

Intelligent Anomaly Detection in Insurance Claims Processing: A Supervised Learning Approach to Fraud Classification

Dr. Nandini Sinha, Associate Professor of Computer Science, Indian Institute

1. Introduction to Fraud Detection in Insurance Claims

The insurance industry is inevitably required to compensate for fraud cases, which have led to significant financial losses. Some of the typical examples of fraudulent activities are exaggerating the value of a claim, intentionally damaging the assets, re-filing the claim by hiding some information, or not revealing the existence of other insurance carriers. Many organizations have recognized that the detection of these fraudulent activities is crucial to reduce the financial impact on insurers and policyholders. Fraud worth over £1 billion was committed in the UK due to insurance claims only. With the advancements in artificial intelligence and the large volume of digital channels, fraudulent activities have expanded from traditional fraud of claim misrepresentation towards technology-driven fraud that requires a more general, modern fraud detection solution.

It is essential to convert from the traditional rule-based methodology to modern methodologies such as machine learning and deep learning, which can capture a more general fraud pattern. With this motivation, this paper proposes to deploy AI-based models to detect fraudulent activities and prevent them from compensating the insurances. The rest of the paper is organized as follows. In Section II, we review various methodologies along with the fraudulent claim prediction challenges. In Section III, we discuss the proposed methodologies. In Section IV, we present case studies and their outcomes. We discuss the results and offer conclusions in Sections V and VI, respectively.

2. Machine Learning Basics for Fraud Detection

Machine learning holds the potential for an AI that can emulate human decision-making or uncover hidden patterns within data, and leverages methods to autonomously adapt to new information and trends. The field can generally be divided based on the amount of human involvement. Supervised machine learning trains a system on a labeled dataset, where inputs and outputs are clearly identified, while unsupervised machine learning is trained on raw data without any human labeling and is expected to classify it into appropriate categories. The relevance of these learning models to insurance claims is that claims data can be structured as either fraudulent or genuine transactions. Furthermore, semi-supervised machine learning encompasses combining elements from each field.

In practical terms, machine learning is applied to fraud detection by training a model on a historic dataset in order to optimize its predictive accuracy on unknown cases in the future. A selection of resampling-based techniques is discussed in relation to this in the practical guidelines. Once a database has been assembled, any instances of conflicting data, duplicates, missing information, incorrect data, and any reported values that fall outside of reasonable ranges can be removed. The training of a machine learning model is an iterative process, wherein a number of methods all attempt to find an optimal combination of inputs designed to maximize the overall predictive performance of the group. There are also a number of different ways that errors can be introduced to the learning process that will have an effect on the performance of the model. The most fundamental of these is overfitting - where the model assumes the random noise within the dataset is representative of the real underlying data trend. Model interpretability has implications for the subsequent decision-making processes once the data has been split into categories, and a distinction can be made between model accuracy and model interpretability when identifying a machine learning model for any specific task. Real-world examples of machine learning for fraud detection relevant to the insurance industry are discussed in greater detail in the applications of machine learning in fraud section.

3. Data Preprocessing and Feature Engineering

Data preprocessing and feature engineering are essential parts when building a machine learning model. Proper data preprocessing usually improves the performance of

machine learning models. One of the key components that guarantee better performance in machine or deep learning models is data quality. Mainly, in order to get the best productivity from our machine learning model, we need to minimize data quality problems such as noise and outliers that usually lead to wrong outputs. In addition to that, missing values are other common data problems that make it even worse. All of these data problems corrupt the data and make the model learn correlated features.

To avoid misinterpretation in modeling, intermediate data preprocessing is needed. Datasets in this project were preprocessed prior to the application of a machine learning model. Some of the common types of data preprocessing in machine learning models are normalization and standardization. These two concepts are implemented to improve the network model's understanding of the particular data. Features can have a huge impact on model performance. Hence, the selection or extraction of subjective domain knowledge is crucial. A combination of these methods can also improve the quality of data input. Feature extraction methods, such as adding synthetic features for the fraud model, can increase model performance. It is important to preprocess the initial data prior to modeling. The subsequent section continues with guidelines to practically implement these phases in preprocessing. It is necessary to reiterate that before the model can be utilized, the data preprocessing step should be done carefully. Therefore, missing values or outliers should be carefully handled. Excluding missing values without a replacement or outlier analysis restricts the functionality of the model. The subsequent section continues with guidelines to practically implement these phases in preprocessing. It is necessary to reiterate that before the model can be utilized, the data preprocessing step should be done with utmost care.

