Exploring Crypto Narratives: Developing Approaches for Opinion Extraction and Question Answering in Social Media Text

Asha Hegde[1], Sharal Coelho[1] and Krishna Tewari[2]

[1]*Department of Computer Science, Mangalore University, Mangalore*
[2]*Indian Institute of Technology (BHU), Varanasi*

Abstract

In the dynamic landscape of cryptocurrency, understanding public sentiment is essential for tracking market trends and investor behavior. Social media platforms like Twitter, Reddit, YouTube and Facebook have become key spaces where users express diverse and often unfiltered opinions about cryptocurrencies. However, the informal and unstructured nature of social media content presents significant challenges for accurate sentiment analysis. Posts vary widely in tone, language, and format, making it difficult to categorize sentiments with precision. To tackle these complexities, the shared task "CryptOQA: Opinion Extraction and Question Answering from Crypto Currency-Related Tweets and Reddit posts" organized at FIRE 2025, calls on the research community to develop innovative solutions. The goal is to enhance the identification and classification of cryptocurrency-related opinions and questions expressed in English across social platforms. To address these challenges, this paper introduces Multi-task Learning (MTL) models that leverage pretrained transformers to effectively capture the hierarchical structure of data for Task 1. For Task 2, we implement question-answering frameworks built upon pretrained transformer models to understand and classify cryptocurrency-related questions from social media posts. Our best-performing model for the opinion classification task achieved macro F1 scores of 0.5694, 0.5570, and 0.7496 on the Reddit, Twitter, and YouTube datasets, respectively, as provided by the task organizers. Additionally, our top model for Task 2 attained a macro F1 score of 0.7940, demonstrating the effectiveness of our approach in handling diverse and unstructured social media content.

Keywords

Multi-task learning, Relevance classification, Opinion mining, Question-Answering

1. Introduction

Cryptocurrency is a paradigm change in the world of finance, defined by digital or virtual money that uses cryptographic methods to secure transactions [1]. In contrast to the standard currencies issued by governments, cryptocurrencies use decentralized networks based on blockchain technology [2]. Most prominent cryptocurrencies like Bitcoin, Ethereum, and Litecoin have gained common knowledge and usage, becoming major players in the world of finance. As the popularity of these digital assets increases, so does the value of recognizing public opinion/sentiments surrounding them [3].

Social media websites have emerged as the primary source of information and public sentiment, giving immediate insights into attitudes and behavior [4]. Websites like Twitter, Reddit, and Facebook are live forums where users post their opinions on a vast array of subjects, including cryptocurrencies. What people say on these sites has the power to immensely impact perceptions and trends in the market, and therefore it is significant to watch and study the opinion being voiced in social media discourse [5]. Sentiment classification in cryptocurrency posts is a reflection of the general opinion of the public and enables stakeholders to make informed decisions. Properly identifying opinions as positive, negative, neutral, or objective can help investors forecast market trends, identify consumer behavior, and create effective marketing campaigns. Sentiment classification in this setting comes with a number of challenges because social media content is diverse and unstructured [6]. One of the greatest challenges stems from the variability and vagueness of social media language. Due to its short

Forum for Information Retrieval Evaluation, December 17-20, 2025, Varanasi, India
hegdekasha@gmail.com (A. Hegde); sharalmucs@gmail.com (S. Coelho); krishnatewari.rs.cse24@itbhu.ac.in (K. Tewari)
0000-0002-6382-4096 (A. Hegde); 0000-0003-2973-5871 (S. Coelho); 0009-0005-6599-9956 (K. Tewari)

