Sarcasm Detection in Dravidian Languages using Transformer Models

Madhumitha.M[1,*], Kunguma Akshatra.M[2], Tejashri.J[3], C.Jerin Mahibha[4] and Durairaj Thenmozhi[5]

,2,3,4[1] *Meenakshi Sundararajan Engineering College, Chennai*
[5] *Sri Sivasubramaniya Nadar College of Engineering, Chennai*

Abstract

Sarcasm in Dravidian languages, including Tamil and Malayalam, is a form of linguistic expression where words and phrases are used ironically to convey a meaning opposite to their literal interpretation. It often involves humor, criticism, or mockery, making it a complex and culturally nuanced phenomenon. Dravidian languages' rich grammatical structures and vocabulary provide ample opportunities for subtle and context-dependent sarcasm. Sarcasm in these languages frequently relies on wordplay, idiomatic expressions, and cultural references. Understanding the intended tone and meaning often requires a deep understanding of the cultural context and linguistic subtleties specific to Tamil or Malayalam. Due to limited annotated data and the unique linguistic features of these languages, developing accurate sarcasm detection models poses a significant challenge, making it an intriguing area of research in the field of natural language processing. We participated in the shared task at DravidianCodeMix@FIRE-2023 and have proposed a model that identifies the sarcasm and sentiment polarity of the code-mixed social media comments and posts in Tamil-English and Malayalam-English using the data set shared for the task. We employed a range of transformer models, including IndicBERT, mBERT, and DistilBERT, to effectively detect instances of sarcasm. The macro F1 score obtained by the proposed model for Tamil is 0.82, 0.73, and 0.81 using the IndicBERT, mBERT, and distilBERT models, respectively. The macro-averaged F1 score for Malayalam obtained using the IndicBERT, mBERT, and distilBERT models is 0.66, 0.58, and 0.48, respectively.

Keywords

Tamil and Malayalam sarcasm, Linguistic subtleties, Transformer models, IndicBERT, mBERT, DistilBERT, Sentiment polarity, Humor and criticism

1. Introduction

The advent of social media has transformed the way we connect, communicate, and disseminate information, presenting numerous advantages. However, an escalating concern revolves around its potential adverse effects on mental well-being, particularly in the realm of identifying sarcasm within Dravidian languages[1] on these platforms. The incessant exposure to meticulously curated online identities, coupled with the prevalence of sarcasm in virtual discourse, can induce sentiments of bewilderment and exasperation. The inherent complexities of discerning sarcasm,

Forum for Information Retrieval Evaluation, December 15-18, 2023, India
*Madhumitha.M
madhumithamurthy2002@gmail.com (Madhumitha.M); kungumaakshatra@gmail.com (K. Akshatra.M); tejashrijayachandran@gmail.com (Tejashri.J); jerinmahibha@gmail.com (C.Jerin Mahibha); theni_d@ssn.edu.in (D. Thenmozhi)

especially within languages like Tamil and Malayalam, may lead individuals to misinterpret remarks, needlessly engaging in disputes that heighten stress and emotional strain. Furthermore, the cloak of anonymity offered by social media platforms can intensify the deployment of sarcasm as a tool for mockery or critique, thereby compounding misunderstandings and emotional turmoil [2]. To alleviate the detrimental impact of sarcasm on mental health in the context of Dravidian languages on social media, it becomes imperative to raise awareness and encourage wholesome online behaviors. Promoting lucid and courteous communication, imparting insights into the subtleties of sarcasm to users, and nurturing a supportive and empathetic digital community can contribute to a more positive and emotionally secure online sphere. Additionally, the development of automated tools for detecting sarcasm in Dravidian languages can play a pivotal role in promptly identifying and addressing sarcastic comments, thereby reducing the potential harm arising from misinterpretations and hurtful interactions [3]. In an ever-evolving social media landscape, recognizing the paramount importance of mental well-being in online interactions obliges us to take proactive measures to establish a more inclusive and emotionally nurturing digital environment that caters to individuals from diverse linguistic backgrounds. The shared task of identifying Sarcasm from social media text considering Dravidian languages was a part of FIRE-2023 [4], which is based on Tamil-English and Malayalam-English comments. The challenge of identifying sarcasm from social media comments is a binary classification task, wherein the model's objective is to classify the text into one of two categories: sarcasm or non-sarcasm.For example,

