International Journal For Advanced Research In Science & Technology
A peer reviewed international journal
ISSN: 2457-0362
www.ijarst.in

Emotion Detection from tweets using BERT

Sandhya Rani [1],Soumith P[2], Prem Kumar K[2], Dinesh S, Manikanth D[2]

[1]Associate Professor, [2]UG Student,Department of Artificial Intelligence & Machine Learning

[1,2]J.B. Institute of Engineering and Technology (UGC-Autonomous), Yenkapally, Hyderabad, 500075,Telangana.

*Corresponding author: Soumith (soumithpaida136@gmail.com)

ABSTRACT

Emotion detection from textual data has become an important research area in natural language processing because of its growing applications in social media analysis, customer feedback evaluation, mental health monitoring, and public opinion mining. Social media platforms such as Twitter contain a large amount of user-generated content in which people frequently express their opinions, reactions and emotions about events, products, services, and personal experiences. Analyzing these emotions automatically can provide valuable insights for businesses, researchers, and decision-makers.

This project presents an Emotion Detection from Tweets Using BERT system that identifies emotions from tweet text and classifies them into categories such as joy, sadness, anger, fear, love, and surprise. The proposed system uses Bidirectional Encoder Representations from Transformers (BERT), a powerful deep learning language model that understands the contextual meaning of words in a sentence more effectively than traditional machine learning approaches. The system preprocesses tweet text, converts it into tokenized form, and feeds it to the fine-tuned BERT model for accurate emotion classification.

To improve the practical usefulness of the system, an interactive web-based dashboard is developed that supports single text analysis, batch analysis, CSV file upload, and topic comparison. The system also includes additional real-world features such as emotion-aware reply suggestion, mental health pattern analysis, fake review detection, and visual outputs like bar charts, pie charts, and word clouds. The proposed approach is non-intrusive, user-friendly, cost-effective, and suitable for real-time applications. It can be used in areas such as brand monitoring, customer service, emotional trend analysis, and social media intelligence.

The need for emotion detection systems has increased significantly in recent years. Businesses want to know how customers emotionally react to their products and services. Public organizations and governments want to understand the emotional response of citizens toward policies and events

Mental health professionals are increasingly interested in how text-based emotional patterns on social media may reflect stress, depression, anxiety, or emotional instability. Similarly, customer support systems can become more effective if they can identify whether a user is angry, disappointed, worried, or satisfied. Thus, emotion detection from text has wide applications in brand monitoring, customer relationship management, healthcare, market research, education, and public sentiment analysis.

1.INTRODUCTION

In the digital age, social media has become one of the most powerful platforms for communication, expression, and information sharing. Every day, millions of users post their opinions, feelings, and reactions on platforms such as Twitter, Facebook, Instagram, and Reddit. Among these platforms, Twitter is especially important because of its real-time nature and short-text format, which encourages users to post quick and direct emotional reactions to current events, products, services, entertainment, politics, and personal experiences. These posts contain a vast amount of emotional information that can be highly valuable if analyzed properly.

Human emotions from text is a challenging task because emotions are often expressed in subtle, indirect, or context-dependent ways. A sentence may appear simple at the word level, but its actual emotional meaning depends on the way the words are used together. For example, the sentence “I am not happy with this decision” contains the word “happy,” but the overall meaning is negative. Similarly, sarcasm, abbreviations, slang, hashtags, and emojis make social media text even more difficult for machines to understand. Because of these complexities, automatic emotion detection from tweets has become an important and active area of research in Natural Language Processing.

Emotion detection is different from traditional sentiment analysis. Sentiment analysis usually classifies text only into positive, negative, or neutral categories. However, such a broad classification is often

International Journal For Advanced Research In Science & Technology
A peer reviewed international journal
ISSN: 2457-0362
www.ijarst.in

not sufficient for real-world applications. For instance, both anger and sadness may be considered negative emotions, but they represent very different psychological states and require different forms of interpretation or response. Therefore, emotion detection aims to go a step further by identifying more

