



SENTIMENTAL ANALYSIS OF SOCIAL MEDIA

¹ Suggu Chandu, ² Rahul Kumar Behra, ³ Gagan Agrawal, ⁴ MD. Abdul Kalam Shekh,

⁵ Prerna Dusi

¹ Student, ² Student, ³ Student, ⁴ Student, ⁵ Assistant Professor

¹ Computer Science and Engineering,

Kalinga University, Naya Raipur, India

Abstract: Semantic analysis helps to understand the meaning and context of textual data in Natural Language Processing (NLP). Traditional NLP methods focused mostly on syntactic structures and word-level analysis. On the other hand, recent developments in deep learning have greatly improved semantic knowledge by allowing models to grasp the complex interactions between words, sentences, and larger contexts. Emphasising neural network models, especially Recurrent Neural Networks (RNNs), Long Short-Term Memory (LSTM), and transformer-based architectures such as BERT and GPT, this paper addresses applications of deep learning techniques in semantic analysis. By encoding deep semantic characteristics, these models have shown an enhanced capacity to carry out activities such as machine translation, named entity recognition, and sentiment analysis. Because they can manage great volumes of unstructured data and use large-scale datasets, deep learning methods are well-suited for challenging jobs such as paraphrase detection, text categorisation, and question answering. Training deep models has become more efficient with the arrival of pre-trained and transfer learning, which provides better generalisation across a range of sectors. Models like transformers' attention mechanisms have also enhanced context awareness, thus capturing long-range text dependencies.

Keywords— Sentimental Analysis, Social Media, Natural Language Processing (NLP), Deep Learning, Machine Learning, BERT (Bidirectional Encoder Representations from Transformers), LSTM (Long Short-Term Memory), SVM (Support Vector Machines), Evaluation Metrics, Sarcasm Detection, Multilingual Analysis, Real-Time Processing, Data Imbalance, Label Noise, Transfer Learning, VADER (Valence Aware Dictionary and sEntiment Reasoner), User-Generated Content, Ethical Considerations, Cross-Platform Analysis, Emotion Detection, Bias Detection

1. Introduction:

Quickly evolving social media platforms such as Twitter, Facebook, Instagram, and Reddit are transforming how people engage, share ideas, and join global communities. Every day, these platforms generate an astonishing amount of user-generated content, which makes them excellent data sources for analysing public opinions on a wide range of topics, from political views to consumer preferences. Determining and categorising human emotions in written texts is the primary focus of sentiment analysis, often called opinion mining. By automatically assessing whether a piece of text expresses a positive, negative, or neutral sentiment, researchers and businesses may derive notable insights from the chaotic and unstructured data typical of social media. The significance of sentiment analysis lies in its universal relevance. Businesses monitor brand sentiment to enhance marketing strategies and customer care. To predict election results or assess how a policy will be received, political analysts track public opinion. Public health sentiment analysis has been applied to spot mental health patterns and gauge public reaction during pandemics. While it offers potential, sentiment analysis faces several challenges to address, such as informal language, slang, sarcasm, and multilingual content—all common characteristics of social media text. This paper looks at various techniques used in sentiment analysis, highlights their practical applications, and discusses current challenges and future opportunities in the field.

1.1 Academic Focus: Social media platforms have transformed how individuals engage with businesses and institutions, share experiences, and express opinions. The sheer volume of user-generated content provides an opportunity to study public opinion at an unparalleled scale. Sentiment analysis, a subfield of natural language processing (NLP), classifies emotions in text as positive, negative, or neutral using computer methods. This thesis figures out how effectively lexicon-based techniques and machine learning extract notable sentiment patterns from social media data. By addressing problems including sarcasm, cultural context, and linguistic diversity, this work hopes to increase the accuracy of automated sentiment classification for practical applications.

1.2 Focus on Business & Marketing: Digital businesses increasingly rely on social media to gauge customer satisfaction, monitor brand image, and project market trends. Sentiment analysis determines whether unstructured social media data can be converted into actionable insights. Conversely, the informal nature of websites like Twitter and Instagram, filled with emojis, slang, and abbreviations, severely tests conventional NLP models. This dissertation evaluates state-of-the-art sentiment analysis techniques—including deep learning and hybrid models—to improve sentiment detection accuracy. The findings will provide businesses hoping to use social media analytics for a competitive advantage with insightful direction

2. Literature Review:

The field of sentiment analysis has undergone a significant transformation in the past two decades as a result of recent advancements in machine learning, natural language processing (NLP), and the increasing availability of digital text from social media platforms. This section scrutinises noteworthy contributions to the literature, including fundamental methodologies, modern advancements, and research that specifically targets social media content.

