

From On-chain to Macro: Assessing the Importance of Data Source Diversity in Cryptocurrency Market Forecasting

Giorgos Demosthenous
University of Cyprus
Nicosia, Cyprus
demosthenous.giorgos@ucy.ac.cy

Chryssis Georgiou
University of Cyprus
Nicosia, Cyprus
georgiou.chryssis@ucy.ac.cy

Eliada Polydorou
University of Cyprus
Nicosia, Cyprus
polydorou.eliada@ucy.ac.cy

ABSTRACT

This study investigates the impact of data source diversity on the performance of cryptocurrency forecasting models by integrating various data categories, including technical indicators, on-chain metrics, sentiment and interest metrics, traditional market indices, and macroeconomic indicators. We introduce the Crypto100 index, representing the top 100 cryptocurrencies by market capitalization, and propose a novel feature reduction algorithm to identify the most impactful and resilient features from diverse data sources. Our comprehensive experiments demonstrate that data source diversity significantly enhances the predictive performance of forecasting models across different time horizons. Key findings include the paramount importance of on-chain metrics for both short-term and long-term predictions, the growing relevance of traditional market indices and macroeconomic indicators for longer-term forecasts, and substantial improvements in model accuracy when diverse data sources are utilized. These insights help demystify the short-term and long-term driving factors of the cryptocurrency market and lay the groundwork for developing more accurate and resilient forecasting models.

VLDB Workshop Reference Format:

Giorgos Demosthenous, Chryssis Georgiou, and Eliada Polydorou. From On-chain to Macro: Assessing the Importance of Data Source Diversity in Cryptocurrency Market Forecasting. VLDB 2024 Workshop: Foundations and Applications of Blockchain (FAB).

VLDB Workshop Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://github.com/gdemos01/FAB-2024>.

1 INTRODUCTION

Emerging in the catastrophic aftermath of the 2008 financial crisis [21], which eroded trust in traditional financial institutions, Bitcoin not only challenged the established financial infrastructure but also pioneered the first pure digital asset class [18], paving the way for what would eventually become a flourishing industry. The cryptocurrency market grew exponentially in recent years, with now more than 10000 assets circulating [1]. This market has experienced rapid mainstream adoption, with an estimated 580 million people holding cryptocurrency by the end of 2023, marking a 34%

year-over-year increase [8]. The unique attributes and operational dynamics of cryptocurrencies, such as 24/7 trading availability, peer-to-peer transactions, minimal transaction fees, borderless nature, and enhanced privacy, make them particularly attractive to all levels of investors, from retail to institutional.

Despite the rising participation in cryptocurrency markets and the increasing volume of work around modeling and trading crypto assets [9], building accurate and resilient forecasting models for this market remains a real challenge. This can be attributed to the non-linear, highly volatile, uncertain, and noisy nature of the cryptocurrency market. The low volume of historical data surrounding this market also contributes to the problem, further highlighting the need to search for alternative information avenues. Most literature surrounding forecasting in cryptocurrency markets, focuses on incorporating technical indicators (e.g., OHLC candles) and sentiment analysis (e.g., social media sentiment) to build models capable of predicting market movements [2, 19]. This approach is problematic because not only it underestimates the noise in this market and the necessity for additional, more valuable data sources, but also does not take advantage of data native to this market.

On-chain data is a new family of metrics extracted from the underlying blockchain technology of cryptocurrencies, which offers a unique insight into investor behavior, network activity, and the supply dynamics of assets in this market. Realizing the importance of on-chain data, recent work has been incorporating it more and more in building their forecasting models and for better understanding how this market operates [4, 13, 23]. In addition to technical, on-chain and sentiment metrics, there are other categories of data (e.g., macroeconomic and traditional market metrics), that when combined could further aid in deriving valuable information from the market, reducing noise, and consequently improving the performance of forecasting models.

Paradoxically, most literature only combines a couple of data categories, usually chosen through intuition and without any prior experimentation (e.g., [13, 19]), which might severely undermine the potential of their forecasting models. In this work, we focus on analyzing a variety of different data categories and showcasing the impact that data source diversity can have on the performance of cryptocurrency forecasting models. We also investigate how the predictive power of each data category, as well as individual indicators, changes as prediction difficulty increases, laying the groundwork for better understanding this market and consequently building more accurate and resilient forecasting models, capable of both short-term and long-term predictions. Specifically, our **contributions** are the following:

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.
Proceedings of the VLDB Endowment. ISSN 2150-8097.

Crypto100 index: We introduce the Crypto100 index, a new market index that tracks the top 100 cryptocurrencies by market capitalization. This index was designed to be directly comparable with any other asset in this market, and a good representation of the entire cryptocurrency market.

Feature reduction algorithm: We present a novel algorithm for creating feature vectors consisting of the most impactful and persistent features across a variety of complementary importance evaluation methods. This methodology ensures the identification of the most important and resilient features from a diverse set of data sources, enhancing model accuracy.

Impact analysis on data source diversity: Our comprehensive experiments illustrate the significant influence of data source diversity on the predictive performance of forecasting models. We examine various time horizons and scenarios, including short-term (1 and 7 days) and long-term (90 and 180 days) prediction windows, highlighting the changing importance of different data categories and individual indicators across these periods.