4. Supervised Learning Models for Fraud Detection

Supervised learning algorithms learn from labeled datasets that contain an input black box and corresponding output that we want to predict. For detecting insurance fraud, the fraud labels are given to the machine learning model, which uses different attributes extracted from the claim to distinguish between fraudulent and genuine claims. After this learning phase, the model can then be used to automatically classify new claims into fraudulent or genuine categories. There is a diverse range of supervised learning techniques that can be utilized to solve fraud problems. These techniques include linear

classification like logistic regression, decision boundaries like decision trees, rule-based models like simple CARTs, and ensemble methods.

Logistic regression works relatively well with real-world datasets, whose attributes are noisy. Decision trees are attractive in that they do not require feature scaling and can vary from being inefficient to effective even when dealing with large feature spaces. Rule-based models create simple descriptions of credit decision processes that are contained in data and thus help keep human involvement low. In particular, ensemble learning-based approaches typically offer the best predictive performance in many applications, even though they require feature scaling. The ensemble methods usually have higher classification accuracy than a single model. However, using a black box model and a complex framework may be disadvantageous from a business perspective. Techniques for optimizing the performance of the chosen classifier include cross-validation methods and searching for the right combination of parameters for tuning the classifier to make better predictions. However, the larger and cleaner the training dataset is, the better the learning performance and the final performance on unseen data.

Data quality attributes might include accurate and consistent data, while data quantity attributes may include having enough data and a sufficient amount of data records. A key challenge in building automated fraud detection systems is to define and use meaningful measures of performance to evaluate and compare the predictive performance of potential model classifiers. Finally, imbalanced outcomes create difficult learning environments for the classification process as the number of samples from each class is naturally imbalanced, with genuine claims heavily outweighing fraudulent ones, and different samples have different interpretations and characteristics regarding the area under the distribution. Several methods exist for overcoming the imbalance problem with particular trade-offs between methods, including over-sampling and under-sampling.

5. Unsupervised Learning Models for Anomaly Detection

Unsupervised learning techniques can be implemented to detect anomalies in a data set lacking labeled data. These algorithms discover regularities in the data and point out the unusual or unexpected patterns. Although there isn't a separate training process, unsupervised learning models are able to make assumptions about the data and determine whether newly seen data should be considered regular or treated as an

outlier. Considering fraud in insurance claims, this type of algorithm could find unusual patterns that could be interpreted as possible instances of fraud. Unsupervised learning models are a group of machine learning systems that extract and find relationships or patterns in the data without predefined outcomes or classifications assigned to it. By finding irregularities in the data during processing, unsupervised learning models could point out fraudulent claims. They could be used to find frequent items and clusters that are noticeably different from the rest of the data by looking for outliers while learning associations between different instances, which could be used to identify unusual connections. The algorithms belonging to this group are k-means, hierarchical clustering, and association rule mining techniques; they could all help detect anomalies and aid in the process of fraud detection.

Moreover, there are some techniques developed to systematize data from a different perspective, analyzing the frequency of the items in the data set compared to the rest of the data. Density-based methods generate a score based on the local density of instances' data and are able to discover anomalies in multi-dimensional datasets. In practice, gaining insight into the underlying structures of the data can help identify what may not actually be wrong, but it certainly doesn't fit into the norm, creating additional alerts that can be used to increase the performance of the fraud detection systems. Anomaly detection offers a different perspective in fraud analysis, working to complement existing detection systems by finding unknown risks and offering insights on causation analysis. Combining unsupervised and supervised learning paradigms would enable insurers to keep pace with the evolving patterns of fraud and continuously provide value by combining analysis of known fraud patterns with the detection of new and emerging threats. Furthermore, there are techniques developed to extract patterns from the data at lower computational costs. In the interest of system scalability, knowledge discovery from association rules and frequency-based analysis can be a lower computational alternative to clustering. Examples of these approaches include the identification of deviation factors in claims using a system to model claims based on what is inflationary or counter-cyclical. Information on these principles is fed into claims systems, where it is used to identify claims from claim handlers' judgment in cases like severity assessment. An outlier score derived from the deviating association rules is then generated.

