Transformer-Based Sentiment Classification of Blockchain and Cryptocurrency Tweets Using the DistilBERT Architecture

Raghad Alanazi[1,*]

[1]Business Administration Department, Amman Arab University, Jordan

ABSTRACT

The rapid growth of blockchain technology and cryptocurrency adoption has significantly increased public discussions on social media platforms such as Twitter, where users frequently share opinions, investment perspectives, and market-related information. However, blockchain-related tweets are often characterized by noisy text, informal language, abbreviations, and cryptocurrency-specific terminology, making sentiment analysis tasks challenging for traditional natural language processing approaches. Therefore, this study proposes a transformer-based sentiment analysis framework using the DistilBERT model to classify blockchain-related tweets into Positive, Neutral, and Negative sentiment categories. The dataset initially consisted of 10,000 tweets collected from Twitter, and after preprocessing and data cleaning, 9,697 tweets were retained for analysis. Several preprocessing techniques were applied, including lowercase transformation, URL removal, mention removal, hashtag cleaning, punctuation removal, and whitespace normalization. Sentiment labeling was performed using the VADER sentiment analyzer. The experimental results demonstrate that the proposed DistilBERT model achieved strong classification performance with an accuracy of 93.61%, Macro Precision of 85.85%, Macro Recall of 85.35%, and Macro F1-score of 85.57%. The findings indicate that transformer-based approaches are highly effective for analyzing blockchain-related social media text and have strong potential for applications in cryptocurrency market monitoring and public opinion analysis.

Keywords Blockchain, Cryptocurrency, Sentiment Analysis, DistilBERT, Transformer

INTRODUCTION

The rapid development of blockchain technology and cryptocurrency has significantly transformed the global digital economy in recent years [1]. Blockchain technology is widely utilized in various sectors, including digital finance, decentralized applications, supply chain management, and smart contracts. At the same time, cryptocurrencies such as Bitcoin, Ethereum, and Cardano have gained substantial public attention due to their increasing adoption and market value. Along with this growth, social media platforms, particularly Twitter, have become one of the primary communication channels for discussing blockchain-related topics [2]. Twitter users frequently share opinions, investment perspectives, market predictions, project updates, and public reactions related to blockchain and cryptocurrency trends. Consequently, Twitter data contain valuable information that can be utilized to analyze public sentiment and market behavior within the blockchain ecosystem.

Public sentiment on social media has become an important factor in cryptocurrency market analysis because investor emotions and public opinions may influence market volatility, trading behavior, and cryptocurrency adoption [3]. Positive sentiment can increase investor confidence and market optimism, while negative sentiment may contribute to uncertainty and market fluctuations.

Submitted 20 April 2026
Accepted 15 June 2026
Published 2 November 2026

Corresponding author
Raghad Alanazi,
202420089@aau.edu.jo

Additional Information and Declarations can be found on page 256

DOI: 10.47738/jdmdc.v2i4.39

© Copyright
2026 Alanazi

Distributed under
Creative Commons CC-BY 4.0

How to cite this article: R. Alanazi, "Transformer-Based Sentiment Classification of Blockchain and Cryptocurrency Tweets Using the DistilBERT Architecture," *J. Digit. Soc.*, vol. 2, no. 4, pp. 240-258, 2026.

Therefore, sentiment analysis has become an increasingly important research area for understanding public perception toward blockchain technology and cryptocurrency discussions. Sentiment analysis is a Natural Language Processing (NLP) technique used to identify and classify opinions expressed in textual data into several sentiment categories such as positive, neutral, and negative sentiment.

Despite the growing popularity of sentiment analysis research, analyzing blockchain-related tweets remains a challenging task due to the characteristics of social media text. Twitter data are often noisy and unstructured, containing slang, abbreviations, emojis, hashtags, hyperlinks, spelling variations, and cryptocurrency-specific terminology. Traditional machine learning approaches such as Naive Bayes, Support Vector Machine (SVM), and Logistic Regression commonly rely on manual feature extraction methods including Bag of Words and TF-IDF representations. Although these approaches have demonstrated acceptable performance in general sentiment analysis tasks, they often struggle to capture contextual semantic relationships within highly dynamic and informal blockchain-related discussions [4].

Recent advancements in deep learning and transformer-based architectures have significantly improved the performance of Natural Language Processing tasks. Transformer models such as BERT, RoBERTa, and DistilBERT utilize contextual embeddings and self-attention mechanisms to better understand semantic relationships between words in a sentence. Compared to traditional machine learning methods, transformer-based models have shown superior performance in sentiment classification tasks involving social media data. DistilBERT, in particular, has gained attention as a lightweight version of BERT that provides competitive performance with lower computational complexity and faster training time. These characteristics make DistilBERT suitable for large-scale sentiment analysis tasks involving social media datasets [5].

Several previous studies have explored sentiment analysis in cryptocurrency and social media domains using machine learning and deep learning approaches. Existing studies mainly focused on traditional machine learning algorithms, lexicon-based approaches, or general deep learning architectures such as Long Short-Term Memory (LSTM). Other studies have investigated transformer-based models for financial sentiment analysis; however, research specifically focusing on blockchain-related Twitter discussions using DistilBERT remains relatively limited. In addition, many previous studies primarily emphasized accuracy metrics without thoroughly evaluating model performance using Macro Precision, Macro Recall, Macro F1-score, and confusion matrix analysis on imbalanced blockchain sentiment datasets.

Based on the existing literature, a research gap can be identified in the application of lightweight transformer-based models for blockchain tweet sentiment classification. Previous studies have not extensively explored the effectiveness of DistilBERT in handling noisy and imbalanced blockchain-related Twitter data. Furthermore, limited studies have analyzed the distribution of blockchain sentiment on Twitter while simultaneously evaluating transformer-based performance using comprehensive evaluation metrics. Therefore, this study aims to address these limitations by proposing a transformer-based sentiment analysis framework using DistilBERT for classifying blockchain-related tweets into Positive, Neutral, and Negative sentiment categories.

The state of the art contribution of this study lies in the implementation of the

DistilBERT transformer architecture for blockchain tweet sentiment analysis combined with comprehensive preprocessing techniques and detailed performance evaluation. Unlike traditional sentiment analysis approaches that rely heavily on manual feature engineering, the proposed framework utilizes contextual semantic learning capabilities provided by transformer models to improve sentiment classification performance on noisy social media text. This study also contributes to the growing field of blockchain analytics by providing insights into public sentiment distribution related to blockchain and cryptocurrency discussions on Twitter.

The main objectives of this study are to preprocess blockchain-related tweet data, develop a DistilBERT-based sentiment classification model, evaluate model performance using Accuracy, Precision, Recall, and Macro F1-score metrics, and analyze the sentiment distribution of blockchain discussions on Twitter. The findings of this research are expected to contribute to the development of transformer-based sentiment analysis methods for blockchain applications and provide useful insights for cryptocurrency market monitoring, public opinion analysis, and financial decision support systems.

