Jurnal Teknik Informatika C.I.T Medicom, 16 (1) (2024) 13-23

Published by: Institute of Computer Science (IOCS)

Jurnal Teknik Informatika C.I.T Medicom

Journal homepage: www.medikom.iocspublisher.org

Advancing fake news detection: a comparative study of RNN, LSTM, and Bidirectional LSTM Architectures

Gregorius Airlangga

Information System Department, Atma Jaya Catholic University of Indonesia, Jakarta, Indonesia

Article Info

Article history

Received : Jan 15, 2024
Revised : Jan 21, 2024
Accepted : Mar 26, 2024

Keywords:

Comparative analysis;
Data preprocessing;
Fake news detection;
LSTM model;
Neural network architectures.

Abstract

In the era of information overload, the exponential growth of digital content has coincided with the proliferation of 'fake news,' posing a critical challenge to online information credibility. This study addresses the pressing need for robust fake news detection systems by conducting a comparative analysis of three neural network architectures: Recurrent Neural Networks (RNN), Long Short-Term Memory (LSTM), and Bidirectional LSTM (BiLSTM). Our primary objective is to assess their effectiveness in identifying fake news in a binary classification setting. To achieve this goal, we employed advanced neural network models and a dataset of news titles. Our applied research method included data preprocessing and the utilization of RNN, LSTM, and BiLSTM models, each tailored to handle sequential data and capture temporal dependencies. we rigorously assessed the performance of RNN, LSTM, and BiLSTM models using a range of metrics, including accuracy, precision, recall, and F1-score. To achieve a comprehensive evaluation, we divided our dataset into training and testing subsets. Specifically, we allocated 67% of the data for training purposes and the remaining 33% for testing. Our research findings reveal that all three models consistently achieved high accuracy levels, approximately 91%, with slight variations in precision and recall. Notably, the LSTM model exhibited a marginal improvement in recall, which is crucial when the consequences of missing deceptive content outweigh false alarms. Conversely, the RNN model demonstrated slightly better precision, making it suitable for applications where minimizing false positives is paramount. Surprisingly, the BiLSTM model did not significantly outperform the unidirectional models, suggesting that, for our dataset, processing information bidirectionally may not be essential. In conclusion, our study contributes valuable insights to the field of fake news detection. It underscores the significance of model selection based on specific task requirements and dataset characteristics.

Corresponding Author:

Gregorius Airlangga,
Information System Study Program
Atma Jaya Catholic University of Indonesia
Jl. Jend. Sudirman No.51 5, RT.004/RW.4, Daerah Khusus Ibukota Jakarta 12930, Indonesia
Emai: gregorius.airlangga@atmajaya.ac.id

1. Introduction

The advent of the digital information era has bestowed upon us an unparalleled wealth of data and news, reshaping the way we access and consume information [1]–[3]. However, this transformative change brings with it a formidable challenge: the rise of 'fake news.' This phenomenon, characterized by the deliberate creation and circulation of false or misleading information, often aims to influence public opinion or obscure the truth [4]–[6]. Its emergence has become a critical issue, demanding urgent and effective computational solutions to identify and mitigate the spread of such deceptive content. The concept of fake news, though historically present, has gained unprecedented momentum and impact with the proliferation of social media and online platforms [7]–[9]. Past studies tackling this issue have varied in approach, spanning from manual fact-checking to automated systems leveraging machine learning [10] and natural language processing (NLP) [11]. Initially, simpler methods like keyword-based filtering [12] were prevalent, but they often faltered in capturing the nuanced and intricate nature of human language.

The emergence of deep learning technologies marked a significant leap forward in the domain of fake news detection [13]–[15]. Recurrent Neural Networks (RNNs) were among the early adopters, favored for their capability to handle sequential data, making them apt for textual analysis. However, their effectiveness was hampered by limitations in processing long-term dependencies in text [16]. This challenge paved the way for more sophisticated architectures such as Long Short-Term Memory (LSTM) networks, which enhanced context capture through innovative memory cell designs [17]. More recent developments have seen the advent of Bidirectional LSTMs (BiLSTMs), enhancing LSTM capabilities by analyzing data in both forward and reverse directions. This bidirectional processing provides a more comprehensive context grasp, significantly improving performance in complex tasks such as text classification and sentiment analysis, which are crucial in discerning fake news [18].

Current research in fake news detection has shown promising results, especially with LSTM and BiLSTM models. These architectures excel over conventional RNNs in their ability to capture long-range dependencies within text data. The integration of attention mechanisms and the use of advanced pre-trained word embeddings like Word to Vector (Word2Vec) or Global Vectors for Word Representation (GloVe) [19] have further refined these models' accuracy, highlighting the potential of deep learning in combating fake news. In light of these developments, our research aims to provide a comprehensive comparison of RNN, LSTM, and BiLSTM models in the realm of fake news detection. We utilize a curated dataset from Word Embedding over Linguistic Features for Fake News Detection (WELFake) dataset [20], comprising news titles that undergo rigorous preprocessing using a suite of NLP techniques, including tokenization, stemming, and stopwords removal. These steps are vital in preparing the textual data for effective analysis by the neural network models.

