Building a Fortress Against Fake News

Harnessing the Power of Subfields in Artificial Intelligence

Nafiz Fahad
Faculty of Science & Technology, American International University-Bangladesh, Dhaka, Bangladesh

Kah Ong Michael Goh
Faculty of Information Science & Technology (FIST), Multimedia University, Melaka, Malaysia

Md Ismail Hossen
Faculty of Science & Technology, American International University-Bangladesh, Dhaka, Bangladesh

Connie Tee
Faculty of Information Science & Technology (FIST), Multimedia University, Melaka, Malaysia

Md Asraf Ali
Faculty of Science & Technology, American International University-Bangladesh, Dhaka, Bangladesh

Abstract: Given the prevalence of fake news in today's tech-driven era, an urgent need exists for an automated mechanism to effectively curb its dissemination. This research aims to demonstrate the impacts of fake news through a literature review and establish a reliable system for identifying it using machine (ML) learning classifiers. By combining CNN, RNN, and ANN models, a novel model is proposed to detect fake news with 94.5% accuracy. Prior studies have successfully employed ML algorithms to identify false information by analysing textual and visual features in standard datasets. The comprehensive literature review emphasises the consequences of fake news on individuals, economies, societies, politics, and free expression. The proposed hybrid model, trained on extensive data and evaluated using accuracy, precision and recall measures, outperforms existing models. This study underscores the importance of developing automated systems to counter the spread of fake news and calls for further research in this domain.

Keywords: machine learning (ML), hybrid model, automated system, accuracy, fake news.

Introduction

In today's digital age, the identification of fake news has emerged as a pressing concern. The proliferation of social media platforms as primary sources of news consumption and sharing has led to the rapid dissemination of both genuine and fake stories, posing serious consequences for society. Differentiating between various forms of false news on platforms like Twitter remains a major obstacle in the effective detection of fake news (Altheneyan & Alhadlaq, 2023). This is why the problem of fake news is a significant issue that puts the credibility of social networks at risk (Rawat *et al.*, 2023). Detecting fake news manually is not possible, which is why an automatic system is required. This system is created using the subfields of artificial intelligence (AI). Similarly, demonstrating the effects of fake news automatically is not feasible; therefore, a systematic literature review (SLR) is necessary to illustrate the impacts of fake news.

Objectives

Our work aims to prevent fake news by using the power of AI's subfields. In essence, our target is to make an efficient automated system to prevent fake news and our objectives are:

1. To create a dependable and precise system that uses machine learning (ML) classifiers to effectively identify fake news.
2. To make a hybrid model with the help of CNNs (convolution neural networks), RNNs (recurrent neural networks), and ANNs (artificial neural networks). No one has done this before.
3. To find out the impacts of fake news with achieve higher accuracy, precision, recall and F_1 scores, which would ensure that the system can detect most fake news accurately.

Literature Review

In the realm of literature, numerous publications exhibit a keen fascination with the identification of fake news.

Altheneyan & Alhadlaq (2023) said that the proliferation of social media platforms as primary sources of news consumption and sharing has led to the rapid dissemination of both genuine and fake stories. The prevalence of misinformation on these platforms has serious consequences for society. Detecting and differentiating between various forms of false news on platforms like Twitter pose a major obstacle to effective fake news detection, which is why the researchers used distributed learning to detect fake news based on ML. In their study, they used a distributed Spark cluster to construct a stacked ensemble model which achieved 93.40% accuracy; the highest among all the approaches (Altheneyan & Alhadlaq, 2023). In

http://doi.org/10.18080/jtde.v11n3.765

contrast, Rawat *et al.* (2023) revealed that using supervised ML algorithms and suitable tools facilitates the differentiation of false information from authentic news by categorising them accordingly (Rawat *et al.*, 2023). They also recognised the potential of using ML techniques as a possible solution for detecting fake news. In contrast, Sharma *et al.* (2023) used Hybrid Ensemble Model with Fuzzy Logic, which achieved 86.8% accuracy for the two datasets in their proposed model.

Singhal *et al.* (2019) developed a framework called SpotFake to identify fake news, which integrates language models with a pre-trained VGG-19 model on ImageNet to incorporate contextual information. The concatenation technique combined text and visual features to create a multimodal fusion module. The investigation revealed accuracy rates of 77.77% and 89.23% on the publicly available Twitter and Weibo datasets, respectively. Building upon the ideas presented in Singhal et al. (2019), the authors introduced SpotFake+ (Singhal *et al.*, 2020), an enhanced version of SpotFake that utilised transfer learning to extract semantic and contextual information from lengthy news articles and images.