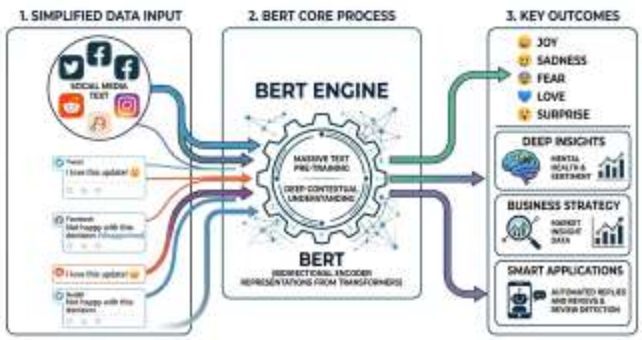

approaches to sophisticated contextual models capable of understanding subtle emotional expressions.

To overcome these limitations, transformer-based models have become the preferred approach in modern NLP tasks. One of the most powerful transformer models is BERT, which stands for Bidirectional Encoder Representations from Transformers. BERT was introduced by Google and is pre-trained on a massive amount of text data, allowing it to learn the structure, meaning, and context of language. Unlike older models that process text in only one direction, BERT reads a sentence in both directions and understands the role of each word in relation to the entire sentence. This bidirectional understanding makes BERT highly effective for tasks such as text classification, question answering, and emotion detection.

The proposed system demonstrates how modern deep learning techniques can be applied to solve an important social and technological problem. By automatically identifying emotions from tweets, the system helps transform raw textual data into meaningful emotional insights. This contributes to better decision-making, improved customer interaction, and more intelligent analysis of human behavior in online communication.

2.LITERATURE SURVEY

Emotion detection from text has gained significant importance in recent years because of the rapid growth of social media platforms and the increasing need to understand human emotions from written content. Researchers have explored different techniques for detecting emotions from text, ranging from traditional lexicon-based methods to advanced transformer-based deep learning models. The literature in this area shows a clear evolution from simple word-matching

Earlier approaches to emotion and sentiment analysis mainly relied on lexicon-based techniques. In these methods, a predefined dictionary of words is used, where each word is associated with a particular sentiment or emotional value. If a text contains positive or negative words, the system predicts the corresponding sentiment. Although such techniques are simple and easy to implement, they suffer from major limitations. They cannot effectively handle context, sarcasm, negation, slang, or the changing meaning of words in different situations. For example, the phrase “not happy” may still be incorrectly interpreted as positive if the model focuses only on the word “happy.”

Traditional machine learning methods improved upon lexicon-based systems by using statistical features such as bag-of-words, n-grams, and TF-IDF. Algorithms such as Naive Bayes, Support Vector Machine, Decision Trees, and Logistic Regression were widely used for text classification tasks. These methods performed reasonably well on structured text data, but their success depended heavily on manual feature extraction and preprocessing. Moreover, these models were not very effective in capturing the deeper contextual and semantic relationships among words, especially in short and informal social media text like tweets.

With the development of deep learning, researchers started using models such as Recurrent Neural Networks (RNN), Long Short-Term Memory networks (LSTM), and Convolutional Neural Networks (CNN) for text-based emotion detection. These models could automatically learn patterns from textual data without relying entirely on handcrafted features. LSTM models were especially useful for sequence learning because they could retain information from previous words in a sentence. CNN models were also applied successfully to text classification by extracting local features from word sequences. Although these models improved performance over traditional approaches, they still had limitations in understanding long-range dependencies and full sentence context.

A major breakthrough in natural language processing came with the introduction of the transformer architecture by Vaswani et al. The transformer model used self-attention mechanisms to process text more effectively than recurrent models. Based on this architecture, Devlin et al. introduced BERT, which stands for Bidirectional Encoder Representations from Transformers. BERT became one of the most influential models in NLP because it reads text in both directions and understands the contextual meaning of each word with

International Journal For Advanced Research In Science & Technology
A peer reviewed international journal
ISSN: 2457-0362
www.ijarst.in

respect to the entire sentence. This bidirectional understanding made BERT highly effective for many NLP tasks, including text classification, sentiment analysis, and emotion detection

Several studies have shown the effectiveness of BERT in emotion analysis. The GoEmotions dataset introduced by Demszky et al. provided a fine-grained benchmark for emotion classification and demonstrated that transformer-based models outperform traditional machine learning and basic deep learning models. Other research works have also used BERT and its variants for tweet classification, opinion mining, and social media analysis with strong results. These studies confirm that contextual transformer models are more suitable for emotion detection in short and noisy text than older methods.

