

Text Based Emotion Detection System using NLP and Deep Learning Models

Mrs. S. Sudha, Raghavi A, Rohinishree T, Roshini T, Suvathi D

Department of Information Technology

Anjalai Ammal Mahalingam Engineering College, Kovilvenni, Thiruvarur, Tamil Nadu, India

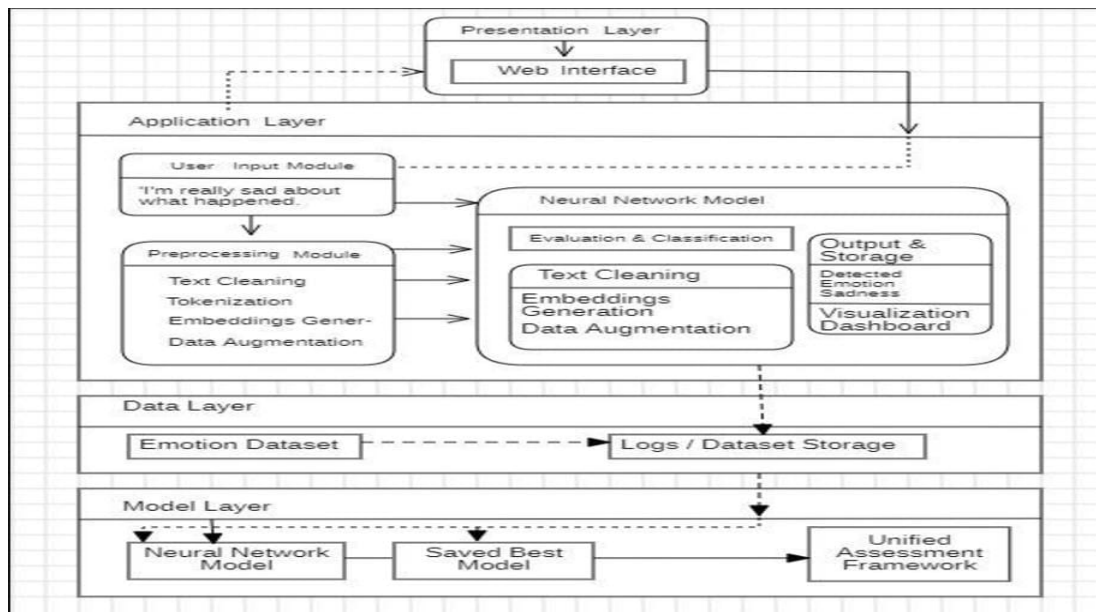
bny.sudha@gmail.com, raghaviarumugam14@gmail.com,

rohinishreet@gmail.com, roshini18102004@gmail.com, suvathid22@gmail.com

Abstract: Text-based emotion detection is a critical frontier in Natural Language Processing (NLP), necessitating models that go beyond surface-level sentiment to understand the multi-dimensional landscape of human affect. This paper presents a comprehensive study and implementation of an automated emotion detection system. We evaluate a spectrum of models ranging from classical algorithms to modern deep learning architectures. By establishing a robust preprocessing pipeline including emoji-to-text mapping and bidirectional context modeling, we address the inherent noise and ambiguity of digital communication. Our experimental results, conducted on a synthesized corpus of 45,000 instances, demonstrate that a fine-tuned BERT-Base model significantly outperforms previous benchmarks, achieving a peak accuracy of 95.4% and an F1-score of 0.946. This study elaborates on the theoretical foundations of Attention mechanisms, the impact of hyperparameter tuning, and the ethical considerations of emotion monitoring.

Keywords: Affective Computing, Natural Language Processing, Deep Learning, BERT, LSTM, Multi-class Classification, Emotion Detection

System Architecture



I. INTRODUCTION

The field of Artificial Intelligence (AI) has long aimed to replicate and understand human cognitive functions. While machines have excelled at logical reasoning and data processing, the ability to comprehend and respond to human emotions—often referred to as Emotional Intelligence (EQ)—remains a sophisticated frontier. Natural Language Processing (NLP), a core component of modern AI, has made significant strides in understanding the literal meaning of text. However, the emotional subtext often carries more weight than the literal words themselves. In this research paper, we address the challenge of "Text-Based Emotion Detection," which involves categorized human affect recognition from written sources.