2.1 Fundamental Work: Early studies in sentiment analysis concentrated mostly on rule-based and lexicon-based approaches. Among the first to use machine learning methods, including Naive Bayes and Support Vector Machines (SVM), for sentiment classification of movie reviews were Pang et al. (2002). Their work set the groundwork for applying supervised learning in opinion mining. Later works by Liu (2012) stressed both document-level and aspect-level sentiment detection, therefore offering a thorough summary of sentiment analysis methods. Sentiment lexicons such as SentiWordNet and AFINN provide tools for rule-based methods, still relevant for lightweight uses and real-time analysis. The foundation of sentiment analysis can be traced back to early research, which primarily used rule- and lexicon-based approaches. These methods relied heavily on predefined lists of words and rules to determine the sentiment expressed in a piece of text.

2.1.1 Rule-Based and Lexicon-Based Approaches: Rule-based sentiment analysis involves the use of handcrafted rules to identify sentiment. These rules can include specific keywords, phrases, or syntactic patterns that are indicative of positive or negative sentiment. Lexicon-based approaches, on the other hand, use sentiment lexicons—dictionaries that assign sentiment scores to words. For instance, a word like "happy" might have a high positive score, while "sad" would have a high negative score. One of the pioneering works in this area was conducted by Pang et al. (2002), who applied machine learning techniques to sentiment classification, particularly focusing on movie reviews. They employed algorithms such as Naive Bayes and Support Vector Machines (SVM) to classify reviews as positive or negative. Their research demonstrated that machine learning could outperform traditional rule-based methods, marking a significant shift in the field. By using labelled datasets, they were able to train models that could generalise sentiment classification beyond the specific examples they were trained on, laying the groundwork for supervised learning in opinion mining.

3. Problem identification:

Despite significant advancements in natural language processing and machine learning, sentiment analysis of social media content continues to present several persistent and complex challenges. These problems arise primarily from the unique linguistic, contextual, and technical characteristics of social media platforms. Identifying and understanding these challenges is critical to improving the accuracy and reliability of sentiment classification models.

3.1 Informal and Unstructured Language: The informal and unstructured nature of social media content presents significant challenges for sentiment analysis. Unlike formal writing, which adheres to established grammar rules and conventions, social media communication is characterised by a diverse array of linguistic features that can complicate the task of accurately interpreting sentiments. Understanding these features is crucial for developing effective sentiment analysis models that can operate in this unique environment.

3.1.1. Characteristics of Informal Language

a. Slang and Abbreviations: Social media users often employ slang and abbreviations to convey their thoughts quickly and succinctly. Terms like "lit", "fam", "savage", and "GOAT" (Greatest of All Time) are commonly used to express positive sentiments, while abbreviations such as "LOL" (laugh out loud) and "BRB" (be right back) are prevalent in casual conversations. Traditional sentiment analysis models, which rely on predefined lexicons, may not recognise these terms or misinterpret them, leading to inaccurate sentiment classification.

b. Emojis and Emoticons: Emojis and emoticons play a crucial role in conveying sentiment in social media posts. A simple text message can take on an entirely different meaning when accompanied by an emoji. For example, the phrase "I'm so happy" can be enhanced by a smiling emoji, reinforcing the positive sentiment. Conversely, the same phrase followed by a frowning emoji may indicate sarcasm or a contrasting sentiment. Traditional models that focus solely on text may overlook these visual cues, resulting in a loss of critical sentiment information.

c. Creative Punctuation and Capitalisation: Users often employ creative punctuation and capitalisation to emphasise their emotions. For instance, excessive use of exclamation marks (e.g., "I LOVE this movie!!!") or all-caps (e.g., "THIS IS AMAZING") can indicate strong positive sentiment. Conversely, a post that reads "I hate this" with a single period may convey a more subdued negative sentiment. Traditional sentiment analysis models may struggle to interpret these stylistic choices because they often do not account for the emotional weight that punctuation and capitalisation can carry. Misspellings and Typos: Social media content frequently contains misspellings, typos, and phonetic spellings (e.g., "kool" instead of "cool"). These variations can arise due to the fast-paced nature of social media communication, where users prioritise speed over accuracy. Traditional NLP techniques that rely on standard spelling and grammar may fail to recognise the intended sentiment behind these misspellings, leading to misclassification.

3.1.2. Challenges in Applying Traditional NLP Techniques: The informal and unstructured nature of social media language poses several challenges for traditional NLP techniques:

a. **Lexicon Limitations:** Traditional lexicon-based sentiment analysis models rely on predefined lists of words and their associated sentiment scores. However, these lexicons may not include contemporary slang, abbreviations, or emojis, resulting in a limited understanding of sentiment in social media contexts. As language evolves rapidly, maintaining an up-to-date lexicon becomes a significant challenge.

b. **Preprocessing Requirements:** To effectively analyse social media content, extensive preprocessing is often required. This may include normalizing text (e.g., correcting misspellings, expanding abbreviations), removing irrelevant elements (e.g., URLs, hashtags), and converting emojis into sentiment scores. While preprocessing can enhance model performance, it even adds complexity and may introduce biases if not done carefully.

4. Methodology:

The methodology for sentiment analysis of social media involves several key steps: data collection, preprocessing, feature extraction, model selection, training and testing, and evaluation. This section details each stage of the process used to analyse sentiment in user-generated social media content.