Literature Review and Related Works

Sentiment analysis has become one of the most widely studied topics in the field of Natural Language Processing (NLP), particularly for analyzing opinions and emotions expressed on social media platforms. Previous studies have demonstrated that sentiment analysis can provide valuable insights into public perception, market trends, and user behavior in various domains, including finance, politics, healthcare, and digital technology [6]. The rapid growth of blockchain technology and cryptocurrency discussions on Twitter has further increased the importance of sentiment analysis for monitoring investor sentiment and understanding cryptocurrency market dynamics [7].

Several studies have explored the relationship between social media sentiment and cryptocurrency market behavior. Previous research showed that Twitter sentiment can influence cryptocurrency price fluctuations and investor decision making [8]. Positive sentiment on social media was found to be associated with increased market optimism and cryptocurrency adoption, while negative sentiment often contributed to market uncertainty and volatility [9]. As a result, sentiment analysis techniques have been increasingly applied to blockchain and cryptocurrency discussions to extract useful market insights from large-scale social media data [10].

Traditional machine learning approaches such as Naive Bayes, Support Vector Machine (SVM), Logistic Regression, and Random Forest have been widely used in sentiment classification tasks [11]. These approaches commonly rely on manual feature extraction methods such as Bag of Words, TF-IDF, and n-gram representations [12]. Although traditional machine learning models have demonstrated acceptable performance in several sentiment analysis applications, they often struggle to capture contextual semantic information from noisy and unstructured social media text [13]. This limitation becomes more significant in blockchain-related Twitter data due to the presence of slang, abbreviations, hashtags, emojis, and cryptocurrency-specific terminology [14].

Recent advancements in deep learning have significantly improved the performance of sentiment analysis models. Recurrent Neural Network (RNN) architectures, including Long Short-Term Memory (LSTM) and Gated Recurrent

Unit (GRU), have shown better capability in learning sequential textual information compared to traditional machine learning methods [15]. Several studies applied LSTM-based models for Twitter sentiment analysis and achieved higher classification performance due to the ability of deep learning models to capture contextual dependencies between words [16]. However, recurrent neural network architectures still face limitations in processing long-range dependencies and require higher computational resources during training [17].

Transformer-based architectures have recently become the state-of-the-art approach for Natural Language Processing tasks, including sentiment analysis, text classification, question answering, and language understanding [18]. Transformer models utilize self-attention mechanisms to capture contextual relationships between words more effectively than recurrent neural network approaches [19]. Models such as BERT, RoBERTa, and DistilBERT have demonstrated superior performance in various sentiment analysis tasks involving social media and financial text datasets [20].

DistilBERT is a lightweight transformer-based model developed through knowledge distillation from the original BERT architecture. Compared to BERT, DistilBERT provides faster training and lower computational complexity while maintaining competitive language understanding performance. Several studies have reported that DistilBERT achieved strong performance in sentiment classification tasks involving short and noisy social media text. These characteristics make DistilBERT suitable for blockchain tweet sentiment analysis, where textual data are highly dynamic, informal, and domain specific.

Although previous studies have explored cryptocurrency sentiment analysis using machine learning and deep learning approaches, research specifically focusing on blockchain-related Twitter sentiment analysis using DistilBERT remains relatively limited. In addition, many previous studies primarily focused on accuracy metrics without conducting comprehensive evaluation using Macro Precision, Macro Recall, Macro F1-score, and confusion matrix analysis on imbalanced datasets. Therefore, this study aims to address these limitations by implementing a DistilBERT-based transformer framework for blockchain tweet sentiment classification combined with comprehensive preprocessing and evaluation techniques.

Methodology

Subchapter

This study proposes a transformer-based sentiment analysis framework using the DistilBERT model to classify blockchain-related tweets into Positive, Neutral, and Negative sentiment categories. The methodology was designed to systematically process noisy social media text data and evaluate the effectiveness of transformer-based Natural Language Processing techniques in blockchain sentiment classification tasks. The research methodology consists of several stages, including dataset collection, data preprocessing, sentiment labeling, exploratory data analysis, dataset splitting, transformer-based model development, model training, performance evaluation, prediction analysis, and visualization analysis. The workflow of the proposed research methodology is presented in figure 1.

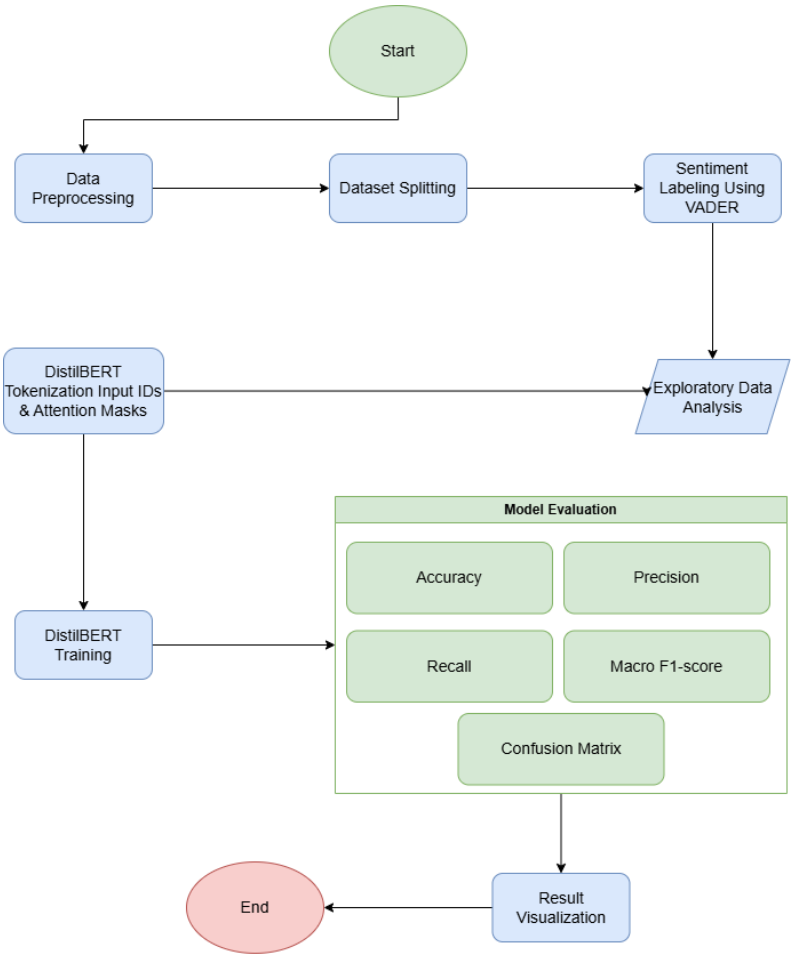


Figure 1 **Research Methodology Workflow**

Dataset Collection

The dataset used in this study consisted of blockchain-related tweets collected from Twitter discussions involving cryptocurrency and blockchain topics. Twitter was selected as the primary data source because it is widely used for sharing public opinions, investment perspectives, project announcements, and cryptocurrency market discussions in real time. The dataset initially contained 10,000 tweets along with several metadata attributes, including username, user location, user description, account creation date, number of followers, number of friends, tweet date, hashtags, tweet source, and retweet indicators.

