# Team Oryu@DravidianLangTech 2026: A Multilingual Transformer Approach for Hope Speech Detection in Code-Mixed Tulu

**Joyeta Barua Moni, Noore Tamanna Orny, Md. Abtahee Kabir, Hasan Murad**
Department of Computer Science and Engineering,
Chittagong University of Engineering and Technology, Bangladesh
{u2104130, u2104092, u2104089}@student.cuet.ac.bd, hasanmurad@cuet.ac.bd

## Abstract

Hope speech detection appears to have an essential role to play in fostering positive and inclusive communication on social media, especially in low-resource multilingual settings. This paper describes the system submitted by Team Oryu for Task 1: Coarse-Grained Hope Tone Classification in Code-Mixed Tulu. The task involves classifying comments in social media texts into one of the four classes: Encouraging, Discouraging, Uninvolved, and Blended Tone. The texts in this task show heavy code-mixing between Tulu, English, and Kannada. In order to overcome this challenge, we employed a fine-tuned multilingual transformer model, code-mixed text processing, data augmentation, and class-weighted loss to handle class imbalance. Our proposed system achieved a Macro F1-score of 63%, securing 3rd position on the shared task. The results demonstrate the efficacy of multilingual transformer models in emotionally nuanced classification in code-mixed environments while underscoring the difficulties in capturing blended emotional tones.

## 1 Introduction

Hope speech detection has emerged as a new direction of research in Natural Language Processing (NLP), especially with respect to creating a positive online discourse and promoting mental well-being. Although a significant amount of research has focused on detecting negative content such as hate speech, abusive language, and offensive expressions, comparatively little work has been done on identifying positive and supportive communication. Hope speech detection aims to identify language that conveys encouragement, optimism, and empathy.

The shared task presents the first benchmark for hope speech detection in code-mixed Tulu. Tulu, a low-resource Dravidian language mainly spoken in the coastal districts of Karnataka and Kerala, has received limited attention in NLP research. On social media platforms such as YouTube, Tulu speakers frequently use code-mixed text consisting of Tulu, English, and Kannada, often written in Roman script.

The first task of the shared task focuses on the classification of coarse-grained hope tones, where each comment is categorized into one of four classes: *Encouraging*, *Discouraging*, *Uninvolved*, and *Blended Tone*. This task differs from conventional sentiment classification as it requires a more nuanced understanding of emotional tone.

To address this challenge, we propose a transformer-based approach for coarse-grained hope tone classification of code-mixed Tulu social media comments. Our method leverages multilingual transformer fine-tuning using the pre-trained *XLM-RoBERTa* model along with targeted preprocessing techniques designed for code-mixed text. Furthermore, we address class imbalance in the dataset using class-weighted loss and apply token-level augmentation to improve model robustness. In addition, we demonstrate the effectiveness of multilingual transformer models specifically for code-mixed Tulu, a highly underexplored and low-resource setting. These strategies help the system capture emotional and contextual information present in informal multilingual user-generated content.

Our major contributions can be summarized as follows:

- We present a multilingual transformer-based model for coarse-grained hope tone classification in code-mixed Tulu social media text.
- We introduce a preprocessing scheme that addresses the challenges of code-mixed data while preserving critical emotional cues such as emojis and colloquial expressions.



*Proceedings of the Sixth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*, pages 403–408

- We incorporate class balancing and data augmentation techniques to improve macro-level performance.

## 2 Related Work

The major research in the field of Natural Language Processing (NLP) has been carried out to detect harmful content on the internet, including hate speech, abusive language, and offensive language. On the other hand, less research has been carried out on detecting constructive and positive content, including hope speech. Hope speech is defined as positive and encouraging content on the Internet. Chakravarthi proposed the concept of Multilingual Hope Speech Detection and the HopeEDI dataset, focusing on the promotion of Equality, Diversity, and Inclusion (EDI) using positive speech on the Internet in English and different Dravidian languages (Chakravarthi, 2022b).