Aslam *et al.* (2021) focused on developing an ensemble model based on deep learning techniques, specifically designed for identifying and detecting fake news. The proposed research aims to stop the spread of rumours and fake news by automatically categorising news articles, enabling people to determine if a news source is reliable or not. They developed an ensemble-based deep learning model to categorise news articles as either fake or real. The dataset underwent pre-processing techniques, and natural language processing (NLP) techniques were specifically applied to the statement attribute. Two different deep learning models were used: a deep learning dense model for nine attributes excluding the statement, and a Bi-LSTM-GRU-dense deep learning model for the statement attribute. The results of the study were highly significant, achieving an accuracy of 0.898 when using the statement feature. Although their proposed study has yielded noteworthy outcomes, there is still room for improvement. Further investigations are necessary to test the model using additional datasets of fake news.

Ahmad *et al.* (2022) worked on developing an improved deep-learning model to create an efficient mechanism for detecting fake news. Their goal was to extensively investigate the challenges associated with automatically detecting rumors on social media. To achieve this, we employ a novel combination of content-based and social-based features specifically designed for identifying rumours. Additionally, we employ a bidirectional LSTM-RNN classifier, a deep learning model, to analyse text data in order to enhance the accuracy of our rumour detection system. Amad *et al*'s bidirectional LSTM-RNN classifiers detect fake news or true news properly. A limitation of their approach is that more comprehensive testing is required to gain a better understanding of how deep learning can effectively detect rumours.

Additionally, the presence of a large amount of unlabelled data on social media poses a challenge, and developing models that can work without relying on labelled data becomes necessary.

Although the previously mentioned studies focused on extensive findings regarding the identification of fake news, there is still a need for further updates and expansion due to the significant increase in the volume of topical publications. It is also evident that no one can ensure that the system can detect most fake news accurately and help humans protect themselves from fake news, whereas our study achieves these issues.

SLR to reveal impacts of fake news

Using a methodical examination of published works, an SLR proves to be a highly efficient strategy for uncovering the impacts of fake news. The following table demonstrates the data we have gathered from diverse research papers to demonstrate the impacts of fake news.

Table 1. Impacts of fake news

Serial Number	**Citation**	**Extracted information**
1	Ahinkorah *et al.*, 2020	People argue that different types of purposeful damage, and different rewards like money, social recognition and political advantages, often motivate the spread of false information.
2	De Oliveira & Albuquerque, 2021	False information can directly impact a person's ability to stay alive.
3	Leeder, 2019	Students are led astray by false information.
4	Shu *et al.*, 2019	Fake news was created to confuse people and cause doubt, making it harder for them to tell what is true.
5	Bakir & McStay, 2018	The use of emotions to grab people's attention and make money for advertisers is a key factor in the issue of fake news. It also highlights the pressures to create automated fake news that caters to the emotions and behaviours of online social groups.
6	Zafarani *et al.*, 2019	False information impacts financial markets and also leads to significant trade disruptions in economies.
7	Sullivan, 2019	People cannot tell if information is trustworthy, so they believe fake news. As a result, people who want to take advantage of the situation spread wrong information online, which can be dangerous and have a frightening impact on real people.
8	Bago *et al.*, 2020	Bigger changes in society and politics, along with the right to express oneself, have encountered difficulties due to the spread of false information.
9	Almenar *et al.*, 2021	The types of wrong information that people receive also differ based on gender. We know that men and women have different behaviours on social media and when consuming news, but these differences are not as noticeable when it comes to false information. The main idea is that, as in other aspects of life, women tend to worry more than men because of the spread of false information.
10	Butler *et al.*, 2018	Fake news is a big issue that makes it hard for people to trust real news sources. It is a sizeable problem because it also makes it harder for the government to be trusted. Fake news hurts real news sources by rendering them seem less reliable, which means

Serial Number	Citation	Extracted information
		people are not equipped with the right information to be involved in a democracy.
11	Stewart, 2021	Fake news can cause harm by spreading violent threats and misleading information that can hurt people mentally or in other ways. It can also be dangerous when it misleads the public about important matters like elections or health.
12	Naeem *et al*, 2021	Reading false information can make people mentally unwell, and it can even put their health in danger.
13	Lakshmanan *et al.*, 2019	False information occasionally discourages individuals.
14	Ghosh & Shah, 2018	Sometimes, individuals read false information and it makes them feel sad and unable to stop doing it repeatedly.
15	Bhatt *et al.*, 2018	False information can damage the democratic system.
16	Pearson, 2017	Fake news destroys public safety.
17	Creech, 2020	False information spread on platforms like Facebook and Twitter also undermines trust.
18	Ho *et al.*, 2022	Fake information and false news on social media usually have negative effects and sometimes annoy everyday people, authorities and/or the government.
19	Khan *et al.*, 2022	Fake news is made to provoke extreme feelings, influence political activities, or create conflict and confusion in society.