2 BACKGROUND

2.1 Forecasting Models

Forecasting models are essential tools for predicting future market behavior based on historical data and current market inputs. These models, which include time series and regression analysis, econometric models, and machine learning algorithms, are crucial for informed decision-making, risk management, resource allocation, and performance evaluation in investment strategies [9]. An additional benefit of building these models is their ability to extract useful information from noisy markets, leading to a better understanding of the factors driving market movements. In the context of cryptocurrency markets, forecasting models are most commonly used for predicting price movements and are integrated into trading strategies [20]. In recent years, these models have also been employed to build more resilient cryptocurrency portfolios [16].

2.2 Categories of Data Sources

Different categories of data, ranging from macroeconomic indicators to on-chain metrics, offer distinct and diverse perspectives on the market, each providing unique insights into market behavior. In the context of this study, we have separated market information into the following categories:

Macroeconomic Indicators: These indicators encompass broader economic metrics and global financial data that can influence a variety of markets, including the cryptocurrency market. Macro indicators help contextualize cryptocurrency market movements within the larger economic climate. Examples of these indicators include central bank interest rates, inflation rates, policy uncertainty indices, and others.

Technical Indicators: Derived from the technical analysis of historical market information such as price and volume, these indicators aid in identifying patterns within the raw market data. Examples of technical indicators include moving averages, the relative strength index (RSI), and Bollinger Bands

Sentiment and Interest Metrics: This category aims to capture the mood, opinions, and interest of market participants from various sources like social media, news articles, and search trends.

Shifts in investor sentiment can provide insights into the upcoming state of the markets. Common metrics include social media posts volume, google trends, social media post sentiments (e.g., negative, neutral, positive), and the fear and greed index.

Traditional Market Indices: These are major financial indices from traditional finance that can provide insights into a variety of markets as well as broader economic conditions. These indices include bonds (e.g., BSV [25]), metals (e.g., gold), exchange rates (e.g., EURUSD), currency strength (e.g., U.S. Dollar index [11]) and stocks (e.g., Nasdaq-100 index - QQQ).

On-chain Metrics: This new family of data points was introduced with the emergence of blockchain technology. These metrics are derived directly from the blockchain and offer unprecedented insights into market behavior by monitoring network activity and statistics. Examples of on-chain metrics include transaction volume, miner revenue, hash rate, active addresses, token distributions, supply information, network fees, and others.

3 METHODOLOGY

3.1 Datasets

To assess the importance of data source diversity, historical data were collected and incorporated into various datasets used to evaluate the significance of individual metrics under different scenarios. For each data category, daily data were collected from sources such as Coinmetrics.io [7], CoinGecko.com [1], the European Central Bank (ECB) [3], and others [14, 24]. Technical indicators were constructed using only BTC historical market information, as we found them to be highly correlated with and influential on the broader cryptocurrency market. Building on this observation, it was assumed and later validated (Subsection 4.1) that Bitcoin’s on-chain metrics could also affect the rest of the market, and were therefore included in the dataset as market representatives. To investigate the importance of stablecoins, on-chain information from the second largest stablecoin by market capitalization, USDC, were also gathered. Although USDT has a much higher marketcap we chose not to include it due to its bad reputation and questionable activity [10], which might negatively influence this study leading to misleading conclusions. To collect on-chain information, we used Coinmetric’s free API instead of running nodes on the blockchain networks to manually collect transaction data and calculate the respective metrics. Although this approach is easier and less costly, future research that requires more advanced on-chain metrics might need to consider directly running nodes on the networks instead. The data collection period spanned from January 2017 to June 2023.

3.1.1 The Crypto100 index. In contrast with most previous work which focuses solely on Bitcoin, Ethereum or a minor collection of popular cryptocurrencies, this study aims to capture a better representation of the entire crypto market. To achieve this, the Crypto100 index was created. This index tracks the top 100 cryptocurrencies by market capitalization, on a daily basis, and is calculated by the following equation:

$$\text{Crypto100} = \frac{\sum_{i=1}^{100} \text{Market Cap}_i}{\left(\log_{10} \left(\sum_{i=1}^{100} \text{Market Cap}_i\right)\right)^7}$$

The top 100 cryptocurrencies constitute the majority of the market, as evidenced by their cumulative market capitalization (cf. Figure 1), making the Crypto100 index a good representation of the entire cryptocurrency market. Inspired by the calculation methodology of the S&P500 index, which uses a divisor to normalize the final value and maintain index continuity over time [12], Crypto100 employs a scaling factor derived from extensive experimentation. This approach ensures the Crypto100 index maintains continuity and resilience to sudden changes in the assets comprising the top 100 list, which is common in a maturing market like the cryptocurrency market. To make the Crypto100 index price directly comparable to other assets in the cryptocurrency market, we used Bitcoin’s price to tune the scaling factor. Figure 2b shows that when the scaling factor is raised to powers lower than 6, it significantly limits its impact, making the Crypto100 index less comparable to other assets, such as Bitcoin. Thus, as illustrated in Figure 2a, the chosen power for tuning the scaling factor ensures the Crypto100 index is appropriately scaled and comparable to other assets in the market. For the index to maintain its resilience in this evolving market as time passes, the scaling factor should be reevaluated periodically and adjusted dynamically if necessary.



Figure 1: Top 100 Cryptocurrencies VS Total Marketcap

3.1.2 *Scenarios.* A significant challenge in a maturing market, such as the cryptocurrency market, is that experiments conducted over different chronological periods can yield varying results [17]. Bearing this in mind, alongside investigating the entire 2017-2023 period, a subset of the original dataset starting from January 2019 was also created. The rationale for selecting January 2019 as the starting date for the subset is two-fold: Firstly, this month marked the bottom of the previous bull market, but more importantly, many critical data points, such as the USDC stablecoin and the fear-and-greed index, were only introduced after the end of 2018.

