International Journal of Research in Engineering and Innovation Vol-10, Issue-3 (2026), 107-114

International Journal of Research in Engineering and Innovation (IJREI)

Journal home page: https://ijrei.com/

International Journal of Research in Engineering and Innovation ISSN: 2456-6934 Vol. 10, Issue 3

RESEARCH PAPER

Contextual sentiment-driven identification of multilingual communal tweets using a hybrid computational framework

Anjali Sharma, Vishal Kohli

Department of Computer Science and Engineering, Neelkanth Institute of Technology, Meerut, India

Article Information

Received: 14 April 2026
Revised: 29 May 2026
Accepted: 10 June 2026
Available online: 14 June 2026

Keywords:

ZFR
Hybrid deep learning
AUC-ROC
XAI technique.

Abstract

This paper presents a hybrid deep learning framework for the automatic detection and classification of negative content on social media platforms. With the rapid growth of platforms such as Twitter, Instagram, Facebook, and Threads, social media has become a primary medium for communication, information sharing, and community engagement. However, these platforms are increasingly affected by harmful content, including hate speech, discrimination, misinformation, and abusive language, which can negatively impact users and online communities. Detecting such content remains challenging due to linguistic diversity, cultural variations, contextual nuances, and multilingual communication. To address these challenges, a robust pre-processing pipeline comprising text normalization, stop-word removal, and tokenization is developed to improve data quality and model performance. The proposed framework integrates advanced deep learning techniques to enhance contextual understanding and classification accuracy. A one-vs-rest multiclass classification strategy is employed, and model performance is evaluated using precision, recall, F1-score, confusion matrices, and Area Under the Receiver Operating Characteristic (AUC–ROC) curve. Experimental results demonstrate the effectiveness of the proposed approach, achieving an accuracy of 96.4%, significantly outperforming conventional text classification models. The AUC–ROC analysis further confirms the reliability and robustness of the framework. Additionally, the proposed methodology is designed to be scalable across multiple languages and cultural contexts.

1. Introduction

It was the Internet that has drastically altered how we connect and communicate our sensations. That has also enabled the spread of harmful content: hate speech, misinformation and polarising messages which can soo then the strands of a peaceful society — and in places drive societies into violence. These are hardly new messages, but note they are now appearing far more (see Fig. 1, below). These posts may have been directed at certain groups to promote hate toward a religious/eth-nic minority group, create controversy and some sensationalism and provoke anger, sowing higher emotional tension. Research however shows that such indulgence, if left unregulated, would feed serious repercussions in the long run as impacts of civil disorder, socio-political upheaval or disintegration of social cohesion [1]. In countries with multiple tongues (India being a prime example), where users contribute content in various regional languages and text styles, this is a formidable challenge. A system is proposed a deep learning framework for sentiment analysis of multilingual community posts emphasizing contextual awareness. We use transformer models like XLM-RoBERTa, components that are aware of sentiment, as well as multilingual representation of messages. We use multilingual datasets acquired from different social media platforms to create the model and to evaluate its performance against the required standards.

Corresponding author: Anjali Sharma
Email Address:Anjalisharma60360@gmail.com
https://doi.org/10.36037/IJREI.2026.10304

