As per UGC guidelines an electronic bar code is provided to secure your paper

International Journal for Modern Trends in Science and Technology
Volume 12, Issue 04, pages 206-214.
ISSN: 2455-3778 online
Available online at: http://www.ijmtst.com/vol12issue04.html
DOI: https://doi.org/10.5281/zenodo.19324805

NLP and Machine Learning Based System for Detecting Fake Product Reviews

P. Mohan, SK. Nahida, I.V. Sai Praneeth, M. Lakshmi, M. Venkata krishna, Ch. Sujith

Department of CSE-Artificial Intelligence, PBR Visvodaya Institute of Technology and Science, Kavali, Andhra Pradesh, India.

To Cite this Article
P. Mohan, SK. Nahida, I.V. Sai Praneeth, M. Lakshmi, M. Venkata krishna & Ch. Sujith (2026). NLP and Machine Learning Based System for Detecting Fake Product Reviews. International Journal for Modern Trends in Science and Technology, 12(04), 206-214. https://doi.org/10.5281/zenodo.19324805

Article Info
Received: 28 February 2026; Revised: 18 March 2026; Accepted: 22 March 2026.

KEYWORDS

ABSTRACT

Online product reviews have become a major factor influencing consumer purchasing decisions, making the authenticity of these reviews extremely important. However, the increasing presence of deceptive or spam reviews threatens the credibility of e-commerce platforms. This research introduces an intelligent fake review detection framework that combines Natural Language Processing (NLP) with supervised machine learning techniques to automatically identify misleading feedback.

The proposed approach focuses on extracting linguistic patterns, sentiment characteristics, and structural features from review text to differentiate between genuine and deceptive content. Several classification algorithms were analyzed, with Support Vector Machine (SVM) selected due to its effectiveness in handling high- dimensional textual data. The system is designed to support large-scale deployment by enabling automated moderation of suspicious reviews.

Experimental evaluation on benchmark datasets demonstrates that integrating NLP preprocessing with machine learning classification improves accuracy and reduces manual moderation effort. The framework contributes toward building more trustworthy online review ecosystems and can be extended to real-time e-commerce applications.

INTRODUCTION

The way people make purchasing decisions has changed dramatically with the growth of online shopping platforms. Instead of depending only on advertisements or brand reputation, customers now rely heavily on the experiences shared by other users through digital reviews. These reviews create a sense of community trust and help buyers compare products quickly. However, the same system that empowers consumers also opens the door for misuse, as businesses or individuals may post misleading or fabricated reviews to influence public opinion Fake reviews are not always easy to identify because they often imitate genuine writing styles and emotional expressions. Some reviews exaggerate product quality, while others intentionally damage a competitor's reputation. As the volume of online feedback grows every day, manually verifying each review becomes unrealistic for platform administrators. This challenge has encouraged researchers and developers to explore intelligent automated solutions that can assist in maintaining the authenticity of online review systems.

Recent advances in Natural Language Processing (NLP) and machine learning provide new opportunities to address this problem. By studying patterns such as word usage, sentiment intensity, repetition, and writing structure, computational models can learn to recognize signals that differentiate genuine opinions from deceptive content. Rather than relying on simple keyword filtering, modern approaches analyze deeper linguistic and behavioral characteristics, allowing more accurate detection of suspicious reviews.In this research, an automated fake product review detection framework is proposed to support trustworthy digital marketplaces. The system integrates text preprocessing, feature extraction, and supervised machine learning techniques to evaluate the authenticity of user-generated reviews. The aim is not only to improve classification accuracy but also to create a scalable and practical solution that reduces manual moderation efforts while promoting fair competition among sellers.

As e-commerce continues to expand, ensuring the reliability of online feedback becomes increasingly important for both consumers and businesses. Developing intelligent tools that can identify misleading information helps protect users from biased decisions and strengthens confidence in digital platforms. This study contributes to that goal by designing a structured approach that combines linguistic analysis with machine learning to build a more transparent and dependable review ecosystem.

OBJECTIVE

The primary objective of this study is to develop an automated and intelligent system capable of identifying fake product reviews using opinion mining and machine learning methods. The system aims to process large volumes of textual feedback, extract meaningful linguistic features, and classify reviews into genuine or deceptive categories with improved accuracy.