nature, carefreeness, and abundance of slang, it is difficult for conventional text classification algorithms to extract the underlying sentiments [7, 8]. Also, the environment in which opinions are presented may vary considerably, adding another level of complexity to the analysis. These considerations underscore the imperative for sophisticated models with the ability to manage social media language complexity. Sophisticated methods must be applied to properly classify opinions and draw meaningful conclusions from the data. Generally speaking, knowledge of social media discourse is imperative for stakeholders seeking to master the volatile cryptocurrency space. In addition, on social media, question-answering discussions surrounding cryptocurrency are often found in the form of users asking for advice, explanations, or views regarding investments, market conditions, or the operations of particular crypto assets [9]. These discussions are mostly casual, unstructured, and immensely varied in terms of tone and language, which makes it difficult to decipher and derive useful insights from them. It is normal for possible investors or curious minds to ask questions, and these are responded to by other users via comments that are either not relevant or wrong. Classification and categorization of these interactions are important because they assist in weeding out false information and finding truly useful answers [10]. Creating learning models that have the ability to effectively identify and map appropriate answers to user questions is crucial for constructing intelligent systems capable of guiding users through the intricate cryptocurrecy world. These models improve information retrieval, aid in the development of knowledge bases related to cryptocurrency, and assist with making informed decisions.

To tackle the challenges of identifying cryptocurrency-related opinion posts on social media, the "CryptOQA: Opinion Extraction and Question Answering from CryptoCurrency-Related Tweets and Reddit posts" shared task was organized at FIRE 2025. This initiative invited the research community to develop models capable of detecting different categories of cryptocurrency posts. The shared task included two components: Task 1, focused on opinion classification from cryptocurrency-related social media posts, and Task 2, centered on Question Answering (QA) from such posts. The dataset provided by the organizers comprised 5,000 posts per platform and was annotated across three hierarchical levels, as illustrated in Figure 1. This structured dataset offered a foundation for developing models to accurately capture the nuances of cryptocurrency-related opinions expressed on social media. To address the challenges provided by the shared task, MTL models that utilize pretrained transformers to capture the hierarchical structure of the given data are implemented. Two separate runs are submitted to the organizers, each based on different transformer models: BERT and CryptoBERT. Among these, the model fine-tuned using CryptoBERT achieved the best performance across the provided datasets, demonstrating its effectiveness in handling cryptocurrency-related content. For Task 2, we developed two question-answering frameworks built on BERT and DistilBERT to identify and classify cryptocurrency-related questions from social media posts. These models performed relevance classification to determine the relationship between questions and candidate answers. Overall, the proposed approaches show promising results in both opinion categorization and question-answering tasks.

The rest of the paper is organized as follows: Section 2 sheds light on related works, and methodology is discussed in Section 3. Experimental results are presented in Section 4 and the paper concludes in Section 5 along with spreading light on future works.

2. Related Work

The market for cryptocurrencies has grown very fast, presenting challenges and benefits for machine learning-style data analysis. To have any success in this extremely volatile market, it is crucial to understand the sentiment of the public and predict market directions. Numerous studies have shown that machine learning, particularly deep learning, can be used to address this issue, with a focus on sentiment analysis of social media posts. A number of research works have investigated the role of social media sentiment in cryptocurrency price forecasting with encouraging outcomes. Among such works, one presented an MAPE of 1.3251 as evidence of its suitability for predicting the price of Bitcoin (BTC). Sahal [11] introduced a method that combines Bidirectional LSTM (BiLSTM) networks and Embeddings from Language Models (ELMo) in examining Twitter data to determine profitable cryptocurrencies

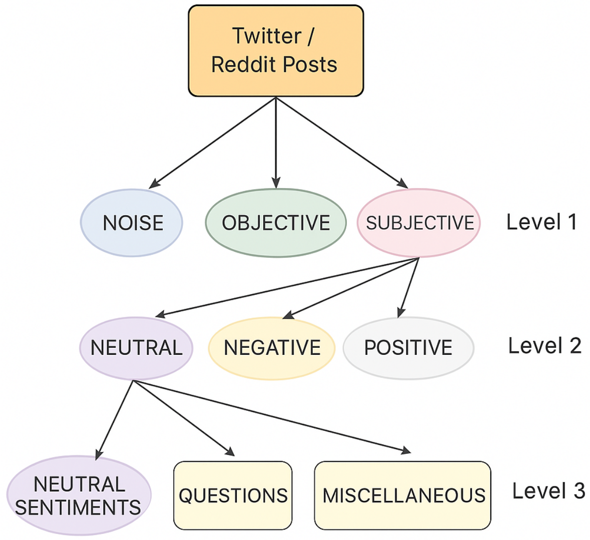

Figure 1: Class hierarchy for Task 1: Categorizing opinions from cryptocurrency-related social media posts.