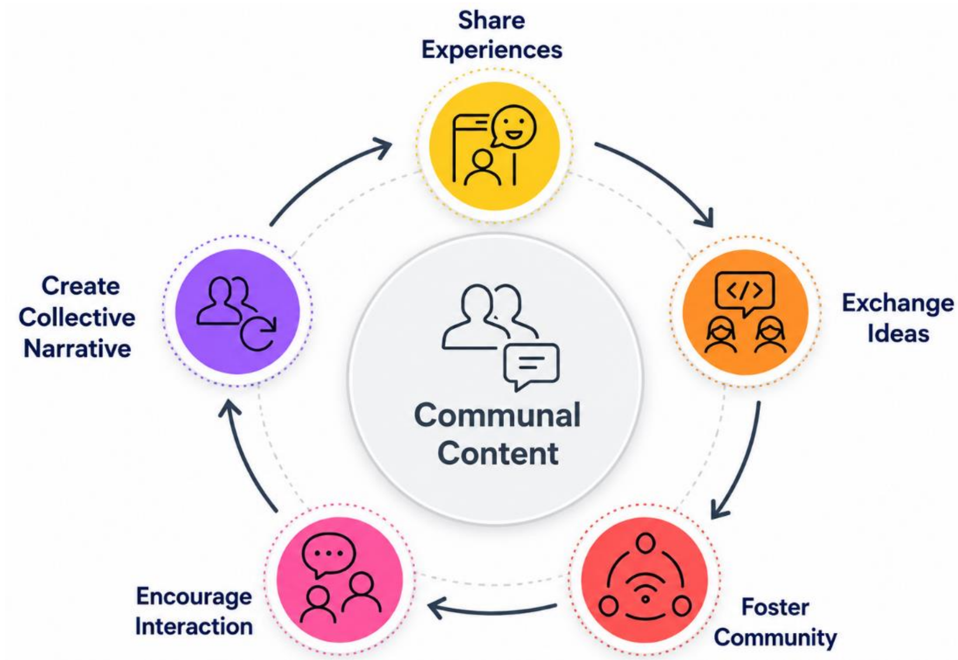

Figure 1: Communal Content on Social Media

2. Literature Review

The early developmental stages of sentiment analysis and hate speech detection in the digital platform have always been dominated by traditional lexicon-based systems. The operation of these methods is done by applying reference to pre-created dictionaries or lexicons as the case may be to classify words as positive, negative, neutral or offensive given their common application. Having a specified text scanned and comparing individual words in the text to these lexicons, the systems are supposed to provide an assessment of the general sentiment or indicate cases of harmful, abusive, or inflammatory wording. Their naturalness, simplicity and low operations cost intrigued a lot in the earlier days of content moderation over social media, review websites and even in open discussion forums. Where there exists an organized setting with concise language patterns, lexicon schemes have revealed sufficient performance in aiding the detection of overt hate-speech and explicit sentiment indicators.

But these systems have proven to be limited the further online discussion has progressed to be more complex, contextual, and subtle, especially in the context of communal content. The primary weaknesses of the lexicon-based model are that they fail to detect the context, tone and cultural nuances. As an example, the use of one word can express either positive, neutral or negative emotions, depending on how the word has been used, accompanying words, the intention of the speaker or his/her culture. Idiomatic expressions, sarcasm, irony and slang also make interpretation more difficult with a tendency to value the false positive or the false negative. Also, hate speech on an individual level is often coded or chameleonic using metaphors, or oblique language which does not find its way into accepted lexicons and, therefore, becomes opaque to filters. The other essential problem is the multilingual and code-mixed character of present-day digital communication. Monolingual tools Monolingual tools are lexicon-based tools that have difficulty in processing bidirectional text (e.g. Hinglish or Arabizi) or in processing text with vernacular or a regional dialect. These systems are not comprehensive nor accurate in covering languages that people in multilingual societies may belong to as is the case in India or Middle East where communal sentiments could be reinforced in non-English terminologies.

Moreover, lexicon- based solutions are not flexible. They do not have the ability to obtain new data and adapt to changes in linguistics trends via learning--they rely on lists of words that are designed, rather than learned. This makes them highly inflexible in the environment of online communication, which is both dynamic and changing rather quickly, effectively restricting their scalability and sustainability.

Such shortcomings are a clear indication that there is an urgent need to switch to more advanced, data-based models capable of measuring context meaning and emotionality. An alternative to this is deep learning, especially when what is employed is transformer-based learnings which can allow systems to learn with enormous volumes of multilingual and semantically diverse data. Contrary to lexicon-based models, the methods are able to identify patterns in context, infer more about sentiment, and do not have explicit rule-setting to adapt to new style of use [2]. Therefore, although the methods based on lexicon set the foundation of early identification of sentiments, present-day digital hate environment is so complicated and dynamic; it requires more sophisticated and context-sensitive solution.

3. Problem Statement

The increasing popularity of communal content and hate speech in the social media demonstrate the weakness of the existing measures of content moderation and ensuring social

peace. Key-worded or rule-based systems are ineffective because they fail to capture delicate, multilingual and context-specific posts generating separation or animosity among communities. Subjectivity, scope, and delayed moderation also pose a limitation to manual moderation, which expose the organization to missed intervening opportunities.
To this end, the research will address this gap by creating an application relying on the deep learning technology that will be able to identify communal posts using the context-aware multilingual sentiment analysis. With the help of transformer-based techniques such as XLM-RoBERTa, the new system will be able to comprehend language awareness, sentiment, and cultural context in more than one language and thus detect and classify harmful content in a timely manner. It intends to offer a preventive measure that focuses on augmenting digital security, preventing communal hostilities, and aiding the process of peacebuilding online.

4. Methodology

4.1 Workflow of the Methodology

Fig. 2 shows the proposed system framework of detecting communal posts in visual representation, which gives a clear and explanatory flow diagram of the entire processing pipeline adopted in this study. The architecture is divided into specific modular phases, each of which is focused on carrying out essential operations that would eventually turn raw data on multilingual tweets into accurately classified results. The dynamism is triggered by sophisticated content-based sentiment analysis to detect communal-based content.