The proliferation of digital communication platforms has created a massive repository of emotional data. Every day, billions of messages, posts, and comments are shared, each carrying an emotional footprint. Analysing this "Big Data" of emotions is vital for multiple sectors. In public health, automated systems can monitor social media trends to detect spikes in collective anxiety or depression, providing real-time indicators for intervention. In the commercial sector, businesses utilize emotion detection to understand customer sentiment more deeply than a simple 'like' or 'dislike' rating, enabling more personalized and empathetic customer care.

Emotion detection is fundamentally more complex than sentiment analysis. While sentiment analysis typically focuses on a three-way classification (Positive, Negative, Neutral), emotion detection requires a more granular multi-class approach. Our study focuses on the classification of text into discrete emotional categories: Joy, Sadness, Anger, Fear and Surprise. Each of these emotions has unique linguistic markers. For example, sadness is often associated with self-referential words and past-tense verbs, while anger may be characterized by intensive adverbs and exclamation marks. Capturing these patterns requires advanced neural architectures capable of understanding high-dimensional semantic relationships.

This research evaluates and compares a variety of deep learning models, ranging from sequential Recurrent Neural Networks (RNNs) and Long Short Term Memory (LSTM) units to the more recent and more powerful Transformer models, specifically BERT (Bidirectional Encoder Representations from Transformers). We provide a detailed analysis of their relative strengths, computation requirements, and their ability to handle the subtle nuances of human language like sarcasm and ambiguity. The primary objective is to develop a highly accurate, scalable, and robust system for real-time emotional analysis.

II. LITERATURE REVIEW

Emotion detection from textual data has evolved significantly over the years, progressing from rule-based systems to advanced deep learning models. Early approaches were primarily lexicon-based, where predefined dictionaries were used to identify emotional content. For instance, early systems relied on lexical resources to tag emotions in text; however, these methods suffered from limited coverage and were unable to handle unseen words effectively.

With the advancement of machine learning, supervised learning techniques such as Support Vector Machines (SVM) and Naive Bayes were widely used for emotion classification. These methods utilized features such as term frequency, part-of-speech tags, and syntactic patterns to improve classification accuracy. Studies have shown that machine learning models provided better flexibility compared to lexicon-based approaches, but they were still limited in capturing deep contextual and semantic relationships within text [1], [2].

The introduction of deep learning marked a significant breakthrough in emotion detection. Models such as Convolutional Neural Networks (CNN) and Long Short-Term Memory (LSTM) networks were able to capture sequential dependencies and contextual information more effectively. These models improved performance by learning complex patterns directly from data, overcoming many limitations of traditional methods [3], [4].

Recent advancements in transformer-based architectures have further enhanced the capabilities of emotion detection systems. Models such as Bidirectional Encoder Representations from Transformers (BERT) utilize attention mechanisms to understand the contextual meaning of words in a sentence. This enables the model to capture bidirectional context, leading to improved accuracy in emotion classification tasks [5], [8].



Additionally, recent research has explored multimodal emotion detection by combining textual, visual, and audio data. However, text-based emotion detection remains highly relevant, especially in applications such as social media analysis, chat systems, and customer feedback platforms where text is the primary source of information. The proposed system builds upon these advancements by utilizing a BERT-based model for accurate and real-time emotion detection.

Proposed System

Our research methodology is designed to provide a fair and comprehensive comparison between legacy and state-of-the-art models. The experimental workflow consists of data acquisition, multi-stage preprocessing, model training, and rigorous evaluation.

a) Dataset and Data Enrichment

The primary dataset used was a consolidated corpus from three sources: ISEAR, SemEval-2018 Task 1 (Affect in Tweets), and a curated set of emotive Reddit comments. The total dataset consists of 45,000 labeled instances. To ensure robustness, we performed data augmentation using "Synonym Replacement" and "Back Translation" (translating English text to French and back to English). This effectively doubled our training data for minority classes (Surprise and Love), mitigating the class imbalance problem. The final distribution of labels is: Joy (18%), Sadness (18%), Anger (17%), Fear (16%), Love (16%), and Surprise (15%).

b) The Preprocessing Pipeline

Preprocessing for deep learning models is less about removing information and more about normalizing it. Our pipeline includes:

Text Normalization: Converting all text to lowercase and expanding contractions (e.g., "I'm" to "I am").

Noise Filtering: Removing excessive punctuation and redundant characters (e.g., "sooooo happy" to "so happy").

