CUET-NLP_MP@DravidianLangTech 2025: A Transformer and LLM-Based Ensemble Approach for Fake News Detection in Dravidian Languages

Md Minhazul Kabir, Md. Mohiuddin
Kawsar Ahmed and **Mohammed Moshiul Hoque**
Department of Computer Science and Engineering
Chittagong University of Engineering & Technology, Chattogram-4349, Bangladesh
{u1904040, u1904103, u1804017}@student.cuet.ac.bd
moshiul_240@cuet.ac.bd

Abstract

Fake news detection is a critical problem in today's digital age, aiming to classify intentionally misleading or fabricated news content. In this study, we present a transformer- and LLM-based ensemble method to address the challenges in fake news detection. We explored various machine learning (ML), deep learning (DL), transformer, and LLM-based approaches on a Malayalam fake news detection dataset. Our findings highlight the difficulties faced by traditional ML and DL methods in accurately detecting fake news, while transformer- and LLM-based ensemble methods demonstrate significant improvements in performance. The ensemble method combining Sarvam-1, Malayalam-BERT, and XLM-R outperformed all other approaches, achieving an F1-score of 89.30% on the given dataset. This accomplishment, which contributed to securing 2^{nd} place in the shared task at DravidianLangTech 2025, underscores the importance of developing effective methods for detecting fake news in Dravidian languages.

1 Introduction

The term fake news describes false or misleading material presented as fact, often shared to deceive and manipulate public opinion. With the rise of social media platforms, the fabrication and rapid spread of fake news have reached unprecedented levels due to the ease of sharing information and the lack of stringent verification processes (Shahi et al., 2021). This phenomenon has significant consequences, including societal division, misinformation, and loss of trust in reliable sources. Fake news can cause serious social, political, and economic issues, including swaying public opinion, upsetting democratic processes, and making it more difficult to handle crises during important occasions like elections or medical crises (Kaliyar et al., 2021). Addressing this issue requires robust approaches to differentiate facts from false narratives, especially in linguistically diverse regions.

Despite significant advancements in detecting fake news across high-resource languages like English, Spanish, and Arabic (Zhou et al., 2023), low-resource languages like Malayalam remain underexplored due to limited datasets and linguistic resources (Thara and Poornachandran, 2022). Malayalam presents unique challenges due to its rich linguistic features, including dialects and idiomatic expressions, making it difficult to process and analyze (Coelho et al., 2023). The rise of code-mixed text, where users mix multiple languages and scripts, further complicates traditional monolingual fake news detection systems (Hegde et al., 2022). This shared task focuses on creating effective approaches for identifying and classifying fake news in Malayalam, emphasizing low-resource and code-mixed scenarios. As a participant in this shared task, our contribution can be summarized as follows:

- Proposed an ensemble approach leveraging two transformer-based models (Malayalam-BERT and XLM-R) and an LLM (sarvam-1) for effective fake news detection in Malayalam.
- Conducted a comprehensive evaluation of various ML models (LR, MNB, SVM, XGBoost), DL models (CNN, LSTM, CNN+BiLSTM), Transformer-based models (Malayalam-BERT, XLM-R, mBERT, DistilBERT), and LLM models (Gemma-2-2b, Llama-3.2-3B, ProjectIndus, sarvam-1) to identify an optimal approach for fake news detection in Malayalam.

2 Related Work

In recent years, a great deal of study has been prompted by the growing ubiquity of fake news across platforms and languages. Coelho et al.

