International Journal of Research Publication and Reviews

Journal homepage: www.ijrpr.com ISSN 2582-7421

AI-Driven Fake News Detection: A Machine Learning-Based Approach

[1]Aniket. Dilip. Domble, [2]Sahil. Manohar. Dhamanda, [3]Bhushan. Dilip. More, [4]Anish. Arvind. Jivtode, [5]Amay. Santosh. Handa, [6]Prof. Vanisha. P. Vaidya

[1,2,3,4,5] Student, Sipna College of Engineering & Technology, Amaravati, Maharashtra, India
[6]Professor, Sipna College of Engineering & Technology, Amaravati, Maharashtra, India

ABSTRACT

The proliferation of fake news across digital platforms has emerged as a critical challenge, undermining the credibility of information, distorting public perception, and threatening societal stability[1]. The rapid dissemination of misinformation, particularly through social media, has necessitated the development of automated tools capable of identifying and mitigating the spread of false content in real time[2]. In this study, we explore the use of machine learning techniques for effective fake news detection, with a focus on both performance and practical deployability.Our research presents a comprehensive evaluation of three prominent supervised machine learning algorithms—Logistic Regression, Gaussian Naive Bayes, and Passive Aggressive Classifier—for binary classification of news articles as real or fake[3]. We employ a robust feature extraction methodology using TF-IDF vectorization, which transforms textual data into informative numerical representations, enabling efficient and accurate classification. The dataset used comprises labeled news articles from publicly available sources, including both real and fake samples, which underwent extensive preprocessing steps such as tokenization, stopword removal, and lemmatization to ensure data quality and consistency.The experimental results indicate that Logistic Regression outperforms the other models, achieving a notable accuracy of 94.95%, demonstrating its ability to model linear decision boundaries effectively in high-dimensional spaces. Passive Aggressive Classifier closely follows, with an accuracy of 93.69%, showcasing its potential for online learning and real-time adaptability. Meanwhile, Gaussian Naive Bayes, a probabilistic model grounded in Bayes' Theorem, achieves 88.30% accuracy, providing a lightweight yet moderately effective solution with faster inference times.The comparative analysis highlights the trade-offs between accuracy, computational efficiency, and interpretability, suggesting that traditional machine learning models, particularly Logistic Regression, remain highly relevant in the context of fake news detection. The lightweight and interpretable nature of these models makes them suitable for real-time applications where rapid decision-making is essential. Moreover, the study emphasizes the importance of feature engineering, demonstrating how optimized vectorization techniques significantly impact classification performance.In conclusion, this work contributes to the field by offering a scalable, interpretable, and efficient machine learning framework for detecting fake news with high precision. Future research directions include the incorporation of deep learning models and transformer-based architectures to capture contextual nuances, as well as the integration of multimodal data (e.g., images, metadata) for a more holistic approach to misinformation detection

Keywords: Fake News Detection, Machine Learning, Logistic Regression, Gaussian Naive Bayes, Passive Aggressive Classifier, TF-IDF, Text Classification, Natural Language Processing (NLP), Supervised Learning, Misinformation Detection.

1. Introduction:

In the digital era, the accessibility and volume of information have increased exponentially due to the widespread use of the internet and social media platforms[4]. While this has facilitated real-time information sharing and global communication, it has also opened the door to the rapid dissemination of **fake news**—intentionally fabricated information aimed at misleading readers for political, financial, or social gain[5]. The proliferation of fake news has become a **serious global concern**, affecting public opinion, influencing elections, inciting social unrest, and eroding trust in legitimate news sources.

Traditional methods of **fact-checking and manual verification**, though accurate, are inherently **time-consuming, labor-intensive**, and incapable of keeping pace with the scale and velocity at which misinformation spreads online[6]. The sheer volume of data generated daily necessitates the development of **automated and scalable solutions** that can detect and filter out fake news quickly and accurately. In this context, **machine learning (ML)** and **natural language processing (NLP)** have emerged as promising tools to address this pressing issue. These technologies can analyze large datasets, identify patterns, and make predictions, thereby offering an efficient approach to tackle the fake news problem.