6. Hybrid Approaches Combining Supervised and Unsupervised Learning

A new trend has emerged in fraud detection systems for dealing with complex and elaborate fraud schemes that cannot be identified by standard methods. This trend is known as hybrid fraud detection because it combines the strengths of both supervised and unsupervised learning to identify advanced fraud patterns. The main strength of supervised learning is its ability to model complex classification problems that are not easily addressed by simple methods. On the other hand, unsupervised techniques can assist in identifying cases of fraud that cannot be modeled effectively in situations where small amounts of fraud exist. When combining both techniques, the fraud detection system can identify both known frauds and unknown frauds running beneath the surface, triggering curiosity through unsupervised techniques.

A number of methodologies have been proposed to integrate the use of unsupervised and supervised learning. Many studies have shown that this approach is successful in the context of insurance fraud. By combining supervised and semi-supervised learning techniques, a new fraud detection system was developed as a replacement for a previous system. The new system significantly increased the number of claims detected on the training set and on the testing set. It was revealed that using a hybrid model is beneficial not only for detecting fraudulent claims but also for handling the issue of data imbalance and ensuring model robustness. A hybrid model-based approach provides more gain than other developed models. Despite differences in the definition, most studies utilize a combination of both techniques to create their hybrid systems. Given these demarcations and the recent attention given to this topic, this approach could be a future trend in the insurance fraud field. There are various challenges to creating a robust hybrid fraud detection system. One of them is explicitly related to the difficulties of managing and maintaining the system, in addition to those that directly influence the fraud detection model itself. That being said, it is also important for adaptive learning, because we specifically focus on cases where it is infeasible to fully encompass and clarify the fraud item. Such cases become increasingly harder to tackle as time progresses due to the possibility of the behaviors and patterns of fraudsters evolving. As such, there is a pressing need for further investigation in establishing hybrid models to combat this.

7. Evaluation Metrics for Fraud Detection Models

'Evaluation Metrics'

Once a model is trained, the next step is to assess its performance. Various evaluation metrics can be used to judge the effectiveness of a fraud detection model. The most popular for supervised models are accuracy, precision, recall, and F1 score, which captures the trade-off between precision and recall. Precision tells us the rate of true positives in all predicted positives, while recall is the share of true positives over all actual positives. It is important to understand the meaning of false positives and false negatives and the trade-offs for the problem at hand. A model predicting all classes as non-fraud will have a 0 precision and 0 F1 score due to its false positives, while a model that flags all claims as fraudulent will have 100% recall, but it is not very useful either due to the high volume of false alarms.

Various metrics for evaluation for unsupervised models also exist. In the case of highly imbalanced data, different metrics are used for evaluation to better measure the performance of the model and select the appropriate evaluation metric or model for the problem. A widely used unsupervised evaluation metric for imbalanced datasets is the F1 score. Predicting all cases as non-fraud due to the imbalance may deliver a high accuracy, but the F1 score provides a remedy by focusing on the minority class. To compare the results of different models, receiver operating characteristic curves can be used, and the area under the curve is a simple summary measure to communicate model effectiveness and is suitable to compare model performance across a range of trade-offs. For an approval process, for example, the actual fraudsters should be ranked higher, so the so-called KS statistics would also be a good evaluation metric as it measures the maximum displacement of the true positive rate from the false positive rate. For industry, a significant evaluation metric used in the evaluation of fraud detection is the percentage of monetary value detected.