To examine how data source diversity impacts forecasting models across short-term and long-term horizons, fine-tuned machine learning (ML) models were trained to predict the Crypto100 index price 1, 7, 30, 90, and 180 days in advance, which are *prediction windows* commonly found in literature [9, 20]. The importance of individual indicators for each prediction window was also extracted and analyzed. Initially, the datasets consisted of 429 metrics across five categories, with on-chain data separated into two subcategories:



(a) Scaling factor raised to the 7th and 8th power vs BTC price



(b) Scaling factor raised to the 6th and 7th power vs BTC price

Figure 2: Comparison of Crypto100 versions using different scaling factors with Bitcoin’s price

On-chain_BTC and *On-chain_USDC*. Following an initial data cleaning and preprocessing phase that included the standard methods used in ML [6] such as filling empty data with interpolation, removing duplicate values, and discarding features that had flat or missing values for very long periods, two new subsets were created from the original dataset: *set 2017*, covering the period between January 2017 and June 2023, and *set 2019*, covering the period between January 2019 and June 2023. Metrics that began recording their values after the initial date of each period (e.g., USDC starting in 2018), were discarded from the corresponding set. As a result, the 2017 and 2019 sets comprised 192 and 283 metrics, respectively.

3.2 Feature Selection

To narrow down the focus of the investigation to the most important indicators for each scenario, a robust feature selection methodology was devised, combining *four* different methods: Pearson correlation, Mean Decrease Impurity (MDI), Permutation Feature Importance (PFI), and SHapley Additive exPlanations (SHAP) [15]. Although the correlation between features and the target variable provides a good initial indication of the importance of individual features, it cannot be used as a standalone method since it fails to capture non-linear relationships. Therefore, feature selection algorithms based on machine learning models were also incorporated into our

methodology. Specifically, Random Forests (RF) and XGBoost (XGB) were fine-tuned using 5-fold cross-validation grid search [22] with minimum mean squared error (MSE) as the objective for each of the 10 different scenarios (2 sets x 5 prediction windows). The search space included parameters relevant to tree structures like number of estimators, maximum depth, sample splits, etc. These models were then used to calculate feature importance using their built-in MDI-based algorithm. Additionally, PFI was also extracted for both models using MSE as the optimization measure. In contrast to MDI, PFI directly measures the effect on each model’s predictive performance, mitigating issues caused by bias during training.

The Feature Reduction Algorithm (Algorithm 1) is a novel methodology for creating feature vectors consisting of the most impactful and persistent features across a variety of different importance evaluation methods. The objective of the algorithm is to iteratively remove features that consistently rank at the bottom of each inner evaluation method until the dimensionality of the vector is reduced below the desired number. Due to the low volume of records in our dataset, we needed to substantially reduce the feature vector’s length to help models distinguish valuable information from noise. To retain important indicators, we set the target length at 100, ensuring we could still examine the feature vector comprehensively, as significant insights might be present throughout, including the tail end. The inner methods used are MDI and PFI with two different models (RF and XGB), as well as Pearson correlation. The correlation value, initially set to 0.5, is used as a threshold for removal of surviving bottom features and increases with each iteration by 0.025 to make the selection process stricter until the algorithm’s objective is met. The strict approach using complementary importance evaluation methods of the Feature Reduction Algorithm (FRA) ensures that each method’s strengths are leveraged while mitigating their individual weaknesses, thereby limiting the possibility of noisy features appearing in the final selection.

Algorithm 1 Feature Reduction Algorithm (FRA) for selecting the most important indicators across all categories

```

1: periods  $\leftarrow$  [2017, 2019]
2: prediction_window  $\leftarrow$  [1, 7, 30, 90, 180]
3: for each period do
4:   for each prediction window do
5:     corr_threshold  $\leftarrow$  0.5
6:     while number_of_features > 100 do
7:       RF_FI  $\leftarrow$  Extract importance using RF (MDI)
8:       XGB_FI  $\leftarrow$  Extract importance using XGB (MDI)
9:       RF_PFI  $\leftarrow$  Extract importance using RF (MSE)
10:      XGB_PFI  $\leftarrow$  Extract importance using XGB (MSE)
11:      if feature_rank = bottom_50% AND feature_cor
< corr_threshold then
12:        Remove corresponding features
13:      end if
14:      corr_threshold  $\leftarrow$  corr_threshold + 0.025
15:    end while
16:    return reduced feature vector
17:  end for
18: end for

```

Table 1: Summary of final feature vectors for all different scenarios (year_prediction window)

Scenario	Number of Features
2017_1	79
2017_7	79
2017_30	81
2017_90	86
2017_180	88
2019_1	100
2019_7	97
2019_30	100
2019_90	91
2019_180	90

To validate the efficacy of FRA and to mitigate possible drawbacks of the algorithm, SHAP was used to compute the contribution of individual features from the original sets, and combine its output with FRA to export the final feature vector. Although SHAP was originally used as a method to help interpret the output of machine learning models, in recent years it has been gaining popularity mainly as a feature importance evaluation approach. It works by computing the Shapley value for each feature, which is derived from cooperative game theory and represents the average contribution of a specific feature across all possible combinations of features. The overlap between the top 100 features extracted by SHAP and the reduced feature vector of FRA (< 100) is on average 78 features, which further validates that the surviving features are indeed the most important ones. To create the final feature vector for each scenario (Table 1), we take the union of the top 75 features ranked by importance from FRA and SHAP.