This research aims to explore and evaluate the application of **supervised machine learning algorithms** for **binary text classification**, i.e., categorizing news articles as either real or fake based on their **textual content**. The study focuses on three widely used ML algorithms—**Logistic Regression**, **Gaussian Naive Bayes**, and **Passive Aggressive Classifier**—chosen for their interpretability, computational efficiency, and proven performance in various text classification tasks. These algorithms are trained and tested using benchmark datasets composed of labeled news articles, and the **features are extracted using Term Frequency-Inverse Document Frequency (TF-IDF)**, which converts textual data into a form suitable for machine learning models.

The primary objective is to determine the **effectiveness and accuracy** of these models in classifying fake news while also evaluating their suitability for **real-time deployment scenarios**, where speed and accuracy are equally critical. Furthermore, the study aims to highlight the **trade-offs between accuracy and computational overhead**, providing insights into which models are most effective in balancing performance and efficiency. The results of this research not only contribute to the growing body of knowledge in **AI-driven fake news detection** but also provide a foundation for developing **scalable, interpretable, and reliable systems** that can assist media platforms, governments, and fact-checkers in combating the spread of misinformation.

In summary, this paper presents a comparative analysis of selected machine learning models for fake news detection, emphasizing their practical applicability and performance. The study underscores the potential of **lightweight yet powerful algorithms** in addressing real-world challenges posed by fake news and sets the stage for future work that may incorporate more complex models, such as deep learning and transformer-based architectures, to further enhance detection capabilities.

1.2 Research Objectives

The overarching aim of this study is to develop an effective and accurate machine learning-based solution for the automatic detection of fake news. Specifically, the research is guided by three primary objectives. First, the study seeks to design and implement a robust machine learning model that can accurately classify news articles as real or fake, thereby contributing to the growing demand for scalable solutions to combat misinformation. Second, it aims to conduct a comparative analysis of several traditional yet widely-used supervised machine learning algorithms, including **Logistic Regression**, **Gaussian Naive Bayes**, and the **Passive Aggressive Classifier**, to evaluate their relative strengths, weaknesses, and practical applicability in real-time fake news detection scenarios. These models are selected for their computational efficiency and interpretability, which are crucial for deployment in environments where speed and transparency are key considerations. Third, the research emphasizes the importance of **feature engineering and data representation**, particularly through the implementation of **TF-IDF (Term Frequency-Inverse Document Frequency) vectorization**, to optimize input features and improve the overall accuracy and generalization capabilities of the models. By addressing these objectives, the study aims to contribute meaningful insights into the development of lightweight yet high-performing fake news detection systems suitable for deployment across a range of platforms and applications.

1.3 Contributions of the Paper

This research paper offers several significant contributions to the growing field of fake news detection, particularly through the application and evaluation of machine learning techniques. First and foremost, the study provides a comprehensive **comparative analysis of three diverse machine learning models**—**Logistic Regression**, **Gaussian Naive Bayes**, and **Passive Aggressive Classifier**—each representing different methodological approaches, from probabilistic to online learning. By rigorously evaluating these models on a benchmark dataset, the paper offers valuable insights into their **respective strengths, limitations, and practical viability** for real-world implementation, especially in dynamic online environments where misinformation rapidly evolves. Additionally, the paper introduces an **optimized feature extraction pipeline using Term Frequency-Inverse Document Frequency (TF-IDF)**, a powerful technique that effectively transforms raw textual content into structured numerical representations suitable for machine learning algorithms. This process not only enhances the model's ability to capture key linguistic patterns but also improves its generalization to unseen data. A notable contribution of this research is the **demonstration of the high efficacy of Logistic Regression**, which, despite its simplicity and computational efficiency, achieves an **impressive classification accuracy of 94.95%**, outperforming more complex models in this context. This finding underscores the potential of lightweight, interpretable models in the accurate and timely detection of fake news, thereby supporting scalable and real-time deployment across various digital platforms. Collectively, these contributions advance the understanding of effective algorithmic approaches to combating misinformation and provide a foundation for future enhancements in automated fake news detection systems.