Proceedings of the Fifth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages, pages 498–504
May 3, 2025

(2023) investigated machine learning models for Dravidian language fake news identification where they trained an ensemble model combining MNB, LR and SVM using TF-IDF of word unigrams that achieved a macro F1-score of 0.83 and third place in the task at DravidianLangTech@RANLP 2023. By utilizing the XLM-RoBERTa base model, renowned for multilingual capabilities, Raja et al. (2023) presented an innovative method that achieved a remarkable macro-averaged F1-score of 87%. Meanwhile, Sujan et al. (2023) employed a multimodal approach by concatenating features from LSTM networks for textual data and VGG16 for image data, achieving a macro F1-score of 0.67. Similarly, Farsi et al. (2024) presented a fine-tuned MuRIL model leveraging parameter tuning that achieved F1-scores of 0.86 and 0.5191 in tasks 1 and 2, respectively, securing 3rd place in task 1 and 1st place in task 2 in DravidianLangTech shared task. In a different research, Devika et al. (2024) curated a dataset specifically for Malayalam fake news detection, achieving an F1-score of 0.3393 with LR trained on LaBSE features while emphasizing the need to address data imbalance. Osama et al. (2024) also contributed to the DravidianLangTech shared task. Their experiments with ML, DL, and transformer-based models revealed that m-BERT achieved the highest macro F1 score of 0.85, securing 4th place in the shared task. Furthermore, Shohan et al. (2024) proposed an intelligent text checkworthiness technique, achieving F1-scores of 75.82% in English (RoBERTa), 52.55% in Arabic (Dehate-BERT), and 58.42% in Dutch (Dutch-BERT). As an instance of notable progress in fake news detection, Kaliyar et al. (2021) presented FakeBERT, which combines BERT with single-layer CNNs, that detected bogus news in English with an impressive 98.90% accuracy. In the context of Arabic fake news detection, Othman et al. (2024) introduced a hybrid model combining 2D-CNN and AraBERT, with F1-scores of 0.6188, 0.7837, and 0.8009 on the ANS, Ara-News, and Covid19Fakes datasets, respectively. Considering the current improvements, this study applies an ensemble method to identify fake news in Dravidian languages.

3 Dataset and Task Description

Task 1 of the shared task competition on "Fake News Detection in Dravidian Languages" (Subramanian et al., 2025, 2024, 2023) focuses on classifying social media posts, specifically YouTube comments into two classes: fake and original in the Malayalam language. The objective is to identify whether a given text is containing misleading or authentic information. The provided dataset (Devika et al., 2024) for this task includes multilingual and codemixed Malayalam data. The given dataset is divided into three sets: train, dev, and test. Each set of data has a nearly equal distribution of fake and original content. Some additional information about the dataset are provided in Table 1.

Set	Class	SC	TW	UW	Avg. Len
Train	original	1658	37229	17472	11
	Fake	1599			
Dev	original	409	8760	5492	7
	Fake	406			
Test	original	512	11266	6587	7
	Fake	507			

Table 1: Dataset distributions, with acronyms SC, TW, UW, and Avg. Len representing sample count, total words, unique words, and average length, respectively.

4 System Overview

This section outlines the methodology for the fake news detection task, which combines traditional machine learning, deep learning, transformer models, and large language models. The diagram of the proposed approach is illustrated in Figure 1. Detailed implementation and source code for this system are accessible on GitHub[1].

In the preprocessing phase, we applied essential text-cleaning techniques such as emoji removal, HTML tag removal, duplicate sample removal, etc. For feature extraction, we used different methods tailored to the type of model being employed. TF-IDF (Takenobu, 1994) and the CountVectorizer method was used extensively for ML models to represent the importance of terms based on their frequency across the documents. For DL models, we employed GloVe embeddings (Pennington et al., 2014), which capture the semantic relationships and contextual meaning of words in the text.