Another key goal is to minimize the influence of spam content on product ratings by integrating NLP preprocessing with advanced classification algorithms. By implementing supervised learning techniques, the model seeks to enhance decision- making for both consumers and platform administrators. Ultimately, the proposed framework aims to strengthen user trust and support the development of reliable online review platforms. The system also aims to reduce manual moderation effort by providing automated review verification support. Another objective is to design a scalable architecture that can adapt to continuously growing review datasets without performance degradation. Additionally, the study focuses on improving model interpretability so that administrators can understand the reasoning behind fake review predictions.

CONTRIBUTIONS OF THIS PAPER

The major contributions of this work are summarized as follows:

Development of an intelligent fake product review detection system using Natural Language Processing (NLP) and Machine Learning techniques to automatically classify reviews as genuine or fake.

Implementation of a hybrid feature engineering approach combining TF-IDF vectorization with additional linguistic and behavioral features such as sentiment polarity, review length, punctuation patterns, and writing complexity.

Comparative analysis of multiple machine learning models including Logistic Regression, Random Forest, and Support Vector Machine (SVM) to identify the most effective model for fake review detection.

Selection and deployment of Logistic Regression as the best-performing model, achieving an accuracy of approximately 89%, ensuring reliable and efficient classification.

Design and development of an interactive Streamlit-based web application that enables real-time fake review detection with user-friendly interface.

Integration of prediction confidence scores and probability-based outputs to provide transparency and improve decision-making reliability.

Implementation of an Explainable AI component that highlights possible reasons for identifying a review as fake, enhancing interpretability of model predictions.

Support for both single review analysis and batch processing using CSV file uploads, enabling scalability for large datasets.

Visualization of prediction results through graphical representation of fake and genuine probabilities, improving user understanding of model outputs.

Creation of a complete end-to-end pipeline including data preprocessing, feature extraction, model training, prediction, and deployment, making the system suitable for real-world e- commerce applications

LITERATURE SURVEY

[5] B. Liu, "Sentiment Analysis and Opinion Mining," Synthesis Lectures on Human Language Technologies, vol. 5, no. 1, pp. 1–167, 2012, presents the fundamental concepts, tasks, and methodologies involved in sentiment analysis and opinion mining. The work discusses various levels of sentiment analysis, including document-level, sentence-level, and aspect-level analysis, and explains how subjective expressions can be extracted from unstructured textual data. This study provides a strong theoretical foundation for understanding customer opinions and is widely used as a reference for designing fake product review detection systems.

[6] A. Mukherjee, B. Liu, and N. Glance, "Spotting Fake Reviewer Groups in Consumer Reviews," Proceedings of the 21st International World Wide Web Conference (WWW), 2012, pp. 191–200, proposes techniques to identify coordinated groups of fake reviewers rather than focusing on individual reviews. The authors analyze reviewer behavior patterns, review similarities, and group-level activities to detect spam campaigns. This work highlights the importance of combining textual features with reviewer behavioral analysis to improve fake review detection accuracy.

[7] A. Mukherjee, V. Venkataraman, B. Liu, and

N. Glance, "Fake Review Detection: Classification and Analysis of Real and Pseudo Reviews," Technical Report, University of Illinois at Chicago, 2013, explores supervised machine learning approaches for classifying genuine and deceptive reviews. The study compares real fake reviews with artificially generated pseudo reviews and analyzes linguistic, statistical, and structural differences between them. The findings emphasize the challenges involved in detecting sophisticated fake reviews and provide insights into effective feature selection strategies.

[8]S.Saumya and J.P.Singh,"Detection of Spam Reviews: A Sentiment Analysis Approach," CSI Transactions on ICT, vol. 6, pp. 137–148, 2018, introduces a sentiment-based spam review detection method that classifies reviews by analyzing opinion polarity and sentiment strength. The authors categorize reviews into genuine, fake, and suspicious classes and demonstrate that sentiment intensity can be an effective indicator of deceptive behavior. The experimental results show improved detection accuracy, making the approach suitable for practical deployment in online review platforms.

