

Leveraging Explainable AI to Support Cryptocurrency Investors

Original

Leveraging Explainable AI to Support Cryptocurrency Investors / Fior, Jacopo; Cagliero, Luca; Garza, Paolo. - In: FUTURE INTERNET. - ISSN 1999-5903. - 14:9(2022). [10.3390/fi14090251]

Availability:

This version is available at: 11583/2971267 since: 2022-09-13T15:32:12Z

Publisher:

MDPI

Published

DOI:10.3390/fi14090251

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)



Article

Leveraging Explainable AI to Support Cryptocurrency Investors

Jacopo Fior *, Luca Cagliero and Paolo Garza *

Dipartimento di Automatica e Informatica, Politecnico di Torino, Corso Duca Degli Abruzzi, 24, 10129 Torino, Italy

* Correspondence: jacopo.fior@polito.it (J.F.); paolo.garza@polito.it (P.G.); Tel.: +39-011-090-7084 (J.F.)

Abstract: In the last decade, cryptocurrency trading has attracted the attention of private and professional traders and investors. To forecast the financial markets, algorithmic trading systems based on Artificial Intelligence (AI) models are becoming more and more established. However, they suffer from the lack of transparency, thus hindering domain experts from directly monitoring the fundamentals behind market movements. This is particularly critical for cryptocurrency investors, because the study of the main factors influencing cryptocurrency prices, including the characteristics of the blockchain infrastructure, is crucial for driving experts' decisions. This paper proposes a new visual analytics tool to support domain experts in the explanation of AI-based cryptocurrency trading systems. To describe the rationale behind AI models, it exploits an established method, namely SHapley Additive exPlanations, which allows experts to identify the most discriminating features and provides them with an interactive and easy-to-use graphical interface. The simulations carried out on 21 cryptocurrencies over a 8-year period demonstrate the usability of the proposed tool.

Keywords: quantitative trading; cryptocurrencies; blockchain



Citation: Fior, J.; Cagliero, L.; Garza, P. Leveraging Explainable AI to Support Cryptocurrency Investors. *Future Internet* **2022**, *14*, 251. <https://doi.org/10.3390/fi14090251>

Academic Editors: Leonilde Varela and Goran D. Putnik

Received: 26 June 2022

Accepted: 20 August 2022

Published: 24 August 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 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

Cryptocurrencies are digital assets whose transfers and accounting are cryptographically established through the blockchain [1]. Even though they are not backed by any physical asset, they have become popular financial assets for online trading. After bitcoin, which is the first and most famous cryptocurrency [2], many different cryptocurrencies have been created, thus increasing the options to invest in cryptoassets significantly.

Following the recent trends on algorithmic trading, many research efforts have been devoted to designing cryptocurrency trading systems based on Machine Learning (ML) and Artificial Intelligence (AI). Existing methods rely on algorithms that span from classical classification and regression methods (e.g., [3–7]) to Deep Learning architectures (e.g., [8–11]). The aim is to learn predictive models from historical data related to cryptocurrency assets (e.g., markets, blockchain-related data, news) and apply them to forecast the future price directions. A recent survey on cryptocurrency trading can be found in [12].

Although Machine Learning-based solutions have shown to achieve better performance than simpler heuristic methods [12], they suffer from the lack of transparency. In fact, most state-of-the-art classification models, including all the Neural Network-based approaches, are not explainable; i.e., domain experts cannot gain insights into the model decisions. Cryptocurrency markets can be influenced by a large variety of factors, including the underlying market trends, the characteristics of the blockchain, and the sentiment of financial investors on the virtual assets. This prompts the need for new approaches aimed at explaining ML reasoning in cryptocurrency trading.

This work focuses on leveraging an established explainable AI (XAI) method, namely SHapley Additive exPlanations (SHAP) [13], to provide domain experts with an effective visualization of the ML reasoning behind cryptocurrency trading. SHAP quantifies the contribution of different features on classifier predictions, thus highlighting the contribution

of different factors to the decisions of the ML-based system. We aim at addressing the following research questions:

- **Q1:** What are the most discriminative features for cryptocurrency price prediction?
- **Q2:** How can cryptocurrency investors be provided with quantitative estimates of the influence of specific features and feature categories on machine learning-based cryptocurrency predictions?
- **Q3:** How can we evaluate the statistical dependency of the Machine Learning (ML) feature ranks returned by SHAP in different time periods and on different cryptocurrencies?

To address Q1, this paper explores a large variety of features computed on the daily price series of 21 different cryptocurrencies. The analyzed features are established for cryptocurrency trading [12] and encompass the price- and volume-related series, the technical indicators summarizing the momentum, volatility, and moving averages of the original price series, and the blockchain-related features. The latter are particularly relevant to the scope of the present study because they are peculiar to cryptoassets.

To address Q2, we use eXplainable AI methods based on the Shapley value [14] to provide cryptocurrency traders with evidence of the main factors influencing algorithmic trading. We present a new eXplainable AI tool for visualizing and monitoring the activities of Machine Learning-based systems, with particular attention paid to the blockchain-based features influencing the decision process.

To tackle Q3, we apply the Rank Biased Overlap similarity measure [15] to quantify the pairwise agreement between the top-10 features shortlisted by SHAP. We also performed the experiment using feature subcategories and categories rather than individual features.

The experiments carried out on a 8-year period produce the following main outcomes:

- **O1:** The high variability of the feature importance across different cryptocurrencies. This confirms the relevance of eXplainable AI solutions for cryptocurrency traders.
- **O2:** A visual eXplainable AI tool, namely Cryptocurrency-based Machine Learning Explainer (CryptoMLE, in short). Some practical examples of use of CryptoMLE are also presented.
- **O3:** The dependency among the feature ranks is weak, whereas those among feature subcategories and categories are stronger.

The paper is organized as follows: Section 2 overviews the related literature. Section 3 details the dataset employed in the study. Sections 4 and 5 introduce the fundamentals of Shapley values and SHAP and presents the Visual Analytics tool, based on SHAP, to support cryptocurrency investors' activities. Section 6 summarizes the main empirical results, whereas Sections 7 and 8, respectively, report a discussion of the main achievements and open issues and draw the conclusions of the present work.

2. Comparison with Prior Works

Table 1 summarizes the main characteristics of the existing approaches to eXplainable AI (XAI) in finance, including the Cryptocurrency-based Machine Learning Explainer (CryptoMLE) presented in this paper. We analyze the current and prior works under the following aspects:

1. The considered assets, which encompass specific stocks, cryptocurrencies, or a combination of the above (e.g., the stocks belonging to the Standard&Poor500 U.S. index).
2. The features under analysis, which describe the environmental and market characteristics considered by the classification models (including the blockchain-related features for cryptocurrency assets).
3. The availability of a user interface to support domain expert decisions.
4. The main model used to explain ML-based decisions (e.g., SHAP [13] for the proposed CryptoMLE tool).
5. The resolution of the analyzed data (typically, one sample per trading day).
6. The goal of the approach (e.g., support decision making with data-driven insights for CryptoMLE).

The main goal of this work is to present a visual analytics tool providing AI-based explanations for cryptocurrency investors. Notice that our goal is not to propose a new, more effective trading system but rather to provide experts with an interactive tool, based on XAI, to explain the decisions of algorithmic trading approaches and make appropriate decisions.

Similar to [16–18], CryptoMLE provides domain experts with a graphical interface. Unlike all the prior works on algorithmic trading, it also allows them to interactively collect, analyze, and compare data models trained in multiple time periods. Analogously to [18], CryptoMLE analyzes a large number of cryptocurrencies. Unlike [18], it also considers blockchain-related data.

CryptoMLE relies on SHAP [13], whereas other approaches (e.g., [17,19]) adopt simpler explainable models such as partitional clustering and decision tree, which are known to be less robust to noise and model bias than SHAP. The work recently proposed in [20] is, to the best of our knowledge, the first attempt to use SHAP in algorithmic trading. Unlike [20], this work (1) Addresses short-term cryptocurrency trading instead of long-term portfolio management. Hence, it compares the outcomes of classification models predicting next-day cryptocurrency price; (2) Presents a graphical tool for supporting decision making. It also allows experts to interact with the tool and gain insights into specific market trends; (3) Analyzes a significantly larger set of cryptocurrencies (21 vs. 8).

Table 1. Comparison with prior works. Legend: crypto = cryptocurrency/cryptocurrencies, BC = blockchain, MA = market data, V = Exchanged volumes, TA = technical analysis, B6 = CME Globex British Pound futures, SPF = S&P E-mini Futures.

| Paper | Asset | Features | User Interface | | XAI Model | XAI Resolution | XAI Goal |
|------------------|----------------|--------------|----------------|-------------|---|-------------------------|---|
| | | | Graphical | Interactive | | | |
| CryptoMLE | 21 crypto | BC, MP, TA | Yes | Yes | SHAP [13] | Daily | Decision making |
| [21] | S&P index | MA | No | No | Ablation, permutation, added noise, integrated gradients [22] | Daily | XAI model comparison |
| [23] | CHES120 China | MA | No | No | Custom LightGBM-based model [24] | 10 s, 30 s, 1 min ticks | Matching testing and real-trading performances |
| [20] | 8 crypto | MA | No | No | SHAP [13] | Daily | Portfolio management approach for crypto |
| [19] | The BTC crypto | BC, MA | No | No | K-means clustering, decision tree classifier [25] | Daily | Valuation method for cryptocurrency markets |
| [16] | The ETH crypto | MA, TA | Yes | No | Adversarial Deep Neural Networks [26] | Daily | Display reversal patterns on candlestick charts [27] |
| [17] | The S&P stocks | MA, TA, News | Yes | No | decision tree classifier | Daily | Identify the most impactful words in business-specific stock market sectors |
| [18] | 18 crypto | MA, Reddit | Yes | No | Ensemble methods, co-occurrence analyses [25] | Daily | Correlation analysis between crypto |
| [28] | B6, SPF | MA, V | No | No | Decision trees [25], SHAP [13] | Daily | Adapt market data to the Machine Learning pipeline. |

3. Data Overview and Categorization

We collect historical data about the 21 most popular cryptocurrencies within the time period from 2011 to 2018 (For the cryptocurrencies whose year of introduction is after 2011, we gathered data from the date they became available.). In the experiments we

sampled cryptocurrency data at a daily granularity. However, the performed analyses can be straightforwardly extended to finer or coarser aggregation levels.

