

A Machine Learning Approach for Bitcoin Forecasting [†]

Stefano Sossi-Rojas ^{1,*}, Gissel Velarde ^{1,2,*} and Damian Zieba ^{3,*}¹ Computational Systems Engineering, Universidad Privada Boliviana, Cochabamba 3967, Bolivia² Vodafone GmbH, 40549 Düsseldorf, Germany³ Faculty of Economic Sciences, University of Warsaw, 00927 Warsaw, Poland

* Correspondence: tenosossi2012@gmail.com (S.S.-R.); gissel.velarde@vodafone.com (G.V.); dzieba@wne.uw.edu.pl (D.Z.)

[†] Presented at the 9th International Conference on Time Series and Forecasting, Gran Canaria, Spain, 12–14 July 2023.

Abstract: Bitcoin is one of the cryptocurrencies that has gained popularity in recent years. Previous studies have shown that closing price alone is not enough to forecast its future level, and other price-related features are necessary to improve forecast accuracy. We introduce a new set of time series and demonstrate that a subset is necessary to improve directional accuracy based on a machine learning ensemble. In our experiments, we study which time series and machine learning algorithms deliver the best results. We found that the most relevant time series that contribute to improving directional accuracy are open, high, and low, with the largest contribution of low in combination with an ensemble of a gated recurrent unit network and a baseline forecast. The relevance of other Bitcoin-related features that are not price-related is negligible. The proposed method delivers similar performance to the state of the art when observing directional accuracy.

Keywords: Bitcoin; forecasting; time series; machine learning

1. Introduction

Bitcoin price forecasting has been the focus of many studies in the literature over the years. Nevertheless, unexpected price movements, smaller and larger bubbles, and different short- and long-term trends mean this task is an ongoing topic for research. In one of the recent studies exploring this area of research [1], the authors verify the performance of different machine learning algorithms and mention the current state of the knowledge in Bitcoin forecasting. One aspect refers to [2] the price of Bitcoin being mainly driven by the spot market rather than the futures market. Another aspect is the division of approaches to forecasting into the “Blockchain approach” and “Financial market’s approach”. The former is based on technical variables such as hash rates or mining difficulty, while the latter uses standard econometric variables such as stocks, bonds, or gold price. We can add to this another approach, which involves taking variables related to investor sentiment such as Google trends data [3], uncertainty indices (VIX, UCRY, see [4]), or the Fear and Greed index, which has not been explored much in the literature. In our study, we used a mix of features from these three approaches and the open, high, low, and close prices to predict the Bitcoin price using a novel approach. Based on this, we were able to verify which variables contribute to the forecast accuracy the most. In fact, the features with the largest contribution are only the price-related ones. The other variables, whether coming from the “Blockchain approach”, the “Financial markets approach”, or the “Sentiment approach” have a negligible impact in terms of improving Bitcoin’s price performance.

We used a method that explores features and machine learning algorithms for Bitcoin’s closing price prediction. In recent forecasting competitions, it was observed that machine learning models and hybrid approaches demonstrated superiority over alternative methods [5]. Therefore, in this study, we evaluated the following machine learning algorithms:



Citation: Sossi-Rojas, S.; Velarde, G.; Zieba, D. A Machine Learning Approach for Bitcoin Forecasting. *Eng. Proc.* **2023**, *39*, 27. <https://doi.org/10.3390/engproc2023039027>

Academic Editors: Ignacio Rojas, Hector Pomares, Luis Javier Herrera, Fernando Rojas and Olga Valenzuela

Published: 29 June 2023



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

Long Short-Term Memory (LSTM), Bidirectional Long-Term Memory (BiLSTM), Gated Recurrent Unit (GRU), Bidirectional Gated Recurrent Unit (BiGRU), and Light Gradient Boosting Machine (LightGBM). In addition we used ensembling. The results were evaluated observing the predicted and actual Bitcoin closing price measuring Root Mean Squared Error (RMSE), Mean Squared Error (MSE), Mean Absolute Error (MAE), and Directional Accuracy (DA).

Although previous studies show that the financial time series closing price, either in the case of stocks [6,7], commodities [8], or cryptocurrencies [9], is not enough for prediction when training deep learning models, we found that ensembling does help to improve the forecasting accuracy. In this work, we demonstrate that the Bitcoin closing price is not sufficient for forecasting and additional features are necessary when using machine learning algorithms. Evaluation on return shows that the method developed in this work presents one of the highest scores, with a directional accuracy score equal to 0.7645, exceeding a baseline by 58.24 percent.

