Overview of the Second Shared Task on Fake News Detection in Dravidian Languages: DravidianLangTech@EACL 2024

Malliga Subramanian[1], Bharathi Raja Chakravarthi[2], Kogilavani Shanmugavadivel[1], Santhiya Pandiyan[1], Prasanna Kumar Kumaresan[3], Balasubramanian Palani[4], Premjith B[5], Vanaja[1], Mithunajha S[1], Devika K[5] Hariprasath S.B[5], Haripriya B[5], Vigneshwar E[5]

[1] Kongu Engineering College, Tamil Nadu, India.
[2] School of Computer Science, University of Galway, Ireland.
[3] Data Science Institute, University of Galway, Ireland.
[4] National Institute of Technology, Tamil Nadu, India.
[5] Amrita School of Artificial Intelligence, Amrita Vishwa Vidyapeetham, Coimbatore, Amrita Vishwa Vidyapeetham, India.

Abstract

The rise of online social media has revolutionized communication, offering users a convenient way to share information and stay updated on current events. However, this surge in connectivity has also led to the proliferation of misinformation, commonly known as fake news. This misleading content, often disguised as legitimate news, poses a significant challenge as it can distort public perception and erode trust in reliable sources. This shared task consists of two subtasks such as task 1 and task 2. Task 1 aims to classify a given social media text into original or fake. The goal of the FakeDetect-Malayalam task2 is to encourage participants to develop effective models capable of accurately detecting and classifying fake news articles in the Malayalam language into different categories like False, Half True, Mostly False, Partly False, and Mostly True. For this shared task, 33 participants submitted their results.

1 Introduction

The Second Shared Task on Fake News Detection in Dravidian Languages, held at DravidianLangTech@EACL 2024[1], is a significant initiative in the field of natural language processing (NLP) and computational linguistics. This research article provides an overview of this event, which focuses on developing and evaluating techniques for identifying fake news specifically in the Dravidian language family. The task aims to address the growing challenge of misinformation in online content by leveraging the linguistic characteristics unique to Dravidian languages. By highlighting the importance of language-specific approaches to fake news detection, this work contributes to advancing the capabilities of NLP models in combating misinformation across diverse linguistic contexts. Fake News Detection (FND) can be categorized as either monolingual or multilingual. Monolingual FND focuses on detecting fake news in a single language, while multilingual FND involves identifying deceptive content that may involve code-mixing or code-switching across two or more languages. Detecting fake news in low-resource languages poses challenges due to limited language resources such as annotated datasets, language models, and pre-trained embeddings. Nevertheless, there are strategies to enhance fake news detection in these languages, including the collection and annotation of data, utilization of cross-lingual models, implementation of transfer learning, and the development of domain-specific models. In the study conducted by Raja et al. (2023), they employed two pre-trained transformer models namely mBERT and XLM-R to assess the feasibility of transfer learning from high-resource languages to low-resource languages in the context of fake news detection within Dravidian languages.

2 Related Work

The authors (Palani and Elango, 2023) utilize contextual word embedding using pre-trained language models like BERT and RoBERTa, and deep learning-based models to identify fake news in Dravidian languages. The authors Chakravarthi et al. (2022) introduced the DL-based FND system, which uses FFN and RoBERTa to identify fake news and extract contextually dependent features, respectively. A Deep learning-based hope speech detection model in which the contextual link between words is captured through word embedding using T5-sentence and Indic-BERT was

[1]https://codalab.lisn.upsaclay.fr/competitions/16055

Proceedings of the Fourth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages, pages 71–78

developed by Chakravarthi (2022). In order to identify the hope speech comments, the contextual elements are fed as input to the CNN model according to Subramanian et al. (2022). The performance of the suggested model is assessed using the HopeEDI multilingual dataset, which is presented in the shared task 2021 (Chakravarthi, 2020).

A multichannel CNN can be used to teach the model a variety of features from multiple perspectives. According to Shanmugavadivel et al. (2022), the model's several channels each extract features from the same input in a different way, resulting in a more accurate representation. A two-stage hope detection approach was proposed by Chinnappa (2021). In this approach, the language detector determines the model's language, and the hope detector categorizes the text as either hope speech, non-hope speech, or not lang. The authors Ahmad et al. (2020) used textual features in conjunction with machine learning-based ensemble algorithms to differentiate between authentic and fraudulent news.