2. Related Work

Extensive research has been undertaken in the domain of fake news detection, employing a wide spectrum of machine learning methodologies ranging from traditional algorithms to sophisticated deep learning models[7]. Traditional machine learning techniques such as **Support Vector Machines (SVM)**, **Logistic Regression**, and **Random Forests** have consistently shown promising results, particularly when integrated with robust feature extraction techniques like **Term Frequency-Inverse Document Frequency (TF-IDF)** and **Word2Vec embeddings**, which convert unstructured textual data into meaningful numerical representations[8]. These models are often favored for their **computational efficiency, interpretability, and ease of implementation**, making them viable for scalable applications. **Probabilistic classifiers such as Naive Bayes** further contribute to this landscape by offering **fast training times and straightforward probability estimations**, although they may sometimes underperform in comparison to more complex models in terms of accuracy, particularly when handling nuanced linguistic patterns[9]. In contrast, recent advancements in **deep learning approaches**, notably **Long Short-Term Memory networks (LSTM)**, **Convolutional Neural Networks (CNNs)**, and **Transformer-based architectures like BERT and RoBERTa**, have revolutionized the field by enabling models to learn **contextual relationships and syntactic structures** within text at a much deeper level[10]. These models excel at capturing **semantic nuances** and **long-range dependencies**, which are often critical in differentiating between genuine and fabricated news content. However, despite the substantial improvements in performance, significant challenges persist, such as the **inability of many models to adapt effectively to evolving misinformation strategies**, **limited access to high-quality, diverse labeled datasets**, and the **difficulty in deploying computationally intensive models in real-time environments[11]**. These limitations underscore the need for continued innovation in developing **accurate, efficient, and adaptive fake news detection systems**.

3. Methodology

3.1 Approach

This research adopts a structured hybrid machine learning framework designed to maximize the efficiency of feature extraction and enable a comprehensive comparative evaluation of three distinct classification algorithms. The approach begins with the collection and preprocessing of news article data, where standard Natural Language Processing (NLP) techniques such as tokenization, stopword removal, and lemmatization are applied to cleanse and normalize the text[12]. The processed textual data is then transformed into numerical representations through **Term Frequency-Inverse Document Frequency (TF-IDF) vectorization**, a powerful technique that quantifies the importance of words relative to the document corpus, thereby enabling the model to focus on discriminative features during training. These TF-IDF vectors form the foundational input for the three selected machine learning classifiers, each representing a unique learning paradigm and contributing varied strengths to the classification task. **Logistic Regression** is employed as a baseline linear classifier, known for its **simplicity, speed, and interpretability**, making it well-suited for high-dimensional text data. **Gaussian Naive Bayes** is included for its **probabilistic approach to classification**, leveraging the assumption of feature independence and offering rapid training and prediction, which is beneficial in real-time scenarios despite its potential limitations in capturing feature dependencies. Lastly, the **Passive Aggressive Classifier**, a powerful **online learning algorithm**, is used to adaptively update its model parameters with each new training instance, making it particularly effective in dynamic environments where fake news patterns evolve continuously. This diversified selection of algorithms allows for a thorough performance comparison and highlights the trade-offs between **accuracy, efficiency, and adaptability** in the context of fake news detection.