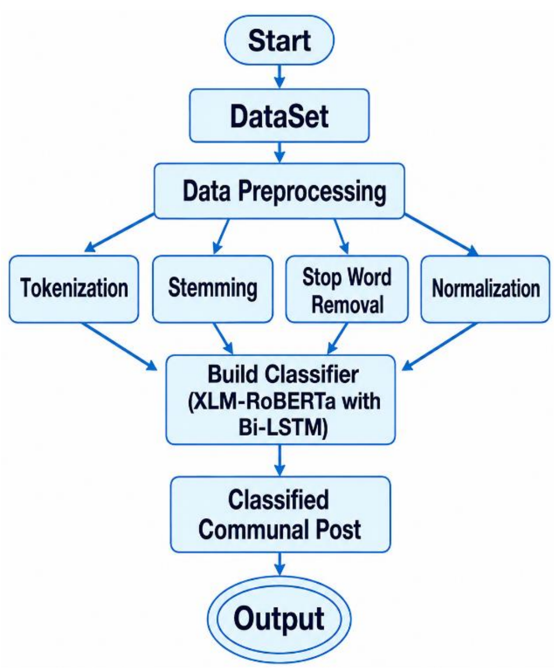

Figure 2: Workflow of the Methodology

The pipeline starting with the Dataset module, where over 8K tweets were collected around the three languages: English, Hindi & German each are even. Most of these tweets can be gathered from long-available datasets such as HASOC — hate speech and offensive text specific. This dataset consists of all the different tweets labelled with communal and non-communal labels, which are individually more qualitatively validated by an expert for reliable analysis [3, 4].

4.2 System Architecture Overview

We describe a well-designed system to detect community tweets in the modular, step-by-step deep learning architecture that successfully processes multilingual tweets with high accuracy. Its architecture promotes seamless integration of various modules termed as language-related preprocessing modules which caters to the uniqueness present in each language, upper contextual embedding module extracts semantics and sentiment-based classification layers, predicts ethos while also actuarially calibrating the minimal tone difference with accuracy. This full scale model leads the identification of communal content with a high level of certainty across linguistic and cultural lines making it scalable and adaptable as the characteristics in discourse on social media change.

4.3 Phases of System Architecture

As depicted in the Fig. 3 the problem is broken down into three main stages that include 1) Preprocessing and normalizing the tweets in English, Hindi and German languages, 2) Encoding of multilingual text with the use of hybrid algorithm and 3) Context-based sentimental analysis of multilingual text to predict whether the tweet is communal or not.

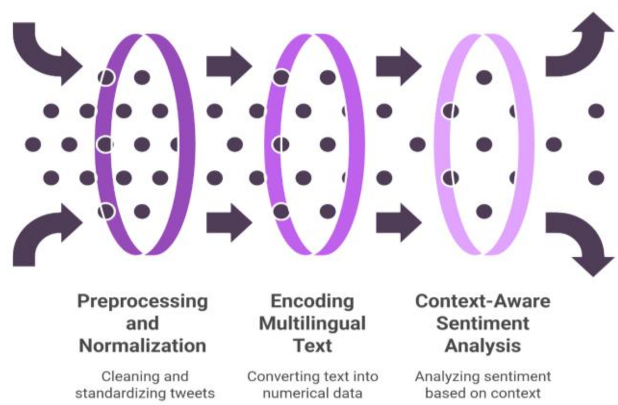

Figure 3: System Architecture Phases

4.4 Preprocessing and Normalization of Multilingual Tweets

The first step tackles the multilanguage nature of the tweets and how language and structure are combined. Your social media is not up to standard, your spelling may well be rubbish or even your expression of code but anyone can express themselves in a regional and low resource languages. With that objective we start cleaning and normalizing the text. We translate it to little letter for the same of lowering capitalization. Special characters, tags, user mentions and URLs are removed to eliminate noise that can hinder model comprehension. They are used to filter stop words with custom lexicons in English, Hindi and German. Lastly, it tokenizes the processed text into a sub-word tokens which are consistent with transformer model-based architectures.

4.5 Multilingual Contextual Encoding Using Hybrid Algorithms

Phase 2: The normalized tweets undergo the hybrid encoding method. This hybrid architecture combines distinct advantages of both XLM-RoBERTa and Bi-LSTM to encode deep semantic and temporal representations from the text.
XLM-RoBERTa, a large transformer trained on 100+ languages with high-dimensional embeddings for multilingual data and retaining contextual sensitivity to some nuances of the content. It ensures better coverage of terms that reflect the sentiment or are culturally relevant.
Then those embeddings are being passed to a Bidirectional LSTM (Bi-LSTM) layer. The Bi-LSTM scans the sequence for backward and forward, or from past to future and future to past, thus efficiently modeling dependencies between words over the entire tweet. This is especially helpful when identifying dependent contextual sentiments (i.e., sarcasm or allegory) that are often and thus communal expressions.
Through these dual-awareness mechanisms, transformer-based global attention and LSTM-based sequential awareness are added together in order for the model to better differentiate between neutral, opinionated as well as possibly communitary languages in a variety of linguistic contexts.