using contextual sentiment analysis. Their model was able to produce a high accuracy of 86.30%, demonstrating its capability in forecasting price variations affected by social media. In another highly cited paper, Huang et al. [12] looked into sentiment analysis on Sina-Weibo, a well-known Chinese social networking site, to forecast cryptocurrency price volatility. They constructed a crypto-oriented sentiment dictionary and used an LSTM-based Recurrent Neural Network (RNN) and autoregressive models, incorporating sentiment with past price information. The LSTM model performed better than conventional autoregressive models, recording a precision of 0.87 and a recall of 0.94, testifying to its performance in reflecting market movements.

Das et al. [13] proposed an explanatory multitasking model for analysis in financial markets by combining sentiment, emotion, and cause extraction from social media text. To this end, Hegde and Shashirekha [14] constructed the FinEMA dataset by merging existing financial news/tweet datasets with 6,500 newly scraped and manually labeled tweets containing sentiment, emotion, and cause annotations. Their methodology utilizes FinBERT embeddings and a hierarchical Emotion-Sentiment Attention Network (ESAN), supplemented with a CentralNet-based multitasking framework, for modeling sentiment classification, emotion detection, and cause identification jointly. The experimental results indicate that ESAN attained 89% accuracy for sentiment classification and 79% for emotion detection, surpassing baselines like BERT, CNN-BiLSTM, and ensemble LSTM-GRU models.

The associated research emphasizes that scholars have used various methods, such as LSTM-GRU ensembles, self-attention mechanisms, and hierarchical graph neural networks, to examine cryptocurrency sentiments. These models make use of varied features and information sources, such as tweets and financial metrics, to analyze intricate market relationships. Although, they have exhibited robust performance in sentiment classification and trend identification, the highly volatile nature of cryptocurrency markets introduces new challenges. In addition, volatile and loud social media content makes the identification of sentiments even more challenging. Therefore, considerable potential exists to develop methodologies and find novel techniques to comprehend cryptocurrency market dynamics.

3. Methodology

The proposed approach introduces two distinct methodologies for Task 1 and Task 2. For Task 1, a hierarchical classification model is built in order to structure labels in a multi-level format, in which there can be more specific and context-sensitive classification by taking into account relationships at various label levels. For Task 2, a relevance classification model is used to train the correspondence between a question and its respective answers, categorizing whether an answer is relevant or not with respect to the question. These approaches collectively try to enhance both answer relevance detection and label organization in the given tasks. It may be noted that common pre-processing for both Task 1 and Task 2 is employed in this paper. The description of each of these is provided in the below subsections (3.2 and 3.3).

3.1. Pre-processing

Pre-processing operation is applied to remove noise from and standardize social media text for opinion extraction and QA [15]. A pre-processing pipeline is implemented that begins with the elimination of user mentions (e.g., *@username*) that appear frequently in social media updates but contribute negligibly to semantic meaning. The text is then converted to lowercase to ensure consistency and eliminate redundancy in word representation. Character-wise normalization is also performed to achieve standardized formatting across the dataset.

3.2. Multi-task Learning Model for Task 1

MTL is a machine learning paradigm in which a model learns to undertake a number of associated tasks simultaneously by sharing some layers and parameters. This approach is particularly beneficial when there is not enough data for individual but similar tasks. MTL allows the model to learn common lower-level features while still acquiring task-specific higher-level features. With an integrated framework, it enhances generalization and increases prediction accuracy [16].