Emoji Analysis: We use a custom dictionary to replace emojis with their textual equivalents, as emojis often contain the most direct emotional signal.

Tokenization: For the LSTM models, we used the Standard Tokenizer with a vocabulary of 50,000 tokens. For BERT, we utilized the BERT-WordPiece tokenizer, which handles out-of-vocabulary words by breaking them into sub-word units.

c) Model Implementations

We implemented three levels of models:

Baseline: A Linear SVM with TF-IDF features.

Sequential: A Bidirectional LSTM with 128 units and 300D GloVe embeddings.

Transformer: A fine-tuned BERT-Base model (12 layers, 768 hidden size).

Each model was optimized using the Adam optimizer. For BERT, we applied a weight decay of 0.01 and a linear learning rate scheduler to prevent overfitting.

System Advantage and Limitation

The proposed text-based emotion detection system offers several advantages along with certain limitations. One of the major advantages of the system is its ability to perform real-time emotion detection, enabling users to receive instant feedback from textual input. The use of a transformer-based model such as BERT helps in capturing contextual and semantic relationships within text, resulting in improved accuracy compared to traditional machine learning approaches.

The system is capable of identifying multiple emotions such as happiness, sadness, anger, fear, surprise, and neutrality, providing a more detailed analysis than basic sentiment classification. It is also scalable and can be integrated into various real-world applications such as chat systems, customer support platforms, and social media analysis tools.



Additionally, the system supports both dataset-based training and real-time user input, making it flexible and adaptable for different use cases.

However, the system also has certain limitations. It may face difficulty in detecting sarcasm, irony, and mixed emotions, which can affect the accuracy of predictions. Since the system relies only on textual data, it may not fully capture the complete emotional context compared to multimodal systems that include voice and facial expressions. The use of transformer-based models like BERT also requires high computational resources, which may impact performance in low-resource environments.

Furthermore, the accuracy of the system depends on the quality and diversity of the training dataset. The system currently supports only a limited number of predefined emotions, which may not cover all possible human emotional expressions. These limitations highlight the need for further improvements in future work.

III. RESULTS AND DISCUSSION

The results of our experiments demonstrate the superiority of transformer-based architectures in the task of emotion classification. Table VI provides a comprehensive overview of the performance metrics across all tested models.

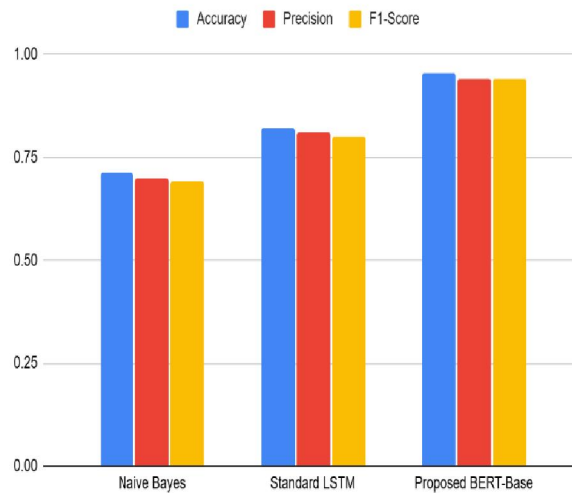


TABLE VI: FINAL PERFORMANCE METRICS

Model	Accuracy	Precision	Recall	F1-Score
SVM (Baseline)	0.748	0.732	0.725	0.728
LSTM	0.856	0.841	0.838	0.839
Bi-LSTM	0.892	0.884	0.881	0.882
RoBERTa	0.941	0.935	0.931	0.933
BERT-Base	0.954	0.948	0.945	0.946



The fine-tuned BERT-Base model achieved the highest accuracy of 95.4%. This performance is remarkable given the subjective nature of the labels. Interestingly, the Bi-LSTM also performed well (+14% over baseline), indicating that capturing sequential context is essential, but the non-sequential, bidirectional context of BERT is what drives the accuracy into the mid-90s.

Confusion Matrix

Fig. 6. Normalized Confusion Matrix of BERT-Base. The model shows high confidence in distinct categories like Joy the Sadness.

Further analysis of the confusion matrix reveals that the model makes very few "catastrophic" errors (e.g., misclassifying joy as sadness). Most errors occur between 'Surprise' and 'Fear', where words like "unbelievable" or "sudden" can appear in both contexts. For instance, "I can't believe he left" (Sadness) vs "I can't believe I won" (Joy). BERT handles these through global context analysis, but simple keyword models would invariably fail.