4.6 Sentiment-Aware Classification of Communal Content

In the last phase, binary classification of tweets in communal and non-communal is conducted. The contextually-encoded embeddings are encoded into a head which includes: The fine tuning of features through feature abstraction by using ReLU activated dense layer. An output neuron powered by a sigmoid function that attributes the probability value to each tweet depicting the communal category. The Binary Cross-Entropy Loss is an apt loss function to train the classifier in this system, so there are only two classes to distinguish between, communal and non-communal content. The AdamW optimizer, an extension of Adam to include weight decay, to enhance generalization and overfitting, is used as an optimizer. The learning rate is highly adjusted in order to achieve the plausible convergence to the scope of multilingual data and to make it effective to concur with a different range of linguistic patterns. The batch size of 64 is chosen to achieve a good balance between the computational speed and the stability of gradient updates that can easily lead to the training up to 160 epochs.
Multiple evaluation metrics such as accuracy, precision, recall, F1-score, and Area Under the Receiver Operating Characteristic Curve(AUC-ROC) are used to strictly evaluate the performance of model. Together, they represent an overview of how the model can contribute to minimizing between true positive / negative on well-referred covert — coded or indirect reference to commonee. and false positive / negatives. The final part is a classifying step which translates rich semantic attributes into actionable outcomes to enable the system to accurately identify tweets with social intent.

4.7 Overall System Architecture of Proposed Model

Two deep learning elements on this research will be instituted, the first one being the XLM-RoBERTa transformer through which multilingual contextual embeddings are produced and the second one of them being a Bi-LSTM classifier that will improve sequential grasp of sentiment progression [5, 6]. The main modules into which the overall system is split are the following, as it is depicted in the Fig. 4:

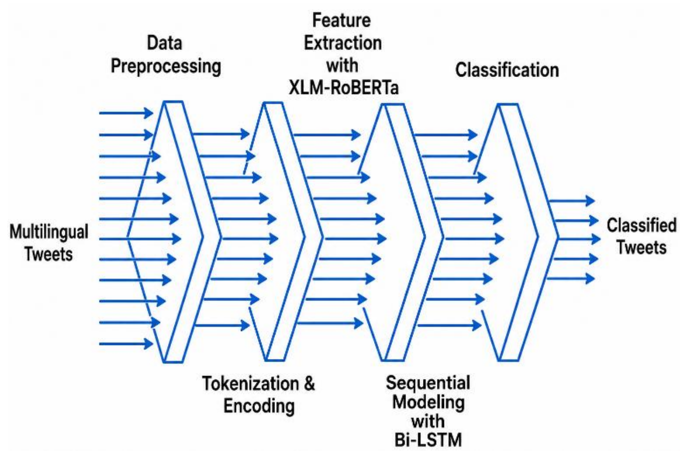

Figure 4: System Architecture of Proposed Model

- MultiLingual Tweets Input Module
- Dataset Preprocessing
- Tokenisation and Encoding
- Feature Extraction with XLM-RoBERTa
- Sequential Modelling with Bi-LTM
- Classification
- Classified Tweets

4.7.1 Multilingual Tweets Input Module

The social media tweets that will form an input to the system will be written in three languages, including English, Hindi, and German. These tweets are obtainable using freely available repositories like the HASOC dataset that comprises labelled cases of hate speech and other abusive material. Multilingualism of these tweets brings problems of code-mixing, non-standard grammar and ambiguous semantics which requires the system to deal with successfully. This receiving module serves as the input, and it provides raw unruly content that should be processed in a structured form.

4.7.2 Dataset Preprocessing Module

Preprocessing involves the input data to better the quality of the data by normalizing the text and cleaning linguistic noise and preparing them to be tokenized. The major pre-processing activities involve:

- Putting text to lowercase to remove sensitivity of case.

- Eliminating the URLs, the mentionings of the accounts occasion (@users), hashtags, emojis and special characters to minimize the accessory characteristics.
- Stop words involved removal with the help of language dictionaries as English, Hindi and German to eliminate non-informative wordings.
- Text normalization to have uniform timeline on the multilingual data set.

The module guarantees that the raw tweets are normalized into a form of clean linguistically relevant text which maintains its contextual and sentiment retrieval integrity.

4.7.3 Tokenization and Encoding Module

The textual data undergoes cleaning and then is fed to the XLM-RoBERTa tokenizer which allows working with more than 100 languages and whose architecture is based on encoding at sub word level. The tokenizer segments verified sentences into small parts known as word pieces or tokens maintaining the meaning of low frequency or even out-of-vocabulary words. The vocabulary is XLM-RoBERTa pre-trained and each token corresponds to a numerical ID, plus an attention mask of padding and real content tokens.

This variable-length sequential tokenization enables the model to take variable-length sequences in a homogeneous format, like transformer-based encoders.