We consider three main feature categories:

- The *Blockchain-related* (BC) features, which describe the underlying characteristics of the distributed ledger technology enabling each cryptocurrency [18].
- The *Market Data* (MD) features, which represent the main cryptocurrency Open–High–Low–Close–Volume (OHLCV) price series as well as a selection of summarized features derived from the candlestick chart [29].
- The *Technical Analysis* (TA) features, which include a variety of momentum indicators, volatility indices, and oscillators that are commonly used in Technical Analysis on both cryptocurrencies and regulated market assets [8].

The features are aggregated into the corresponding category and subcategory according to the hierarchy reported in Table 2. We considered a large variety of features among the most established for cryptocurrency trading (according to [12]). To foster the reproducibility of our work, both the analyzed dataset and the project code are publicly available for research purposes (https://dbdmg.polito.it/dbdmg_web/index.php/leveraging-explainable-ai-to-support-cryptocurrency-investors/, accessed on 1 August 2022). A detailed description of the dataset features is available at https://dbdmg.polito.it/dbdmg_web/wp-content/uploads/2022/08/features.xlsx, accessed on 1 August 2022.

Table 2. Categories and subcategories of the features present in the dataset.

| Category | Subcategory | Description |
|-------------|----------------------|---|
| Blockchain | Addresses | Metrics representing an index of network activity and interest. |
| | Economics | Metrics regarding the ratio of the USD network value divided by the adjusted transfer value (in USD). |
| | Exchange | Metrics representing the currency flow for known centralized exchange addresses for both deposits and withdrawals. |
| | Fees and Revenues | Metrics covering the network's efficiency in terms of transfer costs, representing fees for doing operations on the blockchain such as transactions and smart contract execution. |
| | Market | Metrics covering the economic aspects of cryptocurrency markets such as capitalization, BTC exchange price, ROI and volatility returns. |
| | Mining | Metrics representing protocol-specific parameters. |
| | Network Usage | Metrics covering blockchain activity in the form of mined block and their size. |
| | Supply | Metrics that aim to explain token supply and its distribution among wallets. |
| | Transactions | Metrics addressing transferred value and throughput of the network. |
| Market Data | Prices | Features directly derived from Open, High, Low, Close prices of the current timestamp. |
| | Volume | Features directly derived from the trading volume of the current timestamp. |
| | Volatility | Features directly derived from current volatility of the currency. |
| | History | Features derived from the historical time series of Open, High, Low, Close prices and Volume. |
| | Candlestick Analysis | Features concerning the analysis of the candlesticks shapes. |

Table 2. Cont.

| Category | Subcategory | Description |
|--------------------|-----------------------|---|
| Technical Analysis | Trend Indicators | Trend-following indicators whose values help assess the direction and strength of established trends. |
| | Momentum Indicators | Indicators used to determine the strength or weakness of a stock's price. |
| | Volatility Indicators | Indicators measuring how far the security moves away from its mean price. |
| | Volume Indicators | Indicators representing a security's bull and bear power. |

3.1. Blockchain-Related Features

We gathered BC features containing various specific properties of the enabling blockchain architecture, which are aggregated on a daily basis. The 30 features belonging to the BC category cover different aspects addressed by the following subcategories: *Address*, *Economics*, *Exchange*, *Fees and Revenues*, *Market*, *Mining*, and *Network Usage*. They are likely to show direct or indirect relations with the cryptocurrency bid and ask prices. Hence, they can be deemed relevant by the Machine Learning model to obtain accurate price predictions.

The high variability of the technologies enabling each cryptocurrency makes cross-cryptocurrency analyses of BC features particularly relevant to understand the rationale behind Machine Learning predictions. For example, the in-depth analysis of the blockchain supply and mining features can reveal an increasing/decreasing interest of the cryptocurrency investors in particular virtual assets.

3.2. Market Data Features

MD features characterize temporal trends in cryptocurrency prices [11]. The data we gathered include the raw Open–High–Low–Close–Volume (OHLCV) price series, the residuals from the Seasonal-Trend Decomposition using Loess (STL) [30], and the characteristics of the shapes of the candles in the candlestick chart [29].

3.3. Technical Analysis Features

Technical analysis provides a synthetic description of price- and volume-related trends [27]. They were derived from the historical price and volume series using the TA-Lib Python library (<https://ta-lib.org/>, accessed on 10 January 2022).

The TA feature category describes notable price-related properties of the cryptocurrency such as momentum, volatility, oversold/overbought conditions, etc. Recently, they have shown to be relevant to cryptocurrency trading as well [8].

4. SHapley Additive Explanation Values

SHapley Additive Explanation (SHAP, in short) [13] is a method to explain individual predictions. It is based on the Shapley value, whose applications to eXplainable AI rely on coalitional game theory [14].

4.1. The Shapley Value

Given a set of players $\mathcal{P} = \{P_1, P_2, \dots, P_n\}$, a *player coalition* \mathcal{C} is a \mathcal{P} 's subset cooperating to accomplish a specific task. The *utility* $\mathcal{U}(\mathcal{P})$ evaluates the payoff of the coalition for the task, whereas the *marginal utility* $\mathcal{U}(P_j)$ indicates the additional contribution provided by a new player P_j being added to the coalition \mathcal{P} , i.e.,

$$\mathcal{U}(P_j) = \mathcal{U}(\mathcal{P} \cup P_j) - \mathcal{U}(\mathcal{P})$$

The Shapley value [14] is the expectation of the marginal contribution $\mathcal{U}(P_j)$ in all possible coalitions.

$$SV_i = \frac{1}{n} \sum_{S \subseteq \mathcal{N} \setminus P_i} \frac{\mathcal{U}(P \cup P_i) - \mathcal{U}(P)}{\binom{n-1}{|S|}}$$

Computing the exact Shapley value entails enumerating all the possible coalitions, which is computationally prohibitive in real-world contexts.

4.2. Additive Feature Attribution Methods

Given a training dataset consisting of a set of features $\mathcal{F} = \{F_1, F_2, \dots, F_n\}$, each value of an individual feature F_i acts as a player in a coalition. The number n of considered features can be interpreted as maximum coalition size.

Let f be a complex prediction model, trained on \mathcal{F} instances. For the sake of simplicity, we assume the financial forecasting model f predicts the next-day closing price direction (i.e., *Uptrend* or *Downtrend*) of a specific cryptocurrency based on the past samples observed in the last W days (Hereafter, we will disregard the *Stationary* class (neither uptrend nor downtrend)).

We seek explanations of f clarifying the effects of features in \mathcal{F} . Specifically, we aim at explaining the prediction $f(x)$ of an instance x of \mathcal{F} by computing the contribution of each individual feature.

Within this scope, the Shapley value of feature F_i indicates how to fairly distribute the *payout* among the features; i.e., it quantifies the effect of the individual feature F_i on the outcome of the prediction task. To generalize players as sets of feature values, we exploit the additive feature attribution method to linearly combine the individual Shapley values.

The explanation model g is defined as a linear combination of binary features associated with each feature F_i :

$$g(z') = \phi_0 + \sum_{i=1}^n \phi_i \cdot z'_i, \quad z' \in 0, 1^n$$

where z'_i is a binary variable denoting either the presence of a feature ($z'_i = 1$) or its absence ($z'_i = 0$). ϕ_i is the F_i 's attribution value, which quantifies the effect of F_i on $f(x)$. The explanation model sums the effect of all individual feature attributions approximating the output.

4.3. The SHAP Explanation Model

In [13], Shapley values are leveraged to explain Machine Learning models by applying sampling approximations to the original Shapley expression. Specifically, it approximates the effect of removing a variable from the model by integrating over samples from the training dataset.

The key steps of the SHAP model generation are as follows:

1. Generate random sample coalitions z'' of $m < n$ features in \mathcal{F} , where $z'' \in 0, 1^m$.
2. Sample coalitions to valid instances.
3. Train a regression model on the generated instances, whose target is the prediction for a coalition.

To move from coalitions of feature values to valid data instances (Step 2), instance values are taken from the instance x we want to explain for all features that are present in the coalition ($z'' = 1$), whereas the other features are randomly sampled from the training dataset instances for all the absent features ($z''_i = 0$).

The regression function (Step 3) corresponds to the weighted linear explanation model g previously defined according to the additive feature attribution method.

5. The CryptoMLE Tool

Receiving advice from algorithmic advisors is becoming more and more popular for financial analysts [31]. However, relying on sophisticated Machine Learning models trained on massive datasets is particularly risky in financial market forecasting, because the ML models often act as black boxes and domain experts are not keen to trust.

EXplainable AI models provide insights into ML algorithms by indicating which features are more important and how they could affect ML predictions [32]. They can return either local or global explanations. In the former case, the insight is about a particular instance x . The local model estimates the effect of the features in \mathcal{F} on $f(x)$ [13]. Conversely, global models summarize the main patterns driving ML decisions (on whatever instance). In this work, we conveniently combine the local explanations of the cryptocurrency price predictions provided by SHAP to model the global influence on ML models of individual features, features subcategories and categories.

We present a visual eXplainable AI tool, namely Cryptocurrency-based Machine Learning Explainer (CryptoMLE, in short). It supports cryptocurrency traders and investors in monitoring the performance of quantitative Machine Learning-based cryptocurrency predictions. CryptoMLE consists of an interactive dashboard summarizing the main feature contributions to the ML price predictions.

A snapshot of the dashboard interface is depicted in Figure 1. The plot in the upper side of the dashboard shows the SHAP time series of the 10 most influential features in the prediction of class *uptrend*. The purpose is to explain how ML works within a restricted time period and how ML decisions vary over time. More specifically, a time series value sampled on day d consists of the mean Shapley value of a given feature F_i computed over the W days preceding d (Since historical data are collected at a daily granularity, each time point in the series corresponds to a distinct trading day.). The mean Shapley value of F_i indicates the effect of F_i on the ML model trained on d using a sliding window approach.

For example, according to the SHAP series plot in Figure 1, the MD feature *close_resid* appears to be the most influential one in the period between August 2017 and April 2018, whereas between May 2018 and December 2018, *Close_resid* and *High_resid* are joint winners. The SHAP series plot can be useful, for instance, for discretionary traders who need to select and monitor a relatively small subset of visual features.

