indiDataMiner at SemEval-2025 Task 11: From Text to Emotion: Transformer-Based Models for Emotions Detection in Indian Languages

Saurabh Kumar[1], **Sujit Kumar**[2],
Sanasam Ranbir Singh[1] and **Sukumar Nandi**[1]
[1]Indian Institute of Technology Guwahati, Assam, India
[2]Lee Kong Chain School of Medicine, Nanyang Technological University Singapore
{saurabh1003, ranbir, sukumar}@iitg.ac.in, sujit.kumar@ntu.edu.sg

Abstract

Emotion detection is essential for applications like mental health monitoring and social media analysis, yet remains underexplored for Indian languages. This paper presents our system for SemEval-2025 Task 11 (Track A), focusing on multilabel emotion detection in Hindi and Marathi, two widely spoken Indian languages. We fine-tune IndicBERT v2 on the BRIGHTER dataset, achieving F1 scores of 87.37 (Hindi) and 88.32 (Marathi), outperforming baseline models. Our results highlight the effectiveness of fine-tuning a language-specific pretrained model for emotion detection, contributing to advancements in multilingual NLP research.

1 Introduction

Emotion detection has garnered significant interest among researchers due to its psychological, social, and commercial importance. People express emotions explicitly through words like Happy or Angry and implicitly through context, tone, or figurative language (Garg and Lobiyal, 2020a). The complexity and subjectivity of human emotions make their accurate identification challenging, especially in text-based scenarios. Moreover, emotional expression varies significantly across languages and cultures, making it difficult to develop models that generalize well across linguistic boundaries (Wiebe et al., 2005; Mohammad and Kiritchenko, 2018).

While extensive research has been conducted on emotion detection in high-resource languages such as English, Spanish, German, and Arabic (Strapparava and Mihalcea, 2007; Chatterjee et al., 2019; Nandwani and Verma, 2021; Maruf et al., 2024), many low-resource languages, including those spoken in India, remain largely underexplored. Recognizing this gap, Muhammad et al. (2025a) introduced the BRIGHTER dataset, which covers 28 languages, primarily low-resource, spoken across Africa, Asia, Eastern Europe, and Latin America. This dataset also includes two widely spoken Indian languages, Hindi and Marathi.

Given India's rich linguistic diversity, understanding emotional expression in its various languages is essential for advancing multilingual NLP applications. This paper introduces our proposed approach and provides a comprehensive analysis of the task of SemEval-2025 Task 11[1] (Track A): *Bridging the Gap in Text-Based Emotion Detection* (Muhammad et al., 2025b). Although this task involves multilabel emotion detection across 28 widely used languages from diverse regions of the world, our system is specifically designed for emotion detection in Hindi and Marathi. Since emotions are language-dependent, we leverage a pre-trained model specifically designed for Indian languages. We fine-tune the IndicBERT v2 (Doddapaneni et al., 2023), a pre-trained model specifically trained on 23 Indian languages to detect and classify emotions in Hindi and Marathi. Our contributions focus on enhancing emotion detection for underrepresented languages to advance research in multilingual NLP. Additionally, we investigate the effectiveness of pre-trained multilingual language models in emotion detection for the Indian languages, including Hindi and Marathi languages.

We conduct our experiments using the datasets provided by the organizers of SemEval-2025 Task 11 for Track A. Our results show that the proposed system achieves an F1 score of 87.37 for Hindi and 88.32 for Marathi, outperforming the baseline models. Additionally, we perform an error analysis to evaluate the effectiveness of our system in detecting emotions in Hindi and Marathi.

2 Related Works

Emotion detection in text is a key NLP task, enabling machines to interpret human emotions.

[1]SemEval2025-Task11: https://github.com/emotion-analysis-project/SemEval2025-Task11

Proceedings of the The 19th International Workshop on Semantic Evaluation (SemEval-2025), pages 2020–2027
July 31 - August 1, 2025

Lang	text	anger	disgust	fear	joy	sadness	surprise
hin	अरे वाह! आज तो मेरी बेटी ने अपने कमरे की ही नहीं, पूरे घर की सफाई खुद की है !	0	0	0	1	0	1
	बॉस ने आज मेरी क्लास ले ली यार!!	0	0	0	0	1	0
	मैंने घर के पास एक लम्बे काले साँप को देखा, तब से दिल की धड़कनें बढ़ी हुई हैं।	0	0	1	0	0	0
mar	सरकारच्या निर्णयामुळे जनतेला होत असलेले नुकसान अत्यंत तिरस्करणीय आहे.	1	1	0	0	0	0
	घराची कामे करून माझा दिवस सुरू होतो आणि मला छान वाटते आहे.	0	0	0	1	0	0
	हृदयाच्या खोलीला स्पर्श करणारी संगीताची उदास सूर दु:ख आणखी वाढवते.	0	0	0	0	1	0

Table 1: Samples from the dataset. Here, 1 and 0 represent the presence and absence of a particular emotion.

