

# A Scalable Machine Learning Model for Enhancement of Fake News Detection

M. M. Arinzechi<sup>1</sup>, M. Ibrahim<sup>1</sup>, K. Haruna<sup>1</sup>, M. S. Daniel<sup>2</sup>, A. Ahmed<sup>1</sup>, R. Mustapha<sup>1</sup>, S. Abdulkari<sup>1</sup>  
and A. Muazu<sup>1</sup>

<sup>1</sup>Kaduna state university,

<sup>2</sup>Kaduna State Residents Identity Management Agency.

[arinzechimicheal@gmail.com](mailto:arinzechimicheal@gmail.com)<sup>1</sup>

## Abstract

*The rise of fake news has become a critical issue in the digital era, especially on social media platforms where information spreads rapidly and widely. Fake news, defined as deliberately misleading or false information presented as factual news, poses threats to public health, democracy, and economic stability. While machine learning offers scalable solutions for detecting such content, many existing models face challenges related to scalability, portability, and data quality. This study addresses these limitations by developing a scalable machine learning model to enhance fake news detection across digital platforms. The ISOT Fake News Dataset was employed, and extensive data preprocessing was conducted, including text normalization, lemmatization, and TF-IDF vectorization. Four machine learning algorithms: Logistic Regression, Naive Bayes, Support Vector Machine (SVM), and Random Forest were implemented and evaluated. Exploratory Data Analysis (EDA) was also introduced to uncover data patterns and guide model development. The results showed that SVM and Random Forest achieved superior accuracy (1.00), outperforming Logistic Regression (0.99) and Naive Bayes (0.93). The findings suggest that ensemble-based and kernel-based models are more robust for high-dimensional text classification tasks. Moreover, EDA contributed to improved model interpretability and performance. In conclusion, this study demonstrates that a carefully engineered, scalable approach can significantly improve fake news detection accuracy. Future work will explore deep learning models, cross-lingual datasets, and real-time deployment frameworks to enhance model generalizability and responsiveness.*

**Keywords:** Fake News Detection, Machine Learning Models, Text Classification, Exploratory Data Analysis.

## 1. Introduction

In recent years, the rapid expansion of digital platforms, particularly social media, has drastically reshaped how information is created, shared, and consumed. Social media has democratized information access, allowing people to engage with news in real-time. However, this shift has fueled the spread of fake news, defined as intentionally false or misleading information presented as legitimate news, often to deceive or manipulate audiences (Sadiq et al., 2023). The algorithms of social media platforms have intensified the spread of fake news. These networks prioritize highly engaging content, often sensational or polarizing, to generate more likes, shares, and comments. This structure favors fake news, which elicits strong emotional responses and circulates widely. Additionally, echo chambers reinforce this effect by exposing users primarily to information that aligns with their beliefs, allowing fake news to spread unchecked and deepening societal divides (Chen et al., 2023). The social impacts of fake news are significant, particularly

in politics. For example, during the 2016 U.S. presidential election, fake news widely circulated on social media, causing confusion and raising concerns about the integrity of democratic processes (Allcott et al., 2020). Beyond politics, fake news also endangers public health. During COVID-19, misinformation about the virus and vaccines led to harmful behaviors, such as vaccine hesitancy. This phenomenon, which the World Health Organization termed an “infodemic,” posed a substantial risk to public health efforts (Zarocostas, 2020; Jones et al., 2023; Rahman & Malik, 2024). Economically, fake news can damage businesses by harming reputations or manipulating stock prices through false reports. Companies often incur significant costs managing these repercussions through legal responses and public relations (Gupta et al., 2021; Miller & Torres, 2023). There is a consensus that innovative solutions are needed to combat misinformation. Traditional methods, such as manual fact-checking and media literacy campaigns, have value but are limited in scope. Machine learning (ML) and artificial intelligence (AI) offer scalable alternatives by automating fake news detection and analyzing vast datasets in real-time. Early research demonstrates the potential of techniques such as natural language processing (NLP) and deep learning to identify false information effectively (Zhang et al., 2020; Aslam et al., 2023). Yet challenges remain. The complexity of language, context, and adaptive strategies used by fake news creators often limits current models. Many also struggle with scalability across languages, platforms, and media types (Shu et al., 2021; Kumar & Basu, 2023).