Further research has also been conducted to extend the hope speech detection to social media environments, such as YouTube comments. For this purpose, deep learning techniques, such as CNN-based models, have been used for the classification of hope speech for different languages. The experiments have been conducted to prove the effectiveness of the neural network models for the detection of positive and supportive speech in social media environments (Chakravarthi, 2022a). The importance of the increasing need for analyzing positive speech, as well as negative speech, has also been highlighted.

Besides this, recent research has been carried out on the detection of hope speech in multilingual and code-mixed languages. A multilingual approach using transformer models such as XLM-RoBERTa and mBERT has been proposed for detecting hope speech in low-resource languages (Abdullah et al., 2025). The study showed that transformer-based architectures can effectively capture contextual information for hope speech classification.

Similarly, hope speech detection in code-mixed Roman Urdu has been explored by comparing traditional machine learning models with transformer-based techniques (Ahmad et al., 2025). The results indicate that transformer-based approaches outperform traditional machine learning methods when dealing with informal and orthographically inconsistent social media text.

Significant contributions have also been made in research on code-mixed Dravidian languages. A trilingual code-mixed Tulu dataset for sentiment analysis was introduced to better understand multilingual social media data (Hegde et al., 2022). In addition, a shared task on sentiment analysis for Tamil and Tulu code-mixed text highlighted the challenges of processing multilingual content written in Roman script (Hegde et al., 2023). Code-mixed text often contains multiple languages within the same sentence, making analysis more complex. Related research has also demonstrated the effectiveness of machine learning and deep learning approaches for sentiment analysis in Dravidian code-mixed languages, emphasizing the importance of contextual representation learning (Mishra et al., 2021).

## 3 Data Description

The dataset employed for this shared task consists of code-mixed Tulu social media comments collected from YouTube platforms (Thenmozhi et al., 2026). This is a natural dataset where Tulu is often mixed with English and Kannada, mostly written in the Roman script. Due to the informal nature of social media communication, the dataset contains spelling variations, irregular orthography, elongated characters, emojis, and grammatical errors. The label-wise distribution of the dataset is shown in Table 1.

The dataset is divided into three predefined splits: training, development, and test sets. Each instance contains a comment and a corresponding coarse-grained hope tone label. The task is formulated as a four-class classification problem with the labels *Encouraging*, *Discouraging*, *Uninvolved* and *Blended Tone*.

The class distribution is imbalanced, with the *Uninvolved* category containing more instances than the other classes. Therefore, Macro F1-score is used as the primary evaluation metric to assess model performance across all classes.

| Label | Training Dataset | Development Dataset | Test Dataset |
|---|---|---|---|
| **Blended Tone** | 895 | 191 | 172 |
| **Discouraging** | 711 | 153 | 149 |
| **Encouraging** | 1895 | 406 | 407 |
| **Uninvolved** | 2490 | 534 | 398 |
| **Total** | **5991** | **1284** | **1126** |

Table 1: Label-wise Distribution of Training, Development, and Test Data

## 4 Methodology

The proposed system uses a transformer model as a classification paradigm to detect coarse-grained hope tone in code-mixed Tulu social media text. We have also provided the implementation of our system in a publicly available GitHub repository[1]. The overall workflow of the system includes preprocessing the noisy text, applying a light data augmentation technique, fine-tuning a multilingual transformer model, and addressing class imbalance to enhance macro-level performance. The system architecture is designed to effectively capture emotional content in informal and code-mixed text.

### 4.1 Task and Data Description

The shared task involves coarse-grained hope tone classification of social media comments written in code-mixed Tulu. Each comment is categorized into one of four classes: *Encouraging*, *Discouraging*, *Uninvolved*, and *Blended Tone*. The dataset consists of highly informal text collected from social media platforms such as YouTube, where code-mixing between Tulu, English, and Kannada is prevalent. The dataset contains spelling variations, Romanized script usage, emojis, and grammatical inconsistencies, making the classification task challenging. The organizers provided the dataset pre-split into training, development, and test sets. The final ranking of the systems was determined based on performance on the hidden test set.

