Findings of the Shared Task on Hope Speech Detection in Tulu

Durairaj Thenmozhi[1], Anusha M D[2], Raksha Adyanthaya[2], Rathnakara Shetty P[2], Parameshwar R. Hegde[2], Mohammed Fadhel Aljunid[3], Prasanna Kumar Kumaresan[4], Bharathi Raja Chakravarthi[4]

[1]Sri Sivasubramaniya Nadar College of Engineering, Tamil Nadu, India
[2]YIASCM, Yenepoya (Deemed to be University), Mangalore, India
[3]Moodlakatte Institute of Technology, Karnataka, India
[4]Unit for Inclusive AI, Data Science Institute, University of Galway, Ireland

Abstract

Hope Speech Identification is the process of detecting positive, supportive, and encouraging language in text. It focuses on identifying content that promotes unity, inclusiveness, and resilience. Identification of hope speech helps supports mental well being, create healthier online environments, counter hate speech, and promote positive digital communication. This shared task hope speech detection in code-mixed Tulu language as part of DravidianLangTech @ ACL 2026, focuses on both the coarse-grained hope tone classification and the fine-grained hope type classification tasks. There are 11 teams participated in the tasks and submitted several runs for both the tasks. The teams are ranked based on the macro-F1 score.

1 Introduction

The proliferation of social media platforms has significantly transformed digital communication, enabling individuals to express emotions, opinions, and social support at an unprecedented scale (Chakravarthi et al., 2021). While substantial research in Natural Language Processing (NLP) has focused on detecting harmful content such as hate speech, offensive language, and misinformation, comparatively limited attention has been given to identifying positive and supportive expressions, particularly hope speech (Bali et al., 2014). Hope speech plays a vital role in promoting emotional well-being, strengthening social cohesion, and fostering constructive discourse in online communities (Conneau et al., 2020).

Automatic detection of hope-related expressions presents unique challenges, especially in multilingual and code-mixed environments (Kakwani et al., 2020). Most existing studies on hope speech detection have concentrated on high-resource languages such as English, Hindi, and Tamil (Chakravarthi et al., 2020). However, low-resource Dravidian languages remain underexplored due to the scarcity of annotated datasets and benchmark tasks (Chakravarthi et al., 2021). Tulu, a Dravidian language spoken predominantly in coastal Karnataka, represents one such underrepresented language, particularly in code-mixed contexts involving Tulu, English, and Kannada. The informal and transliterated nature of social media text further complicates computational modeling.

To address this research gap, we introduce the Hope Speech Detection shared task at DravidianLangTech@ACL 2026. The shared task presents the first benchmark dataset for hope speech detection in code-mixed Tulu and proposes a multi-level classification framework comprising two tracks: (i) coarse-grained hope tone classification and (ii) fine-grained hope type classification. The coarse-grained task focuses on identifying the overall emotional orientation of user comments with respect to hope, while the fine-grained task captures nuanced distinctions such as optimistic hope, realistic hope, inspiring hope, fading hope, and hopelessness.

The objective of this shared task is threefold:

1. To encourage research in low-resource Dravidian languages,

2. To benchmark computational approaches for emotionally nuanced language understanding, and

3. To foster the development of robust models capable of handling code-mixed and socially contextualized text.

The task attracted multiple participating teams employing diverse methodologies ranging from traditional machine learning models to transformer-based architectures and ensemble approaches. This

Proceedings of the Sixth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages, pages 80–87
July 3, 2026

paper presents an overview of the dataset, task formulation, participating systems, evaluation results, and key findings from the shared task.

2 Related Work

Hope speech detection, aimed at identifying uplifting, inspiring, and encouraging phrases on the Internet, has become an important problem in Natural Language Processing. Initial studies, as proposed by Chakravarthi, introduced the HopeEDI dataset, a multilingual collection of YouTube comments in English, Tamil, and Malayalam, and provided baseline results with the help of traditional machine learning models, including Support Vector Machines (SVM), Logistic Regression, Decision Trees, and Naive Bayes using the TF-IDF features. The study of (Divakaran et al., 2024) investigated the use of machine learning and transformer-based models, such as SVM, Logistic Regression, BERT, and DistilSpanBERT, for the detection of hope speech in Spanish and English as part of the shared task IberLEF 2024. The transformer models outperformed conventional machine learning techniques, achieving macro F1-scores of up to 0.82 for binary classification and 0.64 for multiclass classification. This improved Hope speech detection in multilingual environments through recent shared task initiatives. Subsequent work by Chakravarthi and Muralidaran explored deep learning-based approaches like Recurrent Neural Networks (RNN), Convolutional Neural Networks (CNN), and CNN-LSTM approaches, which achieved weighted F1-scores of 0.93 in English, 0.84 in Malayalam, and 0.56 in Tamil, better than the traditional machine learning baselines, and found that contextual modeling through deeper neural networks was important to hope speech recognition. Despite these advancements, hope speech detection remains underexplored for low-resource languages like tulu. To address this gap, the present shared task introduces a new dataset and a variety of evaluation frameworks for hope speech detection.

