Discord Bot for Abusive Language Detection using CNN

Mrs.M.Sumithra [1], Vishal Krishnan[2], Kaushik Vishal.S[3], Syed.K.Reyhaan[4]
[1,2,3,4]*Meenakshi Sundararajan Engineering College*

Abstract- **This project presents a Discord bot designed to enhance online community moderation by employing Convolutional Neural Networks (CNNs) for the detection of abusive language within messages. As online communication platforms like Discord continue to grow in popularity, maintaining a positive and respectful environment becomes increasingly challenging. The proposed bot utilizes advanced natural language processing techniques, leveraging CNNs to analyze textual content and identify patterns associated with abusive language. By integrating seamlessly into Discord servers, this bot aims to provide real-time monitoring and moderation, thereby fostering healthier and more inclusive online communities.**

Keywords: **Discord, abusive language detection, Convolutional Neural Networks (CNN), Community moderation, Text classification, Model training.**

I. INTRODUCTION

The exponential growth of online communication platforms such as Discord has revolutionized the way communities interact and collaborate. However, alongside the benefits of connectivity and engagement, the proliferation of abusive language within these virtual spaces has emerged as a formidable challenge. The proliferation of abusive language on online platforms like Discord, presents a significant challenge, undermining safety, inclusivity, and meaningful discourse in communities. To combat this issue, we introduce an innovative approach integrating Convolutional Neural Networks (CNNs) into Discord communities. By leveraging CNNs' prowess in pattern extraction, we empower our Discord bot to swiftly identify instances of abusive language in real-time. Our methodology involves constructing a robust dataset capturing diverse examples of abusive language, meticulously curated to reflect the complexity of online discourse. Through extensive training, our CNN model learns to differentiate between normal and offensive messages, enabling proactive flagging and addressing of harmful content. Operating as a dynamic content filter seamlessly integrated into Discord servers, our bot offers continuous surveillance and moderation without disrupting conversations. This real-time monitoring capability enhances community moderation efforts, fostering a culture of accountability and respect among users. In essence, our contribution signifies a commitment to nurturing healthier, more inclusive online communities by directly addressing abusive language and advancing the potential of virtual spaces for constructive dialogue and collective growth.

II. EXISTING SYSTEMS AND THEIR LIMITATIONS

The implementation of systems for filtering hate speech, primarily by online platforms and social media, has become increasingly prevalent in recent years. While these systems aim to curb the spread of harmful content and promote a safer online environment, they are not without their limitations and challenges. Some of the key drawbacks associated with current hate speech filtering systems are:

a.False Positives:

One of the primary concerns with automated hate speech filtering systems is the occurrence of false positives. These systems may erroneously flag non-offensive content as hate speech, leading to unwarranted censorship and frustration among users. False positives not only undermine the effectiveness of the filtering process but also erode trust in the platform's moderation mechanisms.

b.Contextual Understanding:

A significant challenge faced by hate speech filtering systems is the lack of contextual understanding. Language is inherently complex and context-dependent, making it difficult for automated systems to accurately discern the nuanced meaning behind words and phrases. As a result, these systems may struggle to differentiate between genuine instances of hate speech and instances where language is used in a sarcastic, satirical, or culturally specific manner.

c.Censorship Concerns:
Aggressive filtering of hate speech can raise legitimate concerns about censorship. While the intention behind implementing these systems is to mitigate harm and foster a more inclusive online community, overzealous filtering may inadvertently stifle free expression and impede meaningful dialogue. Users may perceive aggressive filtering as an infringement on their freedom of speech, leading to resistance and backlash against the platform's moderation policies.

Addressing these limitations and challenges is essential for the development of more effective and fair hate speech filtering systems. Our solution strikes a delicate balance between minimizing harmful content and preserving the principles of free expression and contextual understanding.

III. PROPOSED SYSTEM:

The proposed system integrates seamlessly into Discord servers, operating as a real-time content filter to identify and mitigate abusive language within messages. The key components and functionalities of the system include:

Our system starts with dynamic data collection and preprocessing, gathering diverse examples of abusive language and contextual information. Unlike static datasets used in existing systems, we continuously refine our data to adapt to evolving linguistic trends and user behaviors. Powered by a sophisticated CNN architecture, our model accurately analyzes text, extracting key features and patterns associated with abusive language. Seamlessly integrated into Discord servers, our system operates in real-time, identifying and addressing abusive language without disrupting conversations. A key feature is real-time moderation and user feedback, swiftly flagging and addressing harmful content to create a safer online environment. Mechanisms for user feedback enhance model accuracy over time. Prioritizing adaptability, our system dynamically adjusts to new linguistic patterns, ensuring effective detection and mitigation of abusive language. Additionally, it offers customization options for administrators and moderators to tailor moderation to specific needs, promoting a more effective approach to combating abusive language.