Lang	Train	Dev	Test	Total
Hindi	2,556	100	1,010	3,666
Marathi	2,415	100	1,000	3,515
English	2,768	116	2,767	5,651

Table 2: Dataset Statistics.

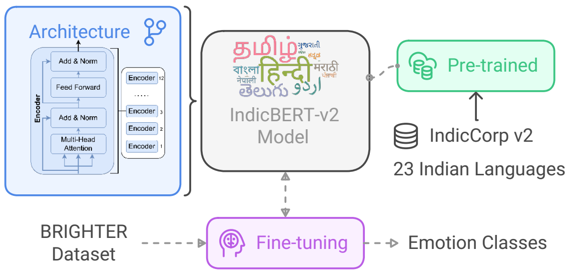

Figure 1: Model architecture and the fine-tuning process of the proposed sytem.

Early studies highlighted the impact of emotions in written communication, similar to face-to-face interactions (Wiebe et al., 2005). Emotion recognition has since been applied in healthcare, social media analysis, and conversational AI (Khanpour and Caragea, 2018; Saffar et al., 2023; Kang and Cho, 2024).

A major challenge in emotion classification lies in distinguishing between expressed and perceived emotions (Mohammad, 2022). Previous work explored multi-label classification of emotions in personal writings (Luyckx et al., 2012), using approaches like the Multi-label Maximum Entropy model (Li et al., 2016) and rule-based methods with affect lexicons (Al Masum et al., 2007). Deep learning advancements, including CNNs with self-attention (Kim et al., 2018) and CNN-LSTM models (Khanpour and Caragea, 2018), have improved fine-grained emotion detection. Transformer-based models have further enhanced performance, especially in textual conversations (Zhong et al., 2019; Jian et al., 2024).

Emotion Detection for Indian Languages: Most early research focused on high-resource languages such as English and Spanish. However, emotion recognition in Indian languages is gaining momentum. Given India's linguistic diversity, this task presents challenges due to script variations and cultural nuances.

For Hindi, lexicon-based methods like Hindi EmotionNet (Garg and Lobiyal, 2020b) and LSTM-based sentiment analysis on tweets (Gupta et al., 2021) have been explored. In Marathi, studies have applied TF-IDF-based supervised classifiers (Patil and Kolhe, 2022) and deep learning models for sentiment classification (Divate, 2021). Other Indian languages, such as Telugu and Bangla, have seen annotation efforts in sentiment analysis (Mukku and Mamidi, 2017; Kabir et al., 2023; Kumar et al., 2024). Despite progress, multi-label emotion detection for Indian languages remains underexplored, motivating our work.

3 System Overview and Experimental Setup

We fine-tune and evaluate IndicBERT-v2 (Doddapaneni et al., 2023) on the BRIGHTER dataset, using F1 Score as the primary performance metric. Additionally, to gain deeper insights into the model's performance across different emotions, we compute Accuracy (Acc), Precision (Prec), Recall, and F1 Score for each emotion.

3.1 Dataset

We focus on multi-label emotion classification for Hindi and Marathi, two Indian languages, using the BRIGHTER dataset—a human-annotated corpus designed for emotion detection across 28 languages, including Hindi and Marathi (Muhammad et al., 2025a). Each text sample in this dataset is labeled for the presence or absence of six emotions: joy, sadness, fear, anger, surprise, and disgust. This means that each snippet is annotated as joy (1) or no joy (0), sadness (1) or no sadness (0), and so on for all six emotions. Notably, the English subset of the dataset includes annotations for only five emotions joy, sadness, fear, anger, and surprise, excluding disgust. A few samples from the Dataset for Hindi (hin) and Marathi (mar) are listed in Table 1. The dataset statistics for Train, Dev, and Test sets are presented in Table 2. We also deep-dived into the dataset to get the count of each emotion present in the dataset. Most of the sample has a single emotion label in both languages. In the training set for Hindi, almost 24% and for Marathi, almost 17% of the sample is not having any emotion. More details are in the Appendix. We have used only the train set to train the model, and all the performance metrics are calculated with the test set.

