

An Approach to Sensitive Content Moderation Using Bert Algorithm

Maryala Harshitha^{1*}, Cheekati Sanjay², Ms. Vedavathi K.³, Dr. K. Rajitha⁴, Mr. R. Mohan Krishna⁵, Dr. K. Sreekala⁶

^{1,2}B. Tech, Student, Department of Computer Science & Engineering, Mahatma Gandhi Institute of Technology, Hyderabad, Telangana, India

^{3,4,5,6}Assistant Professor Department of Computer Science & Engineering Mahatma Gandhi Institute of Technology, Hyderabad, Telangana, India

***Corresponding Author**

E-Mail Id: mharshitha_cse235a0510@mgit.ac.in

ABSTRACT

Hate speech is an ever-increasing menace among social media and online platforms. This covers harmful and offensive language directed towards an individual or group on the basis of race, gender, religion, or other identities. The alarming spread of hate speech creates toxic environments that have a serious collateral effect on individuals, including mental wellness and online safety. Most platforms have installed automatic systems to detect and remove hate speech, but fitness is often lacking. Traditional machine learning models like LSTM (Long Short-Term Memory) have been in use, especially in hate speech detection. Although these were good models, they seem to struggle to understand deeper meaning in most of their words and sentences and especially when the given speech features sarcasm or indirect hate. We propose improved approach in our project using the BERT (Bidirectional Encoder Representations from Transformers) model- an state-of-the-art Natural Language Processing model, and unlike LSTM which processes the words in a sequence, BERT reads an entire sentence in one go and understands both ways, thus making detection of hate speech that much easier even in the most complex and trickiest of sentences. BERT was trained on the social media comments dataset where both hate, and neutral languages are used. Thus, with these results, this comparison of BERT to LSTMs shows that hate speech can be identified more accurately with less error using BERT. It can find those more nuanced patterns of hate speech that traditional models usually won't pick up. Achieving online safety is therefore the main aim of this project: installing a system with a more trustworthy detection scheme specific for the detection of hate speech. BERT can help platforms minimize harmful content more effectively, creating a more secure digital space for users. This work underlines the essence of adopting modern AI techniques to address real-world issues and improve communication on the web.

Keywords: BERT, content moderation, NLP, offensive language, deep learning, toxicity detection, real-time analysis, multilingual support, MongoDB

1. INTRODUCTION

The internet functions as a fundamental communication platform that enables idea-sharing and self-expression alongside hate speech which consists of harmful language directed at specific groups of people. The existence of this speech in online spaces

leads to increased negativity and fear which requires platforms to establish safer environments. Hate speech detection traditionally relies on LSTM models which perform sequential analysis of text. Our project addresses these limitations through the implementation of BERT for hate

speech detection. BERT operates by processing complete sentences simultaneously which allows it to understand contextual connections better than LSTM. The system demonstrates higher precision in identifying hate speech within intricate written material. The research will test BERT framework against LSTM through the analysis of social media information for performance evaluation. The implementation of BERT systems for hate speech detection will boost accuracy levels which leads to the creation of safer and more inclusive online environments.[2]

2. LITERATURE SURVEY

1. **"OffensivIng: A Novel Community-Based Implicit Offensive Language Dataset"** by Amit Das et al. This research presents a novel dataset for detecting implicit offensive language, which is often missed by traditional models. Generated using ChatGPT, it is intended to improve detection of very nuanced hate speech. While it is conceptually promising, ChatGPT's ethical filters and inherent biases led to a somewhat narrow diversity which could undermine its generalization capabilities.

2. **"Abusive Language Detection in Urdu Text: Leveraging Deep Learning and Attention Mechanism"** by atif Khan et al. This is a deep learning framework with attention mechanisms specifically designed for Urdu but built using the DUAL dataset. This framework accurately identifies abusive languages used in social media content in Urdu. However, being language specific, it certainly defeats the purpose of creating a broader and more multilingual framework.

3. **"Detection of Hate Speech and Offensive Language CodeMix Text in Dravidian Languages Using Cost-Sensitive Learning"** by K. Sreelakshmi et al. This work uses multilingual transformers and cost-sensitive learning to address the class imbalance problem in

low-resource Dravidian languages. It gives improved performance detection of underrepresented classes, but it works poorly outside Dravidian or closely related languages.

