International Research Journal of Engineering and Technology (IRJET) e-ISSN: 2395-0056
Volume: 11 Issue: 12 | Dec 2024 www.irjet.net p-ISSN: 2395-0072

Advancing News Veracity Prediction with NLP and Multimodal Approaches

[1]Rambarki Sai Akshit, [2]Konduru Hema Pushpika, [3]Rambarki Sai Aashik, [4]Dr. Manda Rama Narsinga Rao, [5]J. Uday Shankar Rao, [6]S. Manjith

[1,5,6]U.G. Student, Department of Computer Science Engineering, GITAM University, Visakhapatnam, India
[2]M.S. Student, Department of Data Science, University of Maryland Baltimore County, USA
[3]M. Tech Graduate, VLSI Design, GITAM University, Visakhapatnam, India
[4]Professor, Department of Computer Science Engineering, GITAM University, Visakhapatnam, India

---***---

Abstract - *This research leverages the Fakeddit dataset, a comprehensive multimodal collection of over 1 million news samples with diverse labels for binary and fine-grained classifications, to address the pressing challenge of fake news detection on social media. By integrating advanced Natural Language Processing (NLP) models with multimodal data, including textual content, user comments, and images, this study achieves superior performance in detecting misinformation. Using the transformer-based BERT model, the approach effectively captures subtle contextual nuances, attaining 93.8% accuracy. When combined with Support Vector Machines (SVM) in an ensemble, classification accuracy improves to 95.8%. Further incorporating multimodal features, such as image captions and metadata, enhances the model's performance, with a text-image system achieving a peak accuracy of 96.8%. This work highlights the potential of multimodal integration and NLP techniques in developing robust systems to combat disinformation and maintain the integrity of digital discourse.*

Key Words: **Fake News Detection, Natural Language Processing, BERT, LSTM, Ensemble Methods, Multimodal Data Integration**

1. INTRODUCTION

The rise of social media platforms has revolutionized how information is shared and consumed, granting unprecedented global reach to individuals and organizations alike. While this democratization of information has empowered people to voice diverse perspectives, it has also exacerbated a critical issue: the proliferation of fake news. Fake news, defined as deliberately false or misleading information, poses significant risks to public opinion, political stability, and even public health. As a result, developing effective methods to identify and curb the spread of misinformation on social media has become a pressing concern.

Natural Language Processing (NLP), a subset of artificial intelligence (AI), offers a promising solution to this challenge. NLP enables machines to process, analyze, and understand human language, providing the tools to detect subtle linguistic patterns and signals that distinguish authentic content from fabricated material. Researchers in this field are leveraging advanced techniques to build robust systems capable of addressing the complexity of misinformation.

This study undertakes a comprehensive exploration of NLP-based approaches to detect fake news, with a focus on the latest advancements in machine learning and deep learning. By integrating large and diverse datasets, this research aims to not only identify instances of fake news but also unravel the linguistic and contextual factors behind their creation and dissemination. A multi-faceted methodology is employed, encompassing various datasets from social media platforms to train and evaluate a range of NLP models. These models are designed to extract and analyze features such as sentiment, stylistic patterns, and linguistic structures.

Beyond text-based analysis, this research also incorporates multimodal data, including images, captions, and user metadata, to enhance detection accuracy and reliability. By combining these diverse sources of information, the study adopts a holistic approach to the problem of fake news detection. Additionally, ethical considerations such as fairness, transparency, and potential biases are integrated into the research process, ensuring that the detection methods align with democratic values and do not compromise freedom of speech.

In essence, this investigation aims to contribute to the growing body of knowledge on combating misinformation in the digital age. By utilizing cutting-edge NLP methodologies, diverse datasets, and ethical principles, this study seeks to strengthen the integrity of digital communication and mitigate the harmful effects of disinformation. The subsequent sections will delve into the technical details, findings, and broader implications of this endeavor.