Requirement of a hybrid model

After considering the literature review aforementioned in the above table found that the existing classifiers' results were not adequate compared to the proposed models', that the existing classifiers' accuracies were also less than the proposed models', and also found that, due to inaccuracy, no one can properly detect fake or true news and there is still room for improvement, we will first use the existing model followed by making a hybrid model to find an efficient way to detect fake or real news.

Methodology

Our objectives will be fulfilled with the help of ML and deep learning. First, we need to collect a dataset of news so that we can use that data to predict fake news or real news. Basically, we use a dataset consisting of several thousands of news articles and label them as either real news or fake news. Once we have labelled the dataset, we will pre-process the data. A lot of work is involved in this pre-processing step when compared to numerical data because computers and systems don't understand the text or characters. We therefore needs to find a suitable way to convert this text present in the news into meaningful numbers that the machine can understand. Once we pre-process the data or convert the text into meaningful numbers, we then need to split the dataset into training and test data because we need to train our ML model with the training dataset. So, we feed this training data which is pre-processed to our supervised ML classifiers and deep learning.

Proposed method

Figure 1 presents a visual representation of the suggested strategies and sequential actions involved in the proposed method. This diagram illustrates the sequence of these steps in the proposed approach.

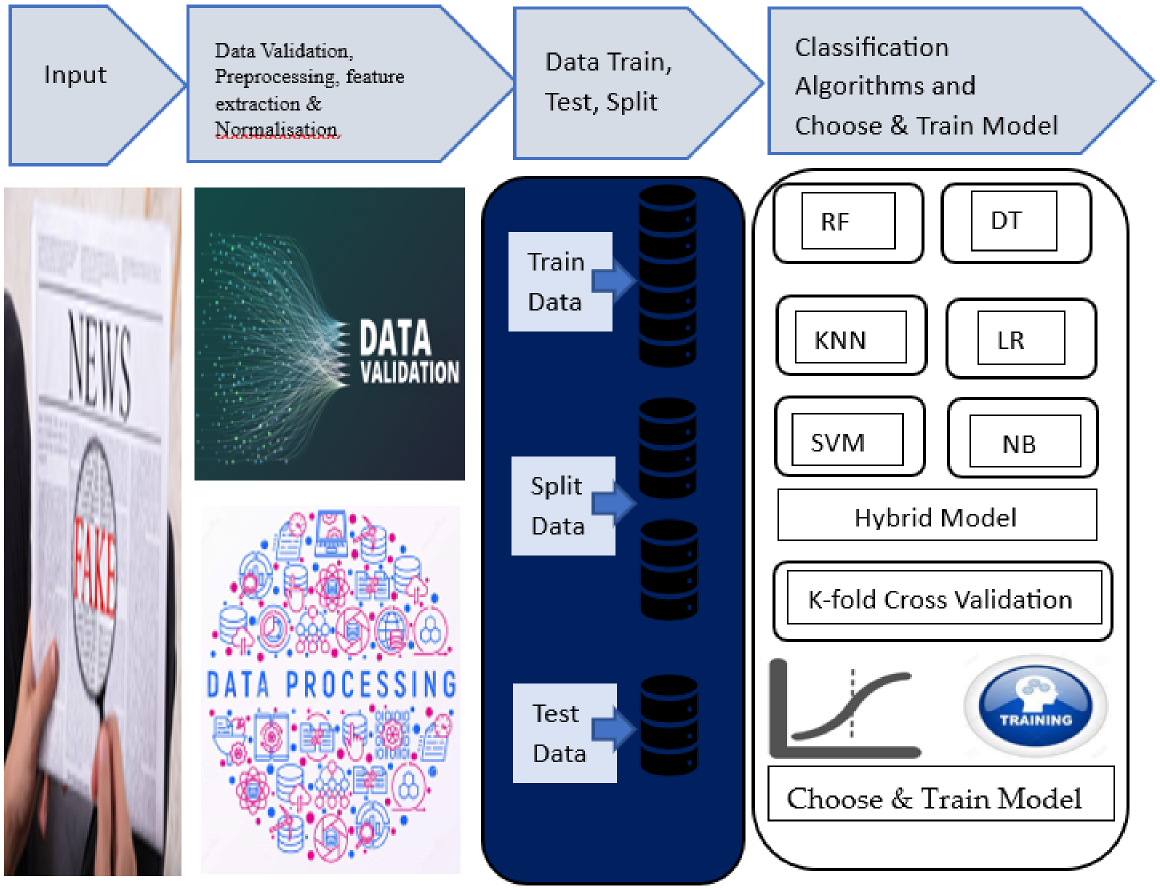

Figure 1. Block diagram

Data collection

The study employs a widely used dataset accessible on Kaggle, comprising 7796 entries and four columns. The initial column acts as a unique identifier for the news, while the second and third columns hold the news article's title and text correspondingly. The fourth column contains labels denoting whether the news is categorised as REAL or FAKE (Mahmoud, 2022).

