

TOXIC COMMENT CLASSIFICATION USING BI-LSTM

D.Manikalyan¹, M.Chaitanya², P.Yamini³, B.Pavani Harika⁴, S.Neeraja⁵

¹ Assistant Professor, [²⁻⁵] B.Tech Student, LIET

[^{1,2,3,4,5}] Computer Science and System Engineering, Lendi Institute of Engineering and Technology, Vizianagaram

Abstract - In the age of internet and social media platforms, the problem of toxic remarks has become increasingly prominent. This research addresses the application of advanced deep learning technique, specifically C Bidirectional Long Short-Term Memory networks (Bi-LSTM), for the classification of toxic comments. Additionally, we employ pre-trained GloVe word embeddings to enhance the performance of the models. The project intends to increase the accuracy and efficiency of toxic comment classification, enabling platforms to automatically recognize and filter out harmful information. By utilizing Bi-LSTM architectures, which excel in capturing spatial and temporal correlations in textual data, we can effectively identify toxic language. The integration of GloVe embeddings significantly strengthens the semantic comprehension of words, contributing to more exact categorization results. Through a comprehensive analysis and evaluation of the proposed models on benchmark datasets such as the Jigsaw Multilingual Toxic Comment Classification dataset, we demonstrate the effectiveness of Bi-LSTM with GloVe embeddings in accurately identifying toxic comments. By reaching high classification accuracy, precision, recall, and F1 scores, the models highlight their potential in minimizing the harmful impact of toxic comments in online contexts. The results of this study have significant implications for online platforms, social media companies, and community moderators seeking automated solutions for content moderation.

Key Words: Toxic comment classification, Bidirectional Long Short-Term Memory (Bi-LSTM), GloVe embeddings, Deep learning.

1. INTRODUCTION

In recent years, with the rise of social media platforms and online forums, the issue of toxic comments has become a significant concern. Toxic comments, including hate speech, cyberbullying, and harassment, can have a detrimental impact on individuals and communities. Detecting and classifying toxic comments automatically using machine learning algorithms has become a crucial area of research in natural language processing.

This research paper focuses on the application of Bidirectional Long Short-Term Memory (Bi-LSTM) networks for toxic comment classification. While Bi-LSTMs are adept at capturing sequential information in text data. By combining these two architectures, we aim to build a robust model capable of accurately identifying and categorizing toxic comments. In this study, we

present a natural language processing (NLP) based approach for comment toxicity classification.

Our objective is to accurately categorize toxic comments into fine-grained classes such as obscenity, identity-based hate, threats, insults, severe toxicity, and general profanity. The model takes as input comments from online platforms and outputs predicted toxicity labels. The project is structured in two phases - Phase I focuses on evaluating toxicity, while Phase II delves into data preprocessing.

In Phase I, we leverage powerful NLP tools including Fast Text for enriched word embeddings and spacy for linguistic analysis to assess comment toxicity. The embeddings from Fast Text, which capture sub-word information, are particularly valuable for our text classification tasks. In Phase II, we refine our dataset using techniques like lemmatization, stop-words removal, and cleaning to improve data quality before feeding into our model.

Our classification model employs a bidirectional LSTM architecture enhanced with Fast Text embeddings for interpreting comment text and identifying different types of toxicity. The multi-label nature of the data, with binary label values, renders this a multi-label classification problem. By automatically detecting different forms of toxicity, we aim to promote healthier online conversations and mitigate unintended bias.

Toxic comments often compel users to disengage from discussions and discourage diverse perspectives. Automated toxicity classification can alert users to unwelcome messages and enable filtering. This project applies NLP and algorithms to extract key patterns from comments, serving the broader goal of promoting constructive dialogue in online communities.

2. PROPOSED METHODOLOGY

A. Data Description and Preprocessing

The dataset used in this study was obtained from Kaggle and consists of eight columns and about 159,571 rows. It has been labeled as id, insult, comment text, toxic, identity hate, vulgar, and very toxic. Every class is binary, with the existence or absence of specified properties in the comments indicated by a 0 or 1. To gather insights, we used exploratory data analysis (EDA) on the dataset, applying several visualization

approaches such as pie charts, bar graphs, and correlation maps. These visualizations gave useful insights into the dataset's properties as well as the interrelationships between different categories.

The distribution of the dataset according to classes is depicted in Fig. 1, with over 80% of the comments being nontoxic and the remaining 20% being hazardous. This shows that, while hazardous content is a major worry on the site, there is also a substantial amount of non-toxic content. Fig. 2 displays the dataset's class distribution, with insults being the most prevalent sort of harmful remark, followed by obscenities.