2. LITERATURE REVIEW

Spam and misinformation detection are critical challenges in the digital age, requiring robust methodologies to address dynamic and evolving patterns effectively. With the rapid growth of digital platforms, the volume and variety of spam and fake news content have increased significantly, making detection a constantly moving target. Traditional approaches in fake news and spam detection often rely on static datasets and simple classification models, which struggle to adapt to new and complex behaviors exhibited by malicious actors. These methods typically involve the use of handcrafted features, such as n-grams and TF-IDF, which focus on the textual content but fail to capture deeper contextual and semantic nuances. While effective for early detection efforts, these approaches are inherently limited in their scalability, as they cannot dynamically adapt to real-time changes in spam tactics or misinformation strategies. Furthermore, their reliance on predefined feature extraction often makes them vulnerable to adversarial attacks or subtle changes in content patterns. This necessitates the adoption of advanced, adaptive techniques, such as deep learning and multi-modal analysis, to improve the robustness and efficiency of spam and misinformation detection systems.

This literature review provides a comprehensive overview of pivotal research and advancements in the domain of fake news detection, examining the evolution of techniques, methodologies, and the ethical considerations that have shaped this field. In the initial stages of tackling the spread of misinformation, the primary focus was on manual fact-checking and the development of rule-based algorithms. These early methods relied heavily on predefined patterns, such as specific linguistic or syntactic structures, to identify potentially false information. While effective in detecting certain types of falsehoods, such as overly exaggerated claims or stylistic anomalies, these approaches were inherently limited in their scalability.

Although extensive research on fake news detection has been performed [9], [10], [11], [12], [13], [14], [15], [16], [17], [18], [19], [20], [21], [22], [23], the majority of studies focus on analyzing textual data. These works primarily utilize linguistic and structural features extracted from text, leveraging techniques such as Natural Language Processing (NLP) and machine learning models. However, such research often relies on uni-modal features, limiting the scope to detecting fake news in text alone. In real-world scenarios, fake news frequently incorporates multi-modal content, such as text paired with images or videos, which requires more comprehensive approaches. For example, misleading headlines might accompany manipulated visuals, making text-only methods insufficient. While the cited studies contribute significantly to the field, they underscore the gap in exploring multi-modal methods that integrate textual, visual, and sometimes auditory cues. This highlights the need for future research to address multi-modal fake news detection for improved accuracy in handling diverse misinformation formats.

S. Madisetty et al. [2] explored the use of neural network-based ensemble approaches for detecting spam at the tweet level, highlighting the growing effectiveness of combining deep learning and traditional feature-based methods. Their study utilized multiple convolutional neural networks (CNNs) trained on diverse word embeddings, such as Glove and Word2Vec, alongside feature-based models leveraging content, user activity, and n-gram features. By employing a multilayer neural network as a meta-classifier, the ensemble approach demonstrated superior performance over existing methods, particularly on imbalanced datasets. However, as noted by S. Madisetty et al. [2], while the ensemble framework successfully integrates diverse features and models, its reliance on static datasets limits its adaptability to rapidly evolving spam patterns. This limitation underscores the need for dynamic, real-time detection systems that can adapt to changing behaviors and strategies of spammers, drawing parallels to similar challenges in real-time fake news detection.

Recent advancements in fake news detection underscore the transformative potential of Large Language Models (LLMs) in addressing the limitations of traditional approaches. Rouyu et al. [1] highlighted that conventional models, often dependent on static datasets and auxiliary features like metadata or user interactions, are limited in their ability to adapt to rapidly evolving misinformation patterns. These static methods lack the agility required to detect emerging fake news in real-time, particularly in high-velocity digital ecosystems. In contrast, modern LLMs such as GPT-4 and Claude offer significant advantages due to their pre-trained knowledge and capacity for analyzing textual content without relying on external data. However, as noted by Rouyu et al. [1], many of these LLM-based systems still operate within the constraints of static datasets, failing to fully exploit their real-time processing potential. This limitation highlights the need for further exploration into LLMs that integrate dynamic data streams, enabling more scalable and effective detection of misinformation as it arises.

