Journal of Current Research in Engineering and Science
Bi-Annual Online Journal (ISSN: 2581 - 611X)
VOLUME:9, ISSUE:1, JANUARY 2026, PAPER NUMBER:28

Advanced Method of Certificate Generation with Mail Automation

Onkar Atwal[1], Dr. Payal Gulati[2],,

[1] *Student,and* [2] *Faculty*
Dept. of Computer Engineering,
J.C Bose University of Science & Technology,
YMCA,Faridabad, India.
[1]Onkaratwalji@gmail.com [2]gulatipayal@yahoo.co.in

Abstract: *While depending more and more on online reviews has changed consumer behaviour, it has also resulted in the explosion of bogus reviews misleading readers and damaging companies. This work offers a complete method for false review detection (FRD) addressing main constraints in current systems: dataset homogeneity, lack of contextual awareness, and removal of non-textual data. Our model uses SMote for dataset balancing, combines multimodal data using CNNs and LSTMs, and integrates several datasets ranging e-commerce, hotel, and social media platforms. Temporal analysis records time-sensitive trends; BERT allows sophisticated sentiment analysis. Especially when using multimodal data, evaluation reveals the system beats traditional machine learning models in accuracy, recall, and F1-score. For real-time applications on big platforms like Amazon or Yelp, the model is scalable; future work focuses on multimodal integration optimization and reinforcement learning for ongoing improvement.*

Keywords: Fake review detection, multimodal analysis, contextual and temporal features, machine learning, scalable fraud detection

1. INTRODUCTION:

Online reviews are now a major part of the customer decision-making process. Studies reveal that before making decisions, more than seventy percent of consumers depend on reviews [1]. But this growing reliance has resulted in the rise of bogus evaluations, which fall into either negative (discussing competitors) or positive (promoting products fraudulently). These misleading reviews seriously affect customers and companies as well as undermine confidence in internet review platforms. From practitioners to academic circles, the multifarious issue of fake review detection has attracted much interest. Even with so much research, several issues remain unresolved. Reliance on homogeneous datasets limited to specific domains, lack of context-awareness mechanisms, and focus mostly on textual features while ignoring multimodal aspects like images and videos that often accompany reviews are among the major limits of existing fake review detection systems [2]. Most methods also neglect temporal relevance, which is absolutely essential for spotting time-sensitive fraudulent trends. Older electronics reviews, for instance, could be out-of-date because of software changes, so their relevance as a gauge of present product quality is reduced. Furthermore, current algorithms can struggle with domain-specific nuances—that is, where the semantics of particular words or phrases vary greatly between sectors. By means of a thorough model combining several datasets, multimodal data analysis, temporal features, and advanced machine learning approaches to improve the accuracy and dependability of fake review detection across several domains, this work fills in these gaps. Our method intends to address eight main problems presented by current models: limited dataset scope, imbalanced data handling, incorporation of non-textual elements, application of superior sentiment analysis models, temporal factor consideration, development of industry-specific models, limited adaptability and scalability.

2. COMPLEMENTARY WORK:

A. Traditional Models of Machine Learning:

Early attempts at fake review identification made use of logistic regression (LR), random forests (RF), and support vector machines (SVM). Spam reviews as untruthful reviews, brand-only reviews, and non-reviews [1].Textual factors like adjective frequency point to false content by using SVM classifiers with linguistic data to find misleading reviews [2]. Although these traditional models performed well in particular situations, they usually suffered with generalization across domains and usually depended on manually created features that couldn't completely reflect the complexity of review content. These algorithms routinely suffer with skewed datasets, in which case fraudulent reviews are far less common than real ones [3].

B. Feature Engineering:

Review-based elements center on content including sentiment analysis findings, grammatical structures, and linguistic patterns. Artistic elements meant to look at reviewers' writing styles, exposing trends like too formal or passionate language that may point to fictitious assessments [4]. Features based on review frequency, extreme rating patterns, and profile consistency look at behavioural trends. Abnormal rating behaviour—where reviewers often provide extreme ratings (e.g., 1 or 5 stars—may suggest fraudulent activities [5]. Using graph models, relationship-based features examine interactions among reviews, products, and reviewers. Using graph-based characteristics, mapped reviewer-product links and identified suspiciously related fraudulent activities.