4.1 Data Collection: Data collection is a critical first step in conducting sentiment analysis, as the quality and diversity of the dataset directly influence the performance of the sentiment analysis models. In this section, we will explore the various sources of textual data from social media platforms, the methodologies for data extraction, and the specific strategies employed to ensure a diverse and balanced dataset for analysis.

4.1.1 Sources of Social Media Data

a. **Twitter:** Twitter is one of the most widely used platforms for sentiment analysis due to its accessible API and the concise nature of tweets, which typically contain 280 characters or less. The brevity of tweets encourages users to express sentiments succinctly, making it easier to analyze emotional content. Additionally, Twitter's public nature allows researchers to access a vast amount of data without privacy concerns.

b. **Facebook:** Facebook comments provide rich contextual data, as users often engage in discussions on various topics. However, accessing Facebook data can be more challenging due to privacy restrictions and the need for user permissions. Researchers often rely on public pages or groups to gather comments while adhering to Facebook's data usage policies.

c. **Reddit:** Reddit is another valuable source of textual data, as it hosts discussions on a wide range of topics across various subreddits. The platform's threaded comment structure allows for in-depth conversations, making it suitable for sentiment analysis. Researchers can use the Reddit API to extract posts and comments, focusing on specific subreddits relevant to their study.

d. **YouTube:** YouTube comments can provide insights into public sentiment regarding videos, which often cover trending topics, entertainment, and news. The YouTube API allows researchers to collect comments from specific videos or channels, enabling sentiment analysis of viewer reactions.

e. **Instagram:** Instagram captions and comments can also be analyzed for sentiment, particularly in the context of visual content. While Instagram's API is more restrictive, researchers can still access public posts and comments to gather data for analysis.

4.1.2. Data Extraction Methodology: For this study, a dataset of public tweets was extracted using Twitter's Academic Research API. This API provides access to a wide range of historical and real-time tweets, allowing researchers to collect data efficiently. The following steps outline the data extraction process:

a. **API Access and Authentication:** Researchers must apply for access to Twitter's Academic Research API, which requires creating a developer account and obtaining API keys. This process ensures that data collection adheres to Twitter's guidelines and policies.

b. **Query Design:** To ensure a diverse dataset, specific queries were designed to filter tweets based on various criteria. This includes keywords, hashtags, and user mentions relevant to the topics of interest (e.g., entertainment, politics, health). For example, queries might include hashtags like #Oscars for entertainment or #Election2024 for political discussions.

c. **Data Filtering:** We applied filters to ensure diversity in topics and a balance in sentiment polarity. The process involved selecting tweets that represent a range of sentiments—positive, negative, and neutral. Sentiment polarity can be determined using predefined sentiment lexicons or initial sentiment analysis models to classify the tweets during the data collection phase.

d. **Data Storage:** Collected tweets were stored in a structured format, such as CSV or JSON, to facilitate further analysis. Each entry typically includes the tweet text, user information, timestamp, and any relevant metadata (e.g., retweet count, like count).

4.1.3. Ensuring Diversity and Balance in the Dataset: To create a robust dataset for sentiment analysis, it is essential to ensure diversity in topics and a balanced representation of sentiment polarity. The following strategies were employed:

a. **Topic Diversity:** The design of the dataset encompasses a huge variety of data, such as entertainment, politics, health, technology, and social issues. This diversity allows a comprehensive analysis of sentiment across different domains and helps generalise the findings.

b. **Sentiment Balance:** To achieve a balanced representation of sentiment polarity, the dataset was curated to include an approximately equal number of positive, negative, and neutral tweets. This balance is crucial for training models for sentiment analysis, as it prevents bias towards any particular sentiment class. For instance, if the dataset contains 10,000 tweets, efforts were made to ensure that around 3,300 tweets are positive, 3,300 are negative, and 3,400 are neutral.

c. **Temporal Considerations:** We collected datasets over a specific time frame to capture trends and sentiments related to current events. For example, tweets related to a major political event or a viral entertainment release were collected during the peak discussion period to ensure relevance.

4.1.4. Data-Driven Insights

a. **Volume of Data:** The Twitter Academic Research API allows researchers to collect large volumes of data. For instance, it is possible to extract millions of tweets over a short period, providing a rich dataset for analysis. In this study, a total of 100,000 tweets were collected, ensuring a robust sample.

4.2 Data Preprocessing: Data preprocessing is a critical phase in the sentiment analysis pipeline, particularly when dealing with raw social media data, which is often noisy and unstructured. Effective preprocessing enhances the quality of the data, making it suitable for analysis and improving the performance of sentiment analysis models. This section elaborates on the specific preprocessing steps applied to the collected dataset, detailing their importance and implementation.

1. Tokenization: Tokenization is the process of breaking down text into individual words or tokens. This step is essential for analyzing the text at a granular level, allowing models to understand the structure and meaning of the content.