We compare our work to related studies for Bitcoin prediction. In [10], LSTM was used in combination with the Empirical Wavelet Transform (EWT) decomposition technique. The authors used the Intrinsic Mode Function (IMF) to optimize and estimate outputs with Cuckoo Search (CS) [10]. In [11], Linear Regression (LR) techniques and particle swarm optimization were used to train and forecast data from beginning of 2012 to the end of March 2020. The best setup for the model was obtained with 42 days plus 1 standard deviation [11]. In [12], Autoregressive Integrated Moving Average (ARIMA) was used for data from 1 May 2013 to 7 June 2019. This model works best for short-term predictions and can be used to predict Bitcoin for one to seven days ahead [12]. Finally, in [13], a BiLSTM with Low-Middle-High features (LMH-BiLSTM) was tested with two primary steps: data decomposition and bidirectional deep learning. The results demonstrate that the proposed model outperforms other benchmark models and achieved high investment returns in the buy-and-hold strategy in a trading simulation [13].

In this work, we tested the previously mentioned machine learning algorithms one by one, and also in ensemble. We fed the algorithms with a new set of 13 time series (see Section 2.1). In addition, we included 11 signals that come from Variational Mode Decomposition (MMD) as proposed in [13]. However, in our experiments, data decomposition did not provide significant improvements. Next, we explain our proposed method.

2. Materials and Methods

An overview of the method is presented in Figure 1. The input data is prepared in three steps. First, it is normalized between 0 and 1. Then, the train and test set partitions are created, where the train set is used for training several machine learning algorithms. The next phase considers algorithm training. The details regarding hyperparameter selection are described in Section 2.4. Next, the evaluation phase is performed as a rolling forecast for 1-step ahead prediction over the test set.

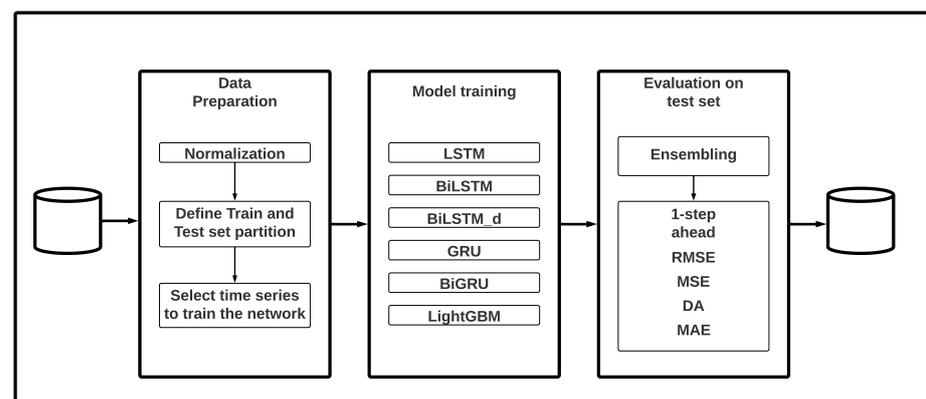


Figure 1. Visual summary of the method.

2.1. Data Collection

We collected the daily Bitcoin closing price from 7 October 2013 to 6 November 2022. We used a public API from the Kraken page [14]; the data collected were the values of close, open, high, low, volume, and date. With the Nasdaq-Data-Link library for Python [15], the following values were obtained: transaction fee, estimated Bitcoin USD transaction volume, Bitcoin USD exchange trade volume, and Bitcoin hash rate. Bitcoin Google trends were obtained with the pytrends library [16]. The gold to USD exchange rate was obtained from the Investing.com page. The Fear and Greed Index was obtained from the Kaggle page [17]. The moving average of the closing value was added, taking the last 30 days. Table 1 presents all features used in this study. A similar set of features was used in [13]: Bitcoin price, Bitcoin transaction fees as Bitcoin miner's revenue divided by transactions, USD trade volume from the top Bitcoin exchanges, Bitcoin transaction volume, USD exchange or trade volume from the top Bitcoin exchanges, gold exchange rate to US dollar, hash rate, and Google trends of Bitcoin.

Table 1. Series, short name, description, and count. Originally some signals had more samples than others; therefore, the dates where no information was recorded for all signals were removed to obtain the same number of samples per signal, down sampling to 3812 in the range 19 November 2013 to 4 November 2022.

Series	Features	Description	Count
Series 1	Close	Daily Bitcoin close price	4379
Series 2	Low	Daily Bitcoin low price	4379
Series 3	High	Daily Bitcoin high price	4379
Series 4	Open	Daily Bitcoin open price	4379
Series 5	Trans_Volume	Bitcoin transaction volume in dollars	3812
Series 6	Volume	Daily quantity of Bitcoins bought or sold	4379
Series 7	Hash_Rate	Number of giga hashes Bitcoin network performed	3812
Series 8	Trans_Fees	Bitcoin miner's revenue divided by transactions	3812
Series 9	XAU_USD	Gold (XAU) Exchange rate to US dollar (USD)	3812
Series 10	Trade_Volume	Bitcoin trade volume in dollars	3812
Series 11	Google_Trend	Bitcoin's Google Trend	3812
Series 12	Fear_Greed	Fear and Greed Index. It is a way to gauge stock market movements and whether stocks are fairly priced	3812
Series 13	Moving_Avg_30	Moving average of Bitcoin's closing price for the last 30 days.	3783

2.2. Variational Mode Decomposition (VMD)

Bitcoin' closing price was decomposed using the VMD method as proposed in [13]. Each decomposed mode was labeled M0 through M10, where M0 has the lowest frequency and M10 has the highest. We can observe the graph of the decomposition in Figure 2.