Our study meticulously constructs three distinct models, each representing one of the neural network architectures mentioned earlier. Each model is equipped with an embedding layer for converting text into vector representations, followed by their respective recurrent layers - RNN, LSTM, and BiLSTM. The final layer in each model is a dense output layer with a sigmoid activation function, categorizing news items as 'fake' or 'real.' We rigorously evaluate the performance of each model based on metrics such as accuracy, loss, and others during both training and validation phases. This comprehensive evaluation seeks to elucidate each architecture's strengths and weaknesses in detecting fake news, thereby not only identifying the most effective model for this task but also enriching the discourse on how different neural network architectures manage language-based classification challenges.

The outcomes of this study are poised to significantly influence the development of robust and accurate fake news detection systems. In an era marked by information overload and the rampant spread of misinformation, such tools are crucial in safeguarding public discourse's integrity. Additionally, our research could catalyze future explorations into incorporating more complex NLP techniques or developing hybrid models that amalgamate the strengths of various neural network architectures. The urgency for conducting this research arises from the critical need to preserve the integrity of information in an increasingly digitized world. The proliferation of fake news poses a significant threat to various aspects of society, including politics, public health, and social harmony. Misinformation can lead to misinformed decisions, polarized public opinion, and even jeopardize democratic processes. In the context of public health, for example, the spread of false information about medical treatments or pandemics can have serious consequences.

Furthermore, the rapid advancement of technology and the ever-increasing sophistication of fake news dissemination methods make it imperative to continuously evolve and improve detection mechanisms. Traditional methods of fact-checking are increasingly insufficient in the face of automated, large-scale production of fake news [21]–[23]. Therefore, there is a pressing need for automated, efficient, and accurate systems capable of identifying and mitigating the spread of false information in real-time. Our research is not just an academic exercise; it addresses a real and pressing challenge in the modern information ecosystem [24]–[26]. By exploring and comparing the capabilities of advanced neural network architectures like RNN, LSTM, and BiLSTM in the context of fake news detection, we aim to contribute to the development of more sophisticated and reliable tools for safeguarding information authenticity. This is particularly crucial in an era where digital platforms have become primary sources of news and information for a significant portion of the global population.

Moreover, we provide remaining explanation structure for this paper, the second section is about 'Problem Formulation", we create this section in order to provide readers the problem that we handle and also its solution. The third section, 'Methodology,' delves into the specifics of our dataset and the preprocessing techniques employed. It also offers a detailed exposition of the architectures of the neural network models we utilized, namely RNN, LSTM, and BiLSTM, thereby laying the groundwork for understanding the subsequent analysis. In the 'Experimental Setup' section, we meticulously outline the design of our experiments. This includes a thorough description of the model training processes, validation procedures, and the metrics chosen for evaluating the performance of each model. These details are crucial for replicability and understanding the context of our results. The 'Results and Discussion' section forms the core of our paper, where we present a detailed analysis of the experimental results. Here, we not only compare the effectiveness of the RNN, LSTM, and BiLSTM models in the realm of fake news detection but also discuss the broader implications of our findings within the fields of Natural Language Processing (NLP) and machine learning. This section aims to bridge the gap between our specific research focus and its relevance to the wider scientific community. Finally, the paper culminates in the 'Conclusion' section, where we summarize the key findings of our study. We emphasize the significance of these findings in the ongoing battle against fake news and propose potential directions for future research in this increasingly vital field. This section aims to inspire further research and highlight the broader impact of our work.

2. Methods

2.1. Problem Formulation

Given a dataset $\mathcal{D} = \{(x_1, y_1), (x_2, y_2), \ldots, (x_N, y_N)\}$ where x_i represents the feature vector extracted from the i^{th} news article and y_i is the corresponding binary label with $y_i \in \{0,1\}$ (0 for real news and 1 for fake news), our objective is to learn a function $f: \mathcal{X} \to \mathcal{Y}$ that maps the feature space $\mathcal{X} \to \mathcal{Y}$, minimizing the prediction error on unseen data. The binary classification task can be modeled as presented in the equation (1).