4.1 Machine Learning Approaches

Several machine learning models such as SVM, XGBoost, MNB, LR, and ensemble methods were explored for the task. LR was used with the regularization parameter set to 0.1 and a maximum of 50000 iterations. SVM was implemented using a linear kernel with a regularization parameter

[1] https://github.com/R1FA7/Fake-News-Detection-Malayalam

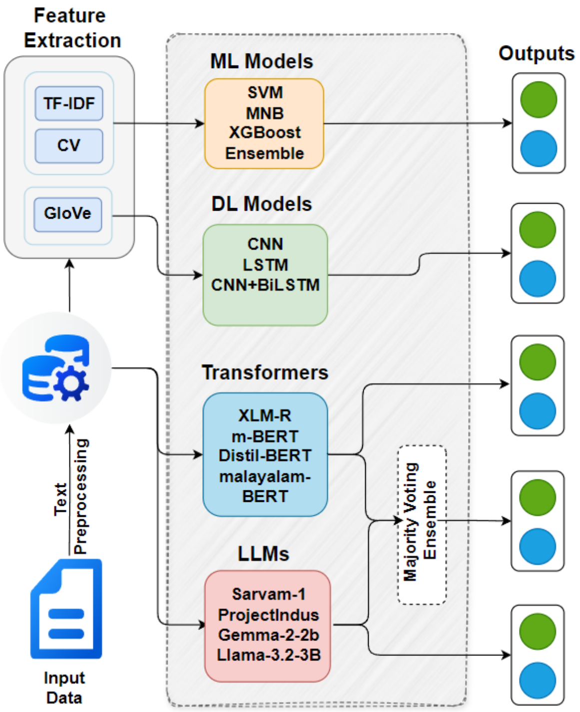

Figure 1: Schematic process for fake news detection

C set to 0.8 and a gamma value of 1. MNB was tested with an alpha value of 1.0, which controls the smoothing of probabilities. Also, XGBoost was implemented with eval metric mlogloss. Finally, an ensemble model combining those was employed, using a majority voting scheme for classification.

4.2 Deep Learning Approaches

For deep learning, we explored CNN, LSTM, BiLSTM architectures to capture the intricate patterns in the text data. A simple CNN was used, consisting of one convolutional layer with 128 filters, followed by GlobalMaxPooling for dimensionality reduction and a dense layer for classification, optimized with Adam and binary cross-entropy loss. LSTM networks were also used to capture sequential dependencies in the text, with a batch size of 64 and a single LSTM layer comprising 200 units. A hybrid CNN and LSTM model was explored combining CNN's feature extraction and BiLSTM's (Schuster and Paliwal, 1997) bidirectional context to improve performance.

4.3 Transformer Based Approaches

We leveraged several pre-trained transformer-based models to address the fake news classification task from Hugging Face (Wolf, 2020). These models included m-BERT (Devlin et al., 2019), XLM-R (Conneau et al., 2020), Distil-BERT (Sanh et al., 2020), and Malayalam-BERT (Joshi, 2023) model. As part of finetune, we concatenated the train and dev set data first. This work randomly took 15% for validation and the rest for training to maximize the data available for training the data-hungry transformer models. Reserving 15% data for validation allowed for hyperparameter tuning and model evaluation without significantly reducing the training data. The text data was tokenized and padded to a maximum sequence length of 90 tokens. The models were trained using a learning rate of $3e^{-5}$, a batch size of 16, and a total of 10 epochs. The performance was optimized based on the macro F1 score.

4.4 LLM Based Approaches

We also explored state-of-the-art Large Language Models (LLMs) such as Gemma-2-2b (Team, 2024), ProjectIndus (Malhotra et al., 2024), sarvam-1 [2] and Llama-3.2-3B (Liu et al., 2024) to further enhance the performance. To fine-tune these models we combined the train and dev sets, randomly kept 20% for validation, and used the rest for training, along with various optimization strategies. First, we leveraged 4-bit quantization to reduce the model's memory footprint without affecting its performance. For model adaptation, we employed Parameter Efficient Fine-Tuning (PEFT), particularly focusing on LoRA (Low-Rank Adaptation). LoRA helps fine-tune large models with fewer parameters, making the process more efficient by updating only key attention layers and freezing other layers. This approach allows the model to adapt to the task without extensive retraining. The LoRA setup used a rank of 4, with a scaling factor (alpha) of 16 and a dropout of 0.15 ensuring that the model's performance improved without overfitting. For the training process, we set the learning rate to $1e^{-4}$ utilizing the AdamW optimizer with a weight decay of 0.01 and batch size 16 in 5 epochs. Additionally, a learning rate scheduler with a reduction factor of 0.5 and patience of 2 epochs was applied to maintain steady improvement.