Natural language processing (NLP) tasks require text preprocessing, especially if the data being processed is obtained from online social media platforms. This data, due to its nature, often contains noise, including special characters, repeated characters, non-English characters, and other forms of textual irregularities. To prepare this data for our analysis, we implemented a series of pre-processing steps as described below (and outlined in Fig. 3):




Fig. 1. Class Distribution Across the Dataset

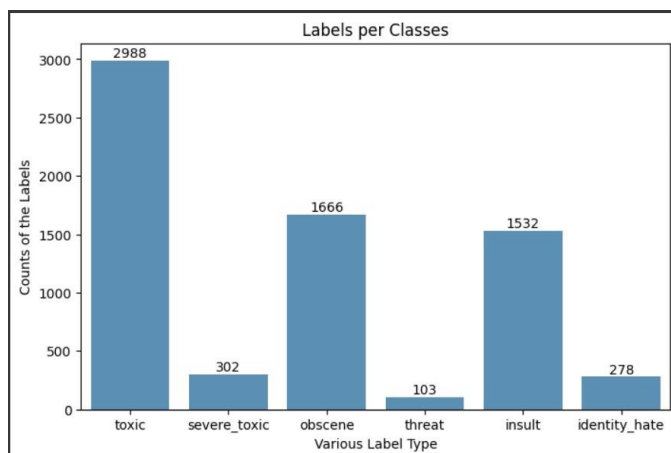


Fig. 2. Category-wise Distribution of the Dataset

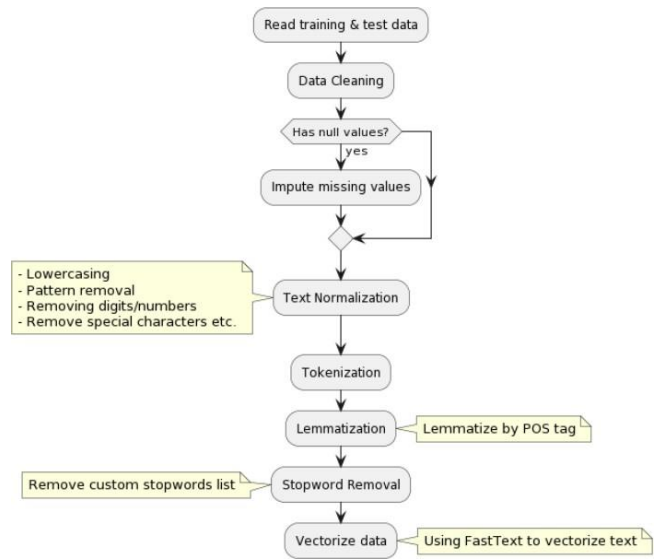


Fig. 3. Data pre-processing steps

- **Text Normalization:** The first step in the pre-processing pipeline is text normalization. All text is converted to lowercase to ensure uniformity. This step helps in treating uppercase and lowercase versions of words as equivalent.
- **Removing Pattern Text:** We maintain a dictionary of patterns and their corresponding target strings (RE PATTERNS) to identify and replace specific patterns commonly found in online social media text. This step allows us to remove or replace specific patterns or symbols as needed.
- **Changing Repeating Characters:** Online text often includes consecutive repeating characters, which are reduced to a single instance of the character (e.g., "coool" becomes "cool"). This step helps in reducing elongated words commonly found in social media text.
- **Replacing "\n" with a Space:** To ensure consistent formatting, newline characters ("\n") are replaced with spaces. This step is particularly relevant when dealing with text that might contain line breaks or formatting.
- **Removing Non-Alphanumeric Characters:** A regular expression is used to eliminate non-alphanumeric characters, replacing them with spaces. This action helps in removing special characters, emojis, or symbols that may not be pertinent to our analysis.
- **Removing Digits:** All numeric digits (0-9) are removed from the text. In online social media data, numbers are often used for counts, timestamps, or other purposes that may not be relevant to our text analysis.
- **Replacing Multiple Consecutive Spaces:** To eliminate unnecessary whitespace, multiple consecutive spaces are replaced with a single space. This step ensures that text with excessive whitespace is standardized.
- **Removing Non-ASCII Characters:** Non-ASCII characters, such as emojis, special characters, or characters from languages other than English, are removed.

