

Non-Contextual BERT or FastText? A Comparative Analysis

Abhay Shanbhag^{1,3}, Suramya Jadhav^{1,3}, Amogh Thakurdesai^{1,3}, Ridhima Sinare^{1,3}
and Raviraj Joshi^{2,3}

¹Pune Institute of Computer Technology, Pune

²Indian Institute of Technology Madras, Chennai

³L3Cube Labs, Pune

Abstract

Natural Language Processing (NLP) for low-resource languages, which lack large annotated datasets, faces significant challenges due to limited high-quality data and linguistic resources. The selection of embeddings plays a critical role in achieving strong performance in NLP tasks. While contextual BERT embeddings require a full forward pass, non-contextual BERT embeddings rely only on table lookup. Existing research has primarily focused on contextual BERT embeddings, leaving non-contextual embeddings largely unexplored. In this study, we analyze the effectiveness of non-contextual embeddings from BERT models (MuRIL and MahaBERT) and FastText models (IndicFT and MahaFT) for tasks such as news classification, sentiment analysis, and hate speech detection in one such low-resource language—Marathi. We compare these embeddings with their contextual and compressed variants. Our findings indicate that non-contextual BERT embeddings extracted from the model’s first embedding layer outperform FastText embeddings, presenting a promising alternative for low-resource NLP.

1 Introduction

Word embedding is a way of representing words into dense vectors in a continuous space such that the vectors capture the semantic relationship between the words for the models to understand the context and meaning of the text. FastText, a context-independent method, basically captures the subword information, enabling it to learn rare words, misspelled words, and out-of-vocabulary words. It is recognized in the NLP community for its efficient performance in tasks like text classification and sentiment analysis. Despite being relatively old, it still remains one of the most effective alternatives when performing tasks on large datasets across various languages due to its subword-based approach.

BERT (Bidirectional Encoder Representations from Transformers) (Devlin et al., 2018) word embeddings understand the meaning of a word based on its context in a sentence. The embeddings extracted just before the first embedding layer of the BERT architecture are referred to as non-contextual embeddings, while those obtained from the last hidden layer of BERT are known as contextual embeddings (Refer Figure 1). Numerous variations of BERT like IndicBERT (Kakwani et al., 2020), MuRIL (Khanuja et al., 2021), AfriBERT (Ralethe, 2020), and mBERT (Devlin et al., 2018) to name a few, are available for experiments.

Recent studies have experimented with both FastText and BERT for various tasks; however, most of them focus on exploring contextual BERT embeddings. Experiments of D’Sa et al. (2020) demonstrated that BERT embeddings outperformed FastText for classifying English text into toxic and non-toxic. Findings of Ahmed et al. (2024) suggested that BERT embeddings outperformed those of FastText with an F1 score of 84% when evaluated for depressive post-detection in Bangla.

While BERT consistently outperforms other word embeddings in various tasks for high-resource languages (HRLs) like English (Malik et al. (2021)), its effectiveness in low-resource languages (LRLs) remains relatively underexplored. This gap is particularly pronounced when balancing model performance with computational efficiency, which becomes a critical factor in low-resource settings.

Previous studies (D’Sa et al. (2020)) have focused on contextual BERT embeddings, which outperform FastText due to their ability to capture contextual information. However, the use of non-contextual BERT embeddings for classification tasks in low-resource languages like Marathi remains unexplored. Unlike contextual embeddings, which require a full forward pass through the model, non-contextual embeddings can be obtained through a simple table lookup. To our knowledge,

no prior work has examined the effectiveness of non-contextual BERT embeddings. We investigate how these embeddings, extracted from the model’s first layer, compare to FastText embeddings for tasks such as news classification, sentiment analysis, and hate speech detection in Marathi.

Additionally, past comparisons often used BERT’s 768-dimensional embeddings against FastText’s 300-dimensional ones, which is unfair since higher dimensions naturally provide better feature extraction. To address this, we ensure a fair comparison by reducing the BERT embeddings to 300 dimensions.