4.5 Ensemble Approaches

In the ensemble approaches, we employed a majority voting scheme to combine the predictions of multiple models in order to enhance the overall performance. The majority voting technique works by aggregating the predictions from different models

[2] https://huggingface.co/sarvamai/sarvam-1

and selecting the class that appears most frequently as the final decision. This approach benefits from the strengths of different models, increasing robustness and accuracy. For our ensemble model, we experimented with different combinations of models from various categories, such as Transformer-based models and Large Language Models. Specifically, we combined models like Malayalam-BERT, XLM-R, mBERT and Distil-BERT from the transformer category with sarvam-1 from the LLM category. These models were selected based on their strong individual performance, ensuring that each contributed unique strengths to the ensemble.

5 Results and Analysis

The results in Table 2 reveal notable differences in the performance of various ML, DL, Transformer, and LLM models. In this study, we used the G-mean score instead of precision and recall to ensure balanced performance across both classes, avoiding bias even in balanced datasets.

Among the ML models, the Ensemble methods with countVectorizer features achieved the highest macro F1-score of 77.11%, outperforming other ML models. In DL, CNN with GloVe embeddings performed the best, achieving an F1-score of 61.00%, followed by CNN+BiLSTM at 61.00%. However, these scores were significantly lower than the top-performing ML models. Among the transformer-based models, Malayalam-BERT outperformed all other models with an F1-score of 88.32%, surpassing XLM-R (86.36%) and mBERT (85.27%) by a notable margin. Distil-BERT scored an F1-score of 84.00%. In the LLMs category, sarvam-1 achieved the highest F1-score of 83.90%, outperforming Google's Gemma-2-2B (82.63%) and Meta's Llama-3.2-3B (83.27%) by a small margin. Finally, our proposed model, an ensemble of sarvam-1, Malayalam-BERT, and XLM-R, achieved the highest performance with an F1-score of 89.30%, which is 1.2% higher than the next best-performing transformer (Malayalam-BERT), and 12.26% higher than the best ML model (Ensemble+CV).

ML and DL models delivered lower scores compared to Transformer and LLM models. Specifically, The score of DL was significantly lower than ML which can be due to their struggle to generalize well on smaller or limited datasets, leading to overfitting on the training data. Malayalam-BERT is a BERT model trained on publicly available Malayalam monolingual datasets, which excelled due to its specialization in the Malayalam language. LLM didn't perform like the transformer-based models possibly due to their less specialized nature for this specific task or less fine-tuning due to resource limitations. However, in our proposed method, We chose the top two transformer models (Malayalam-BERT and XLM-RoBERTa) and the best LLM model (sarvam-1) for a majority voting ensemble, as they showed the best individual performance. This combination excels by combining the complementary strengths of two transformers and one LLM model, delivering superior performance compared to other model ensembles.

