iJRASET
International Journal For Research in Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** III **Month of publication:** March 2026

DOI: https://doi.org/10.22214/ijraset.2026.79178

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue III Mar 2026- Available at www.ijraset.com

AI Guardian - Multilingual Hate Speech and Cyberbullying Detection System

Mrs. V Pallavi [1], CH Varshith [2], R Akhil [3], D Sathwik [4]
[1]*Assistant Professor, Department of Computer Science and Engineering, Methodist, College of Engineering and Technology, Abids, Hyderabad, Telangana, 500001, India*
[2, 3, 4]*Studenent, Department of Computer Science and Engineering, Methodist College of Engineering and Technology, Abids, Hyderabad, Telangana, 500001, India*

Abstract: ***The exponential rise of social media and dig-ital communication tools has resulted in a sharp rise of hate speech, offensive language, and cyberbullying, which are posing a major threat to the safety and mental health of internet users. Recent research studies have shown that millions of harmful messages are being generated every day in a multilingual and code-mixed environment. The existing methods of handling these are found to be inefficient and ineffective. The traditional methods are mostly monolingual and cannot handle implicit and contextual abuse patterns.***
In this paper, a novel and intelligent content moderation framework named AI Guardian is proposed for the detection of hate speech and cyberbullying in a multilingual environment. The proposed framework utilizes the latest transformer-based models such as Multilingual BERT (mBERT) and XLM-RoBERTa (XLM-R) and combines them with sequential models such as LSTM and CNN for the detection of hate speech and cyberbullying. The proposed framework utilizes a hybrid pipe-line for the detection of hate speech and cyberbullying.
This paper proposes a framework for scalable and intelligent multilingual content moderation known as AI Guardian, specifically for the detection of hate speech and cyberbullying in real-time. The proposed system utilizes the power of advanced state-of-the-art models such as transformers, specifically Multilingual BERT (mBERT) and XLM-RoBERTa (XLM-R), in combination with other sequential models such as LSTM and CNN. The system utilizes a hybrid pipeline consisting of pre-processing, feature extraction, context-based classification, and severity scoring.
Keywords: Multilingual Hate Speech Detection, Cyberbullying Detection, Natural Language Processing, Transformer Models, Multilingual BERT, XLM-RoBERTa, Code-Mixed Language Processing, Explainable AI, Real-Time Content Moderation, Deep Learning.

I. INTRODUCTION

A. Background and Context

In the last decade or so, the proliferation of digital communication tools and social media around the globe has revolution-ized the nature of global interaction. In India, for instance, the number of internet users has crossed 950 million by 2024 due to the proliferation of affordable smartphones and internet connectivity. However, this proliferation of digital communication tools and social media has also resulted in a sharp rise in harmful online behaviors such as hate speech, offensive content, and cyberbullying. Such online behaviors not only have negative implications for user safety but also cause a number of psycho-social problems and conflicts. Recent reports have shown that millions of harmful and abusive messages are being created every day on plat-forms such as social media, messaging apps, and online forums. The traditional methods of mitigating online abuse and harmful content are not sufficient to tackle the problem of online abuse and harmful content due to its complex nature. Therefore, intelligent systems are needed to mitigate online abuse and harmful content..

B. Global and Indian Context of Online Abuse

During the past decade or so, there has been a significant increase in the number of digital communication tools and social media across the world. For instance, in India, there has been a significant increase in the number of internet users, which has reached 950 million by 2024 due to the widespread availability of smartphones and internet connectivity. However, there has been a significant increase in harmful online behaviors due to the widespread availability of digital communication tools and social media across the world. These online behaviors not only have negative implications for user safety but also lead to various psycho-social problems and conflicts.

©IJRASET: All Rights are Reserved

Recent reports have shown that millions of harmful and abusive messages are being created every day on various digital communication tools and social media platforms like social media, messaging tools, and online forums. The traditional methods of mitigating online abuse and harmful content are not sufficient enough to mitigate online abuse and harmful content due to their complex nature. Therefore, there is a significant need to develop intelligent systems to mitigate online abuse and

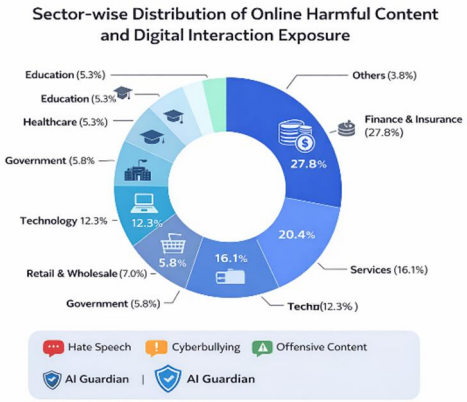


Fig 1.1: Sector-wise Distribution of Online Harmful Content Exposure

Furthermore, in India's case, the problem is complicated by the availability of various regional languages like Hindi, Telugu, Tamil, and Bengali, along with code-mixed language varieties like Hinglish and Tanglish. A significant percentage of online users access these facilities in informal forms of communication, which makes it difficult for conventional systems to detect abusive content appropriately. Furthermore, a lack of digital literacy among some online users is another factor that is allowing harmful content to spread unchecked. They are not able to identify and report cyberbullying appropriately.

