POLITECNICO DI TORINO Repository ISTITUZIONALE

Leveraging the momentum effect in machine learning-based cryptocurrency trading

Original

Leveraging the momentum effect in machine learning-based cryptocurrency trading / Bellocca, Gian Pietro; Attanasio, Giuseppe; Cagliero, Luca; Fior, Jacopo. - In: MACHINE LEARNING WITH APPLICATIONS. - ISSN 2666-8270. - 8:(2022), p. 100310. [10.1016/j.mlwa.2022.100310]

Availability: This version is available at: 11583/2968666 since: 2022-06-26T11:17:32Z

Publisher: Elsevier

Published DOI:10.1016/j.mlwa.2022.100310

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)

Contents lists available at ScienceDirect



Machine Learning with Applications



journal homepage: www.elsevier.com/locate/mlwa

Leveraging the momentum effect in machine learning-based cryptocurrency trading



Gian Pietro Bellocca, Giuseppe Attanasio, Luca Cagliero, Jacopo Fior*

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

ARTICLE INFO

Keywords: Machine learning Cryptocurrency trading Quantitative trading systems

ABSTRACT

Cryptocurrency trading has become more and more popular among private investors. According to recent studies, the momentum effect influences the underlying market. Quantitative trading systems can leverage momentum indicators to open and close trading positions. However, existing approaches that exploit the momentum effect in cryptocurrency trading do not rely on machine learning. Since these systems are based on human generated rules they are not suited to highly volatile market conditions, which are quite common in cryptocurrency markets. This paper proposes to leverage machine learning approaches to automatically detect the momentum effect in cryptocurrency market data. For each cryptocurrency it estimates the likelihood of being affected by the momentum effect on the next trading day as well as the momentum direction. A backtesting session, performed on three very popular cryptocurrencies, shows that the machine learning models are able to predict, to a good approximation, short-term price volatility thus reducing the number of false trading signals and increasing the return on investments compared to state-of-the-art approaches.

1. Introduction

In recent years, cryptocurrencies have gained the attention of financial institutions, and, consequently, the speculative interest in Bitcoin, Ethereum, and other cryptocurrencies has significantly increased (Fang et al., 2020). Cryptocurrencies are assets characterized by peculiar price trends and exchange volumes. This is mainly due to the medium of exchange and the ownership policies, which commonly yield a significantly higher degree of price volatility compared to conventional assets (King & Koutmos, 2021). This poses the question of whether the underlying market adheres to the predictive models that are commonly applied to more traditional markets such as the stock and Forex exchanges.

The models belonging to the traditional finance theory require investors to be risk-averse and able to make rational choices. The purpose is to maximize profits without being influenced by other factors and have complete access to all the information available in the market. These models also require an effective arbitrage mechanism, which plays a critical role in determining the prices of the securities (Sinkala, 2016). The arbitrage mechanism entails giving investors the opportunity to earn by buying/selling the same asset at a lower/higher price respectively, When an arbitrage opportunity arises, it is essential that this opportunity is immediately exploited by investors. In this way the market will allow the prices to return to the right equilibrium immediately, not allowing an asset to be overvalued or undervalued for too long periods. Notably, the aforesaid assumptions are partly unrealistic for the cryptocurrency market (Fang et al., 2020). These reasons are the origin of anomalous phenomena observed in cryptocurrency markets, i.e., the momentum and reversal effects.

The momentum effect refers to the positive autocorrelation of prices or to the tendency for rising asset prices to rise further and falling prices to keep falling. Conversely, the reversal effect refers to the phenomenon whereby asset prices show a negative autocorrelation, and therefore only after a prolonged period of deviation they do revert and move back to their fundamental values. In this paper, we exploit the aforesaid effect to design a profitable cryptocurrency trading strategy.

Recent studies on cryptocurrency markets (Caporale & Plastun, 2020; Plastun et al., 2021) have empirically demonstrated that

- The intraday behavior of the cryptocurrency hourly returns is different on overreaction days compared to normal days.
- There is a momentum effect on the days on which a sharp change in the cryptocurrency prices (i.e., an overreaction day) is observed.
- There is a momentum effect the day after an overreaction day.

The above-mentioned findings have provided the basis for new, profitable cryptocurrency trading strategies based on the momentum indicators. For example, Caporale and Plastun (2020) propose to open a new trading position (i.e., buy or sell) on a cryptocurrency asset when

* Corresponding author.

E-mail addresses: gianpietro.bellocca@studenti.polito.it (G.P. Bellocca), giuseppe.attanasio@polito.it (G. Attanasio), luca.cagliero@polito.it (L. Cagliero), jacopo.fior@polito.it (J. Fior).

https://doi.org/10.1016/j.mlwa.2022.100310

Received 14 December 2021; Received in revised form 4 April 2022; Accepted 22 April 2022 Available online 27 April 2022 2666-8270/© 2022 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY-NC-ND license (http://creativecommons.org/licenses/by-nc-nd/4.0/). the momentum level exceeds a given user-specified cutoff threshold. However, the equities produced by the aforesaid trading strategy tend to have excessive volatility due to the inherent difficulties in detecting the overreaction days in which the strategy must operate.