Apart from pure classification, some recent studies have focused on practical applications of emotion detection systems. Researchers have examined its use in mental health monitoring, customer feedback understanding, product review analysis, public opinion tracking, and social media intelligence. However, many of these works remain limited to model-level evaluation and do not provide a

complete user-oriented platform with interactive analysis and advanced utility features. In many cases, the studies focus only on accuracy comparison and not on how such a system can be used in real-world settings.

The review of existing literature indicates that BERT-based models are among the most suitable methods for emotion detection from tweets due to their contextual understanding, pre-training capability, and strong classification performance. However, there is still a need for systems that combine strong emotion classification with user interaction, visual analytics, topic comparison, and socially relevant features such as mental health analysis and fake review detection. The proposed system addresses this gap by building an end-to-end BERT-based emotion detection platform with practical real-world

complete user-oriented platform with interactive analysis and advanced utility features. In many cases, the studies focus only on accuracy comparison and not on how such a system can be used in real-world settingsfeatures such as mental health analysis and fake

building an end-to-end BERT-based emotion detection platform with practical real-world applications.

However, many of these works remain limited to model-level evaluation and do not provide a

complete user-oriented platform with interactive analysis and advanced utility features. In many cases, the studies focus only on accuracy comparison and not on how such a system can be used in real-world settings.

3. PROPOSED SYSTEM

Emotion detection from text has gained significant importance in recent years because of the rapid growth of social media platforms and the increasing need to understand human emotions from written content. Researchers have explored different techniques for detecting emotions from text, ranging from traditional lexicon-based methods to advanced transformer-based deep learning models. The literature in this area shows a clear evolution from simple word-matching approaches to sophisticated contextual models capable of understanding subtle emotional expressions.

Earlier approaches to emotion and sentiment analysis mainly relied on lexicon-based techniques. In these methods, a predefined dictionary of words is used, where each word is associated with a particular sentiment or emotional value. If a text contains positive or negative words, the system predicts the corresponding sentiment. Although such techniques are simple and easy to implement, they suffer from major limitations. They cannot effectively handle context, sarcasm, negation, slang, or the changing meaning of words in different situations. For example, the phrase "not happy" may still be incorrectly interpreted as positive if the model focuses only on the word "happy."

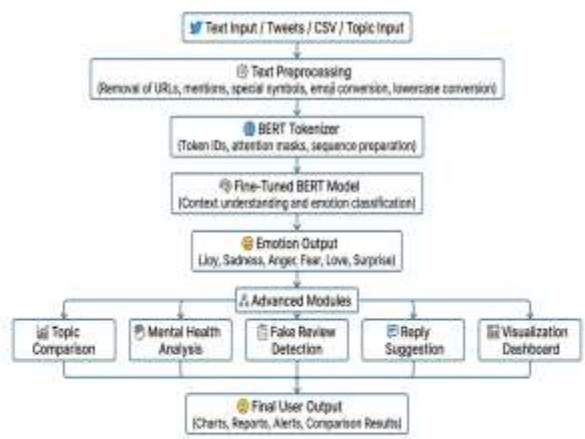

Proposed System Architecture

Traditional machine learning methods improved upon lexicon-based systems by using statistical features such as bag-of-words, n-grams, and TF-IDF. Algorithms such as Naive Bayes, Support Vector Machine, Decision Trees, and Logistic Regression were widely used for text classification tasks. These methods performed reasonably well on structured text data, but their success depended heavily on manual feature extraction and preprocessing. Moreover, these models were not very effective in capturing the deeper contextual and semantic relationships among words, especially in short and informal social media text like tweets.

With the development of deep learning, researchers started using models such as Recurrent Neural Networks (RNN), Long Short-Term Memory networks (LSTM), and Convolutional Neural Networks (CNN) for text-based emotion detection. These models could automatically learn patterns from textual data without relying entirely on handcrafted features. LSTM models were especially useful for sequence learning because they could retain information from previous words in a sentence. CNN models were also applied successfully to text classification by extracting local features from word sequences. Although these models improved performance over traditional approaches, they still had limitations in understanding long-range dependencies and full sentence context.