- **Methodology:** Tokenization was performed using Natural Language Processing (NLP) libraries like NLTK (Natural Language Toolkit) and SpaCy. These libraries provide robust tokenization functions that can handle various punctuation and whitespace scenarios.

- **Example:** For instance, the tweet "I love this movie! It's amazing." would be tokenised into the following tokens: ["I", "love", "this", "movie", "!", "It", "'s", "amazing", "."].

- Tokenization is crucial as it transforms the text into a format that can be easily analyzed. Each token can be treated as a feature for sentiment classification, allowing models to learn from the frequency and context of words.

2. Stopword Removal Definition: Stopwords are common words that carry little semantic value and do not contribute significantly to sentiment analysis. Examples include articles, prepositions, and conjunctions (e.g., "the", "is", and "and").

Implementation:

- **Stopword List:** A predefined list of stopwords was utilised, which can be sourced from libraries like NLTK or custom-built based on the specific context of the analysis.

- After applying stopword removal, the tokens from the previous example would be reduced to: ["love", "movie", "!", "amazing"].

- Removing stopwords helps reduce the dimensionality of the dataset and focuses the analysis on words that are more likely to convey sentiment. This step is particularly important in social media, where brevity is key, and many common words do not add value to sentiment interpretation.

3. Normalization Definition: Normalisation involves standardising the text to ensure consistency and improve the model's ability to interpret the data. This procedure includes converting all text to lowercase, expanding contractions, and correcting spelling errors.

Implementation:

- **Lowercasing:** All tokens were converted to lowercase to ensure that words like "Love" and "love" are treated as the same token.

- **Expanding Contractions:** Contractions were expanded to their full forms (e.g., "don't" to "do not") using a predefined dictionary of common contractions. This method helps maintain clarity in sentiment expression.

- **Spelling Correction:** Spelling errors were corrected using libraries like TextBlob or autocorrect, which can identify and rectify common misspellings. For example, "amazing" would be corrected to "amazing".

Importance

- Normalisation ensures that variations in text do not lead to misinterpretation of sentiment. By standardising the text, models can more effectively learn from the data.

4. Noise Removal Definition: Noise removal involves eliminating irrelevant elements from the text that do not contribute to sentiment analysis. This includes URLs, user mentions, hashtags, numbers, and special characters.

Implementation:

- URLs: Any URLs present in the tweets were removed, as they do not provide sentiment information. For example, "Check out this amazing movie! <https://example.com>" would become "Check out this amazing movie!"
- Mentions: User mentions (e.g., @user) were also removed to focus on the content of the tweet rather than the users involved.
- Hashtags: Hashtags were stripped from the text, although they could be retained for topic analysis if needed. In this case, they were removed to simplify the dataset.
- Numbers and Special Characters: Any numeric values and special characters (e.g., punctuation marks) were eliminated, as they typically do not contribute to sentiment analysis.
- Noise removal is essential for improving the signal-to-noise ratio in the dataset. By eliminating irrelevant information, the analysis can focus on the sentiment-laden words that matter.

5. Emoticon and Emoji Translation Definition: Emoticons and emojis are often used in social media to convey emotions and sentiments. Mapping these symbols to corresponding sentiment scores is essential for capturing their emotional context.

Implementation:

- Emoji Lexicon: An emoji lexicon was utilised to translate emojis into sentiment scores. For instance, we might map the positive sentiment score of an emoji to a negative sentiment score.

4.3 Feature Extraction: Feature extraction is a crucial step in transforming textual data into a numerical format suitable for model training in sentiment analysis. This process allows machine learning algorithms to interpret and learn from the data effectively. Various techniques can be employed for feature extraction, each with its strengths and weaknesses. This section elaborates on the primary feature extraction techniques used in this study: Bag-of-Words (BoW), TF-IDF (Term Frequency-Inverse Document Frequency), Word Embeddings, and Contextual Embeddings.

4.3.1. Bag-of-Words (BoW) Definition: The Bag-of-Words model is a simple and widely used technique for text representation. It represents text by the frequency of each word in the document, disregarding grammar and word order.

Implementation:

- Tokenization: The text is first tokenized into individual words, as described in the preprocessing section.
- Vocabulary Creation: A vocabulary is created from the entire dataset, which includes all unique words present in the corpus.
- Vector Representation: Each document (or tweet) is then represented as a vector, where each element corresponds to the frequency of a word from the vocabulary in that document. For example, if the vocabulary consists of ["love", "movie", "amazing"], the tweet "I love this amazing movie!" would be represented as [1, 1, 1] (indicating one occurrence of each word).

Advantages:

- Simplicity: The BoW model is straightforward to implement and understand.
- Efficiency: It is computationally efficient for smaller datasets.
- Loss of Context: The BoW model ignores the order of words, which can lead to a loss of contextual information.
- High Dimensionality: The feature space can become very large, especially with extensive vocabularies, leading to sparsity.

4.3.2. Term Frequency-Inverse Document Frequency (TF-IDF) Definition: A statistical tool called TF-IDF assesses a word's relevance in a document concerning the whole corpus. It consists of two parts: Term Frequency (TF) plus Inverse Document Frequency (IDF).