3.3 Impact Demonstration

Several approaches were followed to demonstrate the impact of data source diversity. Initially, the contribution of individual data categories to the final feature vector is analyzed across all prediction windows to discover how it changes over time (Section 4.1). We then group scenarios together into short-term and long-term time horizons and proceed to analyze both the top and unique indicators found in each use case (Section 4.2). To investigate the explicit effect of data source diversity on forecasting models, we fine-tune and train ML models for each scenario and explore how their performance changes when using the diverse feature vector compared to only using each individual category (Section 4.3).

4 RESULTS

4.1 The Contribution of Individual Data Sources

The importance of an individual data category in forecasting models can vary significantly depending on the prediction window, with the impact of different indicators changing as the time horizon moves from short-term to long-term predictions. To better understand this changing dynamic between indicator categories, the contribution of each data source to the final feature vector, extracted using the FRA algorithm, across all prediction windows was recorded for both set 2017 (Figure 3) and set 2019 (Figure 4). To make the categories

comparable, we calculate their contribution by dividing the final number of features from the category included in the final vector with the corresponding total number of candidate features in the same category before the feature selection phase took place. In other words, the *contribution factor* is the ratio of features from each category included in the final feature vector.

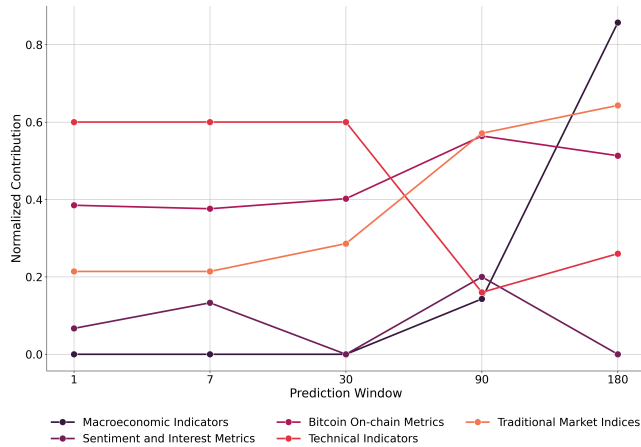


Figure 3: Contribution of individual data sources to the final feature vector in set 2017 across all prediction windows

Looking at Figures 3 and 4 it immediately becomes apparent that *On-chain Metrics* maintain a very high contribution to the final feature vector across all prediction windows. This behavior demonstrates that this new data family not only seems to be very important for cryptocurrency forecasting models, but also appears to be suitable for short-, medium- and long-term predictions.

Technical Indicators on the other hand appear to be very valuable for short-term predictions, but their standalone predictive power quickly diminishes over longer-periods. Their ability to capture evolving market patterns and trends makes them essential for short-term windows, but their decreasing contribution over longer windows suggests that technical indicators need to be supplemented with other data categories to build models capable of long-term forecasting.

Although the contribution of *Traditional Market Indices* is initially at low levels, it rapidly increases for longer prediction windows, eventually becoming the second-highest contributor of indicators to the final feature vector in both 2017 and 2019 sets. The growing relevance of traditional market indices over longer periods highlights the increasing integration and correlation between traditional financial markets and cryptocurrencies. As the cryptocurrency market matures, it becomes more sensitive to broader market conditions and trends, making traditional indices more relevant for long-term predictions.

As depicted in Figure 3, *Macroeconomic Indicators* appear to have no effect in short-term predictions, but their contribution rises in long-term prediction windows. This can be attributed to the delayed effect of economic policies and other macroeconomic factors that gradually permeate all financial markets, including cryptocurrencies. Paradoxically, macroeconomic indicators are completely

absent from the 2019 set (Figure 4), which might be due to any of the following reasons: (i) the decreased size of the 2019 datasets doesn't include significant changes in macroeconomic indicators, which usually become evident over longer time horizons, (ii) the cryptocurrency market in certain time periods (e.g., periods of rapid increase of mainstream adoption), might become more self-contained and independent of broader economic conditions compared to other markets, and (iii) the dominance of other factors such as the inclusion of USDC metrics might simply overshadow macroeconomic influences. The discrepancies observed between the 2017 and 2019 sets further validate the assumption that studying different periods in the cryptocurrency market can sometimes yield varying results. To confirm that these differences are due to changing market behavior and not noise, future research could focus on enhancing FRA by incorporating more dynamic elements, thereby increasing its robustness.

The introduction of stablecoins such as USDC to the cryptocurrency markets aided in increasing market stability and liquidity [5]. Naturally, *On-chain Metrics (USDC)* have a very important role to play in forecasting market movements across all time horizons (Figure 4). Their contribution substantially surpasses even that of Bitcoin's on-chain data, becoming by far the most important contributor of indicators to the final feature vector for 30, 90 and 180 day prediction windows. The reason for this could be that USDC data encapsulate information about macro changes of the crypto market, as well as market sentiment and investor behavior. Inflow/outflow information of the USDC stablecoin can be good indicators of changes in the buying/selling behaviors of investors, profit taking periods, moving funds in and out of the cryptocurrency market, etc. Because of the increasing importance of USDC-based on-chain data and the evolving impact of stablecoins on the cryptocurrency market, further research into the dynamics of individual indicators within this category is needed.