Data validation

To verify the dataset's validity, we employ the langdetect library to determine if the text in each row of the CSV file is written in English. This process involves iterating through the DataFrame, extracting the text from each row, and using the is_english() function to assess its language. If the text is identified as English, a corresponding message is printed. Conversely, if the text is found to be non-English, an appropriate message is printed. Notably, the dataset in question is used by Mahmoud and Kokiantonis for their respective experiments (Mahmoud, 2022; Kokiantonis, 2022).

Data pre-processing

We pre-process the dataset by performing label encoding to convert categorical labels into numerical values. We also remove rows with missing values and eliminate duplicate rows from the DataFrame.

Feature extraction

For the feature extraction process, we use the TF-IDF (Term Frequency – Inverse Document Frequency) technique to convert the text data into numerical features. This is done using the 'TfidfVectorizer' class from the 'sklearn.feature_extraction.text' module. The text data is first combined from the 'title' and 'text' columns, and then the vectorizer is fit on this combined text data. The 'fit' operation calculates the term frequencies and inverse document frequencies. Next, the 'transform' operation converts the text data into TF-IDF feature vectors.

Normalisation

We also use the normalisation process. The feature vectors are normalised using the 'normalise' function from the 'sklearn.preprocessing' module, resulting in normalised TF-IDF vectors that can be used as input for an ML model.

Classification algorithms and hybrid model

Taking into account the characteristics of our dataset, we employ Logistic Regression (LR), Decision Tree (DT), K-Nearest Neighbours (KNN), Random Forest (RF), Support Vector Machine (SVM) and Naïve Bayes (NB) algorithms to investigate the behaviour of the data when subjected to different classifiers. Furthermore, a hybrid model is constructed by combining CNNs, RNNs and ANNs.

LR

Given that we are categorising text or content using a wide range of features and aiming for a binary output (such as true/false or authentic/fake article), the LR model is employed. The LR model offers a straightforward cost-function equation, allowing for classification between binary or multiple classes. In order to achieve optimal results for a specific dataset, we fine-tune the hyperparameters. Various parameters are assessed before obtaining the LR model, which serve as a benchmark for accuracy. Additionally, logistic regression utilises a sigmoid function to transform the output into a probability value. The objective is to attain an optimal probability that minimises the cost function (Pérez-Rosas *et al.*, 2017).

DT

DT, an indispensable tool, exhibits a flow-chart-like structure primarily designed for addressing classification problems. The DT relies on conditions or 'tests' applied to attribute results at internal nodes to determine its branches. As attributes are evaluated, the leaf nodes are assigned class labels. The path from the root to the leaf represents a classification rule, making it versatile for both categorical and dependent variables. Notably, DTs excel in identifying crucial factors and illustrating their relationships, contributing significantly to the development of new variables and insightful data exploration. These tree-based learning algorithms, also known as CART, play a crucial role in constructing accurate predictive models through supervised learning techniques. Their strength lies in effectively capturing non-linear relationships and offering solutions for classification or regression problems (Meel & Vishwakarma, 2021).

KNN

By omitting the need for a dependent variable, KNN demonstrates its ability to make predictions for specific data outcomes. Sufficient training data is provided to enable KNN to accurately determine the exact cluster or category to which a given data point belongs. The value of K determines the number of neighbouring data points considered, and the KNN model calculates the distance between a new data point and its nearest neighbours. If K is set to 1, the new data point is assigned to the class with the closest distance.

RF

RF is a methodology that utilises the amalgamation of numerous DTs or algorithms with similar characteristics in a collection of trees. The RF technique is applicable to both classification and regression tasks.

SVM

SVM is a type of supervised ML model employed by specific classification algorithms. This model excels in solving problems where data is divided into two distinct groups. By training the SVM model with a set of data, it can effectively classify future instances. In scenarios with limited samples, SVM demonstrates superior speed and performance compared to other models. The SVM classifier can be visualised as a straightforward two-dimensional line. It takes data points as input and generates a hyperplane that separates different categories. This line serves as the decision boundary, with one side representing the 'blue' category and the other side representing the 'red' category. The proximity of a data point to the hyperplane determines its assigned tag, with the nearest point having the largest influence and vice versa.

http://doi.org/10.18080/jtde.v11n3.765

Hybrid model

The hybrid model which we name the NFCRA (Nafiz Fahad CNNs, RNNs, and ANNs) model combines CNNs, RNNs, and ANNs to classify news articles as fake or real. It processes the text data by organising it into sequences, then creates separate parts for each architecture. The outputs from each part are combined and passed through a dense layer, using a sigmoid function to classify them into two categories. The model is trained using the training data, and evaluated with the test data.

