



UNIVERSITY  
OF WARSAW



FACULTY OF  
ECONOMIC SCIENCES

## WORKING PAPERS

No. 27/2024 (463)

# INFORMER IN ALGORITHMIC INVESTMENT STRATEGIES ON HIGH FREQUENCY BITCOIN DATA

FILIP STEFANIUK  
ROBERT ŚLEPACZUK

WARSAW 2024

ISSN 2957-0506



## Informer in Algorithmic Investment Strategies on High Frequency Bitcoin Data

Filip Stefaniuk<sup>1</sup>, Robert Ślepaczuk<sup>2\*</sup>

<sup>1</sup> University of Warsaw, Faculty of Economic Sciences

<sup>2</sup> University of Warsaw, Faculty of Economic Sciences, Department of Quantitative Finance and Machine Learning, Quantitative Finance Research Group

\* Corresponding author: [rslepaczuk@wne.uw.edu.pl](mailto:rslepaczuk@wne.uw.edu.pl)

---

**Abstract:** The thesis investigates the usage of Informer architecture for building automated trading strategies for high frequency Bitcoin data. Two strategies using Informer models with different loss functions, Quantile loss and Generalized Mean Absolute Directional Loss (GMADL), are proposed and evaluated against the Buy and Hold benchmark and two benchmark strategies based on technical indicators. The evaluation is conducted using data of various frequencies: 5 minute, 15 minute, and 30 minute intervals, over the 6 different periods. Although the Informer-based model with Quantile loss did not manage to outperform the benchmark, the model that uses novel GMADL loss function turned out to be benefiting from higher frequency data and beat all the other strategies on most of the testing periods. The primary contribution of this study is the application and assessment of the Quantile and GMADL loss functions with the Informer model to forecast future returns, subsequently using these forecasts to develop automated trading strategies. The research provides evidence that employing an Informer model trained with the GMADL loss function can result in superior trading outcomes compared to the buy-and-hold approach.

---

**Keywords:** Machine Learning, Financial Series Forecasting, Automated Trading Strategy, Informer, Transformer, Bitcoin, High Frequency Trading, Statistics, GMADL

**JEL codes:** C4, C14, C45, C53, C58, G13

**Note:** This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors

# 1 Introduction

Developing automated trading strategies is a challenging task that has captivated institutional investors and individual researchers for a long time. With the rise of computing power and machine learning (ML) development, more and more automated algorithms are being developed and deployed on the market. It is estimated that currently 70% to 80% of all market transactions are carried out by automated trading software (Yadav 2015) and this number is expected to increase in the coming years.

Although numerous systems built in the past have handled daily market data, this is insufficient in the age of High-Frequency Trading and advanced computing power. This is especially true in the context of trading assets such as Bitcoin. Bitcoin (BTC) is a decentralized digital currency that allows peer-to-peer transactions over the Bitcoin network, which is a secure, public ledger using blockchain technology. It was first introduced in an anonymous 2009 white paper (Nakamoto 2009), by a person or group of people using the pseudonym Satoshi Nakamoto. Since then, Bitcoin has grown in popularity and many cryptocurrency exchanges that facilitate bitcoin sales and purchases have emerged. This asset is known for its high volatility, for example, in 2017 the Bitcoin price went up nearly 20 times, reaching \$19,497 and then dropping by 84%, in 2021 Bitcoin reached a new all-time high of \$63,314 but then the price fell to \$34,770, almost 55%<sup>1</sup>. Since then, Bitcoin has become somewhat more stable, especially after approval of Bitcoin Spot ETFs, although price movements are still greatly exaggerated compared to more classical investment instruments.

The study explores the idea of building an automated trading strategy for Bitcoin. Four strategies are proposed and evaluated on the historical Bitcoin data of high frequencies: 5 minutes, 15 minutes, and 30 minutes; from 21.08.2019 to 24.07.2024. The first two strategies are treated as benchmarks and are based on classical technical indicators, namely the Moving Average Convergence Divergence (MACD) and Relative Strength Index (RSI). The other two employ the Informer (Zhou et al. 2021), a state-of-the-art attention-based neural network model designed to efficiently handle long time series, to predict the returns and subsequently choose positions according to the model's forecasts.

The work aims to answer the following research questions:

- Q:** *Is it possible to create an algorithmic strategy for trading Bitcoin, that is more efficient than the Buy&Hold approach?*
- Q:** *Does signal from Informer model allow to create strategies that are more efficient on trading Bitcoin than strategies based on technical indicators?*
- Q:** *How does selection of the machine learning model loss function influence the strategy performance?*
- Q:** *Does usage of higher frequency data allow to create more efficient strategies?*

Answering these questions required designing and implementing a method to compare various trading strategies. This comparison is the main contribution of this study. To the best of current knowledge, no other research has yet been performed where an Informer model is trained with the Quantile or GMADL (Michańków et al. 2024) loss function, followed by the utilization of its forecasts in buy/sell signals generation to develop automated trading strategies. The study involves an exhaustive analysis of the approach including the comparison to the benchmark strategies: Buy&Hold and two technical indicator-based strategies, usage of data with different time intervals, and multiple time periods. Finally, a sensitivity analysis is conducted to show how changing the parameters affect the performance of the tested strategies. An additional contribution of the research is an open-sourced implementation of a framework for efficient comparison of trading strategies, that is available on Gitlab and enables reproducing the results<sup>2</sup>.

The structure of the thesis is as follows: Chapter 2 surveys the associated literature and gives a brief overview of prior similar studies. Chapter 3 explains the acquisition, preprocessing, and covers the analysis of the datasets used in the study. In Chapter 4, the notion of trading strategy is formally defined, the metrics for comparison are introduced, and the details of each examined strategy are outlined. Chapter 5 displays the experimental results, detailing the selection of strategy hyperparameters and presenting the outcomes of the strategy evaluation along with sensitivity analysis. Lastly, Chapter 7 draws the conclusions of the study.

<sup>1</sup><https://blog.obiex.finance/6-biggest-bitcoin-crashes-that-have-happened-in-crypto-history/>

<sup>2</sup><https://gitlab.com/FilipStefaniuk/wne-msc-thesis>

## 2 Literature Review

The study of financial market efficiency has a rich history. For years, researchers have been trying to develop algorithms capable of predicting the price movements of financial assets as well as creating automated trading strategies that rely on such signals. The concept of developing an automated system capable of regularly outperforming the market is so attractive that it leads to the publication of thousands of papers on that topic annually.

### 2.1 Efficient Market Hypothesis

These endeavors stand in contrast to the Efficient Market Hypothesis (EMH), also known as the efficient market theory, which argues that stock prices already reflect all accessible information. EMH suggests that it is impossible to consistently achieve higher returns than average market returns, thus the best investing strategy is a passive portfolio. The hypothesis was formulated in three forms: *weak*, *semi-strong*, and *strong*. The *weak* form challenges the foundation of technical analysis by suggesting that stock prices move in a random manner. The *semi-strong* form argues that stock prices immediately incorporate all public information, arguing for the ineffectiveness of the fundamental analysis. The *strong* form extends this argument, implying that the price of the stock includes both public and private information. The EMH, even in its weakest form, implies that the unpredictability of short-term price fluctuations makes efforts to create automated trading systems ineffective (Malkiel 1973).

The EMH was acknowledged in the early review of theoretical and empirical market models by Fama (1970). The author put forth certain evidence indicating the presence of a statistically significant positive correlation in daily returns, suggesting that such patterns might be exploited to develop profitable trading strategies. However, they ultimately concluded that this evidence might not be robust enough to challenge the validity of the EMH. This research was continued in Fama (1991), which was published two decades later. In this publication, the author presented additional evidence for predicting daily and weekly returns based on historical returns, and the findings align with those of the earlier study.

The argument in favor of EMH was presented in Malkiel (2005). The publication demonstrated that over a long period, professional investment managers do not outperform their index benchmarks. The study showed that in a 20-year period 90% of managed index funds was outperformed by the S&P 500 index, asserting that the only valid long-term strategy is the passive portfolio.

Conversely, the most persistent criticisms of the EMH relate to the preferences and actions of market participants. Psychologists and experimental economists recorded numerous specific behavioral biases that are prevalent in human decision-making when faced with uncertainty. They argue that irrational behaviors such as overconfidence (Fischhoff and Slovic 1978; Barber and Odean 2001; Gervais and Odean 2001), overreaction (Bondt and Thaler 1985), loss aversion (Kahneman and Tversky 1979; Odean. 1996), herding (Huberman and Regev 2001), and psychological accounting (Tversky and Kahneman 1981) are predictable and can be utilized to develop profitable trading strategies that exceed passive benchmarks.

Despite extensive research, the validity of the EMH continues to be debated without a definitive conclusion. More modern approaches explore the idea of time-varying weak-form market efficiency (Lim and Brooks 2011) and the Adaptive Markets Hypothesis (Kim et al. 2011) arguing that the predictability of markets can change over time and depend on various conditions such as the composition of market participants.

### 2.2 Early Statistical Approach

Nevertheless, many statistical systems and techniques that attempt to analyze and forecast asset prices have been developed. The Box-Jenkins method applies autoregressive moving average (ARMA) or autoregressive integrated moving average (ARIMA) models to find the best fit of a time series model to past values of a time series (Box and Jenkins 1976). The model combines autoregression, differencing, and moving averages to predict stock prices, and despite its simplicity, it has been widely applied and successful in financial forecasting. Another method called the exponential smoothing model (ESM) employs the exponential window function to smooth the data in which exponential smoothing parameters are estimated (Billah et al. 2006). Furthermore, those approaches were often combined with several other classical Machine Learning non-linear models.