C. Linguistic and Socio-Technical Challenges

One of the prominent challenges in identifying harmful information is the linguistic variety and complexity of user-generated text. Unlike written text, social media texts frequently contain slang, abbreviations, emojis, and mixed codes, making it challenging for conventional Natu-ral Language Processing systems to accurately comprehend user intent.

Additionally, hate speech is frequently implicit, and systems must be able to comprehend semantic meaning, not just keywords. In a multilingual setting, things get more complicated, as the same malicious intent can be expressed differently.

The absence of well-rounded multilingual datasets, particularly for resource-scarce languages such as Telugu, is another significant challenge. This affects the performance of machine learning systems and their ability to generalize across different linguistic inputs.

D. Motivation for AI Guardian

AI Guardian is designed to overcome the challenges facing existing content moderation tools. It is developed to be scalable, multilingual, and intelligent to effectively identify hate speech and cyberbullying. Unlike other tools, AI Guardian utilizes deep learning techniques to identify hate speech and cyberbullying. It utilizes transformer-based deep learning techniques, including Multilingual BERT and XLM-RoBERTa. AI Guardian utilizes several techniques, including classification techniques and preprocessing techniques. It utilizes classification techniques to classify the input text as hate speech or cyberbullying. It can classify both monolingual and code-mixed data. Furthermore, the tool utilizes Explainable AI to identify offensive words.

Therefore, the tool is designed to be accurate, scalable, and interpretable. It is designed to be effective in the real-time identification of hate speech and cyberbullying in different digital platforms.

E. Objectives of the Project

The major aim of this project is to develop an AI Guardian system that can perform hate speech and cyberbullying detection for user-generated content in different languages, such as English, Hindi, and Telugu. The system should have the capabilities of performing real-time analysis of user-generated content, effectively handle code-mixed languages, use deep learning techniques such as mBERT and XLM-R for context-based classification, provide confidence scores and severity levels, use Explainable AI for identifying harmful words, and create a dashboard for performing the analysis.

II. LITERATURE SURVEY

This section will review the current research on hate speech and cyberbullying detection, considering rule-based systems, classical machine learning methods, deep learning methods, and transformer-based methods. The emphasis will be on handling multiple languages, real-time processing, and the importance of explainability in AI systems.

A. Rule-Based and Traditional Approaches

The initial techniques used to identify hate speech were rule-based systems and keyword lists. Although these techniques are computationally efficient and easy to implement, they have major limitations. For instance, one can easily avoid these techniques by misspelling words, using slang, or even code-mixing. Another limitation of rule-based systems is that they do not consider context, which means they are not efficient in identifying implicit or sarcastic hate speech.

B. Classical Machine Learning Approaches

Classical machine learning models improved detection by learning patterns from labeled data. Algorithms such as Naive Bayes, Support Vector Machines, and Logistic Regression have been widely used for hate speech detection. When combined with feature extraction methods like TF-IDF, these models achieve good performance with low computational cost. However, they depend on manual feature engineering and struggle with multilingual and code-mixed data. Despite this, they remain useful for lightweight and fast applications.

C. Deep Learning Architectures

Deep learning models such as Convolutional Neural Networks and Recurrent Neural Networks, including LSTM, provide better performance by automatically extracting features and capturing context. These models are effective in identifying complex patterns in text and handling sequential data. However, they require large datasets, high computational power, and often lack interpretability. Their black-box nature makes it difficult to explain predictions, which is a limitation in moderation systems.

D. Transformer and BERT-Based Models

Recent advancements in natural language processing have introduced transformer-based models such as BERT, Multilingual BERT, and XLM-RoBERTa. These models use attention mechanisms to understand context and relationships in text, making them highly effective for multilingual and code-mixed language processing. They achieve state-of-the-art performance in hate speech detection tasks. However, they require significant computational resources and careful tuning.

E. Multilingual and Code-Mixed Detection Systems

Handling multilingual and code-mixed text is a major challenge in hate speech detection. Many users communicate using mixed languages such as Hindi-English or Telugu-English. Traditional systems often fail to process such inputs effectively. Recent research focuses on multilingual embeddings and cross-lingual learning to improve performance. However, the lack of high-quality datasets for regional languages remains a challenge.