To counteract the negative effects of excessively volatile financial markets, machine learning techniques have shown to be more robust than traditional, rule-based trading systems under challenging market conditions (Bustos & Pomares-Quimbaya, 2020). The main reason is that, based on a deep data exploration, they are able to tailor the trading strategies to the observed market conditions.

This paper proposes to advance existing cryptocurrency trading models based on the momentum effect by leveraging machine learning techniques. Unlike any previous approaches to momentum-based cryptocurrency trading it combines market forecasting with supervised momentum analysis. Specifically, for each cryptocurrency it estimates the likelihood of being affected by the momentum effect on the next trading day as well as the momentum direction. It opens a new trading position on a cryptocurrency asset only when the machine learning model outcome is in agreement with the observed momentum-based signal.

We run experiments on historical data relative to three of the most representative cryptocurrencies. The results achieved by the machine learning-based approach are superior to those of the heuristic approach in terms of both F1-score of the classification model and return of investment of the simulated trades. Therefore, machine learning seems to be particularly effective in counteracting the negative effects of cryptocurrency assets.

A summary of the main scientific contributions is given below.

- It presents a new machine learning-based strategy to cryptocurrency trading based on recent evidence on price momentum (Caporale & Plastun, 2020; Plastun et al., 2021). The purpose is to more effectively manage conditions of excessive price volatility, which are typical of cryptocurrency markets.
- 2. Unlike all existing approaches to cryptocurrency trading, it leverages machine learning to automatically detect the momentum effect in cryptocurrency market data and not only to perform market forecasting (Fang et al., 2020).
- 3. It achieves better performance than state-of-the-art approaches (e.g., Caporale and Plastun (2020)) on real cryptocurrency market data.

The rest of the paper is organized as follows. Section 2 discusses the prior work. Section 3 briefly reviews the fundamentals of the baseline approach presents by Caporale and Plastun (2020), whereas Section 4 presents the model extension based on machine learning. Finally, Sections 5 and 6 report the experimental results and draw the conclusions and future research directions, respectively.

2. Related works

2.1. Machine learning-based approaches to cryptocurrency trading

Although the use of machine learning to forecast the stock market is established (Bustos & Pomares-Quimbaya, 2020; Huang et al., 2019; Ozbayoglu et al., 2020; Rundo & di Stallo, 2019), only few attempts to adapt the current trading systems to the cryptocurrency market have been made. Specifically, Attanasio et al. (2019) have investigated the application of traditional classification techniques such as Support Vector Machines and Decision Trees to forecast the next-day cryptocurrency price, whereas in Lahmiri and Bekiros (2021), Livieris et al. (2020a) the authors have explored the use of Deep Learning techniques. A key aspect is the exploration of different data input types such macro-financial indicators and blockchain information. Moreover, social media data have been used to forecast cryptocurrencies prices such as Twitter data (Kraaijeveld & De Smedt, 2020) or GitHub and Reddit data (Glenski et al., 2019). Another line of works investigates the use of ensemble methods, exploring the use of both standard shallow approaches, like Random Forests and Stochastic Gradient Boosting Machine (Derbentsev et al., 2021), and the use of deep learning models as component learners (Livieris et al., 2020b). Sun et al. (2020) combine daily price data of 42 cryptocurrencies with key economic indicators acquired from the stock and Fiat markets to train a Gradient Boosting Decision Tree algorithm. However, all the aforesaid features are typically characterized by a high degree of noise thus making the inference process complex and not easily explainable. To overcome this issue, Zhang et al. (2021) propose to use an Attentive Memory module that combines a Gated Recurrent Unit with a self-attention component to establish attentive memory for each input sequence. The results confirm that the raw sequence already incorporates most of the relevant information whereas context information is hardly usable to build accurate predictive models. Rather than addressing shortterm price forecasting as in Lahmiri and Bekiros (2021), Livieris et al. (2020a), Zhang et al. (2021), this paper proposes a new machine learning-based strategy aimed at predicting the overreaction effect. Hence, it incorporates the underlying market properties highlighted by recent empirical evidence (Caporale & Plastun, 2020) to leverage the predictive power of machine learning models on historical prices. To the best of our knowledge, this is the first attempt to use machine learning to perform momentum-based cryptocurrency trading.

2.2. Applications of the momentum effect in the financial domain

The cryptocurrency market is a particularly new and relatively unexplored case of market extremely vulnerable to overreactions, given its high volatility compared to the traditional markets such as Forex, commodity and stock etc. Recent studies (Bartos, 2015; Urquhart, 2016) have analyzed its efficiency, long-memory properties and persistence in price (Bariviera, 2017), the existence of price bubbles (Corbet et al., 2018), its competitiveness (Gandal & Halaburda, 2014), the issue of price predictability (Plastun et al., 2021), and the presence of anomalies (Kurihara & Fukushima, 2017).