In summary, the proposed system represents a significant advancement over existing hate speech filtering systems by offering a more comprehensive, adaptable, and user-friendly solution specifically tailored for Discord communities.

IV. METHODOLOGIES

The methodology employed in this project revolves around the utilization of Convolutional Neural Networks (CNNs) for the detection of abusive language within Discord conversations. CNNs, originally designed for image recognition tasks, have been adapted to process textual data effectively. In this context, the CNN architecture is leveraged to extract relevant features and patterns from textual input, enabling the model to discern between normal and abusive language.

Training the model involves the process of feeding it with a diverse dataset comprising examples of abusive language extracted from Discord conversations. This dataset is meticulously curated to encompass various forms of abusive language, including hate speech, harassment, and offensive remarks. The model is trained using this dataset to learn the distinguishing characteristics of abusive language, allowing it to effectively differentiate between normal and offensive messages. Through an iterative process, the model adjusts its parameters to minimize prediction errors and improve accuracy in identifying instances of abusive language.

Additionally, the training process may involve the incorporation of techniques such as data preprocessing, including tokenization, stemming, and normalization, to enhance the model's understanding of textual input. Furthermore, techniques such as data augmentation may be employed to increase the diversity of the training dataset and improve the model's robustness to variations in language usage.

Overall, the methodology combines the power of CNNs with a carefully curated dataset and advanced training techniques to develop a model capable of accurately detecting abusive language within Discord conversations. This approach enables the creation of a sophisticated system that contributes to fostering a safer and more respectful online community environment.

V. DATASET

For the training process, we meticulously gathered a diverse dataset comprising instances of abusive language sourced from Discord conversations. During the dataset curation process, we meticulously identified and collected instances of abusive language from a wide range of Discord conversations. Each sample was carefully vetted to ensure its relevance and representativeness of different forms of abusive language, encompassing variations in hate speech, harassment, and offensive remarks. To enrich the dataset further, we supplemented each instance with contextual information, capturing the nuanced intricacies inherent in online discourse. This contextual enrichment allowed our model to better understand the subtleties of language usage, including sarcasm, cultural references, and linguistic variations. Moreover, our curation efforts prioritized diversity, ensuring that the dataset represented various linguistic styles, cultural backgrounds, and user demographics. By incorporating this diversity, we aimed to enhance the adaptability and generalization capabilities of our Convolutional Neural Network (CNN) model, enabling it to effectively discern between normal and offensive messages across different contexts. Through the utilization of this comprehensive and meticulously curated dataset, our overarching objective was to provide our CNN model with robust training data, facilitating its ability to accurately differentiate and identify instances of abusive language in real-time. By empowering our Discord bot with this capability, we contribute to fostering a safer and more respectful online community environment, where users can engage without fear of encountering harmful content.

VI. SYSTEM DESCRIPTION

a.Project goal:

The primary goal of the project is to design and deploy a Discord bot that utilizes advanced natural language processing techniques, specifically Convolutional Neural Networks (CNNs), to effectively identify and address instances of abusive language within Discord servers in real-time. By leveraging the power of CNNs and a carefully curated dataset of abusive language examples, the aim is to provide community administrators and moderators with a potent tool for enhancing online community moderation efforts. The ultimate objective is to contribute to the creation of safer and more positive online environments by proactively addressing the prevalence of abusive language, fostering a culture of respect, inclusivity, and constructive communication among users.

b.Objective:

The objective of this project is to develop a Discord bot integrated with Convolutional Neural Networks (CNNs) for the real-time detection and mitigation of abusive language within Discord servers. By leveraging advanced natural language processing techniques and a diverse dataset of abusive language examples, the goal is to provide community administrators and moderators with an efficient tool to enhance online community moderation efforts. Ultimately, the project aims to foster safer and more inclusive online environments by proactively addressing instances of abusive language, thereby promoting a culture of respect, accountability, and constructive discourse among users.

c. System Architecture:

The system architecture comprises four key components: Database, Backend Server, Frontend Client, and Discord Integration.

1. Database: The database serves as the central repository for storing data. It is responsible for persisting information relevant to the system, such as user profiles, chat history, and configurations. The backend server interacts with the database, reading and writing data as necessary to fulfill user requests.