Variational Mode Decomposition (VMD) is a completely nonrecursive signal decomposition technique proposed by [18]. VMD is a problem of variational optimization that aims to minimize the total bandwidth of each mode. This work used the vmdpy python library [19], with parameters by default using Bitcoin's close price as the input and a bandwidth of 5000.

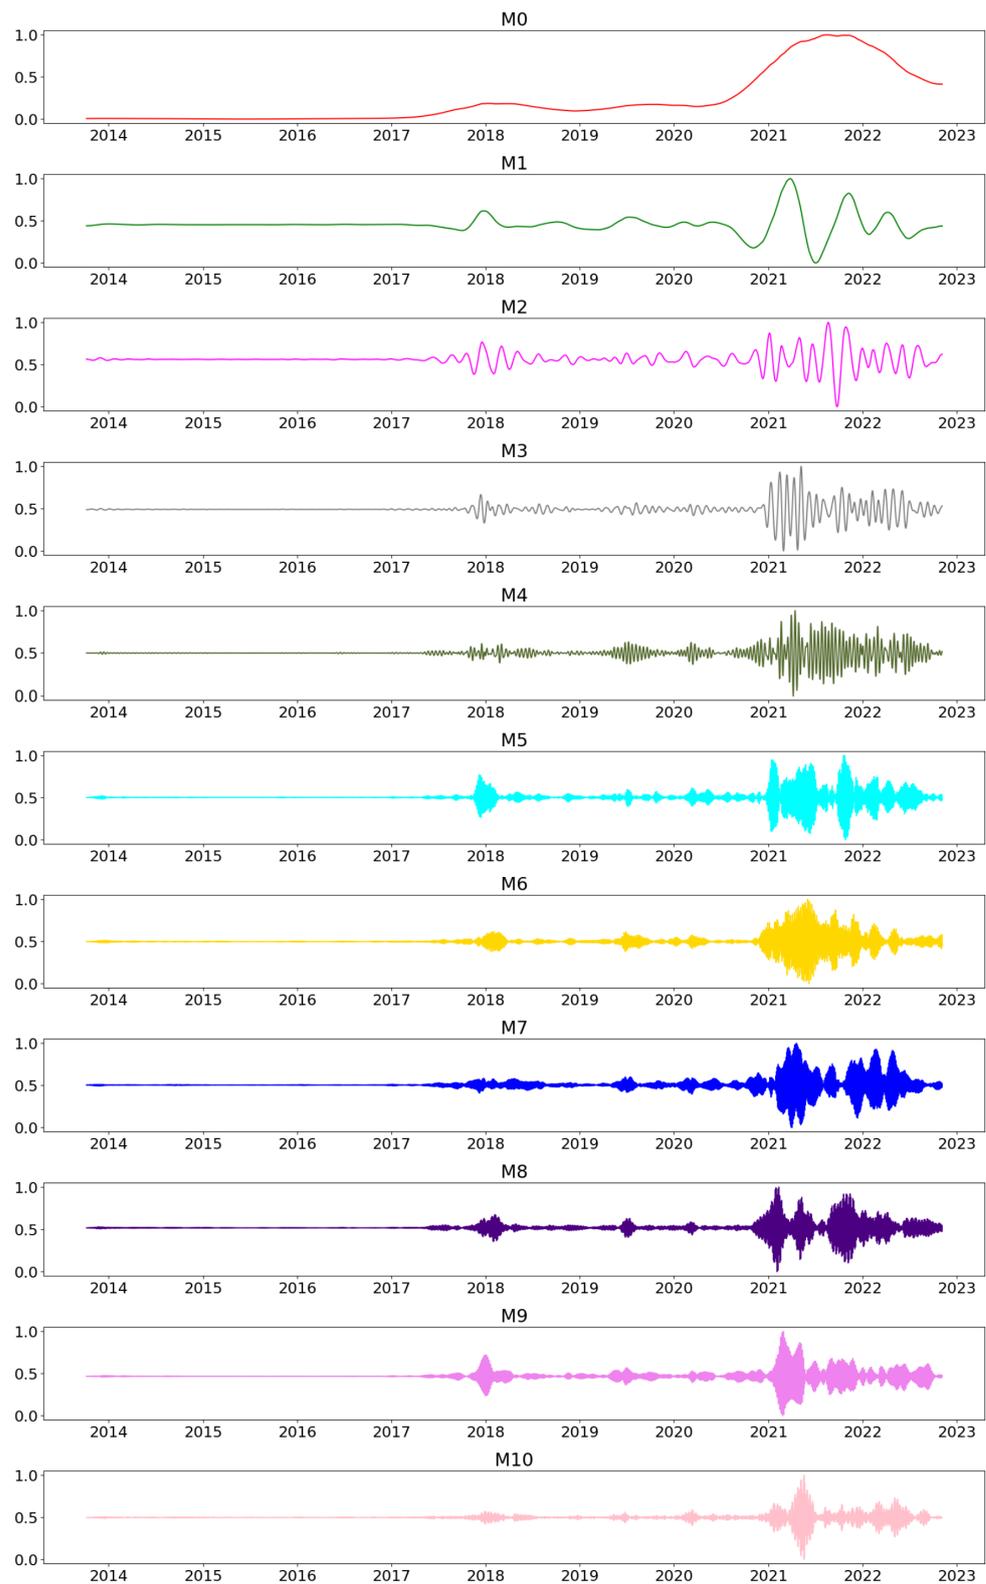


Figure 2. Variational Mode Decomposition (VMD) decomposition of Bitcoin’s closing price, normalized between 0 and 1, from 7 October 2013 to 6 November 2022.

2.3. Data Preparation

For the LSTM, BiLSTM, GRU, and BiGRU models, the data were scaled between 0 and 1 before training, except the LightGBM model for which data scaling was performed for evaluation only. Next, the data was divided into 25-day windows, converting the data

table into 3D lists (arrays), where the first dimension corresponds to the batch size, the second to the number of time-steps, and finally, the third dimension to the number of units of one input sequence [20]. Next to each list, the expected future value was saved. This last value is taken from the closing values of the previous day. For the LightGBM model, the complete data Table 1 was used without modifications. Data partitioning was performed as follows: the set was trained from 7 October 2013 to 8 August 2022, and the set was tested from 9 August 2022 to 6 November 2022, where 90 days were used for testing, and the rest of data were used for training.

2.4. Model Training

We tested five deep learning architectures and one tree boosting method. All deep learning networks present an input layer of 90 units, a set value of 500 epochs, and a batch size of 64 without early stopping.

- **Long Short-Term Memory (LSTM):** The network trains with five layers, an input layer with the activation function seen in [6], the bias initializer glorot uniform, kernel regulator l1, l2, kernel constraint unit norm, and the time_major activated, followed by a dense layer of 90 units and linear activation. Then, an output layer and another dense layer is used. The model uses the Adam optimizer with a learning rate of 0.002;
- **Gated Recurrent Unit (GRU):** The network trains with five layers, an input layer followed by a dropout layer set to 0.3, an output layer followed by another dropout layer, and a dense layer of one unit. The model uses the Adam optimizer with a learning rate of 0.0001. A similar network was used in [21];
