

Article

Pump It: Twitter Sentiment Analysis for Cryptocurrency Price Prediction

Vladyslav Koltun ^{1,*} and Ivan P. Yamshchikov ^{2,3,*} 

¹ Instituto Superior de Economia e Gestao, University of Lisbon, 1200-781 Lisbon, Portugal

² THWS, CAIRO, 97082 Würzburg, Germany

³ CEMAPRE, University of Lisbon, 1649-004 Lisbon, Portugal

* Correspondence: ivan@yamshchikov.info (V.K.); yndx.kroniker@yandex.ru (I.P.Y.)

Abstract: This study demonstrates the significant impact of market sentiment, derived from social media, on the daily price prediction of cryptocurrencies in both bull and bear markets. Through the analysis of approximately 567 thousand tweets related to twelve specific cryptocurrencies, we incorporate the sentiment extracted from these tweets along with daily price data into our prediction models. We test various algorithms, including ordinary least squares regression, long short-term memory network and neural hierarchical interpolation for time series forecasting (NHITS). All models show better performance once the sentiment is incorporated into the training data. Beyond merely assessing prediction error, we scrutinise the model performances in a practical setting by applying them to a basic trading algorithm managing three distinct portfolios: established tokens, emerging tokens, and meme tokens. While NHITS emerged as the top-performing model in terms of prediction error, its ability to generate returns is not as compelling.

Keywords: cryptocurrency; sentiment analysis; twitter



Citation: Koltun, Vladyslav, and Ivan P. Yamshchikov. 2023. Pump It: Twitter Sentiment Analysis for Cryptocurrency Price Prediction. *Risks* 11: 159. <https://doi.org/10.3390/risks11090159>

Academic Editor: Dimitrios Koutmos

Received: 27 February 2023

Revised: 1 August 2023

Accepted: 17 August 2023

Published: 4 September 2023



Copyright: © 2023 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Cryptocurrency assets have garnered significant attention in the finance and investment spheres in recent years. The term “crypto” saw a surge in Google searches prior to mid-2018, when cryptocurrencies offering yield to holders began to gain traction. The subsequent rise in popularity hit an all-time high between mid-2020 and early 2021, coinciding with a major bull run in the crypto market. Although these assets are not entirely novel financial instruments, their foundation on blockchain technology is considered a groundbreaking advancement [Gupta and Chaudhary \(2022\)](#); [Laboure et al. \(2021\)](#).

Predicting the price of cryptocurrencies, like that of any asset, can be approached from several points of view. Instead of looking at it as a one-dimensional problem, we want to factor in additional sources of information, which we believe might contribute to better understanding the price trajectory.

Each asset class exhibits a degree of volatility, which varies from class to class. Cryptocurrencies and crypto assets, in general, exhibit high volatility due to various factors. As a market that operates 24/7 and is easily accessible, it enables a high volume of daily transactions from a vast number of participants. Although cryptocurrencies have been around for some time, their recent surge in popularity has attracted a significant number of first-time traders. This influx of novice participants often introduces a degree of irrationality into the market, as inexperienced traders are more likely to react emotionally, succumbing to the fear of missing out (FOMO) or fear, uncertainty, and doubt (FUD) [Nagel \(2018\)](#). Numerous factors, including major macro events, contribute to the market’s noise and increase volatility. Regulatory news, celebrity or corporate endorsements of a crypto project, and other global news unrelated to crypto can all significantly impact cryptocurrency prices. Thus, the ability to accurately forecast price movements and seize potential opportunities is imperative.

Indeed, the public's perception of an asset significantly influences its value, and social media now plays a pivotal role in shaping this image. Various platforms, such as Reddit and Discord, generate copious amounts of crypto-related content daily. However, Twitter reigns supreme as the primary platform for crypto-related discourse. This social network is the principal source of shared opinions, news releases, and announcements related to cryptocurrency. Twitter propagates information far more rapidly than traditional news outlets, often impacting markets nearly instantaneously. A prime example of this is the significant price fluctuations that Bitcoin and Dogecoin experienced in response to tweets from prominent figures like Elon Musk, depending on the sentiment conveyed.

The Figure 1 depicted above¹ clearly indicates market shifts corresponding to Elon Musk's tweets: prices dip following his negative comments and climb in response to positive ones. This sentiment analysis of Musk's Bitcoin-related tweets could provide insights for trading, offering potential opportunities for creating long or short positions on Bitcoin based on the tweet's sentiment. Although Musk's influence on the crypto market is notably potent due to his high-profile social media presence, other individuals also wield significant influence, albeit in a less overt manner. Regardless, it is apparent that understanding market sentiment towards a particular cryptocurrency is crucial and may be instrumental in predicting its price trends.



Figure 1. BTC price in May–June 2021.

In order to accurately gauge market sentiment, one must determine appropriate metrics to assess the general mood, then quantify these metrics for the input into a learning algorithm. For instance, from a single tweet, one could gather data on the number of likes, comments, and shares. Alternatively, Natural Language Processing (NLP) algorithms could be used to derive sentiment, polarity, and subjectivity from the text. The existing literature predominantly focuses on extracting sentiment from tweets. Our work seeks to evaluate the comprehensiveness of this approach.

In addition to these, we added information gathered about cryptocurrencies, namely daily traded volume and daily average price. To make sure our results were reliable, we applied different models and subsets of features to our data.

Although there are several papers with goals aligning with ours, this work distinguishes itself in several crucial ways. Firstly, the data we use incorporate information on a broader range of cryptocurrencies. Secondly, we focus on gauging the informational value of Twitter, which we regard as the most influential source of news in the cryptocurrency industry, surpassing newspapers and alternative social media platforms in relevance. Additionally, we implement a machine learning model which, to our knowledge, has yet to be employed in the literature for cryptocurrency price prediction. This model is then contrasted with a comprehensive baseline of commonly utilised models. Finally, rather than limiting our evaluation to error metrics, we design a generic trading strategy to explore the performance of the models when deployed as trading tools.