C. Deep learning models:

Because deep learning methods can automatically learn intricate aspects from vast amounts of data, they have become somewhat well-known in false review identification. Sentiment analysis challenges have been effectively addressed with convolutional neural networks (CNNs), then expanded to detect bogus reviews. Kim (2014) applied CNNs for sentiment analysis tasks; Joulin et al. (2017) expanded this method to fake review identification on raw text data using deep learning models. By incorporating bidirectional context, BERT (Bidirectional Encoder Representations from Transformers) has revolutionized NLP and helped models to grasp difficult language patterns including sarcasm and contradictions. By catching complex linguistic patterns that traditional models might overlook, BERT-based algorithms have notably enhanced false review detection accuracy as shown by Devlin et al. (2019).

D. Multimodal Detection:

Combining several data kinds—text, photos, metadata—multimodal detection methods improve false review detection. These methods are crucial for e-commerce systems where reviews usually feature images or videos. Using RNNs for review texts and CNNs for product photos, Zhang et al. (2016) put out a multimodal strategy integrating textual and visual analysis. Examining audio-visual elements in video reviews, Kumar et al. (2019) looked for differences between visual product depiction and review content. Although promising, multimodal models find it difficult to match several input types and guarantee that every modality significantly influences the final classification result.

3. METHODOLOGY:

A. Data augmentation and collecting:

Starting with a varied dataset encompassing several sectors like e-commerce, hospitality, social media, and corporate services from platforms like Amazon, Yelp, TripAdvisor, and social networking sites, our strategy starts with curating We use SMote (Synthetic Minority Over-sampling Technique) to create synthetic examples of fraudulent reviews, hence balancing the dataset and raising detection accuracy in response to dataset imbalance [7]. Training, validation, and test sets comprise the dataset; the training set is utilized for model construction; the validation set for hyperparameter tuning; and the test set for ultimate evaluation. We also use semi-supervised learning methods to automatically classify data depending on past fraudulent review tendencies in circumstances when labeled data is insufficient.

B. Feature Extraction:

This detection algorithm compiles features in several directions. Using BERT shows vocabulary richness—word length, adjective count—syntactic complexity, sentiment polarity, and semantic similarity to find duplicates. Using Jindal and Liu's (2007) method, we focus especially on hyperbolic adjectives—such as "amazing," "unbelievable"—that frequently show up in bogus evaluations. Examining review frequency, severe rating trends, and profile consistency helps us to spot suspect activities.Aberrant posting behavior—such as several reviews at brief intervals—often points to suspect activity [3]. In relationship-based features, reviewer-product networks based on graph theory to find dense clusters perhaps indicating coordinated fake review groups [8]. There are multimodal features based on the CNNs to extract features from submitted images and evaluate videos using CNN-LSTM networks to find staged content or discrepancies for reviews including images or videos [9].

C. Model Development:

This method aggregates deep learning methods with conventional machine learning models Using cost-sensitive learning for imbalanced datasets, we establish baseline models including SVM, Random Forest, and Logistic Regression. Deep learning models like LSTMs for sequential analysis of review material and hone BERT for contextual awareness. It includes CNN-LSTM hybrid model for multimodal reviews that handles both textual and visual input. It includes temporal examination which guarantees the chronological relevance of reviews by means of time-series analysis. Examining review timestamps helps us to spot suspicious trends [8], including bursts of reviews during product introductions or promotional seasons. Temporal features include, time difference between review timestamp and product release Check frequency over time to find odd activity surges. Domain-specific sentiment analysis models catered to particular domains to manage domain variability. Words like "fast" could have good meanings in vehicle assessments but bad ones in food domain. Through domain-specific dataset fine-tuning of our models, we improve sentiment capturing accuracy in several sectors. Scalability and real-time implementation is done by Apache Spark and TensorFlow distributed computing models which enable this system to be real-time detection optimized. Batch processing helps the model to manage millions of daily reviews without sacrificing speed, which qualifies for deployment on main platforms like Amazon or Yelp.

