



iJRASET

International Journal For Research in
Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 13 Issue: VI Month of publication: June 2025

DOI: <https://doi.org/10.22214/ijraset.2025.71986>

www.ijraset.com

Call: ☎ 08813907089

E-mail ID: ijraset@gmail.com

Hate Speech Classification for Text

Ajay Bhagat, Snehal Bhalerao, Radhika Jagtap, Siddhi More, Dr.Dinesh Bhagwan Hanchate, Prof. Swati Mahadev Atole, Dr.Sachin Sukhadeo Bere

Department of Computer Engineering, Dattakala Group of Institution Faculty of Engineering, Bhigwan-413130, Savitribai Phule Pune University, Maharashtra, INDIA

Abstract: Hate speech has turned into an offense that has increased in the recent past, both online and offline. There are several factors that explain why this is so. On one side, people are more inclined to act violently because of the anonymity provided by the internet, and social networks in particular. Conversely, people's need to express themselves on the internet has grown, and with this has been the prevalence of hate speech. Given how detrimental this kind of discriminatory speech is to society, detection and prevention by social media companies and governments can both be useful. With this survey, we provide an overall overview of the work that has been accomplished in the field, which addresses this dilemma. The use of many complex and non-linear models made this challenge possible, and CAT Boost carried out the others because it employed latent semantic analysis (LSA) for dimensionality reduction. Hate speech refers to abusive or discriminatory language directed at an individual or group based on attributes such as race, religion, ethnicity, gender, or sexual orientation. Its proliferation can result in concrete damage and annihilate the security and inclusivity of online environments. This paper presents a machine learning approach to hate speech categorization in text. Utilizing publicly accessible labeled data, we examine various natural language processing (NLP) techniques and supervised learning algorithms like logistic regression, support vector machines, and deep learning models to categorize hate speech, offensive language, and innocuous content. Key features such as n-grams, TF-IDF scores, and word embeddings are exploited to enhance model performance. The results validate the effectiveness of utilizing linguistic features in combination with optimal classifiers to achieve high precision and accuracy for hate speech detection. The findings emphasize the importance of using well-balanced datasets and ethical considerations while developing automated content moderation systems.

I. INTRODUCTION

With the digital era, the Internet platforms are now at the center of public discussion, where people can share reviews and pass on information to large audiences. However, this openness also creates conditions for the spread of harmful material, such as hate speech, words that attack or demean an individual or group based on characteristics such as race, religion, gender, or sexual orientation. Hate speech categorization is the process of automatically detecting and filtering such abusive content based on natural language processing (NLP) and machine learning (ML) methods. Creating effective and precise models for hate speech detection is not only important for ensuring platform integrity but also for safeguarding users from harassment and encouraging respectful communication. This work is specifically on the construction and deployment of a hate speech classifier based on supervised learning methodologies. Through feature analysis of annotated datasets and consideration of multiple linguistic and contextual information, the system will be capable of differentiating between hate speech, offensive terms, and innocent content with accuracy and consistency.

Hate speech has become a growing issue over the past decade, both online and offline. Many things are at play in this case. The internet anonymity makes individuals more likely to behave aggressively but also more likely to post their views online, which fuels hate speech. Governments and social media corporations can gain from detection and prevention methods as this type of discriminatory speech can have a catastrophic effect on society. We wish that our survey might enlighten us a little on the vast amount of research that has been conducted in this area.

Hate speech is any speech that has the potential to cause harm to a person or group and that could lead to violence, insensitivity, or irrational or inhuman conduct. With growth in the usage of social media sites like Twitter and Facebook, the increase in the usage of hate speech. The existence of hate speech has a proven connection with the increase of hate crimes. As hate speech becomes a controversial topic, multiple government-led endeavors are being instituted, such as the Council of Europe's campaign No Hate Speech. EU Hate Speech Code of Conduct that has to be signed and obeyed by all the social networks in 24 hours is one such method by which it has come into implementation. Many problems raised by it have caused serious questions about dataset quality, which this work tries to remedy. This piece also addresses the second issue, that is, prior to making an effective classifier, the optimal features used to classify hate speech need to be studied and determined.

Hate speech is abusive or threatening language that is directed at an individual or a group of people based on features like race, religion, gender, or sexual orientation. Hate speech can have severe adverse effects on users and online conversations.

This thesis tackles the issue of detecting hate speech by creating a machine learning-based classifier to identify hateful messages in text. The research investigates a range of linguistic and contextual features and tests various supervised learning models to determine their performance. The aim is to develop a strong model that will be part of the continued attempt to make digital spaces safer and more tolerant.

The rise of hate speech online poses an urgent challenge to both platform operators and society in general. Hate speech detection through automated means has emerged as a vital research problem in natural language processing (NLP), working towards limiting the spread of offensive content while maintaining freedom of expression.

With the age of fast digital communication, the spread of user-generated content on the web has brought sophisticated challenges in sustaining healthy discussion. Of these challenges, hate speech has been a particularly insidious type of abusive content. Hate speech is any form of communication that insults an individual or a group on the basis of characteristics like race, religion, ethnicity, gender, sexual orientation, nationality, or other identity features. It not only discredits social cohesion but also adds to actual-world violence, discrimination, and exclusion. As social networking sites, forums, and comment fields increase in usage, so does the amount of content to be moderated. Manual moderation is inefficient and labor-intensive at large scale.