Sentiment and Interest Metrics have a high contribution for short-term predictions, with their effect diminishing on average over longer periods. This could be attributed to immediate market reactions driven by social media, news, and investor interest, which can cause short-term price changes. Although some indicators in this category, like monthly search volume measured by Google trends, provide intelligence about longer time horizons (hence the increased contribution for 90-day window), the majority supply information that can drive short-term price volatility but cannot sustain long-term trends.

4.2 Short-term and Long-term Driving Factors

To delve deeper into which individual metrics affect the predictive capabilities of forecasting models, two groups were created based on the final feature vectors. Specifically, features from sets with different prediction windows were merged together to form the *Short-term* (windows 1 and 7) and *Long-term* (windows 90 and 180) groups. Fine-tuned random forests were used to extract an importance value for each feature in the final feature vectors. When creating the groups, the importance of common features within a group was calculated by averaging their importance value before merging. To better understand which individual indicators affect each group and why, their brief definitions are provided in Table 2.

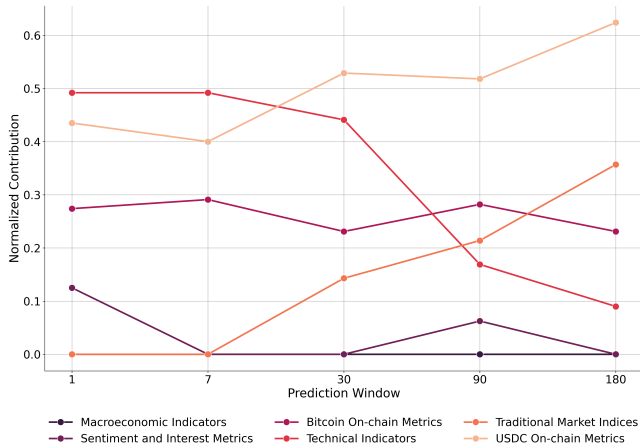


Figure 4: Contribution of individual data sources to the final feature vector in set 2019 across all prediction windows

The *top 5 features* of each group and for each set, ranked by importance, can be seen in Table 3. Following market trends (e.g., through moving averages) seems to be very important for short-term predictions (EMA100/EMA200). A possible reason for this, is that they provide a clear insight about the overall market direction which, when combined with other indicators, can improve prediction performance. The all-time revenue generated by miners (RevAllTimeUSD) is important across all scenarios, present in almost all feature vectors, and even ranks amongst the most important features for short-term and long-term predictions in sets 2017 and 2019 respectively. High miner revenue can signal strong network usage, which not only increases value and confidence in the cryptocurrency market but can also potentially influence price movements. Realized market capitalization (CapRealUSD) provides a better insight into the fair value of a crypto asset based on the actual transaction history rather than the current price, hence the observed increased importance. This indicator shows how much of Bitcoin’s supply has been moved recently offering clues about how the market values the asset and where the price is directed. Although in the 2017 set retail investors are very important (AdrBalUSD100Cnt) in the short-term, medium and larger holders (fish_pct, SplyAdrBalUSD10K, usdc_AdrBalNtv10KCnt) also seem to have growing influence on short-term market movements based on the 2019 set, which could indicate that participation by wealthier and more accredited investors has increased in recent years. Looking at the top features in the Long-term group, it becomes clear that on-chain metrics based on balance and supply dynamics of the market have a lot of predictive power. Wealth distribution among smaller holders (SplyAdrBalUSD100, SplyAdrBalNtv0.01, SplyAdrBalUSD1, etc.) is surprisingly what drives longer-term market movements, despite the gradual increase of larger holders. Metrics about existing and active supply (SplyCur, SplyActEver) are also important long-term indicators because they encapsulate information about inflation and scarcity which do affect an asset’s price in the longer timeframe.

To investigate which features are uniquely important in each group, the 20 features with the highest importance that appeared in one group but not the other were selected and can be seen in

Table 2: Definitions of the features used in our analysis

Feature	Definition
(usdc)_AdrBal1in#Cnt	Sum count of unique addresses holding at least one <i>nth</i> ownership of the current supply of USDC/BTC
(usdc)_AdrBalUSD#Cnt	Sum count of unique addresses holding at least \$# worth of USDC/BTC)
(usdc)_AdrBalNtv#Cnt	Sum count of unique addresses holding at least # USDC/BTC
(usdc)_SplyActPct#yr	Percentage of the current USDC/BTC supply that has been active in the trailing # year
(usdc)_SplyActEver	Sum of unique USDC/BTC held by accounts that transacted at least once
(usdc)_SplyAdrBalNtv#	Sum of all USDC/BTC being held in addresses whose balance is # native units or greater
(usdc)_SplyAdrBalUSD#	Sum of all USDC/BTC being held in addresses whose balance is \$# or greater
(usdc)_SplyCur	The current supply of USDC/BTC
s2f_ratio	The stock-to-flow ratio of BTC
(usdc)_CapAct1yrUSD	Sum USD value of all active USDC/BTC in the last year
CapRealUSD	The sum USD value based on the USD closing price on the day that a native unit last moved
(usdc)_CapMrktFFUSD	Sum USD value of the current free float supply(USDC/BTC)
market_cap	The market capitalization of BTC
EMA#_variable	The exponential moving average of a <i>variable</i> based on the # past days
SMA#_variable	The simple moving average of a <i>variable</i> based on the # past days
fish_pct	The ratio of wallets holding 10-100 BTC
shrimps_pct	The ratio of wallets holding less than 10 BTC
total_balance	Sum of all BTC being held in addresses owned by whales, sharks, fish or shrimp
gt_X_monthly	The monthly search volume of <i>X</i> in Google trends
X_Close	The close price of index/indicator <i>X</i>
RevAllTimeUSD	Sum USD value of all Bitcoin miner revenue (fees plus newly issued BTC) since genesis
RevHashRateUSD	The USD value of the mean daily miner reward per estimated hash unit per second performed during the day
SER	Supply Equality Ratio. The ratio of supply held by addresses with less than one ten-millionth of the current supply of USDC/BTC to the supply held by the top one percent of addresses.
SplyMiner0HopAllUSD	The sum of the balances of all mining entities in USD
VelCur1yr	The ratio of the value transferred in the trailing 1 year divided by the current supply