### 4.2 Data Preprocessing

Because of the noisy and informal nature of code-mixed social media text, preprocessing was performed to improve text quality while preserving emotional information. URLs and unnecessary whitespace were removed. Repeated character sequences were normalized to reduce noise caused by elongated informal expressions. Unlike aggressive normalization techniques, emojis and other emotional indicators were retained as they carry important information useful for identifying hope tones. The text data was tokenized using the *XLM-RoBERTa tokenizer*. Padding and truncation were applied to ensure that all sequences had a uniform length of 128 tokens.

### 4.3 Data Augmentation

A random token deletion technique was employed to improve the robustness of the models trained on the smaller dataset. In this approach, tokens are randomly removed from the input text with a probability of 0.1, which results in the generation of data augmentation versions of the original input data.

### 4.4 Classification Model

For the classification of texts into hope tone, the pre-trained multilingual transformer model *XLM-RoBERTa-base* (Conneau et al., 2019) was fine-tuned. The model was selected due to its strong cross-lingual representation capability and effectiveness in handling code-mixed text classification tasks. A task-specific classification head was added on top of the pre-trained model. Fine-tuning was performed using the AdamW optimizer within the Hugging Face `Trainer` framework. The model was trained with a learning rate of $2 \times 10^{-5}$, a batch size of 16, and for 3 epochs.

### 4.5 Handling Class Imbalance

The dataset exhibits class imbalance across the four hope tone categories. To mitigate bias toward majority classes, class weights were computed from the training data distribution and incorporated into the cross-entropy loss function (Mao et al., 2023). This strategy improves macro-level performance, particularly the Macro F1-score, by enhancing prediction quality for minority classes.

### 4.6 Training Strategy

To ensure reproducibility and stable convergence, seed values were set in NumPy, PyTorch, and CUDA. The general flow of the model pipeline, starting from data preprocessing, to training and predicting, is described in Figure 1 below. The model was trained with the help of the Hugging Face `Trainer` class, evaluated on the validation set. The training was done on a GPU-enabled system.

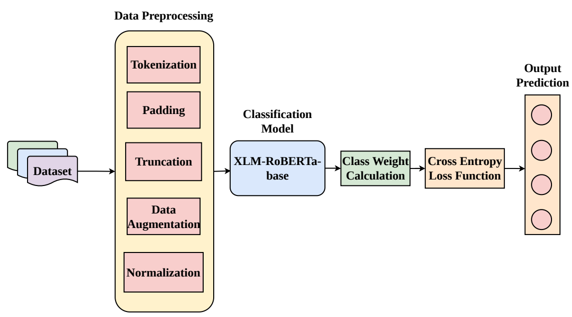


Figure 1: Proposed Transformer-based Hope Speech Detection Architecture

[1]The code repository is available at: GitHub Repository.

## 5 Results and Analysis

### 5.1 Leaderboard Performance

Team Oryu has achieved 3rd position in the task of coarse-grained hope tone classification in the shared task. For our submission, the system has achieved an Accuracy of 67%, Macro Precision of 64%, Macro Recall of 65%, and Macro F1-score of 63%. Considering the complexity of the task and the low-resource nature of the Tulu language, the proposed system has achieved a Macro F1-score higher than 60%, demonstrating the effectiveness of the proposed approach for code-mixed hope tone classification using a multilingual transformer model.