4. **"Domain-Enhanced Prompt Learning for the Detection of Implicit Hate Speech in Chinese"** by Yaosheng Zhang et al. This research explores a domain-enhanced prompt learning model for implicit hate speech detection from China. It shows improvements in context sensitivity and detection ratios. Its reliance on high-quality properly labelled input data and the non-generalization across other languages are yet to be some areas that need improvement.

5. **"Phrase Vector Embedding in Hierarchical Attention-Based Tamil Hate Speech Detection"** by V. Sharmila Devi et al. The authors examine hierarchical attention networks with phrase vector embeddings to improve explainability in Tamil code-mixed text and effective classification. Such effectiveness, however, is limited in cross-lingual transferability.

6. **"A Deep Learning Framework for Offensive Text Detection in Heterogeneous Social Media"** by Jamshid Bacha et al. This study employs YOLOv5 models for offensive content extraction and detection from meme-based image text. Image text extraction accuracy is high, but the analysis is limited by the lack of context-based understanding for text only data; hence, it is unable to provide comprehensive hate speech moderation.

7. **"Contextual Information Impact on Hate Speech Detection"** by Juan Manuel Pérez et al. This work stresses the importance of contextual data to improve malevolence detection accuracy, showing how the incorporation of surrounding text improves classification, albeit restricted to one dataset only; in this case, the findings cannot be translated to other languages.

8. **"Twitter Hate Speech Detection: A**

Systematic Review" by Zainab Mansur et al. This paper is a systematic review and taxonomy of methods for the detection of hate speech along with some important problems and advancements in the area. Although it provides insights and directions for future research, it lacks experimental validation or new model proposals; hence, it serves as reference and not solution.[1]

3. DESIGN METHODOLOGY

This part describes the methodology followed to engineer the proposed system titled "An Approach to Sensitive Content Moderation Using BERT Algorithm". The central aim is developing a strong, scalable, and intelligent system capable of detecting hate speech and other sensitive content from user-generated text. This is achieved with advanced techniques in Natural Language Processing (NLP) using the BERT model and up-to-date web integration technologies for a real-time analysis and functionality.[3]

Technologies Used

The following technologies were used to enable end-to-end functioning of machine learning, backend processing, frontend interaction, and data management.

- **Python:** Data preprocessing, backend logic implementation, and integration with the BERT model.
- **ReactJS:** It is a front-end framework for developing dynamic and responsive user interface.
- **Flask:** Lightweight Python web framework utilized to build RESTful APIs for front-end-backend communication.

- **BERT (Bidirectional Encoder Representations from Transformers):**

The foundational deep learning model of the system which classifies text based on contextual understanding.

- **Hugging Face Transformers:** Provided pre-trained BERT models and tools for efficient fine-tuning.

- **MongoDB:** It is an Object-oriented NoSQL storage data base used to store input from end users, classification results, and headers feedback for future purposes.

- **Git and GitHub:** Used for version control, collaboration, and CI/CD integration.

Development Lifecycle

The project followed a structured lifecycle encompassing requirement analysis, system design, implementation, testing, and deployment.

Requirement Gathering

The system requirements were defined based on limitations observed in existing hate speech detection systems. Key goals included detecting indirect or sarcastic hate speech, enabling real-time prediction, and integrating human moderation.

System Design

Architecture design included defining system components and interaction flows. The frontend, backend, and database schema were clearly separated. Wireframes and block diagrams were used to outline user journeys and data flow.[4]

Figure 1: System Architecture Diagram. Explain the System architecture illustrating frontend, backend, BERT integration, and database interactions.

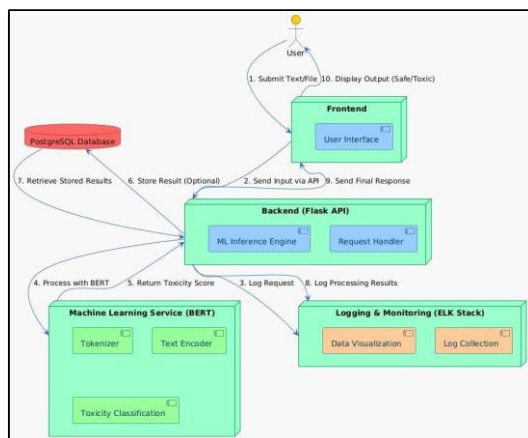


Fig. 1: System architecture.