$$(\hat{y} = 1) if (f(x; \theta) \geq 0.5), and (\hat{y} = 0)\ otherwise. \tag{1}$$

where $\hat{y}$ is the predicted label for the input x, and θ represents the parameters of the model. To solve the problem, we propose a neural network model parameterized by θ. The prediction function $\hat{f}(x; \theta)$ is realized through the network, and the goal is to find the optimal parameters θ^* that minimize the loss function $\mathcal{L}$. The loss function for binary cross-entropy is defined as $\mathcal{L}(\theta) = -\frac{1}{N}\sum_{i=1}^{N}\left[y_i \log\left(\hat{f}(x_i; \theta)\right) + (1 - y_i)\log\left(1 - \hat{f}(x_i; \theta)\right)\right]$. The optimization is typically performed using algorithms like stochastic gradient descent (SGD) or Adam. The update rule for SGD is $\theta_{t+1} = \theta_t - \alpha\nabla_{\theta\mathcal{L}}(\theta_t)$, where α is the learning rate. For architectures like LSTM, the cell state C_t and hidden state h_t at time step t are updated as presented in the equation (2)-(6).

$$f_t = \sigma\left(W_f \cdot [h_{t-1}, x_t] + b_f\right) \quad (2)$$

$$\mathrm{I_t} = \sigma(\mathrm{W_i} \cdot [\mathrm{h_{t-1}}, \mathrm{x_t}] + \mathrm{b_i}) \quad (3)$$

$$\mathrm{o_t} = \sigma(\mathrm{W_o} \cdot [\mathrm{h_{t-1}}, \mathrm{x_t}] + \mathrm{b_o}) \quad (4)$$

$$C_t = f_t * C_{t-1} + i_t * \tanh\,(W_C \cdot [h_{t-1}, x_t] + b_C \quad (5)$$

$$h_t = o_t * tanh(C_t) \quad (6)$$

Where f_t This represents the "forget gate" at time step t. It is a sigmoid activation function (σ) applied to the weighted sum of the previous hidden state (h_{t-1}) and the current input (x_t), with weights W_f and bias b_f. The forget gate determines what information from the previous cell state (C_{t-1}) should be discarded or remembered. For the i_t , this represents the "input gate" at time step t. Similar to the forget gate, it is a sigmoid activation function applied to the weighted sum of the previous hidden state (h_{t-1}) and the current input (x_t), with weights W_i and bias b_i. The input gate determines what new information from the current input should be added to the cell state.

In addition, o_t represents the "output gate" at time step t. Again, it is a sigmoid activation function applied to the weighted sum of the previous hidden state $\left(h_{\{t-1\}}\right)$ and the current input (x_t), with weights W_o and bias b_o. The output gate controls what information from the cell state (C_t) should be passed as the output of the LSTM cell. Furthermore, C_t represents the current cell state at time step t. It is updated using a combination of the forget gate (f_t) and the input gate (i_t). The forget gate determines what part of the previous cell state $\left(C_{\{t-1\}}\right)$ to keep, and the input gate determines what new information to add to the cell state. The tanh activation function is applied to the weighted sum of the previous hidden state $\left(h_{\{t-1\}}\right)$ and the current input(x_t), with weights W_C and bias b_C, to produce a new candidate cell state. This candidate cell state is then combined with the previous cell state using the forget and input gates.

Next, h_t represents the current hidden state at time step t. It is calculated by applying the output gate (o_t) to the hyperbolic tangent (tanh) of the current cell state (C_t). The tanh function squashes the values in the cell state to be between -1 and 1, and the output gate determines which parts of the squashed cell state should be included in the hidden state. In addition, σ is the sigmoid function and $*$ denotes element-wise multiplication. Performance metrics such as accuracy, precision, recall, and F1-score are defined as presented in the equation (7)-(10).

$$\text{Accuracy} = \frac{TP+TN}{TP+FP+FN+TN} \quad (7)$$

$$\text{Precision} = \frac{TP}{TP+FP} \quad (8)$$

$$\text{Recall} = \frac{TP}{TP+FN} \quad (9)$$

$$\text{F1-score} = 2 \cdot \frac{\text{Precision}\cdot\text{Recall}}{\text{Precision}+\text{Recall}} \quad (10)$$

where *TP* is true positives, *TN* is true negatives, *FP* is false positives, and *FN* is false negatives.

2.2. Experiment Setup

As presented in Figure 1, our experiment setup is meticulously crafted to ensure a thorough and comprehensive analysis of neural network architectures in the detection of fake news. The core of our study revolves around a carefully curated dataset obtained from WELFake [20], a well-regarded source known for its comprehensive compilation of news titles. The dataset consists of news titles that span a broad spectrum of topics, sources, and styles, offering a rich foundation for our analysis. The WELFake dataset was assembled using a systematic approach to ensure variety and veracity. The news titles were collected from multiple reputable and diverse sources, including established news agencies, online news portals, and social media platforms. This collection method was designed to encompass a wide range of linguistic styles and content types, from formal reporting to more colloquial forms of news dissemination. Each news title underwent a rigorous verification process, involving cross-referencing with trusted news databases and fact-checking tools to label them as 'fake' or 'real.' This meticulous process ensures the reliability and accuracy of our dataset. In preparing the dataset for our experiment, we divided it into two subsets: training and testing. The training subset, comprising 67% of the total dataset, is used to train the neural network models. This subset is designed to represent the full diversity of the dataset, encompassing a balanced mix of fake and real news titles across various topics and styles. The remaining 33% of the dataset forms the testing subset, which is used to evaluate the performance of the models. This split ensures that the models are exposed to and evaluated against a comprehensive and representative sample of the dataset.