3 Task Description

Shared task on hope speech detection in code-mixed Tulu language consists of two subtasks namely coarse-grained hope tone classification and fine-grained hope type classification.

- Task 1: Coarse-Grained Hope Tone Classification

  The task involves classifying social media comments' emotional tone by examining hope-related expressions. Comments must be categorized based on presence, absence, or ambiguity of hope.

  1. Encouraging: Text that conveys a hopeful, positive, or emotionally supportive tone. This may include affirmations, motivational statements, reassurance, or empathy.
  2. Discouraging: Content that reflects hopelessness, negativity, emotional withdrawal, despair, or defeat.
  3. Uninvolved: Neutral or informational content that does not express emotional engagement related to hope.
  4. Blended Tone: Text containing both hopeful and discouraging elements, resulting in emotional ambiguity or complexity.

  This coarse-grained classification enables systems to capture the general emotional orientation of user-generated content.

- Task 2: Fine-Grained Hope Type Classification

  The second task involves a more detailed classification of hope-related expressions. This fine-grained categorization is applied when hope signals are present (e.g., in Encouraging or Blended Tone categories, and certain Discouraging cases reflecting diminishing hope).

  The following labels are defined:

  1. Optimistic Hope: A confident belief that positive outcomes will occur in the future.
  2. Realistic Hope: Hope grounded in practical expectations and achievable outcomes.
  3. Inspiring Hope: Motivational expressions intended to encourage others toward a better future.
  4. Fading Hope: Expressions indicating that hope is gradually diminishing due to persistent challenges.
  5. Hopelessness: A strong sense of despair characterized by the belief that no positive outcome is possible.

This fine-grained task allows models to move beyond binary sentiment detection and capture subtle emotional distinctions within hope-related discourse.

Participants were given training, development, and test datasets to help them create classification models. The goal of this objective was to promote research in under-resourced languages and investigate novel techniques to offense detection in multilingual and code-mixed texts. Table 1 and Table 2 present illustrative examples from the dataset for Task 1 and Task 2 respectively, along with their English translations to aid understanding of the code-mixed Tulu text and the distinctions between label categories.

4 Data Description

The dataset is organized into two tasks: Task 1 (Coarse-Grained Hope Tone Classification) and Task 2 (Fine-Grained Hope Type Classification). The data for both sub tasks is divided into training, development, and test sets while maintaining representative label distributions.

Task 1 consists of 7,504 instances, comprising 8,559 comments and 9,972 sentences with a total of 69,798 words. Task 2 contains 4,550 instances, including 3,014 comments and 32,511 sentences with 671,420 total words.

Each sample was first annotated by two native Tulu speakers who also have a background in NLP research, giving them both the linguistic intuition and technical understanding needed for the task. The annotated samples were then passed on to linguistic experts from university language faculties, who bring years of teaching experience in the language. These experts reviewed the labels, clarified ambiguous cases, and ensured consistency across all categories. Inter-annotator agreement measured using Krippendorff's Alpha (Krippendorff, 2018) is 0.892 for Task 1 and 0.901 for Task 2, indicating strong and excellent annotation reliability respectively.

5 Participants Methodology

A total of 11 teams participated in Task 1, coarse-grained hope tone classification and 9 teams participated in Task 2, fine-grained hope type classification. They all used methodologies ranging from traditional classifiers with TF-IDF vectorization to transformer models for the above mentioned tasks. The run submissions of each team are explained below.