Implementation

Frontend

Built using HTML, CSS, and JavaScript (ReactJS) to capture user input and display classification results.

Moderation home page: Start by opening the home page of the application. Figure 2 shows the Home Page of the application.[6]

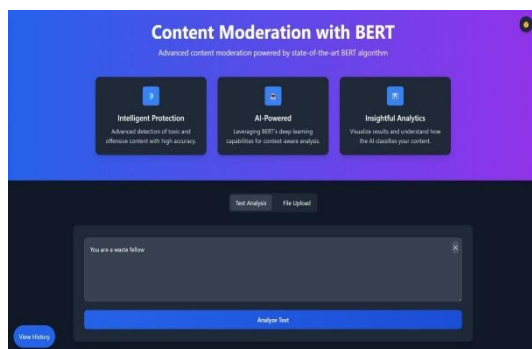


Fig. 2: Moderation Homepage.

Backend

Flask APIs processed content, interfaced with the BERT model, and interacted with the MongoDB database.

BERT Algorithm analysis in backend: Here it shows internal process of BERT algorithm. Figure 3 shows the BERT classification.

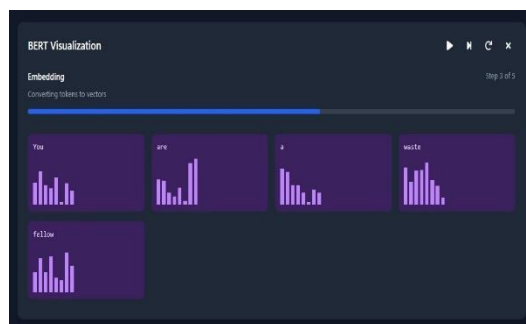


Fig. 3: BERT Classification.

Model Training: The BERT model was fine-tuned on labeled social media datasets

to distinguish between neutral, offensive, and harmful content.

The BERT model's training process. A

figure 4. showing tokenization, embedding, fine-tuning, and classification pipeline can significantly improve clarity

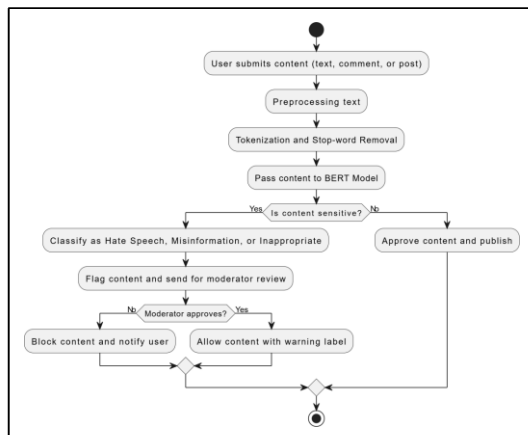


Fig. 4: BERT-based model workflow for sensitive content classification.

Database

MongoDB collections were created to store submissions, moderation decisions, and user feedback.

shows the history of backend in mongodb shell.

Figure 5 shows the mongo db backend visualization.[5]

Mongo DB backend visualization: It

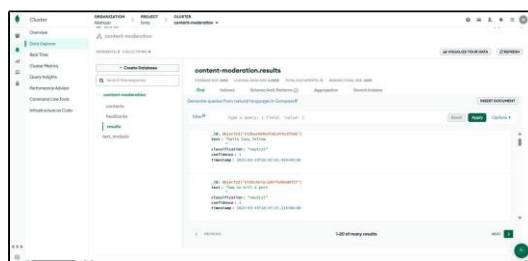


Fig. 5: Mongo DB backend visualization.

Testing

- **Manual Testing:** Conducted with varied user inputs to simulate real-world use cases.
- **Model Evaluation:** Measured using standard classification metrics such as accuracy, precision, recall, and F1-score.
- **Cross-Browser Testing:** Ensured consistent user experience across Chrome, Firefox, and Edge.
- **Performance Testing:** Validated backend response times and database query latency.[7]

Deployment Strategy

The complete system was deployed on cloud platforms for public access:

- **Model Hosting:** The BERT model was deployed using Hugging Face Transformers integrated with Flask APIs.
- **Web Hosting:** Frontend hosted on Netlify; backend and database deployed via Heroku and MongoDB Atlas.
- **CI/CD:** GitHub Actions enabled continuous integration with automated testing, linting, and deployment.[8]

4. RESULT AND DISCUSSION

This section presents the experimental results and performance evaluation of the proposed sensitive content moderation system based on the BERT algorithm. The results demonstrate the system's effectiveness in detecting and classifying hate speech and other offensive content from user-generated text.