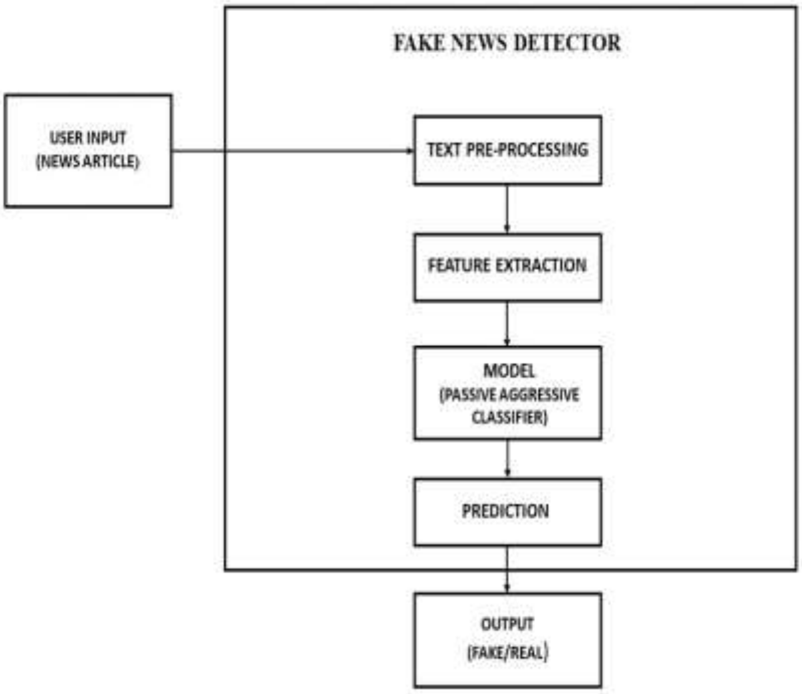

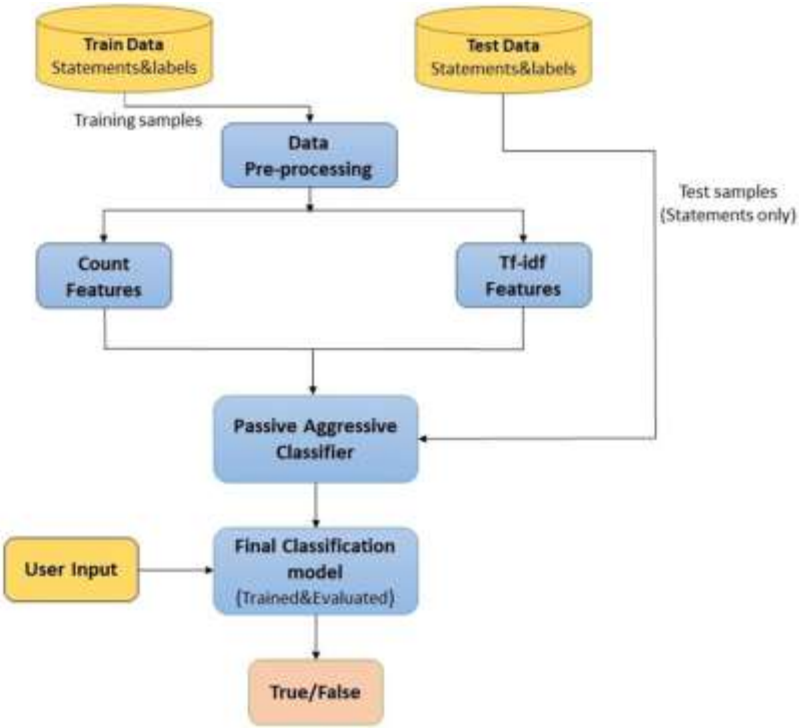

3.2 Algorithms and Models

- **Logistic Regression:** Logistic Regression is a linear classification algorithm that predicts the probability of a data point belonging to a particular class using the logistic (sigmoid) function. It is particularly well-suited for binary classification problems, such as fake versus real news detection, due to its simplicity, computational efficiency, and interpretability. The model operates by finding a linear relationship between the input features and the log-odds of the outcome, allowing it to make decisions based on a weighted combination of input variables. In the context of text classification, Logistic Regression works effectively when paired with robust feature extraction methods like TF-IDF, which convert textual data into numerical vectors capturing term significance. In this study, Logistic Regression emerged as the most accurate classifier, achieving a **highest accuracy of 94.95%**, demonstrating its capability to generalize well over the dataset while maintaining low computational overhead. Its balance of speed and performance makes it highly suitable for deployment in real-time fake news detection systems where both precision and responsiveness are critical.
- **Gaussian Naive Bayes:** Gaussian Naive Bayes is a probabilistic machine learning algorithm based on **Bayes' Theorem**, which predicts class membership probabilities under the assumption that features are statistically independent and normally distributed. This simplification allows for fast training and inference, even on large datasets or in resource-constrained environments. While the independence assumption rarely holds for real-world textual data, Gaussian Naive Bayes often performs surprisingly well, especially when the data is clean and well-structured. The model calculates the likelihood of the data belonging to each class and selects the class with the highest posterior probability. In this research, Gaussian Naive Bayes demonstrated **an accuracy of 88.30%**, highlighting its efficacy as a lightweight solution for fake news detection. Though its accuracy is lower than other models evaluated, its ease of implementation, low memory usage, and fast performance make it valuable for applications requiring quick deployment or operation with limited computational resources.