## 2. Related Works

Fake news, though often perceived as a modern problem, has deep historical roots in misinformation and propaganda. Historically, it was employed to manipulate public opinion and justify political or military actions. The invention of the printing press enabled the mass dissemination of propaganda, shaping early public narratives and reinforcing power structures (Tandoc et al., 2020). The term “fake news” gained modern relevance during the 2016 U.S. presidential election, where fabricated content was widely distributed on social media platforms, significantly influencing public perception (Allcott et al., 2020). Algorithms that favor engagement over accuracy inadvertently boosted the reach of sensational and misleading content (Vosoughi et al., 2020), creating an ecosystem in which fake news thrives. With the ubiquity of digital media, the velocity and reach of fake news have intensified. Social media enables rapid, large-scale dissemination of information, making misinformation difficult to contain (Pennycook et al., 2021). Echo chambers—where users are predominantly exposed to ideologically similar views—further entrench misinformation and reduce exposure to corrective content (Chen et al., 2023). The emergence of sophisticated disinformation tools such as deepfakes adds another layer of complexity to detection, as they blend hyperreal visuals and audio with fabricated narratives (Chesney & Citron, 2019; Sadiq et al., 2023). In response, manual fact-checking platforms like Snopes, Politifact, and Full Fact were developed. However, given the scale of online content, these methods are inherently limited in speed and scalability (Graves, 2020). Automated systems are thus necessary to supplement human efforts and provide real-time response capabilities. The impacts of fake news span several domains. Socially, it undermines trust in the media and polarizes public discourse (Pennycook & Rand, 2021). Politically, it has influenced voter decisions and fueled unrest, notably during the 2016 and 2020 U.S. elections and more recently in global elections including Nigeria and India (Okoro et al., 2024). In public health, misinformation about COVID-19 and vaccination contributed to widespread vaccine hesitancy and misinformation-fueled behaviors, with studies showing direct correlations between social media exposure and lowered vaccine uptake (Jones et

al., 2023; Rahman & Malik, 2024). Economically, false reports have led to stock market volatility, reputational damage, and significant corporate losses (Gupta et al., 2021; Miller & Torres, 2023). Traditional detection approaches, such as rule-based heuristics and symbolic reasoning, are often rigid and unable to keep pace with the evolving nature of misinformation (Rubin et al., 2019). In contrast, machine learning (ML) offers scalable and adaptive mechanisms for fake news detection. By training on large-scale text corpora, ML models extract linguistic, syntactic, and semantic features to classify content as real or fake (Shu et al., 2020). These models can be updated regularly and scaled across platforms, making them ideal for dynamic information environments. Advanced ML methods such as ensemble learning, transformers, and neural networks have improved classification accuracy. The CSI model, for example, integrates content, social context, and user behavior to enhance performance (Ruchansky et al., 2021). More recently, BERT, RoBERTa, and GPT-based models have achieved state-of-the-art performance by capturing contextual semantics in news articles (Aslam et al., 2023; Sahoo & Mandal, 2023). Despite this progress, challenges remain. Many models are trained on English datasets, limiting cross-lingual scalability. Others suffer from data quality issues or class imbalance problems that affect predictive performance (Mehrabani et al., 2021). Additionally, fake news frequently blends facts with fiction, making it harder to identify purely through textual cues (Shu et al., 2020). Ethical concerns such as algorithmic bias, censorship, and transparency also require further exploration (Kumar & Basu, 2023; Li et al., 2022). Recent studies continue to expand the landscape. Ahmad et al. (2020) utilized ensemble classifiers, achieving high accuracy but with poor generalization to new data. Goldani et al. (2021) applied capsule neural networks for fake news detection, though scalability remained an issue. Sahoo & Mandal (2023) proposed an interpretable deep learning approach that fuses explainability with high performance, suggesting that transparency and trust must go hand in hand. This study builds upon these foundations and addresses some of the observed gaps. Notably, our baseline reference, Uma S. et al. (2021), compared multiple machine learning models and concluded that Logistic Regression achieved the best performance. However, their study did not integrate Exploratory Data Analysis (EDA), a crucial step for understanding feature behavior and improving model interpretation. Our approach introduces a scalable model that combines robust preprocessing, EDA, and comparative evaluation to enhance performance and generalizability across platforms.