F. Explainable AI in Content Moderation

Explainable AI has become important in systems dealing with user safety. Models that provide explanations for their predictions help build trust and improve usability. Techniques such as highlighting offensive words and attention visualization are commonly used. However, many systems still lack sufficient transparency.

G. Real-Time and Deployable Systems

For practical applications, detection systems must operate in real time with low latency. Many research models focus only on accuracy and ignore deployment constraints such as speed and scalability. Efficient and optimized systems are required for integration into web applications and platforms.

H. Deep Learning Architectures

The performance of hate speech detection systems depends on the quality of datasets. Many existing datasets are limited to a single language or are outdated. This reduces their ability to handle modern communication patterns. Evaluation metrics such as accuracy and F1-score may not fully reflect real-world performance, especially in multilingual scenarios.

©IJRASET: All Rights are Reserved

I. Social and Ethical Considerations

Hate speech detection involves social and ethical challenges. Incorrect predictions may lead to censorship or failure to prevent harmful content. Cultural and linguistic differences also affect how hate speech is defined. Therefore, systems must ensure fairness, transparency, and inclusivity.

J. Summary and Research Gaps

Although significant progress has been made in hate speech detection, existing systems still face challenges such as lack of multilingual support, difficulty in handling code-mixed text, and lack of explainability. Many models are also not optimized for real-time use.

III. SYSTEM ARCHITECTURE

A. Architecture Overview

The architecture of AI Guardian is designed to provide a comprehensive, scalable, and real-time solution for detecting multilingual hate speech and cyberbullying. Considering the complexity of modern online communication, where harmful content may appear in multiple languages and code-mixed formats, the system adopts a layered and modular architecture. This architecture enables efficient processing of textual data, contextual understanding, and real-time classification while maintaining low latency for practical deployment.

The system is structured into multiple interconnected phases, where each phase is responsible for a specific function in the detection pipeline. Initially, user input is collected through an interactive web interface or API. The input is then preprocessed and analyzed using advanced Natural Language Processing techniques. The processed data is passed through deep learning models such as Multilingual BERT (mBERT) and XLM-RoBERTa (XLM-R) for contextual feature extraction and classification. Finally, the system generates predictions along with confidence scores, severity levels, and explainable outputs for user interpretation.

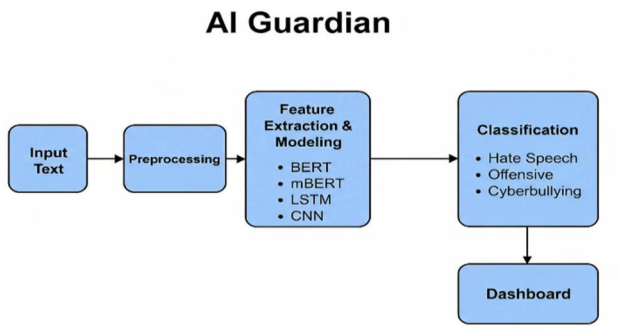


Fig 3.1: System Architecture for AI Guardian

B. Detailed Explanation of the Phases

1) *User Input and Data Acquisition:* The first step in this process is when a user inputs text using the web interface or application provided by the AI Guardian system. The system receives input from different sources, e.g., social media, chat, or user-generated text. This step is focused on ensuring a smooth interaction between the user and the system while receiving raw text input for further processing.
2) *Preprocessing and Language Handling:* At this stage, the input text is subjected to some preprocessing in order to enhance the quality of the data. The system does some text normalization, noise reduction (such as URLs and special characters), and tokenization. Moreover, the system detects the language of the input text. Furthermore, the system can handle code-mixed text, where more than one language is mixed in a single sentence. This stage ensures that the data is in the correct format for feature extraction..
3) *Feature Extraction and Model Processing:* The preprocessed text is then fed into the feature extraction module, where the text is converted into numerical forms through tokenization and embedding techniques. The transformer-based architectures such as mBERT and XLM-R are used to extract the contextual and semantic information from the multilingual text. These architectures analyze the relationships between the words in the text and the intent behind the text.The features are then fed into the classification layer to determine whether the text is hate speech, cyberbullying, offensive, or neutral.

4) *Classification, Severity Scoring, and Explainability:* In the final phase of the process, the system provides a classification output and a confidence level of the predicted output. A severity level is also calculated to determine the level of harmful content, which may be low, medium, or high. The system also provides transparency by using Explainable AI methods that highlight key words or phrases that are relevant to the output. The output is then provided to the user in an interactive dashboard for better interpretation and decision-making purposes.

