International Research Journal of Engineering and Technology (IRJET) e-ISSN: 2395-0056
Volume: 10 Issue: 12 | Dec 2023 www.irjet.net p-ISSN: 2395-0072

A Novel approach to Fake News Detection using Bi-directional LSTM Neural Network Model

Shreyas Anand

B. Tech Student, School of Computer Science and Engineering (SCOPE), Vellore Institute Of Technology, Vellore, Tamil Nadu, India

---***---

Abstract - *In today's digital age, media plays a crucial role in swiftly delivering information about events to the public through the Internet and social networks, becoming an integral part of daily life. The widespread availability of online news caters to diverse age groups, with continuous article publications on various digital platforms.*

However, this surge in online news has given rise to a pressing issue – the intentional dissemination of false information, commonly known as fake news. This poses a significant societal challenge, as misinformation is often crafted for commercial or political motives to mislead readers.

To tackle this challenge, our research introduces an innovative approach to fake news detection. We employ a Bidirectional Long Short-Term Memory (Bi-LSTM) based deep learning model, complemented by sequential attention mechanisms. The Bi-LSTM architecture, adept at analyzing variable-length sequential data, comprehensively examines news articles in both forward and backward directions. The incorporation of sequential attention enhances the model's ability to discern intricate patterns in textual data, aiming to boost accuracy in detecting false news. The proposed model is evaluated using publicly available news article datasets, demonstrating superior accuracy compared to alternative methods like Convolutional Neural Networks (CNN), Recurrent Neural Networks (RNN), and unidirectional LSTMs.

Key Words: Fake News, Deep Learning, BiLSTM, Word Embedding, Neural Networks

1. INTRODUCTION

In the contemporary global landscape, news organizations play a pivotal role in the ubiquitous dissemination of information. The evolution of social media platforms has not only extended the reach of news outlets but has also ushered in a multitude of sources, ranging from established entities to those of questionable reliability. The widespread accessibility of the internet has led to a surge in platform users, enabling the swift sharing of articles by both individuals and news portals. However, this increased accessibility has also given rise to the deliberate propagation of fake news, fueled by motives such as garnering popularity and financial gain.

The repercussions of fake news have garnered substantial research attention. The sheer magnitude of data processed by social media platforms poses a formidable challenge in validating each post, effectively transforming these platforms into breeding grounds for the dissemination of deceptive news [9, 15, 35]. The impacts are far-reaching, encompassing everything from inciting unrest and influencing elections to endangering individuals and fostering misinformation during critical events, such as the ongoing pandemic [4, 16].

Termed as 'rumor,' 'misinformation,' or 'hoax,' fake news spans a spectrum of fraudulent information, from unverified claims to intentional misrepresentation. Despite concerted efforts to combat false information, studies leveraging traditional machine learning techniques have yielded only moderate accuracy. The advent of deep learning has introduced more effective models, with dense neural networks and sophisticated architectures like Recurrent Neural Networks (RNN) and Gated Recurrent Units (GRU) demonstrating superior performance. Hybrid deep learning models, leveraging the strengths of multiple architectures, have emerged as a promising approach for binary classification, effectively distinguishing between genuine and fraudulent news [18, 31].

The research endeavors in this domain aim not only to identify and curb the spread of false information but also to comprehend the dynamics of fake news in the evolving digital landscape. With the continual evolution of deep learning paradigms and the exploration of novel model architectures, the research community is poised to make significant strides in enhancing the accuracy and robustness of fake news detection, contributing to the broader discourse on media reliability and information integrity.

2. RELATED WORK

News classification is a widely studied field due to its potential negative consequences. In this section, we provide an overview of recent work in fraudulent news detection.

Traditional Models

The exploration of fake news detection has been a dynamic field, witnessing substantial research contributions across various dimensions. In a seminal work, Xinyi et al. [36] delved into the root causes of false news dissemination, underscoring the urgency of early identification to curb widespread dissemination. Recognizing the technological underpinnings, several machine learning models have emerged for false news classification.