### 3. Methodology

This chapter describes the research methodology for fake news detection using machine learning. It covers data collection, preprocessing, model selection, training, evaluation, and techniques for managing imbalanced data, detailing the rationale for each step to ensure accuracy and relevance. Uma et al. (2021) developed a fake news detection model using standard machine learning steps with data sourced from social media, search engines, news sites, and fact-checking sources like BuzzFeed News and LIAR. Key steps in their approach included:

- Data Preprocessing: Data was cleaned through punctuation removal, tokenization, stopword removal, and stemming.
- Feature Extraction: Features like Bag-of-Words, TF-IDF, and N-grams were derived from the preprocessed data.

- Model Training: Naive Bayes, Logistic Regression, and Random Forest classifiers were trained on these features.
- Evaluation: Models were evaluated using F1 scores and confusion matrices.
- Model Selection and Tuning: The top-performing models were selected, and Grid Search CV was applied for parameter optimization.
- Deployment: Logistic Regression was chosen as the best model and saved for classifying fake news.

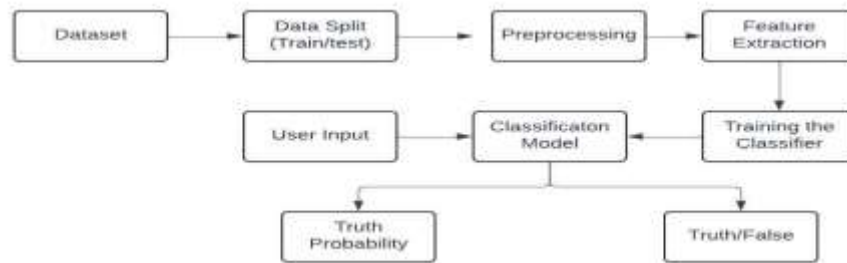


Figure 1: The architectural diagram of the static model

The figure below shows the architectural diagram of the scalable model.

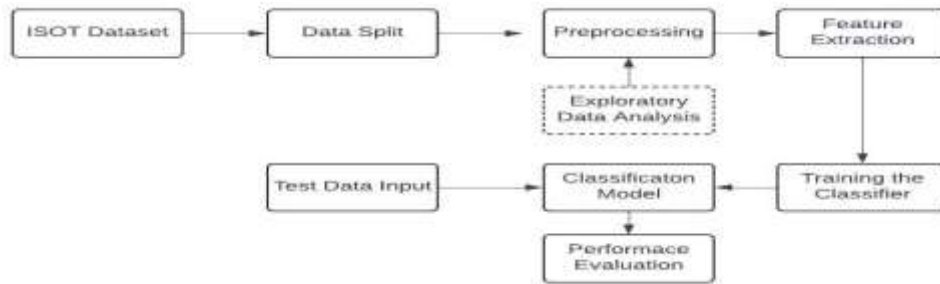


Figure 2: The architectural diagram of the scalable model

The scalable model was developed to mitigate the gap noticed in the existing model stated in fig above. Exploratory data analysis was added to the preprocessing of the dataset to achieve an enhancement on the existing model discussed above. The process shown below were used to achieve this enhancement model. We will discuss below in detail the processes used in this model.

**Data Collection**

The ISOT Fake News Dataset, a widely used resource for misinformation detection, consists of two CSV files: True.csv and Fake.csv. The dataset includes 44,898 articles: 21,417 true and 23,481 fake, offering a balanced representation of both types. Each article contains five columns: title, text, subject, date, and label, with labels indicating whether the article is true (1) or fake (0). The diverse subjects, ranging from politics to global events, and the time-stamped articles, enable temporal analysis and highlight trends in misinformation, such as during elections or crises. This balanced dataset is ideal for training machine

learning models and addressing issues like class imbalance, ensuring unbiased evaluation metrics such as accuracy, precision, recall, and F1-score. The dataset's credibility is reinforced by its use of reliable news sources and fact-checking organizations.