In the study by Zhang (2003), authors combined the ARIMA model with Artificial Neural Networks (ANNs) to predict the prices of pairs of currencies. Similar approach was further explored in Khashei and Bijari (2011), in which authors use the ARIMA-ANN hybrid model to further improve predictions on the aforementioned currencies datasets. In Pai and Lin (2005), the authors conducted a comparison between

the ARIMA model and Support Vector Machines (SVMs) across diverse stock prices, demonstrating that both methodologies successfully forecasted the prices. In Wang et al. (2012) the authors developed a hybrid approach that uses ESM, ARIMA, and ANN to predict daily prices of the Dow Jones Industrial Average and Shenzhen index. Adebisi et al. (2014) utilized the ARIMA model in data sets from the New York and Nigeria Stock Exchanges, illustrating its effectiveness in forecasting short-term prices, and comparing its results with ANN. Azari (2018) illustrated the effectiveness of the ARIMA model in forecasting Bitcoin's price by examining its predictions on a time series over three years. However, they admitted that the approach was unable to predict abrupt price fluctuations, such as the volatility observed in late 2017. Vö and Ślepaczuk (2022) utilized a combination of ARIMA and GARCH family models to predict S&P500 log returns, employing these predictions to develop a trading strategy. Their findings indicated that strategies incorporating hybrid models outperformed those relying only on the ARIMA model.

## 2.3 Early Machine Learning Systems

Advancements in computing power lead to wider adoption of ML models and the development of more complex neural network architectures. Models such as Convolutional Neural Networks (CNNs) (LeCun et al. 2015) or Recurrent Neural Networks (RNNs) (and their variants Long-Short Term Memory Networks (LSTMs) (Hochreiter and Schmidhuber 1997) and Gated Recurrent Units (GRUs) (Cho et al. 2014)) were developed. Initially, they were used mainly in the computer vision and natural language processing domains, but they were quickly adapted and successfully implemented for financial forecasting.

Chen et al. (2015) modeled and predicted China stock market daily returns with LSTM, showing it outperforms previous classic statistical approaches. Honchar and Di Persio (2016) compared how different RNN architectures are able to forecast price movements. They used Google stock price index data and applied multi-layer RNN, LSTM, and GRU, to predict prices at different time horizons. The study showed that LSTM outperformed other architectures for long sequences. In the other publication Persio et al. (2017), the same authors compared Multi Layer Perceptron (MLP), CNN with LSTM architectures for predicting daily S&P500 returns and minute-by-minute FOREX EUR/USD returns. They reported that CNN can model financial time series better than other architectures. Fischer et al. (2017) utilized an LSTM model for selecting S&P 500 stocks in portfolio construction. The data came from 1992 until 2015, and the goal was formulated as a classification problem. The researches reported that the LSTM outperformed classical ML methods like random forest and logistic regression classifiers. Moreover, they attempted to disentangle the black box of LSTMs predictions by analyzing the selected stocks. They discovered that the selected stocks consist mainly of stocks with below-mean momentum, strong short-term reversal characteristics, and high volatility. Grudniewicz and Ślepaczuk (2023) applied several classical Machine Learning algorithms to technical analysis indicators for the WIG20, DAX, S&P 500, and a few selected CEE indices. The findings of the study demonstrate that quantitative methods outperform passive strategies when considering risk-adjusted returns, with the Bayesian Generalized Linear Model and Naive Bayes emerging as the leading models for the examined indices.

Selvin et al. (2017) predicted prices for 2000 NSE-listed companies with LSTM, RNN, and CNN models. They identified CNN as the best model that turned out to be best in predicting the changes in trends. Nelson et al. (2017) used LSTM in combination with technical analysis indicators. They reported an average accuracy of 55.9% when predicting if the price of a particular stock is going to rise or not in the near future. Hossain et al. (2018) proposed model integrating LSTM and GRU architectures to predict S&P 500 prices. In this method, the input data is initially fed into the LSTM network to produce a preliminary prediction, after which the output from the LSTM layer is forwarded to the GRU layer for generating the final prediction. The research demonstrates that the suggested hybrid model surpasses the performance of standalone LSTM and GRUs. Siami-Namini et al. (2018) examined how deep learning-based time series forecasting algorithms such as LSTM compare to traditional-based algorithms such as the ARIMA model. The average reduction in error rates obtained by LSTM was found to be between 84 and 87 percent compared to ARIMA, indicating the superiority of LSTM over ARIMA. Similarly Kijewski et al. (2024) evaluated the effectiveness of investment strategies based on traditional models and an LSTM on the S&P 500 index time series, spanning 20 years of data from 2000 to 2020. The combination of signals from several methods doubled the returns on the same level of risk as the passive strategy benchmark. The study, however, concluded that the LSTM model exhibited a considerably lower robustness to parameter variations compared to traditional methods.

## 2.4 Modern Machine Learning Systems

Although RNNs became widely adopted and proved to be effective for sequential data, the architecture struggled with processing longer sequences due to exploding and vanishing gradient phenomena (Basodi et al. 2020). In order to overcome this limitation, the Attention mechanism (Bahdanau et al. 2014; Luong et al. 2015) was developed.

While initially it was used in combination with the RNNs, Vaswani et al. (2017) introduced a neural network architecture based solely on the Attention mechanism called Transformer. The newly introduced architecture exhibited a significant enhancement in the performance, leading to its widespread adoption across multiple domains. Nowadays, it lies at the heart of state-of-the-art large-scale, multimodal models such as GPT-4 (OpenAI et al. 2024) or Gemini (Team et al. 2024).

The time series forecasting domain also adopted the new architecture. The Benidis et al. (2023) collects general-purpose state-of-the-art time-series forecasting architectures developed in recent years. Among the first of the modern deep forecasting models is DeepAR (Salinas et al. 2019) which still is primarily based on RNNs. The model outputs parameters of a previously chosen family of distributions. Samples from this distribution can be fed back into the model during prediction. However, most of the other modern models fully adapt to the new Transformer architecture. Temporal Fusion Transformer (TFT) developed by Lim, Arik, et al. (2020), incorporates Transformer architecture with novel components for embedding static covariates, performing “variable selection”, and gating components that skip over irrelevant parts of the context. The TFT is trained to predict forecast quantiles and promotes forecast interpretability by modifying self-attention and learning input variable importance. Eisenach et al. (2020) proposed MQ-Transformer, a Transformer architecture that employs attention mechanisms in the encoder and decoder separately, and considers learning positional embeddings from event indicators. The authors discussed improvements not only on forecast accuracy but also on excess forecast volatility where their model improves over the state-of-the-art. Zhou et al. (2021) proposed the Informer, a computationally efficient Transformer architecture, that specifically targets applications with long forecast horizons. These robust novel architectures can be specifically utilized for forecasting financial time series, particularly when dealing with higher-frequency data.

Barez et al. (2023) explored the application of deep learning Transformers architectures for high-frequency Bitcoin-USDT log-return forecasting and compared them to the LSTM model. The results indicate that the model based on the Transformer achieves a higher cumulative PnL than the LSTM when trading with multiple signals during backtesting. Hu and Hu (2021) applied TFT for stock price predictions comparing it with LSTM and again showing the superiority of the Transformer model. Zhao et al. (2022) investigated the prediction capability of the Transformer model on Bitcoin and Ethereum price data, furthermore augmenting the model with sentiment data collected from Twitter. Similarly Hájek and Novotny (2024) integrated news sentiment and investor attention metrics with a Temporal Fusion Transformer framework. The effectiveness of the model was demonstrated by analyzing stock price data for 41 of the largest market capitalization companies over the period from 2010 to 2021.

Penmetsa and Vemula (2023) analyzed the effectiveness of deep learning to predict Bitcoin, Ethereum, and Litecoin prices, by integrating momentum and volatility technical indicators into the network's input. They conducted a study using LSTM and Transformer architectures, concluding that Transformers tend to outperform LSTM models in price prediction and trends in cryptocurrency data. Wang (2023) introduced a hybrid approach combining Transformer architecture with bi-LSTM. The authors argued that Transformer is better at obtaining long-distance information, while LSTM can better capture short-distance signals. They presented the effectiveness of their method using 5 index stocks and 14 Shanghai and Shenzhen stocks.

The Informer architecture was used by Ding et al. (2023) to predict intra-day Chinese stock price. They also did a comparison with the LSTM model showing that Informer outperformed the RNN network. Another such comparison was done by Duan and Ke (2024) who systematically compared the effectiveness of the Informer and LSTM models for price predictions concluding that while in general, Informer outperforms LSTM, the fusion of the two combines the advantages of both models and enhances the prediction accuracy. Lastly Lu et al. (2023) designed three experiments to compare Informer with the commonly used networks of LSTM, Transformer, and BERT (Devlin et al. 2018) on 1-minute and 5-minute frequencies for four different stocks and market indices. The prediction results are measured by three evaluation criteria: MAE, RMSE, and MAPE. Informer obtained the best performance among all the networks on every dataset.

The chapter provided an overview of the historical development of research aimed at predicting financial asset movements. It emphasized several significant publications within this continually expanding



area of study. Although financial forecasting has consistently held popularity, recent years have witnessed an intensified focus on predictions conducted at higher frequencies, aided by the utilization of increasingly sophisticated systems akin to the one presented in this study.

### 3 Data

The data used in the research consider the BTC/USDT cryptocurrency pair from a period of 21.08.2019 to 24.07.2024 (5 years). Bitcoin (BTC) is a decentralized digital currency that allows peer-to-peer transactions. Tether (USDT) is a type of stablecoin, with a value that closely mirrors that of the USD. It actively works to keep its valuation stable through market mechanisms: each Tether issued is backed by one US dollar worth of assets. The pair closely mirrors the behavior of BTC/USD however it has a higher trading volume, as it is used by investors who want to hedge against the inherent volatility of their cryptocurrency investments while still keeping value within the crypto market.