Evolution of Deep Learning

The evolution of deep learning has played a pivotal role, overcoming initial skepticism by proving its effectiveness over complex mathematical models [17]. Basic deep neural networks have outperformed traditional models such as random forest and SVM in classifying text data [34]. Recurrent Neural Networks (RNNs), adept at processing time series data, have been integral to research exploring the efficiency of deep learning models in fake news detection across diverse media channels [18].

Hybrid Models

Hybrid models have garnered attention, with fused CNN-RNN models and intricate recurrent models like LSTM exhibiting promise [3, 26]. Convolutional Neural Networks (CNNs), primarily designed for image classification, have been stacked atop time series models to enhance feature extraction, improving the model's understanding of word meanings [19, 32].

Gated Recurrent Units (GRUs)

Gated Recurrent Units (GRUs) have proven efficient for text data classification, particularly when combined with attention mechanisms and word embeddings [32]. Bidirectional GRU models, processing data in both forward and reverse directions simultaneously, enhance relations between distant words, reducing variance in training and validation sets.

Long Short-Term Memory (LSTM) Cells

The intricate architecture of Long Short-Term Memory (LSTM) cells has demonstrated excellence in classifying complex time series data, especially when stacked with Convolutional Neural Networks (CNN) [19].

Transformer-Based Models

Transformer-based models, including BERT, XLBERT, and XLNET, have showcased remarkable accuracy, with ensemble models surpassing single-model setups [16].

Data Preprocessing Techniques

Efficient data preprocessing and feeding mechanisms have been pivotal to successful classification tasks [8]. Studies incorporating various n-grams and similarity scoring techniques, such as TF-IDF vectorization with cosine similarity, have emerged as effective methods for calculating word similarity scores.

In essence, the landscape of fake news detection is characterized by the interplay of traditional machine learning, evolving deep learning paradigms, and innovative hybrid models. The success of these models is intricately tied to their ability to adapt to the dynamic nature of textual information across diverse media platforms. The present study contributes to this continuum by exploring the efficacy of Bi-LSTM and sequential attention mechanisms, adding nuanced insights to the evolving field of fake news detection.

3. MATERIALS & METHODS

The hybrid Bidirectional LSTM (BiLSTM) and Sequential Attention model, combined with few other layers are proposed in this research work. We propose a novel hybrid BiLSTM model incorporating GloVe embeddings and Sequential Attention, complemented by a dense layer. The rationale behind opting for a Bidirectional LSTM (Bi-LSTM) over a unidirectional LSTM lies in its simultaneous processing of data in both directions, enabling the model to effectively correlate words in textual data. The incorporation of Sequential Attention enhances the model's ability to understand relationships between different words. Additionally, a fully connected dense layer is introduced above the attention layer, contributing to the complexity of the model. The stacked layers facilitate the model in comprehending the intricacies of complex data, enabling better differentiation between different data classes.

The data preprocessing steps, including lemmatization, tokenization, and the word embedding process, are initially explained. Subsequently, the BiLSTM concept, focusing on extracting targets from the embedded text, is introduced. Finally, the Sequential Attention mechanism and the various layers employed for classifying fake news within the input text are detailed.

3.1. DATASET PRE-PROCESSING, LEMMATIZATION, AND TOKENIZATION

We utilized a dataset named Fake News dataset, accessible through Kaggle, consisting of verified news articles from various news organizations. This dataset employs two distinct labels, 0 and 1, for each news article. Class 0 signifies genuine articles, while class 1 corresponds to fraudulent news articles. The dataset comprises over 10,000 articles for each class and includes essential attributes such as article ID, title, author's name, text, and corresponding labels for all news articles. The articles within the dataset exhibit a wide range of lengths, varying from concise to extensive compositions.