### Data Pre-processing

Data preprocessing ensures that raw text is in an optimal form for machine learning models. First, text normalization was applied by converting all text to lowercase to eliminate case sensitivity. Contraction expansion was performed to standardize contractions (e.g., "can't" to "cannot"). Punctuation and digits were removed to reduce noise, and stopwords (common but uninformative words) were eliminated using NLTK. Tokenization split the text into individual words for analysis, followed by lemmatization, which reduced words to their base form (e.g., "running" to "run"). Finally, TF-IDF vectorization was used to convert the text into numerical features, highlighting words specific to each category (true or fake news). These preprocessing steps streamlined the dataset, reducing noise while preserving key information, preparing it for further analysis and model training.

### Exploratory Data Analysis (EDA)

The Exploratory Data Analysis (EDA) phase serves as a crucial step in uncovering patterns and insights within the dataset, helping to guide the modeling process. In this project, EDA was conducted using a range of visual and statistical techniques to better understand the characteristics of both fake and true news articles. One of the first analyses conducted was the label distribution. A simple bar plot was generated to visualize the distribution of fake and true news articles within the dataset. As the ISOT dataset is relatively balanced, the plot showed that approximately half of the articles were true news, and the other half were fake.




Figure 3: Distribution of News Label

This balance is particularly advantageous in machine learning tasks, as it prevents models from becoming biased toward one class. A balanced dataset ensures that the evaluation metrics such as accuracy, precision, and recall provide a more reliable measure of model performance.

Another key analysis focused on text length. By generating histograms for the length of the text column, the project examined whether the length of articles varied between true and fake news.




Figure 4: Distribution of news text length

The analysis revealed that true news articles tend to be slightly longer than fake news articles. This difference in text length might be attributed to the fact that legitimate news outlets often provide more detailed reporting, while fake news articles may rely on brevity to quickly capture attention. This insight into text length could be used to enhance feature engineering for models, as the length of an article may serve as a useful predictor for distinguishing between true and fake news. The subject column was also explored to identify the topics most commonly associated with fake and true news. A count plot revealed that fake news articles were overwhelmingly concentrated in the subject category of “Politics,” while true news articles were more evenly distributed across various categories such as “World News” and “Technology.” This observation aligns with the well-documented trend that political content is a frequent target of misinformation campaigns.




Figure 5: Distribution of news subject by labels

Understanding the subject distribution not only helps contextualize the dataset but also provides a useful feature for model development. For instance, subject categories could be encoded as categorical features and used as inputs to machine learning models. The date column allowed for a temporal analysis of fake news production. Articles were grouped by year and month, and a line plot was created to visualize the

frequency of fake and true news over time. The analysis showed noticeable spikes in the production of fake news during politically significant periods, such as election years.




Figure 6: Distribution of news by date.

This insight suggests that fake news may be more prevalent during times of heightened political activity, which can inform future work on detecting patterns of misinformation during specific events. Temporal features could be added to the model to account for these fluctuations.

### Model Selection

In this study, four machine learning models were selected for fake news detection based on their effectiveness in text classification: Logistic Regression, Naive Bayes, Support Vector Machine (SVM), and Random Forest. These models were chosen for their scalability, suitability for text data, and proven performance.

### Model Training

The dataset was split into 80% training and 20% test sets to train and evaluate the models. The training set was used to fit the models, while the test set evaluated their performance. This split ensures sufficient data for both training and evaluation while maintaining the distribution of true and fake news.

Each model was trained using the training set, with Logistic Regression, Naive Bayes, and SVM using the TF-IDF vectors. Random Forest was trained by aggregating predictions from multiple decision trees. A fixed random seed (123) was used to ensure reproducibility of results.

### Model Evaluation

After training the models, the next step is to evaluate their performance using a set of predefined metrics. The purpose of model evaluation is to understand how well each model performs in classifying news articles as either fake or true. The metrics used in this project include accuracy, precision, recall, F1-score, and a confusion matrix. Each of these metrics provides a different perspective on the model's performance, ensuring that the evaluation is comprehensive and considers various aspects of classification.