This paper focuses on utilizing FastText and non-contextual BERT for the Marathi language for the following tasks: Sentiment Classification, 2-Class and 4-Class Hate Speech Detection, and News Article Classification for headlines, long paragraphs, and long documents. We construct a comprehensive analysis of FastText embeddings, IndicFT (Kakwani et al., 2020) and MahaFT (Joshi, 2022) embeddings, and BERT embeddings, including muril-base-cased (Khanuja et al., 2021) and marathi-bert-v2 (Joshi, 2022). To enhance the comparison, we replicate the experiments using widely utilized contextual BERT embeddings. We also evaluate the impact of compression on both contextual and non-contextual BERT-based embeddings. Our analysis shows that non-contextual BERT embeddings generally perform better than FastText in most tasks. Furthermore, contextual BERT embeddings consistently outperform FastText across all evaluated tasks. However, compressing non-contextual embeddings reduces their performance, making FastText more effective than compressed non-contextual BERT.

The key contributions of this work are as follows:

- We conduct a detailed study comparing non-contextual BERT embeddings and FastText embeddings for Marathi, a low-resource language. The evaluation covers multiple classification tasks, including sentiment analysis, news classification, and hate speech detection.
- To ensure a fair comparison, we compress BERT embeddings from 768 to 300 dimensions using Singular Value Decomposition (SVD). This allows us to analyze how dimensionality reduction impacts BERT’s performance compared to its uncompressed version and FastText.

- We explore the differences between contextual and non-contextual BERT embeddings, examining their impact on classification performance in low-resource settings.

The paper is organized as follows: Section 2 provides a concise review of prior research on FastText and BERT. Section 3 includes the datasets and model embeddings that are utilized for the experiments. Section 4 presents the methodology used. Section 5 presents the results and key insights drawn from the findings along with a comparative analysis of FastText embeddings and BERT. In Section 6, we analyze our results and explain the reasons behind them. In Section 7, we conclude our discussion.

2 Literature Review

The existing literature emphasizes the superiority of contextual BERT embeddings over traditional word embedding techniques like Word2Vec (Mikolov et al., 2013), GloVe (Pennington et al., 2014), and FastText across various natural language processing (NLP) tasks. For instance, Khaled et al. (2023) compare four popular pre-trained word embeddings—Word2Vec (via Aravec (Mohammad et al., 2017)), GloVe, FastText, and contextual BERT (via ARBERTv2)—on Arabic news datasets. They highlight BERT’s superior performance, achieving over 95% accuracy due to its contextual interpretation.

Similarly, Kabullar and Türker (2022) analyzes the performance of embeddings on the AG News dataset, which includes 120K instances across four classes. They conclude that contextual BERT outperforms other methods, achieving 90.88% accuracy. FastText, Skip-Gram, CBOW, and GloVe achieve 86.91%, 85.82%, 86.15%, and 80.86%, respectively.

While traditional embeddings perform reasonably well, the consistent dominance of contextual BERT in complex tasks is also noted in sentiment analysis. For instance, Xie et al. (2024) explores how combining BERT and FastText embeddings enhances sentiment analysis in education, demonstrating that BERT’s contextual understanding, along with FastText’s ability to handle out-of-vocabulary words, improves generalization over unseen text.

In the domain of toxic speech classification, D’Sa et al. (2020) utilize both contextual BERT and FastText embeddings to classify toxic comments

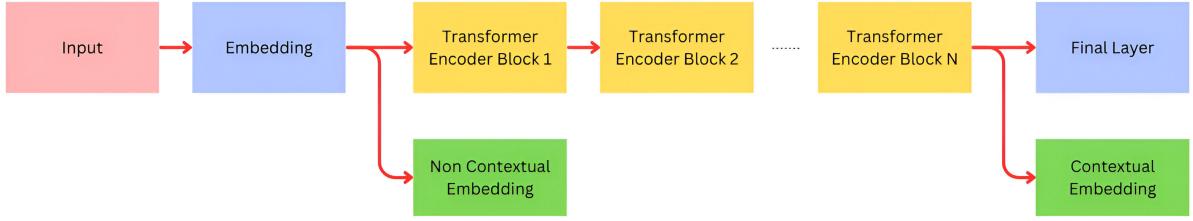


Figure 1: Embedding extraction workflow for contextual and non-contextual representations