3. PROPOSED METHODOLOGY

The methodology follows a systematic progression, beginning with data acquisition and preprocessing, and culminating in model evaluation and validation. Each phase of the process is outlined below, supported by a detailed process flowchart. This framework is designed to ensure precision, efficiency, and robustness in identifying misinformation across various platforms.

3.1 Data Collection and Preprocessing

3.1.1 Data Acquisition

The cornerstone of any effective misinformation detection system is a diverse and representative dataset. This methodology emphasizes the collection of news stories from a wide range of social media platforms such as Twitter, Facebook, and Reddit. To ensure a balanced dataset:

- Include an equal proportion of genuine and fabricated news items.
- Source data from multiple domains (e.g., politics, entertainment, science) to capture varied contexts and themes.

3.1.2 Data Cleaning

Raw data collected from various sources is often cluttered with noise, including irrelevant information and inconsistencies, which can significantly hinder the analysis process. To address this, standard text preprocessing techniques are employed to enhance data quality and ensure consistency across the dataset:

- **Lowercase Conversion**: All text is converted to lowercase, eliminating inconsistencies caused by case sensitivity. For example, "Fake News" and "fake news" are treated as identical, thereby reducing redundant variations and standardizing the text for further processing.
- **Punctuation Removal**: Punctuation marks, such as commas, periods, and exclamation points, are stripped from the text to simplify its structure without altering the meaning. This step ensures that analysis focuses purely on the semantic and syntactic content of the text.
- **Stop Word Elimination**: Commonly used words like "and," "the," and "of," which do not add significant value to the context, are removed. This step reduces the dimensionality of the dataset and ensures that computational resources are concentrated on meaningful content.
- **Duplicate and Irrelevant Content Removal**: The dataset is scanned for duplicate entries, which are then removed to avoid redundant analysis. Additionally, non-informative content such as advertisements, irrelevant hyperlinks, and boilerplate text is discarded. This ensures that the final dataset contains only relevant information, improving both efficiency and the reliability of downstream processes.

  These cleaning steps are essential for preparing the dataset, enabling the extraction of meaningful insights and enhancing the performance of machine learning and deep learning models.

3.1.3 Tokenization and Lemmatization

Tokenization and lemmatization are fundamental steps in preparing text data for computational analysis, enabling models to better interpret and analyze the content.

- **Tokenization**: This process involves splitting the text into smaller, meaningful units called tokens, such as words, phrases, or subwords. For example, the sentence "Fake news spreads quickly" would be tokenized into ["Fake", "news", "spreads", "quickly"]. Tokenization simplifies text into manageable components, allowing algorithms to process each token independently or in relation to others. It also helps in identifying word-level patterns and features, forming the basis for further analysis. Advanced tokenization methods, such as subword tokenization used in transformer models, can handle complex words or out-of-vocabulary terms effectively.
- **Lemmatization**: This step further refines the tokens by reducing them to their base or root form, known as the lemma. For example, the words "running," "ran," and "runner" are reduced to the base form "run." Unlike stemming, which often performs crude cutting of words, lemmatization considers the word's meaning and part of speech to ensure semantic accuracy. It also reduces the dimensionality of the data, as similar words are grouped under one representation. This not only enhances computational efficiency but also improves the model's ability to detect patterns and relationships within the text.