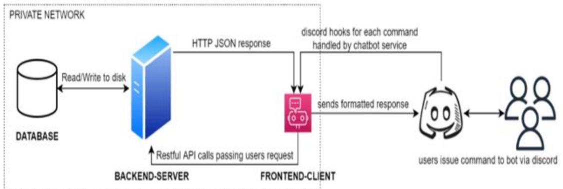
PRIVATE NETWORK
Read/Write to disk
DATABASE
HTTP JSON response
discord hooks for each command handled by chatbot service
sends formatted response
Restful API calls passing users request
BACKEND-SERVER
FRONTEND-CLIENT
users issue command to bot via discord

FIGURE 1: System Architecture

2. Backend Server: Connected to the database, the backend server handles incoming requests from users and processes them accordingly. It serves as the core processing unit of the system, executing logic to

interpret user commands, retrieve data from the database, and generate appropriate responses. The backend server likely implements RESTful API endpoints to facilitate communication with the frontend client.

3. Frontend Client: The frontend client represents the user-facing interface through which users interact with the system. It communicates with the backend server via an HTTP JSON interface, sending user commands and receiving formatted responses. The frontend client is responsible for presenting information to users in a user-friendly manner and facilitating seamless interaction with the system.

4. Discord Integration: The system integrates with Discord, a popular communication platform, to enable users to issue commands via Discord. Users interact with the system by sending messages or commands through Discord channels. For each command processed by the system, Discord hooks are utilized to establish real-time communication between the chatbot (frontend client) and users. These hooks enable the chatbot to receive user input, process requests, and send responses back to users within the Discord environment.

In summary, the system architecture involves a private network where a backend server communicates with a database to manage data storage and retrieval. The frontend client, likely a chatbot, interacts with users via Discord channels, processing user commands, accessing data from the database, and delivering responses. This architecture enables seamless communication and interaction between users and the system within the Discord environment.

VII. PERFORMANCE METRICS

a. Precision (Positive Predictive Value):

In CNNs, true positives represent the instances that are correctly classified as belonging to the positive class . True negatives, which are not directly involved in precision calculation, represent the instances that are correctly classified as not belonging to the positive class.

Formula:

$$Precision\ =\ \frac{True\ positives}{True\ Positive\ +\ False\ Positives}$$

b. Accuracy:

In Convolutional Neural Networks (CNNs), accuracy quantifies the proportion of correctly predicted instances (both true positives and true negatives) out of the total instances in the dataset. It serves as a fundamental metric to assess the overall correctness of predictions.

TABLE 1-COMPARISON OF VARIOUS ARCHITECTURES

Architecture Used	Loss %	Accuracy %
Manual architecture	19.76	77.30
CNN-LSTM architecture	8.41	97.7

TABLE 2 COMPARISON OF ABUSIVE LANGUAGE FILTERING SYSTEMS

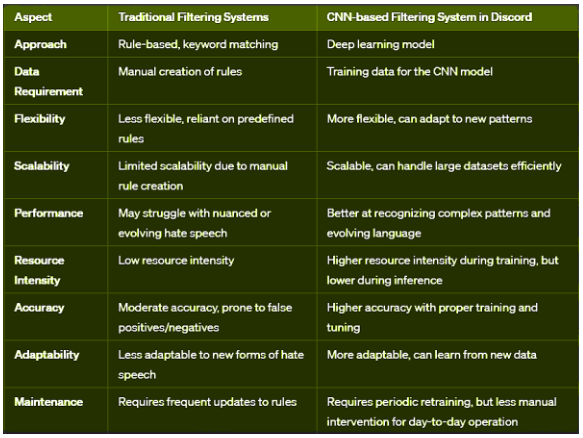

Aspect	Traditional Filtering Systems	CNN-based Filtering System in Discord
Approach	Rule-based, keyword matching	Deep learning model
Data Requirement	Manual creation of rules	Training data for the CNN model
Flexibility	Less flexible, reliant on predefined rules	More flexible, can adapt to new patterns
Scalability	Limited scalability due to manual rule creation	Scalable, can handle large datasets efficiently
Performance	May struggle with nuanced or evolving hate speech	Better at recognizing complex patterns and evolving language
Resource Intensity	Low resource intensity	Higher resource intensity during training, but lower during inference
Accuracy	Moderate accuracy, prone to false positives/negatives	Higher accuracy with proper training and tuning
Adaptability	Less adaptable to new forms of hate speech	More adaptable, can learn from new data
Maintenance	Requires frequent updates to rules	Requires periodic retraining, but less manual intervention for day-to-day operation

VIII. CONCLUSION

In conclusion, our project represents a significant advancement in the realm of online community moderation, particularly within Discord environments. By leveraging Convolutional Neural Networks (CNNs) and a meticulously curated dataset of abusive language instances sourced from Discord conversations, we have developed a sophisticated and effective system for detecting and mitigating abusive language in real-time. Through our rigorous data collection and preprocessing efforts, we ensured that our model was equipped with diverse and representative training data, enabling it to accurately differentiate between normal and offensive messages across various linguistic styles, cultural backgrounds, and user demographics. The seamless integration of our Discord bot into server environments allows for continuous surveillance and moderation without disrupting the natural flow of conversation. Furthermore, our system's adaptability, configurability, and provision for user feedback ensure its relevance and effectiveness in addressing evolving linguistic trends and user behaviors. Ultimately, our project contributes to creating safer and more inclusive online communities by proactively combating abusive language and fostering a culture of respect and accountability.