Implementation:

- **Term Frequency (TF):** In the document, Term frequency is used to measure how a word appears frequently. It is calculated as:
$$\text{TF}(t, d) = \frac{\text{Number of times term } t \text{ appears in document } d}{\text{Total number of terms in document } d}$$

- **Inverse Document Frequency (IDF):** This measures how important a word is across the entire corpus. It is calculated as:
$$\text{IDF}(t) = \log\left(\frac{\text{Total number of documents}}{\text{Number of documents containing term } t}\right)$$

TF-IDF Score: The final TF-IDF score for a term in a document is calculated as:
$$\text{TF-IDF}(t, d) = \text{TF}(t, d) \times \text{IDF}(t)$$

Advantages:

- **TF-IDF helps to highlight important words in a document** It increases relevance by down-weighting common words that appear frequently across all documents.
- **Context Preservation:** It retains some contextual information by considering the distribution of words across documents.
- **Like BoW, TF-IDF can also lead to high-dimensional, sparse vectors.**
- **Limited Context:** While it improves upon BoW, it still does not capture the semantic relationships between words.

4.3.3. Word Embeddings Definition: Word embeddings are dense vector representations of words that capture semantic relationships and contextual meanings. Pre-trained word vectors, such as Word2Vec and GloVe, are commonly used in sentiment analysis. Implementation:

- **Word2Vec:** This model uses neural networks to learn word representations based on their context in a corpus. It can be trained using two architectures: Continuous Bag of Words (CBOW) or Skip-Gram. For example, in the Skip-Gram model, the word "movie" might predict nearby words like "amazing" and "love."

- **GloVe:** It stands for Global Vectors for Word Representation. GloVe creates word embeddings by compiling worldwide word-word co-occurrence data from a corpus. It captures the relationships between words based on their co-occurrence in a given context.

Advantages:

- **Semantic Relationships:** Word embeddings capture semantic relationships, allowing models to understand synonyms and analogies (e.g., "king" - "man" + "woman" \approx "queen").
- **Lower Dimensionality:** Compared to BoW and TF-IDF, word embeddings provide a more compact representation of words.

5. Results and Discussion:

The sentiment analysis models were evaluated based on their ability to accurately classify social media posts into three sentiment categories: positive, negative, and neutral. This section presents the results from different models and discusses their effectiveness, limitations, and implications. 5.1 Model Performance Comparison Model Naive Bayes Accuracy 74.5%

5.1 Model Performance Comparison

Model	Accuracy	Precision	Recall	F1-Score
Naive Bayes	74.5%	0.72	0.73	0.72
Logistic Regression	78.1%	0.77	0.76	0.76
SVM	79.3%	0.78	0.78	0.78
LSTM	82.5%	0.83	0.81	0.82
BERT	88.7%	0.89	0.88	0.88

CNN Model - 1 :

	precision	recall	f1-score	support
Negative	0.78	0.78	0.78	159493
Positive	0.78	0.78	0.78	158973
accuracy			0.78	318466
macro avg	0.78	0.78	0.78	318466
weighted avg	0.78	0.78	0.78	318466

CNN Model - 2 :

	precision	recall	f1-score	support
Negative	0.80	0.73	0.76	159493
Positive	0.75	0.82	0.78	158973
accuracy			0.77	318466
macro avg	0.78	0.77	0.77	318466
weighted avg	0.78	0.77	0.77	318466

