SSN_HopeNetters@DravidianLangTech 2026: Multi-Level Hope Speech Detection using XLM-RoBERTa

Moogambigai A, Nandhika Saravanan, Nikhil Karthik S, Pandiarajan D, Bharathi B
Department of Computer Science and Engineering
Sri Sivasubramaniya Nadar College of Engineering, Chennai, India
moogambigai2370071@ssn.edu.in
nandhika2370028@ssn.edu.in
nikhilkarthik2370024@ssn.edu.in
pandiarajan2370062@ssn.edu.in
bharathib@ssn.edu.in

Abstract

This paper presents our system submission to the Shared Task on Hope Speech Detection in Code-Mixed Tulu Language at DravidianLangTech @ ACL 2026. We introduce a transformer-based approach built on XLM-RoBERTa-base for multilingual hope speech classification. Our system addresses two subtasks: coarse-grained classification of hope versus non-hope speech and fine-grained categorization of different hope expressions. Since hope is often expressed in subtle ways, especially in mixed-language text, our model looks at the full context of a sentence to understand its real meaning rather than just focusing on specific words. Experimental results demonstrate that multilingual transformer models effectively model supportive and encouraging expressions, underscoring their suitability for promoting constructive discourse in low-resource and code-mixed language settings.

1 Introduction

The rapid growth of social media platforms (Chakravarthi, 2022) has transformed the way people express opinions, emotions, and social support online. While significant research has focused on identifying harmful and hateful content, comparatively less attention has been given to detecting positive and encouraging expressions such as hope speech. Identifying such content can help promote inclusivity and constructive communication in online communities.

In this work, we present SSN_HopeNetters, a transformer-based system for hope speech detection in code-mixed social media text. The task is challenging due to multilingual usage, code-mixing, spelling variations, and contextual ambiguity. To address these challenges, we employ XLM-RoBERTa-base, which is effective for multilingual representation learning.

Our approach fine-tunes the pretrained model on the provided dataset using appropriate preprocessing and hyperparameter tuning. Experimental results demonstrate that transformer-based contextual embeddings effectively capture semantic cues associated with hopeful and supportive expressions.

Keywords

Hope Speech, Positive Language Modeling, Cross-lingual Representation Learning, Context-aware Text Understanding, Code-mixed Social Media Text, Inclusive Language Identification, Transformer Fine-tuning

2 Related Work

Research in text classification and predictive modeling has evolved through both traditional machine learning frameworks and task-specific linguistic modeling. Early work on learning predictive structures from multiple tasks and unlabeled data demonstrated the importance of leveraging shared representations for improved generalization(Ando and Zhang, 2005). Such frameworks laid the foundation for handling limited labeled data settings, which is particularly relevant for emerging tasks like hope speech detection.

Traditional approaches to text classification have influenced the development of systems for detecting socially meaningful content, including hope speech (Balouchzahi et al., 2021). Earlier statistical and log-linear modeling techniques demonstrated the feasibility of scalable learning in high-dimensional textual spaces, providing a foundation for later advances in discourse-level classification tasks (Andrew and Gao, 2007). Although hope speech detection involves more nuanced semantic interpretation, these foundational approaches contributed to the broader understanding of efficient text modeling strategies.

From a linguistic and structural standpoint, computational methods for string processing and sequence analysis have shaped modern text analy-

Proceedings of the Sixth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages, pages 366–370
July 3, 2026 ©2026 Association for Computational Linguistics

sis techniques (Gusfield, 1997). Furthermore, advances in syntactic parsing, such as the Yara Parser (Rasooli and Tetreault, 2015), highlight the importance of structural modeling in understanding sentence-level relationships. Such structural insights are relevant for hope speech detection, where meaning often depends on contextual cues, compositional structure, and implicit positive intent.

The task of hope speech detection has been formally introduced in recent years, emphasizing the identification of positive, supportive, and inclusive language in online discourse (Chakravarthi, 2022). Unlike hate speech detection, hope speech identification requires modeling subtle contextual cues and positive intent, making it inherently challenging.