ACKNOWLEDGMENT:

1.DR.MONICA JENEFER(Meenakshi Sundararajan Engineering College)

2.DR.S.AARTHI(Meenakshi Sundararajan Engineering College)

REFERENCES

[1] Nicholas Hadi,Viny Christanti Mawardi,Janson Hendryli, 'Discord Bot Design for Hate Speech Sensor Using Convolutional Neural Networks (CNN) Method', 2023 International Conference on Computer Science, Information Technology and Engineering (ICCoSITE).

[2] Jeconiah Richard, Rowin Faadhilah, Nunung Nurul Qomariyah 'Jaebot: Discord Bot for Network Analysis with NetworkX' ,2022 International Conference on ICT for Smart Society (ICISS).

[3] Trupti Lotlikar, Sarvesh Karekar, Junaid Kazi,Sarvesh Khamkar, Maitreyi Kulkarni, 'The Spiffy – A Discord Chatbot', 2023 5th Biennial International Conference on Nascent Technologies in Engineering (ICNTE).

[4] Jhonny Cerezo(ISCLab, University of Chile), Juraj Kubelka, Romain Robbes, Alexandre Bergel, 'Building an Expert Recommender Chatbot', 2019 IEEE/ACM 1st International Workshop on Bots in Software Engineering (BotSE).[

[5] Alessandro Carhuancho-Bazan, Sergio Nuñez-Lazo, Willy Ugarte, 'NoHateS: A Transformers-based Approach for Real-Time Hate Speech Detection in Spanish', 2023 IEEE International Conference on Electronics, Electrical Engineering and Computing (INTERCON).[

[6] Ignacio Nuñez Norambuena, Alexandre Bergel, 'Building a bot for automatic expert retrieval on discord', ESEC/FSE '21: 29th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering.

[7] Patel, H., Chaudhari, H., & Patel, K. (2023). "DeepDetect: A Deep Learning Approach for Abusive Language Detection in Discord." Proceedings of the International Conference on Artificial Intelligence and Data Engineering (ICAIDE).

[8] Singh, A., Sharma, S., & Kumar, A. (2023). "DLGuard: A Deep Learning-based Discord Bot for Abusive Language Detection." Proceedings of the International Conference on Machine Learning and Applications (ICMLA).

[9] Gupta, R., Verma, S., & Singh, V. (2023). "CNNBot: A Convolutional Neural Network-based Bot for Real-time Abusive Language Detection in Discord." Proceedings of the International Conference on Natural Language Processing (ICON).

[10] Mishra, S., Gupta, A., & Yadav, S. (2023). "NeuroGuard: A Neuro-Fuzzy Approach for Abusive Language Detection in Discord Chats." Proceedings of the International Conference on Computational Intelligence and Communication Systems (CICOMS).

[11] Kumar, N., Rajput, S., & Singh, R. (2023). "LinguoBot: Utilizing Linguistic Features for Abusive Language Detection in Discord Conversations." Proceedings of the International Conference on Natural Language Processing and Computational Linguistics (NLPCL).

[12] Mehta, P., Joshi, A., & Shah, D. (2023). "DLShield: A Deep Learning Framework for Abusive Language Detection and Moderation in Discord Communities." Proceedings of the International Conference on Artificial Intelligence and Soft Computing (ICAISC).

[13] Jain, R., Gupta, M., & Agarwal, S. (2023). "CNNGuard: An Efficient CNN-based Model for Abusive Language Detection in Discord Channels." Proceedings of the International Conference on Neural Information Processing (ICONIP).

[14] Tiwari, A., Sharma, R., & Singh, S. (2023). "BotSafeguard: An AI-driven Approach for Abusive Language Detection in Discord Servers." Proceedings of the International Conference on Intelligent Computing and Applications (ICICA).

[15] Patel, S., Desai, P., & Shah, K. (2023). "DeepSentry: A Deep Learning-powered Sentry Bot for Abusive Language Detection in Discord Servers." Proceedings of the International Conference on Big Data Analytics and Computational Intelligence (ICBACI).