Classification of hate speech is a difficult task for a number of reasons:

- **Subjectivity and Context-Dependence:**

What is hate speech may differ between cultures, communities, and individuals. Sarcasm, contextual subtleties, and implicit bias make it hard for machines to accurately determine intent.

- **Data Imbalance:**

Hate speech occurrences are frequently sparse relative to neutral or non-hateful material, resulting in highly skewed datasets that can skew learning algorithms.

- **Evolving Language:**

Users often adopt coded language, slang, or euphemisms to bypass moderation, necessitating continuous updates to models and lexicons.

A. Terminologies of Hate Speech Classification for Text

Datasets usually fall into one of these categories because of:

1) *Hate Speech:*

Hate speech is any message that insults, discriminates, incites violence or prejudice against, or demeans or demoralizes individuals or groups on the basis of immutable or inherent characteristics like race, religion, ethnic origin, gender, nationality, sexual orientation, or disability. Hate speech can be overt (e.g., slurs or insults) or veiled (e.g., coded language or racist stereotypes). It is uniquely dangerous because it may lead to offline harm such as social exclusion, violence, and institutionalized discrimination.

2) *Offensive Language:*

Offensive language consists of abusive, profane, vulgar, or disrespectful language, but it does not necessarily target protected groups or promote hatred. It usually involves personal insults or rude remarks and can violate community standards without qualifying as hate speech. Examples include general cursing or trolling. Differentiating between offensive language and hate speech is essential for ethical AI implementation and effective moderation.

3) *Toxicity:*

Toxicity is a more general term encompassing any content that is hateful, inflammatory, or likely to degrade the quality of discourse. It includes hate speech, abusive comments, harassment, threats, bullying, and other forms of verbal abuse.

Models such as the Perspective API by Jigsaw/Google use toxicity scores to assess comment harmfulness and support real-time moderation.

4) *Classification:*

Classification is the process of labeling input data with predefined categories. In text, this means analyzing written content and predicting its class—such as “hate speech,” “offensive language,” or “neutral.”

It is a supervised learning task trained on annotated datasets. Classification tasks can be binary or multi- class/multi-label based on the application.

5) *Machine Learning:*

Machine Learning (ML) enables systems to learn from data and make predictions. In hate speech detection, ML models learn linguistic patterns from labeled data. Common ML methods include:

- **LogisticRegression:** Suitable for binary classification tasks.
- **Support Vector Machines (SVM):** Effective for high-dimensional text data.
- **Decision Trees:** Simple and interpretable, but prone to overfitting; often enhanced through ensemble methods like Random Forest or XGBoost.

6) *Natural Language Processing (NLP):*

NLP is a subfield of AI focused on enabling machines to understand and generate human language. In hate speech classification, NLP provides tools for extracting syntactic, semantic, and contextual information from text. Techniques such as named entity recognition, part-of-speech tagging, and semantic similarity are instrumental in creating accurate classification systems.

7) *Text Preprocessing:*

Text preprocessing prepares raw text data for analysis and model training. Typical preprocessing steps include:

- **Tokenization:** Dividing text into words or phrases.
- **Stop Word Removal:** Eliminating common but non-informative words.
- **Stemming and Lemmatization:** Reducing words to their root or base form.
- **Lowercasing and punctuation removal:** To normalize text.

These steps reduce noise and standardize input for machine learning models.

8) *Tokenization:*

Tokenization is the process of breaking a text string into smaller units called tokens, which can be words, subwords, or characters. This is a foundational step in NLP, as most models operate on tokens. Advanced methods like Byte-Pair Encoding (BPE) and WordPiece (used in BERT) help manage rare and compound words more effectively.

9) *Sentiment Analysis:*

Sentiment analysis detects emotional tone or opinion in text and classifies it as positive, negative, or neutral. While not equivalent to hate speech detection, it can act as a supplementary tool. Negative sentiment often coexists with toxic or hateful language, although sentiment analysis alone lacks the specificity to detect targeted or ideological hate.

10) *Precision and Recall:*

These are core metrics in evaluating classification models:

$$\bullet \text{ Precision} = \frac{\text{True Positives}}{\text{True Positives} + \text{False Positives}} :$$

Measures the proportion of predicted hate speech instances that were actually hate speech.

$$\bullet \text{ Recall} = \frac{\text{True Positives}}{\text{True Positives} + \text{False Negatives}} :$$

Measures how many actual hate speech instances were correctly identified by the model.

High precision minimizes false positives, while high recall minimizes false negatives. A good balance is vital for fair and effective hate speech detection.

11) F1Score:

The F1 Score is the harmonic mean of precision and recall, offering a balanced performance metric:

$$F1 = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

It is especially useful in scenarios with imbalanced datasets, where accuracy may be misleading. A high F1 score indicates the model is neither too lenient nor too strict in classifying hate speech.

II. LITERATURE SURVEY

A literature survey of hate speech classification for text would report the current research, methods, datasets, challenges, and advancements in this area.

A. Introduction to Hate Speech Classification:

Hate speech isn't just about harsh words—it's about language that targets people because of who they are. Whether it's based on race, religion, gender, or other core identities, hate speech aims to insult, exclude, or harm.

It's different from offensive or rude language, which might be inappropriate but not necessarily discriminatory. It also overlaps with things like online harassment and abuse, but not every insult is hate speech.

As social media becomes a bigger part of our daily lives, the need to detect and filter hate speech automatically has grown. These platforms allow messages to spread quickly—sometimes with real-world consequences. But building automated tools to identify hate speech is tricky.