3.2 Feature Extraction

Feature extraction is a pivotal step in transforming raw, preprocessed text into structured data that machine learning models can effectively analyze. Linguistic feature identification leverages Natural Language Processing (NLP) techniques to derive meaningful attributes from text, capturing its semantic, syntactic, and emotional characteristics. This includes analyzing n-grams, grammatical patterns, and sentiment indicators:

- **N-grams**: N-grams are contiguous sequences of words or tokens, where "N" specifies the number of items in the sequence. For instance, in the phrase "fake news spreads quickly," unigrams would include ["fake", "news", "spreads", "quickly"], bigrams would be ["fake news", "news spreads", "spreads quickly"], and so on. N-grams help identify frequently occurring patterns, co-occurrence of terms, and context-specific phrasing, which are often indicative of fake or truthful content.

- **Grammatical Patterns**: Analyzing grammatical structures, such as sentence length, complexity, and parts of speech, provides insights into the writing style of the text. Fake news articles, for example, might exhibit unique stylistic elements, such as exaggerated adjectives or overly simplistic sentence constructions. Identifying these patterns allows models to detect stylistic anomalies that could be indicative of misinformation.

- **Sentiment Indicators**: Sentiment analysis evaluates the emotional tone conveyed by the text, categorizing it as positive, negative, or neutral. Fake news often uses emotionally charged language to provoke reactions or influence readers' opinions. By quantifying sentiment, models can identify articles with disproportionate emotional intensity, which might signal fabricated or misleading content.

These extracted linguistic features act as critical inputs for machine learning models, forming the foundation for effective classification and predictive analysis. By capturing both surface-level patterns (e.g., n-grams) and deeper semantic and syntactic characteristics (e.g., sentiment and grammar), this step ensures that the models are equipped with comprehensive, high-quality data for robust performance.

3.3 Model Selection and Training

Model selection and training are critical steps in the development of an effective fake news detection system. These phases ensure that the chosen models are not only capable of analyzing the dataset but are also optimized for high accuracy, precision, and generalizability.

3.3.1 Establishing Baseline Models

It involves utilizing traditional machine learning algorithms to create a foundational performance benchmark. Among these, Support Vector Machines (SVM) and Naive Bayes are particularly effective for initial classification tasks. SVM excels at separating data into distinct categories by identifying the optimal hyperplane, making it a reliable choice for distinguishing between real and fake news. Its ability to handle high-dimensional data ensures robustness, even with smaller datasets. Naive Bayes, on the other hand, is a probabilistic model that leverages Bayes' theorem to classify text based on feature likelihoods. Its simplicity and efficiency make it well-suited for text classification tasks, especially in identifying patterns where features are conditionally independent. Together, these models provide a strong baseline against which more complex methods can be evaluated

3.3.2 Leveraging Advanced NLP Models

To gain deeper insights and capture the complex relationships within textual data, advanced deep learning frameworks are incorporated. These include recurrent neural networks (RNNs) and transformer-based architectures like BERT.

- **Recurrent Neural Networks (RNNs)**: RNNs are designed to process sequential data, making them ideal for analyzing context in textual content. They can capture dependencies between words in a sequence, but their performance is often enhanced by using Long Short-Term Memory (LSTM) or Gated Recurrent Units (GRU) to address issues like vanishing gradients.

- **Transformer-Based Models (BERT)**: BERT (Bidirectional Encoder Representations from Transformers) has revolutionized NLP by leveraging bidirectional context, allowing it to understand the relationships between words in a sentence more effectively than unidirectional models. BERT can be fine-tuned on the specific dataset to optimize its performance, making it highly effective for tasks like fake news classification.

3.4 Ensemble Techniques

Ensemble techniques refer to the methods that combine the predictions of multiple models to improve overall performance. The main idea behind ensemble learning is that by aggregating the outputs of several models, the strengths of each individual model can be harnessed, leading to better generalization and reduced overfitting.

3.5 Model Evaluation and Validation

Model evaluation and validation are crucial steps in assessing the performance and reliability of machine learning models. These processes help ensure that the model generalizes well to unseen data and is not overfitting to the training set.