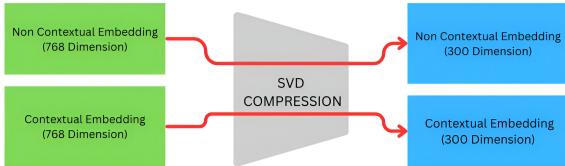


Figure 2: SVD compression of BERT embeddings

in English, with BERT embeddings outperforming FastText. This trend continues in hate speech detection, where [Rajput et al. \(2021\)](#) find that neural network classifiers using contextual BERT embeddings perform better than those with FastText embeddings alone, further supporting BERT’s effectiveness.

Additionally, [Chanda \(2021\)](#) assess contextual BERT embeddings against traditional context-free methods (GloVe, Skip-Gram, and FastText) for disaster prediction, demonstrating BERT’s superior performance in combination with traditional machine learning and deep learning methods.

For low-resource languages (LRLs), [Ahmed et al. \(2024\)](#) examine methods like traditional TF-IDF, contextual BERT, and FastText embeddings within a CNN-BiLSTM architecture for detecting depressive texts in Bangla. Their results show that BERT embeddings yield the highest F1 score (84%), indicating their dominance over other methods. This suggests that BERT’s efficacy extends even to LRLs.

In medical applications, [Khan et al. \(2024\)](#) proposes integrating contextual BERT embeddings with SVM for prostate cancer prediction. By incorporating both numerical data and contextual information from clinical text, they achieve 95% accuracy, far outperforming the 86% accuracy achieved with numerical data alone.

Moreover, [Malik et al. \(2021\)](#) uses both contextual BERT and FastText embeddings to preprocess a dataset of conversations from Twitter and Facebook. Applying various machine learning and deep learning algorithms, they find that CNN yields the

best results, further demonstrating BERT’s capabilities.

Finally, while [Asudani et al. \(2023\)](#) offers a comprehensive analysis of traditional word embeddings alongside more advanced techniques like ELMo and contextual BERT, providing insight into commonly used datasets and models for benchmarking, [Umer et al. \(2022\)](#) highlights the versatility of FastText in various domains, despite BERT’s consistently superior performance.

We note that the reviewed literature highlights the consistent superiority of BERT embeddings across various NLP tasks and domains. However, most existing studies focus mainly on contextual BERT embeddings, but not on non-contextual embeddings. Moreover, these studies predominantly address high-resource languages, leaving low-resource languages like Marathi largely unexplored. In particular, there is a lack of research assessing the effectiveness of non-contextual BERT embeddings for Marathi. Additionally, the impact of dimensionality leveling, i.e. the efficacy of BERT embedding compression, has not been explored.

3 Datasets and Models Used

In our research work, we used 3 Marathi datasets, [MahaSent](#): A 3-class sentiment analysis dataset ([Pingle et al., 2023](#)), [MahaHate](#): A 2-class as well as a 4-class hate classification dataset ([Patil et al., 2022](#)) and [MahaNews](#) is a news categorization dataset consisting of three sub-datasets, each with 12 classes: Short Headline Classification (SHC), Long Document Classification (LDC), and Long Paragraph Classification (LPC) ([Mittal et al., 2023](#)).