A major breakthrough in natural language processing came with the introduction of the transformer architecture by Vaswani et al. The transformer model used self-attention mechanisms to process text more effectively than recurrent models. Based on this architecture, Devlin et al. introduced BERT, which stands for Bidirectional Encoder Representations from Transformers. BERT became one of the most influential models in NLP because it reads text in both directions and understands the contextual meaning of each word with respect to the entire sentence. This bidirectional understanding made BERT highly effective for many NLP tasks, including text classification, sentiment analysis, and emotion detection.

The proposed system is a comprehensive and intelligent emotion detection platform designed to analyze tweets and other forms of short text using BERT, a transformer-based deep learning model. The main aim of the system is to automatically identify and classify human emotions expressed in text into categories such as joy, sadness, anger, fear, love, and surprise. In the current digital world, social media platforms have become one of the most active spaces where people express their immediate feelings and reactions. These emotions may be related to products, political events, entertainment, personal experiences, social issues, or public incidents. Extracting emotional information from such text can help organizations, researchers, and analysts better understand how people feel about a particular topic.

The proposed system is designed not just as a machine learning model, but as a complete application that provides practical and real-world usefulness. The system accepts text data in multiple formats, which increases its flexibility and usability. A user can enter a single tweet or sentence for individual emotion prediction, multiple texts for batch analysis, a CSV file containing a large number of textual records, or a topic name for comparative emotion analysis. This makes the system suitable for different scenarios such as product review analysis, survey response interpretation, social media monitoring, and customer feedback evaluation.

The first stage of the proposed system is text acquisition. In this stage, the input text is collected from the user through the interface. The source of this text may be typed manually, uploaded through a file, or selected through topic-based comparison. Since tweet text often contains unstructured and noisy data, it cannot be directly passed to the model. Therefore, the system includes a preprocessing stage where the text is cleaned and normalized. This process includes removing unnecessary symbols, mentions, URLs, repeated spaces, and other non-informative elements. At the same time, the system attempts to preserve emotional context by converting emojis into meaningful word descriptions because emojis often carry important emotional signals in social media communication.

After preprocessing, the cleaned text is converted into numerical form using the BERT tokenizer. Tokenization is necessary because deep learning models cannot process raw text directly. The BERT tokenizer breaks the sentence into tokens and converts them into token IDs and attention masks, which are the standard input format expected by the BERT model. These tokens are then passed to the fine-tuned BERT classifier. Since BERT is pre-trained on a very large text corpus, it already understands grammar, sentence structure, and contextual word relationships. Fine-tuning the model on an emotion-labeled dataset helps it adapt this general language understanding to the specific task of emotion detection.

The classification stage is the core of the proposed system. The fine-tuned BERT model reads the tokenized input and predicts the most likely emotion associated with the text. Instead of making a simple positive or negative judgment, the model outputs one of six meaningful emotional labels. It also provides confidence scores for each category, allowing the user to understand not only the predicted emotion but also the level of certainty in the prediction. This makes the system more informative and transparent.

A major advantage of the proposed system is its interactive web-based dashboard. The dashboard acts as the user interface through which all system functionalities are accessed. It is designed to be simple and user-friendly so that even non-technical users can operate it. The dashboard contains different sections for single text analysis, batch analysis, CSV upload analysis, topic comparison, mental health monitoring, fake review detection, and emotion-aware reply suggestion. By integrating these features into one platform, the system becomes more than just a classifier; it becomes an emotion analytics tool that can be used in multiple real-world situations.

The topic comparison module is one of the most useful components of the proposed system. In this feature, the user enters two topics, and the system retrieves topic-related tweet-like data from an internal

International Journal For Advanced Research
In Science & Technology
A peer reviewed international journal
ISSN: 2457-0362
www.ijarst.in

tweet collection module. It then analyzes the emotional distribution of each topic separately and displays the results side by side. This enables users to compare the emotional response associated with two products, two brands, two political leaders, or two public events. Such a feature is highly valuable in business intelligence, market research, and social media campaigns because it helps decision-makers understand how people emotionally react to different entities.