Figure 1: Price of the BTC/USDT



Note: Price of BTC/USDT cryptocurrency pair in a period from 21.08.2019 to 24.07.2024. The data was obtained from <https://www.binance.com/en-NG/landing/data>.

The selection of this data for the research was driven by the following motivations:

- Availability of high-quality historical data. Many cryptocurrency exchanges provide free access to historical data, with fine-grained intervals, even up to 1 second. The data is easily accessible through either a dedicated API or can be downloaded as CSV files. This is a significant difference compared to the availability of historical stock prices, which are not easily accessible even for hourly interval data, and usually, the access requires a fee. Additionally, this concrete pair BTC/USDT was one of the earliest pairs available on the exchanges, so it has the longest historical data available.
- Cryptocurrency exchanges operate continuously, 24 hours a day, 7 days a week. Hence, there is no need for special handling of holidays and overnight price fluctuations during stock market closures. The historical price data forms a consistent time series with equally spaced intervals.
- Bitcoin stands as the most established cryptocurrency, with the BTC/USDT pair being the most widely traded and one of the highest in trading volume <sup>3</sup>, which makes it more liquid and stable compared to the other cryptocurrency pairs.

The data was obtained through the Binance API<sup>4</sup>. Binance is one of the largest global cryptocurrency exchanges that provides a platform to trade various cryptocurrencies. It was established in 2017 and offers a comprehensive range of services, including trading, listing, fundraising, and delisting or withdrawal of cryptocurrencies. The API allows downloading the so-called k-line data, the candlestick data of various intervals, consisting of: open time, close time, open price, high price, low price, close price and volume. The research uses k-line intervals of 5min, 15min and 30min.

After downloading, the data was checked for missing data points. Even though the exchange operates continuously, there were maintenance windows and the exchange sometimes was down. 8 gaps in the data were detected, with missing data of over 101 hours (4.2 days) in total. This constitutes 0.002 of all data points in the period, thus it shouldn't have a significant impact on the experiment. For simplicity's

<sup>3</sup><https://coinranking.com/exchange/-zdVbieRdZ+binance/markets>

<sup>4</sup><https://www.binance.com/en-NG/landing/data>

sake and to make the time series consistent, the gaps were filled by copying the price and volume values from the last available data point. Then, returns from each interval are computed as:

$$\text{returns} = \frac{\text{close price} - \text{open price}}{\text{open price}} \quad (1)$$

Table 3 presents the descriptive statistics of the BTC/USDT returns for all three intervals. We can observe, that the standard deviation of the returns grows with the interval length, which is expected. The minimum and maximum returns for each interval are quite alike, suggesting that significant price fluctuations occur abruptly and swiftly. Interestingly, the maximum return was observed for the 15min interval, indicating that the price must first have quickly increased and then dropped (in 30 min). The

Table 1: Descriptive statistics of BTC/USDT

Statistic	BTC/USDT 5min	BTC/USDT 15min	BTC/USDT 30min
count	518400	172800	86400
mean	0.0000060	0.0000176	0.0000346
std	0.0021843	0.0036712	0.0050768
min	-0.1022537	-0.1191688	-0.1662924
25% percentile	-0.0007716	-0.0013018	-0.0017677
50% percentile	0	0.0000094	0.0000280
75% percentile	0.0007855	0.0013443	0.0018575
max	0.1842885	0.2262878	0.1460125
kurtosis	203.12	140.22	58.61
skewness	0.57	0.97	-0.29
KS test stat.	0.49	0.49	0.49
KS test p-value	0.00e+00	0.00e+00	0.00e+00

Note: Descriptive statistics of returns with intervals of 5min, 15min and 30min. The statistics are not annualized. The null hypothesis of the Kolmogorov-Smirnov (KS) test is that the distribution is normal.

distributions are leptokurtic, with kurtosis significantly higher than the normal distributions, which increases with smaller intervals. In all cases, skewness is between -1 and 1 which indicates a relatively symmetrical distribution. Furthermore, in all three cases, Kolmogorov-Smirnov test for normality has a test statistic equal to 0.49 and a p-value of 0.0. Therefore, a null hypothesis that the data is normally distributed is rejected. The given analysis underscores the significance of investigating and creating trading systems that function in shorter time intervals. This is because rapid and abrupt price movements, which are critical for effective trading strategies, would go unnoticed by algorithms that operate on longer intervals.

### 3.1 Additional Data

The research utilizes data collected over several years, encompassing periods with varying market conditions and sentiments. To account for this, the data was enhanced with additional information that attempts to capture this information.

Additional data that was added is composed of:

**The Cboe Volatility Index (VIX Index)**<sup>5</sup>, a benchmark index that measures the market's expectation of future volatility. The VIX Index is based on options of the S&P 500 Index and is considered the leading indicator of the broad U.S. stock market. The VIX Index is recognized as the world's best gauge of the U.S. equity market. The daily close value of the VIX index is used.


**The Federal Funds effective rates**<sup>6</sup>, that is the interest rate at which depository institutions trade federal funds (balances held at Federal Reserve Banks) with each other overnight. The federal funds rate indirectly influences longer-term interest rates such as mortgages, loans, and savings. The data was available at the monthly frequency.

<sup>5</sup>[https://www.cboe.com/tradable\\_products/vix/vix\\_historical\\_data/](https://www.cboe.com/tradable_products/vix/vix_historical_data/)

<sup>6</sup><https://fred.stlouisfed.org/series/FEDFUNDS>




Figure 2: Cboe Volatility Index (VIX)



Note: Cboe Volatility Index (VIX) value in a period from 21.08.2019 to 24.07.2024. Data was obtained from <https://www.cboe.com/tradable-products/vix/vix-historical-data/>


Figure 3: Note: The Federal Funds effective rates



Note: The federal fund's effective rates in a period from 21.08.2019 to 24.07.2024. Data was obtained from <https://fred.stlouisfed.org/series/FEDFUNDS>

**Crypto Fear/Greed index<sup>7</sup>**, which is founded on elements such as volatility, market momentum, Reddit sentiment analysis, Google trends data and public surveys. The index is available at a daily frequency.

Figure 4: The Crypto Fear/Greed Index



Note: The Crypto Fear/Greed Index value in a period from 21.08.2019 to 24.07.2024. Data was obtained from <https://alternative.me/crypto/fear-and-greed-index/>

The aforementioned data was collected at a much lower frequency (daily/monthly) compared to the BTC/USDT data. Consequently, for every observation, a last known value had to be assigned, e.g. a

<sup>7</sup><https://alternative.me/crypto/fear-and-greed-index/>

5min observation with `open` time of 2022-08-17 04:05:00 was enriched with the value of VIX Index and Crypto Fear/Greed Index from 2022-08-16 and effective rates value from 2022-07. Note that the values from the previous day/month are used as it is unknown at what hour/day the value was published. Using the information from the same day/month carries a potential risk of using the information that wasn't available at the time.

### 3.2 Data Windows

The research follows the approach of Michańków et al. 2022, where strategies are independently evaluated on a rolling window that passes through the testing period. A rolling window consists of an *in sample* part of the 24 months (2 years) and *out of sample* part of 6 months. Each next window is moved forward by the length of the *out of sample* part. In the end, the results of each subperiod are combined to provide testing results for the whole testing period. This method allows us to better evaluate the robustness of the trading strategies, as it evaluates them independently on various periods when the market conditions are different.

Figure 5: Rolling data windows



Note: The figure presents how the dataset was split into rolling data windows consisting of *in sample* and *out of sample* parts.

In total six windows with identical number of data points were created. The *out of sample* is used for testing, further referenced as the *test* part. The *in sample* part is split into *train* and *validation* parts, with the *validation* part being 20% of the *in sample* data.

Table 2: Number of data points

Part	BTC/USDT 5min	BTC/USDT 15min	BTC/USDT 30min
train	165888	55296	27648
validation	41472	13824	6912
test	51840	17280	8640

Note: Number of data points in each of the dataset parts for the 5min, 15min and 30min rolling windows.

Tables 3, 3 and 5 further detail the descriptive statistics for each *out-of-sample* windows for 5min, 15min and 30min interval data. The statistics are consistent with the ones computed for the whole period (Table 3). The standard deviation is higher for the first two windows in all frequencies, suggesting that the market was more volatile during those periods. Kurtosis is significantly higher than the normal distributions with the values largest for the 1st and 5th window of 5min data and the 5th window for 15min data, indicating many data points with returns close to zero in those windows. The Kolmogorov-Smirnov tests for normality indicate that none of the distributions is normal.

The differences between *out-of-sample* data windows are best visible in Figure 6. Heatmaps on the left side of Figure 6 show the distances between the returns distributions between each *out-of-sample* window,

Table 3: Descriptive statistics for 5 min out-of-sample data

Statistic	W1	W2	W3	W4	W5	W6
count	51840	51840	51840	51840	51840	51840
mean	0.0000010	-0.0000089	0.0000016	0.0000050	0.0000080	0.0000099
std	0.0021663	0.0023011	0.0015967	0.0013956	0.0013791	0.0016259
min	-0.0989054	-0.0438821	-0.0469833	-0.0263642	-0.0866874	-0.0370693
25% percentile	-0.0010410	-0.0009298	-0.0005471	-0.0005338	-0.0004952	-0.0006930
50% percentile	-0.0000139	-0.0000029	0.0000070	0	0	0.0000064
75% percentile	0.0010127	0.0009137	0.0005638	0.0005350	0.0005185	0.0007262
max	0.0643117	0.0426144	0.0487321	0.0200162	0.0501479	0.0189663
kurtosis	124.60	27.33	72.17	27.74	419.46	18.82
skewness	-1.41	0.61	-0.27	-0.27	-3.84	-0.59
KS test stat.	0.49	0.49	0.50	0.50	0.50	0.50
KS test p-value	0.00e+00	0.00e+00	0.00e+00	0.00e+00	0.00e+00	0.00e+00

Note: Descriptive statistics of returns in data windows of 5min interval data. The statistics are not annualized. The null hypothesis of the Kolmogorov-Smirnov (KS) test is that the distribution is normal.