- **Bidirectional Long Short-Term Memory (BiLSTM):** The network consists of an input layer with the *tanh* activation function, followed by a backward learning layer and a dense layer. The model uses the Adam optimizer with a learning rate of 0.01. A similar network was used in [13];
- **Bidirectional Long Short-Term Memory with dropout (BiLSTM_d):** The network consists of an input layer with *tanh* activation function, followed by a dropout layer, followed by a backward learning layer, followed by a dropout layer and a dense layer. The model uses the Adam optimizer with a learning rate of 0.01 and the dropout set to 0.3;
- **Bidirectional Gated Recurrent Unit (BiGRU):** The network trains with five layers, an input bidirectional layer followed by a dropout layer set to 0.3, an output bidirectional layer followed by another dropout layer and a dense layer. The model uses the Adam optimizer with a learning rate of 0.0001;
- **Light Gradient Boosting Machine (LighGBM):** This presents an early stopping round set to 50, and verbose evaluation set to 30 with 3600 number of boost rounds. The model trains with a gradient booting decision tree, with the objective set to *tweedie* and a variance power of 1.1, and uses an RMSE metric with n-jobs set to -1 . In addition, it uses 42 seeds with a learning rate of 0.2, the bagging fraction is set to 0.85 and the bagging frequency is set to 7. Moreover, colsample by tree and colsample by node are set to 0.85 with a min data per leaf of 30, and the number of leaves is 200 with lambda l1 and l2 set to 0.5. A similar network was used in [5].

2.5. Evaluation

The predicted results were normalized between 0 and 1 before evaluation measurements were made.

We measure the Mean Squared Error (MSE), Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and Directional Accuracy (DA) between the predicted and actual closing price as described in [22], such that n is the number of samples, and y_t and x_t are the predicted and actual closing price at time t :

$$MSE = n^{-1} \sum_{t=1}^n (x_t - y_t)^2. \quad (1)$$

$$RMSE = \sqrt{MSE}. \tag{2}$$

$$MAE = \frac{\sum_{t=1}^n |y_t - x_t|}{n}. \tag{3}$$

$$DA = \frac{100}{n} \sum_{t=1}^n d_t, \tag{4}$$

where

$$d_t = \begin{cases} 1 & (x_t - x_{t-1})(y_t - y_{t-1}) \geq 0 \\ 0 & otherwise. \end{cases}$$