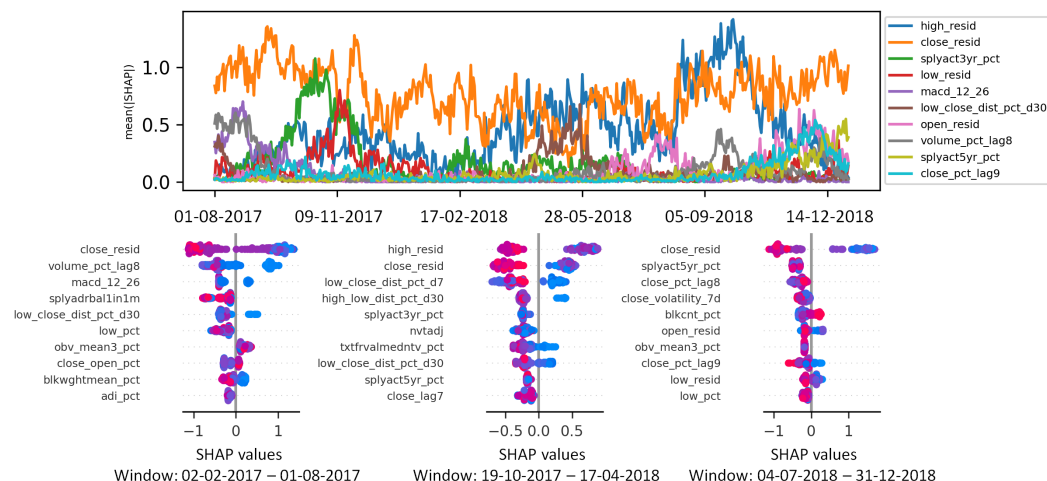


Figure 1. Interactive dashboard snapshot. Uptrend class. LTCUSD. Training window size $W = 90$.

The bee-swarm summary plots in the lower side of the dashboard snapshot are pop-up windows that analysts might activate when they are interested in gaining insights into the characteristics of the ML model trained on a particular day. It shows the Shapley values of all instances belonging to a training window of size W (i.e., W points per feature). For the sake of readability, only the top-10 features in order of decreasing Shapley value are visualized.

For example, according to the left hand-side bee-swarm summary plot in Figure 1, *close_resid*, *volume_pct_tag8* *macd_12_26*, and *low_close_dist_pct_d30* are the only features obtaining a significant number of positive Shapley values during the training window from the beginning of February 2017 to the end of July 2017. Comparing different summary plots over time can be useful, for instance, for detecting temporal changes in the ML decisions. Traders can manually verify and possibly revise the current trading strategy based on the alarms triggered by the eXplainable AI tool.

To generate the plot, we apply the procedure described in Algorithm 1 considering one cryptocurrency at a time. First, the dataset D_c , containing the data of the cryptocurrency c , is split into train and test, and the feature importance scores are computed based on a general-purpose Machine Learning model trained on D_{train} (e.g., XGBoost [33]). Then, we generate a ranked feature list, based on the importance score, and tune the system hyper-parameters. This first phase aims at performing feature selection and parameter tuning before training the following models. To have up-to-date and contextualized models, one model is retrained for each test date/time-step t considering the latest W days preceding t , using the previously defined feature subset and hyperparameters; i.e., we employ a sliding window approach to train ML models tailored to the time-steps t . Finally, the trained ML models (one per test time-step) are analyzed to compute the SHAP series and the summary plots, thus enabling the visual exploration of the ML reasoning at different time points. The procedure is repeated for all cryptocurrencies of interest.

Algorithm 1: CryptoMLE: Procedure of dashboard generation for a cryptocurrency.

```

Input :  $F$ : feature set;
          $D_c$ : dataset associated with cryptocurrency  $c$ ;
          $W$ : sliding training window;
          $f_{pr}$ : chosen Machine Learning model for the prediction step;
          $f_{fs}$ : chosen Machine Learning model for the feature selection step;

output:  $SH$ : time series of average Shapely values per-feature;
          $BS$ : bee-swarm summary plots for each point of the test-set;

/* Train-test dataset split */
 $D_{train}, D_{test} \leftarrow \text{SplitDataset}(D_c)$  */
/* Feature selection */
 $M_{fs} \leftarrow \text{TrainFeatureSelectionModel}(f_{fs}, D_{train}, F)$ 
 $R \leftarrow \text{FeatureImportanceRankingForModel}(M_{fs})$ 
 $F_s \leftarrow \text{SelectFeaturesFromRanking}(R, F)$ 
/* Hyper-parameters tuning */
 $P \leftarrow \text{TuneHyperparameters}(f_{pr}, D_{train}, F_s)$ 
/* Dashboard generation */
foreach time-step  $t \in D_{test}$  do
    |  $M_{pr} \leftarrow \text{TrainPredictionModel}(f_{pr}, D_{(t-W,t)}, F_s, P)$ 
    |  $BS_t \leftarrow \text{ProduceBeeswarmPlot}(M_{pr})$ 
end
 $SH \leftarrow \text{ProduceShapTimeSeries}(BS_t)$ 
return  $SH, BS_*$ 

```

6. Experimental Results

In this section, we simulate a session of Machine Learning-based forecasting of 21 cryptocurrency prices explained by CryptoMLE.

The rest of the section is organized as follows.

- Section 6.1 clarifies the experimental settings and the reproducibility aspects.
- Section 6.2 reports the main findings related to Research Question 1, i.e., *What are the most discriminative features for cryptocurrency price prediction?* Empirical outcome O1 compares the feature importance plots relative to different cryptocurrencies.

- Section 6.3 addresses the Research Question 2, i.e., *How to provide cryptocurrency investors with quantitative estimates of the influence of specific features and feature categories on Machine Learning-based cryptocurrency predictions?* The empirical outcome O2 consists of a selection of SHAP series and bee-swarm summary plots highlighting interesting trends in the analyzed cryptocurrencies.
- Section 6.4 addresses the Research Question 3, i.e., *How can we evaluate the statistical dependence of the ML feature ranks returned by SHAP in different time periods and on different cryptocurrencies?* We address O3 by evaluating the pairwise agreement between the shortlisted feature ranks using the Rank Biased Overlap similarity measure [15].

6.1. Experimental Design

In the following, we describe the hardware used to perform the experiments and the experimental settings to improve reproducibility.

Hardware settings. We run experiments in a single-node setting on an HPC facility. The node runs Ubuntu 20.04.2 LTS, with an 8 CPU threads Intel(R) Xeon(R) Gold 6140 CPU @ 2.30 GHz and 40 GB of RAM.

Experimental settings and reproducibility. The source data described in Section 3 and a detailed per-feature description are available for research purposes. We also release the guidelines for dashboard creation (again for research purposes only).

As a representative ML model for both classification and feature importance estimation, we used the XGBoost classifier available in the SK-Learn library [33]. It is both efficient and accurate. To run SHAP [13], we use the publicly available code released by the paper’s authors.

6.2. Empirical Outcome O1: Feature Importance across Cryptocurrencies

The pie charts in Figures 2–5 show the feature importance scores (returned by the XGBoost ML model) computed over all cryptocurrencies (see Figure 2) and separately for BTCUSD, BCHUSD, and ETHUSD (see Figures 3–5). BTCUSD is, by far, the most famous cryptocurrency. BCHUSD is a fork of BTCUSD, whereas ETHUSD is another extremely popular cryptocurrency.

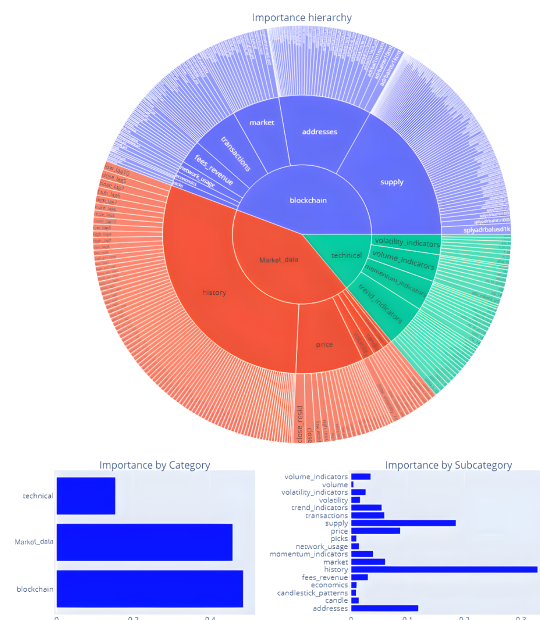


Figure 2. Hierarchical mean feature importance over all the analyzed cryptocurrencies.

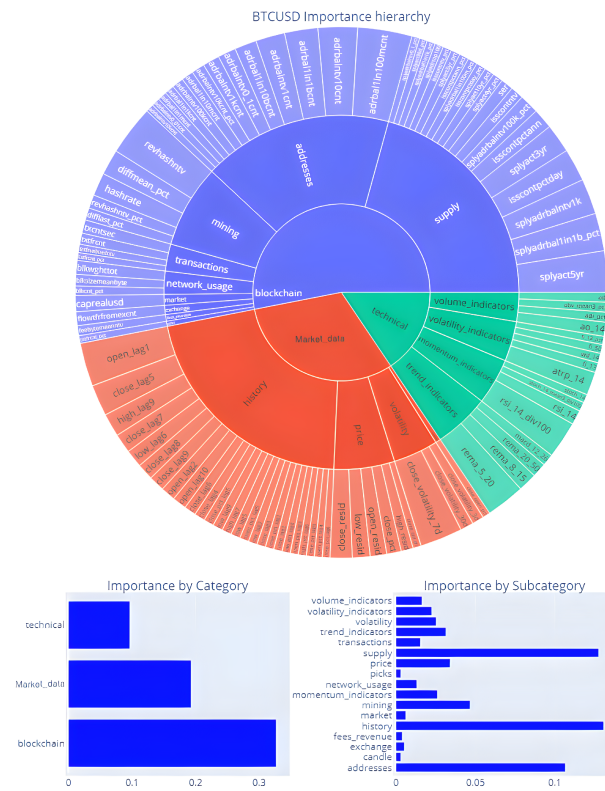


Figure 3. Hierarchical mean feature importance for BTCUSD.