CNN

CNN is a sophisticated ML framework initially developed for interpreting visual information, yet it can also be effectively employed to handle textual information in natural language-processing endeavours. It uses filters or kernels to scan through input text, extracting local patterns or features. The input text is represented as numerical vectors using word or character embeddings. Convolutional layers perform convolutions on the input text, generating feature maps. Pooling layers downsample the feature maps to capture important features and reduce noise. Finally, fully connected layers make predictions or classifications. CNNs can be effectively trained on labelled text datasets by employing optimisation techniques such as stochastic gradient descent. Additionally, they can be further customised to excel in specific NLP tasks through fine-tuning.

RNN

RNN is a type of deep learning model for processing sequential data, including text. It uses feedback loops to capture information from previous time steps and maintain context. RNNs can have different architectures such as simple RNNs, LSTMs and GRUs, with varying memory and capability for handling long-term dependencies. They are widely used in NLP tasks such as text classification and sequence generation. RNNs can be trained using labelled text datasets and optimisation techniques, but have some limitations like vanishing gradients.

ANN

ANN is a computational model used in ML to analyse textual information by employing a network of interconnected nodes. The input text is represented as numerical vectors, which are processed through hidden layers with activation functions. The output layer produces predictions or classifications. ANN learns from labelled text datasets through parameter updates during training. ANN is versatile for various NLP tasks but may have limitations in capturing sequential dependencies and require large amounts of data for training.

Data analysis techniques

We evaluate our data using accuracy, precision, recall and F_1 scores as our analysis metrics. The analysis techniques are depicted below.

Accuracy

Accuracy is the predominant metric used to assess the proportion of correctly predicted outcomes, whether they are true or false. The following equation can be employed to calculate the accuracy of a model.

$$Accuracy = \frac{TP+FP}{(TP+FP+TN+FN)} \tag{1}$$

Precision

Precision is a measure that evaluates how accurately a model predicts positive outcomes. In our study, we determine precision by dividing the count of correctly predicted positive results by the overall count of positive predictions.

$$Precision = \frac{TP}{(TP+FP)} \tag{2}$$

Recall

Recall, which represents the total number of correctly classified instances excluding the true class, is a key measure in our experiment. It specifically refers to the percentage of articles among all accurately predicted articles that were correctly anticipated.

$$Recall = \frac{TP}{TP+FN} \tag{3}$$

F_1

The F_1 score is a single number that combines precision and recall, giving an overall measure of how well a model performs in tasks where it has to classify areas into two categories.

$$F1 = \frac{2\times(Precision \times Recall)}{Precsion+Recall} \tag{4}$$

Result

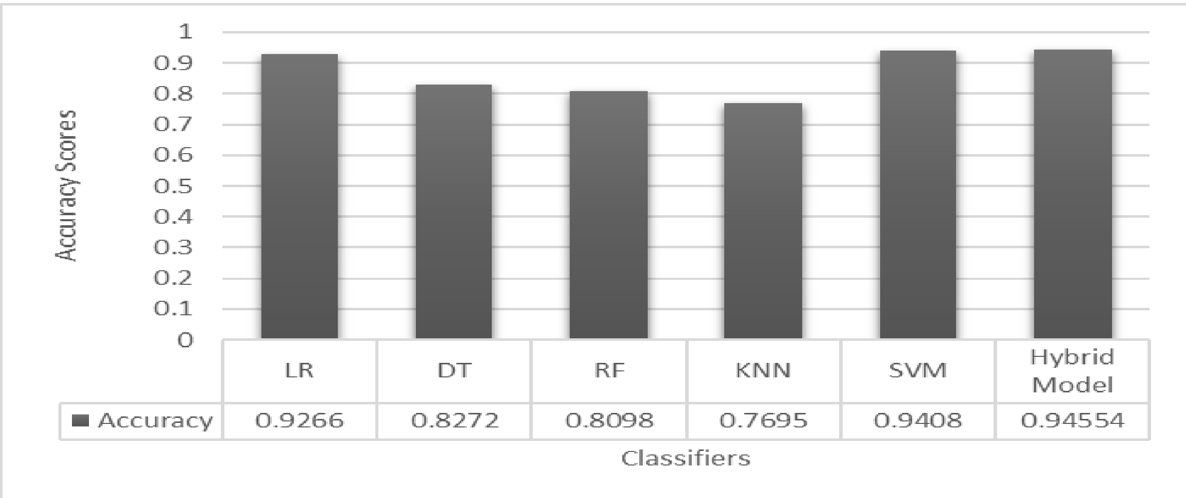

Figure 2. Accuracy vs classifiers

This section provides an overview of the outcomes obtained from the experimental analysis conducted in the research study. The study evaluated various classifiers, including LR, DT, RF, KNN, SVM and a hybrid model, with the aim of assessing the variables of recall, accuracy and precision. The findings regarding accuracy, precision and recall are presented here. Additionally, the performances of the proposed systems for each classifier and the hybrid model are reported. On the test data, the LR, DT, RF, KNN, SVM and hybrid models achieved accuracy scores of 92.66%, 82.72%, 80.98%, 76.95%, 94.08% and 94.5% respectively. Notably, the hybrid model exhibited the highest accuracy score, reaching 94.5%.