Table 4. Indicators around recent market movements dominate the unique driving factors of short-term forecasting markets. More recent SMAs/EMAs around price and market capitalization, with a past window of 5 to 30 trailing days, are the most populant in the Short-term group. Broader economic conditions evidenced by major indices appear to have a long-term effect in the predictability of the cryptocurrency market. These indexes include top

Table 3: Top 5 most important features for short-term and long-term Predictions in 2017 and 2019 sets

Set	Short-term	Long-term
2017	EMA100_market-cap	SplyAdrBalUSD100
	EMA200_close-price	SplyAdrBalNtv0.01
	RevAllTimeUSD	SplyCur
	CapRealUSD	SplyAdrBalUSD1K
	AdrBalUSD100Cnt	SplyActEver
2019	usdc_AdrBalNtv1Cnt	AdrBalNtv0.1Cnt
	fish_pct	SplyAdrBalUSD10
	SplyAdrBalUSD10K	RevAllTimeUSD
	usdc_AdrBalNtv10KCnt	SplyAdrBalNtv0.001
	total_balance	SplyAdrBalUSD1

non-financial companies (QQQ_Close), the strength of the dollar (UUP_close) and the euro (EURUSD_Close), as well as the bond markets (BSV_CLOSE, MBB_close). The unique-feature analysis also further validates that supply dynamics and market activity metrics (SplyActPct1yr, SER, VelCur1yr, s2f_ratio) can substantially affect long-term predictions. It is apparent that stablecoin data, in this case On-chain USDC metrics, have a major role to play in both short- and long-term predictions as evidenced by the sheer number of respective features present in the top 20 unique feature groups. In the short-term, USDC metrics focus on address balances (e.g., usdc_AdrBalNtv10KCnt), which can indicate the recent activity and behavior of market participants, such as investors entering/leaving the market. In the long-term, USDC metrics highlight wealth distribution (e.g., usdc_SplyAdrBalNtv100), supply activity (e.g., usdc_SplyAct2yr), and supply dynamics (e.g., usdc_SplyCur), which are indicators of long-term market stability and direction.

4.3 Model Performance Improvement

The results discussed up to this point have proven that combining metrics from different categories not only provides additional insight into important driving factors of short- and long-term market movements, but is also important in improving the predictive performance of forecasting models. To quantify the impact of data source diversity on the performance of forecasting models, ML models were fine-tuned and trained for each scenario. We then explored how their performance changes when using the diverse feature vector compared to only using data from a single category. *Performance improvement* is defined as the percentage decrease of the mean squared error after evaluating the model on the diverse feature vector. Tables 5 and 6 include the average performance improvements of a RF model fine-tuned, using 5-fold cross-validation grid search, and trained for each respective scenario.

In Table 5, the average model improvement for each prediction window under both 2017 and 2019 sets is presented. With the notable exception of the shortest prediction window (w=1), the impact of data source diversity on the forecasting model seems to be steadily increasing over longer time periods. This is because predicting the future price of Crypto100, as well as any other price metric, naturally becomes more challenging as the prediction window increases, due to larger amounts of uncertainty regarding the future of the market, and therefore requires the additional insight

Table 4: Top 20 unique features for short-term and long-term predictions in 2017 and 2019 sets

Set	Short-term	Long-term
2017	AdrBalUSD100Cnt	QQQ_Close
	EMA14_close-price	AdrBalNtv0.1Cnt
	EMA10_close-price	AdrBal1in1BCnt
	EMA10_market-cap	AdrBalNtv0.01Cnt
	AdrBal1in1MCnt	SplyActPct1yr
	EMA20_close-price	AdrBal1in10KCnt
	EMA20_market-cap	UUP_Close
	EMA14_market-cap	EMA200_volume
	market-cap	gt_Ethereum_monthly
	EMA5_market-cap	AdrBal1in10BCnt
	SMA_20_close-price	EURUSD_Close
	SMA_10_market-cap	SER
	EMA5_market-cap	EMA100_volume
	SMA_20_market-cap	RevHashRateUSD
	SMA_10_close-price	BSV_Close
	EMA30_close-price	VelCur1yr
	CapAct1yrUSD	s2f_ratio
	SMA_5_close-price	shrimps_pct
	SplyMiner0HopAllUSD	MBB_Close
	EMA30_market-cap	SplyAdrBal1in1M
2019	fish_pct	AdrBalNtv0.1Cnt
	SplyAdrBalUSD10K	usdc_SplyAdrBalNtv100
	usdc_AdrBalNtv10KCnt	SplyActPct1yr
	usdc_AdrBalUSD1MCnt	usdc_SplyAdrBalNtv0.001
	usdc_AdrBalUSD100KCnt	usdc_SplyCur
	AdrBalUSD100Cnt	UUP_Close
	SplyAdrBalUSD10M	EURUSD_Close
	usdc_AdrBalUSD10KCnt	SER
	usdc_SplyAct7d	usdc_SplyAdrBal1in100M
	AdrBalUSD1Cnt	usdc_SplyAct2yr
	SplyAdrBal1in100K	SplyAct1yr
	gt_Cryptocurrency_m	gt_Ethereum_monthly
	SplyAdrTop1Pct	usdc_CapMrktFFUSD
	EMA14_close-price	usdc_SplyAdrBalUSD10
	SplyAdrBalUSD1M	EMA200_volume
	SplyAdrBalNtv100	ROI1yr
	EMA10_market-cap	usdc_SplyAdrBalNtv1
	usdc_SplyAdrBalUSD10M	AdrBal1in1BCnt
	EMA20_market-cap	usdc_SplyAct3yr
	EMA14_market-cap	

