Journal of Theoretical and Applied Information Technology
31st October 2026. Vol.104. No.20
© Little Lion Scientific
ISSN: 1992-8645 www.jatit.org E-ISSN: 1817-3195

HYBRID TF-IDF - WORD2VEC FEATURE FUSION WITH STACKING ENSEMBLE LEARNING FOR URDU FAKE NEWS DETECTION

ZAHIR ABBAS KHAN[1], REKHA V[2]

[1]Assistant Professor, Department of Computer Science and Engineering, Maulana Azad National Urdu University, Hyderabad. Polytechnic, Bangalore, India.

[2]Associate Professor, Department of AI and DS Engineering, Christ University, Bangalore, India

[1]E-mail: abbas.zak@manuu.edu.in, Orcid: 0009-0009-5557-5024

[2]E-mail: rekha.v@christuniversity.in, Orcid: 0000-0001-5275-5235

ABSTRACT

The rapid dissemination of misleading information online has become a major social issue, particularly in low-resource languages like Urdu. The ability of existing false news detection techniques to recognize semantic connections in textual data is limited because they often solely employ lexical representations. This study proposes a hybrid fake news detection method that combines Word2Vec embeddings with Term Frequency–Inverse Document Frequency (TF-IDF) to leverage lexical and semantic information present in Urdu news articles. The gathered features are fused and classified using a stacking ensemble that includes Support Vector Machine (SVM), Logistic Regression (LR), and Extreme Gradient Boosting (XGBoost) as foundation learners. For the studies, a combined Hook-and-Bait Urdu fake news dataset including 78,406 news articles 42,179 real and 36,230 fake news samples—was utilized. The proposed framework had an accuracy of 95.17%. Comparative studies show that the hybrid representation and ensemble learning approach works better than conventional machine learning techniques based solely on TF-IDF features. The findings demonstrate that integrating lexical and semantic representations can greatly improve Urdu false news detection ability.

Keywords: *Fake News Detection, Machine Learning, Natural Language Processing, Stacking Ensemble, TF-IDF, Word2Vec*

1. INTRODUCTION

The growth of digital media and online news outlets has accelerated the flow of information, but it has also increased the dissemination of fake news and misinformation, affecting public opinion and trust in digital sources. Hence, automated fake-news detection is an important research problem for NLP and machine learning (ML) especially for low-resource languages such as Urdu where linguistic variation and lack of annotated datasets and computational resources pose additional challenges.

Traditional approaches often use the TF-IDF (Term Frequency–Inverse Document Frequency) which captures discriminative words and phrases but does not explicitly represent semantic relationships. Word2Vec learns distributional semantic associations, but averaging word embeddings to document vectors may lose important lexical distinctions. This combination of complementary representations may improve the news classification. Also, stacking ensemble learning combines the predictions of several classifiers and improves the performance by utilizing their complementary strengths.

The present study addresses the problem of representing and categorizing Urdu news articles by combining lexical and semantic features with ensemble learning. There have been a few studies on machine-learning classifiers and word representations for fake-news detection. However, the effectiveness of TF-IDF and Word2Vec fusion with stacking needs to be systematically evaluated on the Hook-and-Bait Urdu dataset.

To address this gap, this paper proposes a feature-fusion framework by combining TF-IDF features and document-level Word2Vec embeddings for binary Urdu fake-news classification. The base classifiers are Support Vector Machine (SVM), Logistic Regression (LR) and Extreme Gradient Boosting (XGBoost), and the meta-classifier is Logistic Regression to combine their predictions.

The specific aims are to:

1. Clean and normalize Urdu news articles for classification.

2. Adopt TF-IDF and Word2Vec to obtain the complementary lexical and semantic features.

3. Fuse both representations into a shared feature space.

4. Test stacking ensemble of SVM, LR and XGBoost with LR as the meta-classifier.

5. Compare the proposed framework with individual classifiers using accuracy, precision, recall, F1-score, confusion matrices and ROC-AUC where applicable.

The study contributes to the detection of Urdu fake news by evaluating the integration of lexical and semantic representations with stacking ensemble learning. The proposed framework is evaluated against individual classifiers with a common experimental protocol. Rather than claiming this combination is better than any other approaches, the study investigates the effectiveness of this specific combination and paves the way for future research on contextual language models, explainable detection, and cross-domain evaluation.

2. RELATED WORKS

2.1 Conventional Machine Learning Techniques for Fake News Detection

The traditional machine learning methods employ text representations like Bag of Words, Term Frequency-Inverse Document Frequency (TF-IDF) and classifiers like Logistic Regression (LR), Support Vector Machine (SVM), Random Forest and gradient boosting. These approaches provide computationally tractable solutions for classification of labeled news articles. TF-IDF finds words and phrases that discriminate but does not represent explicitly the semantic relationships between words. To solve this problem, we use complementary representations.

2.2. Dataset Preparation for Urdu Fake News Detection

Most of the research on Urdu fake-news detection has been on dataset creation, language-specific processing, and supervised classification. Humayoun [3] addressed the detection of Urdu fake-news in the FIRE task, and Farooq et al. [5] created a manually annotated multi-domain Urdu corpus, and assessed the performance of machine-learning approaches. But there are still some difficulties such as spelling variations, inconsistent normalization, domain-specific vocabulary and limited annotated data. Variations in datasets and evaluation protocols also affect comparability, therefore emphasizing the requirement for well documented preprocessing and experimental procedures.