5.1 CUET_Synthetica

For Task 1, three pre-trained transformers xlm-roberta-base, google/muril-base-cased (MuRIL), and mbert were fine-tuned by this team (Zaman et al., 2026) for classifying text into four hope speech categories. The text was preprocessed by removing unwanted characters, normalizing whitespace, and tokenizing with a 128-token maximum length. Class-weighted loss was used to handle class imbalance. Models were trained using the AdamW optimizer (learning rate: 2e-5, weight decay: 0.01) for up to 8 epochs with early stopping (patience=2) based on the macro F1 score. The predictions from all three models were combined using an ensemble approach, where their logits were averaged element-wise. The final predictions were made using the argmax of the averaged logits. The ensemble of all three models achieved a development macro F1 score of 0.622 on the development set, improving upon individual models (XLM-R: 0.591, MuRIL: 0.591, mBERT: 0.607).

For Task 2, the model used xlm-roberta-base, a pre-trained transformer, for fine-grained classification of text data into five categories: optimistic hope, realistic hope, inspiring hope, fading hope, and hopelessness. The data was preprocessed by cleaning and tokenizing the text, with labels converted to lowercase and mapped to numerical values. The model was fine-tuned with class-weighted loss to address class imbalance, using the AdamW optimizer with a learning rate of 2e-5 and weight decay of 0.01. Training was done for up to 8 epochs with early stopping after 2 epochs of no improvement. The model's performance was evaluated using macro F1 score, and a learning rate scheduler was applied with a warm-up ratio of 0.1.The model achieved a macro F1 score of 0.4349 on the development set.

5.2 SSN_HopeNetters

Team SSN_HopeNetters (A et al., 2026) fine-tuned the multilingual transformer XLM-RoBERTa for Task 1. Input texts are tokenized using the XLM-R tokenizer with a fixed maximum sequence length, and the model is trained to predict one of four tone labels. To address class imbalance across tone categories, they incorporate a class-weighted cross-entropy loss, ensuring improved learning for minority classes such as discouraging and blended hope.

Table 1: Illustrative Examples for Task 1 : Coarse-Grained Hope Tone Classification

Text (Code-Mixed Tulu)	English Translation	Label
Rathish mendon.. thank you so much anna.. and comment maltina prativoriyagla solmelu	Rathish Mendon, thank you so much brother, and my greetings to everyone who commented	Encouraging Hope
Audio days marre madee sari ijja	Even the audio these days is not good at all	Discouraging Hope
Namma tulundadu comedy first epodu	Our Tulu content always puts comedy first	Uninvolved
pudar daane change malthini natakada,..	It will change on its own, this is just drama	Blended Hope

Table 2: Illustrative Examples for Task 2: Fine-Grained Hope Type Classification

Text (Code-Mixed Tulu)	English Translation	Label
Navin padilerna gotanaga Portand serial mast soku ethundu, anchina scripts malpule, jai thulu nadu	If Navin acts in it, the Portland serial will be really great, please write such scripts, long live Tulu Nadu	Optimistic Hope
Very nice. Baari yedded malpuva. Superb. Good Luck	Very nice. You are doing really well. Superb. Good Luck	Inspiring Hope
undu artists.. situation toodhu dakkini, super skit	These artists have understood the situation well, super skit	Realistic Hope
Tuluta nadut kannada borchitund materla tulut paternda yedde ettund.	While walking in Tulu, Kannada is being mixed in; the original Tulu pattern is gradually being lost	Fading Hope
bupper pura paterver mare karma. inchi sav da film manpvar nikulu. thu	Completely wasted effort, what bad karma. Those who make such films should just stop	Hopelessness

Table 3: Dataset Split for Task 1 and Task 2

Task	Train	Dev	Test	Total
Task 1 (Coarse)	5,252	1,126	1,126	7,504
Task 2 (Fine)	3,185	682	683	4,550

Table 4: Task 1 Label Distribution

Label	Count
Encouraging Hope	2,708
Uninvolved	2,657
Blended Hope	1,146
Discouraging Hope	993
Total	**7,504**

Table 5: Task 2 Label Distribution

Label	Count
Inspiring Hope	1,613
Hopelessness	1,338
Realistic Hope	719
Optimistic Hope	543
Fading Hope	337
Total	**4,550**

The model is optimized using AdamW with a low learning rate and extended training epochs.