We used two types of embeddings in our experiments: FastText and BERT embeddings. For FastText, we utilized both [IndicFT](#) ([Kakwani et al., 2020](#)) and [MahaFT](#) ([Joshi, 2022](#)) embeddings. This was because both models included a Marathi corpus as part of their training data. MahaFT, in par-

Type	Model	MahaSent	MahaHate		MahaNews		
		3-class	4-class	2-class	SHC	LDC	LPC
Contextual	MahaBERT	82.27	66.8	85.57	89.83	93.87	87.78
	MahaBERT (Compressed)	82.89	66.15	84.37	89.61	93.53	87.82
	Muril	81.64	64.55	84.00	89.54	93.64	87.33
	Muril (Compressed)	81.91	63.2	83.36	88.38	93.48	87.45
FastText	IndicFT	76.4	58.25	80.13	85.57	92.15	79.19
	MahaFT	78.62	62.75	81.76	85.89	92.62	80.32
Non-Contextual	MahaBERT	77.56	66.5	82.64	86.45	91.69	81.76
	MahaBERT (Compressed)	76.31	63.9	81.57	83.85	91.25	80.08
	Muril	76.58	65.77	81.79	85.95	91.61	81.36
	Muril (Compressed)	75.16	63.25	81.44	82.72	90.39	79.00

Table 1: Performance of model embeddings on MahaSent, MahaHate, and MahaNews datasets using Multiple Logistic Regression. Key: SHC = Short Headline Classification, LPC = Long Paragraph Classification, LDC = Long Document Classification

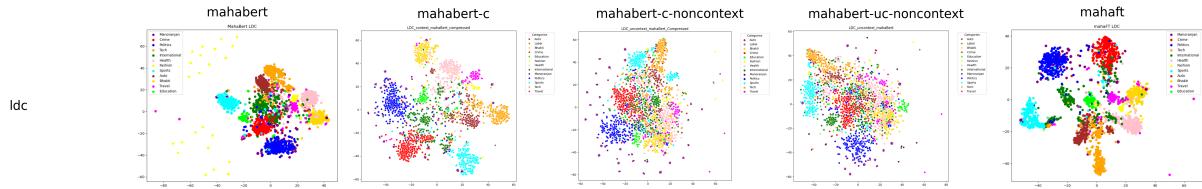


Figure 3: T-SNE Plot For BERT and FastText Embeddings (c stands for compressed) .

ticular, was specifically trained on a Marathi corpus, making it especially relevant for our experiments. For BERT embeddings, we primarily used two BERT-based models: **MahaBERT** (Joshi, 2022) and **MuRIL** (Khanuja et al., 2021). Since both models were trained on Marathi data, we selected them to compare with the FastText embeddings.

4 Methodology

For each sentence, corresponding embeddings were generated and the corresponding categorical labels were encoded into numerical labels. The creation of BERT embeddings was done by first tokenizing the text using the BERT tokenizer, along with padding and truncation. The tokenized input was then passed to the model and the output of the last hidden layer of BERT was taken, which was then averaged to get contextual embeddings for every sentence. Whereas for non-contextual embeddings, the output of the first embedding layer was used. Refer Figure 1 for the embedding extraction workflow for contextual and non-contextual representations.

However, for FastText, which is a non-contextual embedding by default, the process was slightly different due to the lack of a predefined vocabulary. Unlike BERT, which employs a tokenizer capable of processing entire Marathi sentences, FastText

necessitates the creation of a custom vocabulary. To achieve this, the training and validation datasets were concatenated and passed through a text vectorizer, which generated vectors for every word in the dataset. The vectorizer returned the vocabulary as a list of words in decreasing order of their frequency. The FastText model was then loaded using the FastText library, and for each word in the vocabulary, a word vector was retrieved to construct the embedding matrix. For each sentence, the text was split into individual words, and the corresponding embeddings were retrieved from the embedding matrix. These embeddings were then averaged to produce the final sentence embeddings. Additionally, we experimented with compressed embeddings by reducing the dimensionality from 768 (the traditional BERT embedding dimension) to 300. This compression was performed using Singular Value Decomposition (SVD) to select the most relevant features, extracting the top 300 components for all the combinations of contextual as well as non-contextual for MahaBERT as well as Muril. Refer Figure 2 for SVD compression of BERT embeddings.