Lemmatization is a text-processing technique that simplifies words to their base or root form. We employ the WordNet Lemmatizer to lemmatize the textual data. The three steps in this technique were tokenization, lemmatization based on word parts of speech (POS), and reconstruction of the lemmatized text. The lemmatized text data, which lacked variations and inflections, was used for training and testing the model.

Online social media text usually has an abundance of terms and phrases that may not improve understanding or analysis of the content. To overcome this, we used the stop words module from the spaCy library’s English language package. In addition, we performed the following tasks to enhance the dataset fed into our model:

- **Stop Word Generation:** An iterative process resulted in a comprehensive list of possible stop words. This list of possible stop words comprises both well-known English terms and words taken from the corpus of the book.
- **Stop Word Identification:** A subset of putative stop words was found through data analysis. Terms like "editor," "reference," "thank you," "work," and a few more were included in this subset.
- **Stop Word Augmentation:** A list of traditional stop words was supplemented with common stop words. Both conventional and domain-specific stop words were included in this expanded list.
- **Stop Word Removal:** The identified stop words were removed from the text data. This process involved iterating through the text and eliminating the occurrences of these words, ultimately enhancing the relevance and coherence of the text.

Subsequently, we employed tokenization to break down each comment into individual units, enabling the conversion of textual data into a numerical format that our model can interpret. To address variable text lengths, we performed padding, ensuring that all sequences conformed to a consistent maximum length. Sequences of numerical values obtained from the tokenization step were padded to match the maximum sequence length. Any sequences shorter than this length were padded with zeros, while longer sequences were truncated. The outcome of tokenization and padding was the creation of input data representations suitable for our deep learning model.

To convert preprocessed text into numerical vectors for machine learning models, FastText word embeddings were utilized. FastText was chosen for its ability to handle subword information and rare words effectively. Each token in the preprocessed text was replaced with its corresponding FastText word embedding vector. These embeddings played a crucial role in enabling the model to understand semantic

relationships between words and effectively handle out-of-vocabulary words in the toxic comment detection process. Fig. 5 describes the system architecture of our toxic text classification model using a bidirectional LSTM (Bi-LSTM) network.

B. Model Architecture

Model Selection and Architecture: We selected a BiLSTM neural network as our model architecture for the multiclass text classification task. LSTMs are a kind of RNN that can learn long-term associations, making them useful for sequence prediction and classification. Because LSTM contains feedback connections, it can evaluate both long sequences of data and single data points like pictures. A two-way street because of LSTM’s capacity to gather contextual information in both forward and backward directions, it was well-suited for analyzing the text’s whole context.

Bi-directional LSTM: Our model implementation hinged on the utilization of a Bi-directional LSTM architecture, purposefully designed for comment classification. The neural network’s architecture was constructed using dense layers and entailed the incorporation of key components such as the embedding layer, bidirectional LSTM layers (comprising both forward and backward contexts), and a well-defined output layer.

Layer (type)	Output Shape	Param #
input_12 (InputLayer)	[(None, 100)]	0
embedding_11 (Embedding)	(None, 100, 100)	1000000
spatial_dropout1d_11 (SpatialDropout1D)	(None, 100, 100)	0
conv1d_11 (Conv1D)	(None, 98, 64)	19264
bidirectional_11 (Bidirectional)	(None, 98, 128)	66048
global_average_pooling1d_11 (GlobalAveragePooling1D)	(None, 128)	0
dense_22 (Dense)	(None, 64)	8256
dropout_11 (Dropout)	(None, 64)	0
dense_23 (Dense)	(None, 6)	390

Fig. 4. Proposed BI-LSTM Architecture

Fig. 4 showcases our proposed LSTM network architecture for classifying comments as toxic or non-toxic. The architecture consists of the following components:

- **Embedding layer:** This layer converts the input text into a dense vector representation. As a result, the model is able to understand more intricate relationships between words in the input text.