3.2 Model Architecture and Training

One of the key challenges in developing emotion detection models for Indian languages is the limited support offered by most large language models (LLMs) and pre-trained language models, which are predominantly focused on English languages. Although models such as XLM-R (Conneau et al., 2020), mT5 (Xue et al., 2021), MuRIL (Khanuja et al., 2021), IndiSocialFT (Kumar et al., 2023), and IndicBERT-v2 (Doddapaneni et al., 2023) provide some level of support for Indian languages, they often fall short in handling the full spectrum of linguistic diversity. Among these, IndicBERT-v2 stands out as the state-of-the-art for many NLP tasks in Indian languages (Doddapaneni et al., 2023), making it the ideal choice for our emotion detection model.

The model architecture and the fine-tuning process of the proposed system are shown in Figure 1. To adapt IndicBERT-v2 for emotion detection, we fine-tune it using the BERTForSequenceClassification framework, treating the task as multi-label classification where a single text can express multiple emotions simultaneously, such as anger, fear, joy, sadness, or surprise. We employ binary cross-entropy loss (BCEWithLogitsLoss), which enables independent probability estimation for each emotion, ensuring a more flexible and accurate classification. The Adam optimizer is applied with a learning rate of 2×10^{-5}. The model is trained for 15 epochs with a batch size of 32. During training, logits from the classification head are converted to probabilities using the sigmoid activation function, and a threshold of 0.5 is applied during inference to determine emotion labels.

Model	Hindi	Marathi	English
LaBSE	75.25	80.76	64.24
RemBERT	85.51	82.20	**70.83**
XLM-R	33.71	78.95	67.30
mBERT	54.11	60.01	58.26
mDeBERTa	54.34	66.01	58.94
IndicBERT	**87.37**	**88.32**	66.32

Table 3: Performance of models in terms of F1 Score.

4 Result and Analysis

We evaluate the performance of our fine-tuned emotion detection model for both Hindi and Marathi and compare it against several baseline multilingual language models, including LaBSE, RemBERT, XLM-R, mBERT, and mDeBERTa. The evaluation is based on the F1-score, which serves as the primary metric. The baseline model scores are directly taken from the BRIGHTER paper (Muhammad et al., 2025a).

Table 3 presents the overall F1-scores for different models across Hindi, Marathi, and English. The scores, except for the Proposed model, are taken from Paper (Muhammad et al., 2025a). Our proposed model achieved the highest F1-scores for Hindi (87.37) and Marathi (88.32), outperforming all baseline models. For English, the proposed model attained an F1-score of 66.32, which is competitive but slightly lower than RemBERT (70.83). The superior performance of our model for Hindi and Marathi can be attributed to IndicBERT, a language-specific model trained for Indian languages, indicating that emotion recognition is highly influenced by language characteristics.

To gain deeper insights into the performance of our proposed system across different emotions, we analyzed the F1 scores for each emotion. Table 4 provides a detailed breakdown of performance of our system across all emotions in Hindi and Marathi. In Hindi, it achieved the highest F1-score for Fear (91.17), followed by Joy (89.66). Similarly, in Marathi, the best performance was observed for Disgust (92.86), followed by Fear (91.03) and Surprise (90.70). The consistently high accuracy across different emotions demonstrates the robustness of our approach. Since the test set also contains non-emotional samples, we evaluate the model's ability to detect non-emotional sen-

Emotion	Hindi				Marathi			
	Acc	Prec	Recall	F1	Acc	Prec	Recall	F1
Anger	94.75	80.00	89.44	84.46	95.90	**93.01**	81.10	86.64
Disgust	96.83	82.64	**90.09**	86.21	**98.60**	92.86	92.86	**92.86**
Fear	**97.52**	**94.16**	88.36	**91.17**	97.40	92.31	89.80	91.03
Joy	96.14	90.86	88.48	89.66	93.70	77.72	89.71	83.29
Sadness	94.16	78.35	89.94	83.75	93.90	84.76	85.99	85.37
Surprise	96.83	90.21	87.76	88.97	97.60	88.64	**92.86**	90.70
No-Emotion	97.03	87.58	92.41	89.93	94.70	90.80	79.57	84.81

Table 4: Performance of IndicBERT-based system across all emotions for Hindi and Marathi