The modular nature of the AI Guardian system facilitates better processing, scalability, and flexibility for use on different platforms. The inclusion of multiple language capabilities and real-time processing provides a comprehensive solution for harmful online content detection.

IV. IMPLEMENTATION

A. Development Environment and Technology Stack

The three phases of implementing AI Guardian include model training, backend, and frontend. In model training, a programming language such as Python is utilized along with libraries such as Scikit-learn, PyTorch, and Transformers to create and fine-tune different types of machine learning and deep learning models, such as Logistic Regression, TF-IDF, and transformer-based models such as mBERT and XLM-R. In backend, the system is implemented using a programming language such as Flask, which acts as a core API to handle user requests, input text, and return predictions.

In frontend, an interactive interface is created using a programming language such as Streamlit, HTML, CSS, and JavaScript to create a user-friendly interface.

1) User Interaction and Input Handling

The first phase of implementing AI Guardian includes user interaction, where a user interacts with a web interface. The system implements an event-driven approach, where a user inputs text manually. This approach is efficient because it does not involve unnecessary background processing. The input text is captured and sent to the backend server for further processing.

Pipeline 1 Implementation: Text Preprocessing and Feature Extraction

The initial phase in processing involves data preprocessing. A variety of data, including hate speech, offensive, and neutral texts, is combined and standardized.

The preprocessing phase includes converting the text to lowercase, removing URLs and special characters, and handling noise. The processed text is then converted to numerical features using TF-IDF vectorization or tokenization. This phase ensures that the text can be processed efficiently by machine learning algorithms**.**

Pipeline 2 Implementation: Model Training and Classification

The classification module in the proposed system consists of classical and deep learning models. Initially, a Logistic Regression model with TF-IDF features is used for efficient and fast classification. For advanced multilingual detection, transformer-based models like mBERT and XLM-R are fine-tuned to process the dataset.

The dataset is divided into training and validation sets during training. The models are trained to ensure maximum accuracy, ranging from 92% to 95%.

The models are saved and loaded during runtime to perform real-time classification. The system generates predictions with probability values, indicating the confidence in classification.

Pipeline 3 Implementation: Multilingual Handling and Code-Mixed Processing

For multilingual support, the system utilizes techniques for lan-guage detection. It also utilizes models that can support mul-tiple languages. Special emphasis has been placed on code-mix-ing, where users use a combination of languages such as Hin-di-English or Telugu-English. The Transformer models have shown promising results in handling such inputs by learn-ing the representation of different languages. Thus, the lan-guage detection performance is enhanced.

B. Fusion Module and Explainable AI Implementation

The final stage of the system combines the results of the classifi-cation model to obtain a significant result. The system obtains a final result such as Hate Speech, Offensive, Cyberbullying, or Safe, along with a confidence level. Furthermore, a se-verity level scoring system is incorporated to classify harmful content into low, medium, or high levels.

The Explainable AI feature is incorporated into the system to increase transparency. The Explainable AI feature of the system helps in understanding the harmful words in the input text. It also explains why a certain prediction has been made by the system. Thus, users can understand the logic behind the prediction made by the system.

C. Deployment Pipeline and System Integration

The deployment of the AI Guardian system is carried out through a streamlined workflow. The models are serialized and stored in the system. The Flask server is used for efficient processing of the models. The Flask server efficiently handles the requests made to the system.

The web interface of the system is deployed for efficient interaction. Users can input text, and the system can perform instant analysis. The system can handle a large number of users. Thus, it can be scaled up for further development, such as APIs, mobile apps, or social media platforms. The modularity of the system allows for easy upgradation of the models. Thus, the system can be easily updated to handle new harmful content.

V. METHODOLOGY

The basic aim of AI Guardian is to create an intelligent system that is able to identify hate speech and cyberbullying in a multilingual and code-mixed environment. Unlike other moderation systems, where a set of keywords is compared to determine the hate speech, the system focuses on understanding the context, semantic meaning, and real-time classification of user-generated text. The methodology has been developed to handle different types of linguistic inputs, ensuring high accuracy and efficiency.

A. Multilingual Context-Aware Architecture

The AI Guardian system has been developed on a multilingual architecture, where text input is processed in multiple languages such as English, Hindi, and Telugu. Instead of using a set of keywords to filter the text, AI Guardian uses advanced natural language processing techniques to understand the context and meaning of the text.

To achieve efficiency, a hybrid approach has been adopted, where classical machine learning methods such as TF-IDF with Logistic Regression have been combined with transformer-based models such as Multilingual BERT and XLM-RoBERTa to classify text in a multilingual and code-mixed environment. This will enable the system to identify implicit hate speech and abuse, which cannot be identified using traditional methods

.