3 Task Description

Task 1: The goal of this task is to classify a given social media text into original or fake. The data sources are various social media platforms such as Twitter, Facebook, etc. Given a social media post, the objective of the shared task is to classify it into either fake or original news. For example, the following two posts belong to fake and original categories, respectively. This is a comment/post-level classification task. Given a YouTube comment, the systems submitted by the participants should classify it into original or fake news. To download the data and participate, go to the Participate tab.

Task 2: The Fake News Detection from Malayalam News (FakeDetect-Malayalam) shared task provides a platform for researchers to address the pressing challenge of identifying and flagging fake news within the realm of Malayalam-language news articles. In an age of information overload, accurate detection of misinformation is crucial for fostering trustworthy communication. The core objective of the FakeDetect-Malayalam shared task is to encourage participants to develop effective models capable of accurately detecting and classifying fake news articles in the Malayalam language into different categories. Here, we considered five fake categories - False, Half True, Mostly False, Partly False, and Mostly True.

4 Dataset

The dataset described in the provided Table 1 pertains to the Second Shared Task on Fake News Detection in Dravidian Languages, conducted at DravidianLangTech@EACL 2024, focusing on two distinct tasks. For Task A, the objective is to classify news items into 'Fake' and 'Original' categories. The dataset for this task comprises 1,599 training instances for 'Fake' and 1,658 for 'Original', with respective testing sets of 507 and 512 instances, and development sets of 406 and 409 instances. Task B is more granular, aiming to categorize news into 'Half True', 'False', 'Partly False', and 'Mostly False'. The numbers of training samples for these categories are 145, 1,251, 44, and 242, respectively. The test set includes 24 instances for 'Half True', 149 for 'False', 14 for 'Partly False', and 63 for 'Mostly False'. The dataset does not include development sets for Task B. This carefully curated dataset is instrumental for researchers in the field of computational linguistics, particularly for those focusing on the development of automated fake news detection systems within the scope of Dravidian languages, which are less commonly addressed in computational research.

Task	Classes	Train	Test	Dev
Task A	Fake	1,599	507	406
	Original	1,658	512	409
Task B	Half True	145	24	-
	False	1,251	149	-
	Partly False	44	14	-
	Mostly False	242	63	-

Table 1: Dataset statistics for Malayalam

5 Participants Methodology

5.1 Task A

- The team **"CUETDUO"** (Rahman et al., 2024) employed a state-of-the-art text classification approach using a pre-trained Malayalam BERT model (Joshi, 2022). The method involved fine-tuning the BERT model on a labeled dataset consisting of Malayalam text samples with corresponding Fake labels. This approach achieved an F1 score of 0.88.

- **"Punny_Punctuators"** (Tabassum et al., 2024) team has used BERT (Devlin et al., 2018) and XLMRoBERTa Base (Conneau et al., 2019) for task A and obtained an F1 score of 0.87.

- The team called **"TechWhiz"**(M et al., 2024) used transformer models to classify the fake news. This team achieved a score of 0.86.

- The team **"CUET_Binary_Hackers"** (Farsi et al., 2024) team employed many BERT models, including indicBert (Kakwani et al., 2020), mBert, specifically MuRIL (Khanuja et al., 2021), and other multilingual BERT models. Out of all the models considered, the team specifically focused on a fine-tuned MuRIL BERT model for the submission and demonstrated an F1 score of 0.86.

- In an attempt by the team **"CUET_NLP_GoodFellows"** (Osama et al., 2024), they have used transformer-based approaches such as XLM-R and mBERT for task 1. And, this attempt produced an F1 score of 0.85.

- The team **"CUETSentimentSillies"** (Fardaush Tripty et al., 2024) did some preprocessing like emoji removal, punctuation removal, English stopwords removal, url removing, and lowercasing. This team also found some most frequent words removed them from the dataset and finetuned the m-bert transformer using huggingFace trainer API. 0.84 is the F1 score from this work.

- The team **"Habesha"** used transformer-based DistilBERT (Sanh et al., 2019) and a combination of deep learning and transformers. This team used the character-based deep learning approach GRU and achieved an F1-Score of 0.82.

- **"KEC_DL_KSK"** team has applied sampling techniques like SMOTE and Random Oversampler to overcome the class imbalance problem. This work tried different word embedding techniques like TFIDF, Word2Vec, Doc2Vec, BERT, and FastText. Machine learning algorithms like Random Forest, SVM, Logistic Regression, Naive Bayes, and Deep Learning models like LSTM, BiLSTM, GRU, BIGRU, and Pre-trained Model BERT for model building and classification have been used in this work.