The model for Task 1 proposed in this study is a typical MTL design and uses pre-trained transformer models to process input sequences and derive text representations that have meaningful contextual and semantic information [17, 18]. These representations of context and semantics are then pooled across tasks and used for task-specific classification, as shown in Figure 2 (a). By transforming text to feature vectors, the model is able to discern patterns and relationships more effectively, resulting in better accuracy in classification. Two transformer-based models, namely BERT and CryptoBERT, are employed in this work, with Hugging Face's 'BertTokenizer' and 'BertModel' for tokenization and loading pre-trained models. Description of each of these transformer-based models is given below:

- **CryptoBERT** - is a transformer-based model that has been borrowed from BERT and fine-tuned for cryptocurrency text data. The model is specifically tailored to understand the distinctive vocabulary, jargon, and discourse tendencies present in cryptocurrency conversations on social media platforms such as Twitter and Reddit. This targeted training enables CryptoBERT to grasp sentiment and opinion more effectively in the crypto space, enhancing its performance in tasks including cryptocurrency sentiment prediction and opinion mining.
- **BERT** - is a Google-developed pre-trained language model that has revolutionized natural language processing (NLP). It works to comprehend the context of words by taking into account their left and right contexts via a bidirectional attention mechanism. BERT can be used for various applications, including text classification, QA, and sentiment analysis, as it offers rich contextual embeddings.

As illustrated in Figure 2 (a), the MTL architecture under consideration utilizes contextualized representations produced by transformer-based models to process the input sequence. These are then fed to two distinct Recurrent Neural Network (RNN) modules that specialize in processing a given subtask. The embeddings coming out of the RNNs are utilized to produce probability distributions for target labels of a given task. For improved learning, the outputs are fed through a Feed Forward Network (FFN), which adds non-linearity and better picks up the more intricate patterns within the data. The produced logits are further processed using a softmax layer for translation into probability values for the final predictions. The total loss L is the weighted combination of individual losses for each subtask, $L = \sum_{i=1}^{I} w_i L_i$, where I is the number of labels and w_i is the associated loss weight.

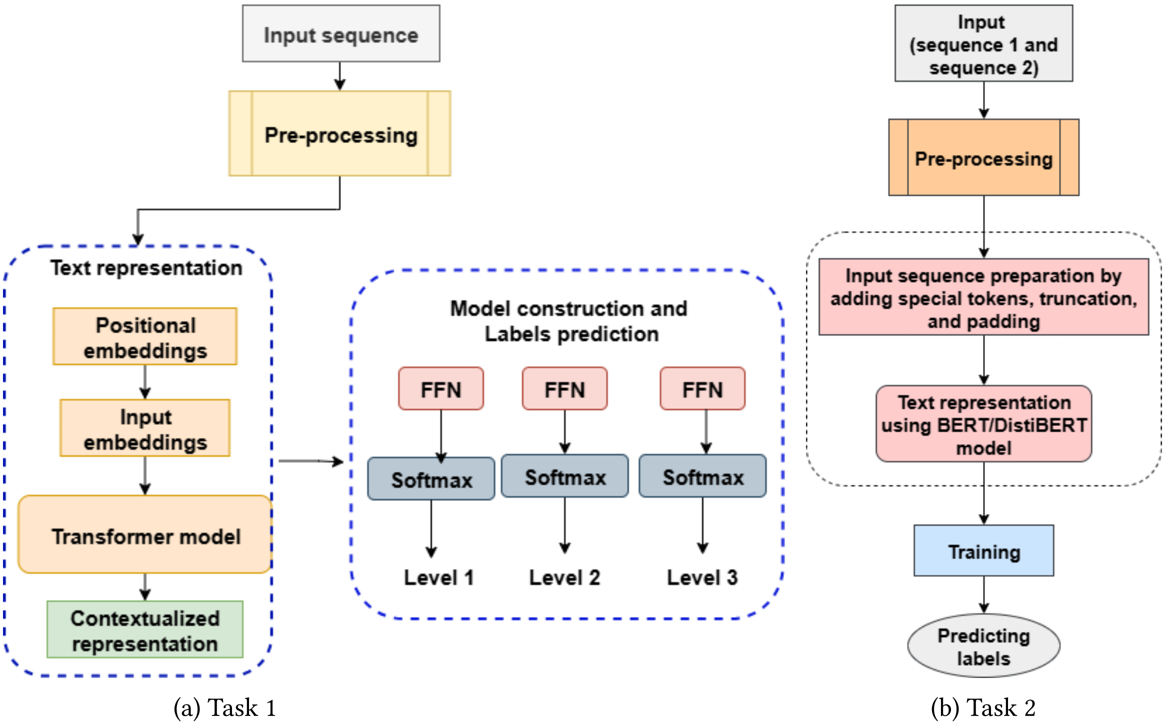

Figure 2: Framework of the proposed multi-task learning model for Task 1 and relevance classification model for Task 2

This architecture allows the model to share feature vectors among tasks to enhance generalization and task-specific performance.