[9]Y.RenandD.Ji,"NeuralNetworksfor Deceptive Opinion Spam Detection: An Empirical Study," Information Sciences, vol. 385–386, pp. 213– 224, 2017, investigates the effectiveness of deep learning models for fake review detection. The study evaluates recurrent neural networks (RNNs) and convolutional neural networks (CNNs) for learning complex semantic patterns in review text. The results indicate that deep learning approaches outperform traditional machine learning classifiers, particularly in capturing subtle linguistic cues present in deceptive reviews.

[10]E.D.WahyunianandA.Djunaidy,"FakeReview Detection from Product Reviews Using Modified Iterative Computation Framework," MATEC Web of Conferences, vol. 58, 2016, proposes a modified iterative computation framework for fake review detection. The model assigns honesty scores to reviewers and repeatedly refines these scores based on review behavior and consistency. This iterative process improves

detection accuracy by gradually isolating deceptive reviewers, making the approach effective for identifying coordinated spam activities over time.

EXISTING SYSTEM

Researchers have proposed a wide range of techniques to address the problem of fake review detection. Advanced models such as the ICF++ framework incorporate reviewer honesty scores and have demonstrated a significant improvement in detection accuracy, reporting an increase of up to 49% [11]. Polarity-based approaches, including VADER sentiment analysis, have been used to categorize reviews into true, false, and suspicious classes by assigning polarity values of +1, −1, and 0, enabling effective identification and elimination of deceptive reviews [12]. Comprehensive surveys conducted over the past decade provide extensive coverage of sentiment analysis techniques and their applications in fake review detection, offering valuable insights into existing methodologies [13]. method achieved 91% of Some studies have focused on spam review detection by analyzing comments associated with reviews to assess their reliability and credibility, achieving F1-scores as high as 91% [14]. Temporal pattern-based methods have also been proposed to detect singleton spam reviews by identifying correlated posting behaviors over time [15].

Additionally, divergence-based techniques such as Kullback–Leibler (KL) divergence have been applied to distinguish fake reviews from genuine ones by exploiting asymmetric distribution differences, though challenges remain in identifying pseudo-fake reviews [16]. Other approaches utilize sequential review analysis with repeated feature extraction to effectively classify reviews as either genuine or deceptive [17].

LIMITATIONS OF EXISTING SYSTEM

The existing approaches for fake product review detection suffer from several limitations, which reduce their effectiveness in real-world applications:

High Dependence on Manual Verification Many systems still require manual inspection of reviews, which is time-consuming, inefficient, and not scalable for large datasets.

Limited Feature Utilization Traditional methods rely mainly on basic textual features such as keywords or n-grams, ignoring deeper linguistic, behavioral, and contextual patterns.

Low Detection Accuracy Existing: models often fail to accurately distinguish between genuine and fake reviews, especially when fake reviews closely mimic real writing styles.

Inability to Handle Large-Scale Data Many: systems are not designed to process millions of reviews in real-time, making them unsuitable for large e-commerce platforms.

Lack of Explain ability Most: models do not provide reasoning or justification for their predictions, making it difficult for users and administrators to trust the results.

Poor Handling of Short and Ambiguous Reviews Short or unclear reviews often lead to incorrect predictions due to insufficient contextual information.

No Real-Time Detection Capability Existing approaches are mostly offline and cannot provide instant predictions for newly submitted reviews.

Ignoring Behavioral and Sentiment Patterns: Many systems fail to consider important indicators such as sentiment intensity, punctuation usage, and writing patterns, which are crucial for detecting fake reviews.

Difficulty in Detecting Sophisticated Fake Reviews:

Advanced fake reviews that mimic human writing styles are hard to detect using simple machine learning techniques.

Limited Scalability and Deployment Issues: Many research models are not designed for practical deployment and lack integration with user-friendly interfaces or web applications.

PROPOSED SYSTEM

The proposed research introduces an intelligent framework for identifying and filtering deceptive product reviews by combining opinion mining techniques with supervised machine learning methods. In this framework, the Support Vector Machine (SVM) algorithm acts as the primary classification model, enabling the system to distinguish between authentic and misleading feedback with improved precision. The main objective is to strengthen the credibility of online review platforms by automatically analyzing textual patterns present in user-generated content.

Initially, review data is gathered from publicly available sources or benchmark datasets and forwarded to a preprocessing module. This stage focuses on cleaning the raw text by eliminating irrelevant components such as stop words, punctuation symbols, URLs, and repeated noise terms. Tokenization, lowercasing, and normalization are applied to standardize the textual input and prepare it for further analytical processing.