ML Models			
Classifier	**G-mean(%)**	**F1(%)**	**Ac(%)**
SVM+TF-IDF	74.94	75.04	75.00
XGBoost+CV	74.50	73.25	74.00
MNB+CV	75.34	75.62	76.00
LR+CV	76.13	76.31	76.00
Ensemble+ CV	**77.46**	**77.11**	**77.00**
DL Models			
Classifier	**G-mean(%)**	**F1(%)**	**Ac(%)**
LSTM(GloVe)	56.92	60.00	63.00
CNN+ BiLSTM(GloVe)	57.71	61.00	64.00
CNN(GloVe)	**58.15**	**61.00**	**64.00**
Transformers			
Classifier	**G-mean(%)**	**F1(%)**	**Ac(%)**
XLM-R	86.49	86.36	86.00
mBERT	84.50	85.27	85.00
Distil-BERT	83.64	83.61	83.63
Malayalam-BERT	**88.50**	**88.32**	**88.00**
LLMs			
Classifier	**G-mean(%)**	**F1(%)**	**Ac(%)**
Gemma-2-2b	82.45	82.63	82.51
Llama-3.2-3B	84.00	83.27	84.00
ProjectIndus	59.97	59.86	60.00
sarvam-1	**83.98**	**83.90**	**84.00**
Ensemble			
Classifier	**G-mean(%)**	**F1(%)**	**Ac(%)**
(mBERT + XLM-R + Malayalam-BERT)	88.12	88.10	88.10
(Distil-BERT +XLM-R + Malayalam-BERT)	88.15	88.13	88.15
(sarvam-1 + XLM-R + Malayalam-BERT) (Proposed)	**89.48**	**89.30**	**89.40**

Table 2: Performance of the different methods on the test set

5.1 Error Analysis

A comprehensive quantitative and qualitative error analysis is conducted to provide detailed insights into the proposed model's performance.

Quantitative Analysis

Figure 2 presents a confusion matrix that classifies text in the test set as either fake or original. The figure indicates that out of 1019 test samples, 910 were correctly identified and 109 were incorrect predictions. Figure 3 shows that the model was trained on Fake sentences with an average length of 14.51, while Original sentences had an average

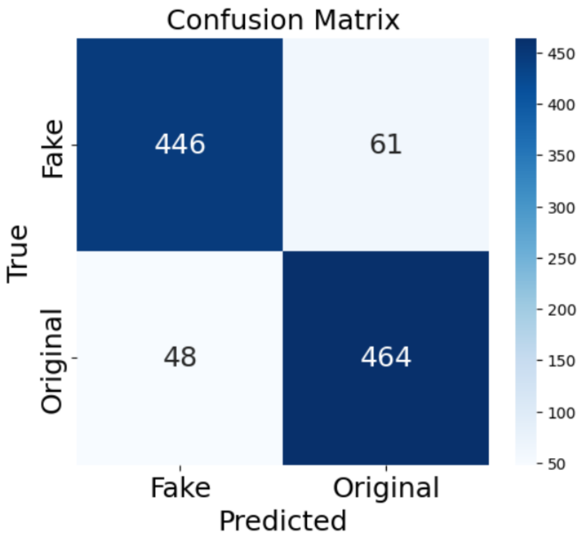

Figure 2: Confusion matrix of Ensemble model

length of 8.46. However, when the model encounters Fake sentences in the test set with an average length of 7.77, it tends to misclassify them. This discrepancy in sentence length suggests that the model may be struggling with shorter Fake sentences, leading to ambiguity and incorrect predictions.

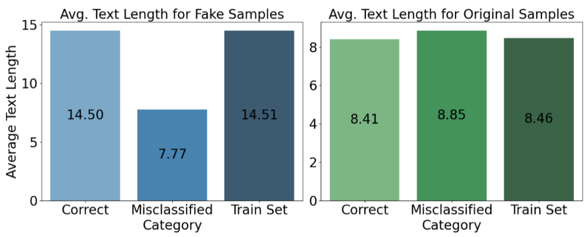

Figure 3: Classified and misclassified average length of classes

Qualitative Analysis

Figure 4 presents some predicted outputs of the developed model. In the first and third texts, the model successfully predicted the class of the text. On the other hand, it failed to do so in the second and fourth texts. The proposed model, which is fine-tuned in this work, is primarily trained on long length fake data compared to fake data in test data. This could be one of the reasons for this model's failure in some samples.

6 Conclusion

This study investigated the performance of several LLM, transformer-based, DL, and ML models on the Malayalam fake news detection dataset. The results demonstrate that the the ensemble model combining sarvamai/sarvam-1, Malayalam-BERT, and XLM-RoBERTa reached the highest F1-score of 89.30%. This finding shows the effectiveness of the transformer and LLM-based ensemble approach, in successfully tackling the challenge of fake news detection. Fine-tuning on shorter fake news samples, and exploring advanced preprocessing methods for code-mixed data could further boost the performance in the future.