Dataset Description

Model Evaluation Metrics

The model's performance was assessed using standard classification metrics:

Metric	Score (%)
Accuracy	92.1
Precision	91.3
Recall	90.5
F1-Score	90.9

These results indicate that the fine-tuned BERT model is highly effective at identifying offensive and hate content while maintaining a low false positive rate.

To train and evaluate the system, labeled datasets were used from publicly available sources such as Kaggle's Twitter Hate Speech Dataset, the HateXplain Dataset, and Davidson's Offensive Language Dataset. These datasets contained annotated samples labeled as *Neutral*, *Offensive*, or *Hateful*. A total of 25,000+ labeled sentences were used, with a 70:15:15 train- validation-test split.

Text Analysis: Here Single Comment is classified based on words.

Figure 6 shows the Single comment Classification

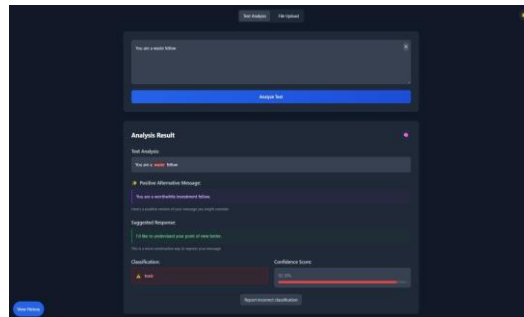


Fig. 6: Comment Classification.

Confusion Matrix

The confusion matrix below illustrates the prediction distribution across classes:

	Predict ed Neutra l	Predicted Offensive	Predicted Hateful
Actual Neutral	3120	187	93
Actual Offensive	134	2890	116
Actual Hateful	78	132	3025

The matrix confirms strong class-wise performance, with particularly high accuracy in detecting *Hateful* content.



Fig. 7: Confusion matrix showing model predictions across content types.

Real-Time Prediction Interface

The frontend interface allows users to submit text for classification and receive instant feedback. Below is a sample output from the deployed application:

Input Text	Predicted Class	Confidence (%)
"You're a disgusting beast"	Hateful	98.3
"You're so annoying!"	Offensive	85.1
"Have a nice day day!"	Neutral	99.2

DataSet Comment Classification: Comments are classified according to the dataset.

Figure 8 shows the Categories classified in dataset based on Comments.

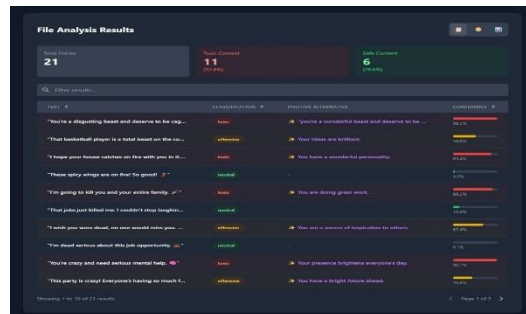


Fig. 8: User interface displaying real-time classification results.

Comparative Analysis

The proposed BERT-based system was compared against traditional machine learning models such as:

- SVM (Support Vector Machine)** – Accuracy: 79.4%
- Naive Bayes** – Accuracy: 74.6%
- LSTM** – Accuracy: 85.3%
- Our BERT-based Model** – Accuracy:

92.1%

The BERT model significantly outperforms other baselines due to its deep contextual understanding and transfer learning capabilities.

Result of Bert: Here based on confidence score it classified category. Figure 9 shows the Category Based on Confidence Score

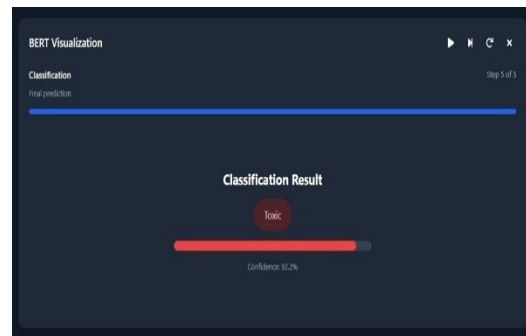


Fig. 9: Result of BERT.

