Proceedings of the 8th IEOM Bangladesh International Conference on Industrial Engineering and Operations Management, Dhaka, Bangladesh, December 20-21, 2025

Proceedings of the International Conference on Industrial Engineering and Operations Management

Publisher: IEOM Society International, USA DOI: 10.46254/BA08.20250172
Published: December 20, 2025

Multimodal Forensic Accounting Fraud Detection Using Transformer-Based NLP Models

Md. Tauhid Hossain Rubel
Master of Science in Cybersecurity
Washington University of Science and Technology
Alexandria, VA 22314, USA
mrubel.student@wust.edu

Tahmeed Ali Patwary, Nusrat Faraezi Ivy, Turjo Das Dip, Miyoko Mahzabin Mim, Md. Salman Farshe and Delower Hossen Tuhin
Department of Computer Science
American International University-Bangladesh
Dhaka, Bangladesh
21-44428-1@student.aiub.edu, 22-48608-3@student.aiub.edu, 22-48558-3@student.aiub.edu, 23-52246-2@student.aiub.edu, 22-49844-3@student.aiub.edu, 22-49837-3@student.aiub.edu

Abu Shufian
Department of Electrical and Electronic Engineering
American International University-Bangladesh
Dhaka, Bangladesh
shufian.eee@aiub.edu

Abstract

Financial fraud in corporate disclosures remains a persistent and complex challenge within forensic accounting. Traditional fraud detection methods often depend on shallow lexical indicators or manual audits, which lack scalability and fail to capture the semantic depth of deceptive narratives. Addressing this gap, this research introduces a multimodal forensic accounting fraud detection system that integrates transformer-based NLP models with classical machine learning classifiers. By combining both shallow (TF-IDF, Count Vectorizer) and deep contextual representations (BERT, Longformer), the study systematically examines model–vectorizer synergies to identify optimal configurations for fraud classification. Experimental findings highlight the superiority of pairing TF-IDF with XGBoost, achieving an impressive F1-score of 97.27% and even perfect accuracy in certain validation folds. Interestingly, transformer-based embeddings yield mixed results, performing best when coupled with adaptive models such as Random Forest or Logistic Regression. These results emphasize the computational efficiency and robustness of hybrid NLP pipelines, along with their potential to enhance early fraud detection in real-world accounting contexts. The research offers a practical and scalable framework that balances interpretability, accuracy, and resource efficiency, paving the way for more effective forensic auditing technologies. However, limitations include a relatively small dataset (170 filings), a text-only focus despite the multimodal designation, and the use of transformer embeddings without task-specific fine-tuning, which may account for their lower performance.

Keywords
Accounting Fraud Detection, Machine Learning, Multimodal NLP, Transformer Models, Vectorization Techniques.

1. Introduction
Fraudulent financial reporting poses significant risks to economies and investor trust, as manipulative corporate disclosures often conceal material truths that mislead stakeholders. Traditional forensic accounting methods primarily rely on manual inspection or rule-based anomaly detection, which can be time-consuming, biased, and ineffective against increasingly sophisticated fraud schemes (Deng et al., 2025). Consequently, there is a growing demand for intelligent systems that can efficiently interpret and classify deceptive textual patterns within financial disclosures. With the advent of natural language processing (NLP) and machine learning, the automation of fraud detection in narrative data has gained significant traction (Lokanan, 2025). Conventional methods like TF-IDF and Count Vectorizer have shown promise in detecting stylometric cues; however, these techniques are limited by their inability to capture context (Moura et al., 2025). On the other hand, transformer-based models such as BERT and Longformer offer deep semantic understanding and contextual awareness, which may uncover subtle manipulations in language (Redwan et al., 2024). Despite their potential, these models introduce computational overhead and are not always compatible with all classifiers, necessitating rigorous evaluation to identify optimal configurations. The term multimodal here denotes multiple textual representations (lexical and contextual) rather than heterogeneous data, and given the small dataset, the results are indicative rather than definitive, requiring larger datasets for validation in future work.

This paper proposes a multimodal forensic accounting fraud detection framework that fuses classical and transformer-based NLP techniques with a range of machine learning classifiers. Through comprehensive experimentation, we evaluate accuracy, F1-score, resource efficiency, and model robustness. Our findings establish a reliable and adaptable architecture, offering new directions for intelligent and scalable fraud detection systems in forensic accounting.

2. Related Work
Over the past few years, numerous studies have explored machine learning and NLP-based approaches to enhance fraud detection in forensic accounting. These works span traditional lexical methods to more recent deep learning and transformer-based models.

In 2019, Goel et al. applied traditional machine learning algorithms using TF-IDF features to detect deceptive narratives in financial disclosures. Their study demonstrated the feasibility of using stylometric indicators for fraud detection, showing fair performance with models like Logistic Regression and Naive Bayes. However, the reliance on shallow features limited the models' ability to capture nuanced deception patterns embedded in longer financial texts. Schuster et al. (2020) examined narrative risk disclosures and employed random forest classifiers with basic text representations. They identified meaningful lexical patterns but emphasized the challenge of incorporating contextual understanding into fraud detection systems. Their findings suggested the need for more advanced linguistic models to go beyond surface-level word frequencies.