Figure 3. Precision vs classifiers

Additionally, precision score experiments were carried out to evaluate the effectiveness of the proposed approach. The precision score measures the proportion of accurate positive predictions out of all positive observations, indicating how often the positive forecasts turn out to be correct. A higher precision score is desirable in this context. The precision scores obtained for LR, DT, RF, KNN, SVM and hybrid model were 90.5%, 81.2%, 76.8%, 91.6% and 94.4% respectively. Figure 3 presents the results of all classifier precision experiments, clearly demonstrating that the hybrid model achieves the most favourable outcomes.

The calculation of the recall score involves determining the number of actual positive predictions in relation to all the actual label classes. Recall, also referred to as sensitivity, represents the percentage of true positive findings. A higher recall score indicates better performance. According to the obtained recall scores, LR, DT, RF, KNN, SVM and hybrid model achieved 94.56%, 83.74%, 85.92%, 95%, 96.35% and 94.56% respectively. Once again, SVM achieved the highest rating in terms of performance (96.35%). The experimental recall scores are presented in Figure 4.

Figure 4. Recall vs classifiers

Providing a comprehensive summary of all the classifiers, Figure 5 displays a chart presenting accuracy, precision, recall and F_1 scores. The chart proves that the hybrid model outperforms other models because accuracy, precision, and F_1 scores are higher than the existing classifiers we used.

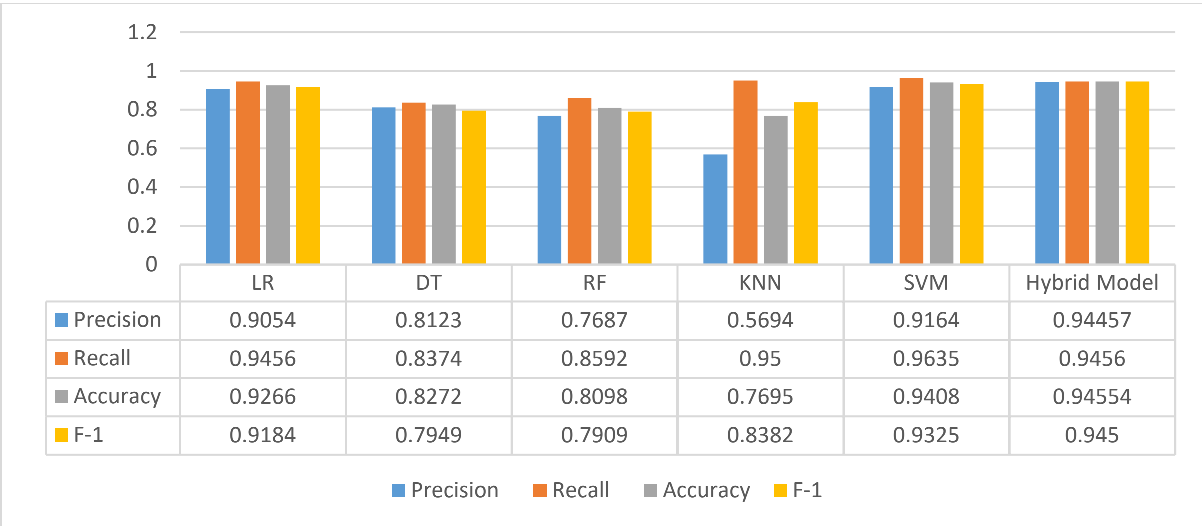

Figure 5. Summary of classifiers and hybrid model with accuracy, precision, recall and F1 scores

Comparison

Table 2 demonstrates that our research paper has been compared to several recently published papers, and that our model's accuracy is superior to the models in those papers.

Table 2. Comparison of models

Best **a**pproach	Accuracy of best approach
Ensemble model (Altheneyan & Alhadlaq, 2023)	93.40%
Hybrid Ensemble Model with Fuzzy Logic (Sharma ***et al***., 2023)	86.8%
Our hybrid model	94.5%

Discussion

It is clear that fake news is a significant problem in today's digital world. As people spend more time on the Internet, they are more likely to come across fake news, which can have serious consequences. Many researchers have tried to develop systems to prevent fake news, and various ML classifiers have been used for this purpose. The literature review also shows that fake news has various effects, including the potential to cause harm to individuals, misguide students, affect financial markets, erode the legitimacy and credibility of traditional news outlets, and hinder people's ability to distinguish between truth and lies.