Text	Actual	Predicted
വിഷമം വരമ്പോൾ ഞാൻ ഓടി വരും (I will run when trouble comes)	Original	Original
നശിപ്പിച്ചു കളയണം ഈ കമ്മ്യൂണിസ്റ്റ് രാജ്യത്തെ (This communist country must be destroyed)	Fake	Original
Corona ye jesus oddikkum (Corona is Jesus' Holy Spirit)	Fake	Fake
വാക്സിൻ എന്തിനാണ്? (Why the vaccine?)	Original	Fake

Figure 4: Sample predictions made by the proposed Ensemble model with actual and predicted label.

Limitations

Despite the effectiveness of our ensemble method, several limitations remain. One major concern is the reliance on transformer-based models and large language models, which may not generalize well to languages beyond Malayalam. Additionally, the model's performance was affected by the limited training data, particularly with shorter fake news samples, leading to misclassifications. Future work can address these challenges by expanding the training dataset with more diverse and representative samples, improving fine-tuning strategies, and exploring cross-lingual transfer learning to enhance generalization across different languages. Integrating more advanced data augmentation techniques and leveraging multimodal approaches may also contribute to improved robustness and accuracy in fake news detection.

Acknowledgments

We thank the DravidianLangTech 2025 shared task organizers for running this task. This work was supported by the Directorate of Research & Extension (DRE), Chittagong University of Engineering & Technology (CUET).

References

Sharal Coelho, Asha Hegde, G Kavya, and Hosahalli Lakshmaiah Shashirekha. 2023. Mucs@ dravidianlangtech2023: Malayalam fake news detection using machine learning approach. In *Proceedings of* *the Third Workshop on Speech and Language Technologies for Dravidian Languages*, pages 288–292.

Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.

K Devika, B Haripriya, E Vigneshwar, B Premjith, Bharathi Raja Chakravarthi, et al. 2024. From dataset to detection: A comprehensive approach to combating malayalam fake news. In *Proceedings of the Fourth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*, pages 16–23.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. *Preprint*, arXiv:1810.04805.

Salman Farsi, Asrarul Eusha, Ariful Islam, Hasan Mesbaul Ali Taher, Jawad Hossain, Shawly Ahsan, Avishek Das, and Mohammed Moshiul Hoque. 2024. Cuet_binary_hackers@ dravidianlangtech eacl2024: Fake news detection in malayalam language leveraging fine-tuned muril bert. In *Proceedings of the Fourth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*, pages 173–179.

Asha Hegde, Shubhanker Banerjee, Bharathi Raja Chakravarthi, Ruba Priyadharshini, Hosahalli Shashirekha, John Philip McCrae, et al. 2022. Overview of the shared task on machine translation in dravidian languages. In *Proceedings of the second workshop on speech and language technologies for Dravidian languages*, pages 271–278.

Raviraj Joshi. 2023. L3cube-hindbert and devbert: Pre-trained bert transformer models for devanagari based hindi and marathi languages. *Preprint*, arXiv:2211.11418.

Rohit Kumar Kaliyar, Anurag Goswami, and Pratik Narang. 2021. Fakebert: Fake news detection in social media with a bert-based deep learning approach. *Multimedia tools and applications*, 80(8):11765–11788.

Zechun Liu, Changsheng Zhao, Igor Fedorov, Bilge Soran, Dhruv Choudhary, Raghuraman Krishnamoorthi, Vikas Chandra, Yuandong Tian, and Tijmen Blankevoort. 2024. Spinquant: Llm quantization with learned rotations. *Preprint*, arXiv:2405.16406.

Nikhil Malhotra, Nilesh Brahme, Satish Mishra, and Vinay Sharma. 2024. Project indus: A foundational model for indian languages. *Tech Mahindra Makers Lab*.