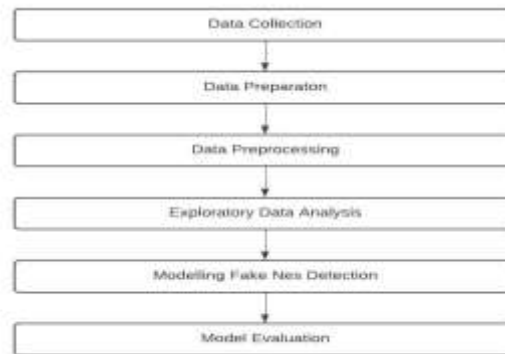


Figure 7: Outline of Methodology

#### 4. Results and Discussion

This section presents the results of four machine learning models, Logistic Regression, Naive Bayes, Support Vector Machine (SVM), and Random Forest in classifying fake and true news. The models are evaluated using accuracy, precision, recall, F1-score, and confusion matrices. The impact of text preprocessing and handling imbalanced data is also discussed. Logistic Regression: Achieved 99% accuracy with high precision and recall (0.99 for true news, 0.98 for fake news) and balanced F1-scores of 0.99 for both classes. It misclassified 67 true news and 35 fake news articles but remains reliable for deployment. Naive Bayes: Delivered 93% accuracy with precision of 0.93 for both classes. Recall for fake news was 0.92. The model misclassified 285 true news articles as fake and 346 fake news articles as true, revealing limitations with text data. SVM: Achieved 100% accuracy with nearly perfect precision (1.00 for true news, 0.99 for fake news) and recall (1.00 for both). Only 24 true news and 20 fake news articles were misclassified, showing strong performance on high-dimensional data. Random Forest: Also achieved 100% accuracy with perfect precision, recall, and F1-scores. Only 12 true news and 11 fake news articles were misclassified, highlighting its ability to handle noise and imbalances. Comparison: SVM and Random Forest were the top performers with near-perfect results. Logistic Regression performed well, but Naive Bayes showed limitations due to its assumptions. Random Forest's ensemble method offered better accuracy, though it's more resource-intensive than SVM. The results support existing literature, with SVM and Random Forest proving to be highly effective for fake news detection. Logistic Regression performed well as a baseline model, while Naive Bayes struggled with the complexities of text classification.

A summary of the model results is provided in Table 1 below:

Table 1: Model Performance Summary

Model	Accuracy	Precision (Class 0)	Precision (Class 1)	Recall (Class 0)	Recall (Class 1)	F1-Score (Class 0)	F1-Score (Class 1)
Logistic Regression	0.99	0.99	0.98	0.99	0.99	0.99	0.99
Naive Bayes	0.93	0.93	0.93	0.94	0.92	0.93	0.93
SVM	1.00	1.00	0.99	0.99	1.00	1.00	0.99
Random Forest	1.00	1.00	1.00	1.00	1.00	1.00	1.00

Support Vector Machine (SVM) and Random Forest emerged as the top performers, both achieving perfect accuracy scores of 1.00. SVM, known for handling high-dimensional data and its resistance to overfitting, achieved a precision of 1.00 for true news and 0.99 for fake news, with near-perfect recall of 0.99 and 1.00, respectively. These results align with the literature, highlighting SVM's ability to find optimal hyperplanes for class separation.

Similarly, Random Forest achieved perfect scores across all metrics, demonstrating its robustness in managing noisy data and reducing overfitting. Its ability to aggregate predictions from multiple decision trees enabled it to capture complex patterns, making it highly effective for fake news detection.

Logistic Regression, with an accuracy of 0.99, also performed well. Its precision and recall of 0.99 for true news and 0.98 for fake news reflect its reliability. While simpler than ensemble methods, its linear decision boundary was sufficient for the task when combined with proper preprocessing techniques.

Naive Bayes showed lower performance, with an accuracy of 0.93. Its precision and recall scores of 0.93 and 0.92 reveal challenges in handling text data, mainly due to its assumption of feature independence, which does not account for word correlations in text. This limitation led to significant misclassifications, highlighting its struggles compared to more sophisticated models like SVM and Random Forest.