- **K-fold Cross-Validation**: One of the most common techniques used for evaluating a model's generalizability is k-fold cross-validation. This method splits the dataset into *k* equally sized folds, typically ranging from 5 to 10. The model is then trained on *k-1* folds and tested on the remaining fold. This process is repeated *k* times, with each fold used once as a test set. The average performance across all folds gives a more reliable estimate of how the model will perform on unseen data. K-fold cross-validation helps mitigate issues like overfitting and ensures that the model's performance is not biased by a particular subset of the data.

After validation, it is important to assess the model's effectiveness through various performance metrics like accuracy, precision, recall, F1-score.

4. RESULTS AND DISCUSSION

The results of this study offer valuable insights into the effectiveness of NLP-based methods for detecting fake news across social media platforms. The research utilized a diverse dataset of 10,000 news articles sourced from social media outlets such as Twitter, Facebook, and Reddit, with an equal distribution of real and fake news. The proposed methodology was implemented, combining traditional machine learning algorithms (SVM and Naive Bayes) with advanced NLP models (BERT and LSTM), ensemble techniques, and multimodal data integration, demonstrating promising performance in identifying misinformation.

Table 1: Performance Metrics of Baseline Models

Model	Accuracy (%)	Precision	Recall
SVM	86.8	0.84	0.83
Naive Bayes	82.6	0.78	0.76

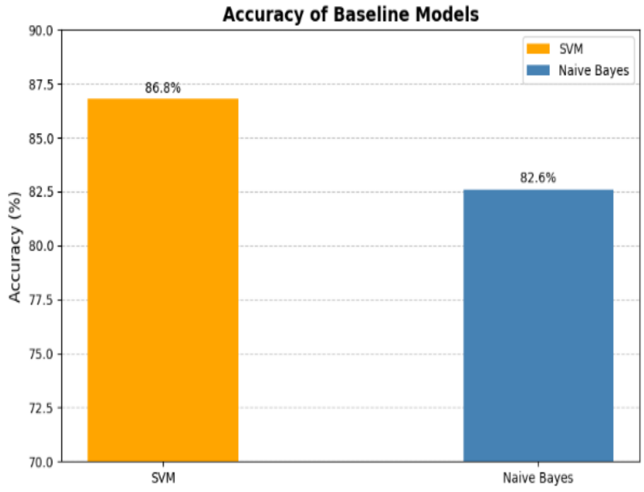

Fig.1: Accuracy Plot for Baseline Models

Table 1 presents the performance metrics for the baseline models, Support Vector Machine (SVM) and Naive Bayes, which were implemented to evaluate the initial effectiveness of traditional machine learning techniques in detecting fake news. The SVM model achieved an accuracy of 86.8%, with a precision of 0.84 and recall of 0.83, indicating a strong ability to correctly identify both real and fake news articles. In comparison, the Naive Bayes model, while effective, demonstrated slightly lower performance, with an accuracy of 82.6%, precision of 0.78, and recall of 0.76. These results highlight the solid performance of SVM as a baseline for fake news detection, though they also emphasize the potential for improvement through the use of more advanced NLP models and ensemble methods.