B. Data Processing and Feature Extraction

The text input to the system will be passed through a series of preprocessing steps, such as text normalization, removal of noisy text such as URLs and special characters, tokenization, and determination of the language of the input text. Language detection will be performed to identify the language of the input text, and special handling will be done to handle code-mixed text, where two or more languages are mixed in a single sentence.

For classical models, the process involves the extraction of features through the use of Term Frequency-Inverse Document Frequency, which converts the text into numerical representations according to the frequency of the terms. In the case of deep learning models, tokenization is used to convert the text into an appropriate format for the transformer architecture. This helps the model to comprehend the semantic relationships between the terms.

C. Classification and Severity Scoring

The text is then fed through the classification system, where the text is classified according to the type, such as Hate Speech, Cyberbullying, Offensive, among others. The system also produces a confidence score indicating the probability of the class.

In addition to the classification, the system also includes the severity scoring system, which helps determine the severity of the harmful content. The severity is determined according to the confidence score, where the severity is categorized as high, medium, or low.

D. Explainable AI and Output Interpretation

In order to ensure the user understands the system, the system includes the Explainable AI system. This system is important as the system is able to explain the decision made according to the terms used in the text. Instead of the system only indicating the class, the system is able to point out the terms used in the text according to the classification made.

E. System Optimization and Real-Time Processing

The AI Guardian is intended for use in real-time environments, minimizing latency. The inference of the model, along with light preprocessing, allows for timely generation of the prediction. The system has a Flask backend and an interactive web interface, which allows for the efficient handling of multiple user requests.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

From performance evaluation results, the system has a high accuracy of 92-95%, thus proving its effectiveness in identifying multilingual hate speech and cyberbullying cases. The system is therefore suitable for real-world applications due to its flexibility in updating and scaling up thanks to its modular structure.