The dataset's balanced nature, with an equal mix of 'fake' and 'real' news titles, is critical for ensuring an unbiased foundation for model evaluation. This balance is maintained in both the training and testing subsets, allowing for a fair and rigorous assessment of the models' capabilities. The diversity and breadth of the dataset are crucial for the robustness and reliability of our analysis, as they enable the models to be tested across various linguistic styles and content types, thereby enhancing the generalizability of our findings.The preprocessing of textual data forms a critical component of our methodology. This multistage process begins with tokenization, where the text is broken down into individual words or tokens, allowing for finer analysis and manipulation. Following this, stopword removal is conducted to eliminate commonly used words that contribute minimal informative value, thereby reducing data noise and focusing on the content that is more likely to be indicative of the veracity of the news. Subsequently, stemming is applied to simplify words to their base or root form. This helps in generalizing the models by grouping different forms of the same word, thus aiding in reducing the complexity of the linguistic analysis. The preprocessing pipeline also includes vectorization, where the processed text is converted into numerical formats such as one-hot encoding or word embeddings. This conversion is essential to make the data interpretable by neural network models. Lastly, padding is employed to standardize the length of the sequences, maintaining uniformity across the dataset, which is necessary for batch processing in neural network architectures.

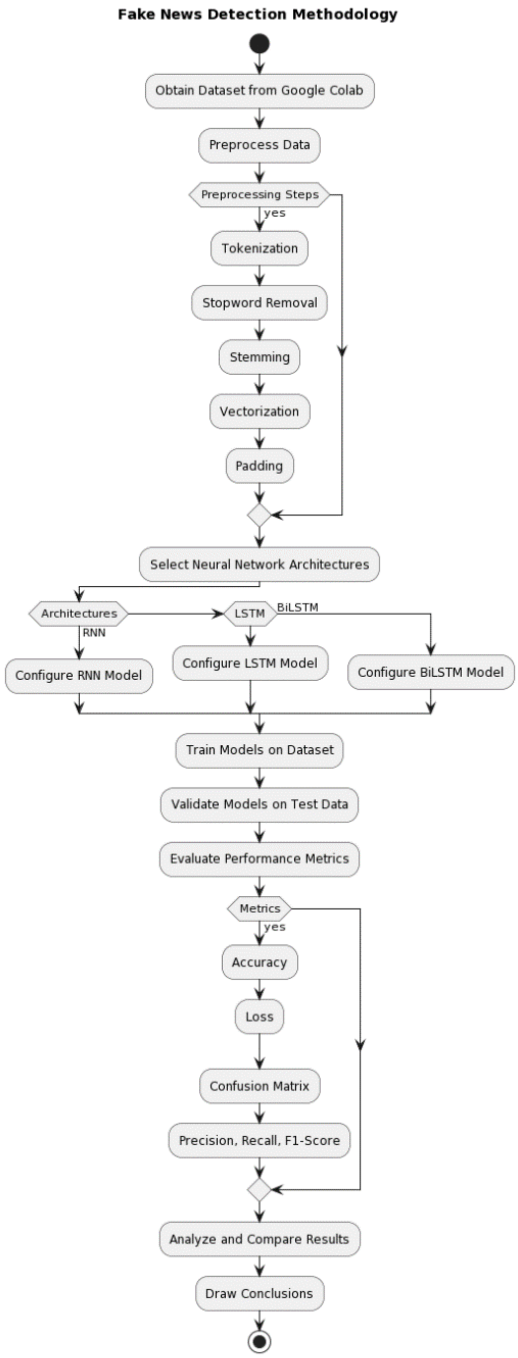

Figure. 1. Experiment Setup

The selection of neural network architectures is a pivotal aspect of our methodology. We have chosen three advanced architectures for our comparative analysis: the Recurrent Neural Network (RNN), Long Short-Term Memory (LSTM), and Bidirectional Long Short-Term Memory (BiLSTM). Each of these models is designed to cater to the nuances of sequential data processing, a characteristic inherent in text analysis. The RNN model, recognized for its ability to process sequential data, consists of an embedding layer for transforming input text into a numerical format, followed by recurrent layers that capture temporal information, and culminates in a dense output layer with a sigmoid activation function for binary classification. The primary limitation of the basic RNN is its struggle with long-term dependencies due to the vanishing gradient problem. The LSTM (Long Short-Term Memory) model, an advanced iteration of the RNN, addresses this limitation. Structured similarly with an embedding layer and a dense output layer, the key distinction lies in its LSTM layers. These layers incorporate memory cells that regulate the flow of information, allowing the model to both retain and discard data over extended sequences. This capability is vital for grasping the nuanced context in news titles, which often requires understanding information spread across the entire title. The BiLSTM (Bidirectional Long Short-Term Memory) model further extends the LSTM's capabilities. While it follows a similar structure with an embedding layer and a dense output layer, the pivotal difference is in its BiLSTM layers that process data in both forward and backward directions. This bidirectional processing enables the model to capture context from both past and future states in the sequence, offering a more comprehensive understanding of the text. This aspect is particularly beneficial in scenarios where the meaning of a word or phrase is heavily dependent on surrounding elements in the sequence. Theoretically, while the RNN, LSTM, and BiLSTM share a foundational structure, they differ significantly in their ability to handle long-term dependencies and contextual understanding. The LSTM improves upon the RNN's handling of long sequences, and the BiLSTM further enhances this by incorporating insights from both directions of a sequence.