- ”Yuvan haters kallu uppu eduthuttu varisayil varavum Yuvaniyans waiting Thaaa puthasalum puthayala varuvanda” represents Non-Sarcastic in Tamil.;
- Kickin it old school ...apdiye konjam Patti tinkering panni..namba ooru masala va.. mixed... Different attempt dhan new illa.” - Sarcasm - Tamil
- ”Enthinaa chettaa ingane thala vetti vekkunnath . Mammukkante thala vallavanteyum kazhuthinu mukalil vekkumbol , pani ariyaavunnorkk koduthaal pore” - Non-Sarcasm - Malayalam Malayalam;

2. Related Works

The use of IndicBERT for the process of detecting sarcasm from social media text has been proposed by Amir et al. [5]. The model had exhibited strong performance by effectively capturing contextual information and sarcasm cues within the text. The IndicBERT had been found to accurately capture language nuances and characteristics specific to Indian languages [6]. A soft-attention-based bi-directional LSTM in combination with a feature-rich CNN approach has been employed for sarcasm detection [7]. Hate speech from Malayalam social media content had been identified using the MBERT transformer model [8], and offensive language from Tamil posts had been identified using various machine learning and deep learning models [9]. Sarcasm detection in social media has been explored using the mBERT model in combination with a Graph Convolution Network (GCN) [10], [11]. Agrawal et al. [12] delve into the intriguing domain of leveraging transitions of emotions to enhance sarcasm detection, shedding light on the dynamic nature of emotional cues in sarcasm identification. In contrast, Plepi (2021) focuses on perceived and intended sarcasm detection using graph attention networks, emphasizing the significance

Table 1
Data Distribution for Tamil

Category	Training Dataset	Evaluation Dataset
Sarcastic	7170	1820
Non-Sarcastic	19866	4939

of graph-based attention mechanisms in discerning the nuances of sarcasm.The use of the mBERT model and supervised learning techniques for textual analysis to effectively identify sarcasm in news headlines has been proposed by Jayaraman (2022). Two notable studies employ BERT-based transfer learning techniques to enhance model performance. Transfer learning with BERT has been used for sarcasm detection [13]. On the other hand, Baruah 2020 Context focuses on context-aware sarcasm detection using BERT, emphasizing the role of context in refining sarcasm identification. The use of cross-lingual embeddings for sentiment analysis from Tamil text has been proposed by Mahibha (2021). In the domain of sarcasm detection, Pandey and Vishwakarma [14] addressed the challenge of multimodal sarcasm detection (MSD) in videos, utilizing deep learning models. Their research focused on leveraging various modalities, including visual, audio, and textual cues, to effectively identify sarcasm in video content [15]. The Gaussian model has also been used to discern and classify different types of irony within textual content [16] using the dataset provided by SemEval-2022. In the context of sarcasm detection, Wicana et al. [17] conducted a comprehensive review, exploring various machine learning approaches for sarcasm detection. This review provides valuable insights into the state-of-the-art techniques and challenges associated with detecting sarcasm in textual content. A pragmatic and intelligent model for sarcasm detection in social media text has also been proposed [18]. Both Ren (2018) and Razali (2021) explored the integration of contextual information and deep learning techniques. The paper[19] explores enhancing social media by promoting positive expression. It introduces a custom deep network using T5-sentence embeddings and various machine learning models, achieving superior results in detecting positivity across English, Tamil, and Malayalam. The paper[20] addresses homophobia and transphobia on social media, particularly targeting the LGBTQIA+ community. It explores improving abusive speech detection by using data augmentation techniques, specifically pseudolabeling and transliteration, with multilingual language models.

3. Data set

The data sets for Tamil and Malayalam that are used to implement sarcasm detection were the training, evaluation, and test datasets that were provided by the organizers of the shared task. Each instance of the training dataset had a label specifying whether the text was sarcastic or non-sarcastic.