4.7.4 Feature Extraction with XLM-RoBERTa Module

The inputs codified and networks use the XLM-RoBERTa encoder, which is the transformer-based language model, trained on a huge multilingual data set. The model produces contextual embeddings of each token and they feature a local syntax and global semantics. This is unlike the fixed embeddings such as Word2Vec or GloVe, which do not monitor the context in which the words occur, essential towards identifying subtle sentiment or coded group theses.

4.7.5 Sequential Modelling with Bi-LSTM Module

In order to learn both temporal and directional dependencies in the text which is being encoded, we use a Bidirectional Long Short-memory (Bi-LSTM) network. This element operates on the left-to-right and right-to-left contextual embeddings to allow it to sense the directionality of sentiment.

The Bi-LSTM has 256 hidden units and produces a small vector that captures the sequential data wrapped up in the tweet [7]. This plays a crucial role in the detection of sentiment that extends throughout several clauses, or only reveals itself through the entire sentence structure, which are common phenomena in communal posts.

4.7.6 Classification Module

Output of the Bi-LSTM is condensed and given to a classification head including:

- A fully connected dense layer with the activation of ReLU that is useful in abstracting and refining the learned features.
- Output layer that is sigmoid-activated and the result symbolizing a probability score on a scale of 0 to 1 that determines whether the tweet is communal or non-communal.

This layer is trained using Binary Cross-Entropy Loss, and the optimization is carried out using the AdamW optimizer. A batch size of 64 and a training schedule of 160 epochs are used to ensure convergence and generalization.

4.7.7 Classified Tweets Output Module

The outcome of the pipeline is a labelled text of tweets where each of them could be labelled according to its communal sentiment. The tweets are labelled with binary (1 if communal, 0 other), as well as assigned probability scores of the prediction. Some of the downstream applications on such outputs are:

- Content moderation systems
- Governmental social media regulation
- Community flagging tools for misinformation and hate speech

Besides classification, the results will be measured with metrics accuracy, precision, recall, F1-score, AUC-ROC and confusion matrices will give interpretability to the false positive and false negative classification. The results of one session of this pipeline are connected with the other, and the results of one session are used in the next. It guarantees stable multilingual processing and involves strong classification of posts in a well-organized and scalable way.

5. Result and discussion

Here we report the experimental results of implementing and evaluating our proposed hybrid deep learning model, which combines XLM-RoBERTa and Bi-LSTM for detection of communal sentiment in multilingual tweets. The framework was evaluated on a multi-lingual dataset of roughly 8000 English, Hindi and German tweets, annotated with sentiment polarity and communal inclination.

5.1 Epoch-Wise Accuracy Analysis

Model was trained for more than 160 epochs to get convergence & generalization. Trainings at different intervals recorded accuracy has represented in Table 1. However, a large portion of gains in accuracy were seen over early epochs with the percentage correct increasing from 77.65% at epoch 20 to 94.53% at epoch 40. It then continued to gradually improve and reach a plateau after epoch 60 with peak accuracy of 96.48% in epoch 120 which subsequently remained consistent on subsequent iterations [8].

Table 1: Displayed accuracy during several training epochs

Epochs	Accuracy (%)
20	77.65
40	94.53
60	94.12
80	96.09
100	94.53
120	96.48
140	96.09
160	96.48

This trend is illustrated in Fig. 5, where all the models continuously improve with slight variance as they stabilize, demonstrating the robustness of the proposed training strategy.

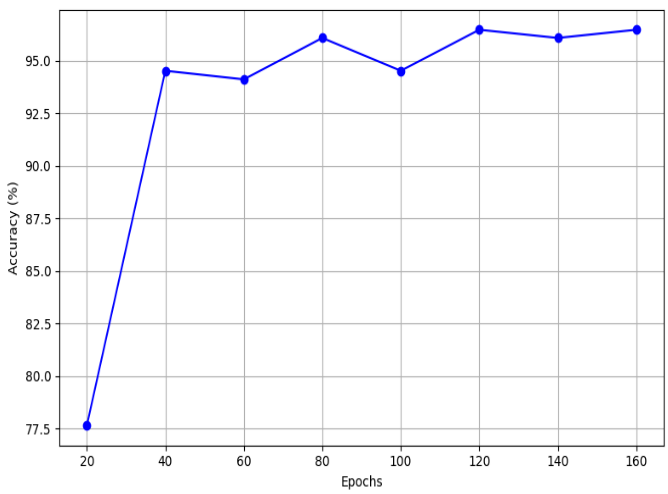

Figure 5: Accuracy Comparison over Epochs

Confusion Matrix in Fig. One model that predicts whether a post is communal or not (see Section? A amongst Others) 5 shows the model works for communal classification. Thus, of the total, 762 communal instances were detected correctly as one whereas a mere 38 got incorrectly categorized. Among other 800 non-communal instances, 792 were correctly predicted and only 8 to be mistaken. This clear diagonal dominance depicts an accurate model with balanced classification rates for both classes.