### 5.2 Performance Analysis

The multilingual transformer backbone model, XLM-RoBERTa, demonstrated strong performance in processing code-mixed Tulu–English–Kannada text by capturing cross-lingual dependencies and subtle emotional cues. The overall performance across evaluation metrics is shown in Table 2. The use of class weighting helped address class imbalance, while the lightweight data augmentation strategy improved generalization and overall macro-level performance.

| Metric | Value (%) |
|---|---|
| Accuracy | 67 |
| Macro Precision | 64 |
| Macro Recall | 65 |
| Macro F1-score | 63 |

Table 2: Performance Metrics of the Proposed System

### 5.3 Ablation Study

To evaluate the contribution of different components, we compare the baseline model with class balancing and the proposed system.

| Model | Macro F1-score |
|---|---|
| XLM-RoBERTa-base | 0.5554 |
| XLM-RoBERTa-base + WeightedSampler | 0.6220 |
| Proposed Model | 0.6276 |

Table 3: Comparison of baseline and proposed model performance

The results in the Table 3 show that class balancing significantly improves performance over the baseline, while the proposed model achieves the best performance.

### 5.4 Error Analysis & Discussion

Despite the fact that the predictive capabilities of the model with regard to accuracy and F1-scores were impressive, several challenging cases remained, as reflected in the confusion matrix in Figure 2. One of them was *Blended Tone*, which, due to the presence of emotions in one statement, led to either *Encouraging* or *Discouraging* classification. In addition, there was the possibility of misclassification of Discouraging as either Blended Tone or Encouraging. For instance, comments where encouraging messages were combined with sarcastic remarks caused errors in classification, which means that the model could not cope with such nuanced changes of emotions in the text. Also, lack of context, semantics of code-mixed language and alternative spellings in Roman alphabets led to errors. Some Uninvolved comments could even be confused with Blended Tone. The results demonstrate the effectiveness of multilingual transformer models for the detection of emotional signals in code-mixing languages. While our system is outperformed by the best-performing model (Macro-F1: 58%), it achieves comparable performance, despite the room left for improvement in terms of the difference in performance metrics. More advanced approaches to increasing the accuracy of emotion recognition include applying focal loss and ensemble learning with more extensive transformer models. It is also possible to increase the performance of the model through improved classification of ambiguous categories like the *Blended Tone*. Further research might involve exploring other forms of data augmentation apart from the existing one.

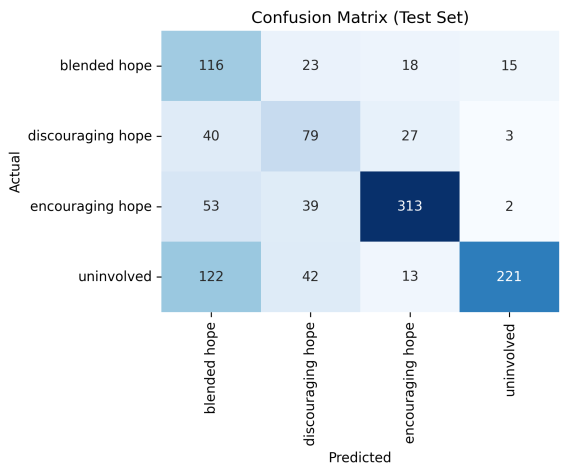


Figure 2: Confusion Matrix on Labeled Test Set

## 6 Conclusion

In this paper, we describe the system developed by Team Oryu for the shared task on Hope Speech Detection in Code-Mixed Tulu. The system employed multilingual transformer fine-tuning, preprocessing, handling of class imbalance, and data augmentation techniques. The system achieved a Macro F1-score of 63% on the test data, which enabled us to secure the 3rd place. The results show the promise of multilingual pre-trained models in detecting emotional cues in code-mixed social media text, and also reveal the difficulties of tone ambiguity. The potential avenues for further research will be to incorporate ensemble methods, data augmentation, focal loss, and larger multilingual pre-trained models to enhance the performance of the systems.