With this research, we aim to answer three key questions:

1. Does Twitter offer valuable information for price prediction? Are some tokens more prone to social media influence than others?
2. Can the NHITS be a new state-of-the-art model for cryptocurrency price prediction?
3. Are the behaviours on the bull and bear markets profoundly different and how should one adjust the models used for price predictions depending on the state of the market?

To provide answers, we compare NHITS with various classical models used for cryptocurrency price prediction. We also assess how each model performs if employed as a trading tool during two separate time intervals that cover a bull and a bear market. Firstly, we provide a brief overview of previous related work. Secondly, we move on to describing the data processing conducted and performance criteria used. Finally, we present the main findings and conclusions drawn.

The reader will find that, through our work, we show a clear dependence between the size of the currency and its susceptibility to manipulations. We hope further research could address this particular issue and that intraday data are too crude to build precise predictions, yet combined with twitter data, they form a ground for speculative trading of smaller cryptocurrencies.

2. Related Works

Price prediction for various financial assets is a well-established task. However, the rise of cryptocurrencies and crypto assets has significantly extended the field. As previously mentioned, the problem is often addressed as a one-dimensional time-series analysis task, where future prices are predicted based on past information [Patel et al. \(2020\)](#); [Peng et al. \(2018\)](#). However, this method often lacks precision and applicability, as it typically neglects other available sources of information. External information not yet fully reflected in the price due to market inefficiencies can be an invaluable asset for price prediction.

Therefore, it is vital to understand what other sources must be factored in. Evidence suggests a tangible link between equity and crypto markets. Market shifts observed in indices like the S&P500 or the DJI are often mirrored, in some form, by Bitcoin (BTC), which frequently sets the tone for other cryptocurrencies. This positive correlation between Bitcoin and the S&P500, as highlighted in [Iyer \(2022\)](#), indicates that movements in traditional markets should be monitored and leveraged as another information source. The authors of this paper include an illustrative example showcasing a striking similarity in the trends of these asset prices, with the correlation between them markedly increasing from 2017 to 2021.

While stocks, as part of the conventional financial market, generally demonstrate more predictable behaviour, research in [Mittal and Goel \(2012\)](#) and [Li et al. \(2014\)](#) suggests that market sentiment remains an important factor in forecasting stock movements as well. Relevant to our work, [Nguyen et al. \(2015\)](#) showed that news data about Ethereum can be helpful in predicting its price fluctuations, and [Abraham et al. \(2018\)](#) demonstrated how Twitter data and Google Trends help make better informed trading decisions. Also, in [Kim et al. \(2023\)](#), it was proven that it is advantageous to incorporate sentiment scores from Twitter when trading Bitcoin. The authors of this work and the ones of [Dwivedi and Vemareddy \(2023\)](#) made compelling arguments for not using a pre-trained model to analyse the tweets by noting that crypto Twitter has a very specific jargon. Although we agree with this statement, we do not believe it is relevant enough to our work.

Although data are highly important, the choice of the model used is also crucial to successfully predict future price fluctuations. Commonly used models include Autoregressive models, Moving Average models, or a combination thereof. Owing to the frequent seasonality observed in most prices, a SARIMA model could be applied, as demonstrated by the authors of [Cabanilla \(2016\)](#).

Reviewing the existing literature, it is apparent that the Long Short-Term Memory (LSTM) model is a commonly utilised tool. First introduced in 1997 in [Hochreiter and Schmidhuber \(1997\)](#), this machine learning algorithm remains a widely used model, often

producing satisfactory results [Vo et al. \(2019\)](#). An interesting application of LSTM is detailed in [Huang et al. \(2021\)](#), where the authors embarked on a research approach somewhat similar to ours, albeit with distinct aspects. The authors analyse the prominent Chinese social media platform, Sina-Weibo, to capture market sentiment, employing these data in conjunction with LSTM to predict the prices of Bitcoin, Ethereum, and Ripple. Their results surpassed those of a leading auto-regressive model by 18.5% in precision and 15.4% in recall. However, as indicated in [Wong \(2021\)](#), LSTM performance, though superior to that of Naive Bayes, was only marginally better than random chance for certain tasks. Furthermore, the field of machine learning has seen significant advancements, with time-series analysis being one of many areas experiencing substantial progress. For example, [Prasad et al. \(2022\)](#) used the Youtube comment section as social media information source and fed these data into a stacked ensemble model consisting of a Decision Tree, K Nearest Neighbours, a Random Forest Classifier and XGBoost and a meta/base classifier—Logistic Regression, which achieved a 94.2% accuracy.

In this study, we suggest the utilisation of a different model that, to our knowledge, was not applied to cryptocurrency price prediction. In [Challu et al. \(2022\)](#), the authors introduced the Neural Hierarchical Interpolation for Time Series Forecasting, or NHITS. This model, inspired by the Transformer model [Vaswani et al. \(2017\)](#), is well-suited to long-horizon time-series prediction problems. Its architecture is specifically designed to address two issues: the volatility of predictions and computational complexity. To tackle these challenges, the authors proposed the incorporation of the following concepts:

1. **Multi-Rate Data Sampling:** This employs sub-sampling layers, thereby reducing memory demands and necessary computations while preserving the model's ability to detect long-range dependencies;
2. **Hierarchical Interpolation:** This mechanism ensures the smoothness of multi-step predictions by reducing the dimensionality of the neural network's prediction.

This dual approach allows each block the focus on predicting its own frequency band of the time-series data. This innovative method delivered state-of-the-art results across six large-scale benchmark datasets, including one comprising exchange rates.

In this paper, we apply NHITS to a large dataset of daily cryptocurrency prices. We compare the resulting model with more frequently used methods; for example, LSTM. We also use aggregated information from the tweets that mention corresponding cryptocurrency and demonstrate how that information could be used for better price prediction. Finally, we demonstrate that this twitter information is far less useful for the currencies with higher daily turnover, such as Bitcoin or Ether; however, it is far more salient for smaller or the so-called "meme currencies".