Md Osama, Kawsar Ahmed, Hasan Mesbaul Ali Taher, Jawad Hossain, Shawly Ahsan, and Mohammed Moshiul Hoque. 2024. Cuet_nlp_goodfellows@ dravidianlangtech eacl2024: A transformer-based approach for detecting fake news in dravidian languages. In *Proceedings of the Fourth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*, pages 187–192.

Nermin Abdelhakim Othman, Doaa S Elzanfaly, and Mostafa Mahmoud M Elhawary. 2024. Arabic fake news detection using deep learning. *IEEE Access*.

Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. volume 14, pages 1532–1543.

Eduri Raja, Badal Soni, and Sami Kumar Borgohain. 2023. nlpt malayalm@ dravidianlangtech: Fake news detection in malayalam using optimized xlm-roberta model. In *Proceedings of the Third Workshop on Speech and Language Technologies for Dravidian Languages*, pages 186–191.

Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2020. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *Preprint*, arXiv:1910.01108.

Mike Schuster and Kuldip Paliwal. 1997. Bidirectional recurrent neural networks. *Signal Processing, IEEE Transactions on*, 45:2673 – 2681.

Gautam Kishore Shahi, Anne Dirkson, and Tim A Majchrzak. 2021. An exploratory study of covid-19 misinformation on twitter. *Online social networks and media*, 22:100104.

Symom Hossain Shohan, Md Sajjad Hossain, Ashraful Islam Paran, Jawad Hossain, Shawly Ahsan, and Mohammed Moshiul Hoque. 2024. Semanticcuetsync at checkthat! 2024: Pre-trained transformer-based approach to detect check-worthy tweets.

Malliga Subramanian, , B Premjith, Kogilavani Shanmugavadivel, Santhia Pandiyan, Balasubramanian Palani, and Bharathi Raja Chakravarthi. 2025. Overview of the Shared Task on Fake News Detection in Dravidian Languages: DravidianLangTech@NAACL 2025. In *Proceedings of the Fifth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.

Malliga Subramanian, Bharathi Raja Chakravarthi, Kogilavani Shanmugavadivel, Santhiya Pandiyan, Prasanna Kumar Kumaresan, Balasubramanian Palani, B Premjith, K Vanaja, S Mithunja, K Devika, et al. 2024. Overview of the second shared task on fake news detection in dravidian languages: Dravidianlangtech@ eacl 2024. In *Proceedings of the Fourth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*, pages 71–78.

Malliga Subramanian, Bharathi Raja Chakravarthi, Kogilavani Shanmugavadivel, Santhiya Pandiyan, Prasanna Kumar Kumaresan, Balasubramanian Palani, Muskaan Singh, Sandhiya Raja, Vanaja, and Mithunajha S. 2023. Overview of the shared task on fake news detection from social media text. In *Proceedings of the Third Workshop on Speech and Language Technologies for Dravidian Languages*, Varna, Bulgaria. Recent Advances in Natural Language Processing.

Adhish S Sujan, Aleena Benny, VS Anoop, et al. 2023. Malfake: A multimodal fake news identification for malayalam using recurrent neural networks and vgg-16. *arXiv preprint arXiv:2310.18263*.

Tokunaga Takenobu. 1994. Text categorization based on weighted inverse document frequency. *Information Processing Society of Japan, SIGNL*, 94(100):33–40.

Gemma Team. 2024. Gemma.

S Thara and Prabaharan Poornachandran. 2022. Social media text analytics of malayalam–english code-mixed using deep learning. *Journal of big Data*, 9(1):45.

Thomas Wolf. 2020. Transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.

Lina Zhou, Jie Tao, and Dongsong Zhang. 2023. Does fake news in different languages tell the same story? an analysis of multi-level thematic and emotional characteristics of news about covid-19. *Information Systems Frontiers*, 25(2):493–512.