3.3. Relevance Classification for Task 2

In QA, relevance classification is responsible for classifying a candidate text (e.g., answer or passage) as relevant or irrelevant to a given question, often applying binary or graded labels [19]. Relevance classification is important in filtering and ranking out content that actually answers the question, thus enhancing answer quality and user experience. For example, in Community QA settings, numerous responses can be of low quality or unrelated, so relevance measures become crucial for useful answers' identification. Work on community QA demonstrates how relevance (as well as measurements such as similarity and completeness) is central to evaluating candidate answers' quality [20].

In this study, we tackle the task of relevance classification between a question and its respective comment or answer in social media text on cryptocurrency and the proposed framework is shown in Figure 2 (b). The dataset has three major columns: MAIN (the question), comment_body (the comment/answer), and relevance (binary labels on whether the question-answer pair is relevant). To preprocess the data, text from both the MAIN and comment_body columns was cleaned (as mentioned in Subsection 3.1). Further, all values are transformed into strings and missing entries are replaced with empty strings. The labels were also standardized into integers for consistency in training. For training and evaluation purposes, the dataset was transformed into Hugging Face Dataset objects to ensure compatibility with the transformers library and enable efficient tokenization and batching. This configuration enabled us to frame the task as a binary classification problem in which the model predicts whether a comment pertains to a particular question.

In order to encode the text pairs, we employed pre-trained transformer-based models (DistilBERT and BERT models), which yields contextual embeddings for sequence classification tasks. Description of DistilBERT is given below (description of BERT model is provided in Subsection 3.2):

- **DistilBERT** - is a lower, quicker, and lighter alternative to BERT, developed through knowledge distillation to capture nearly all of BERT's performance but with much lower computational costs. It has approximately 40% fewer parameters than BERT and processes around 60% more quickly, for which it is particularly suited for real-time or for resource-limited applications. In spite of its small size, DistilBERT retains about 97% of BERT's language comprehension ability, with excellent performance across a broad variety of natural language processing tasks including text classification, sentiment analysis, and question answering. Its trade-off of efficiency and accuracy has earned it a coveted spot among both academic research and industry applications.

We passed MAIN (question) and comment_body (answer) concurrently to the tokenizer, with MAIN being the main sequence and comment_body as the paired sequence. Tokenization, truncation, and padding were accomplished by the Hugging Face tokenizer, converting input text pairs into input_ids and attention_masks that are ingestible by the model. Transformer models took these inputs and processed them to create hidden representations, which were fed into a classification head with a dense layer that produced binary labels.

Table 1
Hyperparameters used for training the MTL model

Hyperparameter	Value
Model	BertMultitaskModel
Optimizer	AdamW
Batch Size	16
Epochs	20
Shuffle	True
Learning Rate	5×10^{-5}
Loss Function	CrossEntropyLoss

Table 2
Hyperparameters used for relevance classification model

Hyperparameter	Value
Number of Epochs	20
Warmup Steps	500
Maximum Sequence Length	128
Batch Size (Train/Eval)	16
Weight Decay	0.01
Learning Rate	5×10^{-5}
Save Strategy	epoch
Evaluation Strategy	epoch
Optimizer	AdamW
Loss Function	CrossEntropyLoss

4. Experimental Results

A set of experiments were conducted across various transformer models to solve Task 1 using the MTL framework and Task 2 using relevance classification models. Hyperparameters and their respective values applied to both models are shown in Tables 1 and 2. Models with excellent performance on the validation sets were also tested on the test sets. Since the datasets are not balanced, the Macro F1 score was adopted as the metric of evaluation, and results on the performance of the proposed models are presented in Table 3.