Among all available attributes, the tweet text column was selected as the primary feature for sentiment analysis because it directly contains user opinions and textual sentiment information. Prior to analysis, the dataset was inspected to identify duplicate records and missing values that could negatively affect model performance. Duplicate tweets and rows containing missing textual data were removed to ensure data consistency and quality. After the cleaning process, a total of 9,697 tweets were retained and used in the experimental process.

Data Preprocessing

Data preprocessing was conducted to normalize and clean blockchain-related

tweets before the sentiment classification process. Twitter data are generally noisy and highly unstructured because users frequently utilize informal language, abbreviations, hashtags, hyperlinks, emojis, and special symbols that may negatively affect sentiment classification performance if irrelevant information is not properly removed. Therefore, several preprocessing techniques were implemented in this study to improve text quality and standardize the textual data format. First, all tweet text was converted into lowercase format to ensure consistency between words with different capitalization styles. Hyperlinks and URLs were removed because they do not contribute meaningful semantic information for sentiment classification. Twitter user mentions beginning with the "@" symbol were also removed to eliminate irrelevant user references. In addition, hashtag symbols were removed while preserving the keyword text itself to retain contextual information related to blockchain and cryptocurrency discussions. Punctuation marks, numbers, and non-alphabetic characters were removed to reduce textual noise, while multiple whitespace characters were replaced with single spaces to standardize sentence formatting. After preprocessing, tweets with empty textual content were discarded from the dataset. The cleaned textual data were then stored in a new feature column named clean_text and used as input for the DistilBERT sentiment classification model.

Sentiment Labeling

Sentiment labeling was performed automatically using the VADER (Valence Aware Dictionary and Sentiment Reasoner) sentiment analyzer. VADER is a lexicon-based sentiment analysis approach specifically designed for social media text and short informal sentences. The sentiment labeling process utilized the compound sentiment score generated by VADER to classify each tweet into Positive, Neutral, and Negative sentiment categories.

The sentiment classification rules used in this study are defined as follows:

Positive sentiment classification:

$$Compound\ Score \geq 0.05 \tag{1}$$

Negative sentiment classification:

$$Compound\ Score \leq -0.05 \tag{2}$$

Neutral sentiment classification:

$$-0.05 < Compound\ Score < 0.05 \tag{3}$$

Each sentiment category was represented using numerical labels:

0 = Negative

1 = Neutral

2 = Positive

Following the labeling process, sentiment distribution analysis was conducted to identify the proportion of each sentiment category within the blockchain tweet dataset.

Exploratory Data Analysis

Exploratory Data Analysis (EDA) was performed to analyze the characteristics of the blockchain tweet dataset before the model training process. Several visualization techniques were implemented to provide insights into textual patterns, tweet characteristics, and sentiment distribution within the dataset. The EDA process included sentiment distribution analysis, tweet length distribution analysis, word frequency analysis using word cloud visualization, confusion matrix visualization, normalized confusion matrix analysis, and prediction probability visualization. Sentiment distribution analysis was conducted to identify the proportion of Positive, Neutral, and Negative sentiment categories in blockchain-related Twitter discussions. Tweet length analysis was performed to determine the average number of words used in blockchain tweets and examine the textual characteristics of social media discussions. In addition, word cloud visualization was generated to identify the most frequently occurring keywords related to blockchain and cryptocurrency topics. Confusion matrix and normalized confusion matrix visualizations were utilized to evaluate classification performance for each sentiment category in detail, while prediction probability visualization was implemented to analyze the confidence level of the DistilBERT sentiment classification model.

Dataset Splitting

The dataset was divided into training data and testing data using the train-test split method. In this study, 80% of the dataset was allocated for training and 20% was allocated for testing. Stratified sampling was applied during dataset splitting to preserve the original sentiment class distribution across both training and testing datasets.

The train-test split process utilized a random state value of 42 to ensure reproducibility and consistency of experimental results. After splitting, the training dataset contained 7,757 tweets, while the testing dataset contained 1,940 tweets.

DistilBERT Model Development

This study utilized DistilBERT as the transformer-based architecture for sentiment classification. DistilBERT is a lightweight version of the Bidirectional Encoder Representations from Transformers (BERT) model developed through knowledge distillation techniques. Compared to the original BERT architecture, DistilBERT provides reduced computational complexity, lower memory usage, and faster training time while maintaining competitive language understanding capability.

The pre-trained model used in this study was distilbert-base-uncased obtained from the Hugging Face Transformers library. Before training, textual input was tokenized using the DistilBERT tokenizer to convert textual sequences into numerical token representations. The tokenization process generated input IDs and attention masks required for transformer model processing.

The tokenization process can be represented as follows:

$$X = \{x_1, x_2, x_3, \ldots, x_n\} \tag{4}$$

Xrepresents the sequence of tokenized input words.

The DistilBERT classification process is represented as follows:

$$y = DistilBERT(X) \quad (5)$$

$y$represents the predicted sentiment label generated by the transformer model.

Several tokenization parameters were applied:

Maximum sequence length: 128 tokens

Padding strategy: max_length

Truncation: enabled

The DistilBERT architecture was configured for multi-class classification with three output labels corresponding to Positive, Neutral, and Negative sentiment classes.

Model Training

The model training process was conducted using the Hugging Face Trainer API. Fine-tuning was applied to adapt the pre-trained DistilBERT model to the blockchain tweet sentiment classification task.

The Adam optimizer was used to optimize model parameters during training. Evaluation was performed at the end of each epoch to monitor model performance and prevent excessive overfitting. GPU acceleration through CUDA was utilized when available to improve training efficiency and reduce computation time.

The optimization process aimed to minimize the loss function:

$$Loss = -\sum_{i=1}^{n} y_i \log(\hat{y}_i) \quad (6)$$

$y_i$represents the actual label and $\hat{y}_i$represents the predicted probability.

The best-performing model was automatically selected based on the highest Macro F1-score obtained during evaluation.

Performance Evaluation

Model performance was evaluated using several classification metrics, including Accuracy, Precision, Recall, and Macro F1-score. These metrics were selected to provide a comprehensive assessment of sentiment classification performance, particularly because the dataset exhibited class imbalance between Positive, Neutral, and Negative sentiment categories.

Accuracy measures the proportion of correctly classified predictions among all predictions.