3. Data

In this study, we utilise the dataset presented in [Garg et al. \(2021\)](#). The data comprise tweets and pricing information related to the top 12 cryptocurrencies, ranked by market capitalisation at the time, with approximately 576 thousand tweets spanning 1255 days. We are confident that this volume of data is ample for our model to adequately capture the general market sentiment for each day. To our knowledge, this dataset represents the most comprehensive publicly accessible resource of cryptocurrency-related tweets. For the purposes of data processing, metric extraction, and model fitting, we rely on Python and its associated libraries.

3.1. Data Processing

Our data encapsulate information about the top 12 cryptocurrencies by trading volume. We examined each cryptocurrency individually, which involved working with twelve separate data frames. We calculated the daily average price and an average of the daily high and low prices. We also incorporated the daily trading volume into our analysis, recognising its significance, particularly in high-volatility contexts.

For each tweet, we extracted three features: sentiment, subjectivity, and polarity. We used DistilBERT [Sanh et al. \(2019\)](#) fine-tuned on SST-2² to extract sentiment. We utilised TextBlob³ to calculate subjectivity and Vader [Hutto and Gilbert \(2014\)](#) to determine polarity. For each feature, we computed a daily average, resulting in our final data frame.

Our analysis spans two time intervals, each reflecting a different market state: one representing a bear market and the other a bull market. Specifically, we used the following time frames:

- Bear market: from October 2017 to August 2020, a period characterised by a sluggish cryptocurrency market and a major downturn in 2018.
- Bull market: from October 2017 to February 2021, encompassing all available data, with the final year and a half marked by significant appreciation in nearly all tokens.

Our aim is to forecast the price at time $t + 1$ using the information available up to time t . Hence, the prices of the cryptocurrencies serve as the dependent variables while all others are independent. For each model, we utilised 371 data points for testing, resulting in less data available for the bear market training. By selecting two distinct time intervals for testing, we are able to assess the sensitivity of our models.

3.2. Performance Evaluation

Assessing the performance of a predictive model solely based on traditional error metrics, such as the difference between the forecasted price and actual price, does not always provide a comprehensive understanding of the model's real-world trading applicability. For instance, an incorrect decision to initiate a long or short position could be detrimental, even if the price prediction slightly deviates from the real price. Conversely, a model that accurately and consistently positions long on a currency that modestly appreciates, and short on a currency set to depreciate, might prove valuable, despite any over-optimistic or over-pessimistic absolute error. Although most papers we encountered conducted their performance analysis using error metrics like MAE, MSE, or R^2 , some also considered overlaying a trading algorithm [Colianni et al. \(2015\)](#). Consequently, this study proposes constructing heuristic trading algorithms based on model predictions and comparing the resulting returns with certain baseline models.

Our proposed approach is straightforward: we endow a bot with a set initial wealth and present it with a choice among three portfolios: established tokens (BTC, ETH, LTC, XRP, XMR), emerging tokens (ADA, BNB, LINK), and meme tokens (DOGE, IOTA, TRX, XLM). The categorisation of tokens into these three groups was based on their use cases, real-world value, and market perceptions. Each day, the bot decides whether to buy (with the amount of wealth allocated to each buy order being proportional to the token's average historical daily volume), hold, or sell a particular coin. Returns below 100% signify a loss, while those above it indicate a winning strategy.

We evaluate the model's performance both in terms of the associated bot's performance and the overall prediction error. Significantly, we expanded our analysis to include twelve tokens, whereas previous studies typically considered a maximum of three.

In this study, we opted to use the Root Mean Square Error Percentage (RMSEP) as our primary error metric, which is defined as follows:

$$\text{RMSEP} = \sqrt{\frac{\sum_{i=1}^n (y_i - \hat{y}_i)^2}{n}} \times 100. \quad (1)$$

We chose this metric because it allows the comparison of our results between cryptocurrencies in terms of relative error.

3.3. Model Description

This section outlines the models employed in our study and the rationale behind each selection.

We initiated our modeling with Ordinary Least Squares (OLS) regressions with a lag of one. Subsequently, we approached the problem as a binary classification task, predicting whether the price will rise or fall the next day. For this, we utilised logit regressions that anticipate the following day's market movement. Due to their simplicity and intuitiveness, these models serve as our benchmark for comparison.

Given the widespread use of Long Short-Term Memory (LSTM) models in the related literature, we also applied this method to our data. Finally, we fitted our data to the Neural Hierarchical Interpolation for Time Series (NHITS) model, as it was explicitly designed to address the volatility and complexity issues associated with long-horizon predictions. It is important to note that these models involve a higher degree of complexity, not only demanding more computational resources, but also necessitating careful selection of a large array of parameters.

4. Results

Our baseline consists of OLS with a lag of one and LOGIT models. Given the availability of multiple distinct features, we initially determined the subset of features that produced the optimal overall scores within each of these two models. For the OLS model with a lag of one, we selected the price, volume, and sentiment features, and we chose to use the same set of features for the LOGIT regression. Refer to Appendix A for details. We marked the best score in each row of the tables presented in this section and in Appendix A in bold.

4.1. Model Comparison

With our baseline established, we compared it against two models: LSTM and NHITS. Although LSTM is commonly employed in the field, we were not successful in reaching either the baseline accuracy or returns using it. We attribute this result to either the high intraday volatility of cryptocurrency prices or an insufficiently large dataset. Conversely, NHITS outperformed all previously considered models, achieving the highest accuracy.