that arises from the relationships between diverse data categories. One explanation about the shortest window exception could be that certain categories like *Sentiment and Interest Metrics* as well as *Technical Indicators* provide information which can mostly aid in predicting immediate market reactions, but the importance of which quickly fades in longer prediction timeframes as discussed in Section 4.1. Table 6 includes the average model improvements for each data category under both 2017 and 2019 sets. Categories that had the least representation in the final feature vector, as well as the initial dataset, such as *Sentiment and Interest Metrics* and *Macroeconomic Indicators*, benefited the most from data source diversity with the improvement level reaching all the way up to 1118.16%. *Bitcoin On-chain Metrics* on the other hand showcased the least

Table 5: Average MSE percentage decrease of the RF model by prediction window for both 2017 and 2019 sets

Prediction Window	Improvement (%)	
	2017	2019
1	855.87%	794.71%
7	189.08%	191.51%
30	218.96%	274.74%
90	378.23%	359.81%
180	636.24%	512.59%

Table 6: Average MSE percentage decrease of the RF model by data category for both 2017 and 2019 sets

Prediction Window	Improvement (%)	
	2017	2019
Macroeconomic Indicators	825.72%	913.55%
Sentiment and Interest Metrics	1118.16%	895.61%
On-chain Metrics (BTC)	12.09%	17.51%
Traditional Market Indices	244.38%	273.13%
Technical Indicators	78.02%	81.72%
On-chain Metrics (USDC)	-	378.52%

improvement, which was 12.09% and 17.51% for 2017 and 2019 sets respectively. One reason for this could be that on-chain metrics include data more relevant to the cryptocurrency market, which capture a variety of both technical and fundamental information (investor activity, supply and demand dynamics, etc.), therefore making some of the other data categories less necessary.

Overall, the performance of the RF forecasting model improved on average by 455.67% and 426.67% in 2017 and 2019 sets, respectively. To validate whether this substantial improvement of performance is also observed in other models, an XGBoost estimator was also fine-tuned, trained and evaluated for every scenario. In the case of XGB, the average overall improvement for the 2017 set was 399.67% and for the 2019 set was 468%, which constitutes more evidence that data source diversity can have a major impact on the predictive power of cryptocurrency market forecasting models.

5 CHALLENGES AND FUTURE WORK

Balanced category representation. A drawback of this study is that some data categories like *Sentiment and Interest Metrics* and *Macroeconomic Indicators* are underrepresented in the original dataset. This might have led to underestimating their importance in certain scenarios. In future work, an effort to achieve a more balanced representation across all data sources in the original dataset should be made. Detailed analysis of isolated categories could also provide additional insight into the impact of individual features within their category.

On-chain data diversification. Both BTC and USDC on-chain data proved very valuable in improving the performance of forecasting models. In addition to these two cryptocurrencies, on-chain data from other targeted assets should also be investigated. We

propose that future work should include on-chain data from cryptocurrencies that constitute good representatives of their market category (e.g., Ethereum for Decentralized Finance) providing more holistic insights about the cryptocurrency market.

Impact on complex models. We have already established that data source diversity is very important for ML models designed for the cryptocurrency market. The next step is to investigate the impact of diversity on more complex models and deep learning architectures, determining whether this diversity is beneficial or introduces unnecessary noise.

Application in finance. The outcome of this research should be used as the baseline for building predictive models for the cryptocurrency markets. Feature engineering techniques could also help discover valuable relationships between data categories, which might further enhance the predictive capabilities of such models. In future work, we plan to use the findings and datasets of this research to build novel portfolio optimization techniques that are resilient to the highly dynamic and uncertain nature of this market.

6 CONCLUSION

In this study, we have explored the significant role that data source diversity plays in enhancing the predictive performance of forecasting models in the cryptocurrency market. Our extensive experiments demonstrate how integrating a variety of data categories—including technical indicators, on-chain metrics, sentiment and interest metrics, traditional market indices, and macroeconomic indicators—can lead to substantial improvements in forecasting accuracy across different time horizons. We also introduced the Crypto100 index, a new market index that effectively represents the overall cryptocurrency market, and developed a novel algorithm to identify the most impactful and resilient features from a diverse set of data sources. Our findings highlight the changing importance of different data categories and individual indicators depending on the time period and prediction window, providing valuable insights for both short-term and long-term market predictions.