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \quad (7)$$

Precision measures the proportion of correctly predicted positive observations among all predicted positive observations.

$$Precision = \frac{TP}{TP + FP} \quad (8)$$

Recall measures the proportion of correctly predicted positive observations among all actual positive observations.

$$Recall = \frac{TP}{TP + FN} \tag{9}$$

F1-score represents the harmonic mean of Precision and Recall.

$$F1 = 2 \times \frac{Precision \times Recall}{Precision + Recall} \tag{10}$$

Macro F1-score was selected as the primary evaluation metric because it provides balanced performance evaluation across imbalanced sentiment classes.

In addition to numerical evaluation metrics, confusion matrix analysis and normalized confusion matrix visualization were also performed to evaluate classification performance for each sentiment category in detail.

Prediction Analysis

To further analyze the effectiveness of the proposed model, a sample tweet prediction experiment was conducted. The trained DistilBERT model generated probability scores for each sentiment class using the softmax activation function. This analysis was used to examine the confidence level of the sentiment classification results and evaluate the semantic understanding capability of the transformer-based model.

The softmax probability function is represented as follows:

$$P(y_i) = \frac{e^{z_i}}{\sum_{j=1}^{K} e^{z_j}} \tag{11}$$

$P(y_i)$represents the probability of each sentiment class and Krepresents the total number of sentiment categories.

Results

Dataset Overview

The blockchain tweet dataset initially consisted of 10,000 tweets collected from Twitter discussions related to blockchain technology and cryptocurrency topics. The dataset contained various tweet characteristics, including informal language, hashtags, hyperlinks, mentions, and duplicated content commonly found in social media data. To improve data quality and prepare the dataset for transformer-based sentiment analysis, several preprocessing techniques were applied. These preprocessing steps included converting all text into lowercase format, removing URLs, eliminating user mentions, cleaning hashtags, removing punctuation and special characters, and normalizing whitespace. After removing duplicate and missing entries, a total of 9,697 tweets were retained for further analysis.

Sentiment labeling was performed using the VADER sentiment analyzer to classify tweets into three sentiment categories: Positive, Neutral, and Negative. The sentiment distribution results indicate that blockchain-related discussions on Twitter were predominantly positive, with 7,299 positive tweets, 1,875 neutral

tweets, and 523 negative tweets. This finding suggests that most Twitter users expressed optimistic opinions regarding blockchain technology, cryptocurrency adoption, and digital financial innovation. Table 1 presents the detailed sentiment distribution of the blockchain tweet dataset.

Table 1 Sentiment Distribution of Blockchain Tweets

Sentiment Class	Number of Tweets
Positive	7,299
Neutral	1,875
Negative	523

Tweet Length Distribution

Tweet length analysis was conducted to examine the textual characteristics of blockchain-related discussions on Twitter. The analysis showed that the average tweet length was 25.44 words, with the shortest tweet containing 3 words and the longest tweet consisting of 53 words. Most tweets were concentrated within the range of 20 to 35 words, indicating that blockchain discussions on Twitter generally consist of short-to-medium text sequences. This distribution suggests that users tend to express opinions, market reactions, and cryptocurrency-related information in concise textual formats. The relatively moderate tweet length is also suitable for transformer-based language models such as DistilBERT, which are effective in capturing contextual semantic information from short social media texts. Figure 2 illustrates the distribution of tweet lengths in the blockchain tweet dataset.

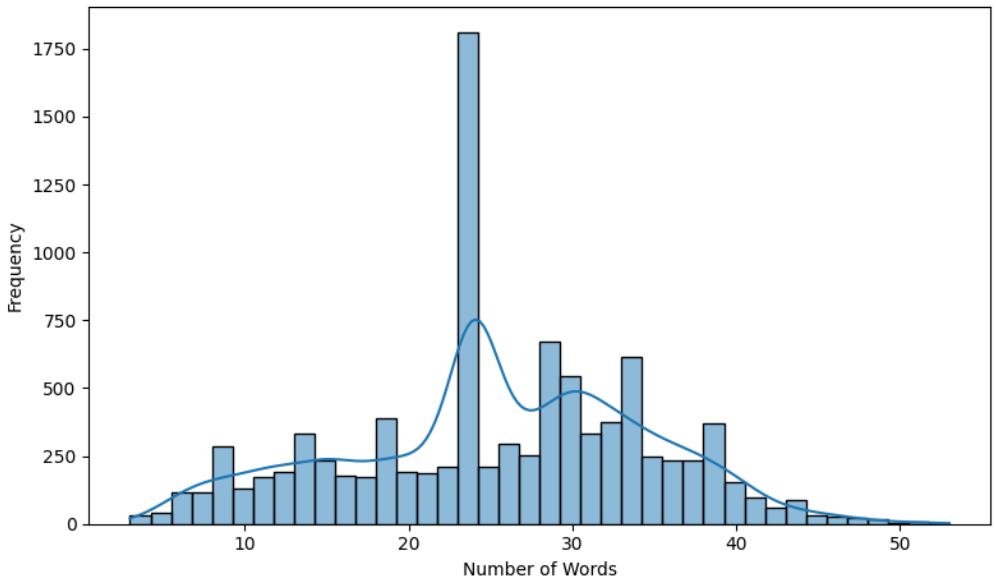


Figure 2 Distribution of Tweet Length

Word Cloud Analysis

Word cloud visualization was generated to identify the most frequently occurring words and dominant discussion topics within the blockchain tweet dataset. The visualization revealed that terms such as blockchain, crypto, cardano, transaction, fintech, project, support, and whale appeared most frequently in Twitter conversations. The prominence of these keywords indicates that public

discussions largely focused on cryptocurrency projects, blockchain-based financial technologies, digital transactions, decentralized finance activities, and investment-related trends. In addition, the frequent appearance of words related to support and projects suggests that many users expressed optimistic opinions and positive engagement toward blockchain innovation and cryptocurrency adoption. The presence of cryptocurrency-specific terminology also demonstrates the informal and domain-specific nature of blockchain discussions on social media platforms. Figure 3 presents the word cloud visualization of blockchain-related tweets.