For Task 2, they performed fine-grained classification of hope expressions into five semantically nuanced categories, focusing on emotionally subtle and imbalanced labels. They employ XLM-RoBERTa as the base encoder to capture multilingual and code-mixed representations of Tulu text. Given the severe label imbalance in the dataset, they apply a weighted loss function that penalizes misclassification of low-frequency hope types such as fading and realistic hope. The model is fine-tuned with a reduced learning rate, longer warm-up, and increased training epochs to stabilize optimization. Predictions are generated using a softmax output layer.

Table 6: Comparison of submitted systems for Task 1: Coarse-Grained Hope Tone Classification.

Team	Preprocessing	Data Balancing	External Dataset	Approach	XLM-RoBERTa	MBert	IndicBert	Muril	DistilBert	Traditional	Prompting
CUET_Synthetica (Zaman et al., 2026)	✓	✓	✗	Ensemble of transformers	✓	✓	✗	✓	✗	✗	✗
SSN_HopeNetters (A et al., 2026)	✗	✓	✗	Transformers with data balancing	✓	✗	✗	✗	✗	✗	✗
Hope_Alchemists (Johnson et al., 2026)	✓	✓	✗	LinearSVC with TFIDF	✗	✗	✗	✗	✗	✓	✗
Team Oryu (Moni et al., 2026)	✗	✓	✗	Transformers with oversampling	✓	✗	✗	✗	✗	✗	✗
SJM_MINDS (Shashirekha et al., 2026)	✓	✓	✗	LinearSVC with TF-IDF	✗	✗	✗	✗	✗	✓	✗
Tensor_Thunders	✓	✗	✗	Fine tuning transformers	✗	✗	✗	✗	✓	✗	✗
CUETEuropa	✗	✓	✗	Fine tuning transformers	✗	✗	✓	✗	✗	✗	✗
cantnlp (Li and Wong, 2026)	✗	✗	✓	Fine tuning transformers with external dataset	✓	✗	✗	✗	✗	✗	✗
MUCS (Shashirekha and A, 2026)	✓	✗	✗	Traditional classifier with TF-IDF	✗	✗	✗	✗	✗	✓	✗
NamDang	✗	✗	✗	DSPy framework with few-shot prompting	✓	✗	✗	✗	✗	✗	✓
PrimeLine (V et al., 2026)	✗	✗	✗	Fine tuning transformers	✓	✗	✗	✗	✗	✗	✗

Table 7: Comparison of submitted systems for Task 2: Fine-Grained Hope Type Classification.

Team	Preprocessing	Data Balancing	External Dataset	Approach	XLM-RoBERTa	MBert	IndicBert	Muril	DistilBert	Traditional
CUET_Synthetica (Zaman et al., 2026)	✓	✓	✗	Fine tuning transformers	✓	✗	✗	✗	✗	✗
SSN_HopeNetters (A et al., 2026)	✗	✓	✗	Transformers with weighted loss	✓	✗	✗	✗	✗	✗
Hope_Alchemists (Johnson et al., 2026)	✗	✓	✗	Fin tuning transformer with oversampling	✓	✗	✗	✗	✗	✗
SJM_MINDS (Shashirekha et al., 2026)	✓	✓	✗	Logistic Regression with TF-IDF	✗	✗	✗	✗	✗	✓
Tensor_Thunders	✓	✗	✗	Fine tuning transformers	✗	✗	✗	✗	✓	✗
CUETEuropa	✗	✓	✗	Fine tuning transformers	✗	✗	✓	✗	✗	✗
cantnlp (Li and Wong, 2026)	✗	✗	✓	Fine tuning transformers with external dataset	✓	✗	✗	✗	✗	✗
MUCS (Shashirekha and A, 2026)	✓	✗	✗	Traditional classifier with TF-IDF	✗	✗	✗	✗	✗	✓
PrimeLine (V et al., 2026)	✗	✗	✗	Fine tuning transformers	✓	✗	✗	✗	✗	✗

5.3 Hope_Alchemists

For Task 1, this team (Johnson et al., 2026) employed the LinearSVC classifier using class_weight="balanced" with a TF-IDF (word 1–2 grams + character 3–5 grams) features to train the model. For Task 2, they fine-tuned XLM-RoBERTa base for the classification. To address label imbalance, they used oversampling to equalize class counts and trained the transformer with focal loss and early stopping.