A confusion matrix was used to visualize the models' performance, breaking down true positives, true negatives, false positives, and false negatives. This helps in understanding the types of errors each model makes and guiding improvements.

#### Key Metrics and Analysis:

- True Positives (TP): 4266 news instances were correctly predicted as "True."
- True Negatives (TN): 4610 news instances were correctly predicted as "Fake."
- False Positives (FP): 67 news instances were incorrectly predicted as "True" (Type I error).

- False Negatives (FN): 35 news instances were incorrectly predicted as "Fake" (Type II error).




Figure 8: Confusion Matrix for Logistic Regression

The Logistic Regression model demonstrates strong performance, with high accuracy, precision, and recall.

Naïve Bayes: The provided confusion matrix evaluates the performance of a Naïve Bayes model in a fake news detection task. It offers a detailed breakdown of the model's predictions, allowing for a comprehensive assessment of its accuracy and effectiveness.

**Key Metrics and Analysis:**

- True Positives (TP): 3955 news instances were correctly predicted as "True."
- True Negatives (TN): 4392 news instances were correctly predicted as "Fake."
- False Positives (FP): 285 news instances were incorrectly predicted as "True" (Type I error).
- False Negatives (FN): 346 news instances were incorrectly predicted as "Fake" (Type II error).




Figure 9: Confusion Matrix for Naïve Bayes

SVM: The provided confusion matrix evaluates the performance of a SVM model in a fake news detection task. It offers a detailed breakdown of the model's predictions, allowing for a comprehensive assessment of its accuracy and effectiveness.

**Key Metrics and Analysis:**

- True Positives (TP): 4281 news instances were correctly predicted as "True."
- True Negatives (TN): 4653 news instances were correctly predicted as "Fake."
- False Positives (FP): 24 news instances were incorrectly predicted as "True" (Type I error).
- False Negatives (FN): 20 news instances were incorrectly predicted as "Fake" (Type II error).




Figure 10: Confusion Matrix for SVM

Random Forest: The provided confusion matrix evaluates the performance of a Random Forest model in a fake news detection task. It offers a detailed breakdown of the model's predictions, allowing for a comprehensive assessment of its accuracy and effectiveness.

**Key Metrics and Analysis:**

- True Positives (TP): 4290 news instances were correctly predicted as "True."
- True Negatives (TN): 4665 news instances were correctly predicted as "Fake."
- False Positives (FP): 12 news instances were incorrectly predicted as "True" (Type I error).
- False Negatives (FN): 11 news instances were incorrectly predicted as "Fake" (Type II error).




Figure 11: Confusion Matrix for Random Forest

Despite the high performance of the models, certain limitations persist. First, the dataset used is focused on specific themes, primarily politics and world news. This limits the generalizability of the models to other domains such as health, entertainment, or technology news. Moreover, while SVM and Random Forest models performed exceptionally well, their computational cost may hinder scalability when dealing with significantly larger datasets or real-time classification systems.

## 5. Conclusion

This research successfully developed a machine learning-based system for fake news detection, achieving high performance with models like Support Vector Machine (SVM) and Random Forest, both attaining perfect accuracy. The study highlighted the importance of data preprocessing techniques, such as exploratory data analysis, text normalization and TF-IDF vectorization, in enhancing model performance for fake news detections.