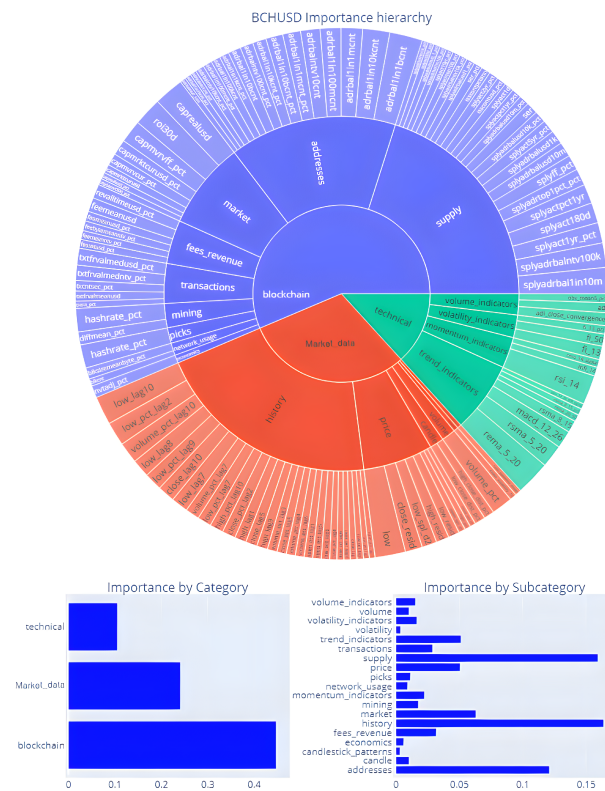


Figure 4. Hierarchical mean feature importance for BCHUSD.




Figure 5. Hierarchical mean feature importance for ETHUSD.

The outermost circular crown of the pie chart in Figure 2 reports the average importance scores per feature by considering all cryptocurrencies. Specific price-related features, such as *close_resid* (i.e., the Seasonal-Trend decomposition using LOESS of the closing price series [30], have shown to be the most relevant to predict future cryptocurrency prices). However, the selected features are not the same for all cryptocurrencies and also include blockchain-related ones. For example, *hashrate_pct*, which indicates the amount of computational operations that a miner or the network of miners is capable of carrying out, is particularly relevant to BitCoin casH (BCH), which has been created to specifically address efficiency issues of the most established BTC cryptocurrency. Conversely, it is not relevant to Ethereum (ETH) because ETH is known to be weakly correlated to BTC.

To have a higher-level view of which features are more discriminating for a given cryptocurrency, we also aggregate the feature importance scores per subcategory and category (see the two inner crowns in Figure 2 and the bar charts). The most relevant features are those belonging to category *Blockchain* (average score 0.48), which is followed by *Market data* (0.46) and *Technical analysis* features (0.16). This means that to drive their investments, cryptocurrency traders should closely monitor blockchain-related features first rather than simply analyzing price-related features (e.g., moving averages, momentum [27]).

Focusing on the most influential subcategories, they encompass the properties of the supply chain, namely *Supply* (BC category), the historical cryptocurrency prices, i.e., *History* (MD category), and the blockchain network activity metrics, namely *Addresses* (BC category). It is worth noticing that restricting the in-depth analysis to these feature subsets allows experts to ignore almost 70% of the original features.

The variability in feature importance across different cryptocurrencies is also quite significant (see Figures 3–5). For example, for ETHUSD, the blockchain-related features turn out to be slightly less significant than for BTCUSD and BCHUSD, which is possibly due to the primary influence of the blockchain architecture on the price movements of the BitCoin-related assets. Ethereum (ETH) is partly uncorrelated with BC and weakly dependent on blockchain-related properties such as hash rate and transaction counts.

6.3. Visual Explanations

We report the dashboard snapshots for three representative combinations of cryptocurrency and prediction class (see Figures 6–11).

From the line chart in the upper side of Figure 6, we can see how the average Shapley value changes over time for the top-10 most influential features for class Uptrend. Some features (e.g., *close_resid*) are always highly relevant regardless of the considered time period, whereas some others show variable influence. The latter can be dynamically included in the trading system models according to the feedback collected from the eXplainable AI tool. Moreover, traders can also use the provided information to assess the reliability of the performed predictions. If the features associated with the highest absolute Shapley values are, based on the prior knowledge of traders, remarkable features, traders will become more confident in the returned predictions and thus will likely use them in the design of the cryptocurrency trading strategy. In a nutshell, the visual explanation of CryptoMLE has a twofold aim: (1) understand the rationale behind ML decisions, and (2) discover potentially interesting (cryptocurrency-specific) patterns that are worth considering in the future trading activities.

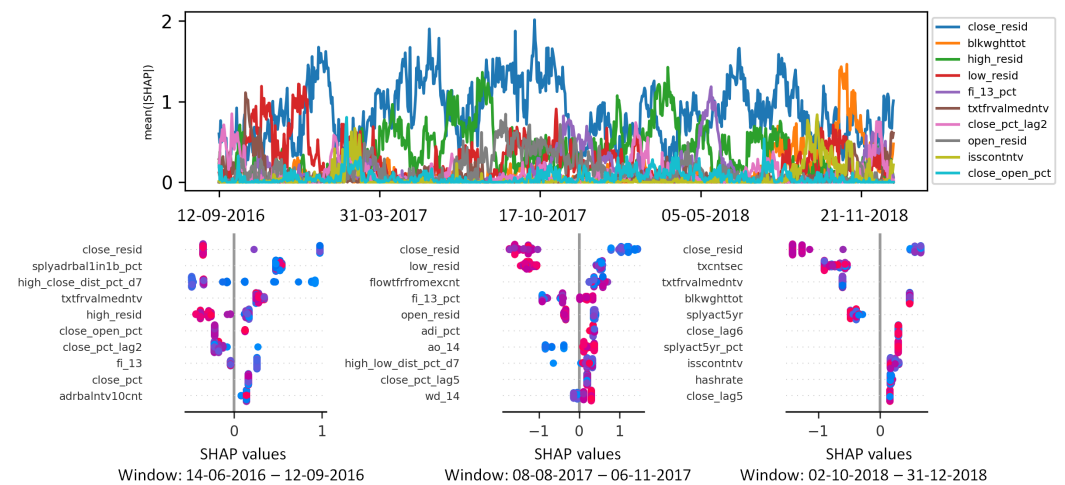


Figure 6. Interactive dashboard snapshot. Uptrend class. BTCUSD. Sliding training window size $W = 90$.

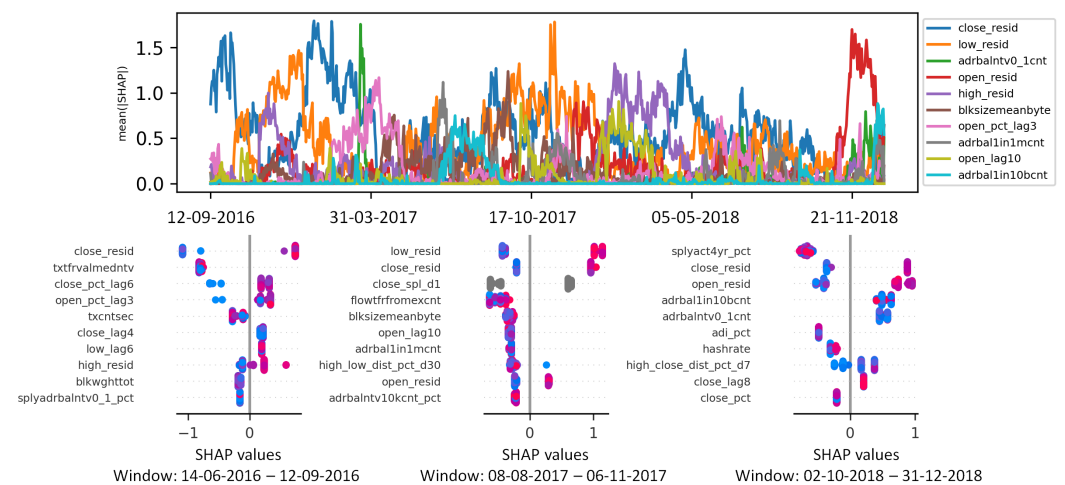


Figure 7. Interactive dashboard snapshot. Downtrend class. BTCUSD. Sliding training window size $W = 90$.

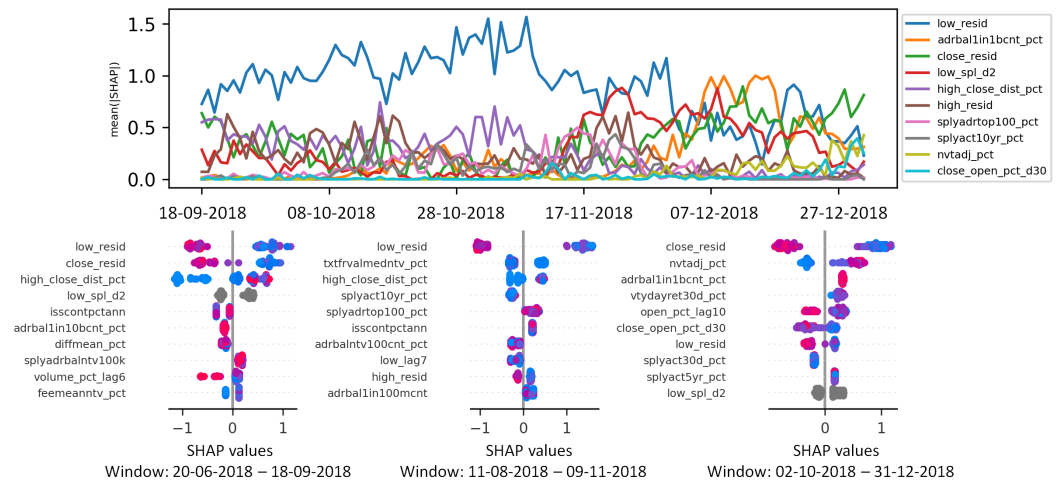


Figure 8. Interactive dashboard snapshot. Uptrend class. BCHUSD. Sliding training window size $W = 90$.

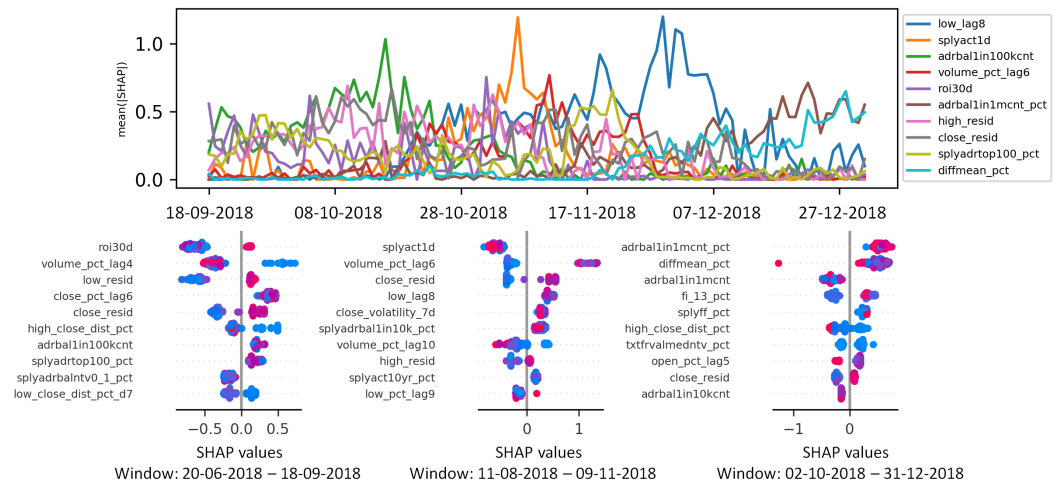


Figure 9. Interactive dashboard snapshot. Downtrend class. BCHUSD. Sliding training window size $W = 90$.