2.3 Semantic representations and word embeddings

Word2Vec learns dense representations of words from their co-occurrence in context, capturing distributional relationships that frequency-based features may miss. Word2Vec and related representations have been used for Urdu fake news in previous work [2, 5]. But the performance of the embedding depends on the application and the coverage of the vocabulary, the training data and the representation of the document. The average of word vectors provides fixed-length representations but loses the order of words and dilutes the discriminative lexical information. TF-IDF combined with embeddings can capture complementary lexical and semantic features.

2.4. Feature level fusion

Feature fusion is the combination of multiple representations to provide complementary information to classifiers. TF-IDF captures informative words and phrases and Word2Vec captures distributional semantic relations. Concatenation of them can be useful for classification, however, the resulting high dimensional feature space may increase the computational requirement. For Urdu fake-news detection, several representations and combinations of features have been explored [5, 10]. Rather than assuming that fusion improves performance, here we evaluate a fused TF-IDF–Word2Vec representation with individual classifiers and a stacking ensemble.

2.5. Fake News Detection Using Ensemble Learning

Ensemble learning is a method of combining several models together to increase the predictive performance. The predictions of base classifiers are used as input to a meta-classifier in stacking. The proposed framework considers SVM, LR and XGBoost as base classifiers and LR as meta classifier. These algorithms employ different learning strategies and can provide complementary predictions. However, the diversity of the classifiers and the training procedures is crucial for the ensemble performance. Reliable training of the meta-classifier requires out-of-fold predictions or some other leakage-free approach.

2.6. Deep learning and Transformer methods

Deep-learning and transformer-based methods can model contextual relationships; static embeddings cannot model. The research has been focused on the pre-trained language models and multilingual methods [4, 6, 12] in the detection of fake news in Urdu. But these approaches may require more computational power and careful fine-tuning. Results are not directly comparable across datasets and experimental setups. Hence, in this paper, we evaluate a feature-engineering and ensemble-learning approach, without claiming that it is better than transformer-based approaches.

2.7 Research Gap and Motivation

The literature indicates three issues: TF-IDF can capture discriminative lexical patterns but cannot capture explicit semantic relationships; averaged Word2Vec embeddings can lose important lexical distinctions; and the benefits of ensemble learning depend on the diversity of individual classifiers and the use of proper evaluation procedures. Differences in datasets and experimental conditions further complicate reproducibility and comparability.

To address these issues, this work explores the combination of TF-IDF and Word2Vec features with a stacking ensemble for Urdu fake-news detection on Hook-and-Bait dataset. The base classifiers are SVM, LR and XGBoost and the meta-classifier is LR. The framework is evaluated based on accuracy, precision, recall, F1-score, and confusion-matrix analysis and ROC-AUC when applicable. In this paper, we present the empirical performance of this particular setting, and discuss its potential limitations. We hope this work will serve as a starting point for future research on misinformation detection in Urdu.

3. PROPOSED METHODOLOGY

3.1 Dataset Description.

Harris et al. [6] 2025, presented the Hook-and-Bait Urdu dataset, a large-scale benchmark corpus intended for the identification of fake news in the Urdu language. News items from several fact-checking websites and reputable news sources were gathered to create the dataset. Six Excel files were combined to produce a single corpus of 78,406 news stories for this investigation, which included 36,230 cases of fake news and 42,179 instances of legitimate news. The dataset offers significant variation in content and writing styles because it spans several disciplines, such as politics, religion, entertainment, social issues, and public affairs. The Hook-and-Bait Urdu dataset is a crucial baseline for creating automated false news detection systems for low-resource languages because of its extensive and genuine Urdu textual content, the class distribution as show in Figure 1.

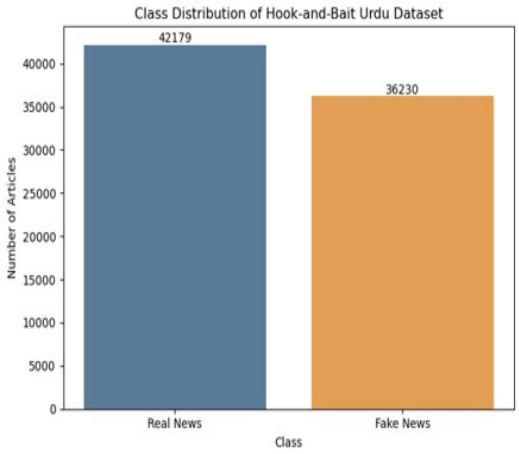


Figure 1. Dataset distribution

3.2 Dataset Preprocessing.

Before feature extraction, the preprocessing pipeline consists of a number of crucial processes to enhance the consistency and quality of the Urdu news articles. Only the pertinent columns with news text and labels were kept after all news records from several Excel files were combined into a single dataset. After that, the class labels were standardized by translating textual labels into numerical representations, giving fraudulent news a label of 1 and authentic news a label of 0 Mansoor et al. in [7] 2019. Duplicate and missing records were then removed in order to preserve data integrity. URLs, online links, superfluous special characters, and excessive white spaces were eliminated during text preprocessing. The Urdu text was further normalized by removing unnecessary symbols that can have a detrimental impact on model performance and formatting it consistently. Lastly, empty records produced during the cleaning procedure were eliminated, leaving a refined dataset that could be used for Word2Vec and TF-IDF feature extraction. The pipeline for preprocessing consists of:

1. URL elimination
2. Removal of numerical characters
3. Removal of special characters

4. Normalization of Urdu text
5. Normalization of white space

3.3 TF-IDF Feature Extraction.

Term Frequency–Inverse Document Frequency (TF-IDF) was used by G. Salton et al. [8] in 1988 to convert textual news items into numerical feature vectors appropriate for machine learning methods, as shown in equation 1. A popular statistical method called TF-IDF measures a word's significance in a document in relation to the corpus as a whole. While the inverse document frequency (IDF) gives discriminative terms more weight and lessens the impact of frequently appearing words, term frequency (TF) quantifies how frequently a word appears within a text. As a result, TF-IDF successfully collects lexical traits and document-specific data that are helpful in differentiating between authentic and fraudulent news pieces.

$$\text{TF-IDF}(t,d) = \text{TF}(t,d) \times \text{IDF}(t) \quad (1)$$

where t is a term, d is a document, TF denotes term frequency, and IDF denotes inverse document frequency.
The TF-IDF method was used in this study to vectorize the cleaned Urdu news corpus with a maximum feature size of 50,000. To capture individual words as well as contextual word pairs, both unigram and bigram representations were taken into consideration. To cut down on noise and enhance feature quality, terms that appeared in less than two papers were eliminated. By maintaining the most instructive lexical patterns in the dataset, the resulting TF-IDF matrix offered a high-dimensional sparse representation of the news articles. A hybrid feature representation that captures both lexical relevance and semantic links among Urdu words was then created by combining these features with Word2Vec embeddings. Setting up:

• 50,000 features is the maximum.
• N-Grams: Bigrams and Unigrams
• Document Frequency Minimum: 2

3.4 Word2Vec Embedding Generation.

The Skip-Gram architecture was used by T. Mikolov et al. [9] in 2013 to create Word2Vec embeddings that captured the semantic links between words in the Urdu news corpus. As shown in equation 2, Word2Vec learns dense vector representations by predicting surrounding context words for a given target word, in contrast to TF-IDF, which concentrates on lexical importance. This allows the model to identify contextual and linguistic similarities in the text by mapping semantically related words to adjacent positions in the embedding space.

$$V_d = \frac{1}{n} \sum_{i=0}^{n} V_{w_i} \quad (2)$$

where, v_d denotes the document vector, v(wi) represents the Word2Vec embedding of the i^th word, and n is the total number of words in the document.

In this study, a Word2Vec model was trained on the training corpus using the Skip-Gram technique after the cleaned Urdu news articles were tokenized. at reduce noise, the embedding dimension was set to 300, the context window size was fixed at 5, and terms that appeared less than twice were removed from the vocabulary. Each phrase was represented as a 300-dimensional dense vector following training. The vectors of every word in a news story were averaged to produce document-level representations, which produced a fixed-length feature vector for every document. In addition to the lexical information supplied by TF-IDF, the resulting Word2Vec embeddings caught the semantic and contextual information found in Urdu news articles. A thorough hybrid representation for the categorization of fake news was then produced by fusing these document vectors with TF-IDF features.
Setting up:
• Dimensions of the vector: 300
• Size of Window: 5
• Two is the minimum count.
• Twenty epochs

3.5 Feature Fusion

Khan et al. [10] used a feature fusion technique in 2023 to combine TF-IDF and Word2Vec features into a single hybrid representation in order to capitalize on the complimentary capabilities of lexical and semantic representations. Word2Vec offers dense semantic embeddings that contain contextual links among words, whereas TF-IDF successfully captures the significance and discriminative power of words inside individual documents. The model can make use of both surface-level lexical patterns and deeper semantic information found in Urdu news articles thanks to the combination of these two representations.

The 300-dimensional document-level Word2Vec vectors produced by the Skip-Gram model were concatenated with the 50,000-dimensional TF-IDF feature matrix in the suggested framework. Each news story has 50,300 features in the resulting hybrid feature space. This fused representation adds semantic information from Word2Vec embeddings while maintaining the sparsity and discriminative power of TF-IDF. As a result, the hybrid feature set offers a more thorough depiction of news content, improving the classifier's capacity to discern between authentic and fraudulent news. For the final classification, the stacking ensemble classifier used the fused features as input.

3.6 Stacking Ensemble Classifier

The framework for detecting bogus news was made more robust and predictive by using a stacking ensemble classifier. Stacking is a sophisticated ensemble learning method that uses a meta-classifier to produce the final prediction, combining the advantages of several basic classifiers Alzubaidi et al. [11] in 2025. In contrast to individual machine learning models, stacking makes use of the variety of learning methods, which lowers model bias and variance while improving generalization capacity.