Definitions of what counts as "hate speech" vary depending on the culture or country, and there's the added challenge of keeping these systems fair and unbiased. We also have to be careful when dealing with sensitive issues related to protected groups, making it vital to approach this work ethically and thoughtfully.

B. Datasets for Hate Speech Classification:

To teach a computer how to recognize hate speech, we first need examples. That's where datasets come in. Some of the most well-known datasets include:

- Hatebase: A multilingual collection of hate-related terms.
- Twitter datasets: Collections of tweets labeled as hate speech, offensive, or neutral.
- Facebook and YouTube datasets: Where comments are annotated based on their toxicity or harmfulness.

These datasets help train and test detection models, but they're not perfect. Many are *platform-specific*, meaning a model trained on Twitter might not work well on Reddit.

Some datasets also suffer from *bias* or a *lack of diversity*, especially when they focus only on English or specific regions. That's why expanding and improving our data sources is an ongoing challenge.

C. Preprocessing Techniques:

Before we feed text into a machine learning model, we need to clean it up—just like editing a rough draft before turning it in. Here's what that usually involves:

- Text Cleaning: Lowercasing all words, removing special characters, punctuation, or filler words like "the" or "is" that don't add much meaning.
- Tokenization, Stemming, and Lemmatization:
 - Tokenization breaks the text into words or phrases.
 - Stemming and lemmatization reduce words to their root forms—so "running" and "ran" both become "run," for example.
- Handling Imbalanced Data: Hate speech is often rare compared to regular, non-harmful content. This imbalance can confuse models, so we use techniques like:
 - Oversampling the minority class.
 - Undersampling the majority class.
 - Data augmentation to create synthetic examples.

D. Feature Extraction Techniques:

Machines don't understand language the way we do—they need numbers. So we have to convert text into numerical formats through *feature extraction*.

- **Bag of Words (BoW) and TF-IDF:** These are traditional methods where each word becomes a number based on how often it appears. TF-IDF adjusts scores by giving less importance to common words.
- **Word Embeddings:** Tools like *Word2Vec*, *GloVe*, and *fastText* create word maps that show how words relate to each other. For example, “king” and “queen” might be closer together than “king” and “banana.”
- **Contextual Embeddings:** Models like *BERT* and *RoBERTa* don't just look at individual words—they understand the meaning based on surrounding words. That means they can tell the difference between “bad dog” (as a scolding) and “bad guy” (as a villain).

E. Machine Learning Models for Hate Speech Classification:

Once we have features, it's time to choose a model. Early on, researchers used simpler methods like:

- **Naïve Bayes, Logistic Regression, and SVM:** These classical models are fast and easy to use, especially with BoW or TF-IDF features.
- **Then came neural networks:**
- **CNNs and RNNs (LSTM, GRU):** These are better at understanding sequences of text and picking up patterns that older models miss.
- **Now, transformers have taken over:**
- **Models like BERT, GPT, and RoBERTa** have been pre-trained on massive amounts of data and fine-tuned for specific tasks like hate speech detection. They're great at capturing context, which makes them highly accurate.

Finally, some researchers are combining different approaches—hybrid models—that mix rule-based filters, traditional ML, and deep learning to better capture the complexity of human language and hate speech.

III. METHODOLOGY

The hate speech classification methodology is a systematic approach that brings together data-driven natural language processing methods with moral considerations.

Presented below is a step-by-step detailed account of the steps used in developing an efficient text-based hate speech detector.

A. Data Collection and Annotation:

The initial step in the methodology is to collect diverse and pertinent text data. This encompasses:

- **Social Media Sites:** Text posts or comments from sites like Twitter, Facebook, Reddit, or YouTube.
- **Online Forums:** Specialized sites that can host ideologically motivated discussions (e.g., extremist or conspiracy forums).
- **Public Datasets:** Datasets found on repositories like Kaggle, GitHub, or research libraries (e.g., Hatebase, Davidson's Twitter dataset).

Annotation Strategy:

- Text samples are categorized into types such as: Hate Speech, Offensive Language, and Neutral.
- Multiple annotators are employed to ensure consistency, and inter-annotator agreement (e.g., Cohen's Kappa) is calculated.
- Clear annotator guidelines are established to help distinguish hate speech from general profanity or sarcasm.

B. Data Preprocessing

Preprocessing is necessary to remove raw text contaminants and prepare data for machine learning models.

- **Cleaning of Text:**
 - Lowercasing text
 - Elimination of punctuation, HTML tags, emojis, and URLs
 - Removal of extra spaces and non-alphanumeric characters
- **Tokenization:**

Splitting sentences into words or subwords using tools like SpaCy, NLTK, or Transformers' Tokenizers.

- **Stemming and Lemmatization:**

Reducing words to their root forms (e.g., “arguing” → “argue”, “better” → “good”) to consolidate lexical variations.

- **Handling Imbalanced Data:**

- Oversampling: Copy or synthesize new examples for the minority class.
- Undersampling: Eliminates samples from the majority class.
- Data Augmentation: Utilize paraphrasing, synonym substitution, or back translation.

C. Feature Extraction

The next step is to represent text in numerical forms that can be understood by machine learning algorithms.

- Classical Techniques:
 - Bag of Words (BoW): Simply counts the frequency of words, irrespective of context.
 - TF-IDF: Assigns weights to words depending on how significant they are in the corpus.
- Word Embeddings: *Word2Vec*, *GloVe*, *FastText*: Pre-trained vectors that capture semantic word similarity.
- Contextual Embeddings: Models such as *BERT*, *RoBERTa*, and *DistilBERT* produce embeddings that take the context of word meaning into account, allowing detection of subtle hate speech.