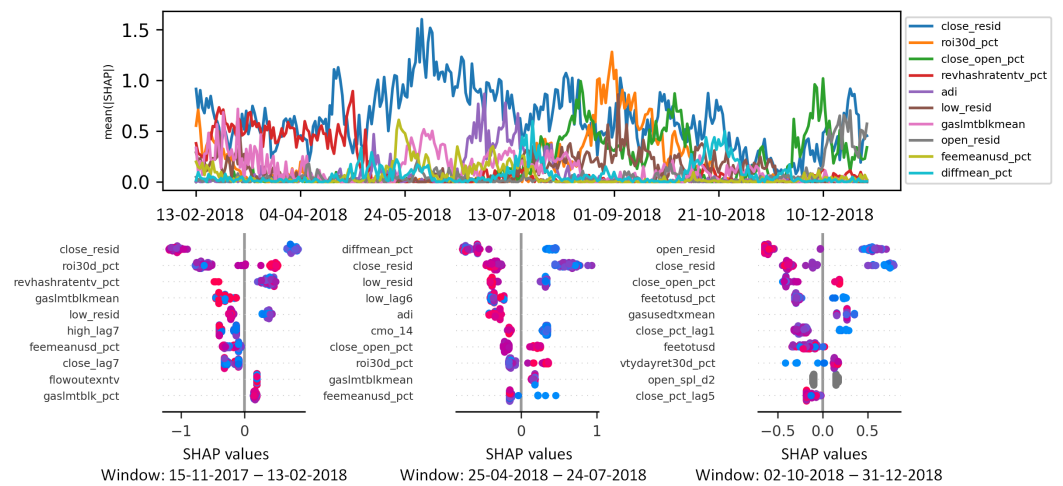


Figure 10. Interactive dashboard snapshot. Uptrend class. ETHUSD. Sliding training window size $W = 90$.

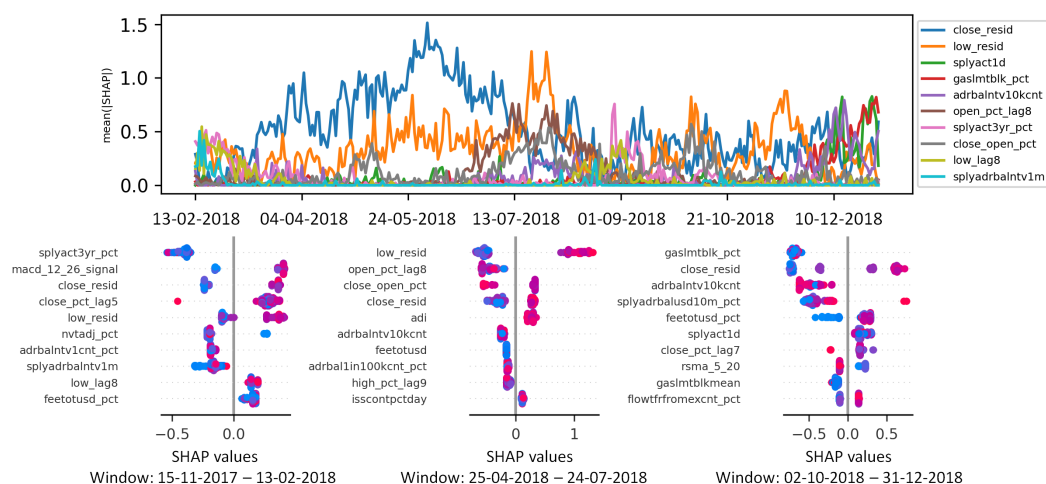


Figure 11. Interactive dashboard snapshot. Downtrend class. ETHUSD. Sliding training window size $W = 90$.

The charts in the bottom part of Figure 6 report the Shapley values computed for three representative windows of size W (i.e., the ones associated with the first, the middle, and last date of the considered evaluation period). Each chart reports, for the top-10 most influential features, the Shapley values associated with the predictions made within the considered time window (i.e., one point in the plot per prediction). The absolute Shapley value indicates the strength of the feature influences. Its sign indicates whether the value of the feature positively or negatively impacts on the prediction of the Uptrend label. If the Shapley values associated with a feature are close to -1 or 1 , it means that the feature affects the prediction more significantly than the other ones. As for the SHAP series plot, the variability in the summary plots over time strongly depends on the underlying market conditions. For example, in the last quarter of 2018, AI model predictions turn out to be primarily influenced by the historical price series, whereas in the previous quarters of 2018, the influence of blockchain-related features is more evident. Based on these results, domain experts can investigate more in depth the reasons behind such a strategy change to judge the reliability of the algorithmic trading approach. More specifically, in the last quarter of 2018, all the BitCoin-related plunged, and such an evident market downtrend is prevailing for algorithmic trading systems.

Figure 7 reports a similar information but the class label Downtrend is considered, i.e., those charts try to explain which features impacted more on the prediction of the label Downtrend. Some features are relevant for the prediction of both class labels, whereas others are specific for each class.

Figures 8–9 and Figures 10–11 report similar pieces of information for BCHUSD and ETHUSD, respectively. We can notice that some of the top features are shared between BTCUSD and BCHUSD, whereas ETHUSD is more affected by other blockchain-related features. Most of the top features categories are shared among all the three considered features.

Tables 3 and 4 report the top-3 most influential features per cryptocurrency and class in terms of average Shapley value. The achieved results confirm that for most of the analyzed cryptocurrencies, the subcategories of the most influential features are independent of the predicted class label.

Table 3. Most influential features for class Uptrend.

| Crypto | Top1_Feature | Top1_Subcategory | Top2_Feature | Top2_Subcategory | Top3_Feature | Top3_Subcategory |
|--------|--------------------|----------------------------------|----------------------|----------------------------------|--------------------|----------------------------------|
| ADA | close_resid | market_data_prices | splyact180d_pct | blockchain_supply | adi_pct | technical_analysis_volume |
| BCH | close_resid | market_data_prices | nvtadj_pct | blockchain_economics | adrbal1in1bent_pct | blockchain_address |
| BNB | close_open_pct_d30 | market_data_candlestick_analysis | splyadrbalntv100k | blockchain_supply | capact1yrusd | blockchain_market |
| BTC | close_resid | market_data_prices | txcntsec | blockchain_transactions | txftrvalmedntv | blockchain_transactions |
| BTG | close_resid | market_data_prices | high_low_dist_pct_d7 | market_data_candlestick_analysis | low_pct_lag4 | market_data_history |
| DASH | close_resid | market_data_prices | adrbal1in1mcent_pct | blockchain_address | low_pct_lag3 | market_data_history |
| DOGE | close_open_pct_d30 | market_data_candlestick_analysis | close_resid | market_data_prices | txftrvalmedntv | blockchain_transactions |
| EOS | close_resid | market_data_prices | open_pct_lag6 | market_data_history | low_pct_lag3 | market_data_history |
| ETC | adrbalntv0_01cnt | blockchain_address | gasmlblk | blockchain_fees | gasmlblk | blockchain_fees |
| ETH | open_resid | market_data_prices | close_resid | market_data_prices | close_open_pct | market_data_candlestick_analysis |
| LINK | txftrcnt | blockchain_transactions | splyadrtop1pct_pct | blockchain_supply | caprealusd | blockchain_market |
| LTC | close_resid | market_data_prices | high_resid | market_data_prices | txftrcnt | blockchain_transactions |
| NEO | adi_pct | technical_analysis_volume | close_resid | market_data_prices | txftrcnt | blockchain_transactions |
| QTUM | low_resid | market_data_prices | volume_pct_lag9 | market_data_history | open_pct_lag5 | market_data_history |
| TRX | high_resid | market_data_prices | close_resid | market_data_prices | high_pct_lag8 | market_data_history |
| WAVE | txcntsec | blockchain_transactions | low_resid | market_data_prices | adi | technical_analysis_volume |
| XEM | cmo_14 | technical_analysis_momentum | high_pct_lag2 | market_data_history | close_resid | market_data_prices |
| XMR | close_resid | market_data_prices | high_lag2 | market_data_history | rema_8_15_pct | technical_analysis_trend |
| XRP | close_resid | market_data_prices | close_open_pct_d3 | market_data_candlestick_analysis | close_pct_lag7 | market_data_history |
| ZEC | high_resid | market_data_prices | close_resid | market_data_prices | close_pct_lag3 | market_data_history |
| ZRX | txftrvalmeanusd | blockchain_transactions | high_pct_lag2 | market_data_history | txftrvaladintv_pct | blockchain_transactions |

Table 4. Most influential features for class Downtrend.