- The team **"MUCS"** has trained LSVC with word+RcharML, Syllable+ensemble, and TL BERT model and produced 0.84 as F1-Score.

- The team **"Quartet_FakeNews"** implemented the Multinomial Naive Bayes model which leverages advanced text preprocessing techniques and the TF-IDF representation of text data to classify news articles as either fake or not. The Naive Bayes algorithm's assumption of conditional independence among features given the class label makes it a suitable choice for text classification tasks, and the model's performance is thoroughly evaluated to ensure its effectiveness in the context of fake news detection.

- In an attempt by the team **"SCOPE"**, the participant implemented the following steps. Firstly, the text or news data was tokenized, which typically involves breaking down the text into individual words or tokens. Subsequently, four machine learning algorithms—Support Vector Machine (SVM), Naive Bayes, Random Forest, and Logistic Regression—were employed. These algorithms were used for classification tasks, such as sentiment analysis or topic categorization. Each algorithm has its strengths and characteristics, and its performance was evaluated to determine which one yielded the best results for the given task.

- The teams **"Tayyab"** (Zamir et al., 2024) and **"TechWhiz"** (M et al., 2024) have implemented CNN and transformer-based models respectively.

- The team **"WordWizard"** (Anbalagan et al., 2024), Support Vector Machines (SVM), and Naive Bayes models were used. These approaches leveraged features extracted from textual content with Bag-of-Words representations and word embeddings. Comparative analysis with baseline models revealed the superiority of the SVM and Naive Bayes ensemble, achieving competitive accuracy, precision, recall, and F1-score metrics. The model used for the submission of task 1 is Support Vector Machine. TF-IDF has been used to vectorize the given text based on the relevancy of the word. For the second task, the predictions given by the Naive Bayes Model have been submitted.

5.2 Task B

A total of twelve teams submitted their models' predictions for problem 2. Each team had the opportunity to submit a maximum of three runs. This section provides a concise overview of the methods employed by the participants to address the issue of classifying fake news in Malayalam.

- **CUET_Binary_Hackers**: The team CUET_Binary_Hackers (Farsi et al., 2024) team employed many BERT models, including indicBert, mBert, MuRIL, and other multilingual BERT models. Out of all the models considered, the team specifically concentrated on a fine-tuned MuRIL BERT model for the submission, which demonstrated superior accuracy and F1 score performance. Throughout the fine-tuning process, the team made modifications and conducted trials with various hyperparameters, including learning rates, batch size, and optimization strategies, to enhance performance.

- **CUETSentimentSilles**: The data was cleaned by excluding punctuation, URLs, emojis, digits, and the most commonly utilized words from the corpus (Fardaush Tripty et al., 2024). Due to the significant corpus imbalance, the authors employed a data augmentation strategy to achieve class balance. The team fine-tuned the Malayalam BERT transformer to construct the models for task 2.

- **Quartet**: The model developed by the team Quartet for this task is constructed via a Multinomial Naive Bayes classifier, and its effectiveness is improved by employing a TfidfVectorizer to transform text data into a numerical representation suited for machine learning. The model development process commences with the application of a sequence of text preparation procedures. The process involves eliminating numerical values and special characters and lemmatizing words using the WordNet lemmatizer. The optimal model, identified through a previous grid search conducted with GridSearchCV, is subsequently employed on the complete training dataset. The model comprises a TfidfVectorizer that utilizes a predefined n-gram range and a Multinomial Naive Bayes classifier that incorporates an ideal alpha parameter. After being trained, the model is then utilized to generate predictions on the test set, generating numerical labels for each news article. In order to simplify the interpretation and reporting process, a mapping dictionary is used to transform the numerical predictions into their appropriate textual labels. This mapping encompasses the five categories of the task.

- **byteSizedLLM**: The team employed embeddings derived from a subset of AI4Bharat's data, comprising 100,000 randomly selected lines (Kodali and Manukonda, 2024). The embeddings were generated using custom-built subword tokenizers for Telugu (7.6 MB) and Tamil (1.3 MB) languages. A Bidirectional Long Short-Term Memory (BiLSTM) classifier was used for classification tasks. The model underwent training using datasets annotated with labels, which were later used to deduce test set outcomes. A customized subword tokenizer was employed on a limited dataset with BiLSTM models, which exhibit exceptional speed and have a small memory footprint (less than 8 MB).