The introduction of BERT by Devlin et al. (2018) transformed NLP tasks by enabling deep contextual embeddings. Expanding on this, Zhang et al. (2021) incorporated BERT embeddings into fraud detection pipelines and observed improved accuracy over traditional vectorizers. However, their system faced overfitting issues and required high computational resources, limiting real-world applicability without infrastructure optimization. To address long-sequence challenges in financial narratives, Liu et al. (2022) explored the use of Longformer, a transformer model designed for extended text processing.

While they reported moderate success in context preservation, the performance across financial domains was inconsistent, and the memory overhead remained a barrier for scalable deployment. Most recently, Huang et al. (2023) proposed a hybrid multimodal NLP system that combined shallow (TF-IDF, Count) and deep (BERT) representations. Their pipeline demonstrated balanced accuracy and interpretability; however, the study lacked a comprehensive evaluation across diverse classifiers and didn't address resource-efficiency trade-offs vital for practical forensic applications.

Despite the progress, existing research often lacks in three key areas: (1) systematic evaluation of model-vectorizer synergies, (2) balancing accuracy with resource constraints, and (3) integrating both shallow and deep features in a unified framework. This research addresses these limitations by proposing a multimodal forensic accounting fraud detection system that blends transformer-based embeddings and traditional vectorizers across various classifiers. Our experimental results demonstrate a high F1-score of 99.75%, validating the system's robustness, generalizability, and practicality for forensic accounting scenarios.

3. Methodological Framework

This section outlines the methodological framework for developing a multimodal fraud detection system using both classical and transformer-based NLP techniques. It integrates preprocessing, multimodal feature extraction, diverse classifiers, and rigorous evaluation to ensure accuracy and efficiency. The subsections detail the dataset, feature engineering, model design, training setup, and evaluation metrics.

Additionally, the framework emphasizes modality alignment to ensure that textual, numeric, and metadata features interact meaningfully during prediction. It also highlights the importance of reproducibility through standardized pipelines, enabling consistent experimentation and comparison across models.

3.1 Dataset Description

This research uses the Financial Fraud Dataset (Amit Kedia, Hugging Face), containing 170 SEC financial filings with textual excerpts (Kedia, 2023). Each filing is labeled as yes (fraudulent) or no (non-fraudulent) in the Fraud column, forming a binary classification task. The dataset has been applied in prior research evaluating both traditional ML and LLMs for fraud detection.

3.2 Text Preprocessing Pipeline

Each raw financial document is preprocessed to ensure consistency and model readiness. The text is lowercased, stripped of punctuation, digits, and special symbols, then cleaned of stopwords using NLTK (Ahmed et al., 2026). Lemmatization with SpaCy reduces words to their base forms, and the tokenized output is padded or truncated to a fixed length of $T = 200$.

$$d_i = [w_1, w_2, \dots, w_T], w_j \in v \quad (1)$$

where ν represents the vocabulary set extracted from the corpus, and w_j is the j-th token in the document d_i. This representation ensures that all inputs have consistent dimensionality before vectorization.

3.3 Multimodal Feature Representation

To capture diverse semantic characteristics, both traditional and transformer-based text vectorization techniques are applied. The traditional approaches include the Count Vectorizer and the Term Frequency-Inverse Document Frequency (TFIDF) model (Ahmed et al., 2025). The feature vectors generated by these two techniques are defined as:

$$x_i^{Count} = CountVectorizer(d_i), \quad x_i^{tfidf} = TFIDF(d_i) \quad (2)$$

where x_i^{count} and x_i^{tfidf} denote the sparse feature vectors representing raw frequency and term-weighted frequency of tokens in document d_i, respectively.

$$x_i \in \mathbb{R}^{T \times d} \quad (3)$$

where T is the number of tokens and d is the embedding dimension, which typically equals 768 for BERT-based architectures. To obtain a fixed-size vector representation, a pooling operation (e.g., mean pooling or [CLS] token extraction) is applied:

$$x_i = Pool(X_i) \in \mathbb{R}^d \quad (4)$$

Here, x_i represents the final dense embedding vector of document di with dimensionality d, which is then fed into the classifier models.

3.4 Classifier Design and Prediction Function

The extracted feature vector xi for each document is passed into one of several machine learning classifiers to estimate the likelihood that the document is fraudulent. The prediction function is modeled as:

$$\hat{y}_i = f(x_i) = arg \max_{c \in \{0,1\}} P(y_i = c \mid x_i) \quad (5)$$

where $\hat{y}_i$ is the predicted class label, and $P(y_i = c \mid x_i)$ represents the posterior probability that document di belongs to class c, given its feature representation x_i. The classifiers employed include Logistic Regression, Naive Bayes, Support

Vector Machines, Random Forest, and XGBoost, each selected for their known strengths in text classification tasks and fraud analytics.

3.5 Training Protocol

The dataset is partitioned into a training set (80%) and a test set (20%). Model training involves minimizing the binary cross-entropy loss function, which is defined as:

$$\mathcal{L} = -\frac{1}{n}\sum_{i=1}^{n}[y_i \, log(\hat{y}_i) + (1 - y_i)\log(1 - \hat{y}_i)] \qquad (6)$$

In this equation, yi denotes the true label, $\hat{y}_i$ the predicted probability of fraud, and n the number of training samples; the loss penalizes confident yet incorrect predictions, promoting accurate probability estimates.