The data distribution of the training and development dataset for Tamil is shown in Table 1, and for Malayalam, it is shown in Table 2. In the case of Tamil, the training dataset comprised 27,036 instances, with 7170 falling under the sarcastic category and 19866 under the non-sarcastic category. Similarly, the development dataset for Tamil contained 6759 instances,

Table 2
Data Distribution for Malayalam

Category	Training Dataset	Evaluation Dataset
Sarcastic	2259	588
Non-Sarcastic	9798	2527

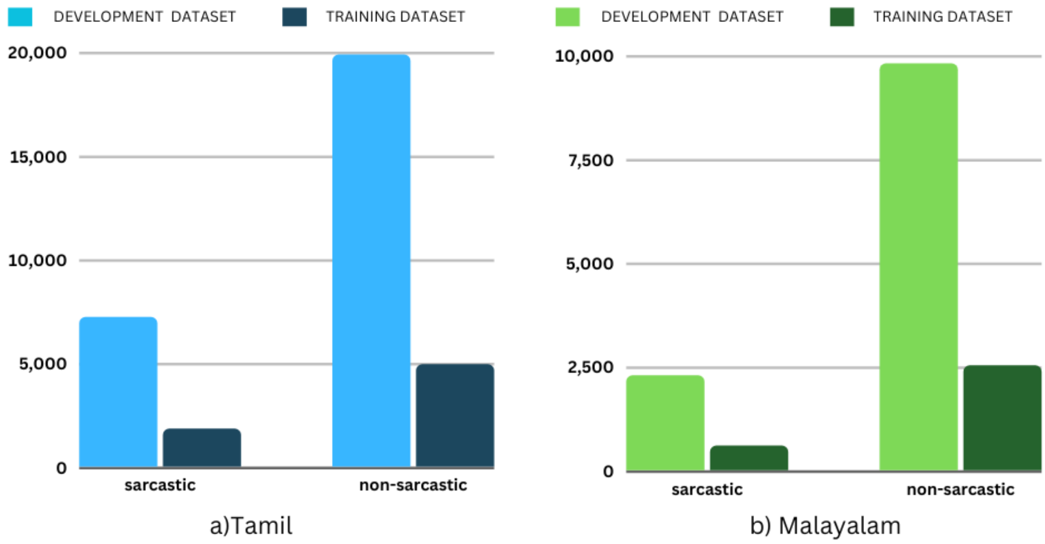

Figure 1: Data distribution

including 1820 in the sarcastic category and the rest in the non-sarcastic category, highlighting the data's imbalance. The test dataset for Tamil consisted of 8449 instances, which were used for evaluating the model's predictions.For Malayalam, the training dataset consisted of 12,051 instances, with 2259 categorized as sarcastic and 9798 as non-sarcastic. The development dataset for Malayalam had 3015 instances, including 588 in the sarcastic category and the remaining in the non-sarcastic category, indicating a similar data imbalance. Lastly, the test dataset for Malayalam encompassed 3768 instances, utilized for assessing the model's predictive performance.

4. System Description

In our approach, we employed a series of processing techniques to address sarcasm detection in Dravidian languages, specifically Tamil and Malayalam. These techniques encompassed various steps, starting with data preprocessing to clean the training dataset. In this stage, we removed unwanted characters, digits, and whitespace from the textual data. Tokenization was then applied to segment the text into subword tokens, and encoding methods were used to prepare the data for model input. Subsequently, we fine-tuned three pre-trained transformer models: IndicBERT, mBERT, and DistilBERT, to adapt them to the task of sarcasm detection in

Dravidian languages. These models were essential components of our approach.