## References

- Ahmed, W., Vidal-Alaball, J., Downing, J., & López Seguí, F. (2020). COVID-19 and the 5G conspiracy theory: social network analysis of Twitter data. *Journal of Medical Internet Research*, 22(5), e19458. <https://doi.org/10.2196/19458>
- Allcott, H. & Gentzkow, M., 2017. Social Media and Fake News in the 2016 Election. *Journal of Economic Perspectives*, 31(2), 211-236.
- Ahmad, I., Yousaf, M., Yousaf, S., & Ahmad, M.O. (2020). Fake News Detection Using Machine Learning Ensemble Methods. *Hindawi*, 1-11. <https://doi.org/10.1155/2020/8885861>
- Al-Ahmad, B., Al-Zoubi, A.M., Abu Khurma, R., & Aljarah, I. (2021). An evolutionary fake news detection method for covid-19 pandemic information. *Symmetry*, 13 (6):1-16.
- Almeida, L.D., Fuzaro, V., Nieto, F., & Santana, A.L.M. (2021). Identificação de “Fake News” no contexto político brasileiro: uma abordagem computacional. In: *Proceedings Workshop sobre as Implicações da Computação na Sociedade (WICS)*, Porto Alegre, Brasil, pp. 78-89. <https://doi.org/10.5753/wics.2021.15966>
- Aslam, N., Ullah Khan, I., Alotaibi, F.S., Aldaej, L.A., & Aldubaikil, A.K. (2021). Fake detect: A deep learning ensemble model for fake news detection. *Hindawi*, 1-8. <https://doi.org/10.1155/2021/5557784>
- Ayoub, J., Yang, X.J., & Zhou, F. (2021). Combat COVID-19 infodemic using explainable natural language processing models. *Information Processing & Management*, 58, 4:1-11. <https://doi.org/10.1016/j.ipm.2021.102569>
- Bahad, P., Saxena, P., & Kamal, R. (2020). Fake News Detection using Bi-directional LSTM-Recurrent Neural Network. *Procedia Computer Science*, 165:74-82.
- Brennen, J. S., Simon, F. M., Howard, P. N., & Nielsen, R. K. (2020). Types, sources, and claims of COVID-19 misinformation. *Reuters Institute*, 7, 3-1. <https://doi.org/10.31235/osf.io/0s98z>
- Chawla, N. V., Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. (2002). SMOTE: synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 16, 321-357. <https://doi.org/10.1613/jair.953>
- Chesney, R., & Citron, D. (2019). Deepfakes and the new disinformation war: The coming age of post-truth geopolitics. *Foreign Affairs*, 98(1), 147-155. Retrieved from <https://www.foreignaffairs.com/articles/world/2018-12-11/deepfakes-and-new-disinformation-war>
- Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)* (pp. 4171-4186). <https://doi.org/10.18653/v1/N19-1423>
- Faustini, P.H.A., & Covões, T.F. (2020). Fake news detection in multiple platforms and languages. *Expert Systems with Applications*, 158:1-9. <https://doi.org/10.1016/j.eswa.2020.113503>
- Gereme, F., Zhu, W., Ayall, T., & Alemu, D. (2021). Combating fake news in “low-resource” languages: Amharic fake news detection accompanied by resource crafting. *Information*, 12, 1:1-9. <https://doi.org/10.3390/info12010020>
- Goel, P., Singhal, S., Aggarwal, S., & Jain, M. (2021). Multi Domain Fake News Analysis using Transfer Learning. *Proceedings 5th International Conference on Computing Methodologies and Communication (ICCMC)*, Erode, India, pp. 1230–1237. <https://doi.org/10.1109/ICCMC51019.2021.9418411>
- Goldani, M.H., Momtazi, S., & Safabakhsh, R. (2021). Detecting fake news with capsule neural networks. *Applied Soft Computing*, 101:1-8. <https://doi.org/10.1016/j.asoc.2020.106991>