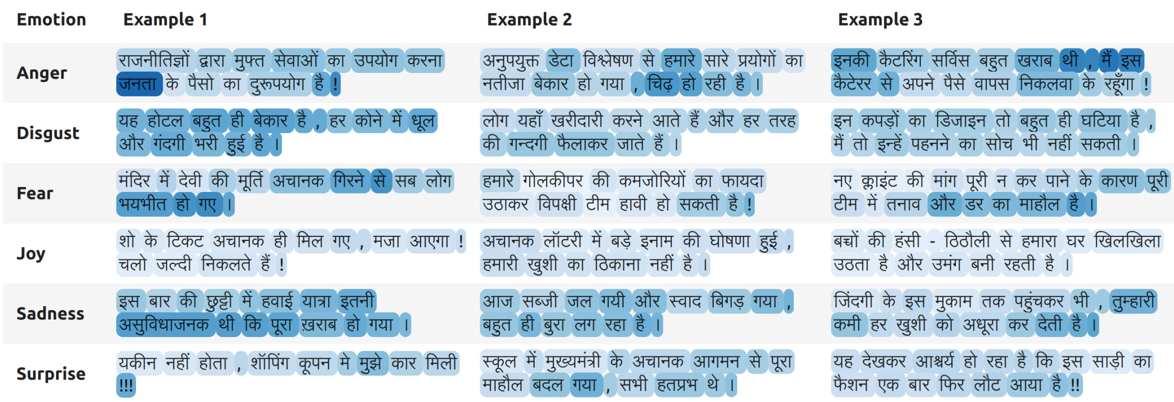

Figure 2: Heat map of the attention while predicting the emotion for some Hindi text.

tences in both languages. Our observations indicate that the model performs slightly better for Hindi in recognizing such neutral sentences.

To determine the interpretability of the system, we performed an attention heatmap analysis. We take three samples from each emotion class and visualize the attention weights to highlight the most influential words in the sentence. This helps in understanding how the model makes predictions by identifying keywords that contribute to each emotion classification. Attention heatmap visualization for the Hindi model is presented in Figure 2. By visualizing these attention weights, we verify that the model focuses on meaningful words while predicting emotions. Due to space constraints, the attention heatmap visualization for the Marathi model is provided in Appendix Figure 5.

4.1 Error Analysis

We conduct several error analyses to understand the error in the prediction of the emotion label by plotting the confusion matrix and manually examining misclassified samples. We plot a confusion matrix to analyze misclassifications between emotions categories. Figure 3 presents the confusion matrix for emotions categories of Hindi and Marathi languages. To construct these matrices, we expand multi-label samples and treat each label independently. The confusion matrix for Hindi reveals that while Joy is detected reliably it confuses with Surprise. Similarly, the model confuses among Anger, Disgust, and Sadness categories. In the case of the Marathi language, frequent misclassifications occur among Anger, Disgust, and Sadness, as well as between Fear and Sadness and Joy and Surprise. Notably, these confusing emotions share similar affective characteristics, explaining the overlap in classification. We also manually examine some of the misclassified samples to understand pattern misclassification. Table 5 presents a few selected samples of misclassified models. Our manual inspections of misclassified samples reveal that the model struggles with context-dependent interpretation, idiomatic expressions, and subtle emotional cues. For instance, in the case of *The Conjuring* movie, the model misclassifies fear as sadness due to a lack of awareness of the genre of the movie and its expected psychological impact. Similarly,

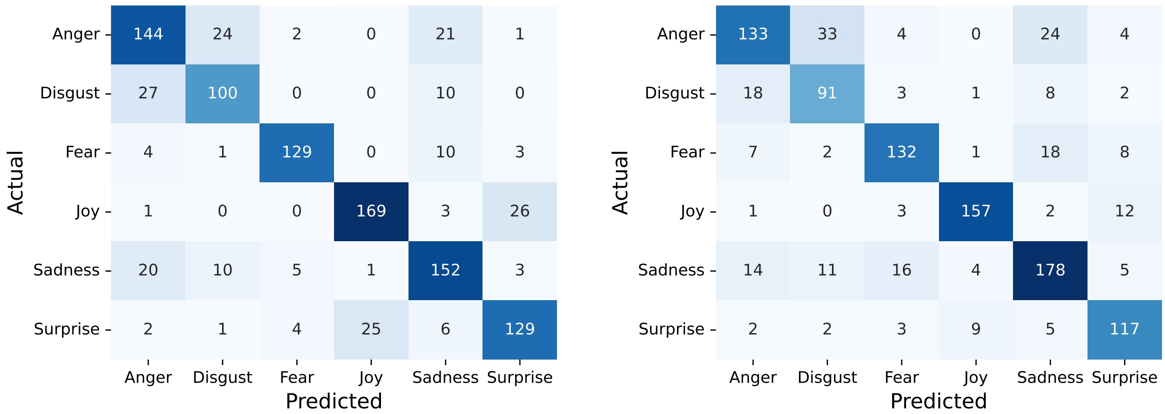

(a) Hindi (b) Marathi

Figure 3: Confusion matrix of the emotion prediction for both Hindi and Marathi model on Test data set