According to the shared task guidelines, only predictions of two models from the Test sets could be submitted by participants. We thus trained our models on the given Train sets and obtained predictions on the Test sets provided by the organizers. The organizers used the macro F1 score as the measure of performance and evaluated our predictions. Our suggested MTL model, which uses CryptoBERT

Table 3
Performance of the proposed models for Task 1 and Task 2, along with the corresponding ranks obtained in the CryptOQA shared task

<table>
<tr><th colspan="8">Task 1</th></tr>
<tr><th rowspan="2">Model</th><th rowspan="2">Dataset</th><th colspan="3">Macro F1 score</th><th rowspan="2">Average score</th><th rowspan="2">Datasetwise Rank</th><th rowspan="2">Taskwise Rank</th></tr>
<tr><th>Level 1</th><th>Level 2</th><th>Level 3</th></tr>
<tr><td rowspan="3">CryptoBERT</td><td>Reddit</td><td>0.8205</td><td>0.5934</td><td>0.2943</td><td rowspan="3">0.6352</td><td>5</td><td rowspan="3">4</td></tr>
<tr><td>Twitter</td><td>0.6081</td><td>0.7269</td><td>0.3361</td><td>5</td></tr>
<tr><td>YouTube</td><td>0.8619</td><td>0.7476</td><td>0.6394</td><td>3</td></tr>
<tr><td rowspan="3">BERT</td><td>Reddit</td><td>0.8419</td><td>0.5243</td><td>0.3278</td><td rowspan="3">0.6295</td><td>-</td><td>-</td></tr>
<tr><td>Twitter</td><td>0.612</td><td>0.7449</td><td>0.2702</td><td>-</td><td>-</td></tr>
<tr><td>YouTube</td><td>0.8452</td><td>0.7636</td><td>0.6321</td><td>-</td><td>-</td></tr>
<tr><th colspan="8">Task 2</th></tr>
<tr><th>Model</th><th colspan="4">Macro F1 score</th><th colspan="3">Rank</th></tr>
<tr><td>DistilBERT</td><td colspan="4">0.7242</td><td colspan="3">-</td></tr>
<tr><td>BERT</td><td colspan="4">0.794</td><td colspan="3">3</td></tr>
</table>

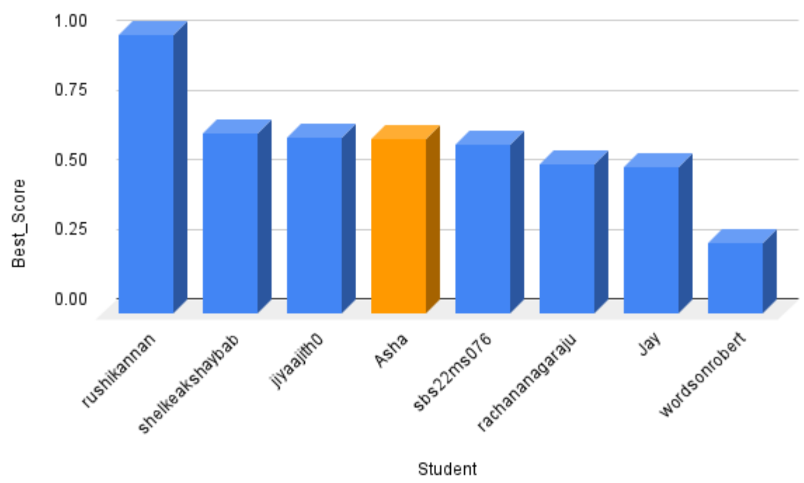

Figure 3: Comparison of macro F1 scores of all participants in Task 1

for text representation, posted competitive performance with macro F1 values of 0.5694, 0.5570, and 0.749, with 5th, 5th, and 3rd ranks on the Reddit, Twitter, and YouTube datasets, respectively, for Task 1. Globally, this model posted the 4th rank in Task 1. For Task 2, the suggested relevance classification model using BERT for text representation posted good macro F1 score of 0.7940 with 3rd rank. A comparison of macro F1 scores achieved by all the participating teams for Task 1 and Task 2 is depicted in Figures 3 and 4, respectively.