D. Model Choice and Training

The transformed features are input into different models for classification. Common models include:

- Traditional Machine Learning Models:
 - Logistic Regression
 - Support Vector Machines (SVM)
 - Naive Bayes
- Ideal for small datasets and rapid baseline testing.
- Deep Learning Models:
 - Convolutional Neural Networks (CNNs): Extract spatial patterns (n-grams).
 - Recurrent Neural Networks (RNNs), *LSTMs*, *GRUs*: Capture sequence and context information.
- Transformer-Based Models:
 - *BERT*, *RoBERTa*, *XLNet*, *GPT*: Fine-tuned on hate speech datasets for state-of-the-art results.
- Hybrid and Ensemble Approaches:
 - Combine rule-based filters (e.g., slur detection) with machine learning classifiers.
 - Use ensemble methods like *Random Forest* + *BERT* for enhanced accuracy and robustness.

E. Evaluation Metrics

The models are evaluated based on their ability to accurately detect hate speech while minimizing mis-classification.

- Accuracy: Overall correctness (less useful in imbalanced datasets).
- Precision: How many predicted hate speech cases were actually hate speech.
- Recall: How many actual hate speech cases were detected.
- F1-Score: Harmonic average of precision and recall.
- Confusion Matrix: Graphical representation of true/false positives and negatives.
- ROC-AUC: Evaluates classifier performance across decision thresholds.

F. Fairness and Bias Analysis

Ethical considerations are crucial. Detection systems must avoid perpetuating existing societal biases.

- Bias Testing: Evaluate model performance across subgroups based on race, gender, religion, etc.
- Fairness Metrics: Measure equity using statistical tests such as equalized odds, demographic parity, and disparate impact.
- Bias Mitigation: Apply methods such as adversarial debiasing, bias-aware loss functions, or data balancing.

G. Deployment Considerations

For real-world application, models must be scalable and continually monitored.

- Real-time pipelines can be deployed using APIs or cloud infrastructure (e.g., AWS, GCP, Azure).
- Continuous monitoring is required to detect data drift and ensure performance does not degrade.

- Incorporate feedback loops from users to improve model accuracy and fairness over time.

IV. IMPLEMENTATION DETAILS

A. Data Collection and Annotation:

Collect text data from social media platforms such as Twitter, Facebook, online forums, or from publicly available hate speech datasets. Examples include various Kaggle datasets that provide labeled data in multiple languages. If constructing a custom dataset, annotate each text instance into predefined categories such as *hate speech*, *offensive language*, and *neutral*.

Annotation should ideally be performed by multiple annotators to ensure consistency and minimize bias. The dataset should be cleaned to remove duplicates, bot-generated texts, and noisy data that could degrade model performance.

B. Data Preprocessing:

- Text Cleaning: Remove special characters, emojis, URLs, and convert all text to lowercase. Optionally, remove stop words that do not contribute to classification.
- Tokenization: Split the text into individual words or subwords using libraries such as SpaCy or NLTK.
- Stemming and Lemmatization: Convert words to their base or root form (e.g., “running” → “run”) to reduce redundancy.
- Handling Imbalance:
 - *Oversampling*: Increase the number of hate speech examples by duplicating or synthesizing new samples.
 - *Undersampling*: Reduce the number of non-hate speech instances.
 - *Data Augmentation*: Use techniques such as synonym replacement, back translation, or paraphrasing to expand the hate speech class.