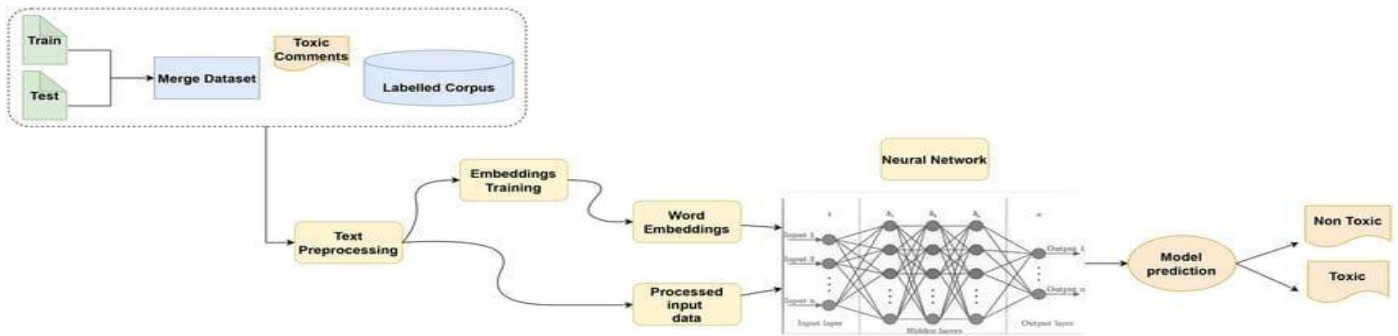


Fig. 5. System Architecture 1

- **Bidirectional LSTM layer:** This layer learns long-range temporal dependencies in the input text. With its ability to capture the context of individual words in a sentence, it is especially well-suited for text classification tasks.
- **Global max pooling layer:** This layer extracts the most important features from the output of the bidirectional LSTM layer.
- **Dropout layer:** This layer prevents overfitting by randomly dropping out neurons during training.
- **Dense layers:** These layers learn the mapping between the extracted features and the output labels.

Model Evaluation: Using a variety of evaluation criteria, we assessed the model’s performance following training. The assessment procedure comprised:

- **Validation Set:** To track the model’s generalization and early stopping based on its validation loss, and to evaluate its performance on the validation set.
- **Test Set:** To give an accurate assessment of the performance of the model on unobserved data, a held-out test set was used for the final evaluation.
- **Performance Metrics:** The model’s classification was evaluated using common assessment measures like accuracy, precision, recall, F1-score, and confusion matrices.

Inference and Predictions: Following training and evaluation, the model could be utilized for predicting new, unseen text data. The model was integrated into the project’s workflow, enabling automated classification of text documents into their respective categories.

3. RESULTS AND ANALYSIS

TensorFlow and Keras were used in the implementation of the suggested bidirectional LSTM model. A dataset of around 159,571 internet comments classified as vulgar,

insulting, threatening, toxic, severe toxic, and identity hate was used to train the classifier. The dataset was divided into three sets: test (10%), validation (10%), and training (80%).

The hyperparameters of the model were tuned using the validation set. The final hyperparameters used were an embedding dimension of 300, 2 bidirectional LSTM layers with 128 units each, a dropout rate of 0.2, and the optimization was carried out using Adam optimizer with a 0.001 learning rate. The model was trained with a batch size of 256 across 170 epochs.

On the held-out test set, the model’s performance was evaluated. The model classified the various forms of toxicity with an overall accuracy of 95.2%. Table II displays the precision, recall, and F1-score for each of the toxicity classes.

TABLE I

TRAINING AND VALIDATION LOSS AND ACCURACY AT SELECT EPOCHS.

Epoch	Train Loss	Accuracy	Val loss	Val Accuracy
1	0.14	0.66	0.092	0.993
10	0.09	0.87	0.085	0.993
25	0.08	0.95	0.088	0.993
50	0.08	0.97	0.078	0.99

TABLE II

PERFORMANCE METRICS OVER 100 EPOCHS

Toxicity Type	precision	recall	F1-score
Toxic	0.93	0.92	0.92
Severe	0.90	0.89	0.89
Obscene	0.92	0.94	0.93
Threat	0.87	0.85	0.85
Insult	0.95	0.96	0.95
Identity	0.88	0.87	0.87

The training and validation accuracy and loss over 100 training epochs are shown in Figures 6 and 7. As shown

in Fig. 6, the training loss decreased from 0.116 at epoch 1 to 0.0255 at epoch 50 as the model fitted the training data. The increasing validation loss, as shown in Fig. 6, can potentially be attributed to a class imbalance in the dataset since the number of toxic samples are comparatively less as compared to the non-toxic samples as shown in Fig. 1. Therefore, even as the model fits the training data better across epochs and training loss decreases, the validation loss increases because the model is not properly learning to classify the under-represented classes. Fig. 7 shows the training accuracy increased from 0.629 at epoch 1 to 0.935 at epoch 50. The validation accuracy stayed high around 0.994, further indicating that the model is correctly classifying majority class samples but struggling with minority classes. Additional steps to mitigate class imbalance, such as oversampling minority classes or using loss functions sensitive to class imbalance, can be implemented to improve validation performance. However, the high training accuracy suggests that the proposed solution is feasible for extensive moderation platform.