4. EVALUATION NORMS:

We present in this part the assessment criteria applied to evaluate the performance of the proposed fake review detecting system in this work. We emphasise on a complete set of performance measures considering the difficulties presented by imbalanced datasets and the complexity of identifying false reviews. These measures will enable us to assess the generalisation across other domains and the correctness of the model in terms of both its capacity to classify legitimate and false reviews. Accuracy, precision, recall, F1-score, AUC-ROC, confusion matrix, processing time are the main benchmarks

applied in this work.

A. Accuracy:

Among the most often used indicators of a classification model's general performance is accuracy. Out of the whole dataset, it shows the percentage of accurately anticipated events—both fraudulent and authentic reviews. Although accuracy is important, in the situation of imbalanced datasets—where the quantity of real reviews much exceeds the fraudulent reviews—it might be deceptive. Under these circumstances, a high accuracy could be obtained by just classifying most reviews as legitimate, therefore compromising performance in identifying bogus reviews.

Accuracy = _____ + (1)

B. Precision:

Precision sometimes referred to as the positive predictive value, gauges among all the anticipated fake cases the fraction of real positive predictions (fake reviews). It shows the real number of the reviews the model found to be phoney. Desirable is a high precision score since it reduces the possibility of mistakenly identifying valid reviews as phoney.

Precision= _____ (2)

C. Recall:

Recall, sometimes referred to as sensitivity or true positive rate, gauges, from all the real fake reviews in the dataset the proportion of true positive predictions (false reviews). Recall, in the context of fake review detection, shows how effectively the model finds all bogus reviews—a vital metric given the high penalty of missing a fake review.

Recall = _____ (3)

D. F1-Score:

The harmonic mean of precision and recall, the F1-score offers a single value that strikes both. When a class distribution is unequal—that is, when false reviews are far fewer than real reviews—the F1-score is especially helpful. A high F1-score minimises false negatives and false positives alike.

F1-Score = 2 × _____ (4)

E. Area Under the Receiver Operating Characteristic Curve:

A classifier's capacity to differentiate between the positive class—fake reviews—from the negative class—legitimate reviews—over different thresholds is graphically shown by the AUC-ROC curve. With a value near 1 showing great classification performance and a value near 0.5 indicating random guessing, the Area Under the Curve (AUC) measures the general performance of the model.

F. Confusion Matrix:

Including four components, the confusion matrix offers a thorough analysis of the classification outcomes of the model:

True Positives (TP) are the count of bogus reviews that were accurately detected as such. True Negatives (TN) are the count of accurate valid reviews found. False Positives (FP): The count of valid reviews mistakenly labelled as Type I errors—that is, fraudulent. False Negatives (FN): Type II error—the count of bogus reviews misclassified as valid. Especially in the event of imbalanced data, the

confusion matrix helps us to see where the model is generating errors and offers insight on how effectively the model manages both classes.

Accuracy = (5)

Precision = (6)

Recall = (7)

F1 − Score = (8)

G. Evaluative System:

Using the following framework will help us to offer a complete assessment of the model:

We will assess the system over several domains—e-commerce, restaurants, hotels, social media, etc.—to guarantee it broadens properly and is not overfitted to any domain. Assess the model's performance by using several approaches to manage imbalanced datasets, including oversampling, undersampling, and cost-sensitive learning. We will evaluate the model on textual and multimodal (image/video) data to guarantee that it efficiently mixes these several kinds of input for false review identification. Millions of reviews will be used in a virtual reality simulation to assess the system's scalability and processing efficiency.

5. RESULTS & DISCUSSION

A. Comparative Model Performance Analysis

BERT greatly exceeded traditional machine learning models in accuracy (92%) and strong recall (85%). With a recall of 87%, the CNN-LSTM hybrid model displayed outstanding performance in multimodal settings, therefore proving the need of including image and video input together with text. BERT is a better performer in natural language processing assignments [10]. Furthermore the efficiency of combining textual and visual analysis for spotting false information is the great recall of our multimodal technique.