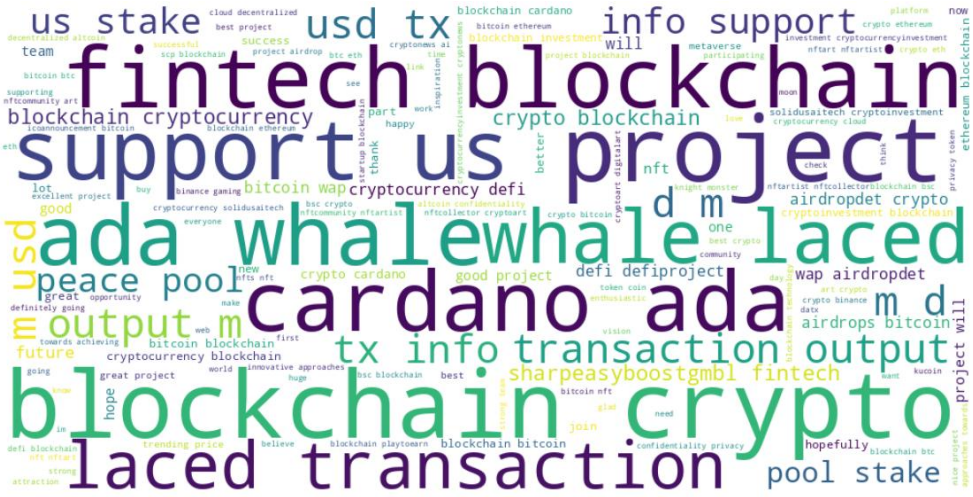


Figure 3 Word Cloud of Blockchain Tweets

Model Training Performance

The proposed sentiment classification model was developed using the DistilBERT transformer architecture to classify blockchain-related tweets into Positive, Neutral, and Negative sentiment categories. The model training process utilized a learning rate of 2e-5, a batch size of 16, and a total of three training epochs. During training, the dataset was divided into training and testing sets using stratified sampling to preserve the original sentiment distribution across classes. The training process aimed to optimize the model's ability to learn contextual semantic representations from noisy and informal Twitter text. Several evaluation metrics, including Accuracy, Macro Precision, Macro Recall, and Macro F1-score, were calculated at the end of each epoch to monitor model performance and generalization capability. The experimental results demonstrate that the DistilBERT model consistently achieved high classification performance throughout the training process. Table 2 presents the detailed model performance results for each training epoch.

Table 2 Model Performance During Training

Epoch	Validation Loss	Accuracy	Macro Precision	Macro Recall	Macro F1-score
1	0.224	92.47%	88.25%	76.20%	80.50%
2	0.209	93.61%	85.85%	85.35%	85.57%
3	0.232	93.97%	86.92%	83.88%	85.30%

The experimental results show that the best model performance was achieved during Epoch 2, where the DistilBERT model obtained a Macro F1-score of 85.57% with an accuracy of 93.61%. The improvement in Macro F1-score from Epoch 1 to Epoch 2 indicates that the model successfully learned contextual semantic patterns from blockchain-related tweets during the training process. Although the overall accuracy slightly increased during Epoch 3, the validation loss also increased from 0.209 to 0.232, while the Macro F1-score slightly decreased. This condition suggests the beginning of overfitting, where the model started to memorize training data patterns rather than improving its generalization capability on unseen data. The use of validation loss and Macro F1-score as evaluation indicators was important to monitor model stability and prevent excessive training. Figure 4 illustrates the evaluation loss trend during the model training process across all epochs.

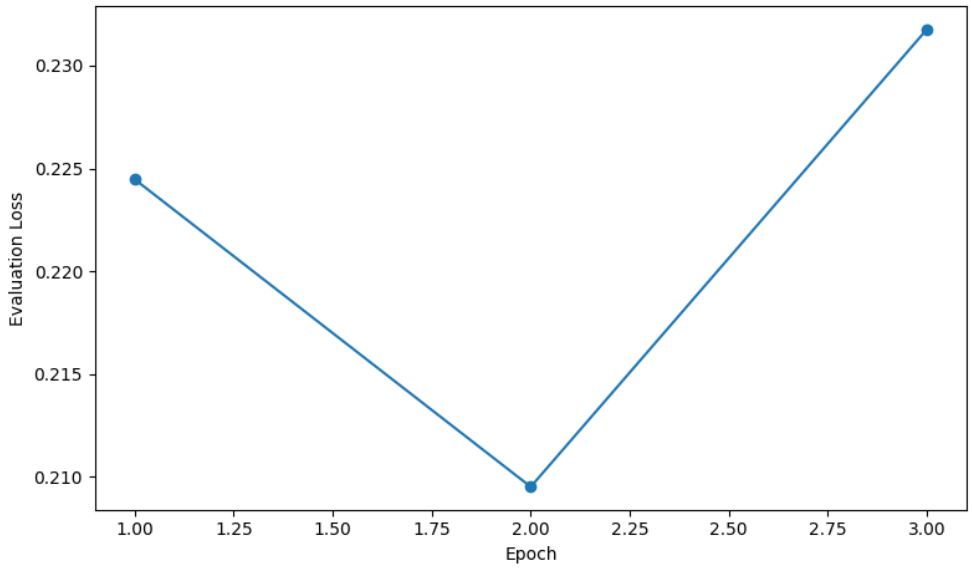


Figure 4 Evaluation Loss per Epoch

Classification Performance

The final evaluation results of the DistilBERT sentiment classification model are presented in table 3. The evaluation process was conducted using the testing dataset to measure the model's ability to classify blockchain-related tweets into Positive, Neutral, and Negative sentiment categories. Several performance metrics were used in this study, including Accuracy, Macro Precision, Macro Recall, and Macro F1-score, to provide a comprehensive assessment of classification performance. The experimental results show that the proposed model achieved an accuracy of 93.61%, a Macro Precision of 85.85%, a Macro Recall of 85.35%, and a Macro F1-score of 85.57%. These results indicate that the model was able to classify sentiment categories with a high level of accuracy and balanced predictive capability across different sentiment classes.

The strong classification performance demonstrates that the transformer-based DistilBERT architecture is highly effective for analyzing noisy and informal social media text related to blockchain and cryptocurrency discussions. The high Macro F1-score further indicates that the model maintained stable performance despite the imbalanced sentiment distribution within the dataset, where positive tweets significantly outnumbered negative tweets. In addition, the contextual

understanding capability of DistilBERT enabled the model to capture semantic relationships and domain-specific terminology commonly used in blockchain-related conversations on Twitter. Overall, the evaluation results confirm that transformer-based approaches are suitable for sentiment analysis tasks involving cryptocurrency and blockchain discussions on social media platforms.

Table 3 Final Evaluation Results

Metric	Score
Accuracy	93.61%
Macro Precision	85.85%
Macro Recall	85.35%
Macro F1-score	85.57%

Confusion Matrix Analysis

The confusion matrix analysis shows that the DistilBERT model achieved the highest classification performance on the Positive sentiment class compared to the Neutral and Negative classes. Most positive tweets were correctly classified, indicating that the model was highly effective in identifying optimistic expressions and favorable opinions related to blockchain technology and cryptocurrency discussions. The confusion matrix also reveals that a small number of Neutral and Negative tweets were incorrectly classified as Positive, which is likely influenced by the dominance of positive sentiment within the dataset. In contrast, the Negative class achieved lower classification performance due to the limited number of negative samples and the complexity of informal crypto-related language commonly used on Twitter. Despite these challenges, the model still demonstrated strong overall classification capability across all sentiment categories. Figure 5 illustrates the confusion matrix results of the proposed DistilBERT sentiment classification model.