More recently, shared task evaluations have expanded hope speech detection to low-resource and underrepresented languages. The findings of the shared task on hope speech detection in Tulu highlight challenges such as limited annotated data, linguistic diversity, and class imbalance. These studies underscore the importance of developing robust models for multilingual and low-resource settings.

3 Dataset

The dataset used in the shared task consists of social media comments annotated for hope speech detection in code-mixed Tulu (Thenmozhi et al., 2026). The data contains multilingual user-generated content, with labels indicating whether a given comment expresses hope speech or not. The overall distribution of the dataset is presented in Table 1.

Training Data: The training set contains labeled comments with fields including *ID*, *TEXT*, and *LABEL*. These annotations distinguish between hope speech and non-hope speech instances.

Test Data: The test set includes comments with *ID* and *TEXT* fields, without labels. Model predictions were generated on this unseen data for evaluation.

The dataset contains informal and user-generated text, introducing challenges such as spelling variations, code-mixing, transliteration, and contextual ambiguity. Class imbalance was observed between categories, making the classification task more challenging.

Basic preprocessing was performed to standardize the input format, including label encoding and stratified splitting into training and validation subsets. This ensured compatibility with transformer-based tokenization and fine-tuning procedures.

Task	Language	Train	Dev	Test
Task 1 (Coarse-Grained)	Tulu	5252	1126	1126
Task 2 (Fine-Grained)	Tulu	3185	682	683

Table 1: Dataset statistics for the Hope Speech Detection shared task.

4 Methodology

Our approach to hope speech detection is based on fine-tuning a pretrained multilingual transformer model. The overall framework of the proposed system is illustrated in Figure 1.

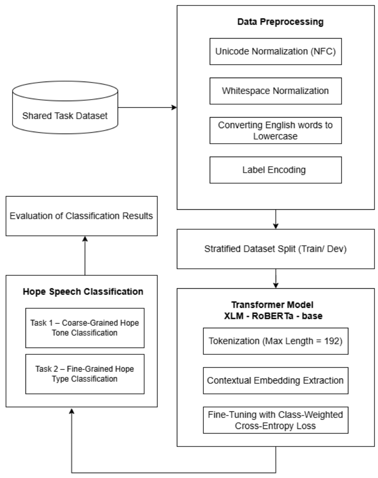

Figure 1: Framework of the Proposed Methodology

4.1 Model Architecture

XLM-RoBERTa-base:

We employed the pretrained XLM-RoBERTa-base model for multilingual text classification (Malik et al., 2023). XLM-RoBERTa is a transformer-based multilingual language model pretrained using masked language modeling on large-scale multilingual corpora. The model was fine-tuned on the labeled hope speech dataset to adapt its contextual representations to the classification task.

A task-specific classification head was added on top of the transformer encoder to predict the probability distribution over the output classes.

XLM-RoBERTa was selected due to its strong cross-lingual representation capability and its effectiveness in handling multilingual and code-mixed social media text (Zhao and Tao, 2021). Its subword tokenization mechanism also enables robust handling of transliterated and out-of-vocabulary words commonly found in informal online communication.

Hyperparameter	Value
Model	XLM-RoBERTa-base
Learning Rate	2e-5
Batch Size	16
Epochs	5
Optimizer	AdamW
Max Sequence Length	128
Loss Function	Cross-Entropy
Scheduler	Linear Warmup

Table 2: Training hyperparameters

4.2 Preprocessing

To ensure data consistency and effective model training, the following preprocessing steps were applied:

- **Text Normalization:** Basic cleaning was performed by removing unnecessary special characters and formatting inconsistencies while preserving semantic information.

- **Label Encoding:** Categorical labels were converted into numerical format for supervised learning.

- **Tokenization:** Text inputs were tokenized using the XLM-RoBERTa subword tokenizer, enabling effective handling of code-mixed and transliterated text.