5.4 Team Oryu

This team (Moni et al., 2026) participated only in Task 1 and they fine-tuned the pretrained XLM-RoBERTa model to leverage its strong multilingual contextual representations for the Tulu dataset. The raw text was first cleaned and tokenized using the XLM-R tokenizer with a maximum sequence length of 192 to preserve contextual cues. To address class imbalance, they applied training data oversampling and optimized the model using focal loss, which improves learning for hard and minority-class examples. The model was trained with the AdamW optimizer, linear learning rate scheduling with warmup.

5.5 SJM_MINDS

This team (Shashirekha et al., 2026) implemented a Linear Support Vector Machine (Linear SVC) classifier for Task 1. Prior to feature extraction, text preprocessing was performed, including lowercasing and removal of URLs, user mentions, hashtags, and punctuation. Text data was transformed using TF-IDF representation with up to 8000 features, considering unigram and bigram combinations and removing rare terms occurring in less than two documents. To address class imbalance, class weights were set to "balanced".

For Task 2, the team used a Logistic Regres-

Table 8: Hope Speech Identification - Task 1 Results

S.no	Team Name	Run	Accuracy	Macro Precision	Macro Recall	Weighted Precision	Macro F1-Score	Rank
1	CUET_Synthetica	Run 1	0.67	0.58	0.58	0.69	0.58	1
2	SSN_HopeNetters	Run 1	0.64	0.56	0.56	0.68	0.55	2
3	Hope_Alchemists	Run 2	0.64	0.55	0.55	0.66	0.55	2
4	Team Oryu	Run 1	0.60	0.55	0.54	0.68	0.53	3
5	SJM_MINDS	Run 2	0.62	0.52	0.51	0.61	0.51	4
6	Tensor_Thunders	Run 1	0.65	0.53	0.51	0.62	0.51	4
7	CUETEuropa	Run 2	0.58	0.54	0.52	0.67	0.51	4
8	cantnlp	Run 1	0.66	0.50	0.51	0.62	0.50	5
9	MUCS	Run 2	0.57	0.47	0.47	0.57	0.47	6
10	NamDang	Run 1	0.50	0.43	0.43	0.52	0.41	7
11	PrimeLine	Run 1	0.60	0.30	0.41	0.44	0.34	8

Table 9: Hope Speech Identification - Task 2 Results

S.no	Team Name	Run	Accuracy	Macro Precision	Macro Recall	Weighted Precision	Macro F1-Score	Rank
1	CUET_Synthetica	Run 3	0.50	0.42	0.42	0.49	0.42	1
2	Tensor_Thunders	Run 1	0.54	0.48	0.41	0.52	0.41	2
3	Hope_Alchemists	Run 1	0.49	0.41	0.40	0.49	0.40	3
4	SSN_HopeNetters	Run 1	0.45	0.40	0.43	0.49	0.38	4
5	SJM_MINDS	Run 1	0.56	0.64	0.39	0.59	0.37	5
6	CUETEuropa	Run 2	0.45	0.36	0.36	0.45	0.36	6
7	cantnlp	Run 1	0.54	0.35	0.36	0.45	0.33	7
8	MUCS	Run 2	0.38	0.29	0.29	0.37	0.29	8
9	PrimeLine	Run 1	0.24	0.19	0.19	0.24	0.19	9

sion classifier with TF-IDF feature representation. The dataset was split into 80% training and 20% validation sets using a fixed random seed (random_state=42) to ensure reproducibility. The Logistic Regression model was trained with a maximum iteration limit of 1000. After validation, the model was retrained on the full training dataset and used to generate predictions for the test set prediction.

5.6 Tensor_Thunders

This team used the DistilBERT model for both Task 1 and Task 2 to perform text classification on the given dataset. The data was preprocessed and then the pre-trained DistilBERT model was fine-tuned by using the optimized hyperparameters.

5.7 CUETEuropa

This team used IndicBERT as the backbone architecture for both Task 1 and Task 2. The hope speech labels were mapped to numerical IDs using a bidirectional label encoding scheme. Text inputs were tokenized and truncated or padded to a maximum sequence length of 128 tokens. A linear classification head was added on top of the pooled <s> token representation for multi-class prediction. To address class imbalance, they applied weighted cross-entropy loss with manually computed class weights along with label smoothing (0.1). The model was fine-tuned using the AdamW optimizer, weight decay of 0.01, and a linear warmup over 10% of training steps, using mixed-precision (FP16) training.