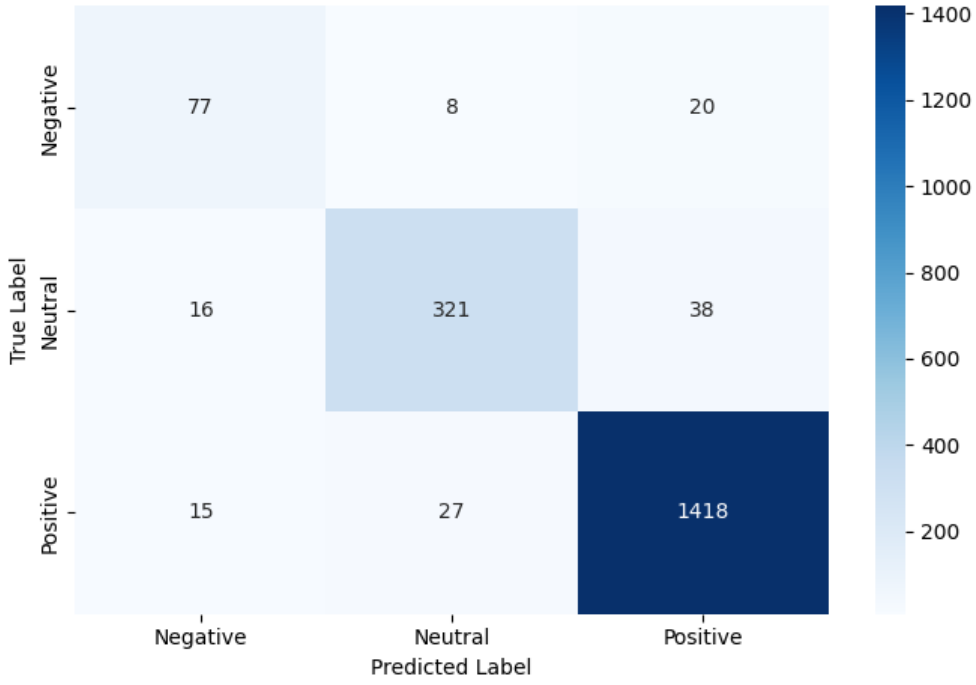


Figure 5 Confusion Matrix of DistilBERT Sentiment Classification

The normalized confusion matrix analysis provides a clearer interpretation of the classification performance for each sentiment category by presenting the recall values in percentage form. The results show that the Positive sentiment class achieved the highest recall score of 97%, indicating that the majority of positive blockchain-related tweets were correctly identified by the DistilBERT model. The Neutral class obtained a recall score of 86%, demonstrating that the model was also effective in recognizing tweets containing balanced or informational opinions. Meanwhile, the Negative class achieved a recall score of 73%, which was lower compared to the other sentiment categories. This lower performance was likely influenced by the imbalance of the dataset and the relatively small number of negative tweet samples available during training. In addition, negative cryptocurrency-related tweets often contain ambiguous expressions, slang, and informal language that increase classification difficulty. Overall, the normalized confusion matrix demonstrates that the proposed model maintained strong and balanced sentiment classification performance across all classes. Figure 6 shows the normalized confusion matrix of the DistilBERT sentiment classification model.

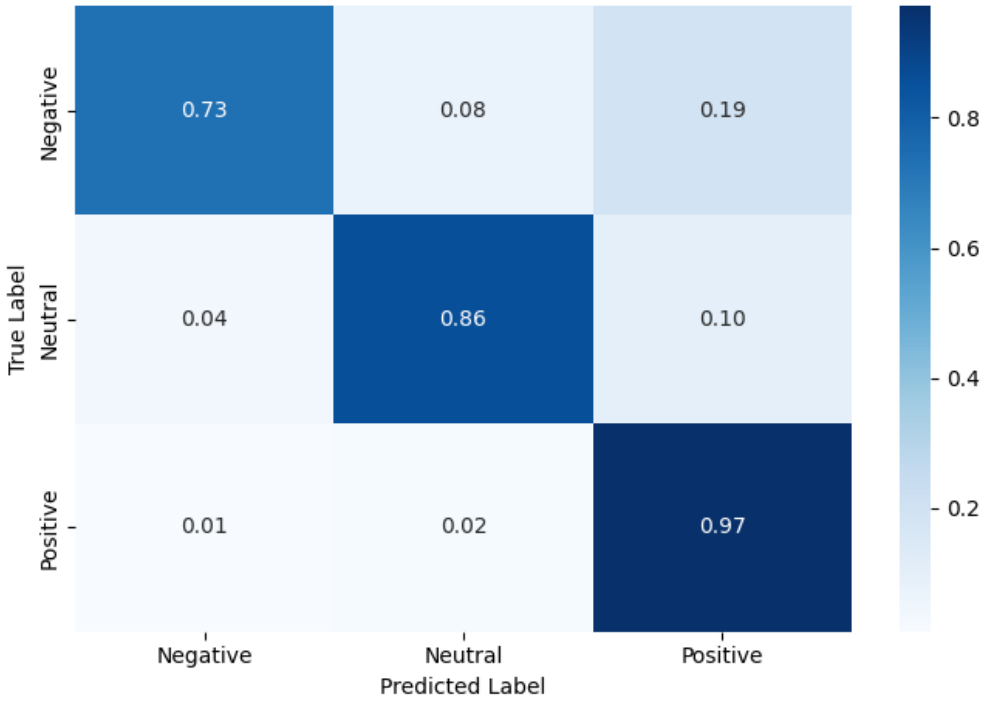


Figure 6 Normalized Confusion Matrix

Prediction Probability Analysis

Prediction probability analysis was conducted using a sample blockchain-related tweet to evaluate the confidence level of the DistilBERT sentiment classification model. The sample tweet, “Blockchain technology is revolutionizing finance and crypto adoption is increasing fast,” was processed through the trained model to obtain sentiment probability scores for each class. The experimental results show that the model predicted a Positive sentiment probability of 77.53%, a Neutral probability of 19.31%, and a Negative probability of 3.15%. Based on these probability values, the model classified the tweet as Positive sentiment. The high positive probability indicates that

DistilBERT successfully captured contextual semantic patterns associated with blockchain innovation, financial technology development, and increasing cryptocurrency adoption. In addition, the low negative probability demonstrates that the model effectively distinguished optimistic blockchain-related expressions from negative or neutral sentiment patterns. These results suggest that the proposed transformer-based model was capable of generating confident and interpretable sentiment predictions for blockchain discussions on social media platforms. Figure 7 presents the sentiment prediction probabilities generated by the DistilBERT model.