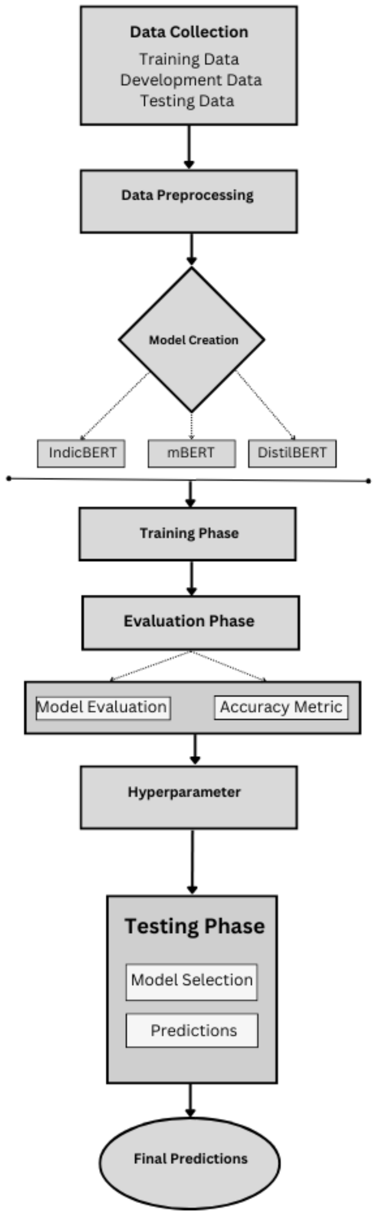

Figure 2: Proposed Architecture

The proposed architecture is represented in Figure 2. The initial phase of our process involved gathering three datasets provided by the task organizers: the training dataset, development dataset, and testing dataset. The training dataset underwent preprocessing, as previously described, resulting in clean and structured data. This prepared dataset was used to train the three models, IndicBERT, mBERT, and DistilBERT. Moving forward, during the training phase, we utilized the preprocessed dataset and the three models. Each of these models was evaluated using the development dataset. The model that exhibited the highest accuracy on this dataset, which happened to be the IndicBERT model in our case, was selected as the final model for submission. This model was then used to make predictions for the testing dataset. The entire process, from data preprocessing to model building and evaluation, is illustrated in the proposed architecture. This architecture underscores the significance of data preprocessing, model training, and fine-tuning in enhancing the accuracy of our sarcasm detection system. During testing, the selected model generated contextual embeddings and predicted labels for the text in the dataset, playing a crucial role in identifying sarcasm.

4.1. IndicBERT Model

IndicBERT [21] is a powerful transformer-based language model that has been purposefully developed to address the intricate linguistic intricacies found in Indic languages. Drawing inspiration from the foundational principles of BERT, IndicBERT follows a similar pre-training approach but incorporates tailored strategies to account for the unique linguistic features of this language group. It leverages fine-tuning techniques and custom tokenization methods to enhance its performance across a diverse range of Indian languages. In the process of fine-tuning IndicBERT, a carefully assembled dataset containing instances of sarcasm in Tamil and Malayalam is meticulously prepared, incorporating annotated data. This dataset undergoes tokenization following the WordPiece scheme specific to IndicBERT, and it is subsequently partitioned into distinct subsets for training, validation, and testing. The fine-tuning phase is executed with predefined hyperparameters and employs a cross-entropy loss function. A comprehensive evaluation is then conducted, utilizing precision, recall, and F1-score metrics to assess the model's proficiency in detecting sarcasm within the contexts of Tamil and Malayalam.

The ”IndicBERT” model is designed with multiple transformer layers, affording it substantial expressive capacity with 768 hidden units and multiple attention heads. This design allows the model to excel in capturing the nuanced aspects of Indian languages while simultaneously optimizing efficiency in both the training and deployment phases. The overall architecture of IndicBERT is engineered with computational efficiency in mind, making it more accessible and facilitating faster training and inference processes. In conclusion, IndicBERT serves as a foundational tool for empowering natural language processing applications in Indic languages, offering robust linguistic understanding without compromising computational efficiency.

4.2. mBERT Model

M-BERT (Multilingual BERT) is a versatile transformer-based language model designed to address multilingual natural language processing challenges. Inspired by BERT's success, M-BERT employs a similar pre-training approach, focusing on effective language processing across multiple languages. It excels in cross-lingual tasks and multilingual applications.The ”c” model boasts a comprehensive architecture, featuring multiple transformer layers with 768 hidden units and various attention heads. This design enables M-BERT to capture linguistic nuances in a wide range of languages while maintaining computational efficiency, making it a valuable tool for multilingual applications.

For fine-tuning M-BERT for Tamil and Malayalam sarcasm detection, a dataset containing labeled examples is prepared, adhering to M-BERT's multilingual approach. The dataset is segmented for training, validation, and testing, and specific hyperparameters govern the fine-tuning process. Ongoing validation checks ensure model performance, with evaluation metrics like precision, recall, and F1-score employed to assess the model's effectiveness in identifying sarcasm within these languages.

4.3. DistilBERT Model

DistilBERT [22], short for Distilled BERT, DistilBERT emerges as a formidable transformer-based language model, offering a streamlined alternative to BERT with a primary emphasis on efficiency and scalability. Drawing inspiration from BERT's pre-training methodology, DistilBERT incorporates inventive techniques to significantly reduce its model size and computational requirements. Despite its more compact architecture, DistilBERT excels in various natural language processing tasks. The "DistilBERT" model is meticulously designed with fewer transformer layers, effectively distilling the complexity of the original BERT. It typically features a reduced number of hidden units and attention heads while retaining crucial linguistic comprehension and capabilities. This reduction in model parameters renders DistilBERT remarkably lightweight and computationally efficient when compared to BERT. The advantages of DistilBERT extend to faster training and inference times, rendering it an appealing choice for a wide range of NLP applications, particularly in scenarios where computational resources are constrained. In essence, DistilBERT represents a distilled yet highly effective iteration of BERT, providing efficient and scalable solutions for natural language understanding and processing.