Irrespective of the choice of performance metric, the applicability of the model or the evaluation metric chosen must be tailored according to the business problem or used in combination with certain other business metrics relevant to the problem. This is because slight changes in the evaluation metric or choice of business metric can sometimes lead to the choice of quite different models, depending on the nature of the trade-off for these evaluation or business metrics. Moreover, all evaluation results should be transparently

documented and interpreted to senior management and stakeholders by clearly communicating the business implications of that metric. Moreover, when the trade-offs are not easily accessible, performance estimation may constitute only an intermediate step, and generalization difficulties call for robust estimates based on testing in scenarios as close as possible to the operational one.

8. Challenges and Future Directions in AI-Based Fraud Detection

While significant progress has been made, AI-based fraud detection is also challenging in many aspects, and currently, many unsolved issues arise. We should pay attention to these issues with distinctive looks. This section provides a future direction of AI in fraud detection as a result of these problems. (1) Data privacy, ethical issues, and regulatory compliance are not only important aspects in the application of AI techniques, but they also influence the fraud detection system. AI systems should be transparent in their decision-making process, and they should also be explainable to be trusted by users and regulators. To the best of our knowledge, this aspect has been mainly overlooked, and it is not clear how to apply this knowledge to the fraud detection models. (2) Fraudulent behaviors continuously evolve, and more efforts in continuous model adaptation are needed, especially in fraudulent insurance claims. Similarly, an exploration of non-conventional machine learning approaches, such as one-shot learning, would also be beneficial. (3) A great deal of research interest has been spent on the development of generic deep learning-based models in fraud detection. Such models are not specifically proposed for the insurance sector. There is potential in proposing models specific to the insurance sector, with insurance-specific feature engineering. (4) We provide case study results validated in the auto insurance fraud classification domain. Results appear to be transferable to other claims fraud detection problems. Besides, there is potential for transferring the current frameworks to more advanced model architectures, like transformer models applied for fraud detection problems. (5) Major model performance improvements in fraud detection are reached by combining the conventional shallow inner-feature-extracting components with the deep learning outer-feature-learning components. In claims text-based fraud detection, the same also applies when matching the shallow models with discriminative pre-processing components extracting insurance-specific domain knowledge. (6) Collaboration between academia and industry has proven valuable in practice. For future work, we encourage the academic world to move in this direction. In our eyes, such ongoing efforts are necessary to repeat the

current operational applications of the proposed frameworks and to allow further fine-tuning to the dynamically changing world of practical insurance fraud management.

9. Conclusion

In this paper, we discussed the role AI and machine learning can play in fraudulent claim detection in the field of insurance. We briefly reviewed the methodologies used for fraud detection based on claims practice characteristics, unsupervised data analytics, social network analytics, and textual data analysis, and then focused particularly on the use of machine learning algorithms in fraud detection. We considered a real-world case wherein an attempt was made to prevent fraud in the automotive insurance field and discussed two successful prototypes that employ machine learning-based concepts and algorithms to effectively identify potentially fraudulent claims. We also elaborated on the challenges associated with these machine learning-based approaches. In remarkable contrast to the manual rule-based systems of the past, the insurance sector is presently developing more intelligent and efficient AI-aided detection techniques. We described how the detection of fraudulent insurance claims has evolved and how technology has significantly impacted this evolution. To further improve and develop more efficient detection models, the integration of a comprehensive approach including data science, interdisciplinary thinking, and collaboration among experts from the domain, academia, and the private sector is necessary. In the future, the provision of more advanced machine learning models will offer insight into more complex data patterns, increase model adaptability, perform predictive analysis, and establish a risk-scoring system. Potential future developments may include the use of Transparent Box Models to describe the main dependencies between dependent and independent variables through interpretable data. From that point forward, firsthand domain insights may be utilized to enhance the objective targeting method. Lastly, associations must be proactive in detecting claims fraud by making use of data-driven methodologies and remaining one step ahead. Overall, fraud detection in the insurance industry necessitates a pragmatic approach and keen-eyed inspection.