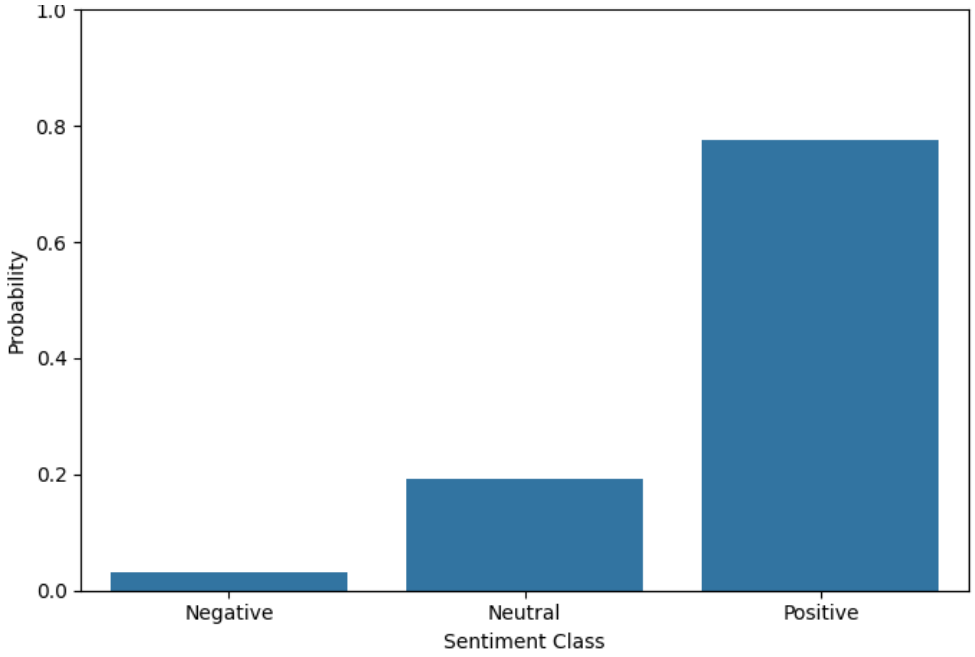


Figure 7 Sentiment Prediction Probability

Discussion

The experimental results demonstrate that the proposed DistilBERT model achieved strong performance in classifying blockchain-related tweets into Positive, Neutral, and Negative sentiment categories. The model obtained an accuracy of 93.61% and a Macro F1-score of 85.57%, indicating that transformer-based architectures are highly effective for sentiment analysis tasks involving noisy and informal social media text. This finding is consistent with previous studies showing that transformer-based models such as BERT, DistilBERT, RoBERTa, and other deep learning architectures can effectively improve sentiment classification performance on Twitter data and other short-text social media datasets [21], [22]. The strong classification performance suggests that DistilBERT successfully learned contextual semantic representations from blockchain-related discussions despite the presence of slang, abbreviations, hashtags, and cryptocurrency-specific terminology commonly found on Twitter. Unlike traditional machine learning approaches that rely heavily on manual feature extraction techniques, DistilBERT utilizes contextual embeddings and self-attention mechanisms to better understand semantic relationships between words within a sentence. This capability allowed the model to effectively interpret short-text discussions related to cryptocurrency investment, blockchain innovation, decentralized finance, and digital asset

adoption. The experimental findings also indicate that transformer-based models are highly suitable for analyzing social media data where textual patterns are often unstructured and highly dynamic [23].

The sentiment distribution analysis revealed that blockchain-related Twitter discussions were predominantly positive, with positive tweets significantly outnumbering neutral and negative tweets. This finding suggests that Twitter users generally express optimistic opinions regarding blockchain technology, cryptocurrency projects, financial innovation, and digital investment opportunities. Similar findings have been reported in previous cryptocurrency sentiment analysis studies, where social media discussions related to blockchain and cryptocurrency markets often showed a dominant positive sentiment orientation [24]. The confusion matrix analysis further demonstrated that the Positive sentiment class achieved the highest recall score of 97%, indicating that the model was highly effective in recognizing supportive and optimistic expressions related to blockchain discussions. In contrast, the Negative class achieved lower classification performance due to the imbalance of the dataset and the relatively limited number of negative tweet samples available during training. In addition, negative cryptocurrency-related tweets often contain sarcasm, ambiguous wording, and informal language patterns that increase classification difficulty. Although the model demonstrated strong overall performance, several limitations were identified in this study. The sentiment labeling process was automatically generated using the VADER lexicon-based approach rather than manual human annotation, which may introduce labeling noise and reduce annotation accuracy. Previous studies have emphasized that lexicon-based sentiment labeling can provide useful silver-standard labels for large-scale Twitter analysis, but manual annotation or hybrid validation is still important to improve label reliability [25]. Future research may improve this framework by using manually annotated datasets, applying data balancing techniques, and comparing multiple transformer architectures such as BERT, RoBERTa, and FinBERT for blockchain sentiment analysis tasks.

Conclusion

This study proposed a transformer-based sentiment analysis framework using the DistilBERT model to classify blockchain-related tweets into Positive, Neutral, and Negative sentiment categories. The experimental results demonstrated that the proposed model achieved strong classification performance, obtaining an accuracy of 93.61% and a Macro F1-score of 85.57%. These results indicate that DistilBERT was highly effective in capturing contextual semantic information from noisy and informal social media text related to blockchain and cryptocurrency discussions. The preprocessing techniques applied in this study, including lowercase transformation, URL removal, mention removal, hashtag cleaning, punctuation removal, and whitespace normalization, contributed to improving the quality of the textual data and enhancing the overall model performance. In addition, the sentiment distribution analysis revealed that blockchain-related discussions on Twitter were predominantly positive, indicating that users generally expressed optimistic opinions regarding blockchain technology, cryptocurrency adoption, decentralized finance, and digital financial innovation. The confusion matrix analysis further showed that the Positive sentiment class achieved the highest classification performance, while the Negative class remained more challenging due to dataset imbalance, limited negative samples, and the complexity of informal cryptocurrency-related

language used on Twitter. Despite these limitations, the proposed DistilBERT model demonstrated stable and reliable sentiment classification capability across all sentiment categories. Overall, the findings confirm that transformer-based approaches are highly suitable for blockchain sentiment analysis tasks and have strong potential for practical applications in cryptocurrency market monitoring, public opinion analysis, blockchain trend analysis, and financial decision support systems. Future research may improve this framework by utilizing manually annotated datasets, implementing data balancing techniques, and comparing multiple transformer architectures such as BERT, RoBERTa, and FinBERT to further enhance sentiment classification performance on blockchain-related social media data.

Declarations

Author Contributions

Conceptualization: R.A.; Methodology: R.A.; Software: R.A.; Validation: R.A.; Formal Analysis: R.A.; Investigation: R.A.; Resources: R.A.; Data Curation: R.A.; Writing Original Draft Preparation: R.A.; Writing Review and Editing: R.A.; Visualization: R.A.; All authors have read and agreed to the published version of the manuscript.