Three heterogeneous classifiers were chosen as basic learners in the suggested framework depicted in Figure 2. Support Vector Machine (SVM), Logistic Regression (LR), and Extreme Gradient Boosting (XGBoost). XGBoost was selected for its capacity to capture complex nonlinear correlations through gradient boosting, Logistic Regression for its simplicity and robust probabilistic interpretation, and SVM for its efficacy in handling high-dimensional textual characteristics. Each base classifier produced prediction outputs during the training phase by individually learning patterns from the hybrid TF-IDF and Word2Vec feature representation.

The meta-classifier then used the base learners' predictions as input features. The final classification decision was generated by combining the results of the individual classifiers using Logistic Regression as the meta-learner. The stacking ensemble successfully caught a variety of decision boundaries and enhanced classification reliability by combining the advantages of several learning techniques.

Base Learner:
•Vector Machine Support
•Logistic Regression
•XGBoost

Meta Learner: Logistic Regression
To increase generalization and lower classification errors, the stacking architecture aggregates predictions from several classifiers.

4. EXPERIMENTAL SETUP

The Hook-and-Bait Urdu dataset, which comprises 78,406 annotated news articles, was used to thoroughly assess the suggested false news detection methodology Harris et al. [12] in 2025. A stratified 80:20 train–test split technique was used to provide reliable and objective model evaluation. In order to mitigate potential sample bias and guarantee that both training and testing sets accurately reflect the dataset's intrinsic imbalance features, stratification was used to maintain the original class distribution across both subsets.

This partitioning technique resulted in the allocation of 62,724 samples for model training and 15,682 samples specifically for performance evaluation that was not yet visible. This division ensures that the model's capacity for generalization is evaluated using entirely distinct data.

A hybrid representation method of Word2Vec embeddings and TF-IDF was used for feature extraction. Word2Vec learned to embed words in a continuous vector space that captured semantic relations, whereas TF-IDF captured discriminative patterns of word frequencies and lexical importance. The extracted features were fused to a common representation that exploits semantic and lexical information. A stacking ensemble classifier was trained on the combined feature space for accurate detection of Urdu fake news. The fused features were divided into training and testing sets to evaluate the generalization ability of the model. Base learners are Logistic regression (LR) and XGBoost (XGB) with meta-classifier combining predictions. Finally, the classification performance of the proposed model was evaluated using accuracy, precision, recall and F1-score.

Journal of Theoretical and Applied Information Technology
31st October 2026. Vol.104. No.20
© Little Lion Scientific
ISSN: 1992-8645 www.jatit.org E-ISSN: 1817-3195

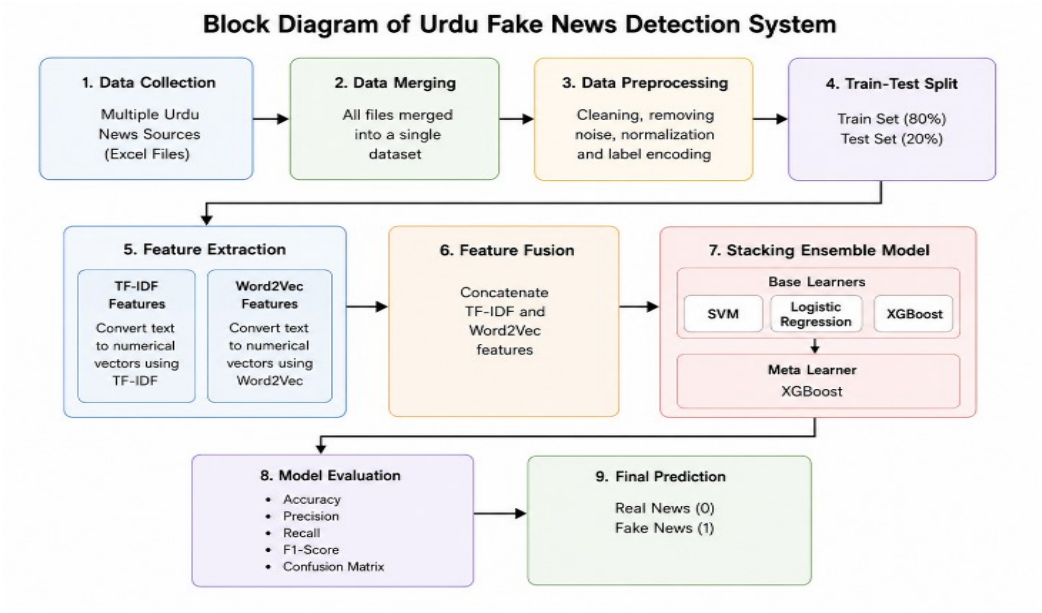


Figure 2. Proposed Framework

Algorithm 1: Urdu Fake News Detection Using TF-IDF, Word2Vec and Stacking

Input: Urdu news corpus D
Output: Fake/Real labels
Step 1: Preprocess text documents.
Step 2: Extract TF-IDF features XTFIDF.
Step 3: Train Word2Vec and generate XW2V.
Step 4: Fuse features XH = [XTFIDF || XW2V].
Step 5: Split XH into train and test sets.
Step 6: Train LR and XGB base learners.
Step 7: Train stacking meta-classifier.
Step 8: Predict labels on the test set.
Step 9: Evaluate using Accuracy, Precision, Recall and F1-score.
Step 10: Return predictions and metrics.

