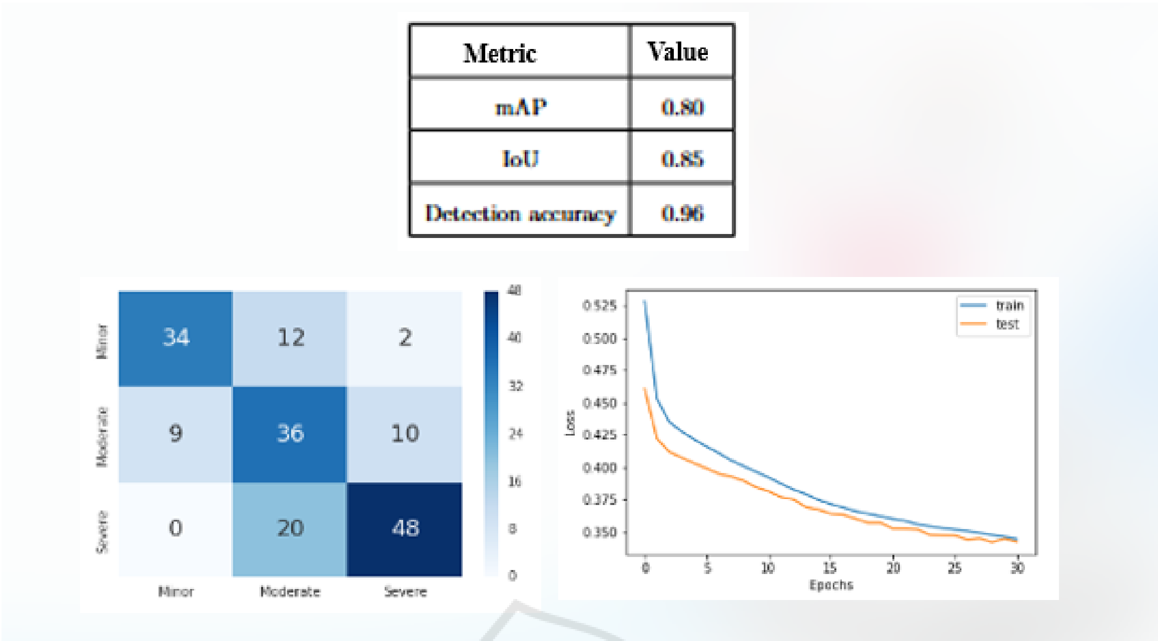

Metric	Value
mAP	0.80
IoU	0.85
Detection accuracy	0.96

Figure 3: Damage detection results.

handling missing values were applied to enhance data quality.

The model's performance was evaluated using key metrics such as accuracy, precision, recall, and F1-score. The results showed that XGBoost achieved an accuracy of 92%, with an F1-score of 0.87, demonstrating its strong ability to distinguish between fraudulent and legitimate claims. The precision-recall curve indicated that the model effectively minimized false positives, reducing the risk of denying legitimate claims. Furthermore, feature importance analysis revealed that factors such as claim amount, previous fraudulent claims, and inconsistencies in reported accident details were the most influential in fraud detection. These findings highlight the effectiveness of XGBoost in improving fraud detection processes and reducing financial losses for insurance companies.

5 DEPLOYMENT

In the final stage, the solution is deployed for operational use in the automobile insurance context, provided it meets the evaluation criteria. During this phase, the models are integrated into the existing insurance system based on the blockchain technology, user interfaces are developed to facilitate interactions, and real-time monitoring is implemented to assess performance in real-world scenarios. Figure 4 illustrates the deployment diagram of our system that integrates AI and blockchain for the management of auto insurance claims. The process can be summarized as follows:

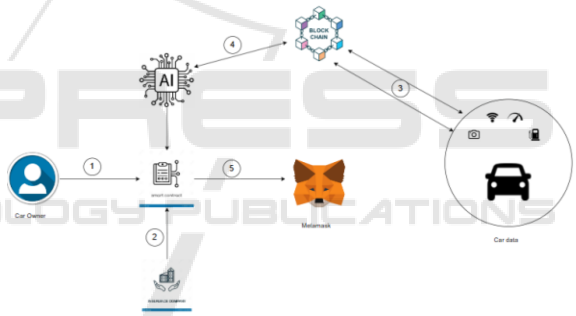

Figure 4: Architecture of the Auto Insurance Claims Management System.

- **Claim Submission.** The vehicle owner initiates a compensation request via a smart contract.
- **Insurance Validation.** The insurance company receives and evaluates the submitted claim.
- **Vehicle Data.** Relevant vehicle information, including images, sensor data, and telematics, is collected and transmitted for processing.
- **AI Analysis.** Our AI model processes the data to assess damage and verify the claim's validity.
- **Blockchain Interaction.** The entire process is recorded and validated on the blockchain, ensuring transparency, traceability, and security. The user can access their information through its Metamask portfolio.

The proposed solution is deployed on the Ethereum blockchain, where a smart contract processes real-time data from the insured vehicle, including sensor readings, camera inputs, and other telematics information. Based on this data, the smart contract autonomously determines whether an accident has occurred. Additionally, the severity of the damage is assessed using an AI-powered damage detection model, which analyzes the collected information to provide an accurate evaluation. The entire process ensures transparency, security, and automation in claims processing, leveraging blockchain for immutability and trust.

To conclude, Blockchain ensures data immutability, transparency, and trust in fraud prevention by securely recording real-time vehicle data and enforcing automated claim validation through smart contracts. Its decentralized nature prevents tampering and internal fraud, enhancing insurance claim verification. Combined with AI, it strengthens fraud detection and ensures fair settlements.