Following preprocessing, the refined review text is converted into structured numerical representations through feature extraction techniques such as Bag-of-Words and Term Frequency analysis. These representations capture important linguistic signals, including unusual word repetition, exaggerated sentiment expressions, and stylistic patterns that are frequently observed in deceptive reviews.

The generated feature vectors are supplied to the Support Vector Machine classifier, which learns to separate genuine and fake reviews by identifying the optimal decision boundary within a high-dimensional feature space. Because of its capability to manage sparse textual data efficiently, the SVM model provides reliable performance in opinion mining applications and reduces the risk of overfitting during training.

After the training phase, the model processes newly submitted reviews and predicts their authenticity in real time. Reviews identified as suspicious are highlighted for further monitoring or automated removal, ensuring that misleading content does not influence customer decisions. Model effectiveness is assessed using evaluation metrics such as accuracy, precision, recall, and F1-score to validate overall performance.

Overall, the proposed framework delivers a scalable and automated solution for fake review detection by integrating machine learning with natural language processing. The system minimizes manual moderation efforts while improving transparency in digital marketplaces. Additionally, the architecture is designed to handle large-scale datasets, making it suitable for deployment in real-world e-commerce environments. Future enhancements may include hybrid learning models and adaptive feature selection techniques to further improve detection accuracy. The modular design of the system also allows easy integration with web-based applications and existing review management tools.

SYSTEM ARCHITECTURE

The proposed fake product review detection system follows a modular and scalable architecture that integrates Natural Language Processing (NLP), feature engineering, and machine learning techniques to automatically classify reviews as genuine or fake.

The architecture consists of multiple interconnected components, each responsible for a specific stage in the processing pipeline.

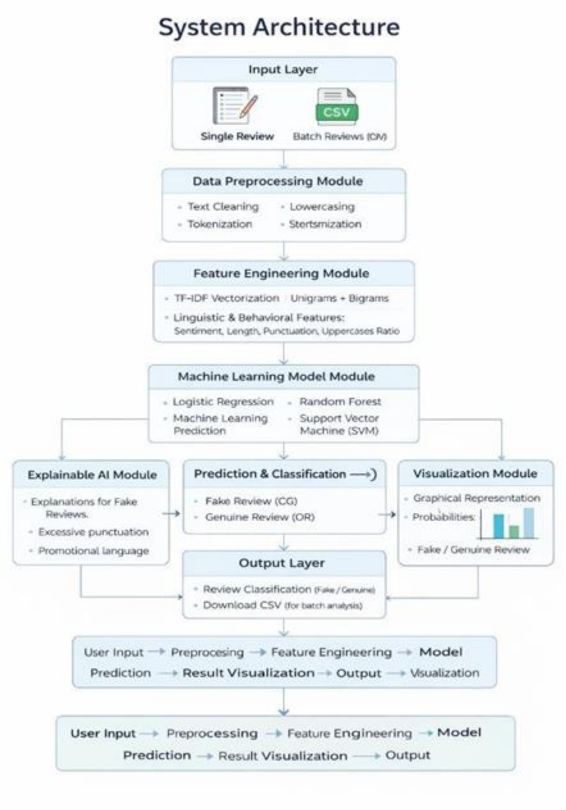

1. Input Layer (User Interface)

The system accepts input in two forms:

• Single review entered manually by the user

• Batch input through CSV file upload This layer is implemented using a Streamlit- based web interface, allowing users to interact with the system easily.

2. Data Preprocessing Module

The preprocessing module cleans and standardizes the raw review text to improve model performance. The following operations are performed:

• Conversion of text to lowercase

• Removal of special characters and punctuation

• Tokenization (splitting text into words)

• Stemming (reducing words to root form) This step ensures that noisy and irrelevant data is removed before feature extraction.

3. Feature Engineering Module

In this stage, the cleaned text is transformed into numerical representations suitable for machine learning models.

• TF-IDF Vectorization: Converts text into numerical feature vectors using unigram and bigram representations.

• Linguistic & Behavioral Features: Additional features such as review length, word count, sentiment polarity, punctuation usage, and uppercase ratio are extracted.

Both types of features are combined to form a high-dimensional feature vector representing each review.