The influence of the momentum effect on the cryptocurrency market is due to its affinity with emerging markets: low regulation, trading barriers, lack of information and asset complexity. Indeed, since large financial institutions do not have the authorization or interest to operate on the cryptocurrency market, small investors often have to access it directly, by creating a wallet and using an exchange. Although these have developed a lot in recent years, becoming more user-friendly, they still remain a strong obstacle to accessing the market. A second huge problem is given by the type of asset being traded, the cryptocurrency, which is much more complex to understand than traditional assets, such as stocks, commodities and Forex. Therefore, investors are generally small players or individuals, attracted by the high volatility of the cryptocurrency market and by the possibility of speculation, which make decisions often driven more by common sentiment, than by reasoning based on pricing analysis.

A small number of studies have focused their attention on momentum and overreactions in the cryptocurrency market. For instance, the research work by Chevapatrakul and Mascia (2019), based on the quantile autoregressive model, reveals that days with deeply negative returns are often followed by periods again characterized by negative returns and that abnormal positive weekly returns will be followed by an increase in prices. More specifically, investors seem to overreact when daily returns are in the lower quantile of the distribution and when weekly returns are in the upper quantile of the distribution. A possible motivation is that investors are quick to exit the market on days of negative feelings when prices drop. Otherwise for the second finding, the results indicate an excessive reaction of investors to favorable news, during the weeks of positive sentiment when prices are rising. At the monthly frequency, no evidence of momentum was found, thus suggesting that the cryptocurrency market has much faster momentum dynamics compared to traditional asset markets, such as the stocks market. Caporale and Plastun (2020) have shown the presence of price patterns after overreactions. Also, they have analyzed the momentum effect in the cryptocurrency market during the overreaction day and the following day. They have observed that, after an overreaction, the price movements are higher than in normal days and, for these reasons, they claim that a trading strategy based on the momentum effect after overreactions is likely to be profitable. The empirical results show that hourly returns during day of positive (negative) overreactions are significantly higher (lower) than the hourly returns achieved during the average positive (negative) day. Moreover, anomalous days can be recognized before the end of the day. In fact, the price trend is likely to follow the direction of the overreaction until the end of the day. As a drawback, applying a posteriori reaction entails yielding potential losses due to the generation of false trading signals. To mitigate such a negative effect, the objective of this work is to develop a machine learning-based system that is able to identify both positive and negative overreaction conditions in cryptocurrency price series.

3. The heuristic approach

The overreaction detection method recently proposed by Caporale and Plastun (2020) is an heuristic method based on hourly inspection of the cryptocurrency price series. The key behind it is to trigger trading operations (buy or sell) only when the current price exceeds predefined thresholds.

The thresholds adopted in the heuristic method are computed using the daily average return of the cryptocurrency and its standard deviation. Daily returns (R_i) are computed as follows:

 $R_i = (Close_i / Open_i - 1) * 100\%$

where:

- R_i: returns on the *i*th day in %
- Open_i: open price on the *i*th day
- Close_i: close price on the *i*th day

The returns calculated are divided into two data sets composed by only positive or negative returns in order to respectively define positive and negative thresholds separately for each trading day.

The *i*th day is characterized by a positive overreaction day if:

$$R_i > (R_n + k * \sigma_n)$$

The *i*th day is characterized by a negative overreaction day if:

$$R_i < (R_n - k * \sigma_n)$$

Where:

- R_n: average daily returns for period n
- σ_n : standard deviation for period n
- k: number of standard deviations

Average daily return and standard deviation of period n (days) are computed as follows:

$$R_n = \sum_{i=1}^n \frac{R_i}{n}$$
$$\sigma_n = \sqrt{\frac{1}{n} \sum_{i=1}^n (R_i - R_n)^2}$$

The heuristic method compares the current price with the thresholds in order to assign the current trading day to one of the following categories: positive overreaction, negative overreaction or normal day. In order to do so, it monitors the hourly price series to detect if the asset price is above any of the two thresholds. If the price has exceeded a given threshold then the current day is labeled with the overreaction label. A prediction is correct if the daily close price is still beyond the threshold level. The trading system proposed by Caporale and Plastun (2020) leverages the momentum effect to open/close trading positions. Specifically, it opens a long-selling (short-selling) position when a positive (negative) overreaction is detected. Every trading position is closed at the end of the day (i.e., intraday trading scenario with no stop loss).

4. The proposed method

We present a new machine learning-based approach to counteracting the main drawbacks of the heuristic method, i.e., the generation of false trading signals.