- **Passive Aggressive Classifier:** The Passive Aggressive Classifier is an **online learning algorithm** specifically designed for large-scale and real-time machine learning tasks. Unlike batch learning algorithms, it updates its model parameters only when it encounters a misclassification (aggressive) and remains passive when the classification is correct. This characteristic makes it particularly efficient for streaming data or scenarios where fake news content evolves continuously, and frequent model retraining is impractical. It operates by minimizing a hinge loss function with an added regularization constraint, allowing it to balance between learning from mistakes and maintaining previous knowledge. In this study, the Passive Aggressive Classifier achieved an **impressive accuracy of 93.69%**, demonstrating its strong performance and adaptability in dynamic data environments. Its ability to handle high-dimensional feature spaces, coupled with rapid updates and minimal computational cost, positions it as a powerful tool for real-time fake news detection where adaptability and speed are essential.

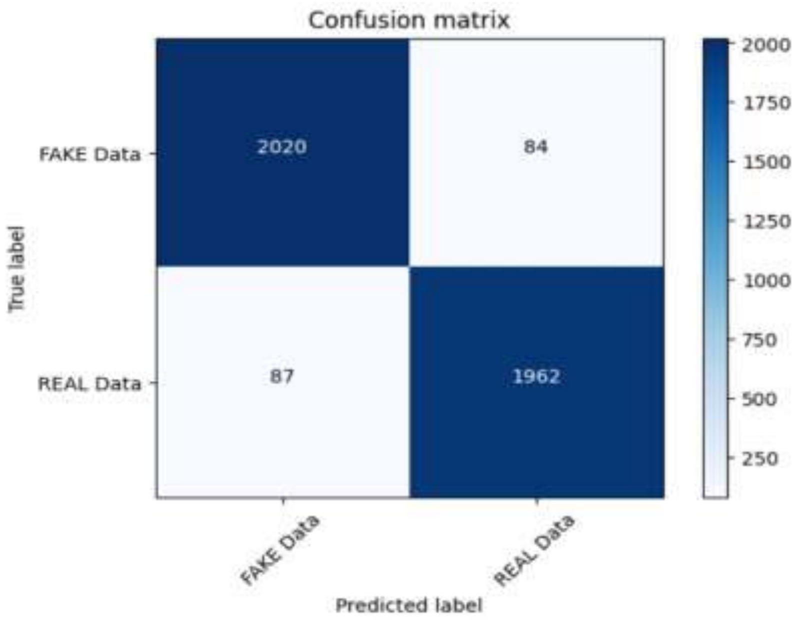

3.3 Data Sources and Preprocessing

In this study, the dataset utilized comprises a balanced collection of labeled news articles categorized as either fake or real, enabling supervised learning for binary classification[13]. The dataset was carefully curated from reputable public sources that specialize in fake news research, ensuring a diverse representation of topics and writing styles[14]. Each article includes headline and body text, providing sufficient context for analysis. The inclusion of both short-form (e.g., headlines) and long-form (e.g., full articles) content allows the models to learn from various linguistic patterns and levels of detail, improving generalizability.