2.6. Return

The return is a financial measure used to assess the efficiency of an asset investment. It is an growth indicator of the value of an investment during a certain period of time. Return On Investment (ROI) is one of the main financial measures used both in the traditional stock market and in the world of cryptocurrencies [23]. The formula can be expressed in terms of the final Value of Inversion (FVI) and the Initial Value of Inversion (IVI):

$$ROI = \left(\frac{FVI - IVI}{IVI}\right)100\%. \tag{5}$$

3. Results

Experiments were conducted to see how different time series influence Bitcoin closing price prediction and how different models perform. The experiments were carried out analyzing the results with different factors. Tables 2–11 show the results obtained in the different experiments. An RMSE, MSE, MAE close to zero and a DA close to one are preferred. The best results are highlighted. The values were compared with values obtained with a Baseline prediction. Baseline prediction means that the predicted value is the last observed value. Series importance for the prediction was obtained according to the LightGBM model in order to classify them later and continue with the experiments, as can be seen in Figure 3.

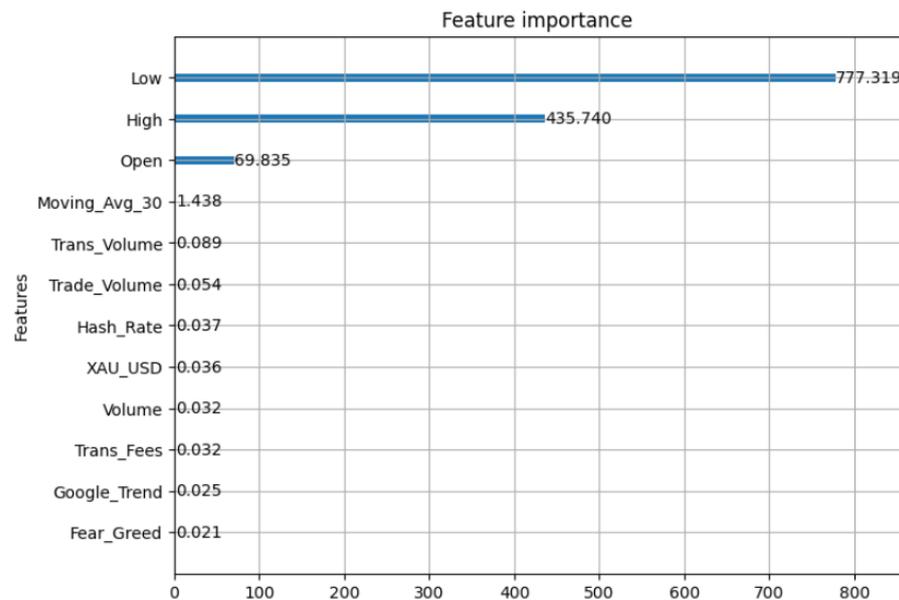


Figure 3. Feature importance found by LightGBM. The most important features are low, high, and open.

We performed the following experiments:

- **Experiment 1.** The first experiment sent a subset of all tested time series, as we assumed these features represent price action over a set period of time, and in combination could be used to predict price movements. The time series used were close, open, high, low, and volume of Bitcoin. These were obtained from the Kraken API [14]. See Table 2 for the results of the prediction. Table 2 shows that BiLSTM was best DA = 0.5056 but BiGRU produced the lowest MAE = 0.0467;
- **Experiment 2.** In the second experiment, we tested all time series presented in Table 1: close, open, high, low, volume, transaction fee, estimated Bitcoin USD transaction volume, Bitcoin USD exchange trade volume, rate of Bitcoin hash, Bitcoin Google trends, and gold to USD exchange rate, Fear and Greed index, and the moving average of the closing value. The main idea was to enrich the input data to help improve the prediction. The results can be seen in Table 3. LightGBM had the best performance and there was an improvement in comparison with Experiment 1 whereby less input data were used;
- **Experiments 3.** For the following experiments, series importance for the prediction was analyzed according to the LightGBM model. The order of importance can be seen in Figure 3. The experiments were performed indicating different combinations of the four most important series: open, high, low, and close values. These time series exhibited a significant correlation with the closing price of Bitcoin and exerted an influence over its behavior. For this experiment, open, high, and low values were used. Table 4 shows that BiLSTM had the best DA = 0.4832, but GRU had the best MAE = 0.0496. The results show a decrease in performance;
- **Experiment 4.** For this experiment, high and low values were used. Table 5 shows that BiLSTM had the best performance with DA = 0.5169. There was an improvement in comparison with Experiment 1, but it did not reach the performance of Experiment 2;
- **Experiment 5.** For this experiment, low values were used. Table 6 shows that LSTM had the best performance with DA = 0.5169. The results show a similar performance in comparison with Experiment 4;
- **Experiment 6.** For this experiment, open and low values were used. Table 7 shows that LSTM had the best DA = 0.5393 and BiLSTM delivered the lowest errors. The results show an improvement in performance compared with previous experiments, but it was not better than Experiment 2;
- **Experiment 7.** Since Experiment 2 had the best performance until this point, we tested all time series used in Experiment 2 and added 11 VMD modes to assess their impact on prediction. That is, the input data included the values of close, open, high, low, volume, transaction fee, estimated Bitcoin USD transaction volume, Bitcoin USD exchange trade volume, rate of Bitcoin hash, Bitcoin Google trends, and gold to USD exchange rate, Fear and Greed index, and the moving average of the closing value and 11 VMD modes. The results are presented in Table 8. This time, BiGRU had the best DA = 0.5730 but LSTM had the lowest errors. The results show a big improvement in performance compared with the previous experiments;
- **Experiment 8.** This experiment consisted of sending the values of all 11 VMD modes as input data, that is, only the 11 VMD modes were added to assess the impact these had on the prediction (see Table 9). BiGRU had the best DA = 0.5618, BiLSTM the best MAE, and LSTM the best MSE and RMSE. Results show an improvement in performance compared with previous experiments, but it was not better than Experiment 7.