Data Availability Statement

The data presented in this study are available on request from the corresponding author.

Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

Institutional Review Board Statement

Not applicable.

Informed Consent Statement

Not applicable.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

[1] M. Swan, *Blockchain: Blueprint for a New Economy*. Sebastopol, CA, USA: O'Reilly Media, 2015.

[2] A. M. Antonopoulos, *Mastering Bitcoin: Programming the Open Blockchain*, 2nd ed. Sebastopol, CA, USA: O'Reilly Media, 2017.

[3] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," 2008. [Online]. Available: https://bitcoin.org/bitcoin.pdf. Accessed: May 21, 2026.

[4] 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, Volume 1 (Long and Short Papers)*,

Minneapolis, MN, USA, Jun. 2019, vol. 2019, no. June, pp. 4171–4186, doi: 10.18653/v1/N19-1423.

[5] V. Sanh, L. Debut, J. Chaumond, and T. Wolf, “DistilBERT, a distilled version of BERT: Smaller, faster, cheaper and lighter,” *arXiv preprint* arXiv:1910.01108, vol. 2019, no. October, pp. 1-5, 2019, doi: 10.48550/arXiv.1910.01108.

[6] B. Liu, *Sentiment Analysis and Opinion Mining*. San Rafael, CA, USA: Morgan & Claypool Publishers, 2012, doi: 10.2200/S00416ED1V01Y201204HLT016.

[7] E. Cambria, D. Das, S. Bandyopadhyay, and A. Feraco, Eds., *A Practical Guide to Sentiment Analysis*. Cham, Switzerland: Springer, 2017, doi: 10.1007/978-3-319-55394-8.

[8] A. Pak and P. Paroubek, “Twitter as a corpus for sentiment analysis and opinion mining,” in *Proceedings of the 7th International Conference on Language Resources and Evaluation (LREC)*, Valletta, Malta, May 2010, pp. 1320–1326.

[9] B. Pang and L. Lee, “Opinion mining and sentiment analysis,” *Foundations and Trends in Information Retrieval*, vol. 2, nos. 1–2, pp. 1–135, 2008, doi: 10.1561/1500000011.

[10] T. Chen and C. Guestrin, “XGBoost: A scalable tree boosting system,” in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD ’16)*, San Francisco, CA, USA, Aug. 2016, vol. 2016, no. August, pp. 785–794, doi: 10.1145/2939672.2939785.

[11] A. Vaswani et al., “Attention is all you need,” in *Advances in Neural Information Processing Systems 30 (NIPS 2017)*, Long Beach, CA, USA, 2017, vol. 2017, no. June, pp. 5998–6008.

[12] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997, doi: 10.1162/neco.1997.9.8.1735.

[13] Y. Kim, “Convolutional neural networks for sentence classification,” in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Doha, Qatar, Oct. 2014, vol. 2014, no. October, pp. 1746–1751, doi: 10.3115/v1/D14-1181.

[14] N. Reimers and I. Gurevych, “Sentence-BERT: Sentence embeddings using Siamese BERT-networks,” in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, Hong Kong, China, Nov. 2019, vol. 2019, no. November, pp. 3982–3992, doi: 10.18653/v1/D19-1410.

[15] M. Abadeer, “Assessment of DistilBERT performance on named entity recognition task for the detection of protected health information and medical concepts,” in *Proceedings of the 3rd Clinical Natural Language Processing Workshop*, Online, Nov. 2020, vol. 2020, no. November, pp. 158–167, doi: 10.18653/v1/2020.clinicalnlp-1.18.

[16] Z. Dai, Z. Yang, Y. Yang, J. Carbonell, Q. Le, and R. Salakhutdinov, “Transformer-XL: Attentive language models beyond a fixed-length context,” in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Florence, Italy, Jul. 2019, vol. 2019, no. July, pp. 2978–2988, doi: 10.18653/v1/P19-1285.

[17] M. T. Ribeiro, S. Singh, and C. Guestrin, “‘Why should I trust you?’: Explaining the

predictions of any classifier," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '16)*, San Francisco, CA, USA, Aug. 2016, vol. 2016, no. August, pp. 1135–1144, doi: 10.1145/2939672.2939778.

[18] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, "Distributed representations of words and phrases and their compositionality," in *Advances in Neural Information Processing Systems 26 (NIPS 2013)*, Lake Tahoe, NV, USA, 2013, vol. 2013, no. October, pp. 3111–3119.

[19] J. Pennington, R. Socher, and C. Manning, "GloVe: Global vectors for word representation," in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Doha, Qatar, Oct. 2014, vol. 2014, no. October, pp. 1532–1543, doi: 10.3115/v1/D14-1162.

[20] A. Rogers, O. Kovaleva, and A. Rumshisky, "A primer in BERTology: What we know about how BERT works," *Transactions of the Association for Computational Linguistics*, vol. 8, no. January, pp. 842–866, 2020, doi: 10.1162/tacl_a_00349.

[21] R. Agrawal, M. Majumder, I. Yadav, N. Taneja, S. Hamdare, and P. Hemnani, "Evaluating sentiment analysis models: A comparative analysis of vaccination tweets during the COVID-19 phase leveraging DistilBERT for enhanced insights," *MethodsX*, vol. 14, no. June, Art. no. 103407, 2025, doi: 10.1016/j.mex.2025.103407.

[22] V. Amiri and M. Ahmadi, "Sentiment analysis in Twitter social network centered on cryptocurrencies using machine learning," *arXiv preprint* arXiv:2501.09777, vol. 2025, no. January, pp. 1-6, 2025, doi: 10.48550/arXiv.2501.09777.

[23] E. K. E. Laison, M. Hamza Ibrahim, S. Boligarla, J. Li, R. Mahadevan, A. Ng, V. Muthuramalingam, W. Y. Lee, Y. Yin, and B. R. Nasri, "Identifying potential Lyme disease cases using self-reported worldwide tweets: Deep learning modeling approach enhanced with sentimental words through emojis," *Journal of Medical Internet Research*, vol. 25, no. October, Art. no. e47014, 2023, doi: 10.2196/47014.

[24] D. Tiwari, B. S. Bhati, B. Nagpal, N. Alturki, and L. Bayisenge, "Attention-augmented hybrid CNN-LSTM model for social media sentiment analysis in cryptocurrency investment decision-making," *Scientific Reports*, vol. 15, no. September, Art. no. 33201, 2025, doi: 10.1038/s41598-025-18245-x.

[25] A. Mardjo and C. Choksuchat, "HyVADRF: Hybrid VADER–random forest and GWO for Bitcoin tweet sentiment analysis," *IEEE Access*, vol. 10, no. September, pp. 101889–101897, 2022, doi: 10.1109/ACCESS.2022.3209662.