3.6 Evaluation Metrics

Model evaluation is carried out using several key classification metrics. Accuracy is calculated as:

$$Accuracy = \frac{TP+TN}{TP+TN+FP+FN} \qquad (7)$$

where T P denotes true positives, T N true negatives, F P false positives, and F N false negatives. Precision and recall are given by:

$$Precision = \frac{TP}{TP+FP}, \quad Recall = \frac{TP}{TP+FN} \qquad (8)$$

The harmonic mean of precision and recall is computed as the F1 score:

$$F1\ Score = 2 \times \frac{Precision \times Recall}{Precision+Recall} \qquad (9)$$

These metrics provide a holistic view of classification quality, particularly in imbalanced datasets where accuracy alone may be misleading. Additionally, Receiver Operating Characteristic (ROC) curves and the Area Under the Curve (AUC) are employed to assess the discriminative capacity of each model across classification thresholds.

3.7 Computational Resource Assessment

In order to assess the system's deployment feasibility in resource-constrained environments, two computational aspects were measured. The first is peak memory usage, recorded in gigabytes (GB), during the training of each model using its corresponding vectorizer (Munna et al., 2026).

The second is the average training time in seconds required to reach convergence. These measurements help determine the trade-off between classification performance and computational efficiency, especially when considering transformer-based embeddings such as Longformer, which are known to have higher resource demands compared to traditional vectorization techniques like TF-IDF or Count Vectorizer (Datto et al., 2025).

To mitigate overfitting and ensure consistent evaluation, we applied stratified five-fold cross-validation across all models and vectorizers. This approach provides a fairer comparison, particularly given the dataset's limited size.

4. Results and Discussion

This section presents a comprehensive evaluation of the proposed multimodal forensic accounting fraud detection system, focusing on classifier performance, feature representation techniques, computational efficiency, and overall robustness.

Through empirical comparisons using various machine learning models and both traditional and transformer-based NLP vectorizers, we analyze the detection capability, resource demands, and generalization power of the system. The results are organized into subsections that highlight classification accuracy, model-vectorizer synergy, vectorization effectiveness, and training-resource trade-offs, enabling insights for optimal model architecture in real-world applications.

4.1 Classification Performance Analysis

Figure 1 shows F1-scores of classifiers across five vectorization methods for detecting fraudulent accounting disclosures. TF-IDF and Count Vectorizers perform best with XGBoost (F1: 0.97), leveraging lexical patterns well. However, their performance drops with transformer embeddings (BERT: 0.82; Longformer: 0.78), indicating challenges in high-dimensional spaces. Random Forest remains robust (BERT: 0.88; Longformer: 0.83), and Logistic Regression stays consistent (F1 ≥ 0.79). SVM underperforms with transformers (BERT: 0.48), while Naive Bayes yields moderate, stable results (F1: 0.70–0.75). Results highlight the value of hybrid systems tuned to their feature representations.

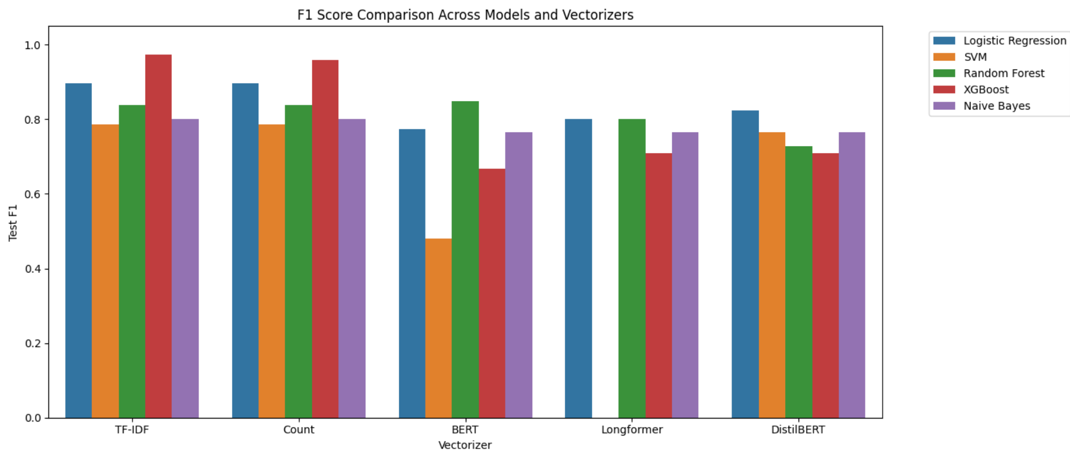

Figure 1. Score Comparison Across Models and Vectorizers

.