In all the experiments, we observed a noticeable improvement in the BiLSTM model when it did not present the dropout layer.

Table 2. Prediction measured values with Bitcoin close, open, high, low, and volume values.

Network	Measured Value			
	MAE	MSE	RMSE	DA
GRU	0.0541	0.0111	0.1053	0.4494
BiGRU	0.0467	0.0107	0.1032	0.4494
LSTM	0.0660	0.0125	0.1120	0.4719
BiLSTM	0.0954	0.0163	0.1277	0.5056
BiLSTM_d	0.1610	0.0667	0.2583	0.4832
LightGBM	0.0513	0.0106	0.1028	0.4157

Table 3. Prediction measured values with all factors.

Network	Measured Value			
	MAE	MSE	RMSE	DA
GRU	0.1426	0.0496	0.2228	0.3371
BiGRU	0.1501	0.0660	0.2569	0.3596
LSTM	0.1731	0.0666	0.2580	0.3708
BiLSTM	0.2019	0.0864	0.2940	0.4382
BiLSTM_d	0.2685	0.1109	0.3330	0.3371
LightGBM	0.0566	0.0154	0.1240	0.5618

Table 4. Prediction measured values with Bitcoin open, high, and low values.

Network	Measured Value			
	MAE	MSE	RMSE	DA
GRU	0.0496	0.0103	0.1014	0.4494
BiGRU	0.0508	0.0110	0.1050	0.4719
LSTM	0.1029	0.0352	0.1875	0.4382
BiLSTM	0.0812	0.0178	0.1334	0.4832
BiLSTM_d	0.1317	0.0571	0.2390	0.4607
LightGBM	0.0864	0.0157	0.1254	0.3483

Table 5. Prediction measured values with high and low values of Bitcoin.

Network	Measured Value			
	MAE	MSE	RMSE	DA
GRU	0.0572	0.0112	0.1059	0.4157
BiGRU	0.0487	0.0112	0.1060	0.4607
LSTM	0.0508	0.0117	0.1083	0.4832
BiLSTM	0.0402	0.0092	0.0961	0.5169
BiLSTM_d	0.1709	0.0599	0.2447	0.4832
LightGBM	0.0796	0.0217	0.1472	0.2472

Table 6. Prediction measured values with low value of Bitcoin.

Network	Measured Value			
	MAE	MSE	RMSE	DA
GRU	0.0558	0.0116	0.1076	0.4719
BiGRU	0.5616	0.0117	0.1080	0.4607
LSTM	0.0378	0.0098	0.0988	0.5169
BiLSTM	0.0460	0.0104	0.1019	0.4719
BiLSTM_d	0.1155	0.0523	0.2286	0.4944
LightGBM	0.1213	0.5991	0.2448	0.1348

Table 7. Prediction measured values with Bitcoin open and low values.

Network	Measured Value			
	MAE	MSE	RMSE	DA
GRU	0.5125	0.0102	0.1011	0.4607
BiGRU	0.0518	0.0104	0.1022	0.4494
LSTM	0.0613	0.0151	0.1228	0.5393
BiLSTM	0.0344	0.0092	0.0958	0.4719
BiLSTM_d	0.1221	0.0623	0.2496	0.5056
LightGBM	0.0801	0.0227	0.1505	0.2809

Table 8. Prediction measured values with all factors plus VMD modes.

Network	Measured Value			
	MAE	MSE	RMSE	DA
GRU	0.1002	0.0245	0.1565	0.5506
BiGRU	0.0751	0.0230	0.1517	0.5730
LSTM	0.0705	0.0206	0.1423	0.5169
BiLSTM	0.0963	0.0222	0.1489	0.4832
BiLSTM_d	0.1081	0.0322	0.1793	0.4832
LightGBM	0.0740	0.0222	0.1490	0.4719

Table 9. Prediction measured values with 11 VMD modes.