Model training and validation are conducted with meticulous attention to detail. Each model undergoes rigorous training using the training subset of our dataset. To ensure adequate learning and to avoid overfitting, we employ a batch size of 32 and train the models over 20 epochs. The choice of the Adam optimizer is due to its efficiency and adaptability, particularly in handling sparse gradients and adaptive learning rate capabilities. Binary cross-entropy is used as the loss function, fitting the binary nature of our classification task. The validation of each model is performed on the testing subset, which is crucial for assessing their generalization capabilities and performance on unseen data. The performance of each model is evaluated using a comprehensive set of metrics. Accuracy measures the proportion of correctly predicted instances, providing a straightforward assessment of model performance. The loss metric indicates the model's performance during the training, with lower values signifying better performance. The confusion matrix provides a detailed view of the classification accuracy for each category, offering insights into the model's ability to correctly identify fake and real news. Additionally, precision, recall, and F1-score are employed to offer deeper insights into the model's performance, especially in handling imbalanced datasets.

Ethical considerations and data privacy form an integral part of our methodology. We adhere to the highest standards of ethical research practices, ensuring that the dataset used does not include any personal or sensitive information. Our models are designed to prevent biases, and all findings are presented with a commitment to impartiality and scientific integrity. This ethical framework is essential not only for the credibility of our research but also for ensuring that the developed models do not perpetuate or amplify existing biases.

The use of software and tools is carefully chosen to optimize the efficiency and effectiveness of our research. Python serves as the primary programming language due to its extensive support for machine learning and data processing tasks. TensorFlow is employed for constructing and training the neural network models, given its robustness and flexibility in handling deep learning algorithms. For data manipulation and numerical computations, we rely on Pandas and NumPy, which offer powerful data structures and mathematical functions that are vital for handling large datasets. Data visualization, an essential component for presenting our findings, is performed using Matplotlib and Seaborn, both of which provide a wide range of capabilities for creating comprehensive and insightful graphical representations of our results. The use of Google Colab as our computational platform provides us with the necessary computational power and flexibility. It enables us to handle the large-scale data and complex model architectures inherent in our study, ensuring that our computational resources are not a limiting factor in our research.

Our research critically examines the effectiveness of neural network models in detecting fake news through a dataset primarily composed of English news titles. Recognizing the dataset's limitations is crucial; it captures the essence of news content but may overlook the complexities found in full articles, potentially introducing language-specific biases and simplification issues due to preprocessing techniques like stemming and stopword removal. Despite these constraints, we ensure our models' reproducibility by detailing our methodology, including dataset preparation, model architecture, and evaluation metrics. The data annotation process, labeling news titles as 'fake' or 'real' based on reputable sources, underpins our models' training, emphasizing the importance of high-quality data for reliable outcomes. Hyperparameter tuning plays a pivotal role in optimizing model performance, meticulously adjusting parameters to prevent overfitting and enhance fake news detection accuracy. In evaluating our models, we conducted a comparative analysis to identify their unique strengths and limitations, offering insights into their suitability for various text classification tasks. This approach not only highlights the individual model's capabilities but also sets the stage for future research expansions in the domain. The findings, while focused on fake news detection, lay a groundwork for broader applications in text classification, encouraging further advancements in the field. Our study, therefore, not only contributes to understanding fake news detection mechanisms but also opens avenues for future explorations, aiming to improve and adapt these models for evolving challenges in digital content analysis.