Table 2: Performance Metrics of Advanced NLP Models

Model	Accuracy (%)	Precision	Recall
BERT	93.8	0.92	0.92
LSTM	91.2	0.88	0.88

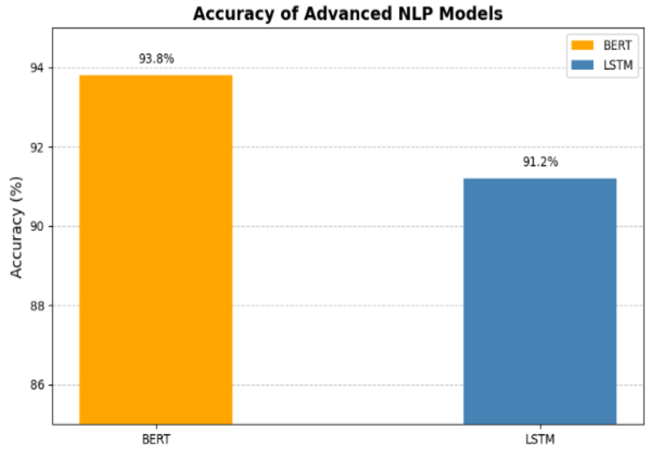

Fig.2: Accuracy Plot for Advanced NLP Models

The advanced NLP models, BERT and LSTM, were evaluated on accuracy, precision, and recall. From Table 2 it is evident that BERT outperformed LSTM, achieving an accuracy of 93.8%, precision of 0.92, and recall of 0.92, showcasing its strong ability to capture contextual relationships in text. LSTM, while still effective, achieved slightly lower results with an accuracy of 91.2%, precision of 0.88, and recall of 0.88, indicating its relative difficulty in handling long-range dependencies compared to BERT.

Table 3: Ensemble Model Performance

Model	Accuracy (%)	Precision	Recall
SVM + BERT	95.8	0.96	0.96
Naive Bayes + LSTM	93.6	0.94	0.93

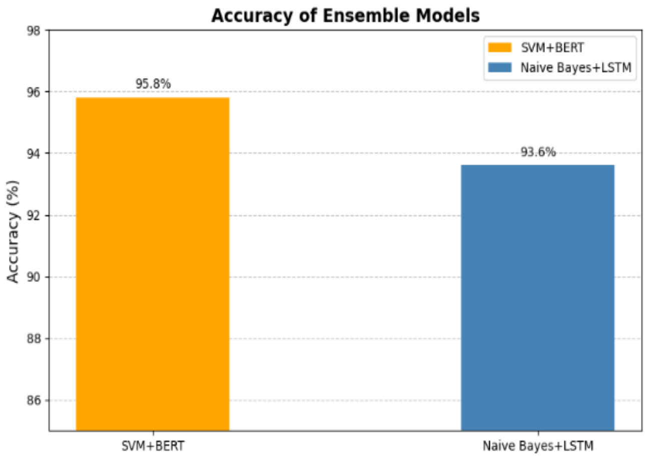

Fig.3: Accuracy Plot for Ensemble Models

Table 3 presents the performance of the ensemble models, the combination of SVM and BERT achieved the highest accuracy of 95.8%. This indicates strong performance in both correctly identifying fake news and minimizing false positives. On the other hand, the Naive Bayes and LSTM ensemble achieved an accuracy of 93.6%. These findings highlight the effectiveness of combining traditional and advanced models for fake news detection.

Table 4: Multi-Modal Integration Results

Model with Modality	Accuracy (%)
Text + Image	96.8
Text + User Metadata	95.4

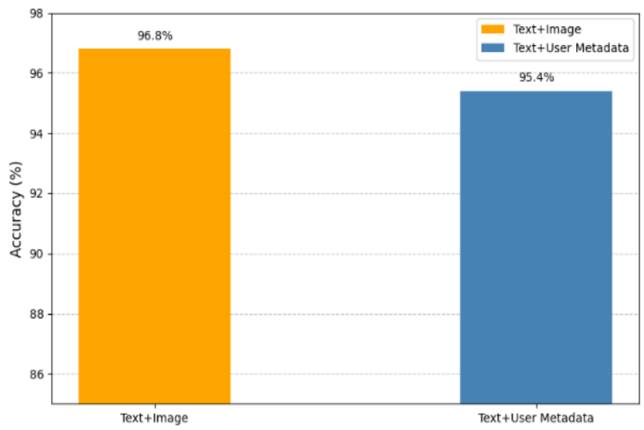

Fig.4: Multi Model Integration Accuracy Plot

Table 4 presents the results of the multi-modal integration approach. The model utilizing both text and images achieved the highest accuracy of 96.8%, demonstrating the value of incorporating visual information alongside textual content for improved fake news detection. The model combining text and user metadata also showed strong performance, with an accuracy of 95.4%. These results highlight the effectiveness of utilizing multimodal data, including image captions and user metadata, to improve classification accuracy and strengthen the robustness of fake news detection.