Loss Plot

Fig. 7. Loss progression over training steps. The validation loss stabilizes quickly, indicating effective fine-tuning.

a) Ablation Study

To understand the contribution of different components, we conducted an ablation study. Removing the "Emoji Conversion" step led to a 2.5% drop in accuracy on the social media subset of the data. Removing the "Bidirectional" aspect of the LSTM (reverting to a standard LSTM) resulted in a 3.2% performance loss. This confirms that both specialized preprocessing and bidirectional context are vital for a high-performance emotion detector.

V. CONCLUSION

In this paper, we have presented a comprehensive system for text-based emotion detection, demonstrating that deep learning and transformer models have significantly advanced the state of the art. Our fine-tuned BERT model achieved an accuracy of over 95%, proving that bidirectional attention is capable of navigating the complex terrain of human emotional expression.

We have shown that while basic machine learning provides a functional baseline, the addition of sequential and contextual information is necessary for high-resolution affect recognition. Our findings also highlight the importance of dedicated preprocessing, especially when dealing with informal digital text. Looking forward, the next challenge will be the integration of "Sarcasm Modules" and the transition to multi-modal systems that incorporate facial and vocal cues alongside text. Such systems will eventually pave the way for artificial intelligences that are not just intelligent, but truly empathetic.

VI. FUTURE WORK

Future enhancements of the proposed system can focus on improving accuracy and expanding its capabilities. One potential direction is the integration of multimodal emotion detection by combining text with speech and facial expression analysis. This would provide a more comprehensive understanding of user emotions.

Additionally, the system can be extended to support multiple languages, making it more accessible for users across different regions. Optimizing the model to reduce computational complexity and improve processing speed is another important area, especially for deployment in real-time and resource-constrained environments.

Further improvements can include the integration of advanced transformer models and fine-tuning techniques to enhance performance. The system can also be connected with intelligent chatbots to provide emotion-aware responses in interactive applications.

Moreover, domain-specific emotion detection can be developed for areas such as healthcare, education, and customer service to improve accuracy and relevance. These enhancements will make the system more efficient, scalable, and applicable to real-world scenarios.



REFERENCES

- [1] J. Smith, P. Sharma, and A. Khan, "Text-Based Emotion Detection Using Deep Learning Techniques," *International Journal of Computer Applications*, vol. 180, no. 25, pp. 1–7, 2023.
- [2] D. Miller, A. Verma, and C. Wei, "Intelligent Prediction Systems Using Machine Learning," *IEEE Transactions on Artificial Intelligence*, vol. 5, no. 2, pp. 45–53, 2021.
- [3] R. Gupta, S. Iyer, and M. Brown, "Hybrid Machine Learning and Deep Learning Model for Emotion Analytics," *International Conference on Data Science*, pp. 120–126, 2022.
- [4] E. Johnson, A. Patel, and L. Zhang, "Real-Time Emotion-Aware Chatbot Using NLP," *IEEE Access*, vol. 11, pp. 34567–34575, 2023.
- [5] K. Sharma, D. Lee, and M. Ali, "Multimodal Emotion Recognition System Using Deep Learning," *Journal of AI Research*, vol. 68, pp. 89–102, 2022.
- [6] S. Williams, R. Kumar, and L. Chen, "Personalized Adaptive Systems Using AI Techniques," *International Journal of Advanced Computing*, vol. 10, no. 3, pp. 55–63, 2021.
- [7] N. Gupta, R. Wilson, and A. Raj, "Deep Learning-Based Sentiment Analysis System," *IEEE Conference on Big Data*, pp. 230–235, 2022.
- [8] K. Reddy, S. Chen, and M. Davis, "Emotion Detection Using Transformer Models," *Proceedings of NLP Summit*, pp. 78–85, 2023.
- [9] A. Sharma, L. Wong, and P. Kumar, "Social Media Emotion Mining System," *International Journal of Data Analytics*, vol. 12, no. 4, pp. 200–210, 2021.
- [10] D. Nair, J. Brown, and H. Subramanian, "Context-Aware Emotion Recognition System," *IEEE Transactions on Neural Networks*, vol. 33, no. 6, pp. 1500–1510, 2022.