To prepare the raw text data for machine learning models, a comprehensive preprocessing pipeline was employed. Initially, the text was cleaned to remove extraneous characters such as punctuation, numbers, and special symbols that do not contribute meaningful information to the classification task. Tokenization was then applied, breaking the continuous text into individual words or tokens, which serve as the basic units for analysis. Following tokenization, stopword removal was conducted to eliminate common words such as "the," "is," and "and," which appear frequently across documents but do not offer significant discriminatory value. This step reduces noise in the data and enhances the quality of the input features.

Subsequently, lemmatization was performed to normalize words to their base or dictionary form. For instance, words like "running," "ran," and "runs" were converted to "run," allowing the model to treat them as a single feature and thus reduce dimensionality. This process ensures that the core meaning of words is preserved while minimizing redundancy in the feature space.

Once the text was preprocessed, TF-IDF (Term Frequency-Inverse Document Frequency) vectorization was used to transform the cleaned textual data into numerical feature vectors suitable for input into machine learning algorithms. TF-IDF assigns weights to words based on their frequency in individual documents relative to their frequency across the entire dataset, emphasizing terms that are both frequent and distinctive. This technique not only preserves essential information about word importance but also helps in mitigating the impact of common, less informative terms. The resulting vectors effectively capture the relevance of words within the context of fake and real news, enabling the models to learn meaningful patterns for classification.

This well-structured preprocessing strategy ensures consistency, reduces overfitting risks, and enhances model performance across different algorithms by feeding them high-quality, informative feature representations. The combination of textual cleaning, normalization, and strategic vectorization plays a crucial role in the success of the machine learning models employed in this research.

3.4 Experimental Setup

The experimental phase of this research was carried out using the **Python programming language**, which offers extensive libraries and frameworks tailored for machine learning and data science applications. Among these, **Scikit-learn** served as the primary toolkit due to its robust support for various machine learning algorithms, preprocessing utilities, model evaluation tools, and seamless integration with other Python libraries such as Pandas and NumPy. This facilitated efficient data manipulation, feature extraction, model training, and performance evaluation.

The experimental workflow began with loading and preprocessing the dataset, including tokenization, stopword removal, lemmatization, and **TF-IDF vectorization** to convert raw textual data into meaningful numerical representations. Once preprocessed, the dataset was divided into **training and testing subsets** to evaluate model generalization. Cross-validation techniques were also employed where necessary to ensure reliable and unbiased performance assessments across different data splits.

Three different **supervised machine learning models**—**Logistic Regression**, **Gaussian Naive Bayes**, and **Passive Aggressive Classifier**—were implemented and trained using the prepared dataset. Each model was selected to represent a distinct learning paradigm: **Logistic Regression** for its linear decision boundaries and high interpretability, **Gaussian Naive Bayes** for its probabilistic foundations and simplicity, and **Passive Aggressive Classifier** for its capability in **online learning scenarios**, where model updates occur incrementally and efficiently.

Model performance was rigorously evaluated using a comprehensive set of **standard classification metrics**, which include:

- **Accuracy**: The ratio of correctly predicted instances to the total number of instances, serving as a basic performance indicator.
- **Precision**: The ratio of true positive predictions to the total predicted positives, reflecting the model's reliability in positive predictions.
- **Recall**: The ratio of true positive predictions to the total actual positives, measuring the model's ability to detect all relevant instances.
- **F1-score**: The harmonic mean of Precision and Recall, offering a balanced metric that accounts for both false positives and false negatives.
- **AUC-ROC (Area Under the Receiver Operating Characteristic Curve)**: This metric evaluates the model's ability to distinguish between classes across different threshold settings, providing insight into overall classification performance beyond a fixed decision boundary.

Among the models evaluated, **Logistic Regression achieved the highest accuracy of 94.95%**, demonstrating its strength in effectively separating fake and real news articles with a relatively simple linear classifier. The **Passive Aggressive Classifier** followed closely, achieving an accuracy of **93.69%**, and stood out due to its **low computational overhead** and suitability for real-time applications where data may arrive continuously. Meanwhile, the **Gaussian Naive Bayes model** reached an accuracy of **88.30%**, offering a **fast and interpretable probabilistic approach**, albeit with slightly reduced performance compared to the other models.