BULL RMSEP Comparison between Models			
RMSEP	OLS-1-S	LSTM-ALL	NHITS-ALL
ada	9.15%	46.2%	5.57%
bnb	5.60%	33.4%	5.14%
btc	3.62%	37.6%	3.60%
doge	10.52%	53.9%	8.15%
eth	4.37%	67.0%	4.36%
iota	7.06%	38.9%	7.53%
link	10.27%	76.0%	5.74%
ltc	4.63%	30.8%	4.52%
trx	8.22%	77.4%	4.75%
xlm	8.40%	305.0%	5.81%
xmr	4.70%	39.3%	4.30%
xrp	15.41%	21.3%	5.99%
BEAR RMSEP Comparison between Models			
RMSEP	OLS1-P	LSTM-ALL	NHITS-ALL
ada	4.78%	46.1%	5.62%
bnb	4.54%	33.4%	4.03%
btc	3.22%	37.6%	3.16%
doge	13.22%	53.8%	4.38%
eth	4.20%	66.9%	3.80%
iota	4.66%	38.9%	5.03%
link	5.03%	75.8%	4.85%
ltc	5.03%	30.8%	3.72%
trx	23.15%	77.3%	5.25%
xlm	11.09%	304.6%	10.82%
xmr	3.69%	39.2%	4.02%
xrp	11.59%	21.3%	3.18%

Regrettably, NHITS did not produce the highest returns, as depicted in Figures 2 and 3. While it performed admirably with certain portfolios, the OLS1 and LOGIT models demonstrated superior outcomes in certain instances. A slight misprediction in the wrong direction could mean missing an opportunity to take a short or long position, while a significant overestimation in the correct direction does not substantially influence the overarching trading strategy. Hence, this disparity could be sufficient to account for the observed variations in returns.

However, we were able to demonstrate that NHITS performed the best according to our selected error metric. The primary aim of price prediction is to stay ahead of the market trends and accrue wealth. Following this line of reasoning, the LOGIT or OLS1 models might be the more suitable choice, as they delivered superior returns. While we utilised a rather simplistic trading algorithm, which was valuable for model comparison in some aspects, we proceeded with NHITS for our further analysis, as subsequent results are reliant on the error metric.

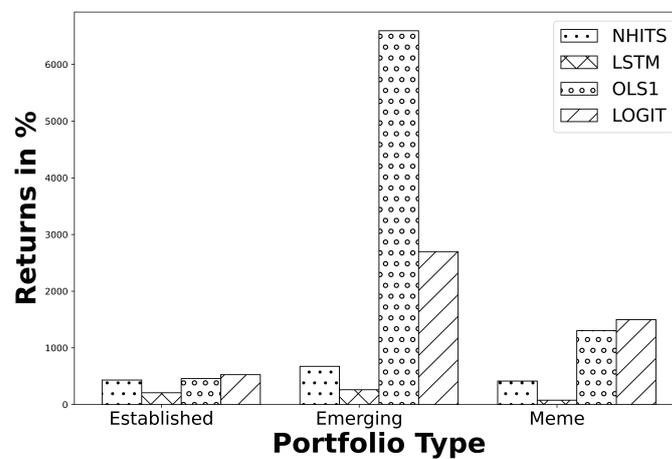


Figure 2. Returns in % achieved by each model separated by portfolio type during the bull market period considered. Note OLS1 and LOGIT outperformed the rest by a large margin and the emerging portfolio generally achieved higher returns.

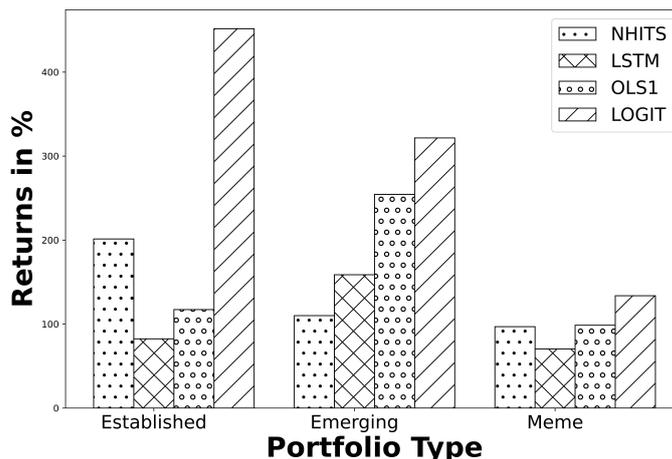


Figure 3. Returns in % achieved by each model separated by portfolio type during the bear market period considered. Again, LOGIT showed an outstanding performance. Some models indicate preference of holding a basket of established tokens, others suggest betting on emerging coins, while none suggest meme tokens as a viable option.

4.2. Further Investigation

Our study took into account both bear and bull market conditions. As anticipated, we recorded lower prediction errors in the bear market, and heightened returns during the bull

market’s peak, as depicted in Figures 4 and 5. Crucially, the disparity in prediction errors did not increase substantially when transitioning from a sluggish market to the brisk pace of a bull market. Additionally, we were able to generate higher returns in a bear market by investing in more stable tokens, while the emerging portfolio proved more beneficial during the bull market phase. Based on this analysis, we found no evidence to support venturing into a portfolio composed of meme coins given the risk–return trade-off was not justified by the returns achieved.

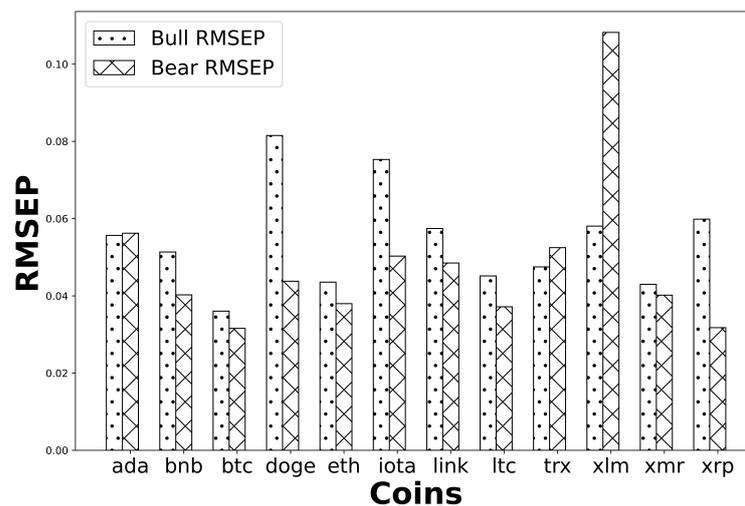


Figure 4. RMSEP % contrast between the bull and bear periods for each coin analysed. It is easier to predict tokens during a bear market, with XLM being the only outlier.