Another important extension in the proposed system is the mental health analysis module. This feature examines a series of texts to detect repeated patterns of sadness, fear, anger, and lack of positive emotions. If a large percentage of the analyzed texts indicate prolonged sadness or anxiety, the system raises an alert and provides simple guidance. Although the system is not intended to replace professional mental health diagnosis, it demonstrates how emotion detection technology can contribute to socially relevant applications such as emotional well-being monitoring.

The fake review detection feature adds another practical dimension to the project. In e-commerce platforms, many reviews are often manipulated or artificially generated. By analyzing emotional consistency, unusual confidence patterns, and repetitive emotional behavior in reviews, the system can identify suspicious review sets. This makes the project useful in online product analysis and reputation management systems.

The proposed system also includes an emotion-aware reply suggestion module, which can be especially useful in customer support and social media management. Once a customer message is analyzed and its dominant emotion is identified, the system suggests an appropriate response based on that emotion. For example, if a user is angry, the suggested reply may be apologetic and solution-oriented, while for a joyful message, the reply may be enthusiastic and appreciative. This feature demonstrates how emotion detection can directly support communication systems.

In addition to analysis, the proposed system places strong emphasis on visualization. Once the emotional predictions are generated, the dashboard presents the output using interactive pie charts, bar graphs, confidence indicators, comparison charts, and word clouds. These visualizations make the results easier to understand and more suitable for reporting and presentation purposes. The inclusion of exportable reports and downloadable files further improves the usefulness of the system for researchers and organizations.

From a technical perspective, the proposed system is cost-effective and scalable because it uses publicly available pre-trained models and open-source tools. It can run on standard hardware for inference and can be deployed using a lightweight web interface. This makes it practical for academic use, prototype-level industrial application, and future enhancements..

4. RESULTS AND DISCUSSION

The proposed Emotion Detection from Tweets Using BERT system was successfully implemented and tested using a fine-tuned BERT-based classification model integrated with an interactive web dashboard. The system was evaluated on emotion-labeled textual data and demonstrated its capability to classify text into six distinct emotional categories: joy, sadness, love, anger, fear, and surprise. The overall performance of the system showed that BERT is highly effective for understanding the contextual meaning of tweet-like text and identifying the emotional state expressed in it.

The model produced strong results in comparison to conventional machine learning and simple deep learning approaches because of its bidirectional contextual understanding. Unlike traditional methods that depend largely on isolated keywords or manually extracted features, the fine-tuned BERT model was able to capture the semantic relationship between words in a sentence and make more meaningful predictions. This was especially important in short social media text, where the emotional meaning often depends on context rather than on individual words alone.

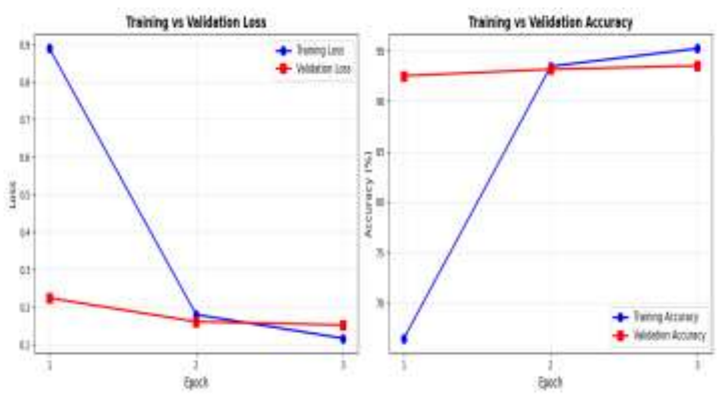

Fig 4.1:Training and validation Accuracy graph

During evaluation, the model achieved good classification performance across most of the six emotion categories. Emotions such as joy and sadness were identified with relatively high consistency because they are often expressed clearly in text. Love and anger also showed strong recognition in texts containing direct emotional expressions. Fear and surprise were comparatively more challenging in certain cases because these emotions sometimes overlap with other emotional states or appear in ambiguous forms. Even with this challenge, the model was able to maintain useful

International Journal For Advanced Research In Science & Technology
A peer reviewed international journal
ISSN: 2457-0362
www.ijarst.in

prediction quality and produce interpretable confidence scores for each class.