About 50,300-dimensional feature vectors made up the final concatenated feature space, which was then fed into the classification model. The representation gains from the complimentary advantages of both statistical and neural embedding techniques because of the high dimensionality and possible feature redundancy.

A stacking ensemble learning architecture was used for classification in order to lower model variance and increase prediction robustness. Three heterogeneous base learners made up the ensemble: Support Vector Machine (SVM), which works well in high-dimensional sparse spaces; Logistic Regression (LR), which offers probabilistic linear decision boundaries; and XGBoost, a gradient boosting framework that uses sequential tree optimization to capture intricate non-linear feature interactions, Asghar et al. [13] in 2021.Logistic regression was used as the second-level learner in a meta-classifier that received the outputs (class probabilities or decision scores) produced by these base learners. By learning weighted decision boundaries over the basis models' outputs, the meta-classifier was trained to optimally integrate their predictions, minimizing the biases of each individual model and enhancing generalization performance. The Figure 3 shows the flowchart of the proposed framework.

The system can take advantage of the complementing capabilities of linear, kernel-based, and boosting models thanks to this hierarchical stacking method, which eventually improves classification stability and detection performance on Urdu fake news data. Several common categorization criteria were used in order to thoroughly evaluate the efficacy of the suggested model Azam et al [14] in 2025. The aggregate percentage of correctly identified cases is known as accuracy. Precision assesses the percentage of anticipated fake news stories that turn out to be fraudulent. Recall quantifies the classifier's capacity to accurately detect instances of bogus news within the dataset. The harmonic mean of precision and recall, or F1-score, offers a fair assessment of

classification performance, especially in situations where there is class imbalance. In order to provide a thorough picture of the classifier's predictive behavior, a confusion matrix was also created to show the distribution of true positives, true negatives, false positives, and false negatives.

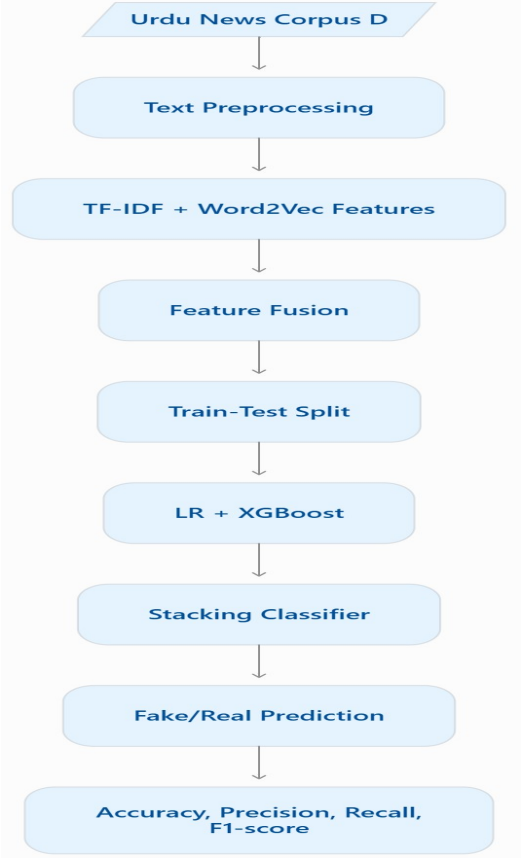


Figure 3. Flowchart of proposed framework

5. RESEARCH METHOD AND EXECUTION PROTOCOL

The proposed method was evaluated on the Hook-and-Bait Urdu fake-news dataset, which consists of real and fake news articles. The dataset was pre-processed by removing duplicate records, null values, empty texts, URLs, numeric characters and special characters. The white space correction and normalization of the Urdu text was done before encoding the real and fake news as 0 and 1 respectively. The data set was randomly split into 80% training and 20% testing set with stratified sampling to preserve the class distribution.

TF-IDF features were extracted with no more than 50,000 features, combinations of unigrams and bigrams and a minimum document frequency of 2. We used the Skip-Gram architecture to train Word2Vec embeddings with a vector size of 300, a context window of 5, a minimum word count of 2, and 20 training epochs. Word2Vec. The word embeddings were averaged to get document level representations. The TF-IDF and Word2Vec features were combined to get the fused feature representation.

The fused features were evaluated using Logistic Regression, Random Forest, Support Vector Machine (SVM), XGBoost and proposed stacking ensemble. The stacking model adopted SVM, Logistic Regression and XGBoost as base classifiers and Logistic Regression as meta-classifier. Its performance was assessed in terms of accuracy, precision, recall, F1-score and ROC-AUC. All experiments were implemented in Python using standard machine learning and NLP libraries. For reproducibility, the final manuscript should specify the particular versions of Python and libraries used, the operating system, processor, RAM and GPU configuration, if applicable.

6. RESULTS AND DISCUSSION

Using the testing set of 15,682 news articles from the Hook-and-Bait Urdu dataset, the effectiveness of the suggested hybrid fake news detection framework was assessed. As illustrated in Figure 4, the suggested model incorporated TF-IDF and Word2Vec feature representations and used a stacking ensemble classifier made up of Support Vector Machine (SVM), Logistic Regression (LR), and XGBoost. The experimental results in Table 1 indicate how well lexical and semantic data may be integrated to detect fake news in Urdu.