Feature scaling was also applied to the outputs. All embeddings were then passed to a multiple logistic regression (MLR) classifier for classification into target labels.

Dataset	Subdataset	Model	Avg	Variance	Std	Test
MahaSent	3 Class	MahaBERT	76.56	0.39843	0.6312	78.01
		MahaBERT-Compressed	74.42	0.8498	0.9218	75.51
		Muril	75.53	0.75268	0.8676	76.53
		Muril-Compressed	72.97	0.48963	0.6997	75.2
		MahaFT	77.28	0.38282	0.6187	78.58
MahaHate	4 Class	MahaBERT	64.92	0.25203	0.5020	66.1
		MahaBERT-Compressed	62.77	0.53875	0.7340	64.1
		Muril	63.51	0.35307	0.5942	65.15
		Muril-Compressed	61.22	0.52378	0.7237	62.9
		MahaFT	62.48	0.22608	0.4755	62.55
MahaNews	2 Class	MahaBERT	84.23	0.37633	0.6135	82.53
		MahaBERT-Compressed	82.3	0.10312	0.3211	81.41
		Muril	83.69	0.39397	0.6277	81.63
		Muril-Compressed	81.67	0.20943	0.4576	81.41
		MahaFT	83.75	0.52153	0.7222	82.61
	SHC	MahaBERT	86.66	0.27687	0.5262	86.64
		MahaBERT-Compressed	84.13	0.36002	0.6000	83.81
		Muril	85.7	0.06973	0.2641	85.66
		Muril-Compressed	82.89	0.11612	0.3408	82.01
		MahaFT	87.25	0.17873	0.4228	85.97
	LDC	MahaBERT	92.47	0.32565	0.5707	91.69
		MahaBERT-Compressed	91.41	0.01637	0.1279	91.57
		Muril	92.03	0.19055	0.4365	91.69
		Muril-Compressed	91.04	0.07753	0.2784	90.39
		MahaFT	92.79	0.15667	0.3958	92.71
	LPC	MahaBERT	81.71	0.18503	0.4302	81.27
		MahaBERT-Compressed	80.03	0.1779	0.4218	80.51
		Muril	81.19	0.17597	0.4195	81.4
		Muril-Compressed	78.82	0.14497	0.3807	79.11
		MahaFT	80.15	1.25257	1.1192	80.32

Table 2: The values were obtained by performing 5-fold cross-validation on the training dataset for **Non-contextual embedding**. The **Avg**, **Variance** and **Std** represent the average, variance and standard deviation respectively performance across the five test subsets (from training) of the 5-fold splits, while the **Test** column reflects the performance on the actual test dataset. Key: SHC = Short Headline Classification, LPC = Long Paragraph Classification, LDC = Long Document Classification

4.1 Experimental Setup

The experiments were conducted on Kaggle notebooks equipped with a P100 GPU accelerator, utilizing 16 GB of GPU memory, 20 GB of storage, and 32 GB of RAM. Accuracy was chosen as the evaluation metric, given the balanced nature of the datasets. For classification, the results obtained from the embeddings were mapped to final labels using a multinomial logistic regression model to maintain methodological simplicity. To determine the validity of the results obtained, 5-fold cross-validation was performed for all tasks, and the results are presented in Table 2.

4.2 Visualisation of Embeddings

To visualize how BERT and FastText embedding can separate the classes, we plotted T-SNE (van der Maaten and Hinton, 2008) graphs for the LDC dataset. We have 5 plots, with 4 plots for MahaBERT and 1 for MahaFT. Refer Figure 3.

5 Results

Table 1 presents the results for various embeddings, including MahaBERT, MuRIL, MahaFT, and IndicFT, across multiple datasets and tasks. It includes both contextual and non-contextual embeddings, as well as the compressed variants of MahaBERT and MuRIL.

In sections 5.1 and 5.2, we have considered the uncompressed versions of Muril and MahaBERT. Further, in section 5.3, we specifically show the effect of compression on Muril and MahaBERT.