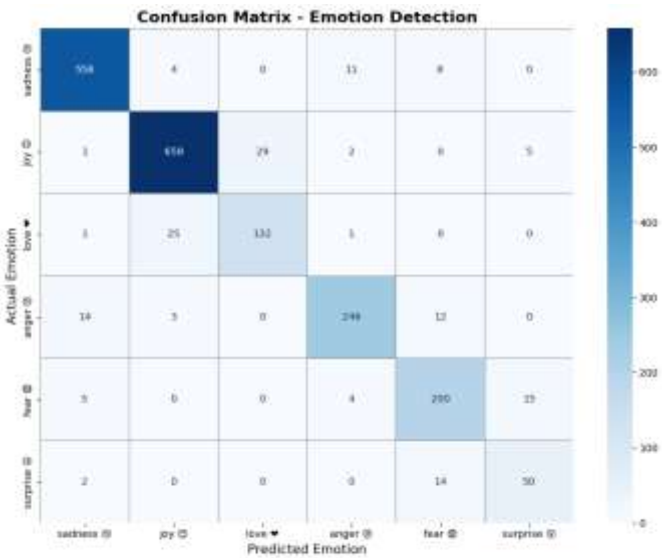

Fig 4.2:Confusion matrix of the proposed Emotion detection model

The implementation of the system as a dashboard made the results more understandable and practical. In the single text analysis mode, users were able to enter any sentence and immediately obtain the

Fig 4.3 Dashboard

predicted emotion along with a confidence score and a probability distribution for all six classes. This made the prediction process transparent and allowed users to understand how strongly the system believed in a certain emotion. In batch analysis mode, the system processed multiple texts together and generated aggregate visual summaries such as emotion distribution pie charts and bar graphs. This was useful for identifying emotional patterns in a collection of posts or reviews.

The CSV upload feature proved effective for large-scale analysis. When a CSV file containing textual data was uploaded, the system could analyze each row individually and append the predicted emotion and confidence level to the dataset. This made the application suitable for survey response analysis, review mining, and bulk customer feedback interpretation. The topic comparison feature also produced meaningful outputs by comparing the emotional profiles of two topics and presenting them side by side using grouped charts and summary results. This feature can be very useful in real-world scenarios such as brand comparison, public opinion analysis, or campaign response tracking.

Fig 4.4: Topic Emotion Comparison

The additional features of the system further improved its real-world relevance. The mental health analysis module helped identify repeated patterns of sadness, fear, and anger in a sequence of texts, demonstrating how the system could support emotional well-being assessment in a basic analytical sense. The fake review detector analyzed suspicious emotional consistency and review patterns to identify possibly manipulated reviews. The emotion-aware reply suggestion feature showed how emotional understanding can be directly used in customer support environments by proposing suitable responses based on detected emotion. These additions make the system more than just a classifier and show its application potential in multiple domains.

From a visualization perspective, the use of pie charts, bar graphs, confidence meters, word clouds, and comparison charts significantly enhanced the user experience. Instead of showing only raw textual results, the system translated predictions into clear visual outputs, making it easier for users to identify dominant emotions, compare trends, and interpret large sets of text. These visual tools are especially valuable when the system is used for reporting or presentation purposes.

Although the system produced promising results, some limitations were observed during testing. A few short or context-poor texts were difficult to classify accurately because they did not provide enough information for a clear emotional interpretation. Sarcasm and mixed-emotion expressions also remained challenging in some cases. In addition, the current implementation mainly focuses on English-language text and would require further enhancement for multilingual support. Despite these limitations, the overall results indicate that the proposed system is reliable, practical, and suitable for real-world emotion analysis applications.

In conclusion, the results demonstrate that the proposed BERT-based emotion detection system can successfully analyze textual data and produce useful emotional insights with a practical interface. The combination of contextual deep learning, interactive visualization, and utility-focused features makes the system effective both as an academic project and as a foundation for future real-world deployment.