5. Conclusion and Future Works

This paper describes the methods of our run submissions for the CryptOQA shared task at FIRE 2025 underscores the efficacy of transformer-based models in dealing with the issues of opinion extraction and QA from social media posts about cryptocurrencies. Through the application of MTL strategies and fine-tuning pre-trained models like BERT and CryptoBERT, we were able to capture the hierarchical nature of opinions expressed on Twitter, Reddit, and YouTube. Our findings illustrate that domain-specific models such as CryptoBERT yield significant improvements over universal models, especially for dealing with subtle and unstructured conversations prevalent in the cryptocurrency community. On Task 1, our models reported competitive macro F1 scores across various platforms, indicating

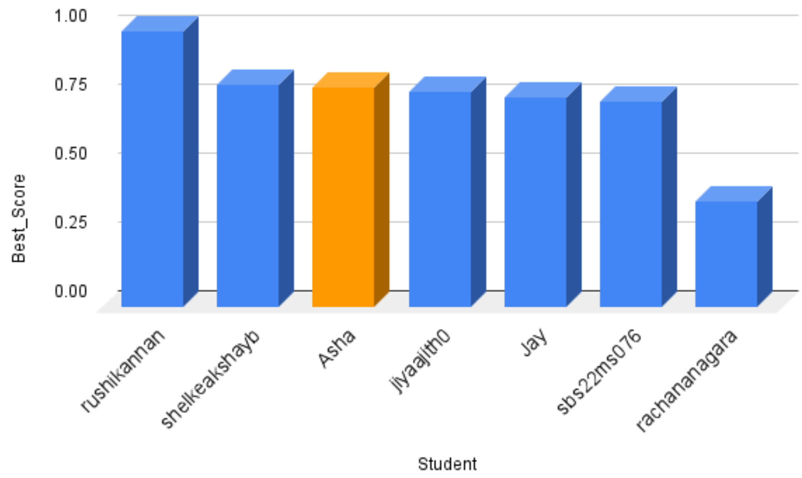

Figure 4: Comparison of macro F1 scores of all participants in Task 2

stability in separating distinct categories of opinion. Analogously, in Task 2, the proposed BERT and DistilBERT-based models achieved good performance in relevance classification, thus further confirming the viability of pretrained models for question-answering tasks in user-generated noisy text. Among the proposed models, our best-performing model for the opinion classification task achieved macro F1 scores of 0.5694, 0.5570, and 0.7496 on the Reddit, Twitter, and YouTube datasets, respectively, as provided by the task organizers. Further, our top model for Task 2 attained a macro F1 score of 0.7940. Although our achievements are promising, the current results still leave much room for improvement in dealing with sarcasm, code-switching, and implicit sentiment, which pervade social media language. Another constraint is the dependency on annotated datasets since annotated data at scale is hard and expensive to attain in the rapidly changing cryptocurrency space. Subsequent work must investigate semi-supervised and unsupervised approaches like self-training, co-training, or weak supervision to take advantage of the vast amounts of unlabeled social media data.

Declaration on Generative AI

During the preparation of this work, the author(s) used ChatGPT-4 in order to:Grammar and spelling check. Using this tool, the author(s) reviewed and edited the content as needed and take full responsibility for the publication's control

References

[1] D.-H. Kim, Effects of social media-based peer opinions on the prices of cryptocurrency options, Journal of Futures Markets (2025).
[2] S. Yu, J. Padfield, Advanced techniques in profiling cryptocurrency influencers: a review, International Journal of Blockchains and Cryptocurrencies 5 (2024) 1–18.
[3] S. Sarkar, G. Sen, K. Rudra, K. Ghosh, Understanding cryptocurrency related opinions and questions from social media posts (cryptoqa 2024), in: Forum of information retrieval (FIRE), 2024.
[4] K. Girish, A. Hegde, F. Balouchzahi, H. L. Shashirekha, Profiling cryptocurrency influencers with sentence transformers., in: CLEF (Working Notes), 2023, pp. 2599–2607.
[5] S. Rouhani, E. Abedin, Crypto-currencies narrated on tweets: a sentiment analysis approach, International Journal of Ethics and Systems 36 (2020) 58–72.
[6] A. Al Suwaidi, Machine Learning Models of Crypto Currency Trends Using Attention-based Sentiment Analysis, Ph.D. thesis, Khalifa University of Science, 2024.