	id	title	author	text	label
0	0	House Dem Aide: We Didn't Even See Comey's Let...	Darrell Lucus	House Dem Aide: We Didn't Even See Comey's Let...	1
1	1	FLYNN: Hillary Clinton, Big Woman on Campus - ...	Daniel J. Flynn	Ever get the feeling your life circles the rou...	0
2	2	Why the Truth Might Get You Fired	Consortiumnews.com	Why the Truth Might Get You Fired October 29, ...	1
3	3	15 Civilians Killed In Single US Airstrike Hav...	Jessica Purkiss	Videos 15 Civilians Killed In Single US Airstr...	1
4	4	Iranian woman jailed for fictional unpublished...	Howard Portnoy	Print \nAn Iranian woman has been sentenced to...	1

Fig. 1: Fake News Dataset

The preprocessing steps for the input text precede its submission to the model. Initially, we eliminate all digits, accent marks, and punctuation. The lemmatizer employed in this study aids in consolidating inflected words, treating them as a singular entity for analysis. Lemmatization, a transparent layer, streamlines customization, including lemma assignments. In the "rule" lemmatization mode, requiring POS (Token. pos) assignment, a Tagger, Morphologiser, or equivalent component that assigns POS must be available and operational before the lemmatizer.

Given the dataset's extensive vocabulary, reducing it becomes imperative. Lemmatization takes precedence over stemming due to stemmers inaccurately categorizing words, resulting in disparate words being treated as one and one word in different forms as distinct entities.

The initial phase of the model involves preprocessing. Extraneous columns are removed, article headings and text are concatenated, and articles with a length below 50 are excluded. Subsequently, all text is converted to lowercase, and any special characters are removed. To equalize the number of articles for each class, articles exceeding a certain threshold are removed. The processed dataset is illustrated in Fig. 2.

For model input, text is tokenized, and padding is applied from the left. If an article contains fewer than 250 words, additional zeros are appended to the text, ensuring uniformity.

0	danish parliament danes should not become mino...	0
1	continuing email flap wont derail clinton but ...	1
2	bomb in istanbul kills 11 near tourist distric...	0
3	democratic house candidates were also targets ...	0
4	russia tests nuclearcapable ballistic missile ...	1

Fig. 2: Dataset after pre-processing

3.2. WORD EMBEDDING

Embedding is a method that transforms high-dimensional vectors into a lower-dimensional space. It incorporates various components such as videos, links, images, gifs, etc., amalgamating them into a cohesive entity, whether as an Instagram post or web media. This embedded content becomes an integral part of the post, enhancing the post's click-through rate and engagement.

Word embeddings are extensively utilized in both machine learning and deep learning models. They contribute to reducing the training time and enhancing the overall performance of a classification model. This layer is typically maintained as static, preventing updates during the gradient descent process [37].

Pre-trained representations come in two types: static and contextual, with significant applicability in natural language processing (NLP). The input is in text format, and the output is in word vector format. The vocabulary is constructed from the trained dataset in text format, enabling an understanding of words in vector format. The primary goal is to identify learned representations that find the closest words for a specified user-given word. For instance, the names of all cricketers will exhibit a smaller distance between them, and similarly, words from similar domains will share a smaller distance [11].

Word2Vec and GloVe [6] currently stand as two of the most widely used word embedding models, effectively transforming words into meaningful vectors. GloVe, a weighted least square model [33], trains using co-occurrence counts of words in input vectors. This unsupervised training model aids in identifying the correlation between two words and their distance in vector space [37], producing vectors known as word embedding vectors.

In the initial phase, we determine the embedding value of available tokens in tokenized text using GloVe embeddings for the fake news prediction model. Subsequently, we construct an embedding matrix, incorporating dropout after adding the embedding layer to reduce model variance.

3.3. TARGET EXTRACTION USING BI-LSTM AND SEQUENTIAL ATTENTION

In this section, we elucidate the extraction of targets from the word embedding tokenized text using BiLSTM and Sequential Attention. LSTM models, being robust recurrent neural network models, excel in processing sequential time series data. They incorporate various sigmoid and tanh activation gates that facilitate effective decision-making on which words to retain in memory and which to discard. Bidirectional LSTM models add an extra layer of LSTMs, with one layer processing input information in the forward direction and the other in the backward direction. This bidirectional flow enhances the model's capability to discern essential data features. Stacking LSTM layers in this manner not only effectively reduces loss but also contributes to improved accuracy.