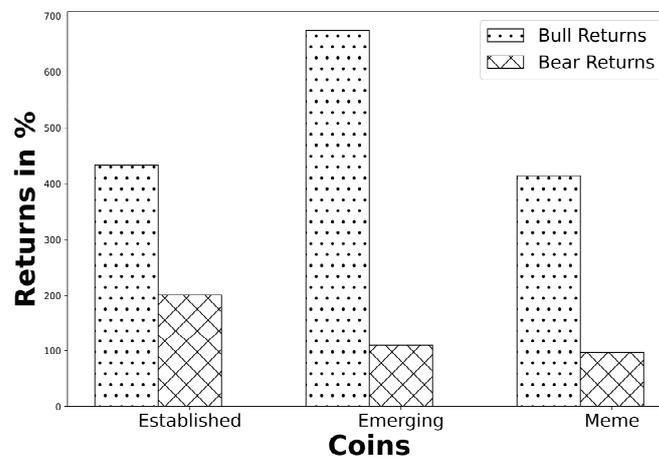


Figure 5. Returns in % comparison during both periods for different portfolios. Even though the market is trickier to predict, returns doubled for established tokens and the impact on emerging and meme portfolios was even more significant.

As illustrated in plots Figures 6 and 7, for comparable levels of daily average trading volume, we registered lower prediction errors for established tokens and higher errors for meme coins, regardless of the prevailing market conditions. As expected, across both periods, BTC and ETH proved to be the easiest to predict, whereas for tokens like DOGE and XLM, we observed a greater degree of error.

Consistent with these findings, upon examination of Figures 8 and 9, we discovered that the lowest prediction errors corresponded to tokens with lower volatility, barring a few outliers. Unsurprisingly, meme coins were scattered across the plots, given that their price movements are largely influenced by hype, fear of missing out (FOMO), and fear,

uncertainty, and doubt (FUD). As anticipated, BTC demonstrated the lowest volatility and RMSEP throughout both periods, reinforcing its stability within the cryptocurrency market.

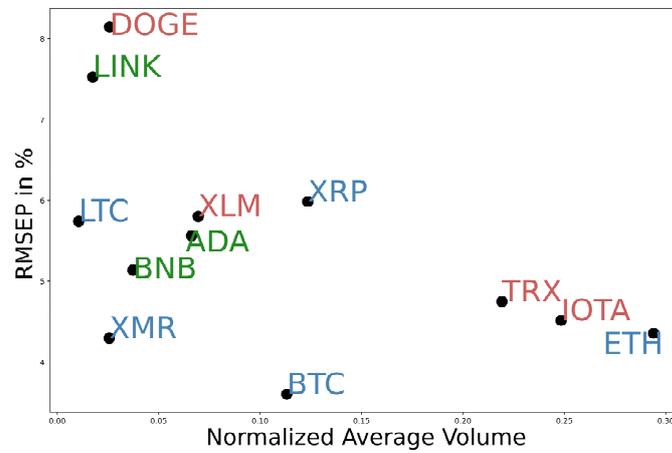


Figure 6. Error metric against the daily average volume of each token during a bull market. Bitcoin, Ethereum and Monero had the lowest errors, while Doge and Link had the worst performance.

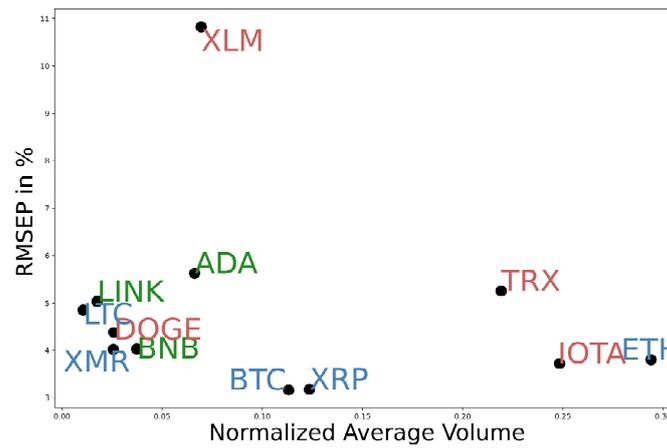


Figure 7. Error metric against the daily average volume of each token during a bear market. Here, BTC, ETH and XMR had better results as well.

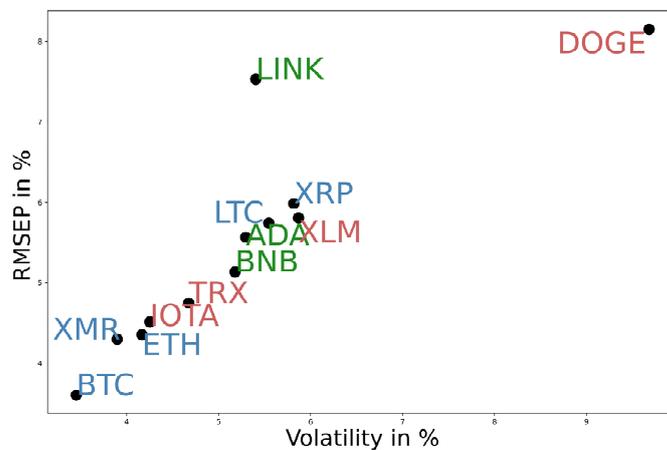


Figure 8. Error and actual volatility comparison during the bull period. An almost straight line relates these, with Bitcoin at the bottom and Doge on top, as expected.