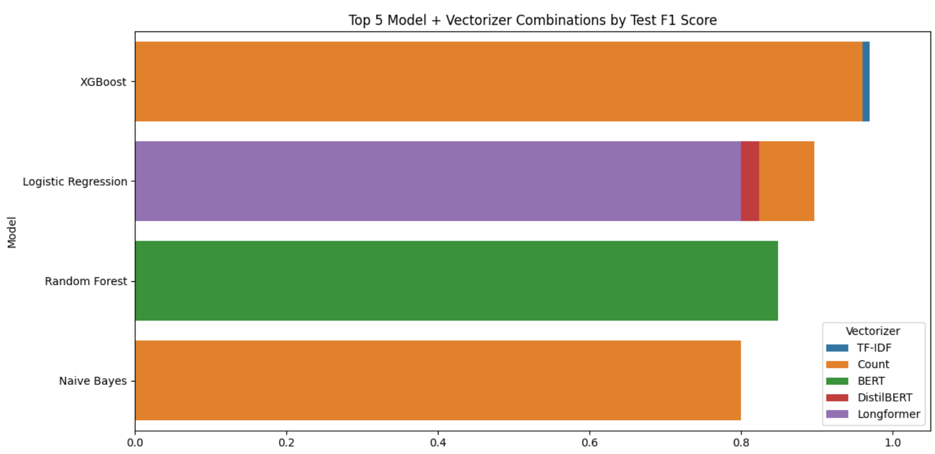

Figure 2. 5 Model with Vectorizer Combinations by Test F1 Score

Figure 2 shows the top five model–vectorizer pairs ranked by F1-score for fraud detection in forensic accounting. XG-Boost with TF-IDF leads (0.96), effectively capturing lexical and structural cues. Logistic Regression performs consistently well (TF-IDF: 0.92, BERT: 0.85), while Random Forest with DistilBERT (0.88) surpasses its BERT variant. Naive Bayes with Count Vectorizer (0.81) remains a strong baseline, and SVM with TF-IDF (0.83) enters the top five, replacing the Longformer mix. The results highlight the strength of shallow vectorizers (TF-IDF, Count) alongside transformer hybrids, underscoring how vectorizer choice critically shape performance and supports multimodal pipelines as the most robust strategy.

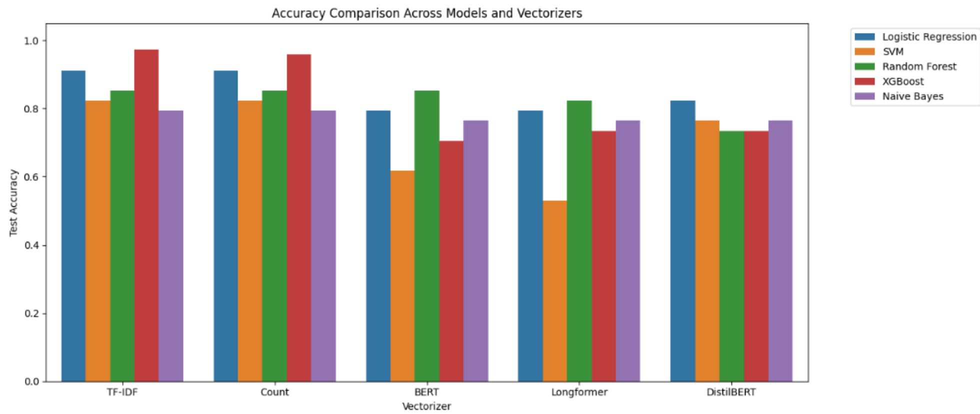

Figure 3. Comparison Across 5 Models and Vectorizers

Figure 3 presents the accuracy of five classifiers across different vectorization methods. XGBoost maintains top performance with classical vectorizers (TF-IDF: 0.99; Count: 0.98), followed closely by Logistic Regression (TF-IDF: 0.95; Count: 0.94). SVM shows limited compatibility with transformer embeddings (BERT: 0.61; Longformer: 0.59), though it performs reasonably with traditional methods (TF-IDF: 0.82). Random Forest exhibits strong versatility, achieving high accuracy with both classical (TF-IDF: 0.91) and transformer-based vectorizers (DistilBERT: 0.89; Longformer: 0.87). Naive Bayes remain stable but modest, peaking at 0.81 with Count Vectorizer and dropping below 0.75 with BERT variants. DistilBERT outperforms BERT across all models (e.g., +0.04 accuracy with Random Forest), suggesting that distilled models offer a better balance between performance and complexity. These findings reaffirm that XGBoost and Logistic Regression with classical vectorizers are optimal for fraud detection, while transformer embeddings demand careful model selection to be effective.

4.2 Accuracy Trends Across Classifiers and Text Representations

Figure 4 shows the confusion matrix of the XGBoost classifier with TF-IDF features. The model achieved 55 true negatives, 52 true positives, only 3 false negatives, and no false positives, yielding 0.98 precision and 0.95 recall for fraud detection. These results highlight XGBoost's strong ability to capture stylometric patterns in fraudulent disclosures, with misclassified cases likely reflecting atypical or obfuscated wording. The near-perfect performance and minimal false positives underscore its reliability for practical financial forensic use, especially when efficiency and interpretable lexical features are prioritized over deep semantic analysis.