4.1. Problem statement

Let O_c be an indicator function defined on trading day d_i for cryptocurrency c as follows:

$$O_c(d_i) = \begin{cases} 1 \ d_i \text{ positive overreaction} \\ -1 \ d_i \text{ negative overreaction} \\ 0 \text{ normal day} \end{cases}$$

We model the relation between the presence of a overreaction condition on the next trading day d_{i+1} and the set of historical feature values S_c describing c on the current trading day d_i and the preceding ones d_{i-1} , d_{i-2} , ..., d_{i-W+1} as an arbitrary classification function f_c :

$$O_c(d_{i+1}) = f_c(S_c(d_i), S_c(d_{i-1}), \dots, S_c(d_{i-W+1}))$$

where $f_c(\cdot)$ is the prediction function we want to find, W is the size of historical time window considered by the classification model, and $O_c(d_{i+1})$ is the value of the target variable.

4.2. Methodology

The main steps of the proposed methodology are summarized below.

- 1. Data acquisition and preparation. This step consist in the gathering of historical data related to one or more cryptocurrencies and performing feature engineering based on established technical analysis indicators (Murphy, 1999) (see Section 4.2.1).
- 2. Dataset labeling. In this second step the dataset samples, each one corresponding to a description of the cryptocurrency price series on a specific trading day, are labeled using the indicator function O_c (see Section 4.2.2).
- 3. **Classification**. This step entails training and applying a classification model to predict the presence and direction of an overreaction condition on the next trading day (see Section 4.2.3).

4.2.1. Data acquisition and preparation

We first retrieve a sufficiently large amount of historical prices relative to the cryptocurrencies under consideration. In our experiments, we retrieved data related to three renowned cryptocurrencies, namely Bitcoin (BTC), Ethereum (ETH) and Litecoin (LTC). Data are downloaded from Crypto Data Download,¹ which is a service gathering cryptocurrency data from the main crypto exchanges. The collected data are related to the Kraken² exchange. Each of the three datasets contains price and volume data at daily granularity. More specifically, each sample is defined by:

- a timestamp;
- open, high, low and close prices;
- · traded volume (in crypto and USD amounts).

www.cryptodatadownload.com (latest access: December 2021).
www.kraken.com.

.....

Feature	Description	Category
Open	Open price of the current day	Candlestick
High	Highest price of the current day	
Low	Lowest price of the current day	
Close	Close price of the current day	
Volume	Trading volume of the current day	
SMA5-20	Relative difference between SMA(5) and SMA(20)	Trend
SMA8-15	Relative difference between SMA(8) and SMA(15)	
SMA20-50	Relative difference between SMA(20) and SMA(50)	
EMA5-20	Relative difference between EMA(5) and EMA(20)	
EMA8-15	Relative difference between EMA(8) and EMA(15)	
EMA20-50	Relative difference between EMA(20) and EMA(50)	
MACD	Moving Average Convergence/Divergence	
AO14	Aroon Oscillator (14 periods)	
ADX14	Average Directional Index (14 periods)	
WD14	Difference between Positive Directional Index (DI+)	
	and Negative Directional Index (DI-) (14 periods)	
PPO12-26	Percentage Price Oscillator (12 and 26 periods)	Volatility
RSI14	Relative Strength Index (14 periods)	
MFI14	Money Flow Index (14 periods)	
TSI	True Strength Index	
SO14	Stochastic Oscillator (14 periods)	
CMO14	Chande Momentum Oscillator (14 periods)	
ATRP14	Average True Range Percentage: ratio between	
	Average True Range and Close (14 periods)	
PVO12-26	Percentage Volume Oscillator (14 and 26 periods)	Volume
ADL	Accumulation Distribution Line	
OBV	On Balance Volume	
FI13	Force Index (13 periods)	
FI50	Force Index (50 periods)	

In compliance with Attanasio et al. (2020), we describe the daily price variations with a set of established technical indicators (Murphy, 1999). The purpose is to summarize the status and trends of the cryptocurrency. Specifically, they signal whether the cryptocurrency price is following a trend, assess the trend strength, and detect a momentum trend is coming to an end due to underlying conditions due to either over-bought or over-sold conditions of the underlying asset. In compliance with Attanasio et al. (2020), in our experiments, we will hereafter consider the set of 22 technical indicators and oscillators reported in Table 1.

4.2.2. Dataset labeling

Since the goal is to exploit the momentum effect produced by overreaction, we label each dataset sample using the indicator function defined in Section 4.1. Analogously to Caporale and Plastun (2020), we exploit the overreaction to selectively open trading positions on the analyzed cryptocurrency.

To determine k and W we follow the indications provided by Caporale and Plastun (2020) and tested K values in the range [0, 2] and W values in the range [50, 360]. Notice that by setting k to 0 the heuristic method generates a high number of wrong trading signals, whereas by setting k to 2 the number of overreaction days significantly decreases. We achieve the best performance by setting k to 1. Finally, the size of the historical window cannot be lower than 50 to allow the computation of the technical indicators. Setting larger W values implies considering also long-term price trends, which turn out to be harmful for daily cryptocurrency trading.

4.2.3. Classification

The machine learning module is in charge of accomplishing a 3class classification task, where predicting either a positive or negative overreaction condition entails generating the corresponding trading signal analogously to the rule-based model described in Caporale and Plastun (2020).