5.8 cantnlp

This team (Li and Wong, 2026) used an approach of fine-tuning XLM-RoBERTa for both Task 1 and Task 2. They have used an external dataset namely "Corpus of Global Language Use" while training the model.

5.9 MUCS

The methodology proposed by the team MUCS (Shashirekha and A, 2026) extracted features including word n-grams, character n-grams, syllable n-grams and subword units using Byte Pair Encoding tokenizer. For Kannada, syllables are generated using Indic syllabifier while English syllables are generated using rule-based vowel grouping method. All extracted features are normalized and merged into a single feature representation. The combined features are transformed into numerical vectors using TF-IDF Vectorization. Several Machine Learning classifiers including MultinomialNB, Logistic Regression, K Nearest Neighbors, Decision Tree are trained on this representation. They used this methodology for both Task 1 and 2.

5.10 NamDang

The team NamDang participated only in Task 1 who employs the DSPy (Declarative Selfimproving Language Programs) framework to move beyond manual prompt engineering toward programmatic optimization. The core logic is defined through declarative signatures, which are then optimized using the BootstrapFewShot teleprompter. This process automatically curates and bootstraps high-quality demonstrations from the training data to guide the model's reasoning. Finally, a robust post-processing layer with fuzzy matching is used to ensure the Large Language Model's outputs are accurately mapped to the four required hope tone categories.

5.11 Prime_Line

Team PrimeLine (V et al., 2026) used XLM-RoBERTa by utilizing model's subword tokenizer and fine-tuned on the provided training data for a small number of epochs to avoid overfitting. This team used the same approach for both Task 1 and Task2.

Tables 6 and 7 summarizes the methodologies used by the teams for the tasks. Some of the teams such as CUET_Synthetica, SJM_MINDS and MUCS applied preprocessing techniques before training the models. A few teams such as CUET_Synthetica, SJM_MINDS and CUETEuropa employed techniques to handle data imbalance problems. One team "cantnlp" used external data set while training the model. Majority of the teams used XLM-RoBERTa to train the models. Team NamDang used prompting with few-shot approach in Task 1.

6 Results

The runs submitted by the teams are evaluated using the metrics namely accuracy, precision, recall and macro-F1 score values. Teams are ranked based on the macro-F1 obtained for their runs. The values for Task 1 are given in Table 8 and the values for Task 2 are provided in Table 9. It is observed that the team CUET_Synthetica who used ensemble of transformers secured the first position in Task 1. The same who fine tuned XLM-RoBERTa secured the top position in the leader board for Task 2.

7 Conclusion

This paper presented the findings of the Shared Task on Hope Speech Detection in Code-Mixed Tulu Language at DravidianLangTech@ACL 2026. It consists of 2 sub tasks namely Task 1: Coarse-Grained Hope Tone Classification and Task 2: Fine-Grained Hope Type Classification. Task 1 is a 4 class and Task 2 is a 5 class classification problems on hope speech identification. The task has 7504 instances for Task 1 and 4550 instances for Task 2. A total of 9 teams participated in both tasks. Additionally, 2 more teams participated only in Task 1. The team CUET_Synthetica secured first position in both the tasks. They used ensemble of transformers for Task 1 and fine tuned XLM-RoBERTa for Task 2. Though, some of the teams tried to address the data imbalancing problem, the methodology they used is not effective. Also, employing language agnostics embeddings may give a good results which was not tried by any of the team.

8 Limitations

The participating teams were able to achieve only 0.58 and 0.42 as the maximum scores for Task 1 and Task 2 respectively. Most of the teams employed traditional classifiers with bag of word features which do not capture any semantics to identify the hope speech. Other teams fine tuned transformers which is also not helped so much to improve the performance.

9 Ethical Considerations

The dataset used in this shared task is collected from social media postings that are publicly available by ensuring compliance with the terms and policies of social media platforms. Special care was taken to avoid collecting sensitive personal data related to religion, ethnicity, gender identity, health status, or political affiliation.

Acknowledgments

Author Bharathi Raja Chakravarthi was funded by a research grant from Research Ireland under grant number SFI/12/RC/2289_P2 (Insight). This work is also supported by the Centre for Research Training in Artificial Intelligence grant number 18/CRT/6223, and a grant from the College of Science and Engineering, University of Galway, Ireland.