BiLSTM and Sequential Attention models are seamlessly integrated, connecting serially with one another in succession. The BiLSTM model uses two LSTM which are connected in forward and backward directions. One LSTM works from top to bottom and other works from top to bottom.

During each time t, the forward hidden layer of LSTM with hidden function h_t is obtained based on the previously hidden value h_{t-1} as well as the current input value x_t. In the backward LSTM with hidden value, h_t is obtained based on the future hidden value, h_{t+1}, and the current input value x_t.

Following target extraction, the opinionated sentences undergo processing through the self-attention layer, Dense layer, and several other layers designed for the classification of fake news.

When dealing with extensive news datasets featuring an extensive vocabulary, a standard LSTM model may encounter challenges associating a primary word with a secondary word. For example in the sentence "The animal didn't cross the street because it was too tired", 'refers to 'animal' and it's very simple for an average human to understand it but it's not so simple for the sequential model. When we are developing a sequential model these minor details can help decrease the accuracy. Hence we need the help of sequential attention (or self-attention) here. In the target extraction output, if all words do not have equal contribution then the attention layer is used to extract the words which are significant and aggregate their representation to form a sentence vector.

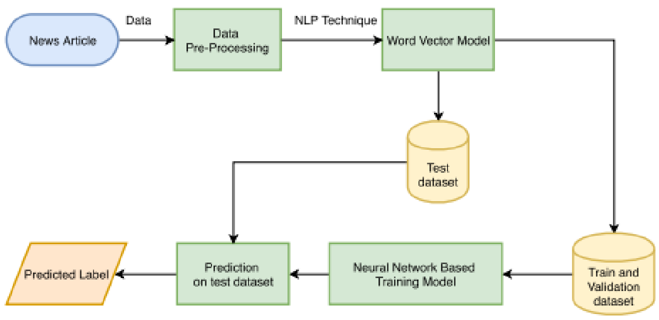

Fig. 3: A General Deep Learning Based Architecture

3.4. DROPOUT LAYER

The Dropout layer, embedded as a regularization technique [22, 24] in our model, operates by selectively deactivating neurons at random intervals during the training process. This regularization strategy introduces an element of randomness, stimulating downstream neurons by temporarily excluding them during the forward pass. It is essential to note that, during this exclusion, no weight updates are applied to the deactivated neurons in the subsequent backward pass. This stochastic deactivation prevents the model from becoming overly reliant on specific neurons, enhancing its overall robustness and preventing overfitting.

In our innovative approach, we take a step further by synergistically integrating the Dropout layer with a sequential attention (or self-attention) model. This additional layer plays a pivotal role in elevating the model's understanding of word meanings, facilitating a more profound comprehension of contextual relationships within the data. The self-attention mechanism allows the model to focus on relevant words and capture intricate dependencies, contributing to a more nuanced interpretation of the input sequence.

To further optimize the model's performance, we introduce both batch normalization and dropout techniques. Batch normalization aids in stabilizing and accelerating the training process by normalizing the inputs at each layer. This contributes to the overall efficiency of the model. The dropout technique, in conjunction with batch normalization, ensures that the model maintains its adaptability and generalizability by preventing overfitting, especially in the presence of complex linguistic patterns.

The integration of these techniques collectively enhances the model's resilience and effectiveness, allowing it to handle intricate linguistic patterns with greater accuracy. This comprehensive approach not only bolsters the model's robustness but also empowers it to excel in the accurate classification and interpretation of textual data, demonstrating its proficiency in the realm of natural language processing.

3.5. SOFTMAX ACTIVATION FUNCTION

The softmax function is employed in the output layer of neural network models to predict a multinomial probability distribution.

Before applying softmax, the neural network produces a set of raw scores called logits. These logits represent the unnormalized outputs from the network, indicating the model's confidence or belief for each class.

The softmax function transforms these logits into probabilities. It achieves this by exponentiating each logit and then normalizing the results. For a given logit z_i the softmax function computes e^{z_i} where e is the mathematical constant (approximately 2.71828).