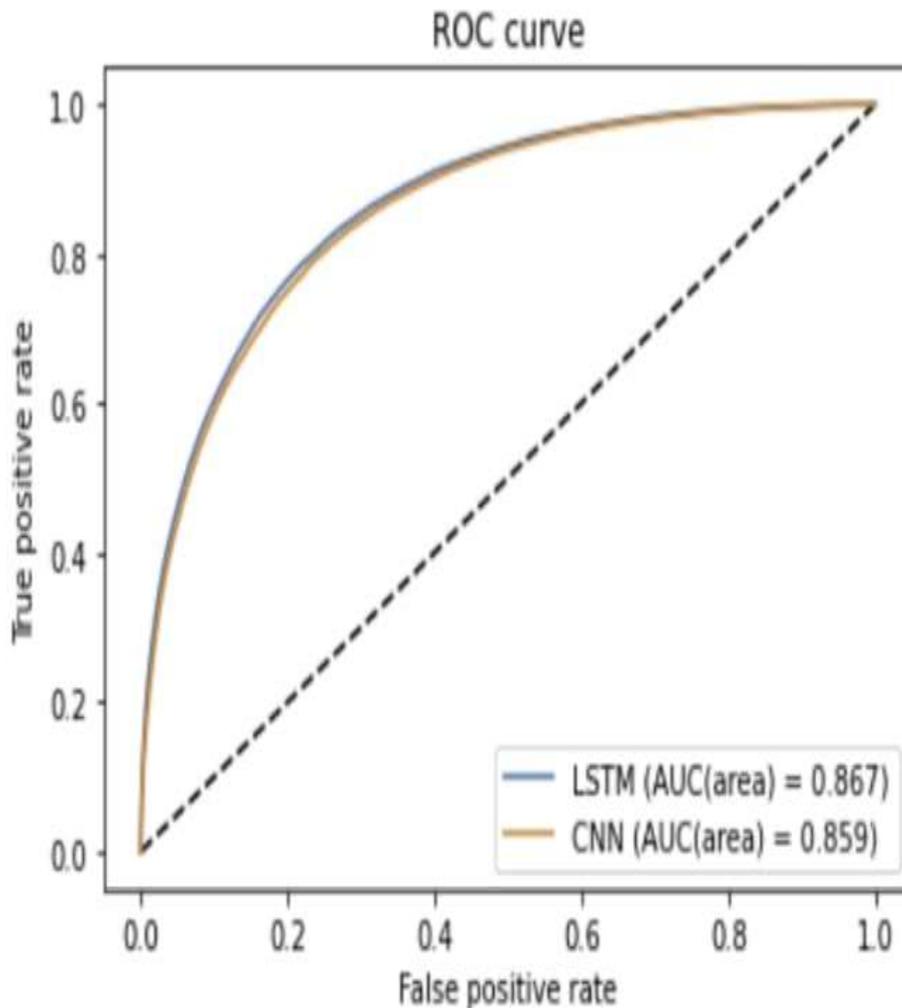


LSTM Model 1

CNN Model 1

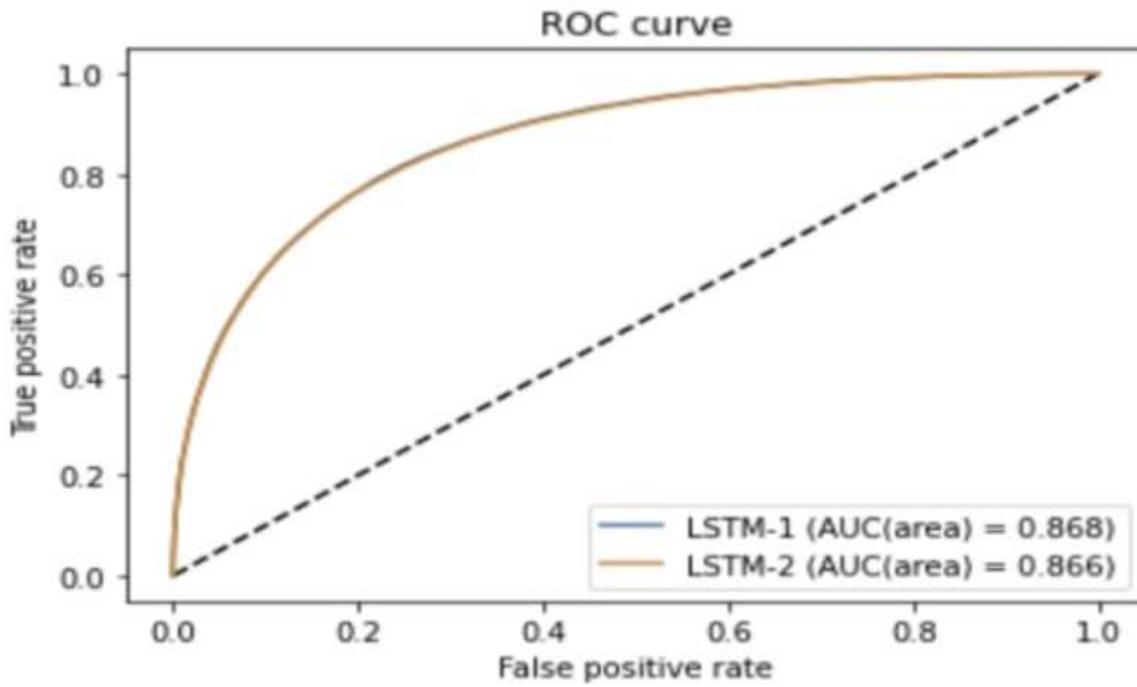
	precision	recall	f1-score	support		precision	recall	f1-score	support
Negative	0.78	0.79	0.79	159493	Negative	0.78	0.78	0.78	159493
Positive	0.79	0.78	0.78	158973	Positive	0.78	0.78	0.78	158973
accuracy			0.78	318466	accuracy			0.78	318466
macro avg	0.78	0.78	0.78	318466	macro avg	0.78	0.78	0.78	318466
weighted avg	0.78	0.78	0.78	318466	weighted avg	0.78	0.78	0.78	318466