- Graves, L. (2020). Understanding the promise and limits of automated fact-checking. *Reuters Institute for the Study of Journalism*. Retrieved from <https://reutersinstitute.politics.ox.ac.uk/our-research/understanding-promise-and-limits-automated-fact-checking>
- Gupta, A., Lamba, H., & Kumaraguru, P., (2021). The Economic Impact of Fake News on Financial Markets. *Journal of Financial Economics*, 10(1), pp. 112-133.
- Jones, M. & Thompson, R., (2023). The Spread of Misinformation During the COVID-19 Pandemic. *Journal of Digital Media Studies*, 12(3). 45-67.
- Julio, C. Reis, A. Correia, F. Murai, A. and Benevenuto, F.; Supervised Learning for Fake News Detection," in IEEE Intelligent Systems, vol. 34, no. 2, pp. 76-81, March-April 2019,doi: 10.1109/MIS.2019.2899143 (2019)
- Li, Y., Mertens, S., Stojanovski, K., & Stalla-Bourdillon, S. (2022). Detection and mitigation of online misinformation: A survey. *ACM Computing Surveys (CSUR)*, 55(4), 1-36. <https://doi.org/10.1145/3524458>
- Loomba, S., de Figueiredo, A., Piatek, S. J., de Graaf, K., & Larson, H. J. (2021). Measuring the impact of COVID-19 vaccine misinformation on vaccination intent in the UK and USA. *Nature Human Behaviour*, 5(3), 337-348. <https://doi.org/10.1038/s41562-021-01056-1>
- Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)*, 54(6), 1-35. <https://doi.org/10.1145/3457607>
- Mohammed, H and Sayeed, A. ;Use of Fake News and Social Media by Mainstream News Channel in India, 16th IEEE International Colloquium on Signal Processing and its Application (2020)
- Pasel, A., Aker, A. Bontcheva, k. Liakata, M. and Procter, R., "Detection and Resolution of Rumours in Social Media": A Survey Article 32.<https://doi.org/10.1145/3161603> (2019)
- Pennycook, G., & Rand, D. G. (2021). The Implied Truth Effect: Attaching Warnings to a Subset of Fake News Stories Increases Perceived Accuracy of Stories Without Warnings. *Management Science*, 66(11), 4944-4957. <https://doi.org/10.1287/mnsc.2019.3478>
- Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?" Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining* (pp. 1135-1144). <https://doi.org/10.1145/2939672.2939778>
- Rubin, V. L., Conroy, N. K., Chen, Y., & Cornwell, S. (2016). Fake news or truth? Using satirical cues to detect potentially misleading news. In *Proceedings of the Second Workshop on Computational Approaches to Deception Detection* (pp. 7-17). <https://doi.org/10.18653/v1/W16-0802>
- Ruchansky, N., Seo, S., & Liu, Y. (2017). CSI: A hybrid deep model for fake news detection. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management* (pp. 797-806). <https://doi.org/10.1145/3132847.3132877>
- Shu, K., Sliva, A., Wang, S., Tang, J., & Liu, H. (2017). Fake news detection on social media: A data mining perspective. *ACM SIGKDD Explorations Newsletter*, 19(1), 22-36. <https://doi.org/10.1145/3137597.3137600>
- Shu, K., Wang, S., Lee, D., & Liu, H. (2020). Mining disinformation and fake news: Concepts, methods, and recent advancements. In *Disinformation, Misinformation, and Fake News in Social Media* (pp. 1-19). Springer, Cham. [https://doi.org/10.1007/978-3-030-42699-6\\_1](https://doi.org/10.1007/978-3-030-42699-6_1)
- Smith, J., Brown, L., & Williams, P., (2022). The Impact of Fake News on Public Opinion: A Survey of U.S. Adults. *Pew Research Center Report*. Available at: <https://www.pewresearch.org/fake-news-impact-2022>.
- Tandoc, E. C., Lim, Z. W., & Ling, R. (2018). Defining "fake news" A typology of scholarly definitions. *Digital journalism*, 6(2), 137-153. <https://doi.org/10.1080/21670811.2017.1360143>
- Uma S., Saran, S., & Patil, S. M. (2021). *Fake News Detection using Machine Learning Algorithms*. *Research & Technology (International Journal of Engineering IJERT)*, 9(3), NTASU-2020 Conference

- Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science*, 359(6380), 1146-1151. <https://doi.org/10.1126/science.aap9559>
- Wang, W. Y. (2017). "Liar, Liar Pants on Fire": A New Benchmark Dataset for Fake News Detection. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)* (pp. 422-426). <https://doi.org/10.18653/v1/P17-2067>
- Zellers, R., Holtzman, A., Bisk, Y., Farhadi, A., & Choi, Y. (2019). Defending against neural fake news. In *Advances in Neural Information Processing Systems* (pp. 9051-9062). Retrieved from <https://proceedings.neurips.cc/paper/2019/file/3e9f0fc9b2f89e043bc6233994dfcf76-Paper.pdf>
- Zhang, Y., Pennycook, G., & Rand, D. G. (2020). Fighting fake news on social media using crowdsourced judgments of news source quality. *Proceedings of the National Academy of Sciences*, 117(5), 2326-2332. <https://doi.org/10.1073/pnas.1912444117>