3. Results and Discussion

The results presented in the Table 1-4 are the performance metrics of three different neural network models applied to a classification task—specifically, fake news detection. As presented in the Table 1 and 2, the RNN model has a precision of 0.93 for class 0 (presumably 'real' news) and 0.89 for class 1 (presumably 'fake' news). This means that when it predicts news as 'real', it is correct 93% of the time, and when it predicts news as 'fake', it is correct 89% of the time. In addition, the recall is 0.91 for both classes, indicating that the model correctly identifies 91% of the actual 'real' and 'fake' news. The F1-score, which balances precision and recall, is 0.92 for class 0 and 0.90 for class 1, suggesting a slightly better performance on class 0. The support is the number of true instances for each class, indicating that the dataset contains 3419 instances of class 0 and 2616 of class 1. Furthermore, the model's overall accuracy is 0.91, meaning it correctly predicts the news class 91% of the time across both classes.

Table 1. Class 0 Detection Comparison Results

Methods	Precision	Recall	F1-Score	Support
RNN	0.93	0.91	0.92	
LSTM	0.94	0.91	0.92	3419
Bi-LSTM	0.92	0.91	0.92	

Table 2. Class 1 Detection Comparison Results

Methods	Precision	Recall	F1-Score	Support
RNN	0.89	0.91	0.90	
LSTM	0.88	0.92	0.90	2616
Bi-LSTM	0.89	0.90	0.89	

Table 3. Macro Average Comparison Results

Methods	Precision	Recall	F1-Score	Accuracy
RNN	0.91	0.91	0.91	0.91
LSTM	0.91	0.91	0.91	0.91
Bi-LSTM	0.90	0.90	0.90	0.90

Table 4. Weighted Average Comparison Results

Methods	Precision	Recall	F1-Score	Accuracy
RNN	0.91	0.91	0.91	0.91
LSTM	0.91	0.91	0.91	0.91
Bi-LSTM	0.91	0.91	0.91	0.91

As presented in the Table 1 and 2, the LSTM model shows a slight improvement in precision for class 0 at 0.94 but a decrease for class 1 at 0.88. The recall for class 0 remains the same at 0.91, but there is a slight improvement for class 1 at 0.92. Furthermore, the F1-scores are 0.92 for class 0 and 0.90 for class

1, identical to the RNN model, which indicates a consistent balance between precision and recall for both classes. In addition, the support numbers are the same as in the RNN results, with 3419 instances of class 0 and 2616 of class 1. The overall accuracy as presented in the Table 3 and 4 is still 0.91, indicating no overall improvement from the RNN model.

Furthermore, as presented in the Table 1 and 2, for bidirectional LSTM method, the precision for class 0 is slightly lower at 0.92, and for class 1, it remains at 0.89. Then, the recall for class 0 is the same at 0.91, while for class 1, it has decreased slightly to 0.90. In addition, the F1-score for class 0 is 0.92, and for class 1, it has decreased slightly to 0.89.The support is consistent with the previous models, with 3419 instances of class 0 and 2616 of class 1. The overall accuracy is 0.91 as presented in the Table 3 and 4, consistent with the other two models.

The comparison of the three neural network models, namely RNN, LSTM, and Bidirectional LSTM, reveals that they all exhibit similar overall accuracies, averaging at 0.91 as presented in the Table 3 and 4. This consistency in performance indicates that all three models are effective in addressing the task of fake news detection. However, delving deeper into the results and their implications provides valuable insights into the strengths and potential applications of each model. The LSTM model, characterized by its ability to capture long-range dependencies in text, showcases a slightly superior recall for class 1 when compared to the RNN. This implies that the LSTM model is more adept at correctly identifying 'fake' news instances. This characteristic is crucial, particularly in scenarios where the consequences of missing 'fake' news can have far-reaching repercussions, such as public opinion manipulation or the spread of false information that may harm individuals or societies.

The Recurrent Neural Network (RNN) model stands out for its high precision in classifying class 1 ('fake' news), surpassing both the LSTM and BiLSTM models. This precision is crucial in contexts where it's imperative to maintain the credibility of 'real' news, such as in journalistic settings or content moderation platforms. The RNN's strength in minimizing the mislabeling of 'real' news as 'fake' suggests its potential as a reliable tool in environments where false positives could have serious reputational impacts. Comparatively, the LSTM model, while not excelling in precision to the same extent as the RNN, demonstrates better recall. This characteristic makes it suitable for scenarios where missing a 'fake' news item is more detrimental than incorrectly labeling a 'real' one. This nuanced distinction between precision and recall in LSTM and RNN highlights the importance of model selection based on specific application needs.