5. Experiments

In this work, we consider the baseline approach proposed in Caporale and Plastun (2020) and several machine learning classifiers to produce trading decisions. We thoroughly test both classification and trading performance via extensive back-testing experiments³.

5.1. Experimental environment

We run experiments in a single-node setting on a HPC facility. The node runs Ubuntu 20.04.2 LTS, with a 8 CPU threads Intel(R) Xeon(R) Gold 6140 CPU @ 2.30 GHz and 40 GB of RAM. For the preliminary LSTM tests we use a NVIDIA V100 GPU with 16 GB of VRAM.

For each machine learning model the training time is in the order of seconds (in the worst cases) whereas the times required for running a grid search are in the order of tens of minutes.

5.2. Experimental design

Our back-testing experiments entail the following steps. First, we collect data available online for a given cryptocurrency. Data is collected at both daily and minutely time granularity. Next, we divide the data in training and back-test (evaluation) data. We then use the training portion to learn the relationship between feature descriptors and overreaction conditions using machine learning models. Finally, we back-test on the remaining days. Note that, similar to Caporale and Plastun (2020), the original heuristic detects overreaction conditions based on hourly data prices, whereas machine learning models predict the overreaction after being trained on daily data.

³ Data and code are available at https://anonymous.4open.science/r/momentum-crypto-trading-3023/.

Table 2

Model	Hyper-parameter	Grid values
RFC	criterion	"gini", "entropy"
	min_samples_split	0.01 , 0.05
	min_samples_leaf	0.005 , 0.01
	max_depth	None, 5, 10, 20
	class_weight	"balanced", "balanced_subsample"
KNN	weights	"uniform", "distance"
	n_neighbors	3, 5, 7
	algorithm	"ball_tree", "kd_tree"
MLP	hidden_layer_sizes	(10,), (30,), (10, 10), (512), (512, 256)
	activation	"relu", "logistic", "tanh"
	solver	"lbfgs", "sgd", "adam"
	learning_rate	"constant", "invscaling"
	learning_rate_init	2e-5, 1e-4, 1e-5, 1e-2, 1e-1
	tol	1e–4, 1e-5
SVC	kernel	"linear", " poly ", "rbf"
	degree	3, 4, 5
	C	0.001, 0.01, 1, 10, 50
MNB	alpha	0.01, 0.1, 1, 10
LSTM	n_layers	2 , 3, 4, 5
	bidirectional	True, False
	sequence length	3, 5, 7, 10
LG	solver	"newton-cg", "lbfgs", "liblinear", "sag",
		"saga"
	penalty	" 11 ", "l2"
	С	1e-4, 1e-3, 1e-2, 1e-1, 1, 10

5.3. Models

In our experiments, we leave the original heuristic approach (HE) from Caporale and Plastun (2020) unchanged while extending the baseline method by using the machine learning-based approach described in Section 4. As discussed in Section 4.1, we frame our problem as a three-class classification task with the goal of predicting positive or negative overreaction situations or neutral days. To address the task, we leverage renowned machine learning models (Tan et al., 2019). Specifically, we test Support Vector Machine (SVM), Gaussian Naive-Bayes (GNB), Multinomial Naive-Bayes (MNB), K-Nearest Neighbors (KNN), Logistic Regressor (LG), Random Forest (RFC), Feed-forward fully-connected Neural Network (MLP). For the traditional classifiers we used the implementations available in the SK-Learn library (Buitinck et al., 2013).

We also run a set preliminary tests using the Deep Learning-based LSTM architecture available in the PyTorch library (Paszke et al., 2019). Specifically, prior to running the empirical validation we verify the presence of data overfitting in the training phase. However, since cryptocurrencies are relatively new financial markets and we are analyzing daily cryptocurrency prices the available data collection appears to be not suited to train Deep Learning models.

5.4. Feature design for ML models

We enrich the cryptocurrencies series using established technical indicators. We use (i) simple and exponential moving averages, at different time scales, to capture price movements, (ii) Moving Average Convergence Divergence, Aroon Oscillator and similar to express trends, (iii) Relative Strength Index, True Strength Index, Stochastic Oscillator and other similar oscillators to capture volatility in price, and four different Volume indicators.

Table 1 reports the full list of features we used in our experiments. Please refer to Murphy (1999) for a more extensive discussion of the meaning and usefulness for predicting the future price movements.

5.5. Back-testing setup

We test our approach on three different cryptocurrencies, namely Bitcoin (BTC), Ethereum (ETH), and Litecoin (LTC). We experiment on the time span between September 2, 2015 and December 31, 2020. We use data from September 2, 2015 and December 31, 2019 as training data for machine learning classifiers. We back-test every model on the remaining year.

We optimize each model's hyper-parameters via exhaustive gridsearch. We validate each configuration using time-aware k-fold crossvalidation, setting the number of folds equal to 5⁴. We then use the best performing configuration in terms of F1-measure weighted by class frequency.