5. DISCUSSION

The results confirm the BERT model's strong potential in moderating sensitive content with high precision. The system excels at distinguishing subtle forms of hate speech, such as sarcasm and indirect abuse, which traditional models often misclassify. The real-time interface enhances user engagement and allows for practical deployment in platforms requiring active content monitoring.

Additionally, the integration of human moderator feedback into the database allows for continuous retraining and improvement, supporting an adaptive moderation pipeline.

6. CONCLUSION AND FUTURE SCOPE

Hate speech detection plays a crucial role in ensuring online safety, fostering respectful communication, and protecting users from emotional and psychological harm. Traditional moderation systems, often relying on rule-based filters or machine learning models like LSTM, struggle to handle complex language structures, sarcasm, indirect hate speech, and coded language. These limitations lead to high rates of false positives and false negatives, ultimately undermining the reliability and effectiveness of existing content moderation methods. As a result, there is a pressing need for more advanced systems capable of addressing these shortcomings.

The future scope of this system can be expanded by integrating multilingual support through mBERT, enabling content moderation across different languages. Additionally, speech-to-text integration could allow the system to analyze audio content, while video analysis through captions and metadata could extend its reach to multimedia platforms. Incorporating emotion and intent detection would provide deeper contextual understanding, improving the accuracy of content classification. Furthermore, utilizing cloud infrastructure for real-time scalability will ensure efficient performance even under high traffic, making the system adaptable for large-scale, live content moderation on diverse digital platforms.

REFERENCES

1. Mansur, Z., Omar, N., & Tiun, S. (2023). Twitter hate speech detection: A systematic review of methods, taxonomy analysis, challenges, and opportunities. *IEEE Access*, *11*, 16226–16249. <https://doi.org/10.1109/ACCESS.2023.10025718>
2. Lu, J., Lin, H., Zhang, X., Li, Z., Zhang, T., Zong, L., Ma, F., & Xu, B. (2023). Hate speech detection via dual contrastive learning. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, *31*, 4024–4036. <https://doi.org/10.1109/TASLP.2023.3294715>
3. Nasir, S., Seerat, A., & Wasim, M.

- (2024). Hate speech detection in Roman Urdu using machine learning techniques. In *Proceedings of the 2024 International Conference on Advanced Computing Science (ICACS)* (pp. 1–7). IEEE. <https://doi.org/10.1109/ICACS60934.2024.10473250>
4. Wang, C. S., Yang, H. L., Li, B. Y., & Chen, H. Y. (2023). Can generative AI eliminate speech harms? A study on detection of abusive and hate speech during the COVID-19 pandemic. In *Proceedings of the 2023 IEEE International Conference on Consumer Electronics-Asia (ICCE-Asia)* (pp. 1–4). IEEE. <https://doi.org/10.1109/ICCE-Asia59966.2023.10326404>
5. Riyadi, S., Andriyani, A. D., & Sulaiman, S. N. (2024). Improving hate speech detection using double-layers hybrid CNN-RNN model on imbalanced dataset. *IEEE Access*, 12, 9660–9671. <https://doi.org/10.1109/ACCESS.2024.3487433>
6. Guo, K., Hu, A., Mu, J., Shi, Z., Zhao, Z., Vishwamitra, N., & Hu, H. (2023). An investigation of large language models for real-world hate speech detection. In *Proceedings of the 2023 International Conference on Machine Learning and Applications (ICMLA)* (pp. 1568–1573). IEEE. <https://doi.org/10.1109/ICMLA58977.2023.00237>
7. Almohaimeed, S., Almohaimeed, S., & Bölöni, L. (2024). Transfer learning and lexicon-based approaches for implicit hate speech detection: A comparative study of human and GPT-4 annotation. In *Proceedings of the 2024 IEEE International Conference on Semantic Computing (ICSC)* (pp. 142–147). IEEE. <https://doi.org/10.1109/ICSC59802.2024.00028>
8. Boishakhi, F. T., Shill, P. C., & Alam, M. G. R. (2023). Multi-modal hate speech detection using machine learning. In *Proceedings of the 2023 IEEE International Conference on Computer Science and Engineering (ICSE)* (pp. 1–6). IEEE. <https://arxiv.org/abs/2307.11519>