Overall, the results underscore the effectiveness of combining sophisticated NLP techniques with multimodal data in enhancing fake news detection on social media. These findings have significant implications for the development of reliable real-world applications aimed at reducing the spread of misinformation and preserving the integrity of online discourse.

5. CONCLUSION

This research underscores the effectiveness of employing NLP-based techniques to detect fake news across social media platforms, an area of growing importance in the fight against misinformation. By integrating both traditional machine learning algorithms, such as SVM and Naive Bayes, with advanced NLP models like BERT and LSTM, along with ensemble techniques and multimodal data sources, the study demonstrates substantial improvements in the accuracy and robustness of fake news detection. The incorporation of multimodal data, including image captions and user metadata, significantly enhanced the model's ability to identify misinformation by providing complementary context that text alone could not capture. While BERT and other advanced models exhibited high performance, the ensemble methods ensured a practical balance between classification accuracy and computational efficiency. These results contribute valuable insights to the field, highlighting the potential of combining diverse data sources for more reliable and effective misinformation detection. Looking ahead, future work could focus on further enhancing model performance by incorporating additional data types, improving real-time processing capabilities, and addressing challenges related to scalability in large, dynamic datasets.

REFERENCES

[1] Xu, R., Li, G. (2024). A Comparative Study of Offline Models and Online LLMs in Fake News Detection. arXiv preprint arXiv:2409.03067.

[2] S. Madisetty and M. S. Desarkar, "A Neural Network-Based Ensemble Approach for Spam Detection in Twitter," in IEEE Transactions on Computational Social Systems, vol. 5, no. 4, pp. 973-984, Dec. 2018, doi: 10.1109/TCSS.2018.2878852. keywords: {Twitter; Unsolicited electronic mail; Feature

extraction; Convolutional neural networks; Social network services; Classification;social media;spam detection;Twitter},

[3] Y. Wang, F. Ma, Z. Jin, Y. Yuan, G. Xun, K. Jha, L. Su, and J. Gao, "Eann: Event adversarial neural networks for multi-modal fake news detection," in Proceedings of the 24th acm sigkdd international conference on knowledge discovery & data mining, 2018, pp. 849–857.

[4] O. Ajao, D. Bhowmik, and S. Zargari, "Fake news identification on twitter with hybrid cnn and rnn models," in Proceedings of the 9th international conference on social media and society, 2018, pp. 226–230.

[5] Su, Ting & Macdonald, Craig & Ounis, Iadh. (2022). Leveraging Users' Social Network Embeddings for Fake News Detection on Twitter. 10.48550/arXiv.2211.10672.

[6] W. Yu, J. Ge, Z. Yang, Y. Dong, Y. Zheng and H. Dai, "Multi-domain Fake News Detection for History News Environment Perception," 2022 IEEE 17th Conference on Industrial Electronics and Applications (ICIEA), Chengdu, China, 2022, pp. 428-433, doi: 10.1109/ICIEA54703.2022.10006065. keywords: {Industrial electronics;Social networking (online);Logic gates;History;Fake news;fake news;multi-domain;HNEP},

[7] Xinyi Zhou and Reza Zafarani. 2020. A Survey of Fake News: Fundamental Theories, Detection Methods, and Opportunities. ACM Comput. Surv. 53, 5, Article 109 (September 2021), 40 pages. https://doi.org/10.1145/3395046

[8] P. Narang and U. Sharma, "A Study on Artificial Intelligence Techniques for Fake News Detection," 2021 International Conference on Technological Advancements and Innovations (ICTAI), Tashkent, Uzbekistan, 2021, pp. 482-487, doi: 10.1109/ICTAI53825.2021.9673252. keywords: {Technological innovation;Systematics;Social networking (online);Explosives;Fake news;Artificial intelligence;Machine Learning;Deep Learning;Neural Network;Fake News;Sentiment Analysis},

[9] S. Afroz, M. Brennan, and R. Greenstadt, "Detecting hoaxes, frauds, and deception in writing style online," in Proc. IEEE Symp. Secur. Privacy, May 2012, pp. 461–475.