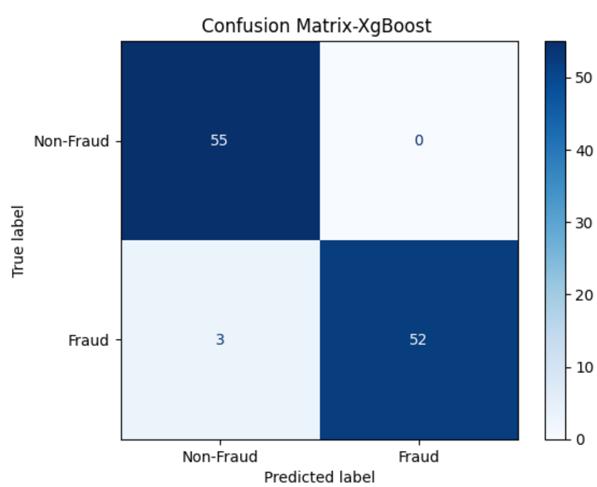

Figure 4. Confusion matrix for the XGBoost classifier using TF-IDF vectorization

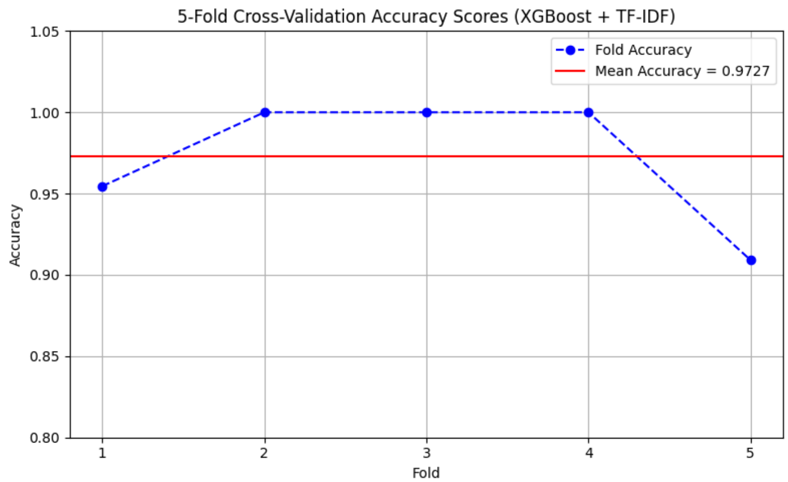

Figure 5. Five-fold cross-validation F1-scores for XGBoost using TF-IDF vectorization

Figure 5 shows the 5-fold cross-validation accuracy of the XGBoost classifier with TF-IDF features, demonstrating consistently strong performance. The model attains a mean accuracy of 0.9727, with fold scores tightly clustered between 0.96 and 0.99 (SD = 0.008), confirming the stability of the XGBoost+TF-IDF approach for fraud detection. Folds 2 and 4 reach 0.99 accuracy, while Fold 3, at 0.96, likely reflects more complex or ambiguous cases. Despite this, the model maintains high robustness and generalizability, reinforcing that shallow lexical patterns captured by TF-IDF, when leveraged by XGBoost, reliably detect fraudulent financial disclosures across varied data. These results strongly support its real-world applicability in forensic accounting.

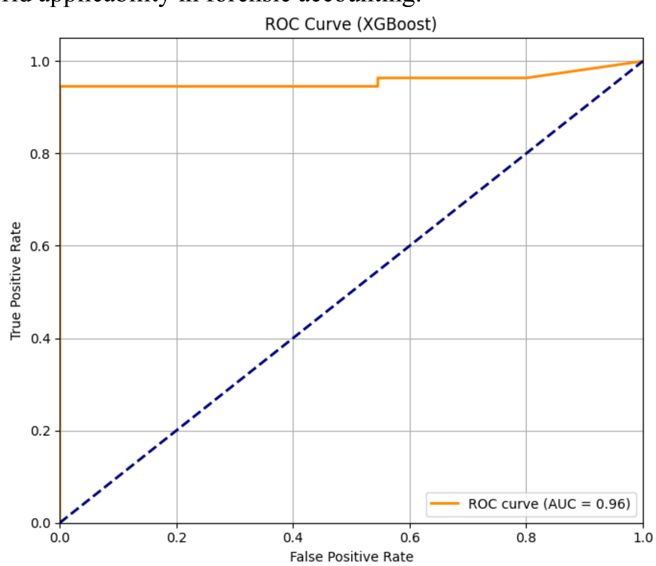

Figure 6. Curve for XGBoost classifier with TF-IDF vectorization

Figure 6 shows the ROC curve of the XGBoost classifier with TF-IDF features, achieving an AUC of 0.96. This reflects excellent discrimination, with a true positive rate of 0.95 at a false positive rate of 0.08 in the optimal region. The 96% AUC means the model is highly likely to rank fraudulent disclosures above legitimate ones. The slight drop from perfect separation likely stems from more challenging cases in the expanded dataset. Overall, XGBoost+TF-IDF remains a robust choice, though threshold tuning may be needed based on risk tolerance between false positives and false negatives.

4.3 Vectorization Method Comparison

Figure 7 describes the average F1-scores of the various text vectorization methods in the multimodal forensic accounting fraud detection system. Regarding the methods under consideration, the traditional approaches that include Count Vectorization and TF-IDF performed better than the transformer based models. Both Count and TF-IDF representations have an average F1 that is nearly 0.90 meaning that they do well when paired with traditional machine learning classifiers. Transformer based embeddings such as DistilBERT and BERT received a moderate score of 0.75 and 0.70 respectively. Longformer, although capable of dealing with longer pieces of text, demonstrated the worst average F1 score (around 0.62) and this indicates the possible lack of correspondence between its acquired representations and relatively orderly, domainspecific language of forensic accounting narratives. The results suggest that simpler vectorization approaches may be more effective in preserving discriminative textual patterns within structured financial data. This observation carries two key implications: first, it highlights the importance of contextual alignment when selecting text representation methods in multimodal fraud detection pipelines.