In the context of DistilBERT fine-tuning, a dataset containing examples of sarcasm in Tamil and Malayalam is collected and tokenized according to DistilBERT's subword scheme. Subsequently, the dataset is partitioned into separate sets for training, validation, and testing. Hyperparameters play a crucial role in guiding the fine-tuning process, with ongoing validation assessments to monitor the model's performance. Evaluation metrics, including precision, recall, and F1-score, are employed to gauge DistilBERT's effectiveness in identifying sarcasm within the context of Tamil and Malayalam. This underscores the adaptability of DistilBERT for specific NLP tasks, such as sarcasm detection in Dravidian languages.

5. Result

The metrics that were studied for the evaluation of the task were the macro-F1 score. The F1 score is an overall measure of a model's accuracy that merges precision and recall. An extreme F1 score means that the classification has happened, accompanied by a reduced number of false positives and low false negatives.Within a classification report, the initial focus is on assessing a binary classification model for two distinct categories, typically labeled as 0 and 1. For class 0, precision provides insights into how accurately our model predicts instances as class 0, while recall sheds light on how effectively it identifies actual instances belonging to class 0. Conversely, concerning class 1, precision aids in evaluating the accuracy of our

Table 3
Performance score for Tamil

Model	F1-Score	Accuracy
IndicBERT	0.70	0.79
mBERT	0.69	0.77
DistilBERT	0.68	0.77

Table 4
Performance score for Malayalam

Model	F1-Score	Accuracy
IndicBERT	0.68	0.84
mBERT	0.45	0.80
DistilBERT	0.49	0.81

predictions for this specific category, while recall delves into the model's ability to pinpoint genuine class 1 instances. These fundamental metrics play a pivotal role in comprehending the model's effectiveness for each class within a binary classification scenario. However, it's crucial to note that in the classification report for class 1 in Malayalam, both M-BERT and Distil-BERT exhibit precision, recall, and F1 score values of 0. This signifies that neither model successfully recognized any instances of this class. This situation underscores a significant issue in the models' performance in class 1 classification, necessitating a thorough examination of both the data and the model's behavior to address this challenge adequately. The performance metrics, particularly the F1 score and accuracy of development dataset, for various models in the case of Tamil are presented in Table 3, while for Malayalam, they are displayed in Table 4.mBERT's lower performance in Tamil sarcasm detection can be attributed to its multilingual training and data imbalance, whereas IndicBERT's language-specific fine-tuning and improved pretraining for Indic languages likely enhance its higher F1 score. Notably, it becomes evident that the IndicBERT model has exhibited superior performance compared to the other models employed in the evaluation. The evaluation of the tasks primarily relied on the macro-F1 score achieved by our proposed model. It is worth emphasizing that our proposed model yielded remarkable results, securing a macro F1 score of 0.70 for Tamil and 0.68 for Malayalam, forming the basis for task evaluation. Furthermore, the accuracy metrics obtained were also notable, with a recorded accuracy of 0.79 for Tamil and 0.84 for Malayalam. These achievements culminated in securing the 11th position in the Tamil language category and the 7th position in the Malayalam language category on the leaderboard. These rankings underscore the effectiveness of our model in the Dravidian CodeMix@FIRE-2023 shared task, highlighting its prowess in sarcasm detection within code-mixed Tamil-English and Malayalam-English social media comments and posts.