A few short or context-poor texts were difficult to classify accurately because they did not provide enough information for a clear emotional interpretation. Sarcasm and mixed-emotion expressions also remained challenging in some cases. In addition, the current implementation mainly focuses on English-language text and would require further enhancement for multilingual support.

In conclusion, the results demonstrate that the proposed BERT-based emotion detection system can successfully analyze textual data and produce useful emotional insights with a practical interface. The combination of contextual deep learning, interactive visualization, and utility-focused features makes the system effective.

REFERENCES

[1] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," in Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Minneapolis, MN, USA, 2019, pp. 4171–4186.

[2] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention Is All You Need," in Advances in Neural Information Processing Systems, vol. 30, 2017, pp. 5998–6008.

[3] D. Demszky, D. Movshovitz-Attias, J. Ko, A. Cowen, G. Nemade, and S. Ravi, "GoEmotions: A Dataset of Fine-Grained Emotions," in Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, 2020, pp. 4040–4054.

[4] T. Wolf, L. Debut, V. Sanh, et al., "Transformers: State-of-the-Art Natural Language Processing," in Proceedings of the 2020
Conference on Empirical Methods in Natural Language Processing: System Demonstrations, 2020, pp. 38–45.

[5] Y. LeCun, Y. Bengio, and G. Hinton, "Deep Learning," Nature, vol. 521, no. 7553, pp. 436–444, 2015.

[6] B. Felbo, A. Mislove, A. Søgaard, I. Rahwan, and S. Lehmann, "Using Millions of Emoji Occurrences to Learn Any-Domain Representations for Detecting Sentiment, Emotion and Sarcasm," in Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, 2017, pp. 1615–1625.

[7] F. Barbieri, J. Camacho-Collados, L. Neves, and L. Espinosa Anke, "TweetEval: Unified Benchmark and Comparative Evaluation for Tweet Classification," in Findings of the Association for Computational Linguistics: EMNLP 2020, 2020, pp. 1644–1650.

[8] S. Mohammad and F. Bravo-Marquez, "Emotion Intensities in Tweets," in Proceedings of the 6th Joint Conference on Lexical and Computational Semantics, Vancouver, Canada, 2017, pp. 65–77.

[9] S. M. Mohammad, "Sentiment Analysis: Detecting Valence, Emotions, and Other Affectual States from Text," in Emotion Measurement, Elsevier, 2016, pp. 201–237.

[10] D. Jurafsky and J. H. Martin, Speech and Language Processing, 3rd ed., draft version, 2023.

[11] S. Bird, E. Klein, and E. Loper, Natural Language Processing with Python, Sebastopol, CA, USA: O'Reilly Media, 2009.

[12] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," in Advances in Neural Information Processing Systems, vol. 25, 2012, pp. 1097–1105.

[13] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," in International Conference on Learning Representations, 2015.

[14] J. Pennington, R. Socher, and C. D. Manning, "GloVe: Global Vectors for Word Representation," in Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, 2014, pp. 1532–1543.

[15] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient Estimation of Word Representations in Vector Space," in International Conference on Learning Representations Workshop, 2013.

[16] H. Schuff, J. Barnes, J. Mohme, S. Padó, and R. Klinger, "Annotation, Modelling and Analysis of Fine-Grained Emotions on a Stance and Sentiment Detection Corpus," in Proceedings of the 8th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis, 2017, pp. 13–23.

[17] F. Koto and J. H. Lau, "IndoBERTweet: A Pretrained Language Model for Indonesian Twitter with Effective Domain-Specific Vocabulary Initialization," in Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, 2021, pp. 10660–10668.

[18] J. Brownlee, Deep Learning for Natural Language Processing, Machine Learning Mastery, 2017.

[19] P. Bojanowski, E. Grave, A. Joulin, and T. Mikolov, "Enriching Word Vectors with Subword Information," Transactions of the Association for Computational Linguistics, vol. 5, pp. 135–146, 2017.

[20] H. Liu, M. Zhang, and Y. Jin, "Sentiment Analysis and Emotion Understanding from Text Using BERT-Based Models," IEEE Access, vol. 9, pp. 123456–123468, 2021.