Table 4: Descriptive statistics for 15 min out-of-sample data

Statistic	W1	W2	W3	W4	W5	W6
count	17280	17280	17280	17280	17280	17280
mean	0.0000024	-0.0000276	0.0000048	0.0000149	0.0000240	0.0000296
std	0.0035625	0.0038529	0.0027021	0.0023567	0.0023828	0.0027726
min	-0.0675796	-0.0470951	-0.0502659	-0.0276761	-0.0897967	-0.0346516
25% percentile	-0.0017774	-0.0015974	-0.0009356	-0.0008720	-0.0008726	-0.0012156
50% percentile	-0.0000367	0.0000399	0.0000138	-0.0000080	0.0000222	0.0000141
75% percentile	0.0017230	0.0015821	0.0009610	0.0009004	0.0009122	0.0012940
max	0.0366072	0.0581246	0.0435405	0.0331951	0.0554152	0.0281460
kurtosis	19.08	19.40	38.81	21.94	176.31	11.54
skewness	-0.30	0.40	-0.64	0.22	-2.28	-0.31
KS test stat.	0.49	0.49	0.49	0.49	0.49	0.49
KS test p-value	0.00e+00	0.00e+00	0.00e+00	0.00e+00	0.00e+00	0.00e+00

Note: Descriptive statistics of returns in data windows of 15min interval data. The statistics are not annualized. The null hypothesis of the Kolmogorov-Smirnov (KS) test is that the distribution is normal.

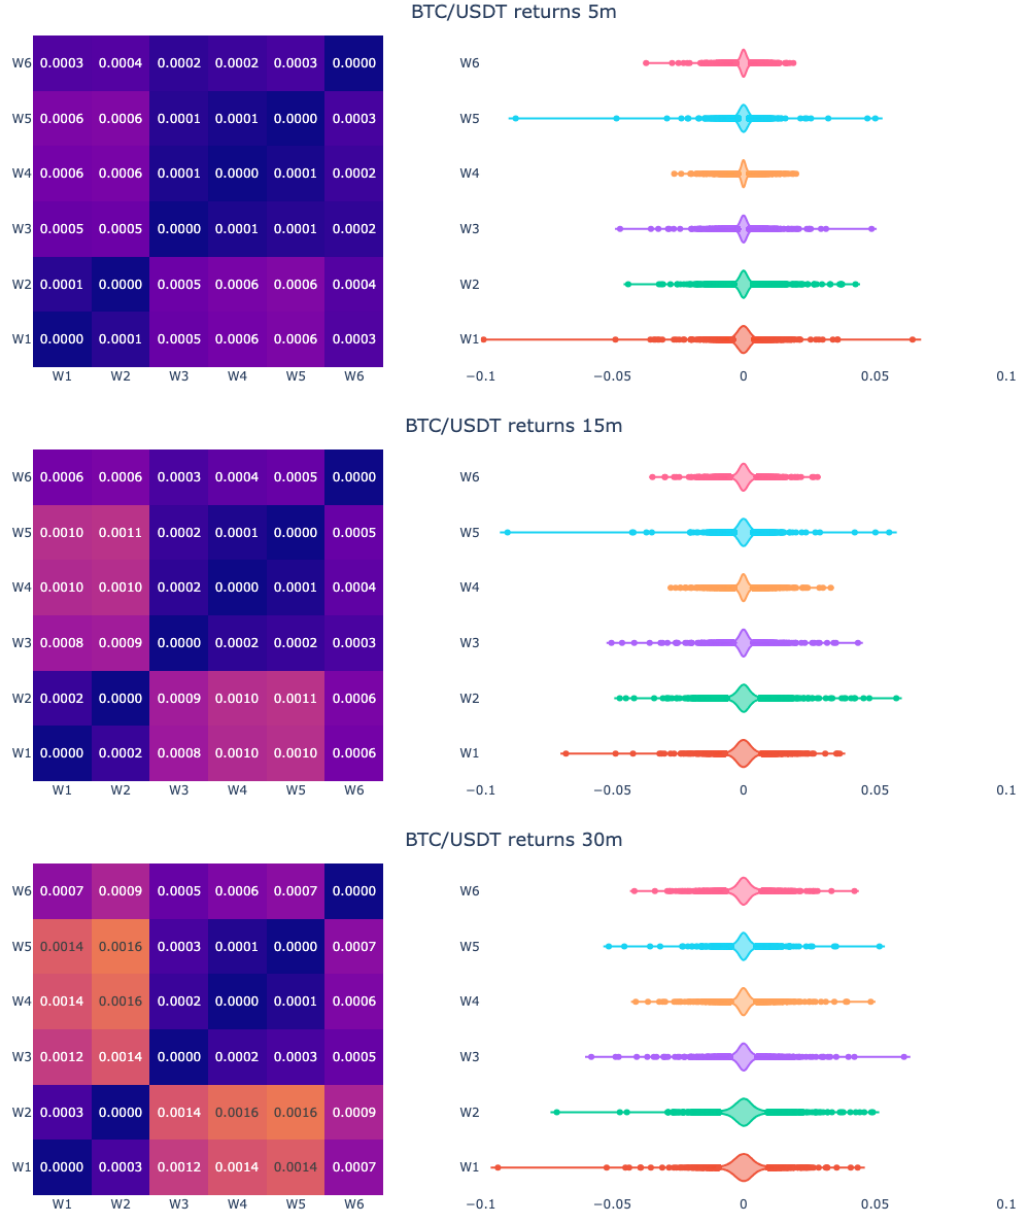
Table 5: Descriptive statistics for 30 min out-of-sample data

Statistic	W1	W2	W3	W4	W5	W6
count	8640	8640	8640	8640	8640	8640
mean	0.0000043	-0.0000540	0.0000101	0.0000296	0.0000473	0.0000595
std	0.0049962	0.0054682	0.0038721	0.0034004	0.0031349	0.0039924
min	-0.0934631	-0.0710468	-0.0579109	-0.0410247	-0.0512786	-0.0415287
25% percentile	-0.0023584	-0.0023627	-0.0012400	-0.0012067	-0.0011561	-0.0016455
50% percentile	0.0000360	0.0000073	0.0000117	-0.0000114	0.0000311	0.0000521
75% percentile	0.0023458	0.0022461	0.0013298	0.0012539	0.0012576	0.0017785
max	0.0434281	0.0492234	0.0610820	0.0484100	0.0517262	0.0422052
kurtosis	22.73	12.11	34.48	25.83	34.78	10.52
skewness	-0.82	0.08	-0.33	0.41	-0.18	-0.10
KS test stat.	0.49	0.49	0.49	0.49	0.49	0.49
KS test p-value	0.00e+00	0.00e+00	0.00e+00	0.00e+00	0.00e+00	0.00e+00

Note: Descriptive statistics of returns in data windows of 30min interval data. The statistics are not annualized. The null hypothesis of the Kolmogorov-Smirnov (KS) test is that the distribution is normal.

according to the Wasserstein metric (Vaserstein 1969). For all the frequencies, the most significant shift in the distributions can be observed between the first two and the next three windows, and then with the last, sixth, window. The violin plots of the distributions on the right-hand side of the figure provide additional information about the shapes of the distributions.

Figure 6: Distributions of returns for out-of-sample data



Note: Visualizations of the returns distributions for the out-of-sample data for different windows from testing period, for 5min, 15min and 30min intervals. On the left, heatmaps visualize distances between each pair of distributions according to Wasserstein distance. On the right distributions are visualized as violin plots.

## 4 Methodology

In this chapter, various trading strategies and the methodology for their evaluation are described. A fair comparison of strategies requires a formal definition of each strategy, clearly established comparison criteria and metrics that compare different aspects of the strategies. Those are detailed in the initial section of this chapter (4.1). The latter section (4.2) describes the strategies that are compared in the study. For each strategy, a mechanism for how the positions are selected is described as well as what are the hyperparameters that influence the strategy behavior.

## 4.1 Evaluation Framework

Strategies are evaluated over certain periods. Consider a time series consisting of  $T$  data points. The data point at each time step  $t \in \{1..T\}$  represents a single time interval. During each time interval, one of the three possible positions can be taken:

- **long position** - The asset is bought. The value of the portfolio increases if the price of an asset increases.
- **short position** - The asset short-sold and will need to be bought back when the position is closed. The value of the portfolio increases if the price of an asset decreases.
- **no position** - No asset in the portfolio. The value of the portfolio does not change, regardless of the price of the asset.

Additionally, the following assumptions are made:

- When the strategy position changes, assets are always bought/sold at the close price of the previous interval. That is, if at the time  $t - 1$  the position was **no position** and at the time  $t$  the position is **long position**, the assets are bought at the **close price** of the interval at time  $t - 1$ .
- Each such position change incurs an exchange fee  $e$ . In the experiments, the value of  $e$  is assumed to be 0.1%.<sup>8</sup>.
- Any fraction of the asset can be traded<sup>9</sup>.
- No splitting of portfolio allocation, the asset is always bought/sold with the full portfolio value.
- At the end of the evaluation period all positions must be closed, that is, at time  $T$  the position must be **no position**.