		text	Actual	Predicted
Hindi	1	"दा कंजूरिंग " फिल्म देखकर घर लौटने के बाद, रात भर नींद नहीं आई !	fear	sadness
		(After returning home from watching the movie "The Conjuring", I couldn't sleep the whole night!)		
	2	जब मैंने आईना देखा तो उसमे मेरा प्रतिबिम्ब ही नहीं था। मेरे तो हाथ-पांव फूल गए।	fear	sadness
		(When I looked in the mirror, my reflection was not there. I was completely terrified.)		
	3	हर चुनौती के बाद, मेरी उम्मीदों का दायरा धीरे-धीरे सिमटता जा रहा था, इस जीत का मिलना किसी चमत्कार सा लग रहा है !	surprise	joy
		(After every challenge, the scope of my expectations was gradually narrowing, getting this victory feels like a miracle!)		
	4	तुम कीचड़ में कैसे खेल लेते हो ??	disgust	anger
		(How do you play in the mud??)		
Marathi	5	कुटुंबातील नातेसंबंध जसेच्या तसे राहिले पाहिजे, अन्यथा कोणत्याही बदलामुळे तणाव वाढू शकतो.	fear	sadness
		(Relationships in the family should remain as they are, otherwise any change may increase tension.)		
	6	शाळेतील सहलींमुळे विद्यार्थ्यांमध्ये सहकार्याची भावना वाढते.	joy	no-emotion
		(School trips increase the sense of cooperation among students.)		
	7	घराच्या नवीन रेमोडेलिंगसाठी खूपच जास्त पैसे गेले, त्यामुळे तो आर्किटेक्ट माझ्या डोक्यात गेलाय.	anger	sadness
		(The new remodeling of the house cost a lot of money, so the architect went over my head.)		
	8	टिफिनमध्ये कुजलेले अन्न सापडले.	disgust	sadness
		(Rotten food found in tiffin.)		

Table 5: Examples of Misclassified Samples.

idiomatic expressions like *haath-pawn fool gaye*, which indicate fear, are misinterpreted as sadness. The model also confuses closely related emotions, such as joy and surprise, where unexpectedness plays a key role. Additionally, implicit emotions requiring external context, such as fear of change in a family setting, are often mistaken for sadness. The misinterpretation of figurative language, sarcasm, and cultural expressions further contributes to the misclassification of emotions, as seen in Marathi sentences, where frustration is classified as sadness. From such observations from misclassification and our annual inspection, we can conclude the need for contextual understanding improvements, idiomatic knowledge incorporation, and better differentiation between overlapping emotions to enhance the model's interpretability and accuracy.

5 Conclusion and Future Work

This study proposes a system for multilabel emotion detection in Hindi and Marathi languages. Our proposed system fine-tuning IndicBERT v2 on the BRIGHTER dataset provided organize of SemEVal-2025 Task 11. Our experimental results suggest that our proposed systems outperformed baseline models. We also have several ablation studies to understand the misclassification of emotions between different emotion categories by plotting a confusion matrix and manual inspections of misclassified samples. Our error analysis revealed that misclassification among emotion categories is due to confusion between similar emotions, misinterpretation of idiomatic expressions, and difficulty in capturing context-dependent emotions. To address these issues, future research will explore integrating knowledge extraction and narrative extraction, improving idiomatic phrase understanding, and incorporating multimodal cues to enhance emotion detection. Additionally, we aim to extend this work beyond Hindi and Marathi to include other widely used Indian languages such as Tamil, Bengali, Telugu, and Kannada.

References

Shaikh Mostafa Al Masum, Helmut Prendinger, and Mitsuru Ishizuka. 2007. Emotion sensitive news agent: An approach towards user centric emotion sensing from the news. In *IEEE/WIC/ACM International Conference on Web Intelligence (WI'07)*, pages 614–620. IEEE.

Ankush Chatterjee, Kedhar Nath Narahari, Meghana Joshi, and Puneet Agrawal. 2019. SemEval-2019 task 3: EmoContext contextual emotion detection in text. In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 39–48, Minneapolis, Minnesota, USA. Association for Computational Linguistics.

Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.

Manisha Satish Divate. 2021. Sentiment analysis of marathi news using lstm. *International journal of Information technology*, 13(5):2069–2074.