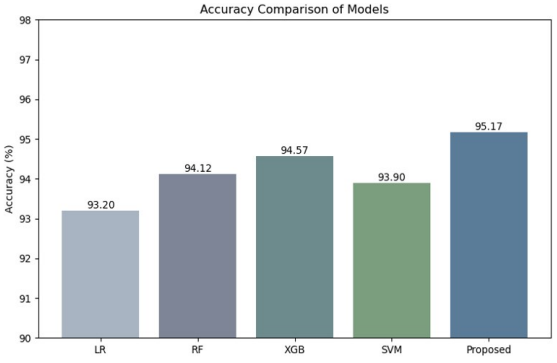


Figure 4. Performance of various models

Table 1. Performance of the various models

Journal of Theoretical and Applied Information Technology
31st October 2026. Vol.104. No.20
© Little Lion Scientific
ISSN: 1992-8645 www.jatit.org E-ISSN: 1817-3195

Model	Accuracy	Precision	Recall	F1-Score
LR	93.20	93.80	91.20	92.48
RF	94.12	94.50	92.80	93.64
XGB	94.57	94.90	93.50	94.19
SVM	93.90	94.20	92.60	93.39
Proposed	95.17	95.77	93.69	94.72

As seen in Figure 5 and Table 2, the suggested framework attained an accuracy of 95.17%, indicating its efficacy in differentiating between false and genuine news pieces. Only 300 actual news pieces were incorrectly classified as false and 457 fake news articles as real by the program, which correctly recognized 8,136 real news cases and 6,789 fake news instances. On the Hook-and-Bait Urdu dataset, the framework achieved a balanced and reliable classification performance with precision of 95.77%, recall of 93.69%, and F1-score of 94.72%.

Table 2. Performance of the Proposed Framework

Metric	Accuracy	Precision	Recall	F1-Score
Value	95.17	95.77	93.69	94.72

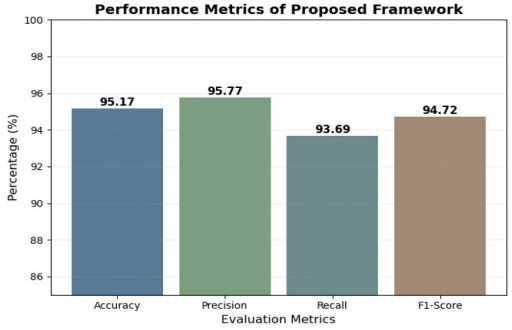


Figure 5. Performance Metrics

As seen in Figure 6, the confusion matrix provides additional evidence of the classifier's efficacy. Of the test samples, 6,789 fraudulent news pieces were successfully identified as fake (True Positives), whereas 8,136 actual news articles were accurately classified as real (True Negatives).

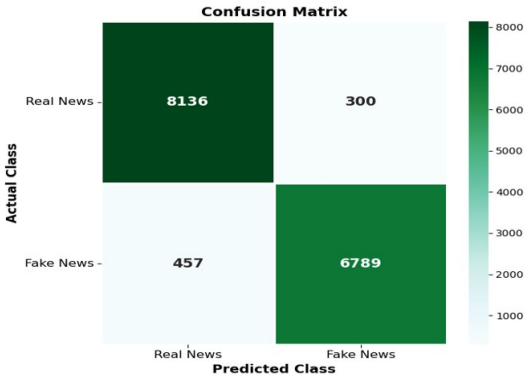


Figure 6. Predicted vs Actual

As shown in Table 3, the model misclassified 457 fraudulent news pieces as real (False Negatives) and 300 actual news articles as fake (False Positives). The very small number of incorrect classifications suggests that the suggested hybrid framework may reliably differentiate between true and deceptive news items.

Table 3. Confusion Matrix of the Proposed Model

Actual Class	Predicted Real	Predicted Fake
Real News	8136	300
Fake News	457	6789

The ROC curve for the suggested TF-IDF + Word2Vec + Stacking Ensemble model is shown in Figure 7. With an AUC value of 0.978, the model demonstrated exceptional capacity to distinguish between authentic and fraudulent news pieces. The ROC curve, which shows the efficacy and resilience of the suggested method for Urdu fake news identification, is located around the top-left corner and stays noticeably above the random classifier line.

The findings show that Urdu news articles' lexical and semantic features are successfully captured by the suggested hybrid framework. The model's capacity to accurately identify instances of bogus news is demonstrated by its high recall, and its balanced performance across classes is confirmed by its great F1-score.

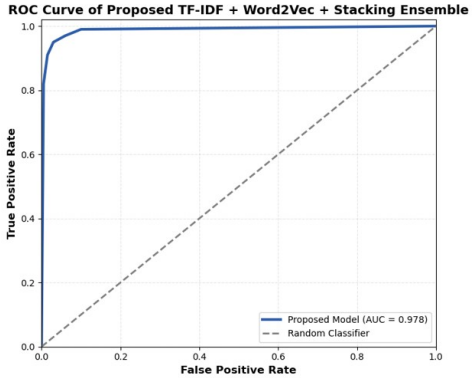


Figure 7. ROC Curve of Proposed TF-IDF + Word2Vec + Stacking Ensemble

This study's three main contributions are: (1) the creation of a hybrid lexical-semantic representation for Urdu fake news detection by combining TF-IDF and Word2Vec features; (2) the application of a stacking ensemble framework that combines the advantages of SVM, Logistic Regression, and XGBoost classifiers; and (3) a thorough assessment on the newly released Hook-and-Bait Urdu dataset, which consists of 78,406 news articles. To the best of our knowledge, only few studies have looked into stacking ensemble learning and hybrid feature fusion for large-scale Urdu fake news detection.