After exponentiation, the softmax function normalizes the exponentiated logits by dividing each exponentiated value by the sum of all exponentiated values across all classes. This normalization ensures that the resulting vector of probabilities sums to 1, creating a valid probability distribution.

The output of the softmax activation function is a vector containing probabilities for all possible outcomes. Importantly, the sum of these probabilities in the output vector equals one.

3.6. LOSS FUNCTION

The loss function within our model serves as a critical metric in assessing the dissimilarity between the model's predicted values and the actual values. It is imperative to recognize the adaptable nature of loss functions, as they dynamically adjust based on the specific objectives assigned to the model. One prominent example of such a loss function is cross-entropy, which plays a pivotal role in gauging the differentiation between two probability distributions originating from the same set of events.

As our model refines its predictive accuracy, the cross-entropy value progressively diminishes. This characteristic makes cross-entropy particularly well-suited as a loss function, especially in the training of classification models. The underlying principle is notable: in an ideal scenario where the model achieves perfect accuracy, the cross-entropy value tends to approach zero. This intrinsic property highlights the efficacy of cross-entropy in steering the model towards optimal performance by quantifying the extent of deviation between predicted and true values.

The versatility of cross-entropy lies in its ability to capture and quantify these deviations, providing a tangible measure of how well the model aligns with the desired outcomes across various tasks. In essence, cross-entropy serves as a guiding force, helping the model fine-tune its parameters and update its internal representations to enhance its proficiency in tasks such as classification. This nuanced understanding of the loss function's role underscores its significance in shaping the learning process and facilitating the model's continual improvement over successive iterations.

4. RESULTS & DISCUSSIONS

The experiment analysis of the proposed Hybrid BiLSTM and Sequential Attention (or Self-Attention) model is conducted in two phases. Initially, training is carried out using the Fake News dataset from Kaggle, exposing the network to all labels. Subsequently, in the validation/testing phase, the model is evaluated with a test dataset. The IMDB dataset comprises 25,514 unique news article statements, with 20,387 news articles used for training and 5127 news articles for validating the proposed hybrid BiLSTM and Sequential Attention model. The performance comparison of our model with previously developed models, focusing on accuracy scores, demonstrates the superior performance of our developed model over existing models, emphasizing the effectiveness of deep learning models. The experimentation results show that the proposed model achieves an accuracy of 98.65% and a loss of 0.0715.

5. EVALUATION METRICS

In a predictive modeling pipeline, assessing the output of a machine-learning model is a crucial step. While a model might achieve high classification accuracy during construction, its ability to address specific problems in various scenarios needs thorough evaluation. Relying solely on classification accuracy is often insufficient for making such judgments. Therefore, a range of assessment metrics becomes essential for a comprehensive evaluation. Constructing a model might be a relatively straightforward task, but creating a truly promising strategy that meets assessment metrics is a more intricate challenge.

Various evaluation metrics are employed to gauge the efficiency of the model, and the evaluation matrix serves as a vital tool for organizing and structuring this evaluation. The confusion matrix, specifically, offers an overview of the model's performance on the testing dataset by presenting the known true values. It encapsulates key metrics such as true positive, true negative, false positive, and false negative, providing valuable insights into the model's success.

The evaluation of the proposed hybrid BiLSTM and Sequential Attention model's performance is done using the analysis of following key metrics: Accuracy, Recall, Precision, F1 Score, ROC CURVE and AUC. We have used Recall, Precision and F1 Score metrics to provide insights into different aspects of our model's effectiveness in classification.

ACCURACY

The accuracy score, commonly referred to as the classification accuracy rating, is calculated as the percentage of correct predictions relative to the total predictions made by the model. The accuracy (A) can be expressed using the formula

ACCURACY = (TRUE POSITIVES + TRUE NEGATIVES) / TOTAL NUMBER OF PREDICTIONS

ROC CURVE AND AUC

The Receiver Operating Characteristics (ROC) curve illustrates the performance of a classification model across various classification thresholds, utilizing True Positive

Rate (Recall) and False Positive Rate (FPR). The acronym AUC stands for "Area Under the ROC curve," measuring the complete two-dimensional area beneath the entire ROC curve. The False Positive Rate (FPR) can be defined as

FPR = FALSE POSITIVES / (FALSE POSITIVES + TRUE NEGATIVES)

RECALL

Recall, also known as Sensitivity or True Positive Rate, measures the ability of the model to correctly identify all relevant instances from the dataset. It is calculated as the ratio of true positives to the sum of true positives and false negatives.