6 RESULTS AND DISCUSSION

This study highlights the transformative potential of AI and ML technologies in the insurance industry. By automating fraud detection and improving damage assessment accuracy, the proposed framework addresses critical inefficiencies in traditional claim processing systems.

The XGBoost classifier and Mask R-CNN model both showed impressive performance in their respective tasks. XGBoost achieved an AUC of 0.89, effectively minimizing false negatives through its regularization techniques and weight updates, while balancing precision and recall. The confusion matrix further confirmed its solid performance, although false positives still remained, indicating room for improvement. On the other hand, the Mask R-CNN model excelled in damage detection, with a remarkable detection accuracy of 96%, an mAP of 0.80, and an IoU of 0.85. These results highlight the model's strong capability to accurately detect and segment vehicle damages, an essential feature for streamlining insurance claims. The stable loss curves throughout training and validation indicate the model's ability to generalize well to new data. Despite some minor misclassifications in damage severity, the overall performance demonstrates its practicality and potential for automating claims processing in the insurance sector.

The findings pave the way for future research exploring advanced techniques, such as deep learning-based anomaly detection and real-time fraud prevention systems, to further enhance the robustness and scalability of these solutions.

7 CONCLUSION

This project focuses on fraud detection in automobile insurance claims and vehicle damage assessment through machine learning and computer vision. The XGBoost model outperformed other algorithms with an AUC of 0.89, effectively minimizing false negatives through regularization techniques and weight updates. This highlights its strong capability in fraud detection. Additionally, the Mask R-CNN model excelled in segmenting and evaluating vehicle damage, achieving a detection accuracy of 96%, a mean average precision (mAP) of 0.80, and an Intersection over Union (IoU) of 0.85. These results underline the model's effectiveness in accurately detecting and segmenting vehicle damages, essential for automating the claims process in insurance. Overall, this project demonstrates the significant potential of combining XGBoost for fraud detection and Mask R-CNN for damage assessment in streamlining insurance operations.

REFERENCES

Aziz, T., Fareedullah, and Mahmood, M. (2022). Insurance fraud detection model: Using machine learning techniques. *Journal Name*, Volume(Issue):Page numbers.

Benedek, B., Ciumas, C., and Nagy, B. Z. (2022). Automobile insurance fraud detection in the age of big data–a systematic and comprehensive literature review. *Journal of Financial Regulation and Compliance*, 30(4):503–523.

Benedek, B. and Nagy, B. Z. (2023). Traditional versus ai-based fraud detection: Cost efficiency in the field of automobile insurance. *Financial and Economic Review*, 22(2):77–98.

Chapman, P., Clinton, J., Kerber, R., Khabaza, T., Reinartz, T., Shearer, C., and Wirth, R. (1999). The crisp-dm user guide. In *4th CRISP-DM SIG Workshop in Brussels in March*, volume 1999. sn.

Jayaseeli, J. D. D., Jayaraj, G. K., Kanakarajan, M., and Malathi, D. (2021). Car damage detection and cost evaluation using mask r-cnn. In *Intelligent Computing and Innovation on Data Science*, volume 248 of *Lecture Notes in Networks and Systems*, pages 279–288, Kattankulathur, India. Springer, Singapore.

Kouach, Y., el Attar, A., and El Hachloufi, M. (2022). Auto insurance fraud detection using unsupervised learning. *Journal Name*, Volume(Issue):Page numbers.

Naidu, G., Zuva, T., and Sibanda, E. M. (2023). A review of evaluation metrics in machine learning algorithms. In

Computer Science On-line Conference, pages 15–25. Springer.

National Insurance Crime Bureau (2024). Staged auto accident fraud. Accessed: 2025-02-01.

of Investigation, F. B. (2023). Insurance fraud. Accessed: 2025-01-30.

PropertyCasualty360 (2009). Fraud causes $30 billion in loss annually: Iii. Accessed: 2025-01-30.

Quan, Y. (2024). Research on detecting auto insurance fraud. *Highlights in Business Economics and Management*, 24:278–284.

Souei, W. B. S., Charroux, B., Sliman, L., and Ben Djemaa, R. (2024). Towards blockchain like soa. In *International Conference on Web Information Systems Engineering*, pages 301–311. Springer.

Souei, W. B. S., El Hog, C., Djemaa, R. B., Sliman, L., and Amor, I. A. B. (2023). Towards smart contract distributed directory based on the uniform description language. *Journal of Computer Languages*, 77:101225.

Souki, O., Djemaa, R. B., Amous, I., and Sedes, F. (2024). Monitoring and analyzing as a service (maaas) through cloud edge based on intelligent transportation applications. *Cluster Computing*, 27(3):3379–3395.

Todevski, D., Malik, H. S., Omkar, S. N., Monis, E. B., Khanna, B., and Samal, S. R. (2021). Fraud detection in insurance with machine learning. *Journal Name*, Volume(Issue):Page numbers.

Widjojo, D., Setyati, E., and Kristian, Y. (2022). Integrated deep learning system for car damage detection and classification using deep transfer learning. In *2022 IEEE 8th Information Technology International Seminar (ITIS)*, Surabaya, Indonesia. IEEE.

Yang, J., Chen, K., Ding, K., and Na, C. (2023). Auto insurance fraud detection with multimodal learning. *Data Intelligence*, 5(2):1–25.

Zhang, Q., Chang, X., and Bian, S. (2020). Vehicle-damage-detection segmentation algorithm based on improved mask rcnn. *IEEE Access*, 8:6997–7004.