6.1 Reasons for the Improved Performance of the Proposed Method

The proposed method is ensemble learning of TF-IDF and Word2Vec representation using stacking. This improvement over the individual classifiers is consistent with the possibility that these components encode complementary information.
The TF-IDF is the importance of the words and word sequences in the corpus. It can capture discriminative terms and short expressions with unigrams and bigrams that are potentially relevant to real or fake news. Sparse representation can preserve informative lexical patterns that would be otherwise diluted in a dense vector representation.

Word2vec is a method for generating dense word embeddings that reflect distributional relations between words. It can be formed by averaging the word vectors, which concisely represents a sense of lexical similarity and contextual use of words. Combining these vectors with TF-IDF could allow the classifier to exploit both the importance information of words and distributed representations of lexical items.
The stacking ensemble merges the predictions of a set of base classifiers using a logistic-regression meta-learner. The base classifiers could learn different decision boundaries from the same representation of features. If the prediction errors are different enough, the meta-learner can exploit the complementary predictions to improve the final classification result. This provides a good explanation for the improvement over the best individual classifier, XGBoost.

The reported comparison however does not isolate the contribution of feature fusion and stacking to the improvement. To evaluate the contribution of each component, an ablation study would be needed for TF-IDF alone, Word2Vec alone, combined features with single classifiers, and the full stacking model. The results thus support the performance of the whole proposed framework, whereas the specific contribution of each component is to be measured separately.

7. CONCLUSION

The proposed study is a hybrid feature-fusion framework for Urdu fake-news detection, which combines TF-IDF and Word2Vec representations with a stacking ensemble classifier. TF-IDF extracts statistically important unigram and bigram features. Word2Vec learns distributed semantic representations by using the Skip-Gram architecture. For document level Word2Vec vectors we average embedding of individual words. The resulting representations are concatenated into a single feature vector that contains complementary lexical and semantic information.

The fused features are evaluated using Logistic Regression, Random Forest, Support Vector Machine (SVM) and XGBoost. We proposed a stacking ensemble classifier for the final prediction with SVM, Logistic Regression and XGBoost as the base classifiers and Logistic Regression as the meta-classifier. Experiments were performed on the Hook-and-Bait Urdu dataset comprising 78,406 news articles; the train-test split was 80:20 stratified ratio.

The proposed ensemble attained 95.17% accuracy, 95.77% precision, 93.69% recall, 94.72% F1-score and a ROC-AUC of 0.978, higher than individual classifiers compared in this study. The findings indicate that the combination of frequency-based and distributed word representations with stacking ensemble learning is a promising way for Urdu fake-news detection. However, its robustness and cross-domain generalizability need to be

established by evaluation on independent datasets. Future work will explore transformer-based contextual embeddings, explainable AI, evidence-based verification and multilingual evaluation to increase detection reliability and practical applicability.

Limitations & Future Work

The evaluation was performed only on the Hook-and-Bait Urdu dataset.

First, Future work should evaluate generalizability using independent datasets, diverse news sources, and different time periods.

Second, TF-IDF and Word2Vec may not fully capture contextual meaning, sarcasm, and complicated patterns of misinformation. Future studies can explore Urdu transformer-based models and compare their performance with the proposed framework.

Thirdly, the contribution of the components and the stability of the results could be assessed using ablation studies, repeated experiments, confidence intervals and statistical tests. Other possible improvements are source credibility, user engagement, social context and evidence of reliable source.

Finally, future work should focus on explainability, computational efficiency and performance under shifting news distributions. In practice we also need to consider false positives, privacy and oversight by humans.

Lastly, the proposed framework is a promising baseline of Urdu fake-news detection with 95.17% accuracy on the tested set. Further validation, context and evidence-based approaches are needed to validate its wider applicability.

Acknowledgement
The author gratefully acknowledges the financial support provided by Maulana Azad National Urdu University (MANUU), Hyderabad, India, under the Minor Research Project (2025).

Conflict of Interest
No conflict of interest

REFERENCES:

[1] Hu, B.; Mao, Z.; Zhang, Y. An Overview of Fake News Detection: From a New Perspective. *Fundam. Res.* **2025**, *5*(1), 332–346. https://doi.org/10.1016/j.fmre.2024.01.018.

[2] Nazar, M.Z.; Shah, S.U.; Khan, M.T. Fake News Detection in Urdu. In *Workshop Proceedings of the 19th International AAAI Conference on Web and Social Media (ICWSM Workshops)*; AAAI Press: Washington, DC, USA, 2025; Vol. 10, No. 51.

[3] Humayoun, M. The 2021 Urdu Fake News Detection Task Using Supervised Machine Learning and Feature Combinations. *arXiv* **2022**, arXiv:2204.03064. Available online: https://arxiv.org/abs/2204.03064.