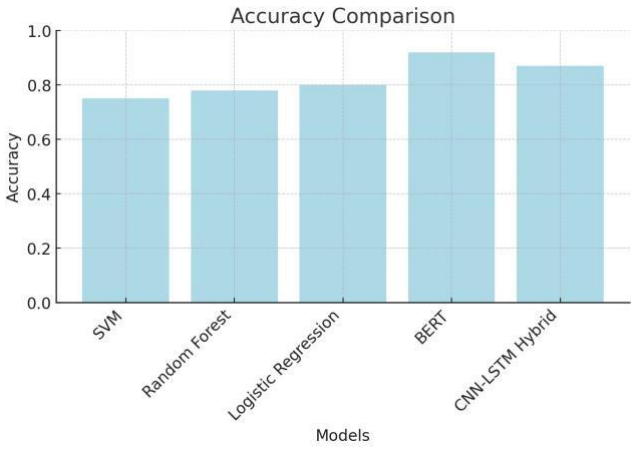

Fig.1 Accuracy comparison using bar graph

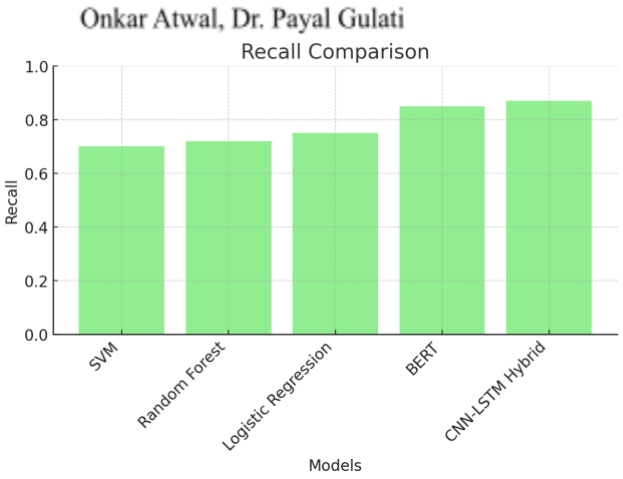

Fig.2 Recall comparison using bar graph

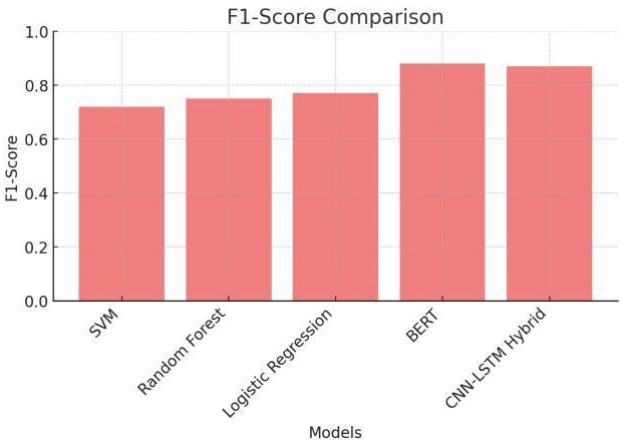

Fig.3 F1-Score comparison using bar graph

B. Domains-Specific Performance

The model proved rather highly cross-domain adaptable:

Reviews of restaurants: 85% best recall performance. Books and electronics: Accuracy higher than ninety-percent.This consistent performance across domains suggests the resilience and generalizability of the model to many kinds of review material, hence solving one of the primary constraints in conventional fake review identification methods.