| Crypto | Top1_Feature | Top1_Subcategory | Top2_Feature | Top2_Subcategory | Top3_Feature | Top3_Subcategory |
|--------|------------------------|----------------------------------|------------------------|-------------------------|------------------------|----------------------------------|
| ADA | splyadrtop100_pct | blockchain_supply | splyact1yr_pct | blockchain_supply | close_resid | market_data_prices |
| BCH | adrbal1in1mcent_pct | blockchain_address | diffmean_pct | blockchain_mining | adrbal1in1mcent | blockchain_address |
| BNB | splyadrbalntylk_pct | blockchain_supply | splyadrbal1in1k_pct | blockchain_supply | low_close_dist_pct_d30 | market_data_candlestick_analysis |
| BTC | splyact4yr_pct | blockchain_supply | close_resid | market_data_prices | open_resid | market_data_prices |
| BTG | close_resid | market_data_prices | low_resid | market_data_prices | txftrvaladjud | blockchain_transactions |
| DASH | close_resid | market_data_prices | isstotl_isstotl365_pct | blockchain_supply | high_resid | market_data_prices |
| DOGE | volume_pct_lag3 | market_data_history | high_resid | market_data_prices | txftrvalmedntv | blockchain_transactions |
| EOS | close_pct_lag8 | market_data_history | close_resid | market_data_prices | open_pct_lag4 | market_data_history |
| ETC | close_resid | market_data_prices | splyactever_pct | blockchain_supply | nvtadj | blockchain_economics |
| ETH | gasmlblk_pct | blockchain_fees | close_resid | market_data_prices | adrbalntv10kcnt | blockchain_address |
| LINK | high_close_dist_pct_d3 | market_data_candlestick_analysis | splyadrbalud1m | blockchain_supply | close_resid | market_data_prices |
| LTC | close_resid | market_data_prices | volume_pct_lag1 | market_data_history | close_open_pct_d30 | market_data_candlestick_analysis |
| NEO | close_resid | market_data_prices | low_pct_lag9 | market_data_history | low_close_dist_pct | market_data_candlestick_analysis |
| QTUM | close_pct_lag10 | market_data_history | low_resid | market_data_prices | open_lag9 | market_data_history |
| TRX | fi_13_pct | technical_analysis_volatility | close_resid | market_data_prices | high_resid | market_data_prices |
| WAVE | close_pct_lag7 | market_data_history | close_resid | market_data_prices | volume_pct_lag3 | market_data_history |
| XEM | low_resid | market_data_prices | volume_pct_lag2 | market_data_history | low_lag4 | market_data_history |
| XMR | close_resid | market_data_prices | txcnt_pct | blockchain_transactions | close_volatility_7d | market_data_volatility |
| XRP | close_resid | market_data_prices | adrbalntv1mcent_pct | blockchain_address | volume_pct_lag4 | market_data_volatility |
| ZEC | close_resid | market_data_prices | close_pct_lag8 | market_data_history | low_spl_d1 | market_data_prices |
| ZRX | high_close_dist_pct_d3 | market_data_candlestick_analysis | splyadrbal1in10k_pct | blockchain_supply | low_resid | market_data_prices |

6.4. Statistical Dependence between Feature Ranked Lists

We evaluate the agreement between the feature ranked lists associated with the 21 cryptocurrencies using the Rank Biased Overlap similarity measure [15]. The goal is to verify whether ML predictions on different cryptocurrencies are influenced by the same features, feature subcategories, or categories.

Tables 5 and 6, respectively, report the pairwise similarity matrices for the classes Uptrend and Downtrend. They allow us to identify specific cryptocurrency clusters characterized by relatively high pairwise similarities. For instance, XMR and ZEC are highly similar, which is probably because they are both focused on privacy aspects.

Table 5. Pairwise similarity among cryptocurrencies. Class Uptrend.

| | ADA | BCH | BNB | BTC | BTG | DASH | DOGE | EOS | ETC | ETH | LINK | LTC | NEO | QTUM | TRX | WAVE | XEM | XMR | XRP | ZEC | ZRX |
|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|
| ADA | 1.00 | 0.84 | 0.98 | 0.84 | 0.89 | 0.94 | 0.94 | 0.72 | 0.80 | 0.78 | 0.68 | 0.94 | 0.85 | 0.72 | 0.76 | 0.90 | 0.81 | 0.90 | 0.88 | 0.91 | 0.86 |
| BCH | 0.84 | 1.00 | 0.85 | 1.00 | 0.74 | 0.89 | 0.84 | 0.61 | 0.66 | 0.94 | 0.84 | 0.84 | 0.69 | 0.61 | 0.59 | 0.74 | 0.63 | 0.71 | 0.77 | 0.81 | 0.78 |
| BNB | 0.98 | 0.85 | 1.00 | 0.85 | 0.90 | 0.96 | 0.96 | 0.74 | 0.82 | 0.80 | 0.70 | 0.96 | 0.84 | 0.74 | 0.78 | 0.88 | 0.79 | 0.88 | 0.89 | 0.93 | 0.87 |
| BTC | 0.84 | 1.00 | 0.85 | 1.00 | 0.74 | 0.89 | 0.84 | 0.61 | 0.66 | 0.94 | 0.84 | 0.84 | 0.69 | 0.61 | 0.59 | 0.74 | 0.63 | 0.71 | 0.77 | 0.81 | 0.78 |
| BTG | 0.89 | 0.74 | 0.90 | 0.74 | 1.00 | 0.84 | 0.91 | 0.83 | 0.65 | 0.79 | 0.52 | 0.91 | 0.79 | 0.83 | 0.87 | 0.79 | 0.83 | 0.96 | 0.96 | 0.94 | 0.76 |
| DASH | 0.94 | 0.89 | 0.96 | 0.89 | 0.84 | 1.00 | 0.95 | 0.77 | 0.76 | 0.84 | 0.74 | 0.95 | 0.80 | 0.77 | 0.75 | 0.84 | 0.73 | 0.82 | 0.87 | 0.92 | 0.89 |
| DOGE | 0.94 | 0.84 | 0.96 | 0.84 | 0.91 | 0.95 | 1.00 | 0.78 | 0.73 | 0.84 | 0.65 | 1.00 | 0.84 | 0.78 | 0.82 | 0.84 | 0.78 | 0.89 | 0.94 | 0.97 | 0.86 |
| EOS | 0.72 | 0.61 | 0.74 | 0.61 | 0.83 | 0.77 | 0.78 | 1.00 | 0.45 | 0.67 | 0.40 | 0.78 | 0.62 | 1.00 | 0.96 | 0.62 | 0.66 | 0.81 | 0.84 | 0.81 | 0.66 |
| ETC | 0.80 | 0.66 | 0.82 | 0.66 | 0.65 | 0.76 | 0.73 | 0.45 | 1.00 | 0.54 | 0.81 | 0.73 | 0.74 | 0.45 | 0.51 | 0.90 | 0.66 | 0.64 | 0.62 | 0.68 | 0.87 |
| ETH | 0.78 | 0.94 | 0.80 | 0.94 | 0.79 | 0.84 | 0.84 | 0.67 | 0.54 | 1.00 | 0.73 | 0.84 | 0.68 | 0.67 | 0.65 | 0.68 | 0.66 | 0.75 | 0.82 | 0.87 | 0.72 |
| LINK | 0.68 | 0.84 | 0.70 | 0.84 | 0.52 | 0.74 | 0.65 | 0.40 | 0.81 | 0.73 | 1.00 | 0.65 | 0.62 | 0.40 | 0.38 | 0.78 | 0.54 | 0.51 | 0.55 | 0.60 | 0.82 |
| LTC | 0.94 | 0.84 | 0.96 | 0.84 | 0.91 | 0.95 | 1.00 | 0.78 | 0.73 | 0.84 | 0.65 | 1.00 | 0.84 | 0.78 | 0.82 | 0.84 | 0.78 | 0.89 | 0.94 | 0.97 | 0.86 |
| NEO | 0.85 | 0.69 | 0.84 | 0.69 | 0.79 | 0.80 | 0.84 | 0.62 | 0.74 | 0.68 | 0.62 | 0.84 | 1.00 | 0.62 | 0.66 | 0.85 | 0.96 | 0.80 | 0.78 | 0.81 | 0.81 |
| QTUM | 0.72 | 0.61 | 0.74 | 0.61 | 0.83 | 0.77 | 0.78 | 1.00 | 0.45 | 0.67 | 0.40 | 0.78 | 0.62 | 1.00 | 0.96 | 0.62 | 0.66 | 0.81 | 0.84 | 0.81 | 0.66 |
| TRX | 0.76 | 0.59 | 0.78 | 0.59 | 0.87 | 0.75 | 0.82 | 0.96 | 0.51 | 0.65 | 0.38 | 0.82 | 0.66 | 0.96 | 1.00 | 0.66 | 0.70 | 0.85 | 0.88 | 0.84 | 0.67 |
| WAVE | 0.90 | 0.74 | 0.88 | 0.74 | 0.79 | 0.84 | 0.84 | 0.62 | 0.90 | 0.68 | 0.78 | 0.84 | 0.85 | 0.62 | 0.66 | 1.00 | 0.81 | 0.80 | 0.78 | 0.81 | 0.96 |
| XEM | 0.81 | 0.63 | 0.79 | 0.63 | 0.83 | 0.73 | 0.78 | 0.66 | 0.66 | 0.66 | 0.54 | 0.78 | 0.96 | 0.66 | 0.70 | 0.81 | 1.00 | 0.85 | 0.82 | 0.79 | 0.75 |
| XMR | 0.90 | 0.71 | 0.88 | 0.71 | 0.96 | 0.82 | 0.89 | 0.81 | 0.64 | 0.75 | 0.51 | 0.89 | 0.80 | 0.81 | 0.85 | 0.80 | 0.85 | 1.00 | 0.93 | 0.90 | 0.74 |
| XRP | 0.88 | 0.77 | 0.89 | 0.77 | 0.96 | 0.87 | 0.94 | 0.84 | 0.62 | 0.82 | 0.55 | 0.94 | 0.78 | 0.84 | 0.88 | 0.78 | 0.82 | 0.93 | 1.00 | 0.96 | 0.79 |
| ZEC | 0.91 | 0.81 | 0.93 | 0.81 | 0.94 | 0.92 | 0.97 | 0.81 | 0.68 | 0.87 | 0.60 | 0.97 | 0.81 | 0.81 | 0.84 | 0.81 | 0.79 | 0.90 | 0.96 | 1.00 | 0.83 |
| ZRX | 0.86 | 0.78 | 0.87 | 0.78 | 0.76 | 0.89 | 0.86 | 0.66 | 0.87 | 0.72 | 0.82 | 0.86 | 0.81 | 0.66 | 0.67 | 0.96 | 0.75 | 0.74 | 0.79 | 0.83 | 1.00 |

Table 6. Correlations among cryptocurrencies. Class Downtrend.