The results underscore the necessity of incorporating a diverse range of data sources to improve the robustness and accuracy of cryptocurrency forecasting models. Future work should focus on balancing the representation of different data categories, diversifying on-chain data, investigating the impact of data diversity on more complex models, and applying these findings to practical financial processes such as portfolio optimization. Overall, our study lays the groundwork for more accurate and resilient forecasting models, enhancing our understanding of the dynamic and volatile cryptocurrency market.

REFERENCES

- [1] 2024. Cryptocurrency Categories. <https://www.coingecko.com/en/categories>. Accessed: 2024-05-14.
- [2] Saúl Alonso-Monsalve, Andrés L Suárez-Cetrulo, Alejandro Cervantes, and David Quintana. 2020. Convolution on neural networks for high-frequency trend prediction of cryptocurrency exchange rates using technical indicators. *Expert Systems with Applications* 149 (2020), 113250.
- [3] European Central Bank. 2024. Harmonised Index of Consumer Prices (HICP). <https://data.ecb.europa.eu/data/datasets/ICP/ICP.M.U2.N.000000.4.ANR>.
- [4] Bruno Casella and Lorenzo Paletto. 2023. Predicting Cryptocurrencies Market Phases through On-Chain Data Long-Term Forecasting. In *2023 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*. IEEE, 1–4.

- [5] Christian Catalini, Alonso de Gortari, and Nihar Shah. 2022. Some simple economics of stablecoins. *Annual Review of Financial Economics* 14 (2022), 117–135.
- [6] Davide Chicco, Luca Oneto, and Erica Tavazzi. 2022. Eleven quick tips for data cleaning and feature engineering. *PLoS Computational Biology* 18, 12 (2022).
- [7] CoinMetrics. 2024. CoinMetrics: Crypto Intelligence for the Future of Finance. <https://coinmetrics.io/>. Accessed: 2024-06-11.
- [8] Crypto.com. 2024. *Crypto Market Sizing 2023*. Technical Report. Crypto.com. https://contenthub-static.crypto.com/wp_media/2024/01/Crypto-Market-Sizing-2023.pdf Accessed: 2024-05-03.
- [9] Fan Fang, Carmine Ventre, Michail Basios, Leslie Kanthan, David Martinez-Rego, Fan Wu, and Lingbo Li. 2022. Cryptocurrency trading: a comprehensive survey. *Financial Innovation* 8, 1 (2022), 13.
- [10] John M Griffin and Amin Shams. 2020. Is Bitcoin really untethered? *The Journal of Finance* 75, 4 (2020), 1913–1964.
- [11] Invesco. 2024. Invesco DB US Dollar Index Bullish Fund (UUP) Profile. <https://www.invesco.com/us/financial-products/etfs/product-detail?audienceType=Investor&ticker=UUP>. Accessed: 2024-06-11.
- [12] Investopedia. 2024. How Is the S&P 500 Index Calculated? <https://www.investopedia.com/ask/answers/05/sp500calculation.asp>. Accessed: 2024-06-11.
- [13] Han-Min Kim, Gee-Woo Bock, and Gunwoong Lee. 2021. Predicting Ethereum prices with machine learning based on Blockchain information. *Expert Systems with Applications* 184 (2021), 115480.
- [14] LunarCrush. 2024. LunarCrush: Social Intelligence for Crypto. <https://lunarcrush.com/>. Accessed: 2024-06-11.
- [15] Scott M Lundberg and Su-In Lee. 2017. A unified approach to interpreting model predictions. *Advances in neural information processing systems* 30 (2017).
- [16] Yechi Ma, Ferhana Ahmad, Miao Liu, and Zilong Wang. 2020. Portfolio optimization in the era of digital financialization using cryptocurrencies. *Technological forecasting and social change* 161 (2020), 120265.
- [17] Khaled Mokni, Ghassen El Montasser, Ahdi Noomen Ajmi, and Elie Bouri. 2024. On the efficiency and its drivers in the cryptocurrency market: the case of Bitcoin and Ethereum. *Financial Innovation* 10, 1 (2024), 39.
- [18] Satoshi Nakamoto. [n.d.]. Bitcoin: A Peer-to-Peer Electronic Cash System. ([n. d.]). www.bitcoin.org
- [19] Marco Ortu, Nicola Uras, Claudio Conversano, Silvia Bartolucci, and Giuseppe Destefanis. 2022. On technical trading and social media indicators for cryptocurrency price classification through deep learning. *Expert Systems with Applications* 198 (2022), 116804.
- [20] Yi-Shuai Ren, Chao-Qun Ma, Xiao-Lin Kong, Konstantinos Baltas, and Qasim Zureigat. 2022. Past, present, and future of the application of machine learning in cryptocurrency research. *Research in Intl Business and Finance* 63 (2022).
- [21] Anna J Schwartz. 2009. Origins of the financial market crisis of 2008. *Cato J.* 29 (2009), 19.
- [22] Scikit-learn. 2024. Grid Search. https://scikit-learn.org/stable/modules/grid_search.html. Accessed: 2024-06-09.
- [23] Helder Sebastião and Pedro Godinho. 2021. Forecasting and trading cryptocurrencies with machine learning under changing market conditions. *Financial Innovation* 7 (2021), 1–30.
- [24] Economic Policy Uncertainty. 2024. Measuring Economic Policy Uncertainty. <https://www.policyuncertainty.com>. Accessed: 2024-06-11.
- [25] Vanguard. 2024. Vanguard Short-Term Bond ETF (BSV) Profile. <https://investor.vanguard.com/investment-products/etfs/profile/bsv>. Accessed: 2024-06-11.