4. Machine Learning Model Module

The processed feature vector is passed to trained machine learning models. Multiple models were evaluated, including:

• Logistic Regression

• Random Forest

• Support Vector Machine (SVM)

Among these, Logistic Regression is selected as the final model due to its superior performance (~89% accuracy) and efficiency in handling high- dimensional text data.

5. Prediction and Classification Module

The trained model predicts whether a given review is:

• Fake (CG – Computer Generated)

• Genuine (OR – Original Review)

The system also provides:

• Prediction confidence score

• Probability of fake and genuine classification

6. Explainable AI Module

To improve transparency, the system provides explanations for predictions, especially for fake reviews. Examples include:

• Excessive punctuation

• Very short review

• Highly promotional language

This helps users understand why a review is classified as fake.

7. Visualization Module

The system presents results using graphical representations such as bar charts showing:

• Fake probability

• Genuine probability

This improves user understanding and decision-making.

8. Output Layer

The final output is displayed to the user through the web interface, including:

• Review classification (Fake/Genuine)

• Confidence score

• Explanation (if applicable)

For batch analysis, results can be downloaded as a CSV file.

9. Workflow Summary

The overall system workflow can be summarized as:

User Input → Preprocessing → Feature Engineering → Model Prediction → Result Visualization → Output

Workflow of the Proposed System:

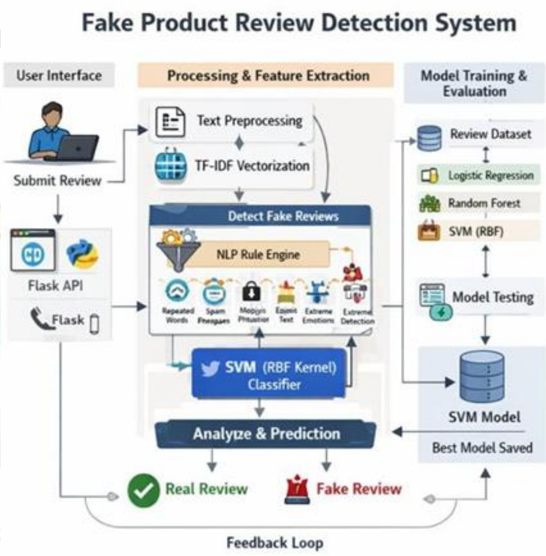

Figure ; Fake Product Review Detection System

The proposed fake product review monitoring and removal system follows a structured workflow that integrates user interaction, text preprocessing, feature extraction, machine learning–based classification, and feedback- driven decision-making to automatically identify deceptive reviews.

1. Review Data Collection and User Interface

The review data collection and user interface module serves as the entry point of the proposed fake review detection system. It provides a simple and structured environment where users can submit product reviews for analysis. Reviews may be gathered from online sources, benchmark datasets, or directly entered through a web-based interface. Each submission typically includes textual feedback along with optional metadata such as user details, timestamps, or product information, which helps maintain contextual relevance during processing.

2. Text Preprocessing

Once a review is received by the system, it is processed through a text preprocessing stage that prepares the raw input for effective analysis. This phase focuses on removing unwanted noise and standardizing the textual structure to improve the reliability of the classification process. Cleaning operations include eliminating stop words, punctuation, hyperlinks, numerical artifacts, special symbols, and repeated tokens that do not contribute meaningful context. Tokenization is then applied to divide the review into individual words or terms, allowing the system to analyze linguistic patterns more precisely.

After tokenization, normalization techniques such as lowercasing, whitespace correction, and basic formatting are performed to maintain consistency across the dataset. Additional steps like stemming or lemmatization may also be applied to reduce words to their base forms, ensuring similar expressions are treated uniformly. By transforming unstructured text into a clean and structured format, the preprocessing module reduces data complexity, improves feature quality, and enhances the overall performance of the fake review detection model

3. Feature Extraction

After preprocessing, meaningful features are extracted from the cleaned review text to prepare it for machine learning analysis. The system applies TF-IDF (Term Frequency–Inverse Document Frequency) vectorization to transform textual data into numerical feature vectors. This approach measures the importance of words based on their frequency within a review and across the dataset. It helps highlight significant terms while reducing the impact of commonly used words. The generated vectors capture linguistic patterns that are useful for detecting deceptive language. Feature extraction also reduces data complexity by converting unstructured text into a structured format. These numerical representations allow the classifier to learn meaningful relationships between words and review authenticity. As a result, the overall accuracy and reliability of the fake review detection model are improved.