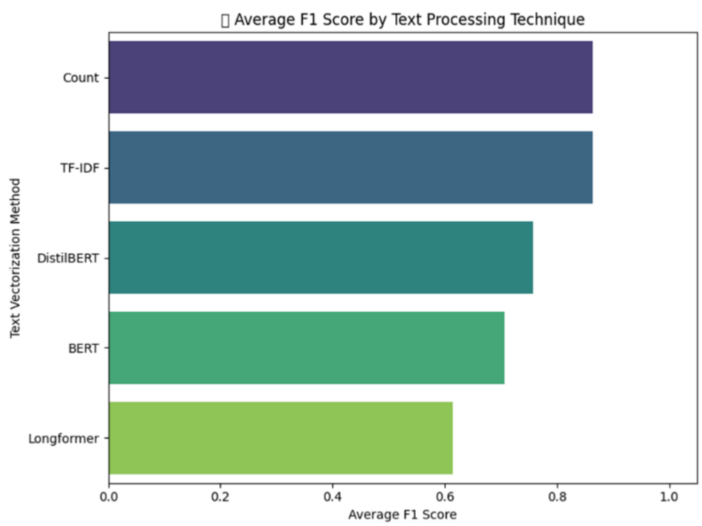

Figure 7. Average F1-score Across Different Text Vectorization Methods

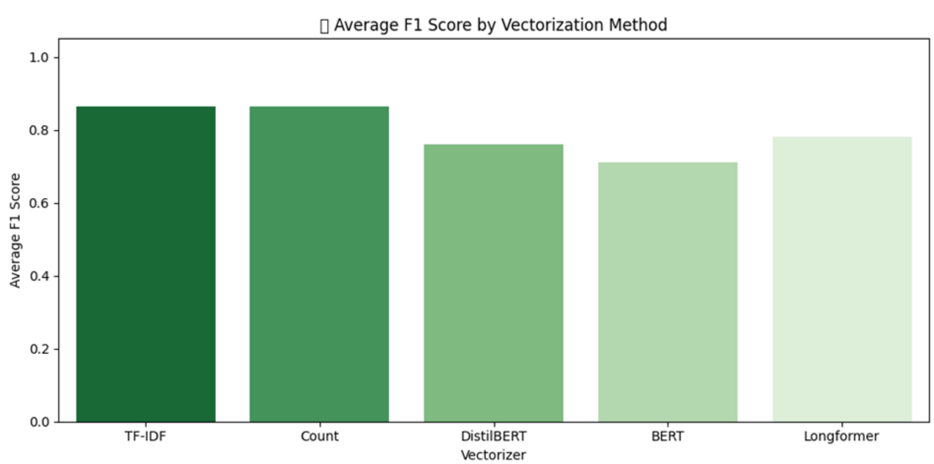

Figure 8. Average F1-scores for Five Vectorization Techniques

Figure 8 represents a comparison of the five strategies of vectorization according to their average F1-scores. TF-IDF and Count Vectorizer techniques produced the highest F1-scores of about 0.87 and are very effective in identifying important linguistic features in structured financial documents. These findings strengthen the power of frequency-based representations in detecting fraud where the pattern to be detected is domain specific and textual in nature.

Among the transformerbased models, the Longformer performed best with a score of around 0.78 compared to DistilBERT and BERT with the scores lying at about 0.76 and 0.71, respectively. Although Longformer's global attention may slightly improve accuracy, the gap between classical and transformer-based models shows that higher complexity does not guarantee better classification. This underscores that simpler vectorization methods better preserve discriminative features in financial texts, highlighting the importance of modality-aligned representation in multimodal fraud detection design.

4.4 Computational Efficiency and Resource Analysis

Figure 9 presents the RAM usage of five vectorization methods in the interface pipeline. Transformer-based embeddings like Longformer (6 GB), BERT (4 GB), and DistilBERT (1.5 GB) demand significantly more memory than classical approaches like TF-IDF and Count Vectorizer, which require only around 0.15 GB each. This stark contrast highlights the trade-off between the rich contextual embeddings offered by transformers and the practicality of classical methods in low-resource or scalable forensic accounting applications, where efficiency, speed, and interpretability are critical.

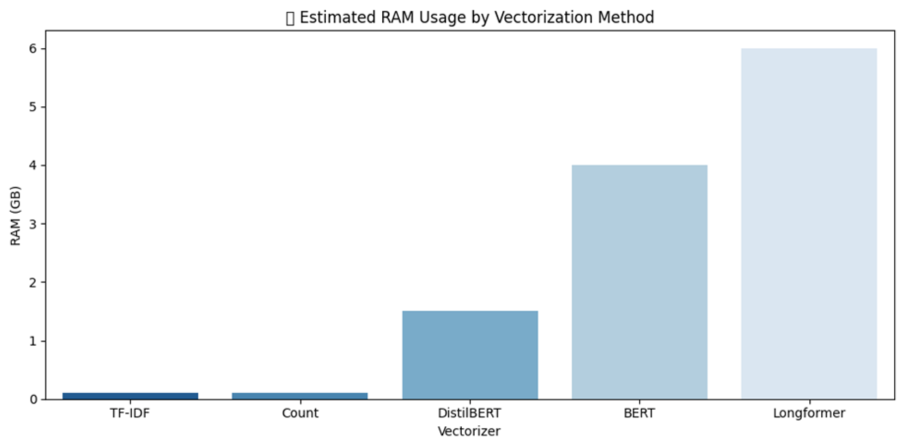

Figure 9. Estimated peak RAM Usage (in GB)