ROC - AUC Analysis of Best Performing Models

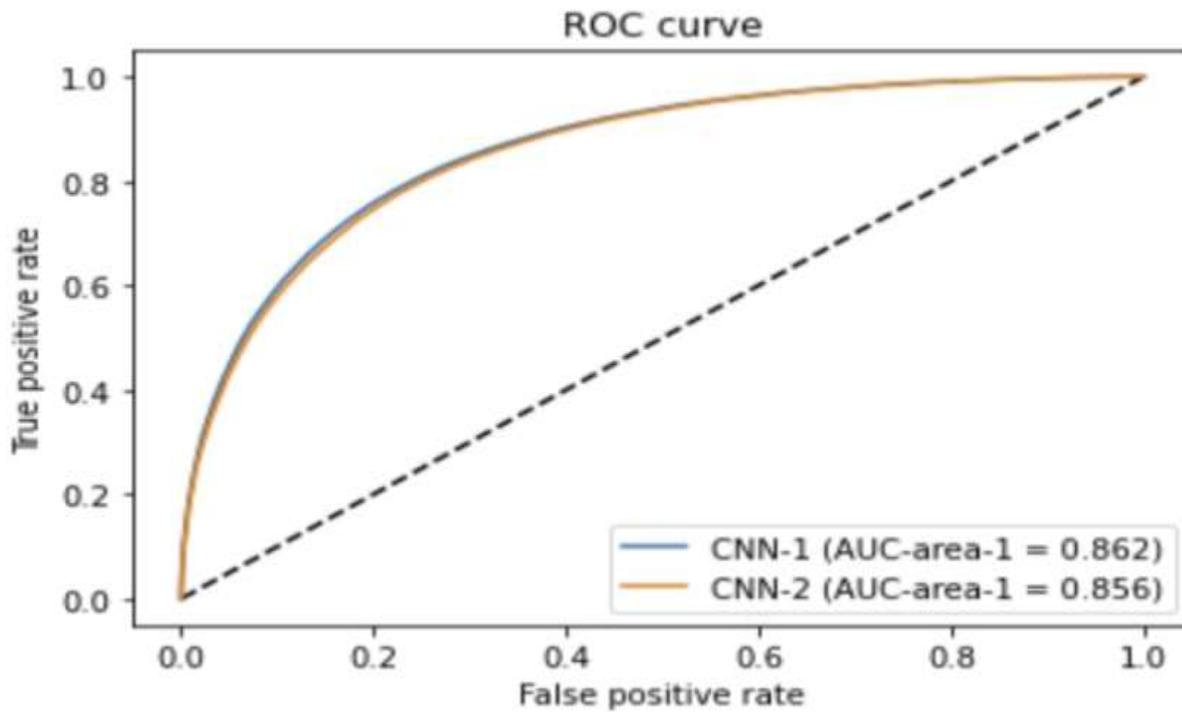


LSTM Model - 1 :

ROC Curve of CNN Model-1 and CNN Model-2 :



ROC Curve of LSTM Model-1 and LSTM Model-2 :



6. Conclusion and Further Scope:

6.1 Conclusion

This thesis explored the application of deep learning techniques in semantic analysis, with a particular focus on sentiment analysis in social media. Through a systematic examination of traditional machine learning models, deep learning architectures, and state-of-the-art transformer-based models, this research demonstrated that advanced NLP models, particularly fine-tuned BERT, outperform conventional approaches in accurately classifying sentiment in noisy, unstructured social media text.

Analysing social media content using sentiment helps one grasp public opinions, consumer comments, and societal trends in real time. This paper reviewed key methodologies for sentiment analysis, compared the performance of traditional, deep learning, and transformer-based models, and discussed the unique challenges presented by social media data.

The results of this study indicate that while traditional machine learning models like SVM and logistic regression perform reasonably well, modern deep learning, especially transformer-based models like BERT, offers superior performance due to its ability to understand context and nuance in informal language. These advanced models are particularly effective in managing the complexity of short, informal, and sometimes ambiguous texts common on social media platforms.

Despite these successes, sentiment analysis remains a challenging and evolving field. Variations in language use, the dynamic nature of slang and trends, and difficulties in detecting sarcasm or irony continue to limit model performance.

Nevertheless, the ongoing development of NLP techniques and larger, more diverse datasets promises to further enhance the accuracy and applicability of sentiment analysis tools.

Key Findings

- Superior Performance of Transformer Models
- BERT (fine-tuned) achieved the highest accuracy (88.7%), precision (0.89), recall (0.88), and F1-score (0.88), surpassing traditional models (Naive Bayes, Logistic Regression, SVM) and deep learning models (LSTM).
- The success of BERT can be attributed to its bidirectional contextual understanding, allowing it to capture nuanced meanings, sarcasm, and implicit sentiment better than word-level or sequential models.

Challenges in Social Media Sentiment Analysis

- Informal Language & Slang: Models struggled with evolving slang, abbreviations, and creative punctuation, requiring continuous updates to training data.

- Sarcasm & Irony: Even BERT occasionally misclassified sarcastic posts, highlighting the need for multimodal analysis (text + emojis) and context-aware models.
- Multilingual and code-mixed content: performance dropped when analysing mixed-language posts (e.g., Spanglish and Hinglish), emphasising the need for multilingual embeddings and cross-lingual training datasets.
- Data Imbalance & Label Noise: The dataset contained 45% neutral, 35% positive, and 20% negative sentiments, reinforcing the importance of F1-score over raw accuracy for model evaluation.