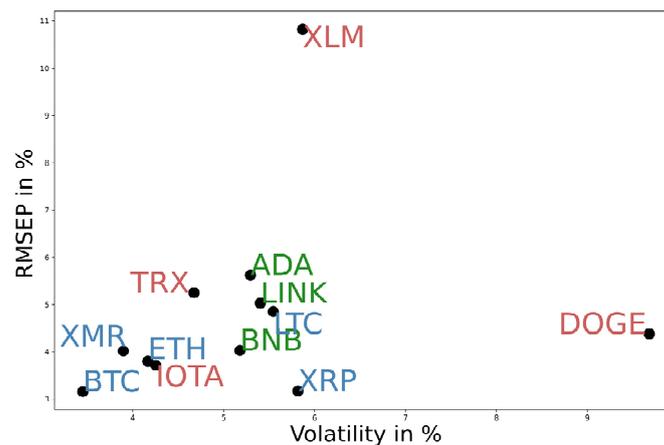


Figure 9. Error and actual volatility during the bull period. There is no clear line formed, with established tokens being in the down-left corner and Doge and Stellar far away from the rest.

5. Conclusions

This study explores the potential utility of social media sentiment as a predictive tool for cryptocurrency prices during bear and bull market phases. Drawing from our results and the existing literature, we conclude that market sentiment enhances the predictive accuracy of our models, with the average daily sentiment serving as the most informative feature. In the context of daily price predictions, “memoryless” models, such as one-day OLS, significantly outshine LSTM, yet fall short of NHITS in accuracy. This innovative model outperforms others in the prediction task, but does not yield the same returns as LOGIT and OLS1. We attribute the higher returns of our baseline models to fortunate market movement predictions rather than accurate price estimations given the simplicity of our trading strategy.

Our findings also suggest that the prices of certain cryptocurrencies appear more sensitive to social media signals, whereas “established” coins seem less influenced by such cues. We found no evidence to endorse investing in meme cryptocurrencies for mid-frequency trading⁴, while emerging tokens consistently offer the best return–risk ratio irrespective of market conditions.

The direct implications of this work are the clear dependence between the size of the currency and its susceptibility to manipulations. We hope further research could address this particular issue. The intraday data are too crude to build precise predictions, yet, combined with twitter data, they form a ground for speculative trading of smaller cryptocurrencies. Sentiment analysis proved to be a worthy source of information as it was, but if the extraction of the sentiment is performed with a more specific model, there is potential to obtain better predictions.

Potential enhancements to our study could include expanding the dataset and collecting information, for instance, on an hourly basis. Such increments would provide a more nuanced understanding of the market sentiment’s dynamism and the price behaviour of cryptocurrencies.

As underscored in our discussion, the volume of data is crucial for this kind of analysis. Additionally, the selection of social media platforms can influence the results. We relied on historical data to differentiate between bear and bull market conditions, which might pose challenges in the event of sudden market shocks or ambiguous market states.

Author Contributions: Conceptualization, I.Y.P.; methodology V.K. and I.Y.P.; software V.K.; validation V.K. and I.Y.P.; formal analysis, V.K. and I.Y.P.; investigation, V.K.; data curation, V.K.; writing—original draft preparation, V.K.; writing—review and editing, I.Y.P.; visualization, V.K.; supervision, I.Y.P.; project administration, I.Y.P. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The raw data can be found on <https://github.com/am15h/CrypTop12>, accessed on 1 January 2023.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A

We fitted the OLS1 and LOGIT on three sets of features: price and volume; sentiment, price and volume; and reply, like, retweet, polarity, subjectivity, sentiment, price and volume. We first analysed the bull and bear cases for OLS1 and concluded that there is no clear best performer, but there is a slight edge to use at least sentiment as a feature. It is important to note that the results across the three sets are similar in both metrics used for comparison.

Table A1. Results for OLS with lag 1.

OLS1 BULL			
RMSEP	OLS1-All	OLS1-P	OLS1-S
ada	10.08%	9.10%	9.15%
bnb	5.76%	5.40%	5.60%
btc	3.73%	3.57%	3.62%
doge	10.85%	11.09%	10.52%
eth	4.53%	4.39%	4.37%
iota	7.41%	5.46%	7.06%
link	10.14%	9.14%	10.27%
ltc	4.85%	4.62%	4.63%
trx	8.82%	8.23%	8.22%
xlm	9.05%	8.32%	8.40%
xmr	4.72%	9.70%	4.70%
xrp	15.88%	15.63%	15.41%
OLS1 BEAR			
RMSEP	OLS1-ALL	OLS1-P	OLS1-S
ada	7.53%	4.78%	6.42%
bnb	5.32%	4.54%	4.96%
btc	3.35%	3.22%	3.20%
doge	14.37%	13.22%	14.83%
eth	4.35%	4.20%	4.11%
iota	8.70%	4.66%	8.33%
link	5.02%	5.03%	5.04%
ltc	8.59%	5.03%	5.22%
trx	28.80%	23.15%	27.94%
xlm	12.99%	11.09%	11.33%
xmr	3.80%	3.69%	3.83%
xrp	12.54%	11.59%	13.10%

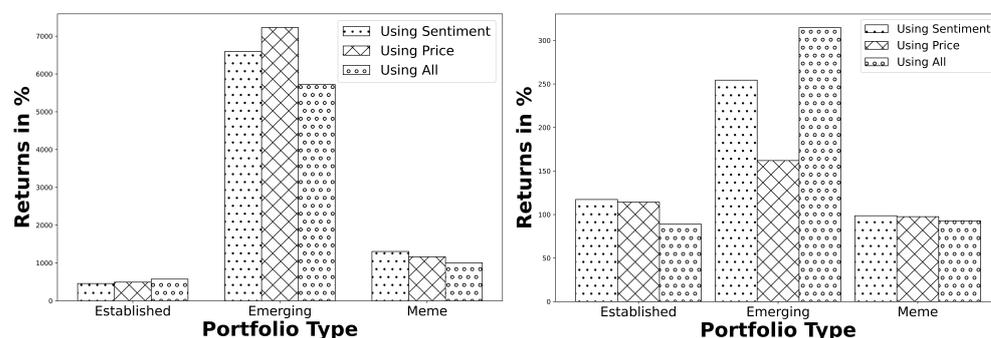


Figure A1. Returns in % using OLS1 for each portfolio type during a bull state on the right and bear state on the left. Results were similar with sentiment slightly improving Established and Meme portfolios during the bull run. Emerging portfolio returns almost doubled in the bear market when we used sentiment on top of volume and price.