The objectives of building a fortress against fake news are also clearly stated in the objective section. One of the primary objectives is to develop a reliable and accurate system for detecting fake news based on accuracy, precision and recall. Achieving a higher level of accuracy, precision and recall would ensure that the system can detect most fake news accurately. Another goal is to construct a hybrid model by using CNNs, RNNs and ANNs through the implementation of deep learning techniques. This approach has not been previously attempted and could lead to more effective fake news detection.

Overall, this research highlights the importance of preventing fake news and the need for automated systems to achieve this goal. It also underscores the potential harm that fake news can cause and the various effects it can have. By developing a reliable and accurate system for detecting fake news, we can reduce the impacts of fake news and help people differentiate truth and false news. The use of ML classifiers and deep learning techniques along with a hybrid model, could also lead to more effective and efficient fake news detection. In essence, this study's accuracy is best and this study's accuracy is superior to the recently published papers.

Conclusion

The proliferation of fake news is a major issue with grave repercussions for individuals, communities and the global community. With the increasing use of the Internet and social media, fake news is spreading at an alarming rate, making it difficult for people to differentiate between true and false news. However, based on a comprehensive analysis of published research, it is evident that fake news exerts numerous detrimental effects on individuals. However, a glimmer of optimism arises from the emergence of automated systems, which employ advanced ML and deep learning algorithms to effectively identify and uncover fake news. By evaluating different ML classifiers and making a hybrid model using deep learning, a reliable and accurate system can be developed that would ensure that most fake news can be detected with high levels of accuracy, which is 94.5%. It is imperative that we continue to

http://doi.org/10.18080/jtde.v11n3.765

research and develop systems to prevent the spread of fake news and protect ourselves from the harmful effects it can have on our lives and society.

Acknowledgements

A version of this paper was presented at the third International Conference on Computer, Information Technology and Intelligent Computing, CITIC 2023, held in Malaysia on 26–28 July 2023.

This work was supported by the TM R&D Fund (Project no. RDTC/221045 and SAP ID: MMUE/220017).

References

Ahinkorah, B. O., Ameyaw, E. K., Hagan Jr, J. E., Seidu, A. A., & Schack, T. (2020). Rising above misinformation or fake news in Africa: Another strategy to control COVID-19 spread. *Frontiers in Communication, 5*, 45. https://doi.org/10.3389/fcomm.2020.00045

Ahmad, T., Faisal, M. S., Rizwan, A., Alkanhel, R., Khan, P. W., & Muthanna, A. (2022). Efficient fake news detection mechanism using enhanced deep learning model. *Applied Sciences*, *12*(3), 1743. https://doi.org/10.3390/app12031743

Almenar, E., Aran-Ramspott, S., Suau, J., & Masip, P. (2021). Gender differences in tackling fake news: Different degrees of concern, but same problems. *Media and Communication*, *9*(1), 229–238. https://doi.org/10.17645/mac.v9i1.3523

Altheneyan, A., & Alhadlaq, A. (2023). Big data ML-based fake news detection using distributed learning. *IEEE Access*, *11*, 29447–29463. https://doi.org/10.1109/ACCESS.2023.3260763

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. *Complexity*, *2021,* 1–8. https://doi.org/10.1155/2021/5557784

Bago, B., Rand, D. G., & Pennycook, G. (2020). Fake news, fast and slow: Deliberation reduces belief in false (but not true) news headlines. *Journal of Experimental Psychology: General*, *149*(8), 1608. https://psycnet.apa.org/doi/10.1037/xge0000729

Bakir, V., & McStay, A. (2018). Fake news and the economy of emotions: Problems, causes, solutions. *Digital Journalism*, *6*(2), 154–175. https://doi.org/10.1080/21670811.2017.1345645

Bhatt, G., Sharma, A., Sharma, S., Nagpal, A., Raman, B., & Mittal, A. (2018). Combining Neural, Statistical and External Features for Fake News Stance Identification. WWW '18: Companion Proceedings of The Web Conference 2018. https://doi.org/10.1145/3184558.3191577

Butler, A. (2018). Protecting the Democratic Role of the Press: A Legal Solution to Fake News. *Washington University Law Review*, *96*(2), 419–440.