| | ADA | BCH | BNB | BTC | BTG | DASH | DOGE | EOS | ETC | ETH | LINK | LTC | NEO | QTUM | TRX | WAVE | XEM | XMR | XRP | ZEC | ZRX |
|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|------|
| ADA | 1.00 | 0.88 | 0.98 | 0.95 | 0.76 | 0.85 | 0.76 | 0.56 | 0.86 | 0.97 | 0.83 | 0.69 | 0.64 | 0.56 | 0.55 | 0.76 | 0.66 | 0.85 | 0.87 | 0.69 | 0.82 |
| BCH | 0.88 | 1.00 | 0.86 | 0.87 | 0.76 | 0.73 | 0.76 | 0.54 | 0.74 | 0.85 | 0.79 | 0.54 | 0.47 | 0.54 | 0.54 | 0.62 | 0.51 | 0.73 | 0.75 | 0.54 | 0.82 |
| BNB | 0.98 | 0.86 | 1.00 | 0.93 | 0.74 | 0.87 | 0.74 | 0.53 | 0.86 | 0.95 | 0.82 | 0.71 | 0.66 | 0.53 | 0.52 | 0.78 | 0.68 | 0.87 | 0.85 | 0.71 | 0.80 |
| BTC | 0.95 | 0.87 | 0.93 | 1.00 | 0.84 | 0.86 | 0.84 | 0.66 | 0.87 | 0.98 | 0.89 | 0.72 | 0.68 | 0.66 | 0.63 | 0.80 | 0.70 | 0.86 | 0.88 | 0.72 | 0.87 |
| BTG | 0.76 | 0.76 | 0.74 | 0.84 | 1.00 | 0.90 | 1.00 | 0.83 | 0.91 | 0.82 | 0.96 | 0.83 | 0.79 | 0.83 | 0.68 | 0.90 | 0.81 | 0.90 | 0.92 | 0.83 | 0.97 |
| DASH | 0.85 | 0.73 | 0.87 | 0.86 | 0.90 | 1.00 | 0.90 | 0.69 | 0.99 | 0.88 | 0.95 | 0.87 | 0.82 | 0.69 | 0.57 | 0.94 | 0.84 | 1.00 | 0.98 | 0.87 | 0.93 |
| DOGE | 0.76 | 0.76 | 0.74 | 0.84 | 1.00 | 0.90 | 1.00 | 0.83 | 0.91 | 0.82 | 0.96 | 0.83 | 0.79 | 0.83 | 0.68 | 0.90 | 0.81 | 0.90 | 0.92 | 0.83 | 0.97 |
| EOS | 0.56 | 0.54 | 0.53 | 0.66 | 0.83 | 0.69 | 0.83 | 1.00 | 0.71 | 0.62 | 0.78 | 0.64 | 0.66 | 1.00 | 0.84 | 0.75 | 0.65 | 0.69 | 0.72 | 0.64 | 0.80 |
| ETC | 0.86 | 0.74 | 0.86 | 0.87 | 0.91 | 0.99 | 0.91 | 0.71 | 1.00 | 0.89 | 0.95 | 0.86 | 0.81 | 0.71 | 0.58 | 0.93 | 0.83 | 0.99 | 0.99 | 0.86 | 0.94 |
| ETH | 0.97 | 0.85 | 0.95 | 0.98 | 0.82 | 0.88 | 0.82 | 0.62 | 0.89 | 1.00 | 0.86 | 0.75 | 0.70 | 0.62 | 0.59 | 0.82 | 0.72 | 0.88 | 0.90 | 0.75 | 0.85 |
| LINK | 0.83 | 0.79 | 0.82 | 0.89 | 0.96 | 0.95 | 0.96 | 0.78 | 0.95 | 0.86 | 1.00 | 0.81 | 0.77 | 0.78 | 0.65 | 0.89 | 0.79 | 0.95 | 0.96 | 0.81 | 0.99 |
| LTC | 0.69 | 0.54 | 0.71 | 0.72 | 0.83 | 0.87 | 0.83 | 0.64 | 0.86 | 0.75 | 0.81 | 1.00 | 0.95 | 0.64 | 0.48 | 0.92 | 0.98 | 0.87 | 0.85 | 1.00 | 0.80 |
| NEO | 0.64 | 0.47 | 0.66 | 0.68 | 0.79 | 0.82 | 0.79 | 0.66 | 0.81 | 0.70 | 0.77 | 0.95 | 1.00 | 0.66 | 0.50 | 0.88 | 0.98 | 0.82 | 0.80 | 0.95 | 0.76 |
| QTUM | 0.56 | 0.54 | 0.53 | 0.66 | 0.83 | 0.69 | 0.83 | 1.00 | 0.71 | 0.62 | 0.78 | 0.64 | 0.66 | 1.00 | 0.84 | 0.75 | 0.65 | 0.69 | 0.72 | 0.64 | 0.80 |
| TRX | 0.55 | 0.54 | 0.52 | 0.63 | 0.68 | 0.57 | 0.68 | 0.84 | 0.58 | 0.59 | 0.65 | 0.48 | 0.50 | 0.84 | 1.00 | 0.59 | 0.49 | 0.57 | 0.59 | 0.48 | 0.67 |
| WAVE | 0.76 | 0.62 | 0.78 | 0.80 | 0.90 | 0.94 | 0.90 | 0.75 | 0.93 | 0.82 | 0.89 | 0.92 | 0.88 | 0.75 | 0.59 | 1.00 | 0.90 | 0.94 | 0.92 | 0.92 | 0.87 |
| XEM | 0.66 | 0.51 | 0.68 | 0.70 | 0.81 | 0.84 | 0.81 | 0.65 | 0.83 | 0.72 | 0.79 | 0.98 | 0.98 | 0.65 | 0.49 | 0.90 | 1.00 | 0.84 | 0.82 | 0.98 | 0.78 |
| XMR | 0.85 | 0.73 | 0.87 | 0.86 | 0.90 | 1.00 | 0.90 | 0.69 | 0.99 | 0.88 | 0.95 | 0.87 | 0.82 | 0.69 | 0.57 | 0.94 | 0.84 | 1.00 | 0.98 | 0.87 | 0.93 |
| XRP | 0.87 | 0.75 | 0.85 | 0.88 | 0.92 | 0.98 | 0.92 | 0.72 | 0.99 | 0.90 | 0.96 | 0.85 | 0.80 | 0.72 | 0.59 | 0.92 | 0.82 | 0.98 | 1.00 | 0.85 | 0.95 |
| ZEC | 0.69 | 0.54 | 0.71 | 0.72 | 0.83 | 0.87 | 0.83 | 0.64 | 0.86 | 0.75 | 0.81 | 1.00 | 0.95 | 0.64 | 0.48 | 0.92 | 0.98 | 0.87 | 0.85 | 1.00 | 0.80 |
| ZRX | 0.82 | 0.82 | 0.80 | 0.87 | 0.97 | 0.93 | 0.97 | 0.80 | 0.94 | 0.85 | 0.99 | 0.80 | 0.76 | 0.80 | 0.67 | 0.87 | 0.78 | 0.93 | 0.95 | 0.80 | 1.00 |

We performed a further experiment to compare the list of categories of the features that are more relevant for predicting Uptrend or Downtrend. We considered three different windows/time periods (P1, P2, P3) to analyze also the impact of the time dimension. Table 7 reports the results. For each cryptocurrency, we report the computed correlations in P1, P2, and P3. For almost all cryptocurrencies, the correlation value is stable with respect to the time slot and is higher than 0.7. Hence, for almost all cryptocurrencies, the decision about the class label is based on the same categories of features independently of the predicted label.

Table 7. Uptrend/Downtrend correlations.

| Crypto | P1 | P2 | P3 |
|--------|------|------|------|
| ADA | 0.72 | 0.62 | 0.80 |
| BCH | 0.82 | 0.90 | 0.66 |
| BNB | 0.92 | 0.76 | 0.80 |
| BTC | 0.90 | 0.83 | 0.78 |
| BTG | 0.96 | 0.90 | 0.93 |
| DASH | 0.77 | 0.99 | 0.97 |
| DOGE | 0.87 | 0.75 | 0.99 |
| EOS | 0.91 | 1.00 | 0.93 |
| ETC | 0.77 | 0.65 | 0.74 |
| ETH | 0.69 | 0.81 | 0.75 |
| LINK | 0.82 | 0.78 | 0.71 |
| LTC | 0.83 | 0.94 | 0.84 |
| NEO | 0.94 | 0.64 | 0.64 |
| QTUM | 0.60 | 0.76 | 0.93 |
| TRX | 0.76 | 0.84 | 0.82 |
| WAVE | 0.89 | 0.75 | 0.76 |
| XEM | 0.59 | 0.82 | 0.67 |
| XMR | 0.95 | 0.98 | 0.79 |
| XRP | 0.80 | 0.82 | 0.86 |
| ZEC | 0.87 | 0.70 | 0.87 |
| ZRX | 0.59 | 0.89 | 0.87 |

7. Discussion

Explainability plays an important role in many Machine Learning-driven applications, including quantitative cryptocurrency trading. Despite their accuracy, ML models are deemed as not reliable enough, as domain experts do not trust the automated solutions. In the financial, in particular, a clear explanation of the rationale behind machine-driven decisions is deemed as unavoidable.

EXplainable AI opens the ML black boxes providing global or local explanations based on the underlying data features. Due to their high dimensionality and multi-faceted nature, cryptocurrencies are particularly suited to eXplainable AI. The main purposes are:

- The enhancement of existing cryptocurrency trading systems based on the collected feedback on the current market trends.
- The online support to discretionary traders, who commonly monitor the financial markets and execute trading operations in real time.

CryptoMLE is designed for supporting the monitoring of ML model performances on cryptocurrency markets. To enhance trading system strategies, CryptoMLE helps cryptocurrency investors verify the predictive rules inferred by the ML algorithms against the domain knowledge. To support online discretionary traders' activities, it shortlists the most influential cryptocurrency features that are worth monitoring. The SHAP series and the SHAP summary plots provide them with a simple, interactive environment to obtain actionable feedback based on the recent ML outcomes.

The main takeaways from the empirical outcomes can be summarized as follows:

- Feature relevance to cryptocurrency price forecasting is either *generalized*, i.e., valid for all cryptocurrencies independently of time periods and market conditions (e.g., for the *Close_rel* feature), or *selective*, i.e., valid only for a subset of features and for specific time periods. For the latter feature subset, CryptoMLE provides experts with an automated way to recognize them and leverage their predictive power for quantitative trading.
- Based on the prediction outcomes, the relevance of the individual features is highly variable (see the dashboard snapshots in Figures 6–11). To drive short-term cryptocurrency investments, it is crucial to monitor the most likely causes of market movements. For example, the percentage variations of the trading volume between current and previous days (namely *volume_pct**) appear to be relevant to predict BCHUSD variations (see Figure 9), whereas they are less influential in the prediction of other cryptocurrency prices.
- The influence of feature subcategories and categories is less sensitive to the market conditions, but they can be tailored to particular cryptocurrencies. For example, a cryptocurrency is more likely to be more influenced by BC features than others. This can be easily verified using CryptoMLE in real trading simulations.
- The discrepancies between the observed results among the target class (e.g., Uptrend, Downtrend) are often negligible. Therefore, traders relying on both long- and short-selling trading strategy can easily and quickly interact with CryptoMLE to gather all the required information.
- Simpler ML models analyzing only the prices of the target cryptocurrency assets appear to be suboptimal because, according to the achieved results, cryptocurrency prices are likely to be relevantly influenced by many other features (see, for instance, BCHUSD and ETHUSD). This confirms the utility of the CryptoMLE graphical interface, which provides human experts with a summary of the main feature contributions to the ML predictors.