Network	Measured Value			
	MAE	MSE	RMSE	DA
GRU	0.0806	0.0179	0.1338	0.5169
BiGRU	0.0762	0.0176	0.1327	0.5618
LSTM	0.0591	0.0167	0.1290	0.4607
BiLSTM	0.0521	0.0176	0.1329	0.4607
BiLSTM_d	0.0837	0.0209	0.1446	0.5169
LightGBM	0.1299	0.0486	0.2204	0.4494

Ensembling and Return Performance

Since we are predicting one day in the future (rolling forecast), the value of IVI seen in Equation (5) is the same value as the baseline, and FVI is the prediction of the models. Ensembling was obtained with the simple arithmetic average, using Equation (5), combining each model with the baseline. We tested each model individually. The results show that the best combination was GRU and the baseline, with the model being trained using the open, high, and low factors.

- **Experiment 9.** In this experiment, we tested ensembling. Return values were calculated over the prediction of all previous experiments. The best results were obtained using open, high, and low values as inputs. See Table 10 for the results. GRU had the best performance with DA = 0.7645. The results show a big improvement in DA compared to the previous experiments;
- **Experiment 10.** The last experiment was executed in order to make a close comparison with LMH-BiLTSM [13]. Therefore, we used approximately the same date range as used in [13]. Return values were calculated over the prediction period. See Table 11 for the results. This time, GRU had the best DA = 0.7865 but BiGRU produced the lowest errors. The DA obtained here is comparable to that reported by LMH-BiLTSM [13]. Indeed, in this experiment, we observed the highest DA prediction out of all the experiments we performed. However, the lowest errors were measured using the BiLSTM network with the open and low series: MAE = 0.0344, MSE = 0.0092, and RMSE = 0.0958, see Table 7.

Table 12 shows a comparison between previous studies and the best performing model presented in this paper. The model with the best performance consists of a GRU trained with open, high, and low values of Bitcoin. The GRU ensemble achieved 0.7865 ± 0.2113 DA, matching the performance obtained by LMH-BiLSTM [13] in a similar time range. Notice that for the period from October 2013 to 6 November 2022, the DA was slightly lower but the standard deviation was also lower. Making a comparison of the GRU network (Table 10) and the baseline, we see an improvement of 58.14 percent in directional accuracy; this was the model with the highest score.

Table 10. Return measured values for prediction with Bitcoin open, high, and low values. Range between 7 October 2013 and 6 November 2022.

Network	Measured Value			DA
	MAE	MSE	RMSE	
Baseline	0.0371	0.0091	0.0952	0.4831
GRU	0.0843	0.0296	0.1716	0.7645
BiGRU	0.2299	0.0735	0.2711	0.7191
LSTM	0.2629	0.1113	0.3337	0.5169
BiLSTM	0.1687	0.0553	0.2352	0.5169
BiLSTM_d	0.3082	0.1072	0.3274	0.7303
LightGBM	0.1730	0.0488	0.2208	0.6547

Table 11. Return measured values for prediction with all factors plus VMD modes from 7 October 2013 to 1 January 2021, which is approximately same range as that used in [13]. Baseline DA was 0.5281.

Network	Measured Value			DA
	MAE	MSE	RMSE	
Baseline	0.0496	0.0138	0.1176	0.5281
GRU	0.1745	0.0548	0.2341	0.7865
BiGRU	0.1137	0.0400	0.2001	0.7303
LSTM	0.2039	0.0860	0.2933	0.5955
BiLSTM	0.2772	0.1044	0.3232	0.6067
BiLSTM_d	0.3134	0.1332	0.3649	0.5056
LightGBM	0.2806	0.1078	0.3283	0.5730

Table 12. Performance comparisons. The methods presented in this table were trained using different datasets in different date ranges. Therefore, this comparison is relative.

Method	DA
GRU (This work Table 10, 7 October 2013 to 6 November 2022)	0.7645 ± 0.1299
GRU (This work Table 11, 7 October 2013 to 1 January 2021)	0.7865 ± 0.2113
LMH-BiLSTM [13] (29 April 2013 to 1 January 2021)	$0.8170 \pm ?$
ARIMA [12] (1 May 2013 to 7 June 2019)	$0.5719 \pm ?$
LSTM [10] (18 July 2010 to 28 March 2019)	$0.5409 \pm ?$
LR [11] (Beginning of 2012 until the end of March 2020)	$0.5155 \pm ?$

4. Conclusions

We confirm the hypothesis that Bitcoin is difficult to predict with the closing price alone, that is, the closing price does not contain enough information to predict Bitcoin, and a set of price-related time series are necessary to improve prediction. We tested 13 series as shown in Table 1, plus 11 modes decomposed using Variational Mode Decomposition (VMD). In addition, we tested various machine learning algorithms and found that a selected set of time series consisting of open, high, and low values and an ensemble

based on a GRU network combined with the value of return, or a baseline prediction, demonstrates a great improvement in the results of the experiments. Our method delivers a comparable DA when compared to the state of the art, which in contrast uses a BiLSTM with Low-Middle-High features (LMH-BiLSTM) [13].

Author Contributions: S.S.-R. Development, implementation, and writing. G.V. Conceptualization, supervision, and writing. D.Z. Advise on the project, and writing. 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.

Acknowledgments: We thank the anonymous reviewers for their feedback.