Our findings highlight the distinct strengths and weaknesses of the Recurrent Neural Network (RNN), Long Short-Term Memory (LSTM), and Bidirectional Long Short-Term Memory (BiLSTM) models in fake news detection, offering valuable insights for their practical application. The RNN model, with its simpler structure, is less computationally intensive and easier to implement. This makes it an accessible choice for scenarios with limited computational resources or where rapid deployment is necessary. However, its major drawback lies in its limited ability to process long sequences effectively, which can be crucial in understanding the context of certain news titles. Hence, while the RNN demonstrates commendable precision, especially in minimizing false positives for 'real' news classification, it may not always capture the full context of longer or more complex news titles.

In contrast, the LSTM model, designed to address the limitations of RNNs, excels in handling long-term dependencies. This makes it particularly effective in scenarios where the context spread across a news title is essential for accurate classification. Our study found that the LSTM model shows improved recall in detecting 'fake' news, suggesting its suitability in applications where missing out on fake news has serious consequences. However, the increased complexity of the LSTM model requires more computational resources and time for training and inference, which can be a limiting factor in certain applications. The BiLSTM model extends the capabilities of the LSTM by processing sequences in both forward and backward directions, offering a more comprehensive understanding of context. This theoretically positions the BiLSTM as the most capable model in terms of contextual understanding. However, our study did not find a marked improvement in performance with the BiLSTM over the LSTM and RNN models in the specific context of fake news detection. This suggests that the additional complexity and computational requirements of the BiLSTM may not always translate to significantly better performance in all scenarios. The BiLSTM model might be more beneficial in cases where the context is highly complex and nuanced, necessitating a deep understanding of both preceding and subsequent information in the sequence. Therefore, while all three models demonstrate robust performance in fake news detection, the choice of model should be informed by the specific strengths and weaknesses of each, aligned with the requirements of the task at hand. Our research sheds light on these nuances, aiding practitioners and decision-makers in the field of misinformation detection in making informed choices about the most suitable model for their specific needs.

4. Conclusion

In summary, our comparative study of Recurrent Neural Network (RNN), Long Short-Term Memory (LSTM), and Bidirectional Long Short-Term Memory (BiLSTM) models in fake news detection has yielded insightful findings. All three models demonstrated high accuracy, precision, recall, and F1-scores, around the 0.91 mark, indicating their effectiveness in classifying textual data in the context of fake news. The RNN, despite its simplicity, proved reliable, particularly in minimizing false positives. The LSTM showed a slight edge over the RNN in terms of recall, suggesting a better capability in identifying fake news. The BiLSTM, while advanced in theory, did not show a significant improvement in performance, suggesting that the added complexity of bidirectional processing might not always translate to better results in this specific application. However, our study is not without limitations. One key constraint is the reliance on a single dataset, which may not fully capture the diversity of linguistic styles and nuances present in real-world fake news. This limitation raises questions about the generalizability of our findings across different datasets and languages. Additionally, the performance of the BiLSTM model indicates that further investigation is needed to understand when and how bidirectional processing contributes to improved outcomes in fake news detection. For future research, we suggest exploring hybrid models that combine the strengths of RNNs, LSTMs, and BiLSTMs, as well as incorporating attention mechanisms and sophisticated contextual embeddings to potentially enhance performance. It would also be valuable to test these models on a more diverse set of datasets, including those in different languages and formats, to better understand their applicability and limitations in various contexts.

References

[1] J. W. Salmon, S. L. Thompson, J. W. Salmon, and S. L. Thompson, "Big data: information technology as control over the profession of medicine," *Corp. Am. Heal. Care Rise Corp. Hegemony Loss Prof. Auton.*, pp. 181–254, 2021.

[2] D. Bawden and L. Robinson, "Information overload: An overview," 2020.

[3] G. Newton, K. Drysdale, M. Zappavigna, and C. E. Newman, "Truth, proof, sleuth: trust in direct-to-consumer DNA testing and other sources of identity information among Australian donor-conceived people," *Sociology*, vol. 57, no. 1, pp. 36–53, 2023.

[4] J. P. Baptista and A. Gradim, "Understanding fake news consumption: A review," *Soc. Sci.*, vol. 9, no. 10, p. 185, 2020.

[5] J.-N. Kim and H. de Zúñiga, "Pseudo-information, media, publics, and the failing marketplace of ideas: Theory," *Am. Behav. Sci.*, vol. 65, no. 2, pp. 163–179, 2021.

[6] M. Freeze *et al.*, "Fake claims of fake news: Political misinformation, warnings, and the tainted truth effect," *Polit. Behav.*, vol. 43, pp. 1433–1465, 2021.

[7] N. A. Al Shehab, "the dark side of social media: spreading misleading information during covid-19 crisis," *Adv. Data Sci. Intell. Data Commun. Technol. COVID-19 Innov. Solut. Against COVID-19*, pp. 277–306, 2022.