Formally, a trading strategy is a function  $s_\theta : \mathbb{R}^{\lambda \times d} \rightarrow \{-1, 0, 1\}$ , where  $\theta$  are strategy hyperparameters,  $\lambda$  is the size of the lookback window and  $d$  is the dimensionality of the vector  $x_t \in \mathbb{R}^d$  that represents new, known and relevant information at each time step  $t$ , which can be used by the strategy. The co-domain values  $\{-1, 0, 1\}$  respectively, represent: **short position**, **no position**, and **long position**. Then, the recommended position  $p_t$  according to the strategy  $s$  at time  $t$  is calculated in the following way:

$$s_\theta([x_{t-\lambda}; \dots; x_{t-1}]) = p_t \quad (2)$$

where  $[\cdot; \cdot]$  is a concatenation operation. Note that the strategy has access only to information before time  $t$ . The positions are computed for each time index  $t$  over a given period. The portfolio value  $E_t$  at any time  $t$  can be calculated using the following formula:

$$E_0 = 1 \quad (3)$$

$$E_t = E_{t-1}(1 + r_t \cdot p_t)(1 - (|p_t - p_{t-1}|e)) \quad (4)$$

where  $r_t$  is a return in an interval  $t$  computed as in 1,  $p_t$  is a position recommended by the strategy and  $e$  is the exchange fee rate. The portfolio value at the end of the evaluation period  $T$  is equal to  $E_T$ .

Furthermore, following the methodology of Michańków et al. 2022, the following metrics are computed:

**Annualized Return Compounded (ARC)** Metric that shows the annualized rate of return for the evaluation period, computed as

$$ARC = \left(\frac{E_T}{E_0}\right)^{\frac{Y}{T}} \quad (5)$$

where  $E_0$  is the portfolio value at the beginning of the evaluation period,  $E_T$  is the portfolio value at the end of the evaluation period,  $T$  is number of the intervals in the evaluation period and  $Y$  is the number of intervals in the year<sup>10</sup>.

<sup>8</sup>The value is equal to the standard spot fee rate on Binance <https://www.binance.com/en/fee/trading>

<sup>9</sup>In practice the smallest fraction of the asset that can be bought on Binance is 0.00001

<sup>10</sup>It is assumed that there are 365 days in a year, so for the datasets used in the research, this value is equal to  $Y_{5m} = 105120$ ,  $Y_{15m} = 35040$  and  $Y_{30m} = 17520$ .

**Annualized Standard Deviation (ASD)** A metric that represents the yearly standard deviation of returns, assessing the strategy's volatility. The metric is determined as

$$ASD = \sqrt{\frac{Y}{T} \sum_{t=0}^T (r_t - \bar{r})^2} \quad (6)$$

where  $r_t$  is the return at time  $t$ ,  $\bar{r}$  is the mean return in the evaluation period,  $T$  is the number of intervals in the evaluation period and  $Y$  is the number of intervals in a year the same as in 5.

**Information Ratio (IR\*)** The Sharpe ratio (Sharpe 1998) is calculated under the assumption of a risk-free rate of zero. The metric becomes the annualized return divided by its annualized standard deviation:

$$IR* = \frac{ARC}{ASD} \quad (7)$$

**Maximum Drawdown (MD)** The portfolio's maximum drawdown represents the greatest decline measurable between the portfolio's highest value and its consecutive lowest point. This metric acts as an indicator of possible percentage loss within a specified period

$$MD = \max_{(t,\tau) \in \{[0..T]^2: t < \tau\}} \left( \frac{E_t - E_\tau}{E_t} \right) \quad (8)$$

where  $T$  represents the total number of intervals during the evaluation timeframe, with  $E_t$  and  $E_\tau$  denoting the portfolio values at times  $t$  and  $\tau$ , respectively.

**Modified Information Ratio (IR\*\*)** is the Information Ratio (7) adjusted by the Maximum Drawdown (8), a metric initially introduced by Kosciuszko et al. 2019:

$$IR^{**} = IR \times \frac{|ARC|}{MD} \quad (9)$$

**Number of trades (N)** The number of strategy positions changes during the evaluation period. Calculated as

$$N = \sum_{t=1}^T (|p_t - p_{t-1}|) \quad (10)$$

where  $p_t$  is the position at the time  $t$ . Note that while it is possible to change the position from **long position** to **short position** in a single time frame, it is counted as two trades: closing the long position and opening the short position.

**Percentage of Long Position (LONG) / Short Position (SHORT)** Metric that measures the percentage of how many time intervals during the entire evaluation period the strategy took the long/short position.

$$LONG = \frac{1}{T} \sum_{t=0}^T \left( \frac{(p_t + 1)p_t^2}{2} \right) \times 100\% \quad (11)$$

$$SHORT = \frac{1}{T} \sum_{t=0}^T \left( \frac{(1 - p_t)p_t^2}{2} \right) \times 100\% \quad (12)$$

## 4.2 Strategies

This study involves the evaluation and comparison of five strategies. Each strategy considered in the study consists of the complete mechanism that converts the signal, whether originating directly from the data, a technical indicator, or a machine learning model, into the actual position. This is something that is often neglected, especially when evaluating the effectiveness of machine learning models. Although the precision of the model's predictions is crucial, a fair assessment of the strategy is feasible only after these predictions are converted into portfolio positions.



#### 4.2.1 Buy&Hold strategy

The first strategy, the only viable one according to the Efficient Markets Hypothesis, is the "Buy&Hold" strategy. It is regarded as a benchmark that other strategies aim to surpass. Framed as a formal function as described in (2), strategy has no hyperparameters and at every time interval takes a **long position**:

$$s^{BuyAndHold}(\cdot) = 1 \quad (13)$$

#### 4.2.2 Moving Average Convergence Divergence strategy

The strategy is based on the technical indicator of the Moving Average Convergence Divergence (MACD). The indicator was proposed by Gerald Appel almost 50 years ago and is still widely used by technical analysts today (Appel 2005)<sup>11</sup>. The standard MACD oscillator is a collection of three time series calculated from daily historical price data, most often the closing price. Time series are calculated using the Exponential Moving Average (EMA)<sup>12</sup> with different window sizes. These are subsequently converted into positions as follows: when the MACD line moves above the signal line, it signifies an uptrend and acts as a buy signal. Conversely, when the MACD line falls below the signal line, it signals a downtrend, and one should exit the position:

$$MACD_t(\cdot) = EMA_t^{fast}(\cdot) - EMA_t^{slow}(\cdot) \quad (14)$$

$$SIGNAL_t(\cdot) = EMA_t^{signal}(MACD_t(\cdot)) \quad (15)$$

where the *fast*, *slow* and *signal* are window sizes for Exponential Moving Average. Typically, these values are 12, 26, and 9 respectively, but in the context of this study, they are considered strategy hyperparameters:

$$s_{(fast,slow,signal,short)}^{MACD}(\cdot) = \begin{cases} 1 & \text{if } MACD_{t-1} \geq SIGNAL_{t-1} \\ 0 & \text{if } MACD_{t-1} < SIGNAL_{t-1} \text{ and } short = 0 \\ -1 & \text{if } MACD_{t-1} < SIGNAL_{t-1} \text{ and } short = 1 \end{cases} \quad (16)$$

where  $short \in \{0, 1\}$  is an additional strategy hyperparameter that controls whether a strategy uses **exit position** or **short position** on the sell signal.

#### 4.2.3 Relative Strength Index strategy

The second strategy based on the technical indicator employs the Relative Strength Index (RSI) indicator. The RSI was developed by J.Welles Wilder Jr. and serves as a momentum oscillator designed to assess whether changes in asset prices suggest an overbought or oversold condition (Wilder 1978). The RSI is calculated as:

$$RSI_t(\cdot) = \left( \frac{100}{1 + RS_t(\cdot)} \right) \quad (17)$$

$$RS_t = \frac{SMMA^{window}(U(\cdot))}{SMMA^{window}(D(\cdot))} \quad (18)$$

$$U_t = \max(0, \text{close price}_t - \text{close price}_{t-1}) \quad (19)$$

$$D_t = \max(0, \text{close price}_{t-1} - \text{close price}_t) \quad (20)$$

where  $RS$  is a relative strength factor computed using Smoothed Moving Average (SMMA)<sup>13</sup> calculated over the "Up period" ( $U$ ) and "Down period" ( $D$ ) time series, using *window* time intervals before  $t$ -th interval. Although usually the window is assumed to be 14, the study treats *window* as a hyperparameter of the strategy.

To translate oscillator values into the positions, usually two thresholds overbought/oversold for buy/sell signal are used. If the RSI crosses the over-sold threshold, it indicates a bullish sign, and

<sup>11</sup> Although the indicator was initially shared in the 1970s within the Systems & Forecasts newsletter, the book detailing this indicator wasn't released until 2005.

<sup>12</sup> [https://en.wikipedia.org/wiki/Exponential\\_smoothing](https://en.wikipedia.org/wiki/Exponential_smoothing)

<sup>13</sup> [https://en.wikipedia.org/wiki/Moving\\_average#Modified\\_moving\\_average](https://en.wikipedia.org/wiki/Moving_average#Modified_moving_average)

when it crosses below the over-bought threshold, it is a bearish sign. Instead, in the study, four thresholds are used to indicate when: enter long position (**enter long**), exit long position (**exit long**), enter short position (**enter short**) and exit short position (**exit short**).

$$s_{\theta_{RSI}}^{RSI}(\cdot) = \begin{cases} 1 & \text{if } RSI_{t-1}(\cdot) > \text{enter long} \\ 0 & \text{if } RSI_{t-1}(\cdot) < \text{exit long and } p_{t-1} = 1 \\ -1 & \text{if } RSI_{t-1}(\cdot) < \text{enter short} \\ 0 & \text{if } RSI_{t-1}(\cdot) > \text{exit short and } p_{t-1} = -1 \\ p_{t-1} & \text{else} \end{cases} \quad (21)$$

where  $\theta_{RSI} = (\text{window}, \text{enter long}, \text{exit long}, \text{enter short}, \text{exit short})$  are strategy hyperparameters and  $p_{t-1}$  is a position taken on the previous time interval.