4. NLP Rule Engine

In parallel with feature extraction, the system employs an NLP rule engine to analyze linguistic patterns. This module detects characteristics such as repeated words, spam phrases, excessive modifiers, explicit sentiment expressions, and abnormal usage of emojis or punctuation. The rule-based analysis enhances the system's ability to detect deceptive behavior by complementing the machine learning model.

5. Machine Learning Model Training and Evaluation

The extracted feature vectors are used to train supervised machine learning models. The system evaluates multiple classifiers, including Logistic Regression, Random Forest, and Support Vector Machine (SVM). Among these, the SVM with RBF (Radial Basis Function) kernel is selected as the final model due to its effectiveness in handling high-dimensional text data.

During training, the SVM learns an optimal decision boundary that separates genuine reviews from fake reviews by maximizing the margin between the two classes. Model evaluation is performed using standard performance metrics such as accuracy, precision, recall, and F1-score. The best-performing SVM model is saved for deployment.

6. Review Classification and Prediction

Once the model is trained, incoming reviews are After the training phase is completed, the finalized Support Vector Machine model with the RBF kernel is deployed as the prediction engine of the system. Each newly submitted review passes through the same preprocessing and feature extraction pipeline used during training to maintain consistency in data representation. The transformed feature vector is then provided as input to the saved SVM model, which evaluates the linguistic characteristics of the review and determines its likelihood of being genuine or deceptive. By analyzing learned patterns such as unusual word frequency, exaggerated sentiment expressions, and structural inconsistencies, the classifier generates a prediction label along with a confidence score.

The classification output plays a central role in the system's decision-making workflow. Reviews predicted as genuine are allowed to remain visible on the platform, while those identified as fake are flagged for administrative monitoring or automatic filtering based on predefined thresholds. The prediction results are also logged in the system database to support further analysis and continuous improvement of the model. In addition, the framework can provide real-time feedback to users or moderators, enabling faster response to suspicious activities. This automated classification process reduces manual review effort, improves moderation efficiency, and ensures that only trustworthy content influences customer decisions on the platform.

RESULT ANALYSIS

The proposed fake product review monitoring and removal system was evaluated using benchmark review datasets containing both genuine and deceptive reviews, simulating real- world e-commerce review scenarios. The system performance was analyzed in terms of text classification accuracy, precision, recall, F1- score, and overall effectiveness in identifying fake reviews.

The text preprocessing module successfully normalized raw review data by removing noise such as stop words, punctuation, URLs, and redundant tokens. This preprocessing significantly improved feature consistency and reduced dimensionality, enabling efficient learning during the classification stage.

Feature extraction using TF-IDF vectorization proved effective in capturing the importance of discriminative words across reviews. Words with exaggerated sentiment, repeated promotional phrases, and abnormal frequency distributions were assigned higher relevance, aiding the detection of deceptive review patterns. The NLP rule engine further enhanced detection accuracy by identifying linguistic indicators such as repeated words, spam phrases, excessive modifiers, and abnormal emoji usage.

The Support Vector Machine (SVM) classifier with RBF kernel demonstrated strong performance in separating genuine and fake reviews within the high-dimensional feature space. Compared to baseline classifiers such as Logistic Regression and Random Forest, the SVM model achieved better generalization and reduced misclassification, particularly in borderline cases where fake reviews closely resembled genuine ones.

During testing, the proposed model effectively classified incoming reviews into genuine and fake categories. Reviews identified as fake were correctly flagged for monitoring or removal, preventing deceptive content from influencing customer decisions. The feedback loop enabled continuous refinement of the model by incorporating newly classified data.

Experimental observations indicate:

• Improved classification accuracy using TF-IDF and SVM (RBFkernel)

• Reduced false positives compared to traditional

sentiment-based filtering

• Effective identification of spam reviews with

exaggerated language patterns

•Scalable performance suitable for large review

datasets

Overall, the proposed model demonstrates superior performance compared to conventional review moderation techniques and basic machine learning approaches. By integrating NLP-based rule analysis with SVM classification, the system provides a robust, automated, and reliable solution for fake product review detection, thereby enhancing trust and credibility in online review platforms.