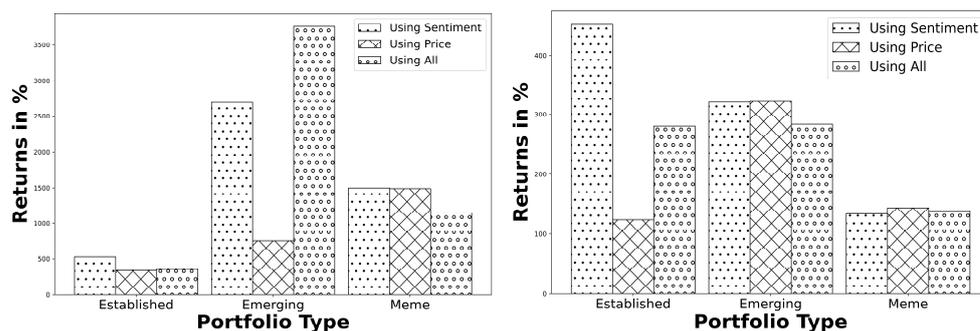


Figure A2. Returns for each portfolio type in % for LOGIT for bull on the left and bear on the right. There is a clear advantage of using at least sentiment as a feature with the Emerging Portfolio and a slight edge in the other 2 during the bull market. For the bear market, we obtained fairly similar returns with Emerging Portfolio and Meme Portfolio, but it was far more advantageous to use at least sentiment for the Established Portfolio.

In a similar fashion, we conducted the same analysis for LOGIT and were able to conclude that sentiment improves the accuracy of our model in most cases and in both markets. It also generates higher returns if we select an appropriate portfolio based on the condition of the market (established for bear and emerging for bull).

Table A2. Accuracy for LOGIT with lag 1 for both market types.

LOGIT BULL			
Accuracy	LOGIT-All	LOGIT-P	LOGIT-S
ada	59.95%	54.30%	59.41%
bnb	56.99%	56.99%	56.18%
btc	50.81%	52.69%	52.69%
doge	52.42%	54.30%	55.38%
eth	58.60%	55.91%	58.60%
iota	55.38%	54.57%	55.11%
link	54.57%	56.45%	55.11%
ltc	53.23%	51.61%	53.49%
trx	54.84%	54.03%	53.76%
xlm	56.18%	55.65%	57.26%
xmr	57.53%	56.72%	59.68%
xrp	55.65%	52.42%	54.84%
LOGIT BEAR			
Accuracy	LOGIT-All	LOGIT-P	LOGIT-S
ada	60.00%	58.92%	60.54%
bnb	60.54%	61.62%	61.08%
btc	56.22%	58.11%	58.11%
doge	56.22%	54.59%	54.59%
eth	60.27%	60.00%	59.73%
iota	54.59%	54.59%	54.32%
link	56.49%	54.05%	55.41%
ltc	54.32%	59.73%	59.19%
trx	52.43%	54.05%	54.32%
xlm	57.84%	55.68%	57.57%
xmr	56.76%	58.38%	59.73%
xrp	55.95%	51.89%	56.49%

Table A3. R² comparison for OLS and NHITS.

Coin	R ² comparison across models and market states			
	Bull		Bear	
	NHITS	OLS1	NHITS	OLS1
ada	98.68%	99.74%	99.14%	99.16%
bnb	97.48%	99.85%	97.67%	97.25%
btc	99.55%	100.00%	96.80%	96.92%
doge	96.27%	99.31%	91.79%	73.39%
eth	99.57%	100.00%	97.62%	98.07%
iota	98.45%	99.74%	96.84%	96.85%
link	99.11%	92.18%	98.87%	98.74%
ltc	99.08%	99.99%	98.06%	97.80%
trx	97.96%	99.51%	95.18%	69.07%
xlm	98.51%	99.45%	96.42%	93.45%
xmr	98.99%	99.79%	97.11%	97.32%
xrp	96.39%	96.82%	97.20%	88.86%

Table A4. MAE comparison for OLS and NHITS.

Coin	MAE comparison across models and market states			
	Bull		Bear	
	NHITS	OLS1	NHITS	OLS1
ada	0.60%	1.13%	0.17%	0.19%
bnb	0.45%	0.23%	0.15%	0.18%
btc	0.72%	0.11%	0.28%	0.27%
doge	0.75%	0.62%	0.11%	0.39%
eth	0.76%	0.06%	0.25%	0.25%
iota	0.36%	0.25%	0.14%	0.27%
link	1.24%	0.85%	0.31%	0.31%
ltc	0.67%	0.26%	0.40%	0.58%
trx	0.34%	0.44%	0.25%	1.64%
xlm	0.73%	0.86%	0.77%	0.72%
xmr	0.65%	0.15%	0.38%	0.32%
xrp	0.31%	0.93%	0.13%	0.70%

Notes

- ¹ Sourced from <https://www.vox.com/recode/2021/5/18/22441831/elon-musk-bitcoin-dogecoin-crypto-prices-tesla>, accessed on 1 August 2023.
- ² <https://huggingface.co/distilbert-base-uncased-finetuned-sst-2-english>, accessed on 12 January 2022.
- ³ <https://textblob.readthedocs.io/en/dev/>, accessed on 23 February 2022.
- ⁴ Mid-frequency trading involves adjusting the portfolio on a daily basis.