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

References

1. Chevallier, J.; Guégan, D.; Goutte, S. Is it possible to forecast the price of bitcoin? *Forecasting* **2021**, *3*, 377–420. [CrossRef]
2. Baur, D.G.; Dimpfl, T. Price discovery in bitcoin spot or futures? *J. Futur. Mark.* **2019**, *39*, 803–817. [CrossRef]
3. Kristoufek, L. BitCoin meets Google Trends and Wikipedia: Quantifying the relationship between phenomena of the Internet era. *Sci. Rep.* **2013**, *3*, 3415. [CrossRef] [PubMed]
4. Lucey, B.M.; Vigne, S.A.; Yarovaya, L.; Wang, Y. The cryptocurrency uncertainty index. *Financ. Res. Lett.* **2022**, *45*, 102147. [CrossRef]
5. Makridakis, S.; Spiliotis, E.; Assimakopoulos, V. M5 accuracy competition: Results, findings, and conclusions. *Int. J. Forecast.* **2022**, *38*, 1346–1364. [CrossRef]
6. Velarde, G.; Brañez, P.; Bueno, A.; Heredia, R.; Lopez-Ledezma, M. An Open Source and Reproducible Implementation of LSTM and GRU Networks for Time Series Forecasting. *Eng. Proc.* **2022**, *18*, 30.
7. Velarde, G. *Forecasting with Deep Learning*; White Paper; Technical Report 2(8); Vodafone, The Data Digest: Duesseldorf, Germany, 2022.
8. Ben Ameur, H.; Boubaker, S.; Ftiti, Z.; Louhichi, W.; Tissaoui, K. Forecasting commodity prices: empirical evidence using deep learning tools. *Ann. Oper. Res.* **2023**, 1–19. [CrossRef] [PubMed]
9. Lamothe-Fernández, P.; Alaminos, D.; Lamothe-López, P.; Fernández-Gámez, M.A. Deep learning methods for modeling bitcoin price. *Mathematics* **2020**, *8*, 1245. [CrossRef]
10. Altan, A.; Karasu, S.; Bekiros, S. Digital currency forecasting with chaotic meta-heuristic bio-inspired signal processing techniques. *Chaos Solitons Fractals* **2019**, *126*, 325–336. [CrossRef]
11. Cohen, G. Forecasting Bitcoin trends using algorithmic learning systems. *Entropy* **2020**, *22*, 838. [CrossRef] [PubMed]
12. Wirawan, I.M.; Widiyaningtyas, T.; Hasan, M.M. Short term prediction on bitcoin price using ARIMA method. In Proceedings of the 2019 International Seminar on Application for Technology of Information and Communication (iSemantic), Semarang, Indonesia, 21–22 September 2019; pp. 260–265.
13. Li, Y.; Jiang, S.; Li, X.; Wang, S. Hybrid data decomposition-based deep learning for Bitcoin prediction and algorithm trading. *Financ. Innov.* **2022**, *8*, 1–24. [CrossRef]
14. Kraken. 2022. Available online: <https://docs.kraken.com/rest/> (accessed on 6 November 2022).
15. Nasdaq. Nasdaq Data Link. 2021. Available online: <https://data.nasdaq.com/tools/python> (accessed on 6 November 2022).
16. Pytrends. 2015. Available online: <https://pypi.org/project/pytrends/> (accessed on 6 November 2022).
17. de Araujo, A. Crypto Fear and Greed Index. Available online: <https://www.kaggle.com/datasets/adelsondias/crypto-fear-and-greed-index/code> (accessed on 6 November 2022).
18. Dragomiretskiy, K.; Zosso, D. Variational Mode Decomposition. 2013. Available online: (accessed on 28 January 2023). [CrossRef]
19. Carvalho, V.R.; Moraes, M.F.; Braga, A.P.; Mendes, E.M. Evaluating Five Different Adaptive Decomposition Methods for EEG Signal Seizure Detection and Classification. 2020. Available online: (accessed on 6 February 2023). [CrossRef]
20. Verma, S. Input and Output Shape in LSTM (Keras). 2019. Available online: <https://www.kaggle.com/code/shivajbd/input-and-output-shape-in-lstm-keras> (accessed on 28 January 2023).
21. Rasifaghihi, N. LSTM-GRU-BiLSTM-in-TensorFlow-for-Predictive-Analytics. 2020. Available online: <https://github.com/NioushaR/LSTM-GRU-BiLSTM-in-TensorFlow-for-predictive-analytics> (accessed on 6 January 2023).

22. Wang, J.J.; Wang, J.Z.; Zhang, Z.G.; Guo, S.P. Stock index forecasting based on a hybrid model. *Omega* **2012**, *40*, 758–766. [[CrossRef](#)]
23. Phemex. Cómo Calcular el Retorno de la Inversión (ROI) de las Criptomonedas? 2021. Available online: <https://phemex.com/es/academy/como-calcular-el-roi-de-las-criptomonedas> (accessed on 6 February 2023).

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