#### 4.2.4 Informer based strategies

The last two strategies use the predictions of the state-of-the-art machine learning model for sequence modeling: Informer (Zhou et al. 2021). The informer uses the modified Transformer architecture (Vaswani et al. 2017) to better handle long sequences. Several studies have shown that it surpasses other deep learning architectures in predicting financial time series (Ding et al. 2023, Lu et al. 2023).

**Informer Architecture** follows the encoder-decoder structure (Chollet 2017, Bahdanau et al. 2014) where the encoder maps an input sequence of vectors  $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$  to a sequence of continuous representations  $\mathbf{Z} = (\mathbf{z}_1, \dots, \mathbf{z}_n)$ . Given  $\mathbf{Z}$ , the decoder then generates an output sequence  $\mathbf{Y} = (\mathbf{y}_1, \dots, \mathbf{y}_m)$  one element at a time. Both the encoder and decoder are built wiserveral of identical layers. Let  $\mathbf{X}_0 = \mathbf{X}$  be an input to the first encoder layer, each  $i$ -th encoder layer transforms the sequence of  $\mathbf{X}_{i-1}$  input sequence to  $\mathbf{X}_i$  sequence, where  $i \in \{1, \dots, L^{enc}\}$  and  $L^{enc}$  is the number of encoder layers. The encoder output is equal to  $\mathbf{X}_{L^{enc}} = \mathbf{Z}$ . Similarly, each decoder layer transforms the sequence of previously predicted targets  $\mathbf{Y}_0$  and the encoder output  $\mathbf{Z}$  into the sequence  $\mathbf{Y}_j$ , which is then consumed by the next decoder layer. The final output of the decoder layer is the output of the decoder  $\mathbf{Y} = \mathbf{Y}_{L^{dec}}$ , where  $j \in \{1..L^{dec}\}$  and  $L^{dec}$  is the number of decoder layers. *Number of the encoder layers* and *the number of decoder layers* are model hyperparameters. The main building block of both those layers is *multi-head probe-sparse self-attention*.

**Multi-head probe-sparse self-attention:** The canonic self-attention of the Transformer is defined based on the tuple inputs, i.e. query (Q), key (K), and value (V), which performs the scaled dot-product as:

$$\mathcal{A}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d}}\right) \mathbf{V} \quad (22)$$

where  $\mathbf{Q} \in \mathbb{R}^{L_Q \times d}$ ,  $\mathbf{K} \in \mathbb{R}^{L_K \times d}$ ,  $\mathbf{V} \in \mathbb{R}^{V_Q \times d}$  and  $L_Q$ ,  $L_K$ ,  $L_V$  are lengths of the query, key and value sequences respectively and the  $d$  is the dimensionality of the model. Softmax<sup>14</sup> is a function that converts a vector of real numbers, so it can be interpreted as a probability distribution. In this formulation, computing attention requires the quadratic time dot-product computation. The authors of Informer proposed a measurement that transforms  $\mathbf{Q}$  into a sparse  $\bar{\mathbf{Q}}$  matrix, that contains only the top  $u = c \cdot \ln(L_Q)$  queries under the sparsity measurement, where  $c$  is a constant sampling factor. The change makes the computation much more efficient and the architecture more applicable to the longer sequences.

$$\mathcal{A}'(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax}\left(\frac{\bar{\mathbf{Q}}\mathbf{K}^T}{\sqrt{d}}\right) \mathbf{V} \quad (23)$$

In each layer, attention is performed  $h$  times with different linear projections of inputs. This mechanism is called multi-head attention and  $h$  is the model hyperparameter that controls the number so-called attention heads.

$$\mathcal{H}_i(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \mathcal{A}'(\mathbf{Q}\mathbf{W}_i^Q, \mathbf{K}\mathbf{W}_i^K, \mathbf{V}\mathbf{W}_i^V) \quad (24)$$

$$\mathcal{MH}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = [\mathcal{H}_1(\cdot); \dots; \mathcal{H}_h(\cdot)] \quad (25)$$

where the projections  $\mathbf{W}_i^Q \in \mathbb{R}^{d \times L_Q}$ ,  $\mathbf{W}_i^K \in \mathbb{R}^{d \times L_K}$ ,  $\mathbf{W}_i^V \in \mathbb{R}^{d \times V_Q}$  are learned parameters.

<sup>14</sup>[https://en.wikipedia.org/wiki/Softmax\\_function](https://en.wikipedia.org/wiki/Softmax_function)

The self-attention means that the layer "attends" to itself, that is, all queries, keys, and values consider the same input sequence i.e.  $\mathbf{Q} = \mathbf{K} = \mathbf{V} = \mathbf{X}$ , where  $\mathbf{X}$  is an input to the *multi-head probe-sparse self-attention layer* ( $\mathcal{MHSA}$ ).

$$\mathcal{MHSA}(\mathbf{X}) = \mathcal{MH}(\mathbf{X}, \mathbf{X}, \mathbf{X}) \quad (26)$$

**Encoder:** The encoder layers of Informer consist of a multi-head probe-sparse self-attention layer with a distilling operation that trims the input's time dimension. Given  $\mathbf{X}_{i-1}$  input to the  $i$ -th Informer layer, the output of the  $i$ -th layer is calculated as:

$$\mathbf{X}_i = \text{MaxPool}(\text{ELU}(\text{Conv1d}(\mathcal{MHSA}(\mathbf{X}_{i-1})))) \quad (27)$$

where *Conv1d* performs one-dimensional convolutional filter over the time dimension (Kiranyaz et al. 2019), with *ELU* activation function (Clevert et al. 2016) and max-pooling layer (Nagi et al. 2011) that down-samples the layer output.

**Decoder:** Informer uses a standard Transformer decoder, which consists of two  $\mathcal{MHSA}$  layers in which the first takes an input of past target outputs and the second output of the encoder. Those are followed with a position-wise fully connected layer ( $\mathcal{FFN}$ ). It consists of two linear transformations with a ReLU (Agarap 2019) activation in between, and is applied to vectors at each position separately:

$$\mathcal{FFN}(\mathbf{x}) = \max(0, \mathbf{x}\mathbf{W}_1 + \mathbf{b}_1)\mathbf{W}_2 + \mathbf{b}_2 \quad (28)$$

where  $\mathbf{W}_1 \in \mathbb{R}^{d \times f}$ ,  $\mathbf{W}_2 \in \mathbb{R}^{f \times d}$ ,  $\mathbf{b}_1 \in \mathbb{R}^f$ ,  $\mathbf{b}_2 \in \mathbb{R}^d$  are trained parameters and  $d$  is the model dimensionality and  $f$  is the size of the fully connected layer, both of these values are model hyperparameters. Around each of those sub-layers, there is a residual connection (Kolen and Kremer 2001), dropout regularization (*Dropout*), and a layer normalization (*LayerNorm*) (Ba et al. 2016).

$$\mathbf{Y}'_j = \text{LayerNorm}(\text{Dropout}(\mathcal{MHSA}(\mathbf{Y}_{j-1})) + \mathbf{Y}_{j-1}) \quad (29)$$

$$\mathbf{Y}''_j = \text{LayerNorm}(\text{Dropout}(\mathcal{MHSA}(\mathbf{Z}) + \mathbf{Z}) \quad (30)$$

$$\mathbf{Y}'''_j = \mathbf{Y}'_j + \mathbf{Y}''_j = [\mathbf{y}'''_{1j}; \dots; \mathbf{y}'''_{mj}] \quad (31)$$

$$\mathbf{Y}_j = \text{LayerNorm}(\text{Dropout}([\mathcal{FFN}(\mathbf{y}'''_{ij}); \dots; \mathcal{FFN}(\mathbf{y}'''_{mj})]) + \mathbf{Y}''') \quad (32)$$

**Model input:** The model input  $\mathbf{X}_0$  consists of a sequence of  $n$   $d$ -dimensional input vectors. These vectors are derived by adding together three input components:

- **Real input variables:** All real-value input variables are concatenated and linearly transformed into  $d$ -dimensional vector.
- **Categorical variables:** All categorical variables are mapped into the embeddings - real-value vectors that are model parameters. Then they are concatenated and linearly transformed into  $d$ -dimensional vector.
- **Positional embeddings:** Since the Informer model contains no recurrence and convolution, the ordering of input vectors needs to be modeled explicitly. To achieve this, a method known as positional encodings is employed. For each position a vector of sine and cosine functions of different frequencies is computed, that is, each dimension of the positional encoding corresponds to a sinusoid.

**Loss function:** Commonly, financial data forecasting is done using the Mean Squared Error (MSE) loss function or is framed as a classification problem with the logistic loss function. However, optimizing the model using these loss functions does not align with the objective of devising a strategy aimed at maximizing returns. Consider an asset, which price does not move for 99 out of 100 intervals and then suddenly drops. A model that predicts the same price every time would have a 99% accuracy, however a strategy based on such model predictions would have very poor results. This example illustrates that while the model using MSE/Logistic loss functions may be trained to have high accuracy of predictions, those may not necessarily be *meaningful* predictions in terms of building the trading strategy. Moreover, with point-wise predictions, such as the ones coming from MSE loss, there is no information on the certainty of the prediction.