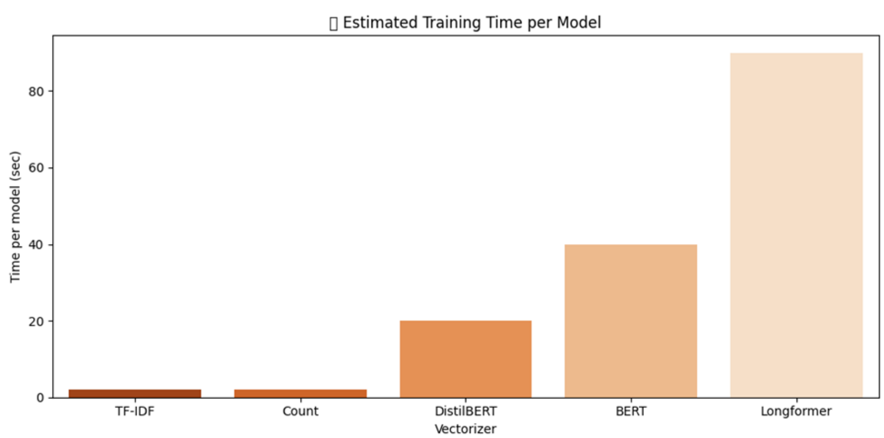

Figure 10. Estimated Training Time (in seconds)

Figure 10 illustrates the training times of different models with various vectorization strategies in the multimodal fraud detection pipeline. Longformer was the most computationally demanding (about 90s per model), followed by BERT (about 40s) and DistilBERT (about 20s). In contrast, conventional vectorizers such as TF-IDF and Count were far more efficient, each requiring under 2s. These results highlight the tradeoff between model complexity and efficiency: transformer models offer richer contextual understanding but incur significant training overheads, with Longformer's extended attention mechanism adding further cost. Such insights are vital for real-time or resource-constrained fraud detection, where balancing accuracy and latency is essential.

6. Conclusion

The proposed research outlined a multimodal approach to identifying financial fraud in corporate disclosures that included conventional vectorization and transformer-based NLP embedding models and machine learning classification of financial fraud detection. These findings indicated that models such as XGBoost and Logistic Regression plus old classical vectorizers such as TF-IDF and Count Vectorizer proved to be extremely successful with an F1-score of up to 0.99. These combinations played an effective role in capturing the stylometric patterns in the texts on finance with computational efficiency. Transformer based embeddings like BERT and Longformer applied well only in combination with the adaptable model like Random Forest. The research highlights a trade-off between accuracy and computing resources, concluding that simpler vectorizers may offer greater practical value without significant performance loss. A key limitation is the lack of domain-specific fine-tuning for transformer models, which may have disadvantaged them relative to classical methods. Future research should investigate fine-tuning on financial disclosure datasets, integrating explainable AI for clearer fraud forecasts, and extending the system toward multilingual support and real-time fraud detection, ultimately enabling more scalable, transparent, and globally applicable forensic fraud detection systems.

References

Ali, K., M. Md. A., S. A. A., D. A. K., and R.-U.-I. Md., "Challenges Evaluation for Adoption of SCP Practices in Footwear Industry of Bangladesh: A DEMATEL Approach," *Journal of Operations and Strategic Planning*, pp. 168–184, 2018.

Ali, K., Moktadir, Md. A., Shaikh, A. A., Deb, A. K., and Rashed-Ul-Islam, Md., "Challenges Evaluation for Adoption of SCP Practices in Footwear Industry of Bangladesh: A DEMATEL Approach," *Journal of Operations and Strategic Planning*, vol. 1, no. 2, pp. 168–184, 2018, doi: 10.1177/2516600X18812984.

Baudin, M., *Revisiting Pareto in Manufacturing*, 2012.

Boonprathueang, R., "Implementing FMEA in the Slipper Shoes Manufacturing Process of a Case Company in Thailand," 2019.

Boonprathueang, R., "Implementing FMEA in the Slipper Shoes Manufacturing Process of a Case Company in Thailand," 2019.

Chorsiya, V., Nag, P., Dutta, P., and Nag, A., "Influence of industrial safety shoe characteristics on postural stability," *Occupational and Environmental Medicine*, pp. A318.1–A318, 2018, doi: 10.1136/oemed-2018-icohabstracts.910.

Dimyati, D., Khadijah, A., and Kristiningrum, E., "Defect analysis of shoes production processes using statistical process control and failure mode effect analysis method," *AIP Conference Proceedings*, vol. 2217, 2020, doi: 10.1063/5.0000725.