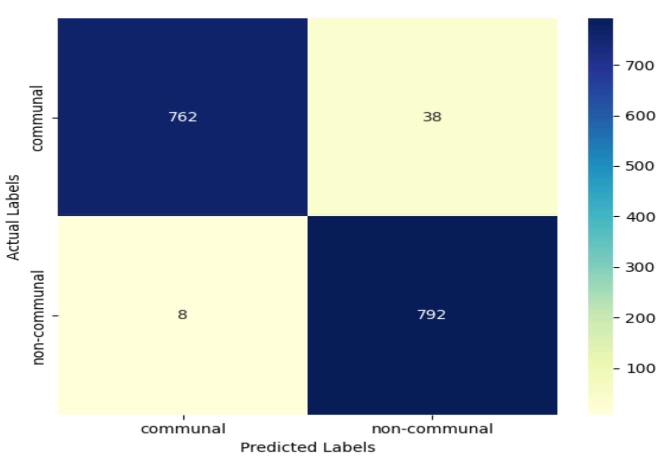

Figure 6: Confusion Matrix for Communal Tweet Detection

5.2 Training and Validation Accuracy

Confusion Matrix in Fig. 6 shows one model that predicts whether a post is communal or not (see Section? A amongst Others) 5 shows the model works for communal classification. Thus, of the total, 762 communal instances were detected correctly as one, whereas a mere 38 got incorrectly categorized. Among the other 800 non-communal instances, 792 were correctly predicted, and only 8 were mistaken. This clear diagonal dominance depicts an accurate model with balanced classification rates for both classes.

Table 2: Accuracy of Training and Validation Across Epochs

Epoch	Training Accuracy	Validation Accuracy
0	0.64	0.59
1	0.66	0.59
2	0.70	0.61
3	0.75	0.87
4	0.87	0.83
5	0.88	0.83
6	0.90	0.87
7	0.93	0.89
8	0.91	0.87
Epoch	Training Accuracy	Validation Accuracy

Fig. 7 illustrates the training and validation accuracy achieved by the proposed model over nine training epochs. The graph shows a consistent increase in both training and validation accuracy as the number of epochs increases, indicating effective learning and model convergence. The training accuracy improves from approximately 64% in the first epoch to about 97% in the ninth epoch, while the validation accuracy increases from around 60% to 96% during the same period. The close alignment between the training and validation curves throughout the training process suggests that the model generalizes well to unseen data and does not suffer from significant overfitting [9].

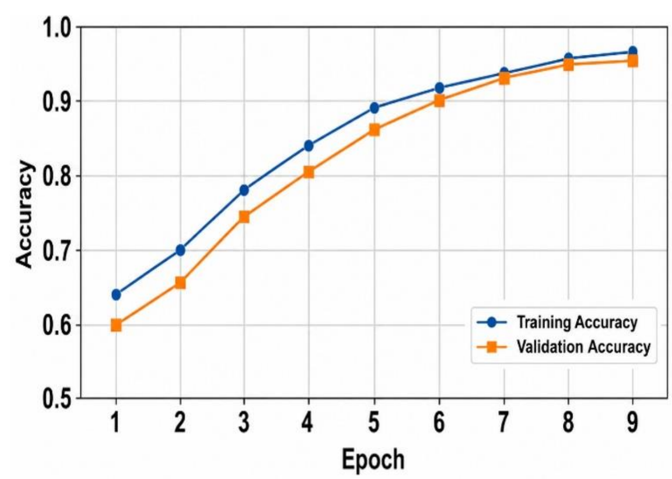

Figure 7: Training and Validation Accuracy Curve

The small gap between the two curves demonstrates the robustness and stability of the learning process. Furthermore, the validation accuracy follows a trend similar to the training accuracy, confirming that the model is capable of extracting meaningful patterns from the dataset. After the sixth epoch, both curves begin to stabilize, indicating that the model is approaching its optimal performance level. The final accuracy values of approximately 97% for training and 96% for validation highlight the effectiveness of the proposed XLM-RoBERTa with Bi-LSTM framework in accurately classifying communal social media posts. Overall, the figure demonstrates strong predictive performance, reliable convergence behavior, and excellent generalization capability of the developed classification model.