5.1 Contextual vs FastText

From Table 1, we observe the following trend when comparing contextual embeddings with FastText embeddings: MahaBERT > MuRIL > MahaFT > IndicFT.

5.2 Non-Contextual vs FastText

The trend of comparing non-contextual embeddings with FastText typically follows this order:

MahaBERT > MuRIL > MahaFT > IndicFT. However, there are exceptions for the MahaSent and LDC datasets.

For these two datasets, FastText tends to perform slightly better. However, the difference is minimal, so we refer to Table 2 to determine whether this deviation is significant or simply random noise. We observe a high variance in MahaSent, suggesting that its deviation from the usual trend when comparing non-contextual embeddings with FastText may be attributed to noise and is unlikely to be significant.

In contrast, the LDC dataset also deviates from the trend but exhibits relatively low variance. As a result, for the LDC dataset, the performance trend when comparing non-contextual embeddings with FastText is as follows: MahaFT > IndicFT > MahaBERT > MuRIL.

5.3 Effect of Compression

From Table 1, it can be inferred that compression negatively impacts non-contextual embeddings, as uncompressed versions generally perform better. This is evident from MahaFT outperforming the compressed non-contextual MahaBERT embeddings in all datasets except MahaHate-4c, suggesting that compression lowers the performance of non-contextual BERT embeddings.

However, the effect of compression on contextual embeddings varies across datasets, making it challenging to derive a consistent conclusion.

6 Inference

In this section, we explain why the non-contextual MahaBERT embeddings outperform FastText (MahaFT and IndicFT) embeddings. Both MahaBERT and MahaFT embeddings have been trained on the same corpus of 752 million tokens Joshi (2022). The superior performance of non-contextual MahaBERT embeddings can be attributed to its larger embedding size, training data size, and contextual training objective. Specifically, the embedding size for Marathi-BERT-v2 is 152M ($197,285 \times 768$), compared to MahaFT, which is 132M ($439,247 \times 300$).

IndicFT performs worse than MahaFT, likely due to its smaller dataset size of 551 million tokens (Kakwani et al., 2020). On the other hand, contextual BERT achieves better results because its hidden layers are effectively utilized.

Additionally, we observe a negative impact when

compressing MahaBERT non-contextual embeddings. Reducing the embedding size from 152M ($197,285 \times 768$) to 59M ($197,285 \times 300$) leads to a decrease in performance, likely due to the loss of representational capacity.

7 Conclusion

In our research, we analyzed the effectiveness of various BERT and FastText-based embeddings on three key NLP tasks for Marathi: news classification, hate speech classification, and sentiment classification focusing primarily on non-contextualised BERT embeddings.

Our results show that contextual BERT embeddings perform better than non-contextual ones, including both non-contextual BERT embeddings and FastText. Among non-contextual embeddings, BERT generally outperforms FastText in most tasks. However, when non-contextual BERT embeddings are compressed, their performance drops, and FastText performs better than compressed non-contextual BERT.

References

Saad Ahmed, Mahdi H Sazan, Miraz A B M Muntasir, Rahman, Saad Ahmed Sazan, Mahdi H. Miraz, and M Muntasir Rahman. 2024. [Enhancing depressive post detection in bangla: A comparative study of tf-idf, bert and fasttext embeddings](#). *ArXiv*, abs/2407.09187.

Deepak Suresh Asudani, Naresh Kumar Nagwani, and Pradeep Singh. 2023. [Impact of word embedding models on text analytics in deep learning environment: a review](#). *Artificial Intelligence Review*, pages 1 – 81.

Ashis Kumar Chanda. 2021. [Efficacy of bert embeddings on predicting disaster from twitter data](#). *ArXiv*, abs/2108.10698.

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.

Ashwin Geet D’Sa, Irina Illina, and D. Fohr. 2020. [Bert and fasttext embeddings for automatic detection of toxic speech](#). *2020 International Multi-Conference on: “Organization of Knowledge and Advanced Technologies” (OCTA)*, pages 1–5.