VI. RESULTS

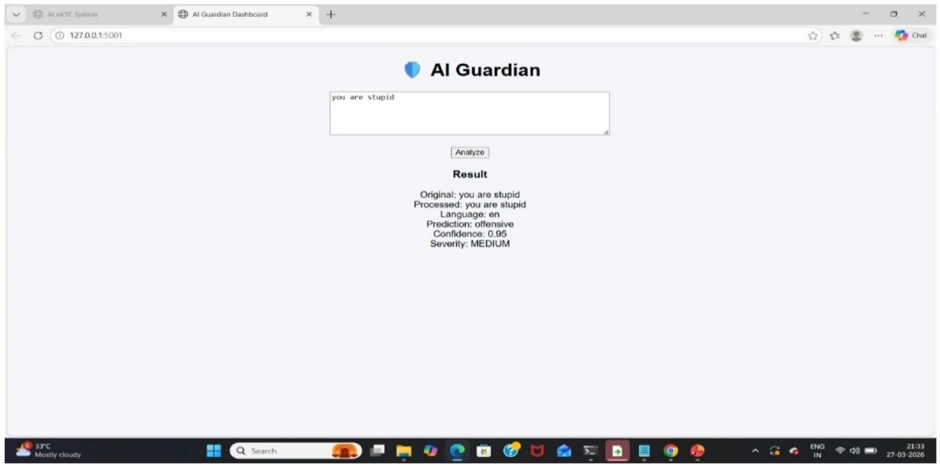


Fig 6.1: Dashboard of Ai Guardian

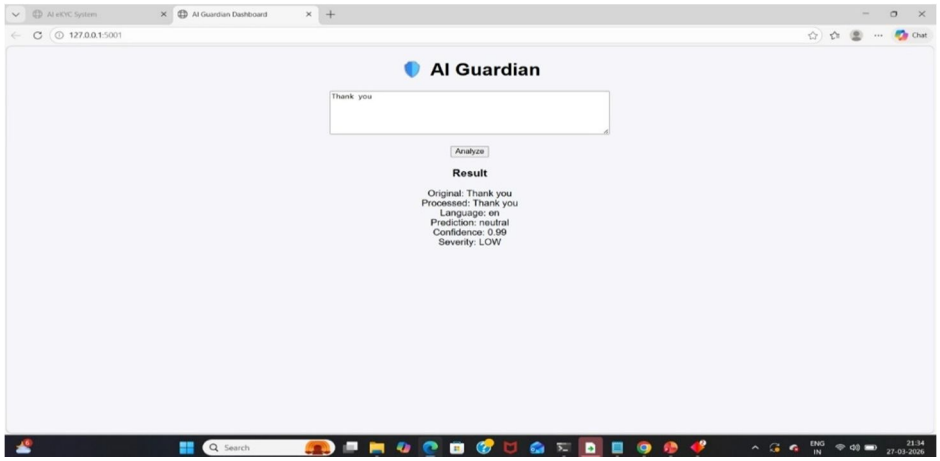


Fig 6.2: Detection of English Neutral language