[8] H. N. Chua, Q. Khan, M. B. Jasser, and R. T. K. Wong, "Problem Understanding of Fake News Detection from a Data Mining Perspective," in *2023 IEEE 13th International Conference on Control System, Computing and Engineering (ICCSCE)*, 2023, pp. 297–302.

[9] N. Zaidi, M. Maurya, S. Grima, and P. Tyagi, "Unveiling AI's Ethical Impact in Marketing Through Social Media's Darker Influence," in *Building AI Driven Marketing Capabilities: Understand Customer Needs and Deliver Value Through AI*, Springer, 2023, pp. 173–193.

[10] Z. Guo, M. Schlichtkrull, and A. Vlachos, "A survey on automated fact-checking," *Trans. Assoc. Comput. Linguist.*, vol. 10, pp. 178–206, 2022.

[11] Q. Su, M. Wan, X. Liu, C.-R. Huang, and others, "Motivations, methods and metrics of misinformation detection: an NLP perspective," *Nat. Lang. Process. Res.*, vol. 1, no. 1–2, pp. 1–13, 2020.

[12] S. Adak *et al.*, "Mining the online infosphere: A survey," *Wiley Interdiscip. Rev. Data Min. Knowl. Discov.*, vol. 12, no. 5, p. e1453, 2022.

[13] R. Varma, Y. Verma, P. Vijayvargiya, and P. P. Churi, "A systematic survey on deep learning and machine learning approaches of fake news detection in the pre-and post-COVID-19 pandemic," *Int. J. Intell. Comput. Cybern.*, vol. 14, no. 4, pp. 617–646, 2021.

[14] K. Choraś Michałand Demestichas, A. Giełczyk, Á. Herrero, K. Ksieniewicz Pawełand Remoundou, D. Urda, and M. Woźniak, "Advanced Machine Learning techniques for fake news (online disinformation) detection: A systematic mapping study," *Appl. Soft Comput.*, vol. 101, p. 107050, 2021.

[15] A. Agarwal, M. Mittal, A. Pathak, and L. M. Goyal, "Fake news detection using a blend of neural networks: An application of deep learning," *SN Comput. Sci.*, vol. 1, pp. 1–9, 2020.

[16] B. Jang, M. Kim, G. Harerimana, S. Kang, and J. W. Kim, "Bi-LSTM model to increase accuracy in text classification: Combining Word2vec CNN and attention mechanism," *Appl. Sci.*, vol. 10, no. 17, p. 5841, 2020.

[17] I. Al-Nader, A. Lasebae, R. Raheem, and A. Khoshkholghi, "A Novel Scheduling Algorithm for Improved Performance of Multi-Objective Safety-Critical Wireless Sensor Networks Using Long Short-Term Memory," *Electronics*, vol. 12, no. 23, p. 4766, 2023.

[18] A. Altheneyan and A. Alhadlaq, "Big data ML-based fake news detection using distributed learning," *IEEE Access*, vol. 11, pp. 29447–29463, 2023.

[19] S. Shreyashree, P. Sunagar, S. Rajarajeswari, and A. Kanavalli, "BERT-Based Hybrid RNN Model for Multi-class Text Classification to Study the Effect of Pre-trained Word Embeddings," *Int. J. Adv. Comput. Sci. Appl.*, vol. 13, no. 9, 2022.

[20] P. K. Verma, P. Agrawal, I. Amorim, and R. Prodan, "WELFake: word embedding over linguistic features for fake news detection," *IEEE Trans. Comput. Soc. Syst.*, vol. 8, no. 4, pp. 881–893, 2021.

[21] M. Tajrian, A. Rahman, M. A. Kabir, and M. R. Islam, "A review of methodologies for fake news analysis," *IEEE Access*, 2023.

[22] W. Ceron, M.-F. de-Lima-Santos, and M. G. Quiles, "Fake news agenda in the era of COVID-19: Identifying trends through fact-checking content," *Online Soc. Networks Media*, vol. 21, p. 100116, 2021.

[23] H. Liao *et al.*, "MUSER: A MUlti-Step Evidence Retrieval Enhancement Framework for Fake News Detection," in *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2023, pp. 4461–4472.

[24] X. Zhou and R. Zafarani, "A survey of fake news: Fundamental theories, detection methods, and opportunities," *ACM Comput. Surv.*, vol. 53, no. 5, pp. 1–40, 2020.

[25] W. Ansar and S. Goswami, "Combating the menace: A survey on characterization and detection of fake news from a data science perspective," *Int. J. Inf. Manag. Data Insights*, vol. 1, no. 2, p. 100052, 2021.

[26] J. Zeng, Y. Zhang, and X. Ma, "Fake news detection for epidemic emergencies via deep correlations between text and images," *Sustain. Cities Soc.*, vol. 66, p. 102652, 2021.