FUTURE SCOPE

The proposed fake product review detection system can be further enhanced by incorporating advanced deep learning techniques such as LSTM, RNN, and transformer-based models like BERT, which are capable of capturing deeper contextual and semantic relationships in text. These models can significantly improve the system's ability to detect complex and well-crafted fake reviews, including those that mimic genuine writing styles or contain sarcasm. In addition, integrating multilingual support would allow the system to analyze reviews in multiple languages, making it more suitable for global e- commerce platforms.

Another important direction for future work is the inclusion of user behavioral analysis, such as reviewer history, posting patterns, and rating consistency, which can provide stronger indicators of deceptive activity. The system can also be extended into a real-time monitoring framework that automatically detects and filters fake reviews as they are posted. Furthermore, implementing continuous learning mechanisms and enhancing explainable AI features would improve model adaptability and transparency. With these improvements, the system can be deployed at scale and integrated directly into real- world e-commerce platforms to ensure more reliable and trustworthy user feedback systems.

CONCLUSION

The results discussed in this article are the comparison of two models developed to justify the model performance for this "Amazon's Yelp" dataset and their relevance for deployment in real-time software applications. Hence, the Random Forest model performed significantly better compared to the Naïve Bayes algorithm by a large margin. The fake review detection problem is addressed effectively and provides meaningful insight into its practical importance and ethical considerations, with the main objective being the selection of a suitable algorithm for identifying and eliminating deceptive reviews. In future work, hybrid architectures and advanced learning models can be explored to further enhance detection accuracy. By utilizing platforms such as Google Colab and NVIDIA GPU resources, the computational efficiency and training speed of the research can be improved. Additionally, incorporating larger and more diverse datasets may help improve model generalization across different domains.

Conflict of interest statement

Authors declare that they do not have any conflict of interest.

REFERENCES

[1] W. Hameed, R. Allami, and Y. H. Ali, "Fake Review Detection Using Machine Learning," Revue d'Intelligence Artificielle, vol. 37, no. 5, 2024.

[2] D. Zhang, "A Deep Learning Approach for Detecting Fake Reviewers," Elsevier, 2023,

[3] J. Liu, "A Study on Fake Review Detection Based on RoBERTa," Elsevier, 2024,.com/science/article/ pii/S1877050924018507

[4] A. Mewada, "Research on False Review Detection Methods: A State-of-the-Art Review," Elsevier, 2022

[5] P. Hajek, L. Hikkerova, and J.-M. Sahut, "Fake Review Detection in E-Commerce Platforms Using Aspect-Based Sentiment Analysis," Elsevier,2023.

[6] R. Mohawesh et al., "Fake Review Detection Using Transformer-Based Enhanced Models," 2024,

[7] P. Phukon et al., "Fake Reviews Detection Using Aspect-Based Sentiment Analysis and GCN," 2025,

[8] "Fake Review Detection Using Machine Learning Algorithms," IJPREMS, 2023,

[9] "Fake Review Detection Using NLP and Neural Networks," IJERT, 2023.

[10] A. Yousif, "Impact of Sentiment Analysis on Fake Review Detection," arXiv, 2022.

[11]

[12] "Machine Learning and Deep Learning for Fake Review Detection," IJSRST, 2024,

[13] "Detecting Fake Reviews Using Random Forest and SVM," Sinkron Journal, 2023.

[14] A. Jnoub and W. Klas, "Declarative Programming Approach for Fake Review Detection," IEEE Workshop, 2020.

[15] R. Gupta, "Recent State-of-the-Art of Fake Review Detection: A Comprehensive Review," Cambridge Knowledge Engineering Review, 2024.

[16] M. Ennaouri, "Machine Learning Approaches for Fake Reviews Detection," JWE, 2023.

[17] "Data Analytics for Fake Reviews Using TF-IDF and N-grams," Techscience CMC, 2024.

[18] "Detecting Fake Review Intentions in Review Context," Elsevier, 2025.

[19] "Fake Review Product Detection Using Opinion Mining," IJARSCT, 2024.

[20] "Multilingual Sentiment Analysis in E- Commerce," ScienceDirect, 2025.

[21] "Detect Fake Reviews: A Survey, IEEE Access," IEEE, 2021.