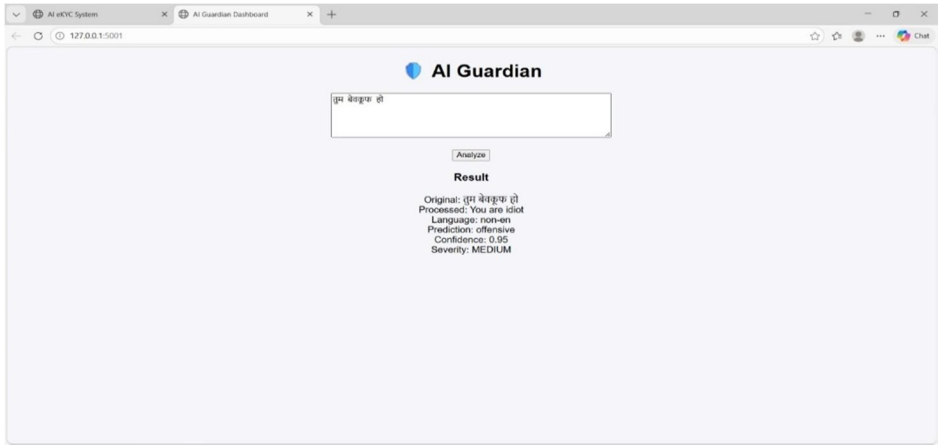


Fig 6.3: Detection of Hindi Offensive language

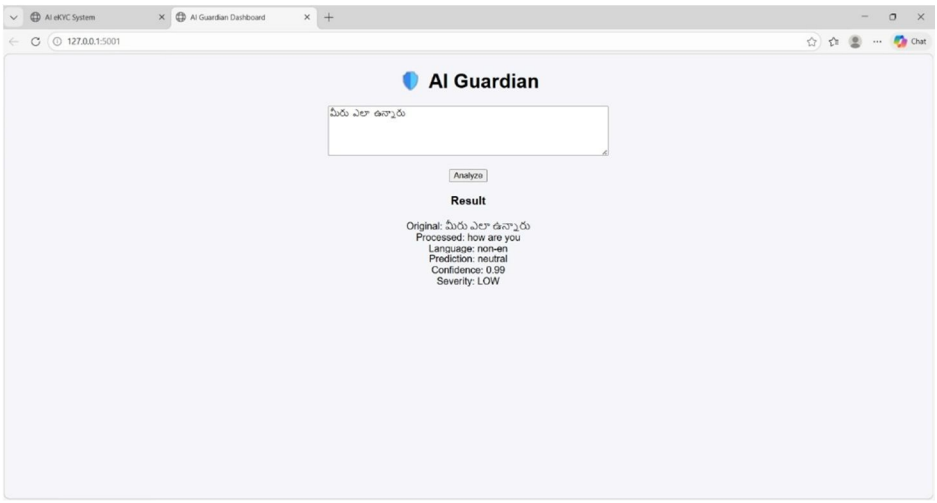


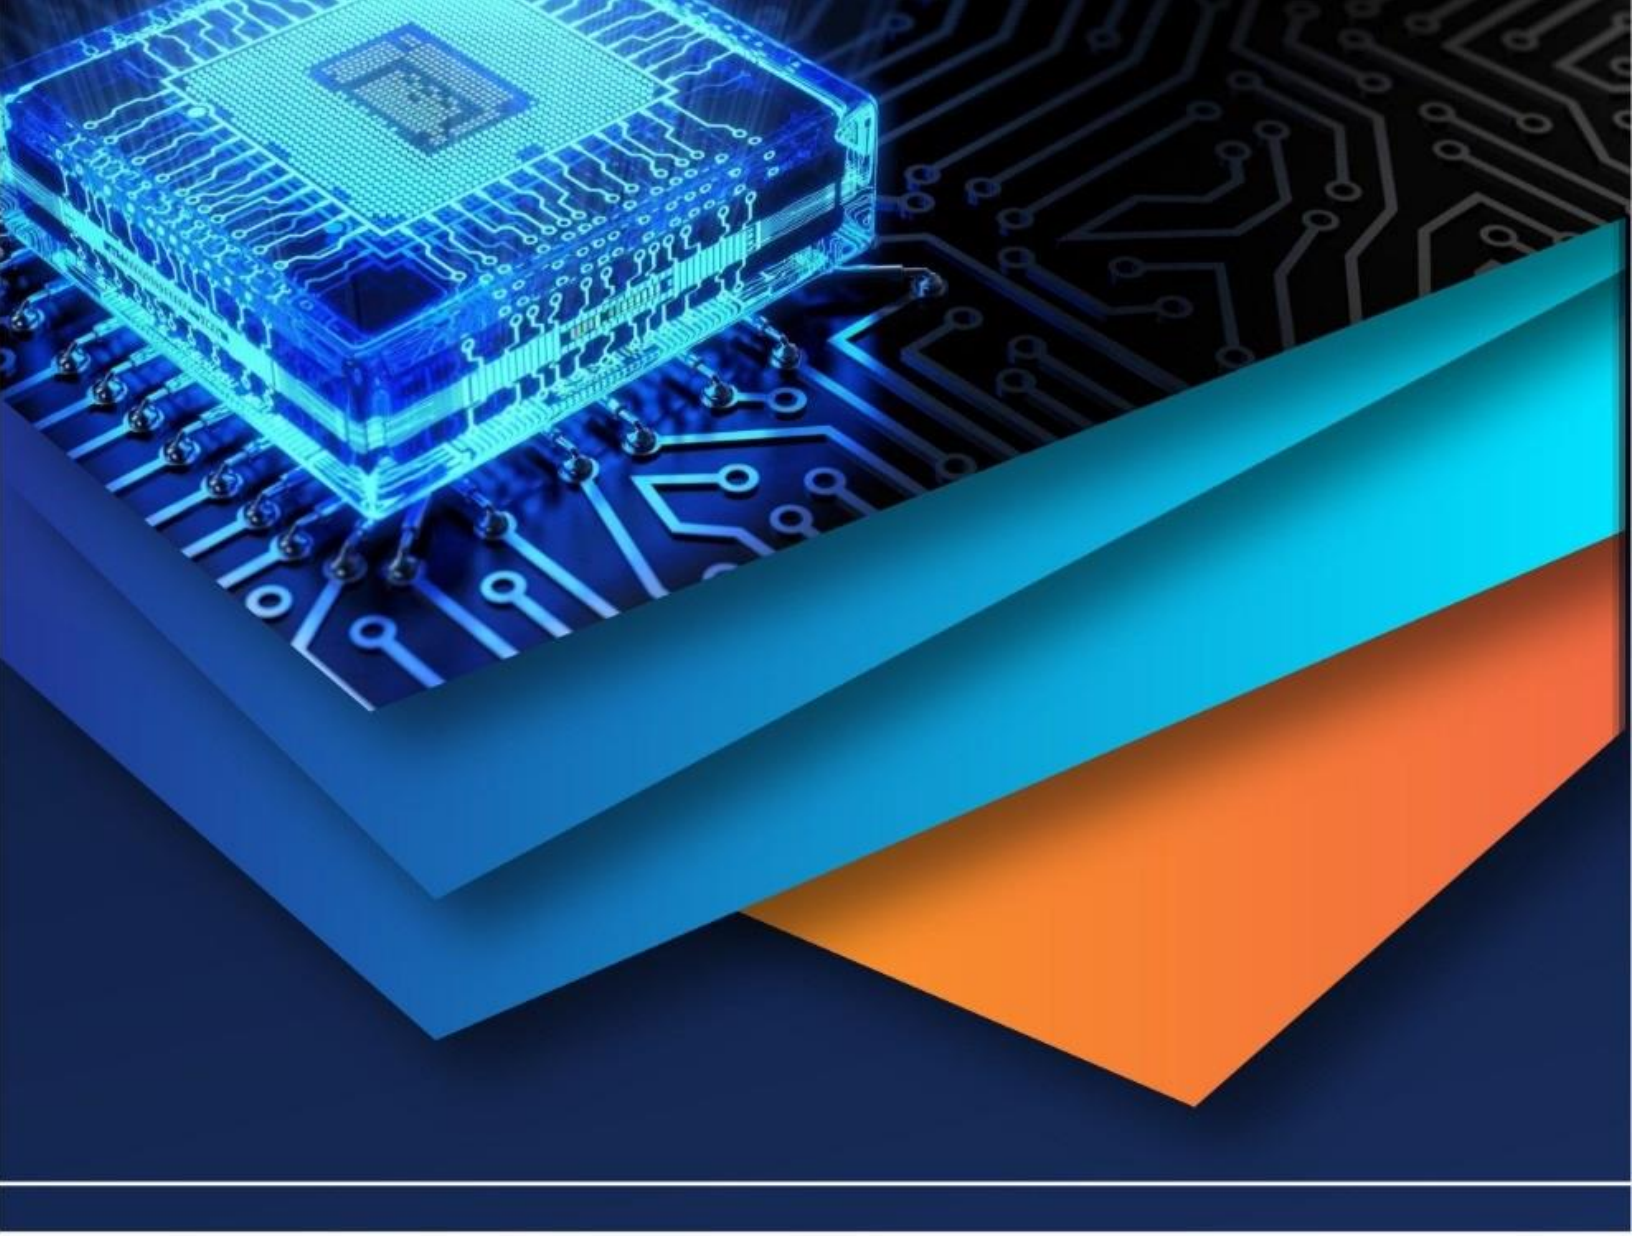
Fig 6.4: Detection of Telugu Neutral language