To address the aforementioned issues the study uses two different loss functions that may be better suited for training models that are later used in trading strategy. The first loss is a Quantile Loss, that is,

instead of predicting the point, the model predicts a distribution, more precisely quantiles of the returns distribution. With such a signal, a strategy can be constructed that considers confidence in the prediction of the model. The Quantile Loss is defined as:

$$QuantileLoss = \frac{1}{N} \sum_{i=1}^N \sum_q \max(q \cdot (y - \hat{y}), (1 - q) \cdot (\hat{y} - y)) \quad (33)$$


where  $y$  is observed value,  $\hat{y}$  is model prediction,  $N$  is number of observations.  $q \in [0, 1]$  is a quantile, in the study the following quantiles were used  $q \in \{0.01, 0.02, 0.03, 0.05, 0.1, 0.25, 0.5, 0.75, 0.9, 0.95, 0.97, 0.98, 0.99\}$ <sup>15</sup>.

The other loss function considered is a Generalized Mean Absolute Directional Loss (GMADL) (Michańków et al. 2024). It puts more emphasis on the direction of the returns, i.e. whether they were positive or negative rather than precision, and rewards the model for correctly predicting larger return values. The loss function is defined as

$$GMADL = \frac{1}{N} \sum_{i=1}^N (-1) \cdot \left( \frac{1}{1 + e^{-a \cdot y \cdot \hat{y}}} - \frac{1}{2} \right) \cdot (|y|)^b \quad (34)$$

where  $y$  is observed value,  $\hat{y}$  is model prediction,  $N$  is number of observations.  $a$  and  $b$  are loss function parameters that control the steepness of the function slope. In the study, they are considered to be equal to  $a = 100$  and  $b = 2$ .

Figure 7: GMADL vs MAE loss functions.



Note: Comparison of the two loss functions: Mean Absolute Error (MAE) on the left and Generalized Mean Absolute Directional Loss (GMADL) on the right.

**Strategies** Leveraging the predictions generated by the Informer model, one can formulate a trading strategy. Given that the models were trained using two distinct approaches, the Quantile Loss and GMADL, each method provides the opportunity to develop a unique strategy that optimally utilizes the model's predictions.

**Quantile Informer strategy:** A strategy based on the predictions of the model trained with Quantile Loss is constructed in the following manner: probabilities are defined to *enter long*, *exit long*, *enter short*, and *exit short* positions, as well as the *threshold* value. If the return predicted by the model is above the *threshold* with a certain probability, i.e.  $1 - \text{enter long}$ <sup>16</sup> quantile is above the threshold, a strategy

<sup>15</sup>A decision to train the model with many different quantiles was motivated by the fact, that then a strategy that bases on this model predictions can be fine-tuned with different quantile values without the need for retraining the model.

<sup>16</sup>Model prediction  $\hat{y}_{t,q}$  for a quantile  $q$  means that the value is below  $\hat{y}_{t,q}$  with  $q$  probability, and above  $\hat{y}_{t,q}$  with  $1 - q$  probability.

enters **long position**. Conversely, if the *enter short* quantile is below  $-\text{threshold}$  a **short position** is opened. The interpretation is that the model prediction has the confidence of *enter long/enter short* that the returns are above/below the *threshold*.

$$s_{\theta_{\text{Quantile}}}^{\text{Quantile}}(\cdot) = \begin{cases} 1 & \text{if } \hat{y}_{t,1-\text{enter long}} > \text{threshold} \\ 0 & \text{if } \hat{y}_{t,\text{exit long}} < -\text{threshold} \text{ and } p_{t-1} = 1 \\ -1 & \text{if } \hat{y}_{t,\text{enter short}} < -\text{threshold} \\ 0 & \text{if } \hat{y}_{t,1-\text{exit short}} > \text{threshold} \text{ and } p_{t-1} = -1 \\ p_{t-1} & \text{else} \end{cases} \quad (35)$$

where  $\hat{y}_{t,q}$  is a model prediction at time  $t$  for the  $q$  quantile and  $p_{t-1}$  is a position on the previous timestep. The strategy hyperparameters are  $\theta_{\text{Quantile}} = (\Theta_{\text{Quantile}}, \text{enter long}, \text{exit long}, \text{enter short}, \text{exit short}, \text{threshold})$ . and  $\Theta_{\text{Quantile}}$  are parameters of the Informer model trained with Quantile loss.

**GMADL Informer strategy:** Strategy that uses predictions of the Informer trained with GMADL is defined as the following: four threshold values are defined when to *enter long*, *exit long*, *enter short* and *exit short*. Those thresholds are compared directly with the return predicted by the model. While in this case, one does not know explicitly the confidence of those predictions, because of how GMADL is constructed they should be more reliable than the predictions of the model trained with MAE loss.

$$s_{\theta_{\text{GMADL}}}^{\text{GMADL}}(\cdot) = \begin{cases} 1 & \text{if } \hat{y}_t > \text{enter long} \\ 0 & \text{if } \hat{y}_t < \text{exit long} \text{ and } p_{t-1} = 1 \\ -1 & \text{if } \hat{y}_t < \text{enter short} \\ 0 & \text{if } \hat{y}_t > \text{exit short} \text{ and } p_{t-1} = -1 \\ p_{t-1} & \text{else} \end{cases} \quad (36)$$

where  $\hat{y}_t$  is the model prediction for time step  $t$  and  $p_{t-1}$  is position on the previous timestep. The strategy hyperparameters are  $\theta_{\text{GMADL}} = (\Theta_{\text{GMADL}}, \text{enter long}, \text{exit long}, \text{enter short}, \text{exit short})$  and  $\Theta_{\text{GMADL}}$  are parameters of the Informer model trained with GMADL loss.

## 5 Experiments

The evaluation framework and strategies described in Chapter 4 were implemented in Python<sup>17</sup> and are publicly available in the Gitlab repository<sup>18</sup>. Technical indicators were computed using TA-lib library<sup>19</sup>, Informer model was trained using PytorchForecasting<sup>20</sup> and the Informer implementation was taken from the following Github repository<sup>21</sup>. Section 5.1 describes the hyperparameter selection process for each strategy and presents the values that were selected. The evaluation of the test data is then presented and analyzed in Section 5.2. The last Section 6 presents the sensitivity analysis of how the selection of strategy hyperparameters influences the strategy performance.

### 5.1 Training and hyperparameter selection

Each strategy was independently instantiated for each data window. Strategies that were based on the machine learning model were trained using the *training* part of the data window and the hyperparameters were selected using *validation* part. Strategies that do not require training i.e. MACD MACD-based strategy and RSI-based strategy had their hyperparameters selected using only the *validation* part of the dataset. An alternative would be to use all in-sample (*train* and *validation* parts) to select hyperparameters for non-ml strategies, however a decision to use only *validation* part was motivated by the fact that to make the comparison fair, the same data should be used for similar strategy hyperparameter selection (thresholds) for ml and non-ml based strategies. The best hyperparameter combinations were selected by comparing the IR\*\* (4.1) metric. The remainder of this section details the hyperparameter search space for each strategy and lists which hyperparameter values were selected for the strategies for each data window.

<sup>17</sup><https://www.python.org/>

<sup>18</sup><https://gitlab.com/FilipStefaniuk/wne-msc-thesis>

<sup>19</sup><https://ta-lib.org/>

<sup>20</sup><https://pytorch-forecasting.readthedocs.io/en/stable/>

<sup>21</sup><https://github.com/martinwhl/Informer-PyTorch-Lightning>

### 5.1.1 MACD Strategy hyperparameters

As described in Section 4.2.2, the MACD strategy has four hyperparameters, three of which (*fast*, *slow*, *signal*) control the size of the window for moving averages and the *short* whether a sell signal should be interpreted to close **long position** or to additionally open **short position**. Table 6 presents the values tested for the hyperparameters. A *short* hyperparameter takes binary values. Window sizes take values of the 16 values of the Fibonacci sequence, the sequence has been chosen as it grows exponentially, thus allowing to search efficiently the space of possible windows - using the equally distributed values would be less efficient as large window sizes are less sensitive to small window size changes. Note that even though combinations were searched exhaustively, an additional constraint that prevents the *fast* window from being longer than *slow* window has been applied.

Table 6: Hyperparameter search space for MACD strategy.

Parameter	Values
<i>fast</i>	[2, 3, 5, 8, 13, 21, 34, 55, 89, 144, 233, 377, 610, 987, 1597, 2584]
<i>slow</i>	[2, 3, 5, 8, 13, 21, 34, 55, 89, 144, 233, 377, 610, 987, 1597, 2584]
<i>signal</i>	[2, 3, 5, 8, 13, 21, 34, 55, 89, 144, 233, 377, 610, 987, 1597, 2584]
<i>short</i>	[0, 1]

NotConsideringred values of hyperparameters for MACD strategy, all 8192 combinations were searched exhaustively.

Hyperparameters selected for each window for every data interval are presented in Table 7. A bias towards larger window sizes can be observed, suggesting that the MACD indicator might not be appropriate for high frequencies such as 5, 15, or 30 minutes. Additionally, for all three frequencies, the strategies for the last two windows do not use short positions. Setting the short position on the downtrend should always be beneficial, so this can suggest that the data in the last two windows may be much less predictable.

Table 7: Selected hyperparameters for MACD strategy.