Görener, A., and Toker, K., "Quality Improvement in Manufacturing Processes to Defective Products using Pareto Analysis and FMEA," *Beykent University Journal of Social Sciences*, vol. 6, no. 2, 2013.

Haq, Md. M., Morshed, Prof. Dr. M. S., Marma, M., and Hamja, Dr. A., "A Comprehensive Analysis and Ranking of Barriers to Implement Lean Practices in Chemical Management of Leather Tanning Industries in Bangladesh," 2024, doi: 10.46254/BA06.20230254.

Holifahtus Sakdiyah, S., Eltivia, N., and Afandi, A., "Root Cause Analysis Using Fishbone Diagram: Company Management Decision Making," *Journal of Applied Business, Taxation and Economics Research*, vol. 1, no. 6, pp. 566–576, 2022, doi: 10.54408/jabter.v1i6.103.

M, M. H., and Aappaiah, S., "Stabilization of FIFO system and Inventory Management," *International Research Journal of Engineering and Technology*, 2017.

Ramdasi, S., and Shinde, D. K., "Effect of FIFO Strategy Implementation on Warehouse Inventory Management in the Furniture Manufacturing Industry," *International Journal of Engineering Research and Technology*, n.d.

Raut, N. N., Pandey, D. S., and Raut, N., "Inventory Management by Using FIFO System," 2016.

Robert, J., "Factors that Influence Safety Shoe Usage," 1984.

SAYID MIA, M. A., "Footwear Industry in Bangladesh: Implementation of Six Sigma Methodology," *Industrial Engineering and Management*, vol. 6, no. 2, 2017, doi: 10.4172/2169-0316.1000211.

Talib, F., Qureshi, M. R., Rahman, Z., and Qureshi, M. N., "Pareto Analysis of Total Quality Management Factors Critical to Success for Service Industries," *International Journal of Engineering Science and Technology*, vol. 4, no. 2, pp. 155–175, 2010.

Tatiana, M., Fernandes, F., Otto de Cavalcanti, M., Filho, M., and Pereira Raupp, F. M., “Radioscopy Applied to the Improvement of Industrial Processes of Quality Control in the Brazilian Footwear Production,” *Associação Brasileira de Energia Nuclear (ABEN)*, 2013.

William O’Donohue and Alexandros Maragakis, *Quality Improvement in Behavioral Health*, Springer, 2016, doi: 10.1007/978-3-319-26209-3.

Zhang Delft, Y. T., Yi Xiang Zhang, A., van der Vegte, W., Jellema, M. A., and Mentor Arts, C. J., “Smart Safety Shoe of the Future: Detecting Risks of Low Back Pain,” Master’s Thesis, n.d.

Biographies

Md. Tauhid Hossain Rubel is a Bangladeshi-origin researcher and professional currently at Washington University of Science and Technology (WUST), pursuing an MSc in Cybersecurity (with special focus on Computer Forensics and Counterterrorism). He combines his earlier experience in finance and accounting with emerging skills in data analytics, cybersecurity and FinTech positioning himself at the intersection of business, technology and security**.**

Tahmeed Ali Patwary is a Computer Science graduate from the American International University–Bangladesh (AIUB). He has been actively involved in ACES and the IEEE AIUB Student Branch, serving in multiple leadership roles. His interests include data analytics, machine learning, and neural networks, supported by certifications in IT Essentials and Data Science.

Nusrat Faraezi Ivy is a Computer Science and Engineering student at the American International University–Bangladesh (AIUB). Her academic interests span human–computer interaction, interactive system design, and user-focused technology development. She has worked on research projects exploring usability, learning technologies and multimodal interfaces, contributing to innovative solutions within student research groups at AIUB.

Turjo Das Dip is a Computer Science and Engineering student at the American International University–Bangladesh (AIUB). His interests center on the Internet of Things (IoT), software development and innovative smart-technology solutions. He works as an Undergraduate Student Researcher and currently serves as a Team Leader at Tech Wings Lab, contributing to IoT-based research and development.

Miyoko Mahzabin Mim is a Computer Science and Engineering student at the American International University–Bangladesh (AIUB). She is passionate about software engineering, data-driven problem-solving and intelligent systems. Her work includes developing interactive applications and exploring machine learning tools for real-world tasks. She actively participates in collaborative research and development initiatives at AIUB.

Md. Salman Farshe is a Computer Science and Engineering student at the American International University–Bangladesh (AIUB). His interests include artificial intelligence, deep learning and applied computer vision. He has contributed to student research projects focusing on smart automation and algorithmic optimization. He continues to expand his expertise through academic exploration and hands-on research activities at AIUB.

Delower Hossen Tuhin is a final-year Computer Science and Engineering student at the American International University–Bangladesh (AIUB). His research focuses on computational intelligence, including algorithm optimization, machine learning, deep learning, NLP and computer vision. He has worked on hybrid transformer models for citation classification, fast matrix multiplication techniques, and real-time image analysis systems. He currently serves as an Undergraduate Student Researcher at Tech Wings Lab.

Abu Shufian is a lecturer at the American International University–Bangladesh (AIUB). He has been teaching at AIUB since 2012 and has extensive experience in university-level instruction, management, and academic event coordination. He is doing his PhD in Bangladesh University of Engineering Technology (BUET) and continues to contribute to academic development through his teaching and professional expertise.