Sumanth Doddapaneni, Rahul Aralikatte, Gowtham Ramesh, Shreya Goyal, Mitesh M. Khapra, Anoop Kunchukuttan, and Pratyush Kumar. 2023. Towards leaving no Indic language behind: Building monolingual corpora, benchmark and models for Indic languages. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12402–12426, Toronto, Canada. Association for Computational Linguistics.

Kanika Garg and D. K. Lobiyal. 2020a. Hindi emotionnet: A scalable emotion lexicon for sentiment classification of hindi text. *ACM Trans. Asian Low-Resour. Lang. Inf. Process.*, 19(4).

Kanika Garg and DK Lobiyal. 2020b. Hindi emotionnet: A scalable emotion lexicon for sentiment classification of hindi text. *ACM Transactions on Asian and Low-Resource Language Information Processing (TALLIP)*, 19(4):1–35.

Vedika Gupta, Nikita Jain, Shubham Shubham, Agam Madan, Ankit Chaudhary, and Qin Xin. 2021. Toward integrated cnn-based sentiment analysis of tweets for scarce-resource language—hindi. *Transactions on Asian and Low-Resource Language Information Processing*, 20(5):1–23.

Zhongquan Jian, Ante Wang, Jinsong Su, Junfeng Yao, Meihong Wang, and Qingqiang Wu. 2024. EmoTrans: Emotional transition-based model for emotion recognition in conversation. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 5723–5733, Torino, Italia. ELRA and ICCL.

Mohsinul Kabir, Obayed Bin Mahfuz, Syed Rifat Raiyan, Hasan Mahmud, and Md Kamrul Hasan. 2023. BanglaBook: A large-scale Bangla dataset for sentiment analysis from book reviews. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 1237–1247, Toronto, Canada. Association for Computational Linguistics.

Yujin Kang and Yoon-Sik Cho. 2024. Improving contrastive learning in emotion recognition in conversation via data augmentation and decoupled neutral emotion. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2194–2208, St. Julian's, Malta. Association for Computational Linguistics.

Hamed Khanpour and Cornelia Caragea. 2018. Fine-grained emotion detection in health-related online posts. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1160–1166, Brussels, Belgium. Association for Computational Linguistics.

Simran Khanuja, Diksha Bansal, Sarvesh Mehtani, Savya Khosla, Atreyee Dey, Balaji Gopalan, Dilip Kumar Margam, Pooja Aggarwal, Rajiv Teja Nagipogu, Shachi Dave, et al. 2021. Muril: Multilingual representations for indian languages. *arXiv preprint arXiv:2103.10730*.

Yanghoon Kim, Hwanhee Lee, and Kyomin Jung. 2018. AttnConvnet at SemEval-2018 task 1: Attention-based convolutional neural networks for multi-label emotion classification. In *Proceedings of the 12th International Workshop on Semantic Evaluation*, pages 141–145, New Orleans, Louisiana. Association for Computational Linguistics.

Saurabh Kumar, Ranbir Sanasam, and Sukumar Nandi. 2024. IndiSentiment140: Sentiment analysis dataset for Indian languages with emphasis on low-resource languages using machine translation. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 7689–7698, Mexico City, Mexico. Association for Computational Linguistics.

Saurabh Kumar, Sanasam Ranbir Singh, and Sukumar Nandi. 2023. IndiSocialFT: Multilingual word representation for Indian languages in code-mixed environment. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 3866–3871, Singapore. Association for Computational Linguistics.

Jun Li, Yanghui Rao, Fengmei Jin, Huijun Chen, and Xiyun Xiang. 2016. Multi-label maximum entropy model for social emotion classification over short text. *Neurocomputing*, 210:247–256.

Kim Luyckx, Frederik Vaassen, Claudia Peersman, and Walter Daelemans. 2012. Fine-grained emotion detection in suicide notes: A thresholding approach to multi-label classification. *Biomedical informatics insights*, 5:BII–S8966.

Abdullah Al Maruf, Fahima Khanam, Md. Mahmudul Haque, Zakaria Masud Jiyad, M. F. Mridha, and Zeyar Aung. 2024. Challenges and opportunities of text-based emotion detection: A survey. *IEEE Access*, 12:18416–18450.

Saif Mohammad and Svetlana Kiritchenko. 2018. Understanding emotions: A dataset of tweets to study interactions between affect categories. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).

Saif M. Mohammad. 2022. Ethics sheet for automatic emotion recognition and sentiment analysis. *Computational Linguistics*, 48(2):239–278.