RECALL = TRUE POSITIVES / (TRUE POSITIVES + FALSE NEGATIVES)

A high Recall indicates that the model is effective at capturing and classifying instances of a particular class.

PRECISION

Precision evaluates the accuracy of positive predictions generated by the model. It is determined by the ratio of true positives to the sum of true positives and false positives.

PRECISION = TRUE POSITIVES / (TRUE POSITIVES + FALSE POSITIVES)

High Precision signifies that when the model predicts a positive instance, it is likely to be correct.

F1 SCORE

The F1 Score is the harmonic mean of Precision and Recall. It offers a balanced measure that takes into account both false positives and false negatives. The F1 Score is particularly useful when there is an imbalance between classes.

F1 = 2 X (PRECISION X RECALL) / (PRECISION X RECALL)

The F1 Score ranges from 0 to 1, with higher values indicating a better balance between Precision and Recall.

By evaluating the model using these metrics, one can gain a comprehensive understanding of its performance, considering factors such as the ability to correctly identify relevant instances (Recall), the precision of positive predictions (Precision), and the overall balance between Precision and Recall (F1 Score).

The precision, Recall, and F1 score obtained using our proposed model for the classes "real news" (class 1) and "fake news" (class 0) showcase the effectiveness of our approach and are shown in table given in Fig. 4. Our model's performance across these metrics emphasizes its capability to accurately classify both genuine and fake news, contributing to the advancement of news classification methodologies.

Category	Precision	Recall	F1 Score
Fake (0)	0.99	0.98	0.99
Real (1)	0.99	0.99	0.98
Micro Average	0.99	0.99	0.99
Weighted Average	0.99	0.99	0.99
Sample Average	0.99	0.99	0.99

Fig. 4: Performance metrics of the proposed method for fake news classification

5. CONCLUSION

In this research, we have presented a novel hybrid model combining BiLSTM and Sequential Attention for the accurate classification of genuine and fake news articles. Leveraging Lemmatization and tokenization for data pre-processing, we employ GloVe and Word2Vec word embeddings to integrate context from both left and right in all layers. The hybrid BiLSTM model, featuring Sequential Attention layers and additional components, is applied to extract targets and discern fake news from authentic sources. Throughout this research, various model combinations and embedding layers are explored, emphasizing the adaptability of natural language processing techniques. Comparative analysis with existing models reveals that our proposed model achieves an impressive validation accuracy of 98.65%, surpassing reported models. Integration of this model into social media or any platform with news articles enhances user reading experiences by mitigating the presence of fake news. This has the potential to reduce reliance on fact-checker websites, and in complex scenarios, both can coexist. Our model facilitates tracking the frequency of fraudulent news articles across publications and authors. While our model is highly effective in delivering genuine news, its primary limitation lies in handling imbalanced datasets.

Future directions include incorporating a more extensive array of hyperparameters to extend the applicability of the model, particularly in scenarios involving multiclass classification. Our work lays a robust foundation for further advancements in the realm of fake news detection, contributing to the ongoing efforts to foster a more reliable and trustworthy information landscape.