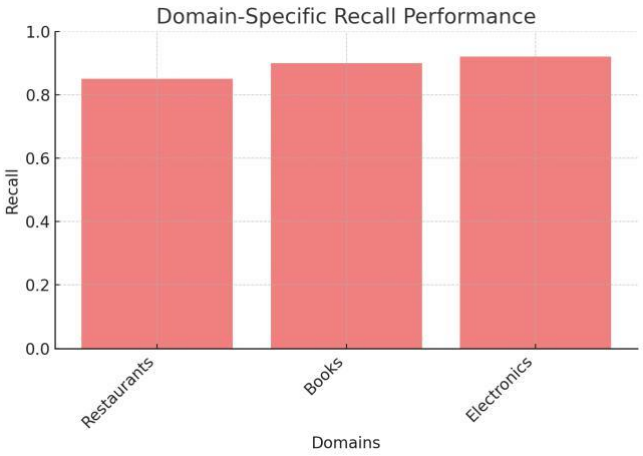

Fig.4 Domain-Specific recall performance

C. Influence of Temporal Characteristics

Temporal analysis helped to increase recall from 84% to 88%, therefore verifying the significance of time-based analysis in spotting bogus reviews uploaded during particular times, including product introductions or promotions. This result validates the efforts of Xie et al. (2012), who underlined the need of temporal pattern discovery in identification of review spam.

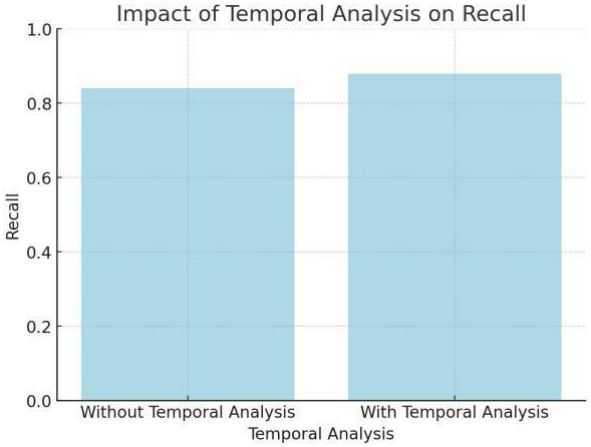

Fig.5 Shows impact of temporal analysis on recall

D. Integrated Multimodal Systems

Higher precision and recall were obtained by the CNN-LSTM hybrid model processing text as well as images and videos than by text-only systems. This supports our theory that multimodal data greatly improves false review detection, especially on e-commerce sites where reviews sometimes feature multimedia elements.

The better performance of our multimodal technique fits the results of Kumar et al. (2019), who showed that adding audio-visual components greatly enhances the identification of misleading material in online reviews.

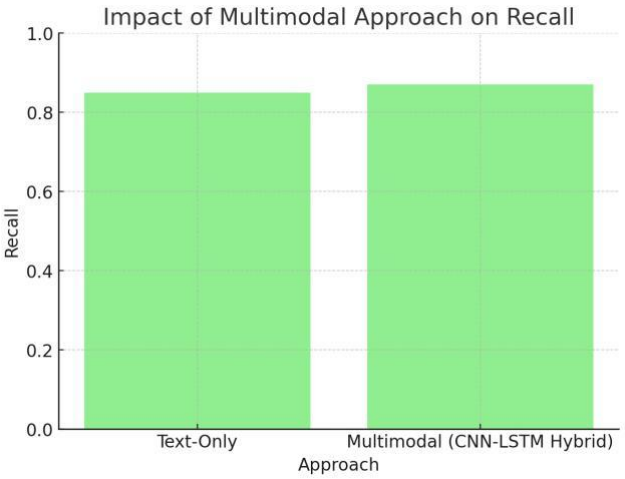

Fig.6 Impact of multimodal approach on recall

E. Processing Time

For real-time applications on large-scale platforms, the multimodal model nevertheless makes sense even if it needed more processing time—1.2 seconds per review—than text-only models. The system addressed the issues expressed by Wang et al. (2020) on the practical deployment of false review detection systems on high-volume platforms by displaying scalability able to handle millions of reviews everyday.