Overall, the experimental setup underscored the importance of proper **feature engineering, algorithm selection**, and **evaluation criteria** in developing efficient and accurate fake news detection systems.

4. Results and Discussion

The experimental evaluation of the machine learning models—**Logistic Regression**, **Gaussian Naive Bayes**, and **Passive Aggressive Classifier**—was conducted to assess their effectiveness in detecting fake news based on textual features extracted using **TF-IDF vectorization[15]**. Each model's performance was measured using a range of evaluation metrics, with **accuracy** serving as the primary indicator due to its relevance in binary classification tasks.

Model	**Accuracy**
Logistic Regression	**94.95%**
Gaussian Naive Bayes	88.30%
Passive Aggressive Classifier	93.69%

The results clearly indicate that **Logistic Regression outperformed the other two models**, achieving an accuracy of **94.95%**, making it the most effective and reliable classifier in this study. The strength of Logistic Regression lies in its ability to create **linear decision boundaries** that accurately separate fake and real news classes when trained on **TF-IDF-transformed data**. Its efficiency, interpretability, and generalization capabilities make it an ideal choice for **production-level fake news detection systems**, especially in scenarios where both **accuracy and computational efficiency** are paramount.

The **Passive Aggressive Classifier** ranked second with an accuracy of **93.69%**, offering a compelling trade-off between performance and **computational efficiency**. This model is specifically designed for **online learning**, updating itself only when a classification error occurs. This characteristic makes it well-suited for **real-time systems** that handle **streaming data** or large-scale content. While slightly less accurate than Logistic Regression, the model's ability to rapidly adapt to new data with minimal overhead provides a strategic advantage in dynamic environments, such as **social media platforms** or **news aggregators**, where new information continuously emerges.

The **Gaussian Naive Bayes** model, achieving an accuracy of **88.30%**, delivered the lowest performance among the three. However, its simplicity and **probabilistic nature** make it useful in scenarios where **interpretability** and **speed** are more critical than absolute accuracy. Its **assumption of feature independence**, while often violated in natural language processing tasks, allows for **fast training and inference**, making it suitable for **resource-constrained devices** or **preliminary filtering systems** where rapid decisions are needed before engaging more computationally intensive models.

In summary, the **comparative analysis** of these models underscores the **importance of selecting an algorithm** based on **application-specific requirements**:

- For scenarios where **high accuracy** is critical and computational resources are available, **Logistic Regression** is the preferred choice.
- For **real-time, large-scale, or streaming environments**, the **Passive Aggressive Classifier** provides a robust and adaptive solution.

- In contexts where **speed and interpretability** are prioritized over accuracy, such as **low-power environments** or **explanatory systems**, **Gaussian Naive Bayes** remains a viable option.

The study illustrates that while all three models have distinct advantages, the **combination of optimized feature extraction (TF-IDF)** and a well-chosen algorithm can significantly enhance the performance of fake news detection systems. Furthermore, these findings suggest that **lightweight machine learning models** can effectively address the challenge of misinformation without the need for **resource-intensive deep learning architectures**, making them suitable for **real-world deployment** across a range of platforms.

5. Conclusion

This research underscores the efficacy of traditional machine learning algorithms in the domain of fake news detection, with a particular focus on the comparative performance of Logistic Regression, Gaussian Naive Bayes, and Passive Aggressive Classifier[16]. Through rigorous experimentation and performance evaluation, it is evident that Logistic Regression, when combined with optimized feature extraction techniques like TF-IDF vectorization, delivers remarkably high accuracy (94.95%) in identifying fake news. Its simplicity, efficiency, and interpretability make it not only effective for detection tasks but also suitable for real-world deployment across various platforms, including web applications and browser extensions.

Moreover, the study highlights the unique strengths of online learning models, particularly the Passive Aggressive Classifier, which achieved 93.69% accuracy while maintaining real-time adaptability. This model's capability to update incrementally with new data instances makes it highly valuable in dynamic environments where information is continuously generated, such as social media feeds or news aggregators. Despite a slightly lower accuracy than Logistic Regression, its scalability and low latency are major assets in high-throughput systems[17].