We test over-sampling for getting class balance in training sets using SMOTE (Chawla et al., 2002) and ADASYN (He et al., 2008). No classification model benefited from over-sampling.

Table 2 reports the remaining hyperparameters tested during validation. Concerning the heuristic method, we run manual grid-search over two core parameters: the number of past days and standard deviations used to compute the high and low thresholds. We test the number of past days in [365, 50] and the number of standard deviations k in [0, 1, 2]. We achieve best results with 50 past days as threshold definition period and 2 standard deviations.

5.6. Classification results

Tables 3–5 report the classification performance in detecting oneday ahead overreaction conditions. We report the scores in terms of F1 measure weighted by class frequency, separately by number of standard deviations (k). The first row reports the baseline heuristic approach by Caporale and Plastun (2020). It is evident from the results that almost any machine learning method outperform the heuristic approach in all circumstances. The K-NN model is on average the best performer. However, there is no a clear written over all the three cryptocurrencies.

⁴ We use the scikit-learn Python implementation as found in *sklearn.model_selection.TimeSeriesSplit.*

Table 3

BTC back-testing: F1 (weighted) scores for each classification model and number of classes. Best performers for each setup are reported in bold. Baseline results are italicized.

Model	# of labels	;
	3	2
HE	0.70	0.74
GNB	0.67	0.71
KNN	0.77	0.77
LG	0.71	0.68
MLP	0.70	0.74
MNB	0.77	0.77
RFC	0.75	0.75
SVC	0.73	0.75

We also evaluate the classification performance using a binary class (i.e., overreaction or normal day) instead of a 3-class (positive overreaction, negative overreaction, normal day). Since the 2-class classification problem is inherently simpler than the 3-class one, the F1-measure is expected to be higher. However, based on the binary classification outcome we are unable to distinguish between positive and negative overreaction. Hence, opening directional trading positions would require additional actions (see Section 3).

5.7. Trading results

In order to test further the applicability of the machine learningbased momentum detection methods, we carried out a set of trading simulations. Specifically, we performed an hourly inspection of the prices using the heuristic-based trading strategy. When the price reaches the positive (negative) threshold a corresponding long (short) position is opened. Conversely, the machine learning-based trading strategy does not adopt hourly price monitoring. This second strategy exploit the labels produced by the classifier in order to open a long (short) position if the label is positive (negative) at the beginning of each trading day. Both strategies close their positions at the end of each trading day.

Tables 6-8 report the scores achieved by heuristic-based strategy and by the two machine learning-based best performer strategies in terms of percentage of profitable trades, total return and average return per single trade. The results show that the machine learning-based strategies open less position than the heuristic strategy, with an higher percentage of profitable trade. These results confirm that the detection method based on machine learning is more effective in preventing excessive signal generation. Furthermore, the total return obtained by the machine learning strategies and the average return per trade are, on average, higher. This is due to the fact that machine learning allows to open a position at the beginning of each overreaction day, without waiting for a match with the threshold level. Hence, it enters the market from a better spot compared to the heuristic trading strategy. The KNN model seems to be the most promising. Even if it is not the top performer for all the considered cryptocurrencies, KNN is always in the best performing groups, with better results than HE. It also provides a consistent total return on BTC, which appears to be the asset over which the momentum effect strategies are not particularly effective.

6. Conclusions and future work

The paper investigated the use of machine learning techniques to overcome the limitations of state-of-the-art momentum-based cryptocurrency trading systems. Specifically, based on the empirical observation that the momentum effect is likely to influence the series of cryptocurrency prices, we designed a methodology that predicts the presence and direction of an overreaction condition. The takeaways of the research are summarized below:

Table 4

ETH back-testing: F1 (weighted) scores for each classification model and number of classes. Best performers for each setup are reported in bold. Baseline results are italicized.

Model	# of labels				
	3	2			
HE	0.68	0.72			
GNB	0.74	0.76			
KNN	0.69	0.70			
LG	0.74	0.76			
MLP	0.68	0.68			
MNB	0.74	0.74			
RFC	0.71	0.76			
SVC	0.73	0.76			

Table 5

LTC back-testing: F1 (weighted) scores for each classification model and number of classes. Best performers for each setup are reported in bold. Baseline results are italicized.

Model	# of labels				
	3	2			
HE	0.69	0.73			
GNB	0.73	0.77			
KNN	0.77	0.78			
LG	0.78	0.69			
MLP	0.71	0.76			
MNB	0.77	0.77			
RFC	0.77	0.78			
SVC	0.75	0.78			

Table 6

BTC trading results. Best performers are reported in boldface. Baseline results are italicized.

Model	Trades	Profitable trades	Profitable trades %	Total return	Return per trade
HE	40	19	47.50%	2.47%	0.06%
KNN	11	6	54.55%	8.63%	0.78%
SVC	48	26	54.17%	3.61%	0.08%

Table 7

ETH trading results. Best performers are reported in boldface. Baseline results are italicized.