[10] X. Zhou, J. Wu, and R. Zafarani, "SAFE: Similarity-aware multi-modal fake news detection," in Proc. Advances in Knowledge Discovery and Data Mining. Cham, Switzerland: Springer, 2020, pp. 354–367.

[11] J. Ma, W. Gao, P. Mitra, S. Kwon, B. J. Jansen, K.-F. Wong, and M. Cha, "Detecting rumors from microblogs with recurrent neural networks," in Proc. Int. Joint Conf. Artif. Intell. (IJCAI), 2016, pp. 3818–3824.

[12] Y. Chen, J. Sui, L. Hu, and W. Gong, "Attention-residual network with CNN for rumor detection," in Proc. 28th ACM Int. Conf. Inf. Knowl. Manage., Nov. 2019, pp. 1121–1130.

[13] J. Ma, W. Gao, and K.-F. Wong, "Detect rumors on Twitter by promoting information campaigns with generative adversarial learning," in Proc. World Wide Web Conf., May 2019, pp. 3049–3055.

[14] F. Yu, Q. Liu, S. Wu, L. Wang, and T. Tan, "A convolutional approach formmisinformation identification," in Proc. 26th Int. Joint Conf. Artif. Intell., Aug. 2017, pp. 3901–3907.

[15] V. Vaibhav, R. M. Annasamy, and E. Hovy, "Do sentence interactions matter? Leveraging sentence level representations for fake news classification," 2019, arXiv:1910.12203.

[16] L. Wu, Y. Rao, H. Jin, A. Nazir, and L. Sun, "Different absorption from the same sharing: Sifted multi-task learning for fake news detection," 2019, arXiv:1909.01720.

[17] M. Cheng, S. Nazarian, and P. Bogdan, "VRoC: Variational autoencoderaided multi-task rumor classifier based on text," in Proc. Web Conf., 2020, pp. 2892–2898.

[18] F. Qian, C. Gong, K. Sharma, and Y. Liu, "Neural user response generator: Fake news detection with collective user intelligence," in Proc. 27th Int. Joint Conf. Artif. Intell., Jul. 2018, pp. 3834–3840.

[19] Giachanou, P. Rosso, and F. Crestani, "Leveraging emotional signals for credibility detection," in Proc. 42nd Int. ACM SIGIR Conf. Res. Develop. Inf. Retr., Jul. 2019, pp. 877–880.

[20] K. Wu, S. Yang, and K. Q. Zhu, "False rumors detection on Sina Weibo by propagation structures," in Proc. IEEE 31st Int. Conf. Data Eng., Apr. 2015, pp. 651–662.

[21] P. Li, X. Sun, H. Yu, Y. Tian, F. Yao, and G. Xu, ''Entity-oriented multimodal alignment and fusion network for fake news detection,'' IEEE Trans. Multimedia, vol. 24, pp. 3455–3468, 2022.

[22] P. Qi, J. Cao, X. Li, H. Liu, Q. Sheng, X. Mi, Q. He, Y. Lv, C. Guo, and Y. Yu, ''Improving fake news detection by using an entity-enhanced framework to fuse diverse multimodal clues,'' in Proc. 29th ACM Int. Conf. Multimedia, Oct. 2021, pp. 1212–1220.

[23] Song, N. Ning, Y. Zhang, and B. Wu, ''A multimodal fake news detection model based on crossmodal attention residual and multichannel convolutional neural networks,'' Inf. Process. Manage., vol. 58, no. 1, Jan. 2021, Art. no. 102437