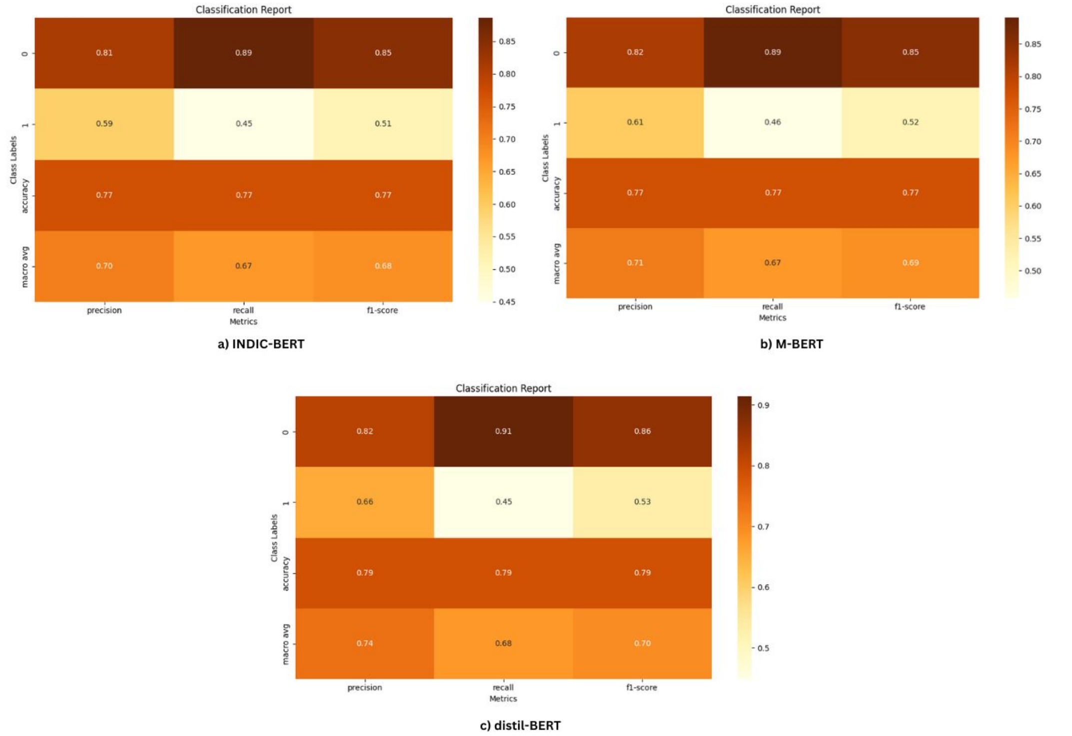

Figure 3: Classification report - Tamil

6. Error Analysis

The F1 score obtained for the task using the proposed Indic-BERT model shows that more false positive and false negative classifications have occurred. One reason for this could be considered the data imbalance nature of the dataset. The classification report of the proposed system considering the languages Tamil and Malayalam is represented in figures 3 and 4, respectively. Considering the number of instances for the class labeled severe is higher, and the F1 score, precision, and recall associated with this class are high when compared to the class, it is not sarcastic. This represents that the number of misclassifications increases when the number of instances for training is lower, which is associated with data imbalance. Data augmentation could be considered to improve the model's performance. Examples of Tamil texts that are misclassified due to the above reasons are shown in Table 5. Considering the first text of the table, it has the specific sarcasm marker ”Marana mass” and is classified as sarcastic instead of the correct label of non-sarcastic. The third text of the table does not have any specific sarcasm marker, ”fucking,” and is classified as non-sarcastic instead of the correct label of sarcastic. The table also shows texts that are sarcastic and misclassified. All the example texts show that sarcasm markers and sarcasm play a major role in the process of classification and identifying whether the text is associated with sarcasm. The binary classification of the