Model	Trades	Profitable trades	Profitable trades %	Total return	Return per trade
HE	84	45	53.57%	59.49%	0.71%
KNN	52	33	63.46%	97.44%	1.87%
LG	30	27	90.00%	126.23%	4.21%

Table 8

LTC	trading	results.	Best	performers	are	reported	in	bold.	Baseline	results	are	italiciz	ed

Model	Trades	Profitable trades	Profitable trades %	Total return	Return per trade
HE	83	33	39.40%	14.46%	0.17%
KNN	24	12	50.00%	36.36%	1.52%
SVC	39	24	61.54%	68.13%	1.75%

- The return per trade of the machine learning-based approach is significantly better than those achieved by the heuristic approach (e.g., on BTC KNN 0.78% vs. Heuristic approach 0.06%, on ETH KNN 1.87% vs. Heuristic approach 0.71%).
- KNN is on average the most performing classifier on all the tested cryptocurrencies and for all the considered settings.

• The use of machine learning is beneficial for trading purposes especially when the number of overreaction days available in the historical data is significant (k = 1).

To balance the excessive reactivity of the heuristic approach and the relatively low speculative predisposition of the machine learning approach, as future work we plan to design an hybrid approach combining heuristic rules based on technical analysis with machine learning approaches.

CRediT authorship contribution statement

Gian Pietro Bellocca: Methodology, Software, Investigation, Writing – original draft. Giuseppe Attanasio: Conceptualization, Methodology, Software, Investigation, Writing – original draft, Supervision. Luca Cagliero: Conceptualization, Methodology, Writing – original draft, Writing – review & editing, Supervision. Jacopo Fior: Conceptualization, Methodology, Investigation, Writing – original draft, Writing – review & editing, Supervision.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

- Attanasio, G., Cagliero, L., & Baralis, E. (2020). Leveraging the explainability of associative classifiers to support quantitative stock trading. In Proceedings of the sixth international workshop on data science for macro-modeling (pp. 1–6).
- Attanasio, G., Cagliero, L., Garza, P., & Baralis, E. (2019). Quantitative cryptocurrency trading: exploring the use of machine learning techniques. In *Proceedings of the 5th* workshop on data science for macro-modeling with financial and economic datasets (pp. 1:1–1:6). ACM, http://dx.doi.org/10.1145/3336499.3338003.
- Bariviera, A. F. (2017). The inefficiency of bitcoin revisited: A dynamic approach. *Economics Letters*, 161, 1–4. http://dx.doi.org/10.1016/j.econlet.2017.09.013, URL: https://www.sciencedirect.com/science/article/pii/S0165176517303804.
- Bartos, J. (2015). Does bitcoin follow the hypothesis of efficient market? International Journal of Economic Sciences, IV, 10–23. http://dx.doi.org/10.20472/ES.2015.4.2. 002.
- Buitinck, L., Louppe, G., Blondel, M., Pedregosa, F., Mueller, A., Grisel, O., Niculae, V., Prettenhofer, P., Gramfort, A., Grobler, J., Layton, R., VanderPlas, J., Joly, A., Holt, B., & Varoquaux, G. (2013). Api design for machine learning software: experiences from the scikit-learn project. In *ECML PKDD workshop: languages for* data mining and machine learning (pp. 108–122).
- Bustos, O., & Pomares-Quimbaya, A. (2020). Stock market movement forecast: A systematic review. *Expert Systems with Applications*, 156, Article 113464. http: //dx.doi.org/10.1016/j.eswa.2020.113464.
- Caporale, G. M., & Plastun, A. (2020). Momentum effects in the cryptocurrency market after one-day abnormal returns. *Financial Markets and Portfolio Management*, 34, 251–266.
- Chawla, N. V., Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. (2002). Smote: synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 16, 321–357.
- Chevapatrakul, T., & Mascia, D. V. (2019). Detecting overreaction in the bitcoin market: A quantile autoregression approach. *Finance Research Letters*, 30, 371–377. http://dx.doi.org/10.1016/j.frl.2018.11.004.
- Corbet, S., Lucey, B., & Yarovaya, L. (2018). Datestamping the bitcoin and ethereum bubbles. *Finance Research Letters*, 26, 81–88, URL: https://EconPapers.repec.org/ RePEc:eee:finlet:v:26:y:2018:i:c:p:81-88.
- Derbentsev, V., Babenko, V., Khrustalev, K., Obruch, H., & Khrustalova, S. (2021). Comparative performance of machine learning ensemble algorithms for forecasting cryptocurrency prices. *International Journal of Engineering*, 34, 140–148.
- Fang, F., Ventre, C., Basios, M., Kong, H., Kanthan, L., Li, L., Martínez-Rego, D., & Wu, F. (2020). Cryptocurrency trading: A comprehensive survey. CoRR abs/2003.11352. URL: https://arxiv.org/abs/2003.11352.