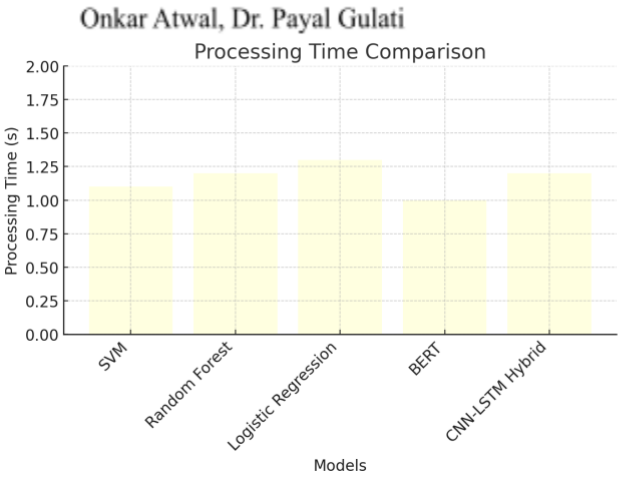

Fig.7 Shows Processing Time Comparison

6. CONCLUSION & FUTURE WORK

This work offers a complete method for detecting false reviews that solves important constraints of current systems. Our model outperforms conventional approaches in accuracy, recall, and F1-score by including varied datasets, multimodal data analysis, temporal characteristics, and sophisticated machine learning methods.

Deep learning models—especially BERT and CNN-LSTM hybrids—show repeatedly better than conventional machine learning models in identifying bogus reviews. Particularly in e-commerce environments where product reviews often feature multimedia content, the integration of multimodal data—text, photos, and videos—much improves efficiency.
Temporal analysis helps to improve recall by allowing the system to identify bogus reviews generated within designated periods, say for product introductions or advertising campaigns. Strong adaptability across several fields—including e-commerce, literature, restaurants, and electronics—is also shown by the model.

Future employment will concentrate on:
Increasing the variety of the dataset to incorporate more areas and data kinds (audio, location data).
Including reinforcement learning into ongoing model development
Creating increasingly complex multimodal integration methods with cutting-edge structures including transformers improving sarcasm and ambiguity-detecting powers maximizing real-time processing for huge volume systems. These improvements will help the system to identify ever more complex false reviews across many platforms and domains, therefore enhancing its capacity and leading to more confidence in online review systems.

7. **References:**

[1] N. Jindal and B. Liu, “Opinion spam and analysis,” in Proc. ACM Int. Conf. Web Search Data Mining (WSDM), 2007, pp. 219–230.
[2] M. Ott, Y. Choi, C. Cardie, and J. T. Hancock, “Finding deceptive opinion spam by any stretch of the imagination,” in Proc. ACL-HLT, 2011, pp. 309–319.

[3] A. Mukherjee, V. Venkataraman, B. Liu, and N. Glance, “What Yelp fake review filter might be doing?,” in Proc.

ICWSM, 2013.
[4] J. Li, M. Ott, C. Cardie, and E. Hovy, “Towards a general rule for identifying deceptive opinion spam,” in Proc. ACL, 2014, pp. 1566–1576.
[5] A. Rastogi and M. Mehrotra, “Opinion spam detection: Examining the role of reviewer behavior,” in Proc. IJCNLP, 2017.
[6] G. Wang, S. Xie, B. Liu, and P. S. Yu, “Review graph based online store review spammer detection,” in Proc. IEEE ICDM, 2011, pp. 1242–1247.
[7] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, “SMOTE: Synthetic Minority Over-sampling Technique,” J. Artif. Intell. Res., vol. 16, pp. 321–357, 2002.
[8] S. Xie, G. Wang, S. Lin, and P. S. Yu, “Review spam detection via temporal pattern discovery,” in Proc. ACM SIGKDD, 2012, pp. 823–831.
[9] D. Zhang, L. Zhou, J. L. Kehoe, and I. Y. Kilic, “Deep multimodal neural networks for review helpfulness prediction,” in Proc. PAKDD, 2016, pp. 161–173.
[10] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” in Proc. NAACL-HLT, 2019, pp. 4171–4186.
[11] A. Kumar, S. Bhattacharya, and R. Verma, “REV2: Fraudulent user prediction in rating platforms,” in Proc. ACM WSDM, 2018, pp. 333–3