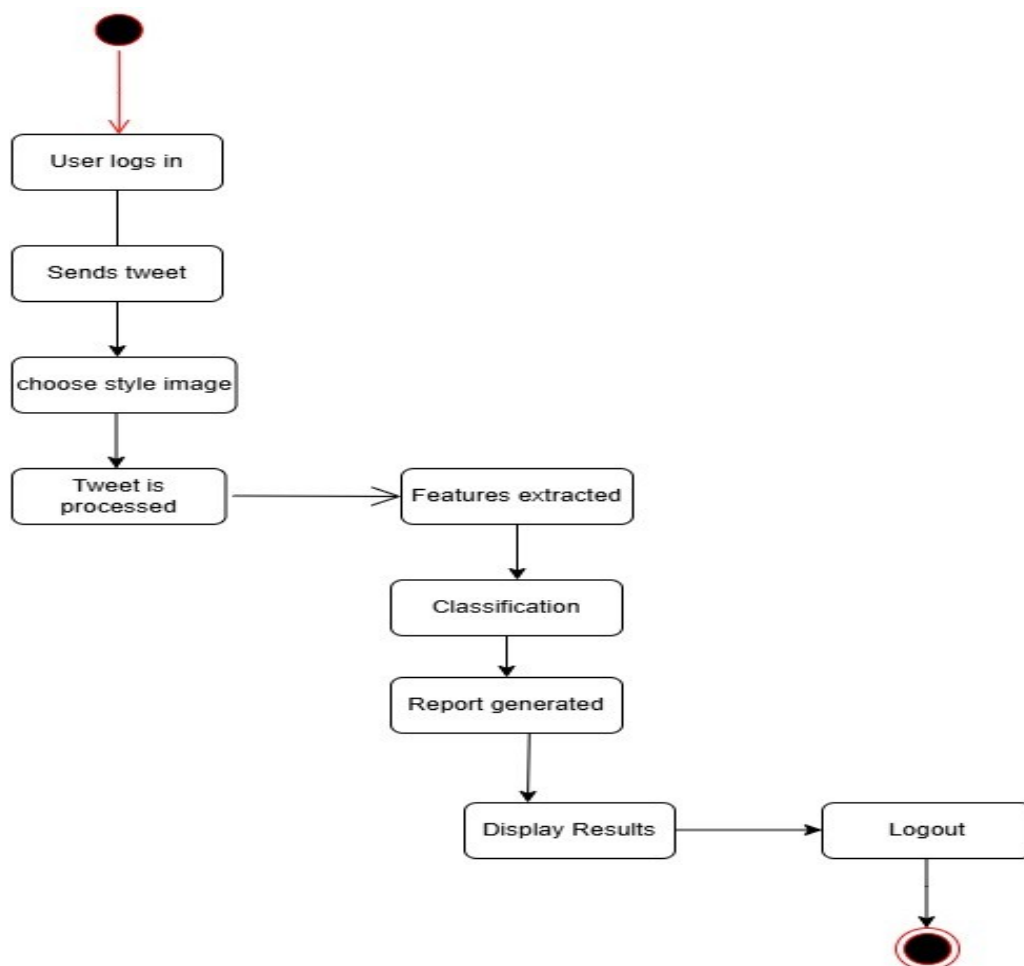


Figure: Activity Diagram

C. FeatureExtraction:

- BagofWords(BoW):Representstextasamatrixofwordoccurrencecounts.
- TF-IDF (Term Frequency-Inverse Document Frequency):A refinement of BoW that gives higher weight to less frequent but more informative words.
- WordEmbeddings:Usepre-trainedwordvectorslikeWord2Vec, GloVe, orFastText tocapturese- mantic relationships between words.
- Transformer-basedEmbeddings: LeveragemodelssuchasBERT,Roberta,orDistilBERTthatgen- erate context-aware word embeddings.

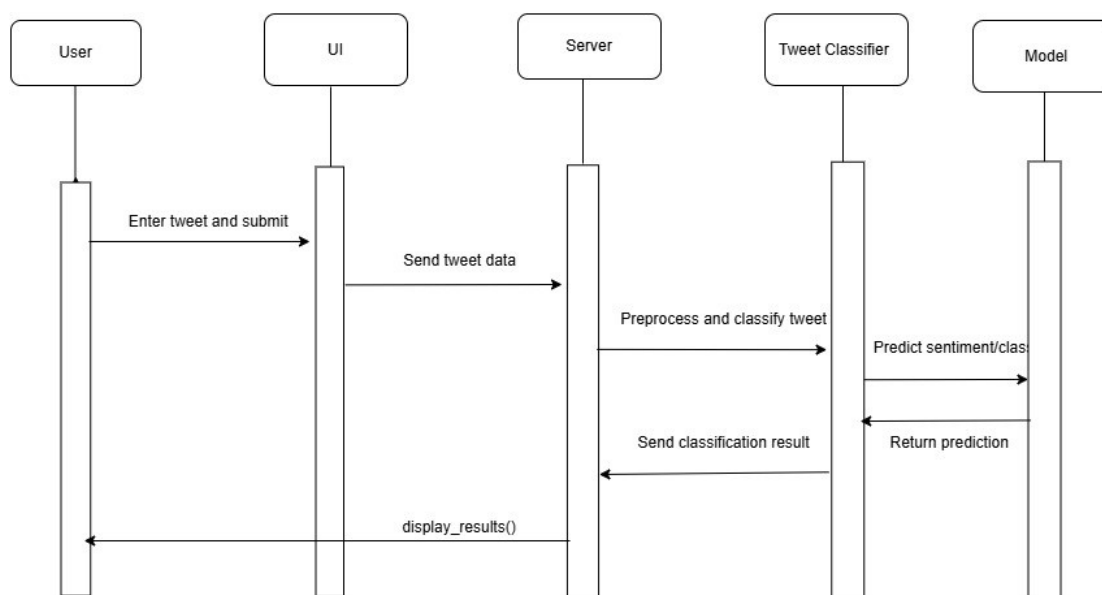


Figure:SequenceDiagram

D. EvaluationMetrics:

- Precision: Proportionofpredictedhatespeechsamplesatareactuallyhatespeech.
- Recall: Proportionofactualhatespeechsamplescorrectlyidentifiedbythemodel.
- F1Score: Harmonicmeanofprecisionandrecall,especiallyvaluableinclass-imbalanceddatasets.
- ConfusionMatrix: Visualizes true positives, false positives, true negatives, and false negatives, allowing for a betterunderstanding of model errors.

E. HandlingBiasandFairness:

- BiasDetection: Evaluate model performance across diverse demographic subgroups to ensure consistent accuracy and recall across races, genders, and other protected categories.
- FairnessMetrics: Applystatisticalfairnesstestssuchasequalizedoddsordisparateimpacttoassessandmitigate potential biases in the model.