- **Dataset Splitting:** The dataset was divided into training and validation subsets using stratified sampling to preserve class distribution.

4.3 Training Setup

The model was fine-tuned using cross-entropy loss for multi-class classification. Optimization was performed using the AdamW optimizer with learning rate scheduling to improve convergence stability.

Table 2 summarizes the training configuration used in our experiments.

Performance was evaluated using Accuracy, Precision, Recall, and Macro-F1 score on the validation set.

5 Experimental Results

This section presents the official leaderboard rankings and detailed evaluation results obtained by our team, **SSN_HopeNetters**, for both coarse-grained and fine-grained hope speech detection tasks.

5.1 Leaderboard Performance

Table 3 summarizes the official shared task leaderboard rankings.

5.2 Quantitative Analysis

Table 4 presents the validation performance of our proposed system for both tasks.

For Task 1, the model achieved an accuracy of 68.69% with a Macro-F1 score of 51.44%. Better performance was observed for majority classes such as *encouraging hope* and *uninvolved*, while lower performance on minority categories was mainly due to class imbalance.

For Task 2, the model achieved an accuracy of 44.43% and a Macro-F1 score of 38.52%. The lower performance reflects the increased difficulty of distinguishing fine-grained hope categories with overlapping semantic meanings.

Overall, the results indicate that multilingual contextual embeddings effectively capture hope-related expressions in code-mixed text, although fine-grained classification remains challenging.

5.3 Error Analysis

To better understand model behavior, we analyzed commonly occurring misclassification patterns observed in the validation predictions.

Error analysis revealed several challenges associated with code-mixed hope speech detection. The model frequently confused semantically similar hope categories due to overlapping emotional expressions. Additionally, transliterated spelling variations, noisy social media syntax, and implicit expressions of hope contributed to classification ambiguity.

Minority classes such as *fading hope* exhibited lower recall because of insufficient training examples. The findings indicate that fine-grained hope speech detection remains particularly difficult in low-resource multilingual settings.

5.4 Confusion Matrix Analysis

Figures 2 and 3 present the confusion matrices for Task 1 and Task 2.

The confusion matrices show that most errors occur between semantically related classes. In Task 2, minority categories such as *fading hope* experienced comparatively higher misclassification rates, highlighting the difficulty of fine-grained hope speech detection in low-resource code-mixed settings.

6 Conclusion

In this work, we presented a transformer-based approach for hope speech detection using the pre-

Task 1: Coarse-Grained Hope Tone Classification							
Team Name	Run	Acc	Macro-P	Macro-R	W-P	Macro-F1	Rank
SSN_HopeNetters	Run 1	0.64	0.56	0.56	0.68	0.55	2

Task 2: Fine-Grained Hope Type Classification							
Team Name	Run	Acc	Macro-P	Macro-R	W-P	Macro-F1	Rank
SSN_HopeNetters	Run 1	0.45	0.40	0.43	0.49	0.38	4