- **KEC_DL_KSK**: The team utilized sampling techniques such as SMOTE and Random Oversampler to address the issue of class imbalance. Various word embedding algorithms were employed, including TFIDF, Word2Vec, Doc2Vec, BERT, and FastText. The team employed various machine learning methods, including Random Forest, Support Vector Machine (SVM), Logistic Regression, and Naive Bayes, as well as deep learning models such as Long Short-Term Memory (LSTM), Bidirectional LSTM (BiLSTM), Gated Recurrent Unit (GRU), Bidirectional GRU (BIGRU), and a pre-trained BERT model, to construct the classifier.

- **WordWizard**: The team introduced a methodology for classifying fake news using the Naive Bayes model (Anbalagan et al., 2024). The proposed approach utilizes features derived from textual content by merging embeddings from the LaBSE model with TF-IDF features. Data preprocessing is a compelling component of the system. The dataset was cleaned to exclude commonly used Malayalam stop-words, ensuring that only the most relevant terms are retained. Additional preprocessing techniques, such as stemming and lemmatization, were employed. Compared with baseline models, the Naive Bayes model demonstrates superiority by obtaining competitive accuracy, precision, recall, and F1-score metrics.

- **Habesha**: The team employed character-based Recurrent Neural Network (RNN), Gated Recurrent Unit (GRU), and DistilBERT models to build the classifier.

- **SCOPE**: The team employed a series of sequential steps. Initially, the input text data underwent tokenization. Following that, four machine learning techniques, namely Support Vector Machine (SVM), Naive Bayes, Random Forest, and Logistic Regression, were utilized for classification.

- **Tayyab**: This team used Convolution Neural Networks (CNN) for classification (Zamir et al., 2024).

- **MUCS**: The team employed a linear support vector machine classifier trained using word-level TF-IDF features. An oversampling strategy was utilized to address the class imbalance issue in the data. Subsequently, a logistic regression classifier was employed for classification. A collection of Siamese networks was also considered for constructing the classifier.

- **Punny_Punctuators**: The team employed a multilingual BERT-cased model and a CNN for categorization. The class imbalance problem was handled by implementing a data augmentation technique known as back translation (Tabassum et al., 2024).

- **TechWhiz**: The team used transformer models to classify the fake news (M et al., 2024).

6 Results

The rank list for Task A and Task B with Macro F1-Score is shown in Table 2 and Table 3. The models proposed by the participant teams are evaluated using the macro F1-Score metric. The evaluation results of task A are presented in Table 2 which represents the macro F1-score of each team rank-wise. 18 teams are participating in Task A and submitted their runs.

From the submissions, it is interpreted that transformer-based models outperform machine learning and deep learning models with better F1-Score. A few methods based on RNN and CNN gave slightly lesser F1-Score than transformer models. Following the RNNs and CNNs, the machine learning models such as SVM, Naive Bayes, Random Forest, and Logistic Regression gave F1-Score ranging from 0.80 to 0.71

7 Conclusion

This paper presents an overview of the Second Shared Task on Fake News Detection in Dravidian Languages - DravidianLangTech@EACL 2024. The task attracted participation from eighteen teams for Task A and Twelve teams for Task B. They employed methods varied among the teams, ranging from traditional TF-IDF vectorizers with machine learning to contemporary pre-trained transformer models for data representation. Upon analyzing the methodologies, a consistent trend emerged: transformer-based approaches consistently outperformed other techniques, as evidenced by evaluation metrics like classification accuracy and confusion matrices. This underscores the effectiveness of transformer models in capturing the performance of fake news detection. To sum up, the paper summarizes the DravidianLangTech@EACL 2024 fake news detection shared task for Malayalam, emphasizing diverse strategies and highlighting the prevalence of transformer-based methods for enhanced performance.

Acknowledgements

This work was conducted with the financial support of the Science Foundation Ireland Centre for Research Training in Artificial Intelligence under Grant No. 18/CRT/6223, supported in part by a research grant from Science Foundation Ireland (SFI) under Grant Number SFI/12/RC/2289_P2(Insight_2).

References

Iftikhar Ahmad, Muhammad Yousaf, Suhail Yousaf, and Muhammad Ovais Ahmad. 2020. Fake news detection using machine learning ensemble methods. *Complexity*, 2020:1–11.