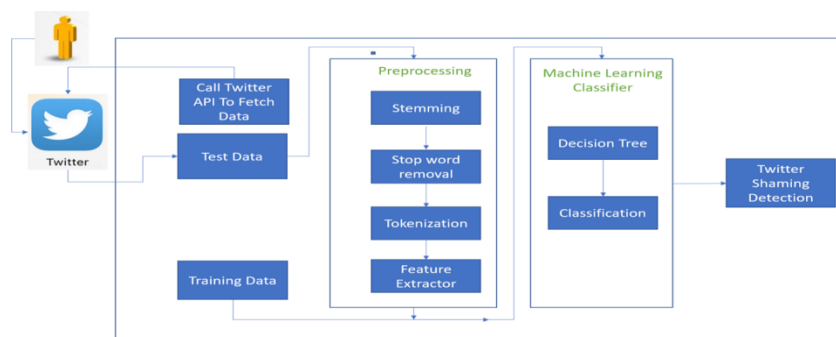


Figure: System Architecture

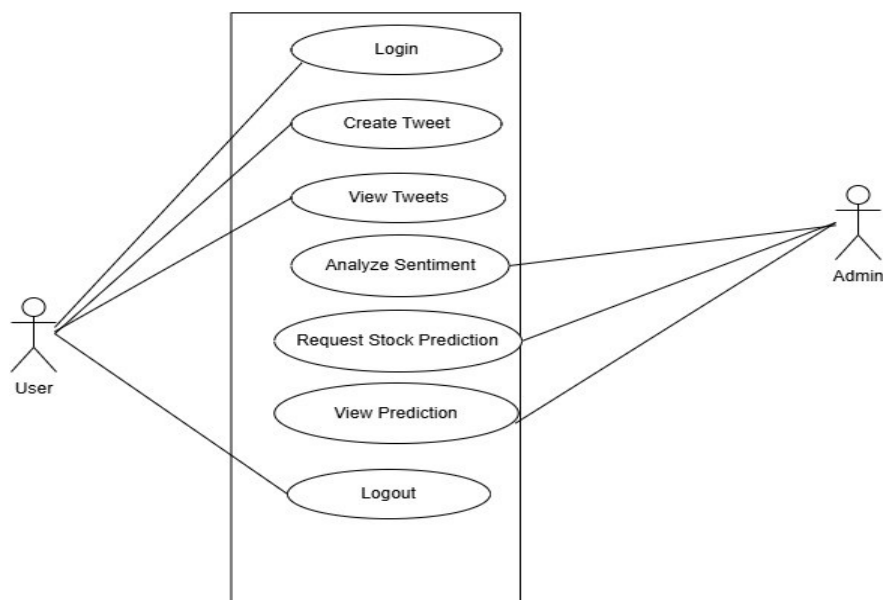


Figure:UseCaseDiagram

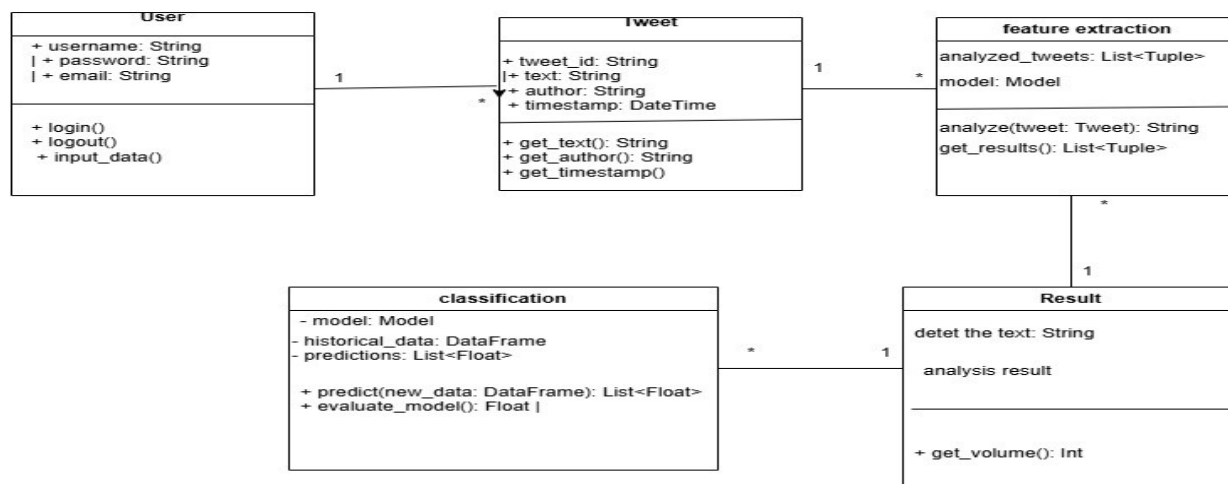


Figure:ClassDiagram




Figure:DFDLevel0

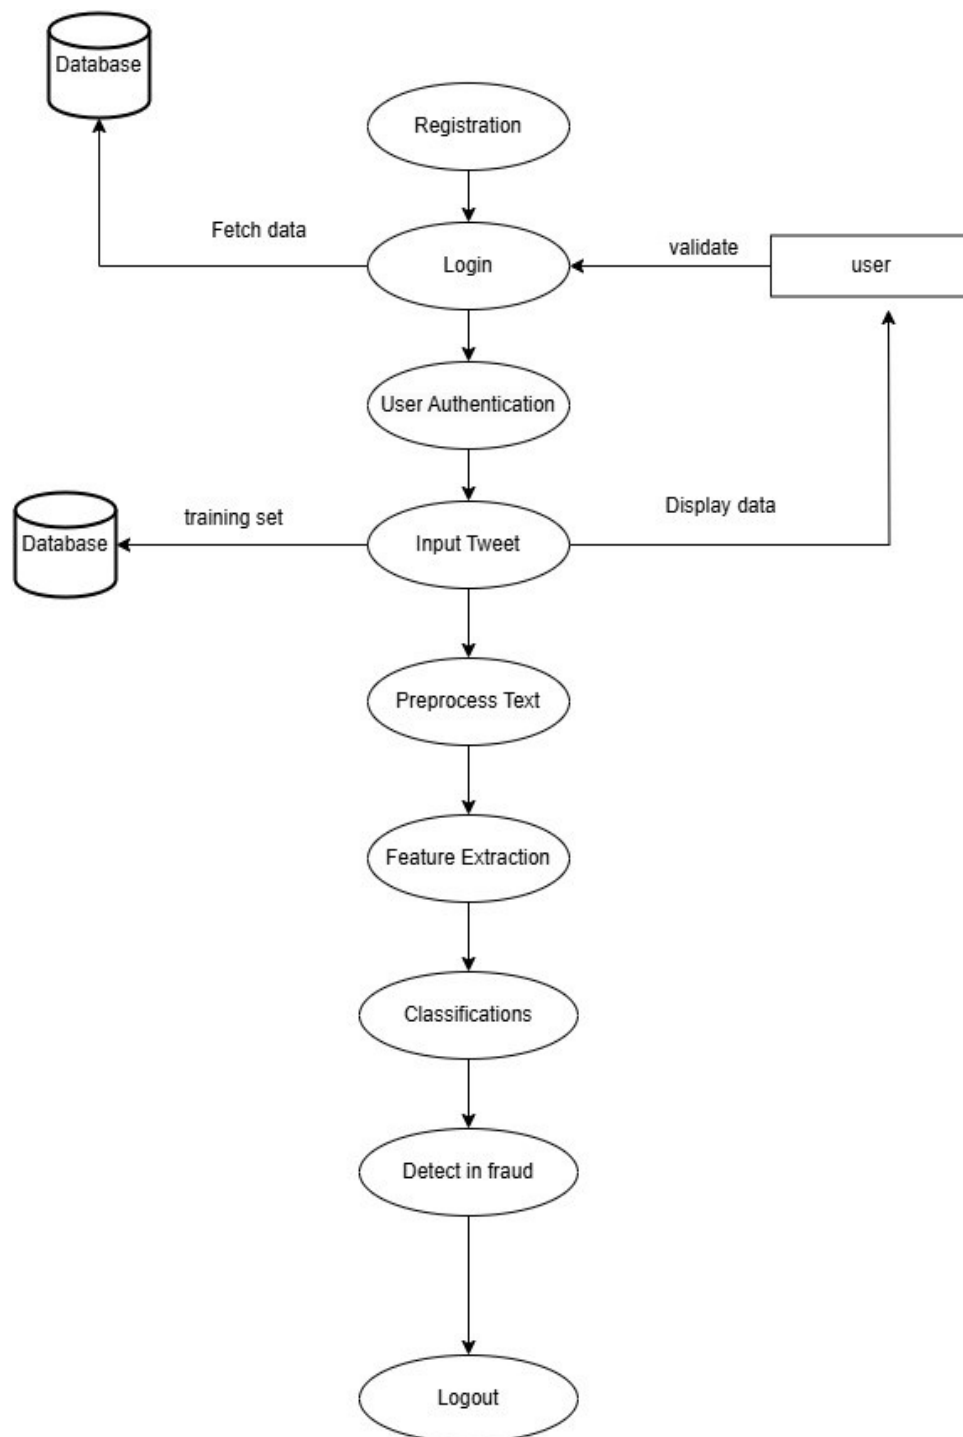


Figure:DFDLevel1

V. RESULTS & DISCUSSION

Detection of hate speech in Tex was overlooked in earlier technology as there was no survey of automatic detection. In the White Supremacy Forum, there are a lot more sentences which are not used for hate speech compared to 'hateful' sentences. There is a good chance that the boost in the F1-score in the two datasets was moderated by the single feature (count of) 'Followers', which also boosted the subset improvement. These patterns and unigrams may be applied as pre-compiled dictionaries not part of the proposed hate speech detection dictionaries as pre-existing dictionaries to be used in future research projects.

VI. CONCLUSION

The conclusion emphasizes the success of this project in overcoming such challenges with a multi-class classification approach. The success factor was the development and utilization of ten separate binary datasets, each dealing with a particular type of hate speech. Rather than grouping everything together, this fine-grained approach allowed models to concentrate on distinctive features of each hate category. Each dataset was annotated with great care by domain experts strictly adhering to guidelines, so labeling was highly consistent and accurate. Taking such care enhanced training and testing of the models, resulting in improved generalization and practical application. Also, the datasets were balanced, which in machine learning is important to avoid bias towards the majority class. Hate speech has been underrepresented in most of the existing datasets, and such bias can occur in model outputs. Equal representation here allowed classifiers to perform better on all the classes.



VII. ACKNOWLEDGEMENTS

This paper would not have been written without the support and encouragement of Asst. Prof. Dr. D. B. Hanchate, guide of BE Dissertation work. Author's special thanks go to all the professors of Computer Engineering Department of DGOI FOE Swami-Chincholi, for their guidance and for giving her an opportunity to work on Hate Speech Classification on Text.