VII. CONCLUSION

This paper has presented a novel, intelligent, and scalable framework, namely AI Guardian, which is capable of effectively detecting multilingual hate speech and cyberbullying in real-time digital communication environments. The proposed system effectively overcomes the major shortcomings of existing content moderation techniques, which are incapable of dealing with multilingual content and lack contextual understanding, and are not transparent in their decision-making process. The proposed system effectively leverages advanced Natural Language Processing techniques along with transformer-based models, namely Multilingual BERT and XLM-RoBERTa, to effectively comprehend semantic and contextual relationships in text-based user-generated content. The severity scoring module further boosts the capabilities of the proposed system, which categorizes hate speech and cyberbullying based on their severity level, thus prioritizing their moderation process. The proposed system, namely AI Guardian, possesses a significant advantage in terms of its Explainable AI module, which provides clear and interpretable outputs in terms of offensive words and their reasoning, thus promoting higher levels of user trust and ease of usage. The proposed system is designed with a modular and efficient architecture, which enables its real-time usage with minimal latency levels. The experimental results have proved the high accuracy levels of the proposed system, which are in the range of 92-95%, thus validating its effectiveness in dealing with multilingual and code-mixed hate speech and cyberbullying. The proposed system, namely AI Guardian, is a robust, interpretable, and scalable solution in dealing with harmful online content, which contributes significantly to promoting a safe digital communication environment.

REFERENCES

[1] OpenAI / Google Research, “Large Language Models for Toxicity Detection and Content Moderation,” 2025.
[2] A. Bhardwaj et al., “Hinglish Offensive Language Detection using Deep Learning Techniques,” 2024.
[3] R. Kumar et al., “Cyberbullying Detection using Hierarchical Transformer Networks,” 2024.
[4] M. Ribeiro et al., “Explainable Content Moderation using Attention-Based Models,” 2024.
[5] T. Ranasinghe et al., “Cross-lingual Hate Speech Detection using XLM-RoBERTa,” 2023.
[6] T. Mandl et al., “Overview of the HASOC 2023 Shared Task: Hate Speech and Offensive Content Identification in Indo-European Languages,” 2023.
[7] A. Saha et al., “Contextual Hate Speech Detection in Social Media Conversations,” 2023.
[8] A. Mozafari et al., “Hate Speech Detection with BERT and CNN: A Study on Explainability,” 2023.
[9] F. Ousidhoum et al., “Multilingual and Multi-Aspect Hate Speech Analysis Dataset,” 2022.
[10] I. Glavaš et al., “Zero-Shot Cross-Lingual Hate Speech Detection,” 2022.
[11] Z. Zhang et al., “Multimodal Hate Speech Detection in Memes using Vision-Language Models,” 2022.
[12] P. Fortuna and S. Nunes, “A Survey on Automatic Detection of Hate Speech in Text,” ACM Computing Surveys, 2018.
[13] M. Zampieri et al., “SemEval-2019 Task 6: Identifying and Categorizing Offensive Language in Social Media (OffensEval),” 2019.
[14] Z. Waseem and D. Hovy, “Hateful Symbols or Hateful People? Predictive Features for Hate Speech Detection on Twitter,” NAACL, 2016.
[15] A. Das et al., “Hate Speech Detection in Code-Mixed Hindi-English Text,” 2021.
[16] J. Devlin et al., “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding,” NAACL, 2019.
[17] A. Conneau et al., “Unsupervised Cross-lingual Representation Learning at Scale (XLM-R),” 2020.
[18] T. Wolf et al., “HuggingFace Transformers: State-of-the-Art Natural Language Processing,” 2020.
[19] S. Hochreiter and J. Schmidhuber, “Long Short-Term Memory,” Neural Computation, 1997.
[20] Y. Kim, “Convolutional Neural Networks for Sentence Classification,” EMNLP, 2014.



10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089 (24*7 Support on Whatsapp)