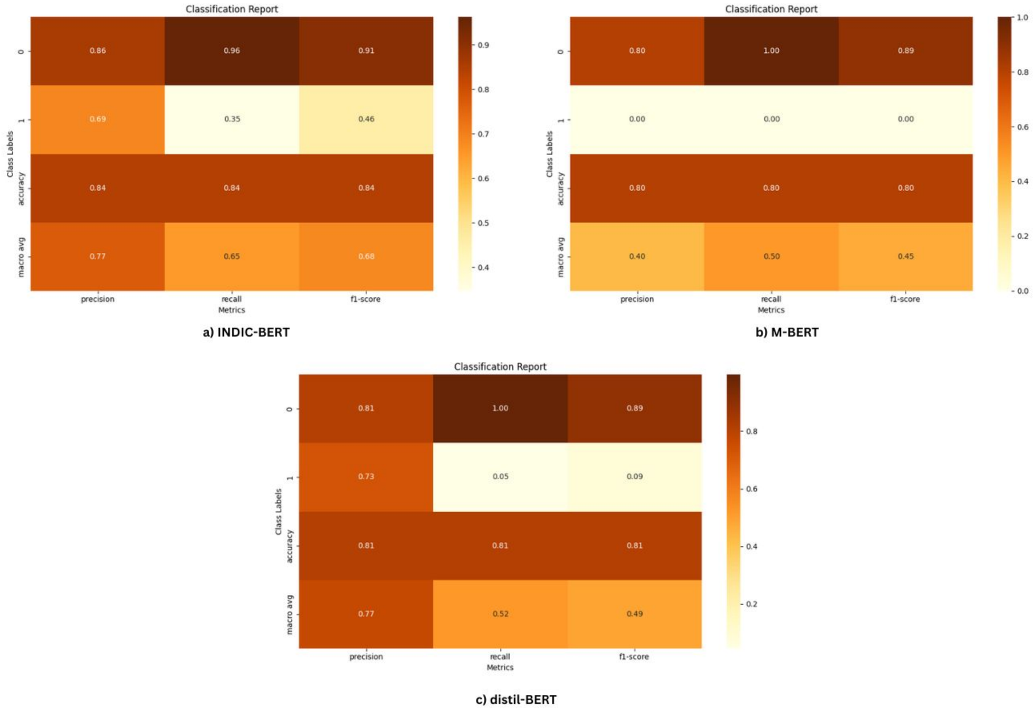

Figure 4: Classification report - Malayalam

task is implemented using IndicBERT, mBERT, and distillBERT transformer models for training and testing the dataset. The task was implemented considering the languages Tamil-English and Malayalam-English.

7. Conclusion

The endeavor to detect sarcasm in Dravidian languages has gained considerable importance due to its interconnectedness with various application domains. Recognizing this significance, FIRE-2023 introduced a pivotal task focused on sarcasm detection within Dravidian languages, classifying text into either sarcastic or non-sarcastic categories. The exploration of this task, conducted under the banner of Dravidian CodeMix@FIRE-2023, underscores the value of harnessing advanced natural language processing techniques for the analysis of mental health aspects. In our pursuit, we employed state-of-the-art language models, namely IndicBERT, mBERT, and DistilBERT, to tackle the challenge of detecting signs of sarcasm within social media text. This approach allowed us to categorize social media content sarcastic and non-sarcastic category, thereby providing a more comprehensive insight into the mental states of individuals. Such granularity in classification empowers us with the potential for tailored interventions and support mechanisms.

Table 5
Error Analysis

S.No.	Text	Model	Predicted Label	Actual Label
01.	Ngk trailer available Marana mass Surya anan	Indic-BERT	Sarcastic	Non-Sarcastic
02.	Eda kazhiv kettavan mare view koottan motham kallatharavumayi erangi alle koothara teaserum olekkada s 3 yum	Indic-BERT	Non-Sarcastic	Sarcastic
03.	Hit like if you think Rajinikanth is a fucking joke	m-BERT	Non-Sarcastic	Sarcastic
04.	Maranaaaaa maaasssssss trailer in the history	m-BERT	Sarcastic	Non-Sarcastic
05.	August 8 tamilnadu fulla terika vidrom.....	Distil-BERT	Sarcastic	Non-Sarcastic
06.	Iethuku yethuku da naadodigal 2 nu vachinga	Distil-BERT	Non-Sarcastic	Sarcastic

The outcomes derived from our utilization of the IndicBERT, mBERT, and DistilBERT models contribute significantly to the burgeoning field of sarcasm detection in social media discourse. These models have exhibited remarkable effectiveness in capturing subtle linguistic cues and contextual information indicative of sarcasm. Looking ahead, we envision opportunities for further enhancement in this domain. One avenue for improvement involves the incorporation of more extensive contextual information, which could enhance the accuracy of sarcasm detection. Additionally, the adoption of hybrid approaches, leveraging the strengths of various deep learning models, may further bolster the efficiency and precision of detecting sarcasm and related mental health markers within text data. This ongoing research trajectory underscores our commitment to advancing the understanding and detection of sarcasm, thereby contributing to the broader field of mental health analysis.

References

[1] P. E. Hook, O. N. Koul, A dedicated sarcasm construction in kashmiri as a feature of the south asian linguistic area, Bhas· (2022).

[2] A. Verghese, Irony's Subversive Bite: Destabilising Colonial Legacies in Contemporary Art, Ph.D. thesis, ResearchSpace@ Auckland, 2022.

[3] W. Hariri, Unlocking the potential of chatgpt: A comprehensive exploration of its applications, advantages, limitations, and future directions in natural language processing, arXiv preprint arXiv:2304.02017 (2023).