## Limitations

Despite the achievement of competitive performance, the proposed system has some limitations. First, the dataset employed in the current research has limited size and is in a low-resource language setting, which limits the ability of the model to learn complex linguistic patterns. Second, the employed dataset has highly informal code-mixed text, including spelling variants, inconsistent Romanization, and mixing of languages such as Tulu, English, and Kannada, which makes the interpretation of the text challenging for the model. Although the multilingual transformer-based model such as XLM-RoBERTa can effectively interpret the context of the given sentence, it is trained on large multilingual datasets that may not effectively cover the linguistic patterns of the Tulu language. Moreover, the employed dataset has class imbalance, which may affect the predictions of the model towards the classes that have the most occurrences despite the application of class weighting to the model. Lastly, because of the computational and time limitations of the current research, the model was not able to perform extensive experiments using other transformer-based models such as large transformer variants. In future research, we plan to improve the model further by utilizing a larger multilingual transformer model and/or ensemble model to improve the classification performance. Also, a larger dataset and better data preprocessing techniques for code-mixed text might improve the model's generalization performance.

## Ethical Considerations

The work conforms to ethical guidelines for responsible research in the field of Natural Language Processing. The data set employed in the research for this paper is publicly available YouTube comments provided by the organizers of the shared task, and no efforts were made to retrieve any personal information of the users of the YouTube platform. Hence, the work maintains the privacy of the users and does not handle any sensitive personal information of the users of the YouTube platform. Since social media data may have bias in terms of language or culture, it should be noted that bias may be present in the predictions made by the model. To avoid any bias in the predictions made by the model, balanced measures were used to evaluate the model's performance, such as Macro F1-Score. The proposed system aims to be used for research to facilitate constructive online discourse and should not be used for moderation. We also recognize that the data is from a low-resource language environment and may not adequately represent the linguistic spectrum of the Tulu-speaking people. Future work could include a larger and more representative data set to promote fairness and inclusiveness.

## References


Ahmed Abdullah, Sana Fatima, and Haroon Mahmood. 2025. Ghalib: A multilingual framework for hope speech detection in low-resource languages. *arXiv preprint arXiv:2512.22705*.

Muhammad Ahmad, Muhammad Waqas, Ameer Hamza, Ildar Batyrshin, and Grigori Sidorov. 2025. Hope speech detection in code-mixed roman urdu tweets: A positive turn in natural language processing. *arXiv preprint arXiv:2506.21583*.

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

Bharathi Raja Chakravarthi. 2022b. Multilingual hope speech detection in english and dravidian languages. *International Journal of Data Science and Analytics*, 14(4):389–406.

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.

Asha Hegde, Mudoor Devadas Anusha, Sharal Coelho, Hosahalli Lakshmaiah Shashirekha, and Bharathi Raja Chakravarthi. 2022. Corpus creation for sentiment analysis in code-mixed tulu text. In *Proceedings of the 1st Annual Meeting of the ELRA/ISCA Special Interest Group on Under-Resourced Languages*, pages 33–40.

Asha Hegde, Bharathi Raja Chakravarthi, Hosahalli Lakshmaiah Shashirekha, Rahul Ponnusamy, Subalalitha Cn, Lavanya SK, Martha Karunakar, Shreya Shreeram, Sarah Aymen, et al. 2023. Findings of the shared task on sentiment analysis in tamil and tulu code-mixed text. In *Proceedings of the Third Workshop on Speech and Language Technologies for Dravidian Languages*, pages 64–71.

Anqi Mao, Mehryar Mohri, and Yutao Zhong. 2023. Cross-entropy loss functions: Theoretical analysis and applications. *Preprint*, arXiv:2304.07288.

Ankit Kumar Mishra, Sunil Saumya, and Abhinav Kumar. 2021. Sentiment analysis of dravidian-codemix language. In *FIRE (Working Notes)*, pages 1011–1019.

Durairaj Thenmozhi, Rathnakar Shetty P, Parameshwar R. Hegde, Anusha M D, Raksha Adyanthaya, Mohammed Fadhel Aljunid, Prasanna Kumar Kumaresan, and Bharathi Raja Chakravarthi. 2026. Findings of the Shared Task on Hope Speech Detection in Tulu. In *Proceedings of the Sixth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.