8. Conclusions and Future Works

This paper introduced an explainable AI tool for cryptocurrency price forecasting. It presented a visual interface based on which domain experts can infer actionable dependencies among input data features and Machine Learning predictions. The interactive dashboard consists of an SHAP series plot, showing the temporal variation of the mean Shapley values associated with the most recent ML predictions, and a selection of pop-up summary plots, which are snapshots of the main features' influences at given time points. The empirical simulation, which was run on a 8-year period, showed the variability of the model explanations across 21 cryptocurrencies and three reference time periods in terms of selected features, feature subcategories and categories.

As future work, we plan to leverage the Shapley values in quantitative intraday trading. Specifically, we aim at dynamically adapt algorithmic decisions in cryptocurrency trading based on the relevant feedback provided by domain experts through the graphical interface.

Author Contributions: Conceptualization, J.F. and L.C.; Methodology, J.F., L.C. and P.G.; Software, J.F.; Supervision, L.C. and P.G.; Writing—original draft, J.F., L.C. and P.G.; Writing—review and editing, J.F., L.C. and P.G. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Not applicable.

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

References

- Giudici, G.; Milne, A.; Vinogradov, D. Cryptocurrencies: Market analysis and perspectives. *J. Ind. Bus. Econ.* **2019**, *47*, 1–18. [CrossRef]
- Nakamoto, S. Bitcoin: A Peer-to-Peer Electronic Cash System. Available online: <https://bitcoin.org/bitcoin.pdf> (accessed on 1 June 2022).
- Alessandretti, L.; ElBahrawy, A.; Aiello, L.M.; Baronchelli, A. Anticipating Cryptocurrency Prices Using Machine Learning. *Complexity* **2018**, *2018*, 8983590:1–8983590:16. [CrossRef]
- Sun, J.; Zhou, Y.; Lin, J. Using machine learning for cryptocurrency trading. In Proceedings of the IEEE International Conference on Industrial Cyber Physical Systems, Taipei, Taiwan, 6–9 May 2019; pp. 647–652. [CrossRef]
- Koker, T.E.; Koutmos, D. Cryptocurrency Trading Using Machine Learning. *J. Risk Financ. Manag.* **2020**, *13*, 178. [CrossRef]
- Catania, L.; Grassi, S.; Ravazzolo, F. Forecasting cryptocurrencies under model and parameter instability. *Int. J. Forecast.* **2019**, *35*, 485–501. [CrossRef]
- Shahbazi, Z.; Byun, Y. Machine Learning-Based Analysis of Cryptocurrency Market Financial Risk Management. *IEEE Access* **2022**, *10*, 37848–37856. [CrossRef]
- Ortu, M.; Uras, N.; Conversano, C.; Bartolucci, S.; Destefanis, G. On technical trading and social media indicators for cryptocurrency price classification through deep learning. *Expert Syst. Appl.* **2022**, *198*, 116804. [CrossRef]
- Lahmiri, S.; Bekiros, S.D. Deep Learning Forecasting in Cryptocurrency High-Frequency Trading. *Cogn. Comput.* **2021**, *13*, 485–487. [CrossRef]
- Sridhar, S.; Sanagavarapu, S. Multi-Head Self-Attention Transformer for Dogecoin Price Prediction. In Proceedings of the 2021 14th International Conference on Human System Interaction (HSI), Gdańsk, Poland, 8–10 July 2021; pp. 1–6. [CrossRef]
- Dâ€™Amato, V.; Levantesi, S.; Piscopo, G. Deep learning in predicting cryptocurrency volatility. *Phys. A Stat. Mech. Its Appl.* **2022**, *596*, 127158. [CrossRef]
- Fang, F.; Ventre, C.; Basios, M.; Kong, H.; Kanthan, L.; Li, L.; Martínez-Rego, D.; Wu, F. Cryptocurrency Trading: A Comprehensive Survey. *arXiv* **2020**, arXiv:2003.11352. Available online: <https://arxiv.org/abs/2003.11352> (accessed on 1 June 2022).
- Lundberg, S.M.; Lee, S.I. A Unified Approach to Interpreting Model Predictions. In *Advances in Neural Information Processing Systems*; Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R., Eds.; Curran Associates, Inc.: Red Hook, NY, USA 2017; Volume 30.
- Shapley, L.S. A Value for n-Person Games. In *Contributions to the Theory of Games (AM-28), Volume II*; Kuhn, H.W., Tucker, A.W., Eds.; Princeton University Press: Princeton, NJ, USA, 2016; pp. 307–318. [CrossRef]
- Webber, W.; Moffat, A.; Zobel, J. A Similarity Measure for Indefinite Rankings. *ACM Trans. Inf. Syst.* **2010**, *28*, 1–38. [CrossRef]
- Chen, J.H.; Wu, C.H.; Tsai, Y.C.; Chen, S.Y.-C. Explainable Digital Currency Candlestick Pattern AI Learner. In Proceedings of the 2022 14th International Conference on Knowledge and Smart Technology (KST), Chon Buri, Thailand, 26–29 January 2022; pp. 91–96. [CrossRef]
- Carta, S.M.; Consoli, S.; Piras, L.; Podda, A.S.; Recupero, D.R. Explainable Machine Learning Exploiting News and Domain-Specific Lexicon for Stock Market Forecasting. *IEEE Access* **2021**, *9*, 30193–30205. [CrossRef]
- Todorovska, A.; Spirovska, E.; Angelovski, G.; Peshov, H.; Rusevski, I.; Marojevik, J.; Vodenska, I.; Chitkushev, L.T.; Trajanov, D. Analysis of Cryptocurrency Interdependencies. In Proceedings of the Blockchain in Kyoto 2021 (BCK21), Kyoto, Japan, 17–18 February 2021; Physical Society of Japan: Kyoto, Japan, 2021. [CrossRef]
- Liu, Y.; Zhang, L. Cryptocurrency Valuation and Machine Learning. *SSRN Electron. J.* **2020**, 1–13. [CrossRef]
- Babaei, G.; Giudici, P.; Raffinetti, E. Explainable artificial intelligence for crypto asset allocation. *Financ. Res. Lett.* **2022**, *47*, 102941. [CrossRef]
- Freeborough, W.; van Zyl, T. Investigating Explainability Methods in Recurrent Neural Network Architectures for Financial Time Series Data. *Appl. Sci.* **2022**, *12*, 1427. [CrossRef]
- Meyes, R.; Lu, M.; de Puiseau, C.W.; Meisen, T. Ablation Studies in Artificial Neural Networks. *arXiv* **2019**, arXiv:1901.08644.
- Yin, T.; Du, X.; Zhang, W.; Zhao, Y.; Han, B.; Yan, J. Real-Trading-Oriented Price Prediction with Explainable Multiobjective Optimization in Quantitative Trading. *IEEE Access* **2022**, *10*, 57685–57695. [CrossRef]

24. Ke, G.; Meng, Q.; Finley, T.; Wang, T.; Chen, W.; Ma, W.; Ye, Q.; Liu, T.Y. LightGBM: A Highly Efficient Gradient Boosting Decision Tree. In Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS'17), Long Beach, CA, USA, 4–9 December 2017; Curran Associates Inc.: Red Hook, NY, USA, 2017; pp. 3149–3157.
25. Tan, P.; Steinbach, M.S.; Karpatne, A.; Kumar, V. *Introduction to Data Mining*, 2nd ed.; Pearson: London, UK 2019.
26. Jabbar, A.; Li, X.; Omar, B. A Survey on Generative Adversarial Networks: Variants, Applications, and Training. *ACM Comput. Surv.* **2022**, *54*, 157. [[CrossRef](#)]
27. Murphy, J. Technical Analysis of the Financial Markets: A Comprehensive Guide to Trading Methods and Applications. 1999. Available online: https://books.google.it/books/about/Technical_Analysis_of_the_Financial_Mark.html?hl=it&id=5zhXEqr_IcC&redir_esc=y (accessed on 10 January 2019).
28. Sokolovsky, A.; Arnaboldi, L.; Bacardit, J.; Gross, T. Interpretable ML-Driven Strategy for Automated Trading Pattern Extraction. Available online: https://www.researchgate.net/publication/350341294_Interpretable_ML-driven_Strategy_for_Automated_Trading_Pattern_Extraction (accessed on 1 June 2022).
29. Hu, W.; Si, Y.W.; Fong, S.; Lau, R.Y.K. A formal approach to candlestick pattern classification in financial time series. *Appl. Soft Comput.* **2019**, *84*, 105700. [[CrossRef](#)]
30. Gamper, J.; Dignös, A. Processing Temporal and Time Series Data: Present State and Future Challenges. In *Advances in Databases and Information Systems*; Darmont, J., Novikov, B., Wrembel, R., Eds.; Springer International Publishing: Cham, Switzerland, 2020; pp. 8–14.
31. Carta, S.; Podda, A.S.; Recupero, D.R.; Stanciu, M.M. Explainable AI for Financial Forecasting. In *Machine Learning, Optimization, and Data Science, Proceedings of the 7th International Conference, LOD 2021, Grasmere, UK, 4–8 October 2021*; Revised Selected Papers, Part II; Lecture Notes in Computer Science; Nicosia, G., Ojha, V., Malfa, E.L., Malfa, G.L., Jansen, G., Pardalos, P.M., Giuffrida, G., Umeton, R., Eds.; Springer: Berlin/Heidelberg, Germany, 2021; Volume 13164, pp. 51–69. [[CrossRef](#)]
32. Molnar, C.; Casalicchio, G.; Bischl, B. Interpretable Machine Learning—A Brief History, State-of-the-Art and Challenges. In Proceedings of the ECML PKDD 2020 Workshops—Workshops of the European Conference on Machine Learning and Knowledge Discovery in Databases (ECML PKDD 2020): SoGood 2020, PDFL 2020, MLCS 2020, NFMCP 2020, DINA 2020, EDML 2020, XKDD 2020 and INRA 2020, Ghent, Belgium, 14–18 September 2020; Communications in Computer and Information Science; Koprinska, I., Kamp, M., Appice, A., Loglisci, C., Antonie, L., Zimmermann, A., Guidotti, R., Özgöbek, Ö., Ribeiro, R.P., Gavaldà, R., et al., Eds.; Springer: Berlin/Heidelberg, Germany, 2020; Volume 1323, pp. 417–431. [[CrossRef](#)]
33. Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; et al. Scikit-learn: Machine Learning in Python. *J. Mach. Learn. Res.* **2011**, *12*, 2825–2830.