REFERENCES

[1] Agarwal A Dixit (2020) Fake news detection: An Ensemble Learning approach. 4th international conference on intelligent computing and control systems, pp. 1178-1183

[2] Ahmed H, Traore I, Saad S (2017) Detection of online fake news using n-gram analysis and machine learning techniques, In: International conference on intelligent, secure, and dependable systems in distributed and cloud environments, Springer, pp. 127–138

[3] Ajao O, Bhowmik D, Zargari S (2018) Fake news identification on twitter with hybrid CNN and RNN models. SM society '18: proceedings of the 9th international conference on social media and society, pp. 226–230 https://doi.org/10.1145/3217804.3217917

[4] Allcott H, Gentzkow M (2017) Social media and fake news in the 2016 election. J Econ Perspect 31(2), pp. 211–236

[5] Al-Zaman MS (2021) COVID-19-related social media fake news in India. J Med 2:100–114

[6] Asparouhov T, Muthen B (2010) Weighted least squares estimation with missing data. Published to Mplus Technical Appendix 2010, pp. 1–10

[7] Berinsky A (2017) Rumours and health care reform: experiments in political misinformation. Br J Polit Sci 47(2):241–262

[8] Bhutani B, Rastogi N, Sehgal P, Purwar A (2019) Fake News Detection Using Sentiment Analysis. 2019 Twelfth International Conference on Contemporary Computing (IC3), pp. 1–5. https://doi.org/10.1109/IC3. 2019.8844880

[9] Castillo C, Mendoza M, Poblete B (2011) Information credibility on twitter. In: Proceedings of the 20th international conference on world wide web, pp. 675–684

[10] Celliers M., Hattingh M. (2020) A Systematic Review on Fake News Themes Reported in Literature. In: Proceedings of the conference on e-Business, e-Services and e-Society I3E 2020: Responsible Design, Implementation and Use of Information and Communication Technology pp 223–234

[11] Cerisara C, Kral P, Lenc L (2018) On the effects of using word2vec representations in neural networks for dialogue act recognition. Comput Speech Lang 47:175–193

[12] Conroy NJ, Rubin VL, Chen Y (2015) Automatic deception detection: methods for finding fake news. Proc Assoc Inf Sci Technol 52(1):1–4

[13] Del Vicario M, Bessi A, Zollo F, Petroni F, Scala A, Caldarelli G, Eugene Stanley H, Quattrociocchi W (2016) The spreading of misinformation online. Proc Natl Acad Sci 113(3):554–559

[14] Girgis S, Amer E, Gadallah M (2018) Deep Learning Algorithms for Detecting Fake News in Online Text. 2018 13th International Conference on Computer Engineering and Systems (ICCES), pp. 93–97, https://doi. org/10.1109/ICCES.2018.8639198

[15] Gorrell G, Kochkina E, Liakata M, Aker A, Zubiaga A, Bontcheva K, Derczynski L (2019) SemEval2019 task 7:" RumourEval, determining rumour veracity and support for rumours." In: Proceedings of the 13th International Workshop on Semantic Evaluation, pp. 845–854

[16] Gundapu S, Mamidi R (21 Jan 2021) Transformer based Automatic COVID-19 Fake News Detection System " arXiv:2101.00180v3 [cs.CL]

[17] Hiramath CK, Deshpande GC (2019) Fake News Detection Using Deep Learning Techniques. 2019 1st International Conference on Advances in Information Technology (ICAIT), pp. 411–415. https://doi.org/10. 1109/ICAIT47043.2019.8987258

[18] Islam MR, Liu S, Wang X, Xu G (2020) Deep learning for misinformation detection on online social networks: a survey and new perspectives. Soc Netw Anal Min 10:Article number: 82

[19] Kaliyar RK (2018) Fake News Detection Using A Deep Neural Network," 2018 4th International Conference on Computing Communication and Automation, pp. 1–7. https://doi.org/10.1109/ CCAA.2018.8777343

[20] Kulkarni P, Karwande S, Keskar R, Kale P, Iyer S (2021) Fake news detection using machine learning. ITM Web Conf 40:030003. https://doi.org/10.1051/itmconf/20214003003

[21] Kumar S, Shah N (23 Apr 2018) False Information on Web and Social Media: A Survey. arXiv: 1804.08559v1 [cs.SI]

[22] Kyriakides G, Margaritis K (2019) An Introduction to Neural Architecture Search for Convolutional Networks. issued December 31 arXiv:2005.11074