References

Moogambigai A, Nandhika Saravanan, Nikhil Karthik S, Pandiarajan D, and Bharathi B. 2026. SSN_hopeNetters@DravidianLangTech 2026: Multi-level hope speech detection using XLM-RoBERTa. In *Proceedings of the Sixth Workshop on*

Speech, Vision, and Language Technologies for Dravidian Languages. Association for Computational Linguistics.

Kalika Bali, Jatin Sharma, Monojit Choudhury, and Yogarshi Vyas. 2014. Borrowing or code-mixing? an analysis of english-hindi code-mixed social media text. *Proceedings of EMNLP Workshop on Computational Approaches to Code Switching*.

Bharathi Raja Chakravarthi. 2020. Hopeedi: A multilingual hope speech detection dataset for equality, diversity, and inclusion. In *Proceedings of the third workshop on computational modeling of people's opinions, personality, and emotion's in social media*, pages 41–53.

Bharathi Raja Chakravarthi and Vigneshwaran Muralidaran. 2021. Findings of the shared task on hope speech detection for equality, diversity, and inclusion. In *Proceedings of the first workshop on language technology for equality, diversity and inclusion*, pages 61–72.

Bharathi Raja Chakravarthi, Vigneshwaran Muralidaran, D Thenmozhi, and John P McCrae. 2021. Hope speech detection in under-resourced languages. In *Proceedings of the 5th Workshop on Online Abuse and Harms (WOAH)*. Association for Computational Linguistics.

Bharathi Raja Chakravarthi, Ruba Priyadharshini, Vigneshwaran Muralidaran, and John P McCrae. 2020. Sentiment analysis for dravidian languages. In *Proceedings of FIRE*.

Alexis Conneau, Kartikay Khandelwal, Naman Goyal, and et al. 2020. Unsupervised cross-lingual representation learning at scale. In *Proceedings of ACL*. Association for Computational Linguistics.

Sonith Divakaran, Kavya Girish, and Hosahalli Lakshmaiah Shashirekha. 2024. Hope on the horizon: Experiments with learning models for hope speech detection in spanish and english. In *IberLEF@ SEPLN*.

Joel Johnson, Meclin A Francis, Jyoti Kumari, Malavika Sreekumar, and Vinay Babu Ulli. 2026. Hope_Speech_Alchemists@DravidianLangTech 2026: TF-IDF SVM and XLM-RoBERTa with focal loss for hope speech detection in tulu. In *Proceedings of the Sixth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.

Divyanshu Kakwani, Anoop Kunchukuttan, Satish Golla, and et al. 2020. Indicnlpsuite: Monolingual corpora, evaluation benchmarks and pre-trained multilingual language models for indian languages. In *Findings of EMNLP*. Association for Computational Linguistics.

Klaus Krippendorff. 2018. *Content Analysis: An Introduction to Its Methodology*. Sage Publications.

Andrew Li and Sidney Wong. 2026. cantnlp@DravidianLangTech 2026: Organic domain adaptation improves multi-class hope speech detection in tulu. In *Proceedings of the Sixth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.

Joyeta Barua Moni, Noore Tamanna Orny, Md. Abtahee Kabir, and Hasan Murad. 2026. Team Oryu@DravidianLangTech 2026: A multilingual transformer approach for hope speech detection in code-mixed tulu. In *Proceedings of the Sixth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.

Hosahalli Lakshmaiah Shashirekha and Rachana A. 2026. MUCS@Dravidianlangtech@ACL2026: Hope speech detection in code-mixed tulu language using multiple features. In *Proceedings of the Sixth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.

Hosahalli Lakshmaiah Shashirekha, Manjula Suvarna, and Jayashree Krishna. 2026. SJM_MINDS@DravidianLangTech@ACL 2026: Machine learning approaches for hope speech detection in code-mixed tulu. In *Proceedings of the Sixth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.

Rithikaa V, Sanjay Krishnan K, Nithya Varshini C N R, and S. Sumathi. 2026. PrimeLine@DravidianLangTech 2026: Hope speech detection in tulu using XLM-RoBERTa for coarse and fine-grained classification. In *Proceedings of the Sixth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.

Sumaiya Zaman, Miftahul Jannat Rishta, Shiti Chowdhury, and Hasan Murad. 2026. CUET_SYNTHETICA@DravidianLangTech 2026: Multilingual transformer based hope speech detection for coarse and fine-grained classification in tulu. In *Proceedings of the Sixth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.