- Gandal, N., & Halaburda, H. (2014). Competition in the cryptocurrency market: Working papers 14-17, NET Institute, URL: https://EconPapers.repec.org/RePEc:net:wpaper: 1417.
- Glenski, M., Weninger, T., & Volkova, S. (2019). Improved forecasting of cryptocurrency price using social signals. arXiv:1907.00558.
- He, H., Bai, Y., Garcia, E. A., & Li, S. (2008). Adasyn: Adaptive synthetic sampling approach for imbalanced learning. In 2008 IEEE international joint conference on neural networks (IEEE world congress on computational intelligence) (pp. 1322–1328). IEEE.
- Huang, B., Huan, Y., Xu, L. D., Zheng, L., & Zou, Z. (2019). Automated trading systems statistical and machine learning methods and hardware implementation: A survey. *Enterprise Information Systems*, 13, 132–144. http://dx.doi.org/10.1080/17517575. 2018.1493145.
- King, T., & Koutmos, D. (2021). Herding and feedback trading in cryptocurrency markets. Annals of Operations Research, 300, 79–96. http://dx.doi.org/10.1007/ s10479-020-03874-4.
- Kraaijeveld, O., & De Smedt, J. (2020). The predictive power of public twitter sentiment for forecasting cryptocurrency prices. *Journal of International Financial Markets, Institutions and Money*, 65, Article 101188. http://dx.doi.org/10. 1016/j.intfin.2020.101188, URL: https://www.sciencedirect.com/science/article/ pii/S104244312030072X.
- Kurihara, Y., & Fukushima, A. (2017). The market efficiency of bitcoin: A weekly anomaly perspective. Journal of Applied Finance & Banking, 7, URL: https:// EconPapers.repec.org/RePEc:spt:apfiba:v:7:y:2017:i:3:f:7_3_4.
- Lahmiri, S., & Bekiros, S. D. (2021). Deep learning forecasting in cryptocurrency highfrequency trading. *Cognitive Computation*, 13, 485–487. http://dx.doi.org/10.1007/ s12559-021-09841-w.
- Livieris, I. E., Pintelas, E. G., Stavroyiannis, S., & Pintelas, P. E. (2020a). Ensemble deep learning models for forecasting cryptocurrency time-series. *Algorithms*, 13, 121. http://dx.doi.org/10.3390/a13050121.
- Livieris, I. E., Pintelas, E., Stavroyiannis, S., & Pintelas, P. (2020b). Ensemble deep learning models for forecasting cryptocurrency time-series. *Algorithms*, 13, http: //dx.doi.org/10.3390/a13050121, URL: https://www.mdpi.com/1999-4893/13/5/ 121.
- Murphy, J. (1999). New York institute of finance series, Technical analysis of the financial markets: a comprehensive guide to trading methods and applications. New York Institute of Finance.
- Ozbayoglu, A. M., Gudelek, M. U., & Sezer, O. B. (2020). Deep learning for financial applications : A survey. *Applied Soft Computing*, 93, Article 106384. http://dx.doi. org/10.1016/j.asoc.2020.106384.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Chintala, S. (2019). Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d' Alché, E. Fox, & R. Garnett (Eds.), Advances in neural information processing systems, vol. 32 (pp. 8024–8035). Curran Associates, Inc..
- Plastun, A., Bouri, E., Gupta, R., & Ji, Q. (2021). Price effects after one-day abnormal returns in developed and emerging markets: Esg versus traditional indices. *The North American Journal of Economics and Finance*, 59, Article 101572. http://dx. doi.org/10.1016/j.najef.2021.101572.
- Rundo, Trenta, & di Stallo, Battiato (2019). Machine learning for quantitative finance applications: A survey. *Applied Sciences*, 9, 5574. http://dx.doi.org/10.3390/ app9245574.
- Sinkla, W. (2016). On the generation of arbitrage-free stock price models using lie symmetry analysis. Computers & Mathematics with Applications, 72, 1386–1393. http://dx.doi.org/10.1016/j.camwa.2016.07.003.
- Sun, X., Liu, M., & Sima, Z. (2020). A novel cryptocurrency price trend forecasting model based on lightgbm. *Finance Research Letters*, 32, Article 101084. http://dx.doi.org/10.1016/j.frl.2018.12.032, URL: https://www.sciencedirect.com/ science/article/pii/S1544612318307918.
- Tan, P., Steinbach, M. S., Karpatne, A., & Kumar, V. (2019). Introduction to data mining (2nd ed.). Pearson, URL: https://www-users.cse.umn.edu/%7Ekumar001/dmbook/ index.php.
- Urquhart, A. (2016). The inefficiency of bitcoin. *Economics Letters*, 148, http://dx.doi. org/10.1016/j.econlet.2016.09.019.
- Zhang, Z., Dai, H.-N., Zhou, J., Mondal, S. K., Garcia, M. M., & Wang, H. (2021). Forecasting cryptocurrency price using convolutional neural networks with weighted and attentive memory channels. *Expert Systems with Applications*, 183, Article 115378. http://dx.doi.org/10.1016/j.eswa.2021.115378, URL: https://www. sciencedirect.com/science/article/pii/S0957417421008046.