Shamsuddeen Hassan Muhammad, Nedjma Ousidhoum, Idris Abdulmumin, Jan Philip Wahle, Terry Ruas, Meriem Beloucif, Christine de Kock, Nirmal Surange, Daniela Teodorescu, Ibrahim Said Ahmad, et al. 2025a. Brighter: Bridging the gap in human-annotated textual emotion recognition datasets for 28 languages. *arXiv preprint arXiv:2502.11926*.

Shamsuddeen Hassan Muhammad, Nedjma Ousidhoum, Idris Abdulmumin, Seid Muhie Yimam, Jan Philip Wahle, Terry Ruas, Meriem Beloucif, Christine De Kock, Tadesse Destaw Belay, Ibrahim Said Ahmad, Nirmal Surange, Daniela Teodorescu, David Ifeoluwa Adelani, Alham Fikri Aji, Felermino Ali, Vladimir Araujo, Abinew Ali Ayele, Oana Ignat, Alexander Panchenko, Yi Zhou, and Saif M. Mohammad. 2025b. SemEval-2025 task 11: Bridging the gap in text-based emotion detection. In *Proceedings of the 19th International Workshop on Semantic Evaluation (SemEval-2025)*, Vienna, Austria. Association for Computational Linguistics.

Sandeep Sricharan Mukku and Radhika Mamidi. 2017. ACTSA: Annotated corpus for Telugu sentiment analysis. In *Proceedings of the First Workshop on Building Linguistically Generalizable NLP Systems*, pages 54–58, Copenhagen, Denmark. Association for Computational Linguistics.

Pansy Nandwani and Rupali Verma. 2021. A review on sentiment analysis and emotion detection from text. *Social network analysis and mining*, 11(1):81.

Rupali S Patil and Satish R Kolhe. 2022. Supervised classifiers with tf-idf features for sentiment analysis of marathi tweets. *Social Network Analysis and Mining*, 12(1):51.

Alieh Hajizadeh Saffar, Tiffany Katharine Mann, and Bahadorreza Ofoghi. 2023. Textual emotion detection in health: Advances and applications. *Journal of Biomedical Informatics*, 137:104258.

Carlo Strapparava and Rada Mihalcea. 2007. SemEval-2007 task 14: Affective text. In *Proceedings of the Fourth International Workshop on Semantic Evaluations (SemEval-2007)*, pages 70–74, Prague, Czech Republic. Association for Computational Linguistics.

Janyce Wiebe, Theresa Wilson, and Claire Cardie. 2005. Annotating expressions of opinions and emotions in language. *Language resources and evaluation*, 39:165–210.

Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. mT5: A massively multilingual pre-trained text-to-text transformer. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.

Peixiang Zhong, Di Wang, and Chunyan Miao. 2019. Knowledge-enriched transformer for emotion detection in textual conversations. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 165–176, Hong Kong, China. Association for Computational Linguistics.

A Insights of the Dataset

The majority of samples in the training set for both Hindi and Marathi contain a single emotion label. Figure 4 provides a visualization of the dataset statistics for both languages using an UpSet plot. From this, we observe that the most frequent multi-label emotion combinations in Hindi include (joy, surprise), (anger, disgust), (anger, sadness), and (sadness, disgust). Additionally, some samples exhibit less common combinations such as (joy, sadness), (anger, surprise), and (fear, anger, sadness). Similarly, in Marathi, multilabel instances frequently appear in (disgust, anger), (sadness, anger), (fear, sadness), (anger, sadness), (surprise, joy), and (disgust, anger, sadness).

Similar to the training set, most samples in the test set for both Hindi and Marathi contain a single emotion label. For Hindi, frequent multilabel emotion combinations include (anger, disgust), (anger, sadness), (sadness, disgust), and (joy, surprise). Additionally, a few instances exhibit rare combinations such as (anger, surprise), (joy, sadness), and (fear, anger, sadness). In Marathi, the most common multilabel occurrences are (disgust, anger), (sadness, anger), (fear, sadness), (anger, sadness), (surprise, joy), and (disgust, anger, sadness). These findings indicate that overlapping emotions in both languages often share affective similarities, making classification more challenging.