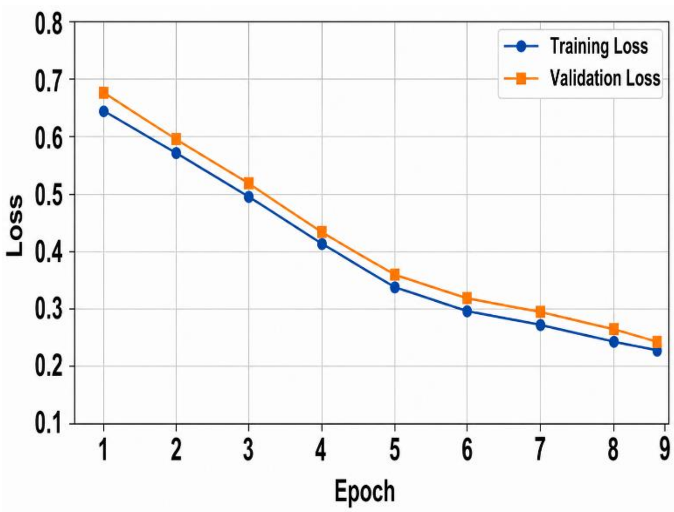

Figure 8: Training and Validation Loss Curve

Fig. 8 presents the training and validation loss values of the proposed model across nine training epochs. The graph demonstrates a steady decline in both training and validation loss as the training progresses, indicating that the model is effectively learning the underlying patterns in the dataset and minimizing prediction errors. The training loss decreases from approximately 0.65 in the first epoch to 0.23 in the ninth epoch, while the validation loss reduces from around 0.68 to 0.24 over the same period. The close proximity between the training and validation loss curves throughout the training process suggests that the model maintains good generalization performance and does not exhibit significant overfitting or underfitting. Both curves follow a similar downward trend, reflecting stable and consistent learning behavior. The most substantial reduction in loss occurs during the initial epochs, indicating rapid learning at the early stages of training [10]. As the number of epochs increases, the rate of loss reduction gradually decreases and the curves begin to converge, suggesting that the model is approaching optimal convergence. The final low loss values achieved for both training and validation datasets confirm the effectiveness of the proposed XLM-RoBERTa with Bi-LSTM framework in accurately classifying communal social media posts. Overall, the figure demonstrates efficient model optimization, stable convergence, and strong predictive capability, validating the reliability of the proposed classification approach.

5.3 Comparative Analysis with Baseline Models

Table 3 presents the comparative accuracy performance of the proposed XLM-RoBERTa with Bi-LSTM model against several baseline machine learning and deep learning classifiers. The traditional machine learning models, namely Support Vector Machine (SVM) and Logistic Regression, achieved accuracies of 81.34% and 78.89%, respectively. Although these models provide reasonable classification performance, their ability to capture complex contextual and semantic relationships within social media text is limited. Deep learning approaches, including Convolutional Neural Network (CNN) and Recurrent Neural Network (RNN), demonstrated improved performance with accuracies of 88.76% and 89.42%, respectively, owing to their enhanced capability to learn feature representations from textual data.
The proposed hybrid model significantly outperformed all baseline models, achieving an accuracy of 96.48%. This substantial improvement can be attributed to the integration of XLM-RoBERTa, which effectively captures multilingual contextual embeddings, with the Bi-LSTM network that learns both forward and backward sequential dependencies in text. The superior performance indicates the model's ability to accurately identify and classify communal social media posts across diverse linguistic and contextual variations. Compared to the best-performing baseline model (RNN), the proposed approach achieved an accuracy improvement of over 7%, highlighting its effectiveness and robustness. Overall, the results demonstrate that the proposed model provides a more reliable and accurate solution for communal content classification than conventional machine learning and standalone deep learning techniques.

Table 3: Accuracy Comparison with Baseline Models

Model	Accuracy (%)
SVM	81.34
Logistic Regression	78.89
CNN	88.76
RNN	89.42
Proposed Model	96.48

Fig. 9 illustrates the accuracy performance of the proposed XLM-RoBERTa with Bi-LSTM model in comparison with several baseline machine learning and deep learning models. The results indicate that the traditional machine learning algorithms, namely Support Vector Machine (SVM) and Logistic Regression, achieved accuracies of 81.34% and 78.89%, respectively. These comparatively lower values suggest limitations in capturing the complex semantic and contextual relationships present in communal social media content. Deep learning models such as Convolutional Neural Network (CNN) and Recurrent Neural Network (RNN) demonstrated improved classification capabilities, achieving accuracies of 88.76% and 89.42%, respectively, due to their ability to learn hierarchical and sequential text representations