[4] Ali, M.Z.; Wang, Y.; Pfahringer, B.; Smith, T.C. Detection of Human and Machine-Authored Fake News in Urdu. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*; Association for Computational Linguistics: Vienna, Austria, 2025; pp. 3419–3428. https://doi.org/10.18653/v1/2025.acl-long.170.

[5] Farooq, M.S.; Khan, W.; Ahmad, A.; Alhakami, W.; Alturise, F. Fake News Detection in Urdu Language Using Machine Learning. *PeerJ Comput. Sci.* **2023**, *9*, e1353. https://doi.org/10.7717/peerj-cs.1353.

[6] Harris, S.; Khan, M.R.; Junaid, M.; Hussain, S.; Ahmad, J. Benchmarking Hook and Bait Urdu News Dataset for Domain-Agnostic and Multilingual Fake News Detection Using Large Language Models. *Sci. Rep.* **2025**, *15*(1), 15553. https://doi.org/10.1038/s41598-025-99511-8.

[7] Manzoor, I.; Singla, S.; Nikita, N. Fake News Detection Using Machine Learning Approaches: A Systematic Review. In *Proceedings of the International Conference on Outcome Based Engineering Education (ICOEI)*; IEEE: Chennai, India, 2019; pp. 230–234. https://doi.org/10.1109/ICOEI.2019.8862770.

[8] Salton, G.; Buckley, C. Term-Weighting Approaches in Automatic Text Retrieval. *Inf. Process. Manag.* **1988**, *24*(5), 513–523. https://doi.org/10.1016/0306-4573(88)90021-0.

[9] Mikolov, T.; Chen, K.; Corrado, G.; Dean, J. Efficient Estimation of Word Representations in Vector Space. *arXiv* **2013**, arXiv:1301.3781. Available online: https://arxiv.org/abs/1301.3781.

[10] Khan, Z.A.; Rekha, V. Fake News Detection Using TF-IDF Weighted with Word2Vec: An Ensemble Approach. *Int. J. Intell. Syst. Appl. Eng.* **2023**, *11*(3), 1065–1076.
[11] Alzubaidi, A.A.; Alawady, A.A. SBERT-Based Stacking Ensemble Model for Fake News Detection. Int. J. Adv. Comput. Sci. Appl. 2025, 16(11). https://doi.org/10.14569/IJACSA.2025.0161149.
[12] Harris, S.; Khan, M.R.; Junaid, M.; Hussain, S.; Ahmad, J. Multi-Domain Urdu Fake News Detection Using Pre-Trained Ensemble Model. *Sci. Rep.* **2025**, *15*(1), 8705. https://doi.org/10.1038/s41598-025-92985-7.
[13] Asghar, S.; Habib, A.; Habib, S.; Khan, A.; Ali, R. Fake News Detection Using Machine Learning: A Survey. *IEEE Access* **2021**, *9*, 57613–57639. https://doi.org/10.1109/ACCESS.2021.3071401.
[14] Azam, A.; Batool, F.; Lo Re, G.; Morana, M.; Khan, M.U.G. Generative-AI vs. Traditional Methods for Fake News Detection. In *Proceedings of the International Conference on Computing, Management and Telecommunications (ComManTel)*; IEEE: New York, NY, USA, 2025; pp. 158–163. https://doi.org/10.1109/ComManTel68363.2025.11368534.
[15] Tanvir, A.A.; Mostafiz, E.S.M.; Choudhury, M.A. Detecting Fake News Using Machine Learning and Deep Learning Algorithms. In *Proceedings of the 2nd International Conference on Smart Systems and Inventive Technology (ICSSIT)*; IEEE: Tirunelveli, India, 2019; pp. 133–137. https://doi.org/10.1109/ICSSIT46314.2019.8987839.
[16] Zhao, Z.; Zhao, J.; Sano, Y.; Levy, O.; Takayasu, H.; Takayasu, M.; Li, H. Fake News Propagation Modeling: A New Framework. *IEEE Access* **2020**, *8*, 123–135.
[17] Asghar, S.; Mahmood, S.; Kamran, H. Fake News Detection Using Machine Learning: A Survey. *IEEE Access* **2021**, *9*, 57613–57639. https://doi.org/10.1109/ACCESS.2021.3075392.
[18] Mondal, S.; Srinivasan, P.; Ahammed, M.E. Fake News Detection: A Comparative Study of Machine Learning Techniques. In *Proceedings of the International Conference on Computational Intelligence, Networks and Security (ICCINS)*; IEEE: Mylavaram, India, 2023; pp. 1–6. https://doi.org/10.1109/ICCINS58907.2023.10450148.
[19] Ma, L.; Zhang, Y. Using Word2Vec to Process Big Text Data. In *Proceedings of the IEEE International Conference on Big Data*; IEEE: Santa Clara, CA, USA, 2015. https://doi.org/10.1109/BigData.2015.7364114.
[20] Evrard, M.M.A.M.; Wang, Q. A Multi-Stage Ensemble Framework for Fake News Detection: Integrating Traditional Machine Learning, Deep Learning, and Advanced Feature Engineering. *Int. J. Sci. Res. Manag.* **2026**, *14*(3). https://doi.org/10.18535/ijsrm/v14i03.ec05.