References

- Abraham, Jethin, Daniel Higdon, John Nelson, and Juan Ibarra. 2018. Cryptocurrency price prediction using tweet volumes and sentiment analysis. *SMU Data Science Review* 1: 13–22.
- Cabanilla, Kurt Izak M. 2016. The Future of Cryptocurrency: Forecasting the Bitcoin-Philippine Peso Exchange Rate Using Sarima Through Tramo-Seats. Available online: https://www.academia.edu/31926493/The_Future_of_Cryptocurrency_Forecasting_The_Bitcoin_Philippine_Peso_Exchange_Rate_Using_SARIMA_Through_TRAMO_SEATS (accessed on 1 January 2023).
- Challu, Cristian, Kin G. Olivares, Boris N. Oreshkin, Federico Garza, Max Mergenthaler-Canseco, and Artur Dubrawski. 2022. N-hits: Neural hierarchical interpolation for time series forecasting. *Proceedings of the AAAI Conference on Artificial Intelligence* 37: 6989–97. [CrossRef]
- Colianni, Stuart, Stephanie Rosales, and Michael Signorotti. 2015. Algorithmic trading of cryptocurrency based on twitter sentiment analysis. *CS229 Project* 1: 1–4.
- Dwivedi, DwijendraNath, and Anilkumar Vemareddy. 2023. Sentiment analytics for crypto pre and post covid: Topic modeling. Paper presented at Distributed Computing and Intelligent Technology: 19th International Conference, ICDCIT 2023, Bhubaneswar, India, January 18–22; pp. 303–15.

- Garg, Amish, Tanav Shah, Vinay Kumar Jain, and Raksha Sharma. 2021. Cryptop12: A dataset for cryptocurrency price movement prediction from tweets and historical prices. Paper presented at 2021 20th IEEE International Conference on Machine Learning and Applications (ICMLA), Pasadena, CA, USA, December 13–16; pp. 379–84.
- Gupta, Hemendra, and Rashmi Chaudhary. 2022. An empirical study of volatility in cryptocurrency market. *Journal of Risk and Financial Management* 15: 513. [CrossRef]
- Hochreiter, Sepp, and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation* 9: 1735–80. [CrossRef] [PubMed]
- Huang, Xin, Wenbin Zhang, Xuejiao Tang, Mingli Zhang, Jayachander Surbiryala, Vasileios Iosifidis, Zhen Liu, and Ji Zhang. 2021. Lstm based sentiment analysis for cryptocurrency prediction. Paper presented at International Conference on Database Systems for Advanced Applications, Tianjin, China, April 11–14; pp. 617–21.
- Hutto, Clayton, and Eric Gilbert. 2014. Vader: A parsimonious rule-based model for sentiment analysis of social media text. Paper presented at International AAAI Conference on Web and Social Media, Ann Arbor, MI, USA, June 1–4; vol. 8, pp. 216–25.
- Iyer, Tara. 2022. *Cryptic Connections: Spillovers between Crypto and Equity Markets*. Washington, DC: International Monetary Fund.
- Kim, Gyeongmin, Minsuk Kim, Byungchul Kim, and Heuseok Lim. 2023. Cbits: Crypto bert incorporated trading system. *IEEE Access* 11: 6912–21. [CrossRef]
- Laboure, Marion, Markus H.-P. Müller, Gerit Heinz, Sagar Singh, and Stefan Köhling. 2021. Cryptocurrencies and cdbc: The route ahead. *Global Policy* 12: 663–76. [CrossRef]
- Li, Xiaodong, Haoran Xie, Li Chen, Jianping Wang, and Xiaotie Deng. 2014. News impact on stock price return via sentiment analysis. *Knowledge-Based Systems* 69: 14–23.
- Mittal, Anshul, and Arpit Goel. 2012. *Stock Prediction Using Twitter Sentiment Analysis*. Stanford: Stanford University. p. 2352. Available online: <http://cs229.stanford.edu/proj2011/GoelMittal-StockMarketPredictionUsingTwitterSentimentAnalysis.pdf> (accessed on 1 January 2023).
- Nagel, Peter. 2018. Psychological Effects during Cryptocurrency Trading. Available online: <https://space.nurdspace.nl/~buzz/MasterThesisPeter.pdf> (accessed on 1 January 2023).
- Nguyen, Thien Hai, Kiyooki Shirai, and Julien Velcin. 2015. Sentiment analysis on social media for stock movement prediction. *Expert Systems with Applications* 42: 9603–11.
- Patel, Mohil Maheshkumar, Sudeep Tanwar, Rajesh Gupta, and Neeraj Kumar. 2020. A deep learning-based cryptocurrency price prediction scheme for financial institutions. *Journal of Information Security and Applications* 55: 102583.
- Peng, Yaohao, Pedro Henrique Melo Albuquerque, Jader Martins Camboim de Sá, Ana Julia Akaishi Padula, and Mariana Rosa Montenegro. 2018. The best of two worlds: Forecasting high frequency volatility for cryptocurrencies and traditional currencies with support vector regression. *Expert Systems with Applications* 97: 177–92. [CrossRef]
- Prasad, Gaurav, Gaurav Sharma, and Dinesh Kumar Vishwakarma. 2022. Sentiment analysis on cryptocurrency using youtube comments. Paper presented at 2022 6th International Conference on Computing Methodologies and Communication (ICCMC), Erode, India, March 29–31; pp. 730–33.
- Sanh, Victor, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv arXiv:1910.01108*.
- Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in Neural Information Processing Systems*. Available online: <https://mitpress.mit.edu/9780262561457/advances-in-neural-information-processing-systems/> (accessed on 1 January 2023).
- Vo, Anh-Dung, Quang-Phuoc Nguyen, and Cheol-Young Ock. 2019. Sentiment analysis of news for effective cryptocurrency price prediction. *International Journal of Knowledge Engineering* 5: 47–52. [CrossRef]
- Wong, Eugene Lu Xian. 2021. Prediction of Bitcoin Prices Using Twitter Data and Natural Language Processing. Available online: <https://dukespace.lib.duke.edu/dspace/bitstream/handle/10161/24081/Prediction%2520of%2520Bitcoin%2520prices%2520using%2520Twitter%2520Data%2520and%2520Natural%2520Language%2520Processing.pdf?sequence=2> (accessed on 1 January 2023).

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.