S. No.	Team Name	Run	Macro F1-Score	Rank
1	CUET_DUO (Rahman et al., 2024)	1	0.88	1
2	Punny_Punctuators (Tabassum et al., 2024)	1, 2 & 3	0.87	2
3	TechWhiz (M et al., 2024)	xlmr	0.86	3
4	CUET_Binary_Hackers (Farsi et al., 2024)	1 & 2	0.86	3
5	CUET_NLP_GoodFellows (Osama et al., 2024)	1	0.85	4
6	CUETSentimentSilles (Fardaush Tripty et al., 2024)	1	0.84	5
7	DLRG	1	0.84	5
8	MUCS	2 & 3	0.84	5
9	CUET_DASH	3	0.83	6
10	Habesha	1	0.82	7
11	KEC_TECH	1	0.82	7
12	Quartet_FakeNews	1	0.81	8
13	KEC_HAWKS (Subramanian et al., 2024)	2	0.80	9
14	KEC_DL_KSK	2	0.79	10
15	SCOPE	1	0.78	11
16	Tayyab (Zamir et al., 2024)	1	0.78	11
17	WordWizard (Anbalagan et al., 2024)	1	0.78	11
18	Fango	1	0.71	12

Table 2: Rank list for the Task A

S. No.	Team	Run	macro F1 score	Rank
1	CUET_Binary_Hackers (Farsi et al., 2024)	1	0.5191	1
2	CUETSentimentSilles (Fardaush Tripty et al., 2024)	1	0.4964	2
3	Quartet	1	0.4868	3
4	byteSizedLLM (Kodali and Manukonda, 2024)	1	0.4797	4
5	KEC_DL_KSK	2	0.4763	5
6	WordWizard (Anbalagan et al., 2024)	1	0.3517	6
7	Habesha	1	0.3153	7
8	SCOPE	1	0.3039	8
9	Tayyab (Zamir et al., 2024)	1	0.2393	9
10	MUCS	1	0.1867	10
11	Punny_Punctuators (Tabassum et al., 2024)	3	0.1747	11
12	TechWhiz (M et al., 2024)	2	0.1733	12

Table 3: Rank list for the Task B

Akshatha Anbalagan, Priyadharshini T, Niranjana A, Shreedevi Seluka Balaji, and Durairaj Thenmozhi. 2024. Wordwizarddravidianlangtech@eacl 2024-fake news detection in dravidian languages. In *Proceedings of the Fourth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*, Malta. European Chapter of the Association for Computational Linguistics.

Bharathi Raja Chakravarthi. 2020. Hopeedi: A multilingual hope speech detection dataset for equality, diversity, and inclusion. In *Proceedings of the Third Workshop on Computational Modeling of People's Opinions, Personality, and Emotion's in Social Media*, pages 41–53.

Bharathi Raja Chakravarthi. 2022. Hope speech detection in youtube comments. *Social Network Analysis and Mining*, 12(1):75.

Bharathi Raja Chakravarthi, Vigneshwaran Muralidaran, Ruba Priyadharshini, Subalalitha Cn, John Philip McCrae, Miguel Ángel García, Salud María Jiménez-Zafra, Rafael Valencia-García, Prasanna Kumaresan, Rahul Ponnusamy, et al. 2022. Overview of the shared task on hope speech detection for equality, diversity, and inclusion. In *Proceedings of the second workshop on language technology for equality, diversity and inclusion*, pages 378–388.

Dhivya Chinnappa. 2021. dhivya-hope-detection@ lt-edi-eacl2021: multilingual hope speech detection for code-mixed and transliterated texts. In *Proceedings* *of the First Workshop on Language Technology for Equality, Diversity and Inclusion*, pages 73–78.

Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Unsupervised cross-lingual representation learning at scale. *CoRR*, abs/1911.02116.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: pre-training of deep bidirectional transformers for language understanding. *CoRR*, abs/1810.04805.

Zannatul Fardaush Tripty, Md. Arian Al Nafis, Antu Chowdhury, Jawad Hossain, Shawly Ahsan, and Mohammed Moshiul Hoque. 2024. Cuetsentimentsillies@dravidianlangtech-eacl2024: Transformer-based approach for detecting and categorizing fake news in malayalam language. In *Proceedings of the Fourth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*, Malta. European Chapter of the Association for Computational Linguistics.

Salman Farsi, Asrarul Hoque Eusha, Ariful Islam, Hasan Mesbaul Ali Taher, Jawad Hossain, Shawly Ahsan, and Mohammed Moshiul Hoque. 2024. Cuet_binary_hackers@dravidianlangtech 2024: Malayalam fake news detection by fine-tuned muril bert. In *Proceedings of the Fourth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*, Malta. European Chapter of the Association for Computational Linguistics.

Raviraj Joshi. 2022. L3cube-hindbert and devbert: Pre-trained bert transformer models for devanagari based hindi and marathi languages. *arXiv preprint arXiv:2211.11418*.