Raviraj Joshi. 2022. [L3Cube-MahaCorpus and MahaBERT: Marathi monolingual corpus, Marathi BERT language models, and resources](#). In *Proceedings of the WILDRE-6 Workshop within the 13th Language Resources and Evaluation Conference*, pages

97–101, Marseille, France. European Language Resources Association.

Elif Kabullar and İlker Türker. 2022. Performance comparison of word embedding methods in text classification for various number of features.

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 the Association for Computational Linguistics: EMNLP 2020*, pages 4948–4961, Online. Association for Computational Linguistics.

M Moneb Khaled, Muhammad Al-Barham, Osama Ahmad Alomari, and Ashraf Elnagar. 2023. Arabic news articles classification using different word embeddings. In *International Conference on Emerging Trends and Applications in Artificial Intelligence*, pages 125–136. Springer.

Asma Sadia Khan, Fariba Tasnia Khan, Tanjim Mahmud, Salman Karim Khan, Nahed Sharmin, Mohammad Shahadat Hossain, and Karl Andersson. 2024. Integrating bert embeddings with svm for prostate cancer prediction. In *2024 6th International Conference on Electrical Engineering and Information & Communication Technology (ICEEICT)*, pages 01–06. IEEE.

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*. Preprint, arXiv:2103.10730.

Pranav Malik, Aditi Aggrawal, and Dinesh Kumar Vishwakarma. 2021. *Toxic speech detection using traditional machine learning models and bert and fasttext embedding with deep neural networks*. 2021 5th International Conference on Computing Methodologies and Communication (ICCMC), pages 1254–1259.

Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. *Efficient estimation of word representations in vector space*. Preprint, arXiv:1301.3781.

Saloni Mittal, Vidula Magdum, Sharayu Hiwarkhedkar, Omkar Dhekane, and Raviraj Joshi. 2023. L3cube-mahanews: News-based short text and long document classification datasets in marathi. In *International Conference on Speech and Language Technologies for Low-resource Languages*, pages 52–63. Springer.

Abu Bakr Mohammad, Kareem Eissa, and Samhaa El-Beltagy. 2017. *Aravec: A set of arabic word embedding models for use in arabic nlp*. *Procedia Computer Science*, 117:256–265.

Hrushikesh Patil, Abhishek Velankar, and Raviraj Joshi. 2022. L3cube-mahahate: A tweet-based marathi hate speech detection dataset and bert models. In *Proceedings of the Third Workshop on Threat, Aggression and Cyberbullying (TRAC 2022)*, pages 1–9.

Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. *Glove: Global vectors for word representation*. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.

Aabha Pingle, Aditya Vyawahare, Isha Joshi, Rahul Tangsali, and Raviraj Joshi. 2023. L3cube-mahasentmd: A multi-domain marathi sentiment analysis dataset and transformer models. *arXiv preprint arXiv:2306.13888*.

G. K. Rajput, Narinder Singh Punn, Sanjay Kumar Sonbhadra, and Sonali Agarwal. 2021. *Hate speech detection using static bert embeddings*. *ArXiv*, abs/2106.15537.

Sello Ralethe. 2020. *Adaptation of deep bidirectional transformers for Afrikaans language*. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 2475–2478, Marseille, France. European Language Resources Association.

Muhammad Umer, Zainab Imtiaz, Muhammad Ahmad, Michele Nappi, Carlo Maria Medaglia, Gyu Sang Choi, and Arif Mehmood. 2022. *Impact of convolutional neural network and fasttext embedding on text classification*. *Multimedia Tools and Applications*, 82:5569–5585.

Laurens van der Maaten and Geoffrey Hinton. 2008. *Visualizing data using t-sne*. *Journal of Machine Learning Research*, 9(86):2579–2605.

Pan Xie, Hengnian Gu, and Dongdai Zhou. 2024. Modeling sentiment analysis for educational texts by combining bert and fasttext. 2024 6th International Conference on Computer Science and Technologies in Education (CSTE), pages 195–199.