Creech, B. (2020). Fake news and the discursive construction of technology companies' social power. *Media, Culture & Society*, *42*(6), 952–968. https://doi.org/10.1177/0163443719899801

De Oliveira, D. V. B., & Albuquerque, U. P. (2021). Cultural evolution and digital media: Diffusion of fake news about COVID-19 on Twitter. *SN Computer Science*, *2*(6), 1–12. https://doi.org/10.1007/s42979-021-00836-w

Ghosh, S., & Shah, C. (2018). Towards automatic fake news classification. *Proceedings of the Association for Information Science and Technology*, *55*(1), 805–807. https://doi.org/10.1002/pra2.2018.14505501125

Ho, K. K., Chan, J. Y., & Chiu, D. K. (2022). Fake news and misinformation during the pandemic: What we know and what we do not know. *IT Professional*, *24*(2), 19–24. https://doi.org/10.1109/MITP.2022.3142814

Khan, A., Brohman, K., & Addas, S. (2022). The anatomy of 'fake news': Studying false messages as digital objects. *Journal of Information Technology*, *37*(2), 122–143. https://doi.org/10.1177/02683962211037693

Kokiantonis, A. (2022) *News.csv*. https://www.kaggle.com/datasets/antonioskokiantonis/newscsv/code. Accessed on 26 April 2023.

Lakshmanan, L. V., Simpson, M., & Thirumuruganathan, S. (2019). Combating fake news: a data management and mining perspective. *Proceedings of the VLDB Endowment*, *12*(12), 1990–1993. https://doi.org/10.14778/3352063.3352117

Leeder, C. (2019). How college students evaluate and share "fake news" stories. *Library & Information Science Research*, *41*(3), 100967. https://doi.org/10.1016/j.lisr.2019.100967

Mahmoud, A. (2022) *News.csv*. https://www.kaggle.com/code/ahmedxmahmoud/fake-news-detection/input. Accessed on 2 April 2023.

Meel, P., & Vishwakarma, D. K. (2021). A temporal ensembling based semi-supervised ConvNet for the detection of fake news articles. *Expert Systems with Applications*, *177*, 115002. https://doi.org/10.1016/j.eswa.2021.115002

Naeem, S. B., Bhatti, R., & Khan, A. (2020). An exploration of how fake news is taking over social media and putting public health at risk. *Health Information & Libraries Journal*, *38*(2), 143–149. https://doi.org/10.1111/hir.12320

Pearson, M. (2017). Teaching Media Law in a Post-truth Context: Strategies for Enhancing Learning about the Legal Risks of Fake News and Alternative Facts. *Asia Pacific Media Educator*, *27*(1), 17–26. https://doi.org/10.1177/1326365x17704289

Pérez-Rosas, V., Kleinberg, B., Lefevre, A., & Mihalcea, R. (2017). Automatic detection of fake news. https://doi.org/10.48550/arXiv.1708.07104

Rawat, G., Pandey, T., Singh, T., Yadav, S., & Aggarwal, P. K. (2023, January). Fake news detection using machine learning. 2023 International Conference on Artificial Intelligence and Smart Communication (AISC), pp. 759–762, IEEE. https://doi.org/10.1109/AISC56616.2023.10085488

Sharma, D. K., Singh, B., Agarwal, S., Pachauri, N., Alhussan, A. A., & Abdallah, H. A. (2023). Sarcasm detection over social media platforms using hybrid ensemble model with fuzzy logic. *Electronics*, *12*(4), 937. https://doi.org/10.3390/electronics12040937

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*, pp. 312–320. https://doi.org/10.1145/3289600.3290994

Singhal, S., Shah, R. R., Chakraborty, T., Kumaraguru, P., & Satoh, S. (2019). SpotFake: A Multi-modal Framework for Fake News Detection. IEEE International Conference on Multimedia Big Data (BigMM), pp. 39–47. https://doi.org/10.1109/BigMM.2019.00-44

Singhal, S., Kabra, A., Sharma, M., Shah, R. R., Chakraborty, T., & Kumaraguru, P. (2020). SpotFake+: A Multimodal Framework for Fake News Detection via Transfer Learning (Student Abstract). *Proceedings of the AAAI Conference on Artificial Intelligence*, *34*(10), 13915–13916. https://doi.org/10.1609/aaai.v34i10.7230

Stewart, E. (2021). Detecting Fake News: Two Problems for Content Moderation. *Philosophy & Technology*, *34*, 923–940. https://doi.org/10.1007/s13347-021-00442-x

Sullivan, M. C. (2019). Libraries and fake news: What's the problem? What's the plan?. *Communications in Information Literacy*, *13*(1), 7. https://doi.org/10.15760/comminfolit.2019.13.1.7

Zafarani, R., Zhou, X., Shu, K., & Liu, H. (2019). Fake News Research. Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, pp. 3207–3208. https://doi.org/10.1145/3292500.3332287