Divyanshu Kakwani, Anoop Kunchukuttan, Satish Golla, Gokul N.C., Avik Bhattacharyya, Mitesh M. Khapra, and Pratyush Kumar. 2020. IndicNLPSuite: Monolingual Corpora, Evaluation Benchmarks and Pre-trained Multilingual Language Models for Indian Languages. In *Findings of EMNLP*.

Simran Khanuja, Diksha Bansal, Sarvesh Mehtani, Savya Khosla, Atreyee Dey, Balaji Gopalan, Dilip Kumar Margam, Pooja Aggarwal, Rajiv Teja Nagipogu, Shachi Dave, Shruti Gupta, Subhash Chandra Bose Gali, Vish Subramanian, and Partha Talukdar. 2021. Muril: Multilingual representations for indian languages.

Rohith Gowtham Kodali and Durga Prasad Manukonda. 2024. bytesizedllm@dravidianlangtech 2024: Fake news detection in dravidian languages - unleashing the power of custom subword tokenization with subword2vec and bilstm. In *Proceedings of the Fourth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*, Malta. European Chapter of the Association for Computational Linguistics.

Madhumitha M, Kunguma Akshatra M, Tejashri J, and C Jerin Mahibha. 2024. Techwhiz@eacl 2024: Fake news detection using deep learning models. In *Proceedings of the Fourth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*, Malta. European Chapter of the Association for Computational Linguistics.

Md Osama, Kawsar Ahmed, Hasan Mesbaul Ali Taher, Jawad Hossain, Shawly Ahsan, and Mohammed Moshiul Hoque. 2024. Cuet_nlp_goodfellows@dravidianlangtech 2024: 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*, Malta. European Chapter of the Association for Computational Linguistics.

Balasubramanian Palani and Sivasankar Elango. 2023. Bbc-fnd: An ensemble of deep learning framework for textual fake news detection. *Computers and Electrical Engineering*, 110:108866.

Tanzim Rahman, Abu Bakkar Siddique Raihan, Md. Tanvir Rahman, Jawad Hossain, Shawly Ahsan, Avishek Das, and Mohammed Moshiul Hoque. 2024. Cuet_duo@dravidianlangtech 2024: Fake news classification using malayalam-bert. In *Proceedings of the Fourth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*, Malta. European Chapter of the Association for Computational Linguistics.

Eduri Raja, Badal Soni, and Samir Kumar Borgohain. 2023. Fake news detection in dravidian languages using transfer learning with adaptive finetuning. *Engineering Applications of Artificial Intelligence*, 126:106877.

Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *ArXiv*, abs/1910.01108.

Kogilavani Shanmugavadivel, Sai Haritha Sampath, Pramod Nandhakumar, Prasath Mahalingam, Malliga Subramanian, Prasanna Kumar Kumaresan, and Ruba Priyadharshini. 2022. An analysis of machine learning models for sentiment analysis of tamil code-mixed data. *Computer Speech & Language*, 76:101407.

Malliga Subramanian, Ramya Chinnasamy, Prasanna Kumar Kumaresan, Vasanth Palanikumar, Madhoora Mohan, and Kogilavani Shanmugavadivel. 2022. Development of multi-lingual models for detecting hope speech texts from social media comments. In *International Conference on Speech and Language Technologies for Low-resource Languages*, pages 209–219. Springer.

Malliga Subramanian, Jayanth J R, Muthu Karuppan P, KeerthiBala A T, and Kogilavani Shanmugavadivel. 2024. Kec hawks@dravidianlangtech-2024 : Detecting malayalam fake news using machine learning models. In *Proceedings of the Fourth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*, Malta. European Chapter of the Association for Computational Linguistics.

Nafisa Tabassum, Sumaiya Rahman Aodhora, Rowshon Akter, Jawad Hossain, Shawly Ahsan, and Mohammed Moshiul Hoque. 2024. Punny_punctuators@dravidianlangtech-eacl2024: Transformer-based approach for detection and classification of fake news in malayalam social media text. In *Proceedings of the Fourth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*, Malta. European Chapter of the Association for Computational Linguistics.

M. T. Zamir, M. S Tash, Z. Ahani, A. Gelbukh, and G. Sidorov. 2024. Tayyab@dravidianlangtech 2024:detecting fake news in malayalam lstm approach and challenges. In *Proceedings of the Fourth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*, Malta. European Chapter of the Association for Computational Linguistics.