[4] B. R. Chakravarthi, N. Sripriya, B. Bharathi, K. Nandhini, S. Chinnaudayar Navaneethakrishnan, T. Durairaj, R. Ponnusamy, P. K. Kumaresan, K. K. Ponnusamy, C. Rajkumar, Overview of the shared task on sarcasm identification of Dravidian languages (Malayalam and Tamil) in DravidianCodeMix, in: Forum of Information Retrieval and Evaluation FIRE - 2023, 2023.

[5] S. Amir, B. C. Wallace, H. Lyu, P. C. M. J. Silva, Modelling context with user embeddings for sarcasm detection in social media, arXiv preprint arXiv:1607.00976 (2016).

[6] K. Jain, A. Deshpande, K. Shridhar, F. Laumann, A. Dash, Indic-transformers: An analysis of transformer language models for indian languages, arXiv preprint arXiv:2011.02323 (2020).

[7] D. Jain, A. Kumar, G. Garg, Sarcasm detection in mash-up language using soft-attention based bi-directional lstm and feature-rich cnn, Applied Soft Computing 91 (2020).

[8] P. Poornachandran, V. Sujadevi, G. Rajendran, V. Ks, V. Vijayan, A. Ram, et al., Malhate: Hate speech detection in malayalam regional language, in: 2022 IEEE 7th International Conference on Recent Advances and Innovations in Engineering (ICRAIE), volume 7, IEEE, 2022, pp. 110–115.

[9] J. Mahibha, S. Kayalvizhi, D. Thenmozhi, Offensive language identification using machine learning and deep learning techniques (2021).

[10] C.-T. Phan, Q.-N. Nguyen, C.-T. Dang, T.-H. Do, K. Van Nguyen, Vicgcn: Graph convolutional network with contextualized language models for social media mining in vietnamese, arXiv preprint arXiv:2309.02902 (2023).

[11] M. Marreddy, S. R. Oota, L. S. Vakada, V. C. Chinni, R. Mamidi, Multi-task text classification using graph convolutional networks for large-scale low resource language, in: 2022 International Joint Conference on Neural Networks (IJCNN), IEEE, 2022, pp. 1–8.

[12] A. Agrawal, A. An, M. Papagelis, Leveraging transitions of emotions for sarcasm detection, in: Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, 2020, pp. 1505–1508.

[13] E. Savini, C. Caragea, Intermediate-task transfer learning with bert for sarcasm detection, Mathematics 10 (2022).

[14] A. Pandey, D. K. Vishwakarma, Multimodal sarcasm detection (msd) in videos using deep learning models, in: 2023 International Conference in Advances in Power, Signal, and Information Technology (APSIT), IEEE, 2023, pp. 811–814.

[15] S. Sangwan, M. S. Akhtar, P. Behera, A. Ekbal, I didn't mean what i wrote! exploring multimodality for sarcasm detection, in: 2020 International Joint Conference on Neural Networks (IJCNN), IEEE, 2020, pp. 1–8.

[16] D. Krishnan, T. Durairaj, et al., Getsmartmsec at semeval-2022 task 6: Sarcasm detection using contextual word embedding with gaussian model for irony type identification, in: Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022), 2022, pp. 827–833.

[17] S. G. Wicana, T. Y. İbisoglu, U. Yavanoglu, A review on sarcasm detection from machine-learning perspective, in: 2017 IEEE 11th International Conference on Semantic Computing (ICSC), IEEE, 2017, pp. 469–476.

[18] M. Shrivastava, S. Kumar, A pragmatic and intelligent model for sarcasm detection in social media text, Technology in Society 64 (2021).

[19] B. R. Chakravarthi, Hope speech detection in youtube comments, Social Network Analysis and Mining 12 (2022).

[20] B. R. Chakravarthi, A. Hande, R. Ponnusamy, P. K. Kumaresan, R. Priyadharshini, How can we detect homophobia and transphobia? experiments in a multilingual code-mixed setting for social media governance, International Journal of Information Management

Data Insights 2 (2022).
[21] D. Kakwani, A. Kunchukuttan, S. Golla, G. N.C., A. Bhattacharyya, M. M. Khapra, P. Kumar, IndicNLPSuite: Monolingual Corpora, Evaluation Benchmarks and Pre-trained Multilingual Language Models for Indian Languages, in: Findings of EMNLP, 2020.
[22] V. Sanh, L. Debut, J. Chaumond, T. Wolf, Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter, arXiv preprint arXiv:1910.01108 (2019).