The Gaussian Naive Bayes model, while yielding a moderate accuracy of 88.30%, provided valuable insights into probabilistic classification. Its fast processing speed and low resource consumption make it appropriate for resource-constrained environments or as a baseline model in multi-stage classification pipelines. The results collectively affirm that with appropriate feature engineering and model selection, even lightweight traditional algorithms can offer competitive performance in detecting misinformation[18].

Reference

1. **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.
2. **Shu, K., Wang, S., & Liu, H. (2019).** *Beyond News Contents: The Role of Social Context for Fake News Detection*. Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining, 312-320.
3. **Shu, K., Cui, L., Wang, S., Lee, D., & Liu, H. (2019).** *dEFEND: Explainable Fake News Detection*. Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, 395-405.
4. **Shu, K., Mahudeswaran, D., Wang, S., Lee, D., & Liu, H. (2020).** *FakeNewsNet: A Data Repository with News Content, Social Context, and Spatiotemporal Information for Studying Fake News on Social Media*. Big Data, 8(3), 171-188.
5. **Wang, Y., Yang, W., Ma, F., Xu, J., Zhong, B., Deng, Q., & Gao, J. (2019).** *Weak Supervision for Fake News Detection via Reinforcement Learning*. arXiv preprint arXiv:1912.12520.
6. **Yuan, C., Ma, Q., Zhou, W., Han, J., & Hu, S. (2020).** *Early Detection of Fake News by Utilizing the Credibility of News, Publishers, and Users Based on Weakly Supervised Learning*. arXiv preprint arXiv:2012.04233.
7. **Roy, A., Basak, K., Ekbal, A., & Bhattacharyya, P. (2018).** *A Deep Ensemble Framework for Fake News Detection and Classification*. arXiv preprint arXiv:1811.04670.
8. **Saikh, T., De, A., Ekbal, A., & Bhattacharyya, P. (2020).** *A Deep Learning Approach for Automatic Detection of Fake News*. arXiv preprint arXiv:2005.04938.
9. **Zhou, X., & Zafarani, R. (2018).** *Fake News: A Survey of Research, Detection Methods, and Opportunities*. arXiv preprint arXiv:1812.00315.
10. **Ahmed, H., Traore, I., & Saad, S. (2017).** *Detecting Opinion Spams and Fake News Using Text Classification*. Security and Privacy, 1(1), e9.
11. **Ruchansky, N., Seo, S., & Liu, Y. (2017).** *CSI: A Hybrid Deep Model for Fake News Detection*. Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, 797-806.
12. **Monti, F., Frasca, F., Eynard, D., Mannion, D., & Bronstein, M. M. (2019).** *Fake News Detection on Social Media Using Geometric Deep Learning*. arXiv preprint arXiv:1902.06673.
13. **Zhang, X., & Ghorbani, A. A. (2020).** *An Overview of Online Fake News: Characterization, Detection, and Discussion*. Information Processing & Management, 57(2), 102025.

14. **Kaliyar, R. K., Goswami, A., & Narang, P. (2021).** *FakeBERT: Fake News Detection in Social Media with a BERT-Based Deep Learning Approach*. Multimedia Tools and Applications, 80(8), 11765-11788.
15. **Zhou, X., & Zafarani, R. (2020).** *Network-Based Fake News Detection: A Pattern-Driven Approach*. ACM SIGKDD Explorations Newsletter, 22(1), 48-55.
16. **Kumar, S., & Shah, N. (2018).** *False Information on Web and Social Media: A Survey*. arXiv preprint arXiv:1804.08559.
17. **Ghanem, B., Rosso, P., & Rangel, F. (2020).** *An Emotional Analysis of False Information in Social Media and News Articles*. ACM Transactions on Internet Technology (TOIT), 20(2), 1-18.
18. **Zhou, X., & Zafarani, R. (2019).** *Fake News Detection: An Interdisciplinary Research*. Proceedings of the 13th International Conference on Web and Social Media, 436-439.