REFERENCES

- [1] HateSpeechExplained: A Toolkit, vol. 19, London, U.K., 2015.
- [2] K. Saha, E. Chandrasekharan, and M. De Choudhury, "Prevalence and psychological effects of hateful speech in online college communities," in Proc. 10th ACM Conf. WebSci., Jun. 2019, pp. 255–264.
- [3] M. Bilewicz and W. Soral, "Hate speech epidemic: The dynamic effects of derogatory language on intergroup relations and political radicalization," Political Psychol., vol. 41, no. S1, pp. , Aug. 2020.
- [4] E. Blout and P. Burkart, "White supremacist terrorism in Charlottesville: Reconstructing unite the Right," Stud. Conflict Terrorism, pp. 1–22, Jan. 2021.
- [5] R. McIlroy-Young and A. Anderson, "From 'welcome new gabbers' to the Pittsburgh synagogue shooting: The evolution of gab," in Proc. Int. AAAI Conf. Web Social Media, vol. 13, 2019, pp. 651–654.
- [6] A. Warofka, "An independent assessment of the human rights impact of Facebook in Myanmar," Facebook Newsroom, vol. 5, Nov. 2018.
- [7] T. H. Paing, "Zuckerberg urged to take genuine steps to stop use of FB to spread hate in Myanmar," Irrawaddy.
- [8] Z. Waseem and D. Hovy, "Hateful symbols or hateful people? Predictive features for hate speech detection on Twitter," in Proc. NAACL Student Res. Workshop, 2016, pp. 88–93.
- [9] T. Davidson, D. Warmusley, M. Macy, and I. Weber, "Automated hate speech detection and the problem of offensive language," Proc. Int. AAAI Conf. Web Social Media, vol. 11, no. 1, pp. 512–515, May 2017.
- [10] A. Schmidt and M. Wiegand, "A survey on hate speech detection using natural language processing," in Proc. 5th Int. Workshop Natural Lang. Process. Social Media, 2017, pp. 1–10.
- [11] P. Fortuna and S. Nunes, "A survey on automatic detection of hate speech in text," ACM Comput. Surv., vol. 51, no. 4, pp. 1–30, 2018.
- [12] V. Basile, C. Bosco, E. Fersini, D. Nozza, V. Patti, F. M. R. Pardo, P. Rosso, and M. Sanguinetti, "SemEval-2019 task 5: Multilingual detection.
- [13] hate speech against immigrants and women in Twitter," in Proc. 13th Int. Workshop Semantic Eval. Vancouver, BC, Canada: Association for Computational Linguistics, 2019, pp. 54–63.

Author Bibliography

	<p>Ajay Suresh Bhagat Birth Place: Bhavaninagar (2002). Diploma in Computer Technology from S. B. Patil Polytechnic, Indapur. Pursuing B.E. in Computer Technology from DGOI FOE. Certified in Robotic Process Automation from Pro Azure Technology, Pune.</p>
	<p>Snehal Jagdish Bhalerao Birth Place: Bhiwandi (2001). Diploma in Computer Technology from Shivajirao S. Jondhale Polytechnic, Asangaon. Pursuing B.E. in Computer Technology from DGOI FOE. Certified in Robotic Process Automation from Pro Azure Technology, Pune.</p>

	<p>Radhika Suryakant Jagtap BirthPlace: Barshi (2004). Diploma from Government Polytechnic, Pune. Pursuing B.E. in Computer Technology from DGOI FOE. Certified in Robotic Process Automation from Pro Azure Technology, Pune.</p>
	<p>Siddhi Vijay More BirthPlace: Bhiwandi, Thane (2002). Diploma in Information Technology from K.K. Wagh Polytechnic, Nashik. Pursuing B.E. in Computer Technology from DGOI FOE. Completed internship in RPA at Pro Azure Company. Certifications in Java, Advanced Java, and Cyber Security.</p>
	<p>Dr. Dinesh Bhagwan Hanchate Birth- Solapur (MH- India), B.E. Comp. (Walchand College of Engg., Sangli, MH, (India)), M. Tech. Comp. (Dr. Babasaheb Ambedkar Technological University, Lonere, MH (India)). Ph.D. Comp. (SGGSIT, Nanded and SRTMU, Nanded, MH (India)). Former HOD of Comp. and IT. Head PG section, PG teacher. Did STTP, QIP programs sponsored by IIT, Kanpur, AICTE, ISTE, SPPU and UG. Positions: PhD Guide, Dean Industry Institute Interaction Cell. R & D coordinator, Head Media Cell, Editor Art Magazine, Students Development Officer, System Officer, Pune Division Head, Career Katta, GoM DTE, Author, Poet, Currently Professor in Dattakala Group of Institutions, Daund. Interest in ML, Software Engineering, AI, IR, Math Modelling, Usability Engg., Optimization, Soft Computing.. Email: dineshbhanchate@gmail.com</p>
	<p>Swati Mahadev Atole Birthplace: Korti (Karmala) Education: BE (Computer science and Engineering) from Karmaveer Bhaurao Patil College of Engineering, Satara (2013) Currently pursuing M.E. in Computer Technology from DGOI FOE. Currently working as Assistant Professor (Computer & IT) dept of Dattakala Group of Institutions Faculty of Engineering, Bhigwan. Email: swati.atole03@gmail.com</p>
	<p>Dr. Sachin Sukhadeo Bere Currently working as Associate Professor in Dattakala Group of Institutions Faculty of Engineering Bhigwan. Completed M.Tech (CSE) with First class & Distinction from JNTU Hyderabad affiliated college. Completed Ph.D. in Computer Science and Engineering, from the Shri Jagdishprasad Jhabarmal Tibrewala University, Rajasthan India. 14 years of Teaching experience and 7 years of Research experience. Maharashtra. Currently working as Associate Professor (Computer & IT) dept of Dattakala Group of Institutions Faculty of Engineering, Bhigwan. Email: sachinbere@gmail.com</p>



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)