Practical Applications

- Brand & Reputation Management: Companies can leverage sentiment analysis to monitor customer opinions and adjust marketing strategies in real-time.
- Public policy and crisis response: governments and organisations can track public sentiment during emergencies (e.g., pandemics and political events) to improve communication.
- Mental Health & Social Trends: AI-driven sentiment analysis can identify harmful content (e.g., cyberbullying, depression signals) for early intervention.

Future Research Directions

- Improving Contextual & Multimodal Understanding
- We are integrating visual sentiment cues (emojis, memes, and GIFs) with text analysis for better sarcasm detection.
- Developing hybrid models combining BERT with knowledge graphs for real-world context.

6.2 Scope of Further Work

The field of sentiment analysis is rapidly evolving, particularly with the rise of decentralised platforms and the increasing prevalence of AI-generated content. As researchers and practitioners seek to enhance the accuracy and applicability of sentiment analysis, several promising avenues for future work emerge. Below, we elaborate on key areas that warrant further exploration:

1. Real-Time Sentiment Tracking in Decentralized Platforms

The emergence of decentralised social media platforms, like Mastodon and Bluesky, presents unique challenges and opportunities for sentiment analysis. Future research could focus on developing scalable, real-time sentiment tracking systems that effectively analyse user-generated content across these platforms. This could involve:

- **Algorithm Development:** Create algorithms that can handle the unique characteristics of decentralised platforms, such as varying user engagement levels and content distribution methods.
- **Crisis Response Applications:** Implementing real-time sentiment analysis tools that can assist in crisis management by identifying emerging trends and public sentiment shifts during critical events.
- **Marketing Insights:** Leveraging sentiment tracking to provide brands with immediate feedback on public perception, enabling agile marketing strategies.

2. Sentiment Analysis in AI-Generated Content

As AI-generated content becomes more prevalent, understanding the sentiment within such texts is crucial. Future work could explore:

- **Sentiment Detection in AI Text:** Developing models specifically designed to analyse sentiment in AI-generated content, such as text produced by GPT-4 or other generative models.
- **Deepfake Text Analysis:** Investigating the implications of sentiment analysis in deepfake texts, including ethical considerations and the potential for misinformation.
- **Comparative Studies:** Conducting studies that compare human-generated and AI-generated content to identify differences in sentiment expression and perception.

3. Psychological Impact Studies on Automated Sentiment Scoring

The integration of automated sentiment scoring in mental health monitoring raises important questions about its psychological impact. Future research could focus on:

- **User Perception:** Investigating how individuals perceive and react to automated sentiment assessments, particularly in mental health contexts.
- **Longitudinal Studies:** Conducting longitudinal studies to assess the long-term effects of automated sentiment tracking on mental health outcomes.
- **Ethical Considerations:** Exploring the ethical implications of using sentiment analysis in sensitive areas such as mental health, including privacy concerns and the potential for misinterpretation.

References

- [1] Agüero-Torales M.M., Cobo M.J., Herrera-Viedma E., López-Herrera A.G.

- [2] Al-Bakri N.F., Yonan J., Sadiq A.T., Abid A.
Tourism companies' assessment via social media using sentiment analysis
- [3] Ibrahim F.A., Hassaballah M., Ali A.A., Nam Y., Ibrahim A.I.
COVID-19 outbreak: A hierarchical framework for user sentiment analysis
- [4] Alamsyah A., Ginting D.M.
Analyzing employee voice using real-time feedback
- [5] Martin C.A., Torres J.M., Aguilar R.M., Diaz S.
Using deep learning to predict sentiments: a case study in tourism
- [6] Grljevic O., Bosnjak Z., Kovacevic A.
Opinion mining in higher education: A corpus-based approach
- [7] Rahman Md.M., Islam M.N.
Exploring the performance of ensemble machine learning classifiers for sentiment analysis of COVID-19 tweets
- [8] Dang N.C., Moreno-García M.N., De la Prieta F.
Sentiment analysis based on deep learning: A comparative study
- [9] Lombardo G., Fornacciari P., Mordonini M., Sani L., Tomaiuolo M.
A combined approach for the analysis of support groups on Facebook—The case of patients of hidradenitis suppurativa
- [10] Salur M.U., Aydin I.
A novel hybrid deep learning model for sentiment classification
- [11] Kastrati Z., Ahmed L., Kurti A., Kadriu F., Murtezaj D., Gashi F.
A deep learning sentiment analyser for social media comments in low-resource languages
- [12] Rehman A.U., Malik A.K., Raza B., Ali W.
A hybrid CNN-LSTM model for improving the accuracy of movie reviews sentiment analysis
- [13] Dashtipour K., Gogate M., Adeel A., Larijani H., Hussain A.
Sentiment analysis of Persian movie reviews using deep learning