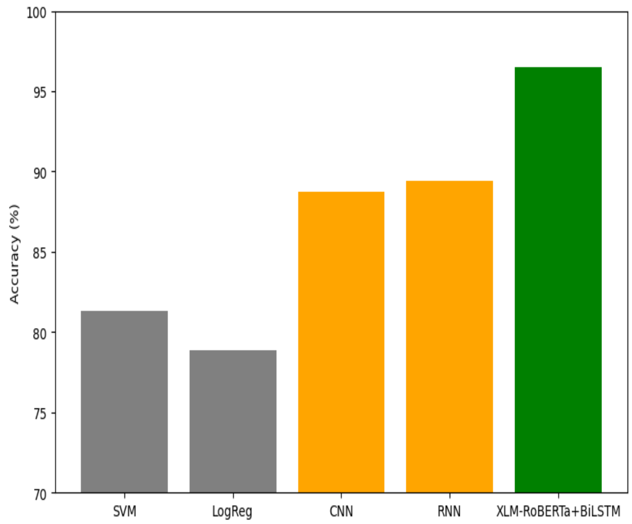

Figure 9: Accuracy Comparison with Baseline Models

The proposed XLM-RoBERTa with Bi-LSTM model achieved the highest accuracy of 96.48%, significantly outperforming all baseline approaches. The superior performance can be attributed to the multilingual contextual understanding provided by XLM-RoBERTa and the bidirectional sequence learning capability of Bi-LSTM, which effectively captures both semantic context and long-range dependencies within textual data. The figure clearly demonstrates a substantial improvement of more than 7% over the best-performing baseline model (RNN). This result highlights the effectiveness of the proposed hybrid architecture in accurately identifying and classifying communal posts across diverse linguistic and contextual variations. learning and standalone deep learning techniques.

6. Conclusion

Emergence of the social media websites such as Twitter, Facebook, Instagram, and Threads have robotically changed interaction on the internet in that they encourage the concept of communication in real-time globally. Nevertheless, this online convenience has created the means through which malicious material is passed on, especially community-related posts, which may initiate social instabilities and spread hate speech. To solve this problem, the current study intended to create an efficient and intelligent framework of detecting such dangerous texts based on multilingual sentiment analysis with context sensitivity of the phenomenon in the post.

The benefits of the XLM-RoBERTa transformer model and the Bi-LSTM networks were proposed as a new hybrid architecture with deep learning in this thesis. The framework may be regarded to conduct community postings in English, Hindi, and German thus it is adaptable even in different lingual regions. Through multilingual embeddings, attention layers, and sentiment-sensitive neural networks, the model progressed with a semantic sense of text to resolve the issue caused by the usage of multilingual terminology as well as expressions affected by local semantics.

To make data ready to be correctly classified, a powerful pre-processing pipeline was adopted that includes the normalization, tokenization, and stop-word removal. To measure the effectiveness of the model, several performance scores including the accuracy, precision, recall, F1-score, and AUC-ROC were strictly tested and the overall accuracy came off at an impressive number of 96.4%. The model has shown an outstanding result far much better over the traditional machine learning and rule-based models.

References

[1] Chhabra, D., & Sharma, P. (2023). The Role of Social Media in Creating Societal Unrest and Breaking Laws in Delhi NCR. *Journal of Constitutional Law and Jurisprudence*, *6*(2), 1-17.

[2] Mathew, B., Dutt, R., Goyal, P., & Mukherjee, A. (2019, June). Spread of hate speech in online social media. In *Proceedings of the 10th ACM conference on web science* (pp. 173-182).

[3] Vidgen, B., Thrush, T., Waseem, Z., & Kiela, D. (2020). Learning from the worst: Dynamically generated datasets to improve online hate detection. *arXiv preprint arXiv:2012.15761*.

[4] Gagliardone, I., Gal, D., Alves, T., & Martinez, G. (2015). *Countering online hate speech*. Unesco Publishing.

[5] Aluru, S. S., Mathew, B., Saha, P., & Mukherjee, A. (2020). Deep learning models for multilingual hate speech detection. *arXiv preprint arXiv:2004.06465*.

[6] Levy, R., Bilu, Y., Hershcovich, D., Aharoni, E., & Slonim, N. (2014, August). Context dependent claim detection. In *Proceedings of COLING 2014, the 25th International conference on computational linguistics: Technical Papers* (pp. 1489-1500).

[7] Tasnim, N., Gupta, S. S., Shihab, M. I. H., Juee, F. I., Tahsin, A., Ghum, P., ... & Sushmit, A. (2024). Mapping violence: Developing an extensive framework to build a bangla sectarian expression dataset from social media interactions. *arXiv preprint arXiv:2404.11752*.

[8] Rudra, K., Sharma, A., Ganguly, N., & Ghosh, S. (2018). Characterizing and countering communal microblogs during disaster events. IEEE Transactions on Computational Social Systems, 5(2), 403-417.

[9] Sevani, N., Soenandi, I. A., & Wijaya, J. (2021, October). Detection of hate speech by employing support vector machine with word2vec model. In 2021 7th international conference on electrical, electronics and information engineering (ICEEIE) (pp. 1-5). IEEE.

[10] Shin, Y., Qadir, H. A., Aabakken, L., Bergsland, J., & Balasingham, I. (2018). Automatic colon polyp detection using region based deep CNN and post learning approaches. IEEE access, 6, 40950-40962.

Cite this article as: Anjali Sharma, Vishal Kohli, Contextual sentiment-driven identification of multilingual communal tweets using a hybrid computational framework, International Journal of Research in Engineering and Innovation, 10 (3), (2026), 107-114. https://doi.org/10.36037/IJREI.2026.10304