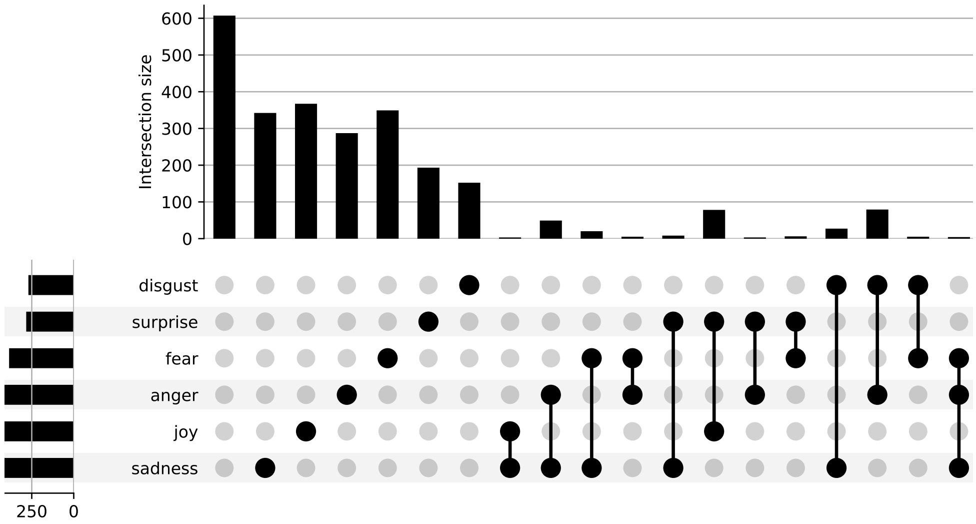

(a) Hindi

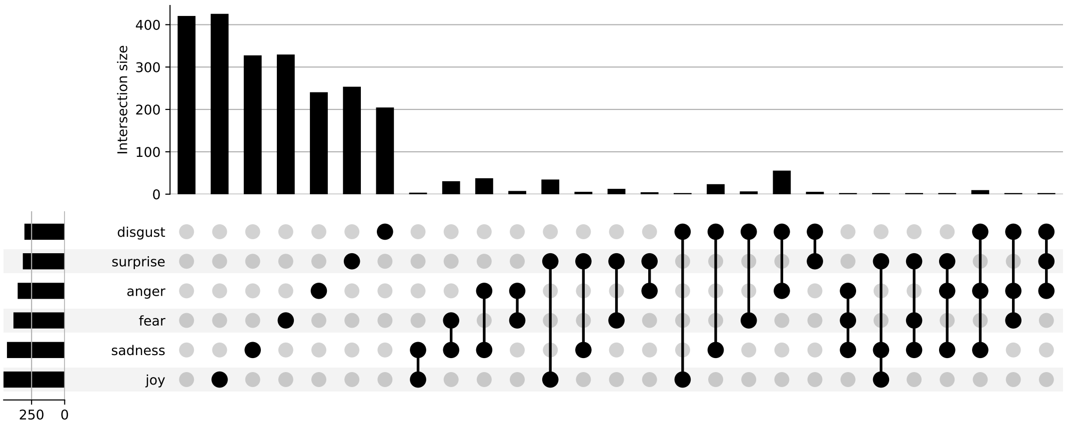

(b) Marathi

Figure 4: Emotion Wise Distribution of Training Data

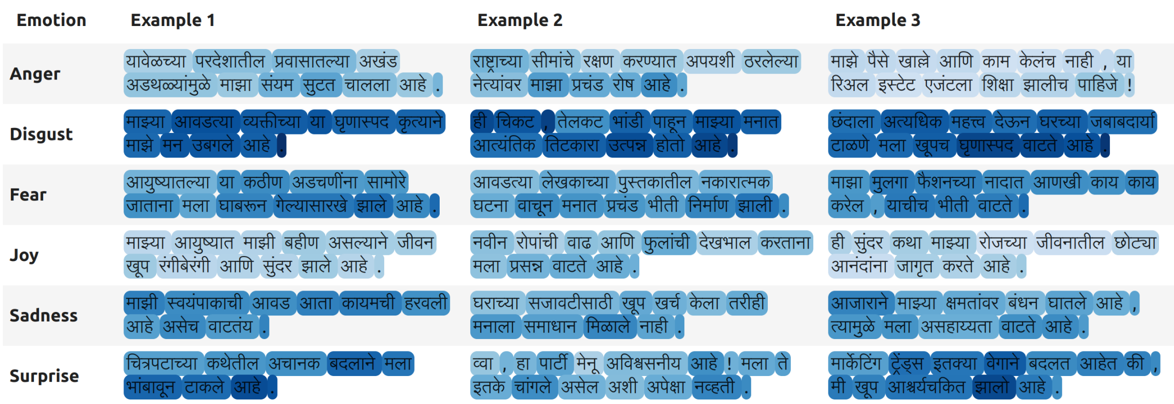

Figure 5: Heat map of the attention while predicting the emotion for some Marathi text.