[23] Lazer D, Baum M, Benkler Y, Berinsky A, Greenhill K, Menczer F, Metzger M, Nyhan B, Pennycook G, Rothschild D, Schudson M, Sloman S, Sunstein C,

Thorson E, Watts D, Zittrain J (2018) The science of fake news. Science." Published to Science 09 Mar 2018 359(6380):1094–1096

[24] Li Y, Yang Y (2017) Convergence analysis of two-layer neural networks with relu activation. 31st conference on neural information processing systems (NIPS 2017), Long Beach, CA, USA, pp. 597–607

[25] Mo Jang S, Kim JK (2018) Third person effects of fake news: Fake news regulation and media literacy interventions. Comput Hum Behav 80:295–302. https://doi.org/10.1016/j.chb.2017.11.034

[26] Nasir JA, Khan OS, Varlamis I (2021) Fake news detection: A hybrid CNN-RNN based deep learning approach. Int J Inf Manag Data Insights 1(1):100007, ISSN 2667–0968. https://doi.org/10.1016/j.jjimei.2020.100007

[27] Nida A. Irfan UK, Farah Salem A, Lama Abdulaziz A, Asma Khaled A (2021) Fake detect: A deep learning ensemble model for fake news detection, Vol 2021, 5557784, https://doi.org/10.1155/2021/5557784

[28] Pennycook G, Rand DG (2020) Who falls for fake news? The roles of bullshit receptivity, overclaiming, familiarity, and analytic thinking. J Pers 88:185–200. https://doi.org/10.1111/jopy.12476

[29] Qawasmeh E, Tawalbeh M, Abdullah M (2019) Automatic Identification of Fake News Using Deep Learning. 2019 Sixth international conference on social networks analysis, Management and Security (SNAMS), pp. 383–388, https://doi.org/10.1109/SNAMS.2019.8931873

[30] Qi Y, Sachan D, Felix M, Padmanabhan S, Neubig G (2018) When and why are pre-trained word embeddings useful for neural machine translation. In: proceedings of the 2018 conference of the north American chapter of the association for computational linguistics: human language technologies, vol 2 (short papers), pp 529–535

[31] Reddy P, Roy D, Manoj P, Keerthana M, Tijare P (2019) A Study on Fake News Detection Using Naïve Bayes, SVM, Neural Networks and LSTM. J Adv Res Dyn Control Syst 1:942–947

[32] Saikh T, De A, Ekbal A, Bhattacharyya (2019) A deep learning approach for automatic detection of fake news. In: Proceedings of the 16th International Conference on Natural Language Processing, International Institute of Information Technology, Hyderabad, pp. 230–238

[33] Seide F, Li G, Chen X, Yu D (2011) Feature engineering in context-dependent deep neural networks for conversational speech transcription. 2011 IEEE Workshop on Automatic Speech Recognition & Understanding, 24–29, https://doi.org/10.1109/ASRU.2011.6163899

[34] Singh LK, Pooja GH, Khanna M, Bhadoria RS (2021) An enhanced deep image model for glaucoma diagnosis using feature-based detection in retinal fundus. Med Biol Eng Comput 59(2):333–353

[35] Vosoughi S, Neo Mohsenvand M, Roy D (2017) Rumour gauge: predicting the veracity of rumours on twitter. ACM Trans Knowl Discov Data 11(4):1–36

[36] Zhou X, Zafarani R (2018) A Survey of Fake News: Fundamental Theories, Detection Methods, and Opportunities. ACM Comput Surv 53(5) October 2020 Article No.: 109 pp 1–40. https://doi.org/10.1145/3395046

[37] Young T, Hazarika D, Poria S, Cambria E (Aug. 2018) Recent trends in deep learning based natural language processing [review article]. IEEE Comput Intell Mag 13(3):55–75. https://doi.org/10.1109/MCI.2018.2840738

BIOGRAPHIES

Shreyas Anand is currently pursuing his Bachelors in Technology from the School of Computer Science & Engineering, Vellore Institute of Technology, Vellore. His main areas of research include Image Processing, Machine Learning, Data Mining and Visualization.