Table 3: Official leaderboard performance of SSN_HopeNetters

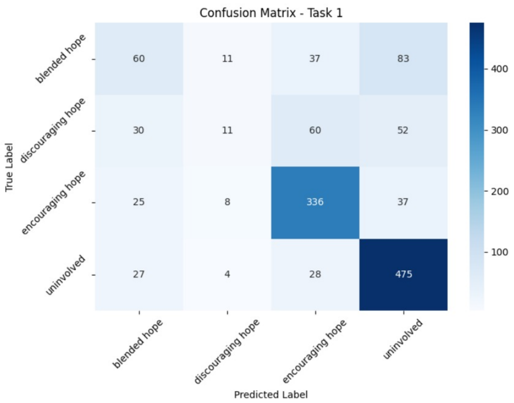

Figure 2: Task 1 Confusion Matrix

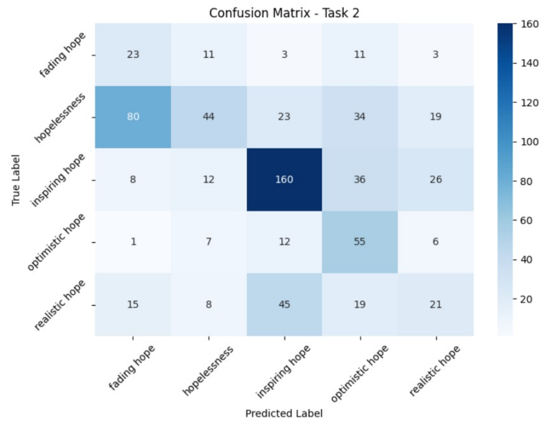

Figure 3: Task 2 Confusion Matrix

Task	Accuracy	Macro F1
Task 1 (Coarse)	0.6869	0.5144
Task 2 (Fine)	0.4443	0.3852

Table 4: Validation performance of SSN_HopeNetters

Actual	Predicted	Reason
Inspiring Hope	Optimistic Hope	Semantic overlap
Hope Speech	Non-Hope	Implicit positivity
Fading Hope	Uninvolved	Minority class sparsity

Table 5: Common misclassification patterns

trained XLM-RoBERTa model. By fine-tuning the model on labeled social media data, the system effectively learned contextual and semantic patterns associated with hopeful and supportive expressions. The multilingual pretraining of XLM-RoBERTa enabled robust handling of informal language, transliteration, and code-mixed content commonly found in user-generated text.

The experimental results demonstrate that contextualized transformer representations significantly enhance the ability to distinguish hope speech from non-hope speech. Overall, the study highlights the effectiveness of large-scale multilingual language models for socially relevant text classification tasks.[1]

7 Limitations

Despite achieving competitive performance, several limitations remain. First, the dataset size is relatively limited, which may restrict the model's ability to generalize to unseen domains or emerging linguistic patterns. Hope speech often depends on subtle contextual cues and implicit meaning, which can be challenging even for transformer-based models.

Additionally, the presence of informal language, spelling variations, and code-mixed expressions increases classification complexity. The model may occasionally misclassify borderline cases where hope is expressed indirectly or sarcastically. Future work should explore larger annotated datasets, improved contextual modeling strategies, and domain adaptation techniques to further enhance robustness and generalization.

[1] Code and implementation details are available at the project repository: GitHub Repository.

References

Rie Kubota Ando and Tong Zhang. 2005. A framework for learning predictive structures from multiple tasks and unlabeled data. *Journal of Machine Learning Research*, 6:1817–1853.

Galen Andrew and Jianfeng Gao. 2007. Scalable training of L1-regularized log-linear models. In *Proceedings of the 24th International Conference on Machine Learning*, pages 33–40.

Fazlourrahman Balouchzahi, Aparna B K, and H L Shashirekha. 2021. MUCS@LT-EDI-EACL2021:CoHope-hope speech detection for equality, diversity, and inclusion in code-mixed texts. In *Proceedings of the First Workshop on Language Technology for Equality, Diversity and Inclusion*, pages 180–187, Kyiv. Association for Computational Linguistics.

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

Dan Gusfield. 1997. *Algorithms on Strings, Trees and Sequences*. Cambridge University Press, Cambridge, UK.

Muhammad Shahid Iqbal Malik, Anna Nazarova, Mona Mamdouh Jamjoom, and Dmitry I Ignatov. 2023. Multilingual hope speech detection: A robust framework using transfer learning of fine-tuning roberta model. *Journal of King Saud University-Computer and Information Sciences*, 35(8):101736.

Mohammad Sadegh Rasooli and Joel R. Tetreault. 2015. Yara parser: A fast and accurate dependency parser. *Computing Research Repository*, arXiv:1503.06733. Version 2.

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.

Yingjia Zhao and Xin Tao. 2021. Zyj@ lt-edi-eacl2021: Xlm-roberta-based model with attention for hope speech detection. In *Proceedings of the first workshop on language technology for equality, diversity and inclusion*, pages 118–121.