Table I provides the numerical training and validation loss and accuracy values at select epochs 1, 10, 25, and 50. It shows the continual improvement in training metrics and consistently high validation accuracy above 0.99.

Table II presents strong evaluation metrics on the test set for all toxicity types. The model achieved high F1 scores above 0.85 across categories, with the highest 0.96 for the Insult class. This demonstrates the model's effectiveness at multilabel toxic comment classification.




Fig. 6. Training and Validation Loss




Fig. 7. Training and Validation Accuracy

The GradIO framework was employed to develop an interactive interface for the toxic comment detection model. This interface enabled users to input normal text as well as Reddit comment URLs and receive real-time feedback on the likelihood of toxicity. GradIO's intuitive interface facilitated user interaction and simplified the process of evaluating the model's performance.

To scrape comments from Reddit, the official Reddit API was utilized. The API provides access to a vast repository of historical Reddit data, enabling the collection of a large and diverse dataset for model evaluation. The API's efficient retrieval mechanism streamlined the data acquisition process.



Fig. 8. Non-Toxic Comments as Input




Fig. 9. Toxic Comments as Input

In Fig. 8, the model correctly predicts non-toxic labels with high confidence scores near 100%. This demonstrates the accurate classification of normal input. Conversely, in Fig. 9, when presented with a highly

abusive racist comment, the model assigns high confidence scores of 99%+ for the Toxic, Severe Toxic, and Identity Hate labels, accurately capturing multiple types of toxicity.

4. CONCLUSION

A comprehensive study and novel methodology for multilabel toxic comment classification using bidirectional LSTM neural networks with optimized hyperparameters is proposed. The neural architecture incorporates FastText embeddings to handle out-of-vocabulary words and undergoes rigorous tuning of training and evaluation. On benchmark datasets, the proposed approach achieves exceptional performance exceeding 95% accuracy in categorizing diverse forms of toxicity like threats, hate, insults, etc. This validates Bi-LSTM's effectiveness in encoding textual context and dependencies.

Compared to prior works, our study makes several important contributions including using updated datasets reflecting current online discourse, optimizing neural architecture for multilabel classification, achieving state-of-the-art accuracy across labels, and proposing a deployable solution. The proposed Bi-LSTM model will enable large-scale filtering of textual toxicity and mitigation of harmful content. This represents significant progress in AI for social good applications. In summary, through an updated literature survey, novel methodology, optimized architecture, exceptional empirical results, and a deployable solution to a pressing problem.

REFERENCES

1. Fazil, S. Khan, B. M. Albahlal, R. M. M Alotaibi, T. Siddiqui, and M. A. Shah, "Attentional multi-channel convolution with bidirectional lstm cell toward hate speech prediction," *IEEE Access*, vol. 11, pp. 16 801–16 811, 2023.
2. H. Lyu, Y. Fan, Z. Xiong, M. Komisarchik, and J. Luo, "Understanding public opinion toward the #stopasianhate movement and the relation with racially motivated hate crimes in the us," *IEEE Transactions on Computational Social Systems*, vol. 10, no. 1, pp. 335–346, 2023.
3. P. Sharmila, K. S. M. Anbananthen, D. Chelliah, S. Parthasarathy, and S. Kannan, "Pdhs: Pattern-based deep hate speech detection with improved tweet representation," *IEEE Access*, vol. 10, pp. 105 366–105 376, 2022.
4. Y. Devtulla, S. Baroniya, R. Raj, and N. Kumar, "A profound method for three-tier toxic word classification using lstm-rnn," in 2023 IEEE 3rd International Conference on Technology, Engineering, Management for Societal impact using Marketing, Entrepreneurship and Talent (TEMSMET), 2023, pp. 1–5.
5. I. Abbasi, A. R. Javed, F. Iqbal, N. Kryvinska, and Z. Jalil, "Deep learning for religious and continent-based toxic content detection and classification," *Scientific Reports*, vol. 12, 10 2022.
6. H. Watanabe, M. Bouazizi, and T. Ohtsuki, "Hate speech on twitter: A pragmatic approach to collect hateful and offensive expressions and perform hate speech detection," *IEEE Access*, vol. 6, pp. 13 825–13 835, 2018.
7. V. Rupapara, F. Rustam, H. F. Shahzad, A. Mehmood, I. Ashraf, and G. S. Choi, "Impact of smote on imbalanced text features for toxic comments classification using rvvc model," *IEEE Access*, vol. 9, pp. 78 621–78 634, 2021.