[7] L. S. Kumar, A. Hegde, B. R. Chakravarthi, H. Shashirekha, R. Natarajan, S. Thavareesan, R. Sakuntharaj, T. Durairaj, P. K. Kumaresan, C. Rajkumar, Overview of second shared task on sentiment analysis in code-mixed tamil and tulu, in: Proceedings of the Fourth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages, 2024, pp. 62–70.

[8] A. Hegde, G. Kavya, S. Coelho, P. Lamani, H. L. Shashirekha, Munlp@ dravidianlangtech2023: Learning approaches for sentiment analysis in code-mixed tamil and tulu text, in: Proceedings of the Third Workshop on Speech and Language Technologies for Dravidian Languages, 2023, pp. 275–281.

[9] S. Sarkar, A. Badwal, A. Roy, K. Rudra, K. Ghosh, Cryptopiqa: A new opinion and question answering dataset on cryptocurrency, in: Proceedings of the 31st International Conference on Computational Linguistics, 2025, pp. 11107–11120.

[10] A. Deroy, S. Maity, Cryptollm: Unleashing the power of prompted llms for smartqna and classification of crypto posts, arXiv preprint arXiv:2411.07917 (2024).

[11] R. Sahal, Predicting Optimal Cryptocurrency using Social Media Sentimental Analysis, Ph.D. thesis, Dublin, National College of Ireland, 2023.

[12] X. Huang, W. Zhang, X. Tang, M. Zhang, J. Surbiryala, V. Iosifidis, Z. Liu, J. Zhang, Lstm based sentiment analysis for cryptocurrency prediction, in: International conference on database systems for advanced applications, Springer, 2021, pp. 617–621.

[13] S. Das, U. Chowdhury, N. Lijin, A. Deep, S. Saha, A. Maurya, Investigate how market behaves: Toward an explanatory multitasking based analytical model for financial investments, IEEE Access 12 (2024) 30928–30940.

[14] A. Hegde, H. Shashirekha, Transformer-driven multi-task learning for fake and hateful content detection, in: Proceedings of the 21st International Conference on Natural Language Processing (ICON): Shared Task on Decoding Fake Narratives in Spreading Hateful Stories (Faux-Hate), 2024, pp. 29–35.

[15] S. Coelho, A. Hegde, P. Lamani, H. L. Shashirekha, et al., Mucsd@ dravidianlangtech2023: Predicting sentiment in social media text using machine learning techniques, in: Proceedings of the Third Workshop on Speech and Language Technologies for Dravidian Languages, 2023, pp. 282–287.

[16] Y. Zhang, Q. Yang, A Survey on Multi-task Learning (2021) 5586–5609.

[17] B. R. Chakravarthi, P. Kumaresan, R. Priyadharshini, P. Buitelaar, A. Hegde, H. Shashirekha, S. Rajiakodi, M. Á. García, S. M. Jiménez-Zafra, J. García-Díaz, et al., Overview of Third Shared Task on Homophobia and Transphobia Detection in Social Media Comments, in: Proceedings of the Fourth Workshop on Language Technology for Equality, Diversity, Inclusion, 2024, pp. 124–132.

[18] A. Hegde, G. Kavya, S. Coelho, H. L. Shashirekha, MUCS@ LT-EDI2023: Homophobic/Transphobic Content Detection in Social Media Text using mBERT, in: Proceedings of the Third Workshop on Language Technology for Equality, Diversity and Inclusion, 2023, pp. 287–294.

[19] X. Hu, Enhancing answer selection in community question answering with pre-trained and large language models. corr abs/2311.17502 (2023), 2023.

[20] G. A. Boroujeni, H. Faili, Y. Yaghoobzadeh, Answer selection in community question answering exploiting knowledge graph and context information, Semantic Web 13 (2022) 339–356.