Window	<i>fast</i>	<i>slow</i>	<i>signal</i>	<i>short</i>
W1-5min	8	2584	987	1
W2-5min	377	610	610	1
W3-5min	987	2584	987	1
W4-5min	610	2584	987	1
W5-5min	987	1597	987	0
W6-5min	1597	2584	377	0
W1-15min	8	377	1597	1
W2-15min	144	987	2584	1
W3-15min	377	2584	2584	1
W4-15min	377	2584	610	1
W5-15min	233	610	233	0
W6-15min	144	377	2584	0
W1-30min	1597	2584	2584	1
W2-30min	233	2584	2584	1
W3-30min	13	2584	2584	1
W4-30min	233	987	233	1
W5-30min	144	233	233	0
W6-30min	1597	2584	1597	0

Note: Table presents hyperparameter values selected for the MACD strategy for each window of 5min, 15min, and 30min interval data.

### 5.1.2 RSI strategy hyperparameters

As detailed in Section 4.2.3 the RSI strategy has five hyperparameters. The *window* parameter sets the window size for computing the moving averages. The other four set thresholds when **long position** and **short position** should be opened or closed. For the window size, similarly as with MACD strategy values of the Fibonacci sequence were used. Combinations of window sizes and various thresholds listed



in Table 8 were searched exhaustively, "-" symbolizes that the position will never be opened/closed regardless of the RSI oscillator value. Note that *exit long* and *exit short* hyperparameter values make sense to be different from "-" only if corresponding *enter long* and *enter short* hyperparameters are different than "-". However, the reverse situation can occur as the long/short position is automatically closed as the next short/long position is entered, that is e.g. a combination where *enter long* is not equal to "-" and *exit long* is equal to "-" can occur if *enter short* is present as the long position will be closed if *enter short* threshold is reached.

Table 8: Hyperparameter space for RSI strategy.

Parameter	Values
<i>window</i>	[2, 3, 5, 8, 13, 21, 34, 55, 89, 144, 233, 377, 610, 987, 1597, 2584]
<i>enter long</i>	[-, 70, 75, 80, 85, 90, 95]
<i>exit long</i>	[-, 5, 10, 15, 20, 25, 30]
<i>enter short</i>	[-, 5, 10, 15, 20, 25, 30]
<i>exit short</i>	[-, 70, 75, 80, 85, 90, 95]

Note: Considering values of hyperparameters for RSI strategy, all 38416 combinations were searched exhaustively.

The table lists selected hyperparameters. Contrary to the MACD strategy, the RSI strategy seems to benefit from higher frequencies, and the selected values for the *window* hyperparameter were rather small. Only in few cases *exit long* and *exit short* hyperparameters were set, meaning that in most cases when the long/short position was closed, the short/long position was immediately opened.

Table 9: Selected hyperparameters for the RSI strategy.

Window	<i>window</i>	<i>enter long</i>	<i>exit long</i>	<i>enter short</i>	<i>exit short</i>
W1-5min	21	80	-	25	-
W2-5min	34	75	-	25	-
W3-5min	5	95	-	15	-
W4-5min	8	90	5	-	-
W5-5min	13	90	20	-	-
W6-5min	21	85	15	-	-
W1-15min	5	80	-	15	-
W2-15min	34	75	-	30	-
W3-15min	5	95	-	10	-
W4-15min	34	75	-	30	-
W5-15min	13	70	-	5	-
W6-15min	21	85	15	-	-
W1-30min	5	95	-	5	-
W2-30min	21	70	-	30	-
W3-30min	5	95	-	15	85
W4-30min	8	95	5	-	-
W5-30min	13	70	-	5	-
W6-30min	34	80	25	-	-

Note: Table presents hyperparameter values selected for the RSI strategy for each window of 5min, 15min, and 30min interval data.

### 5.1.3 Informer Strategies training and hyperparameters

The process of selecting the hyperparameters for Informer-based strategies was divided into two steps. First, model hyperparameters were selected by training the Informer with different combinations of model hyperparameters, then the strategy hyperparameters were selected for the best model. In other words, first, the model that can produce the best predictions was selected, and then, based on the predictions of the models, the best strategy was constructed. The reason for this split was motivated mostly by the heaviness of the training process and the fact that if both model and strategy parameters are changed at the same time, it is difficult to assess which has the biggest impact on the final strategy result. This divide carries the risk that the best model, in terms of the loss function, does not necessarily mean the

best strategy, in terms of  $IR^{**}$ . However, this is addressed, to some extent, by selecting a loss function that aligns more closely to maximize returns (4.2.4).

**Training and model hyperparameters** A separate instance of Informer was trained for each data window and for each Quantile and GMADL loss function. The target variable in each case was **returns**. Each training run consisted of 40 epochs with an early stopping<sup>22</sup>. The input to the model consisted of a sequence whose length was controlled by the *past window* hyperparameter. The sequence consisted of two types of variables: **real** and **categorical**. The **real** variables were normalized before they were passed to the model. The **categorical** variables were mapped into real-value embeddings.

Table 10: Input variables to Informer model.

Variable type	Variable
real	open price
	high price
	low price
	close price
	volume
	returns
	vix
	fed rates
	fear and greed
	open to close price
	high to close price,
	low to close price
	vol 1h
	vol 1d
	vol 7d
	sma 1h to close price
	sma 1d to close price
	sma 7d to close price
	ema 1h to close price
	ema 1d to close price
	macd
	macd signal
	rsi
	low bband to close price
	up bband to close price
	mid bband to close price
categorical	hour
	weekday

Table 10 presents the variables that were passed to the model as an input. The **open price**, **high price**, **low price**, **close price** and **volume** come directly from the dataset described in Chapter 3. The **returns** were computed according to (1) and the **vix**, **fed rates**, **fear and greed** are additional data that the dataset was augmented with (Section 3.1). The other variables were derived from the variables in the data set: **open to close price**, **high to close price**, **low to close price**, **high to low price** are the ratios of the corresponding price values. **vol 1h**, **vol 1d**, **vol 7d** variables are volatility computed on the 1 hour, 1 day, and 7 days past data, respectively. **sma 1h to close price**, **sma 1d to close price** and **sma 7d to close price** are ratios of Simple Moving Average computed on 1 hour, 1 day, and 7-day windows to close price. **ema 1h to close price** and **ema 1d to close price** are ratios of the Exponential Moving Average computed on data from the past 1 hour and 1 day to close price. The **macd** and **macd signal** comes from the MACD technical indicator computed using the

<sup>22</sup>The number of observations in each dataset (5min, 15min 30min) varied and was rather large for a 5 min dataset, so instead of running validation at the end of each epoch, it was run after every 300 batches. The *patience* parameter that controls early stopping was set to 15, meaning that if in the next 15 consecutive validations, the validation loss does not improve, the training is stopped.

`close price` value and default values for window sizes<sup>23</sup>. The `rsi` is a RSI technical indicator, computed using the 14 past data points. Finally `low bband to close price`, `up bband to close price` and `mid bband to close price` are ratios of Bollinger Bands (Bollinger 2002) to `close price`.

The `categorical` variables have values of 0 to 23 for `hour` and 0 to 6 for `weekday` and represent an hour-of-the-day and a day-of-the-week of the interval (in respect to interval's `close time`)<sup>24</sup>.

The hyperparameters of the model were determined by randomly sampling the hyperparameter space. The random search was carried out due to a large number of possible combinations and the training time, which took approximately 1 hour per training run<sup>25</sup>. A sample of 30 combinations (out of 1166400) was tested and the combination of hyperparameters with the lowest validation loss value was selected. Note that since two values of the two different loss functions cannot be directly compared, this procedure had to be carried out separately for each loss function and for every frequency, totaling in  $num\ frequencies \times num\ loss\ functions \times sample\ size = 3 \times 2 \times 30 = 180$  training runs. Furthermore, as this process is extremely heavy, the hyperparameter search was done only on the first data window, and the exact selected combination of training and model hyperparameters was later used to train the model on other data windows for the given time interval. The space of Informer model hyperparameters is presented in Table 11. where *past window* is a length of the input sequence, *batch size* and *learning rate* are the standard machine learning training parameters. Other hyperparameters are specific to Informer architecture and are described in Section 4.2.4.

Table 11: Informer model hyperparameter space.

Parameter	Values
past window	[20, 21, ..., 119, 120]
batch size	[64, 128, 256]
learning rate	[0.001, 0.0005, 0.0001]
model dimensionality ( <i>d</i> )	[256, 512, 1024]
fully connected dimensionality ( <i>f</i> )	[256, 512, 1024]
attention heads ( <i>h</i> )	[1, 2, 4, 6]
dropout	[0.05, 0.1, 0.2, 0.3]
number of encoder layers	[1, 2, 3]
number of decoder layers	[1, 2, 3]

Note: Table presents all hyperparameter values that were considered. The dimension of the grid is 1166400, so the search was done by randomly sampling 30 combinations.

Table 12 lists all the hyperparameter combinations selected for the Informer model for different data frequency and loss functions. Interestingly, it seems that the better models were the smaller ones, having shorter input sequences (20-30) and smaller dimensionality. This might have been caused by the still relatively small size of the training dataset (Fig 2). The different values of *batch size* and *learning rate* from the tested range did not seem to have a significant impact on the model training.

**Quantile Informer strategy hyperparameters** For each data window, once the model was trained, its predictions were used to create a trading strategy. To select the optimal hyperparameters, *validation* part of the data window was used. For the model with Quantile loss, the parameters consisted of *enter long*, *exit long*, *enter short*, *exit short*, and *threshold* hyperparameters described in Section 4.2.4. Here the possible hyperparameter values correspond to the quantiles of the loss function, with "-" meaning that the position won't be opened/closed regardless of the model prediction. Three values for the *threshold* were tested with the smallest one equal to the *exchange fee*. All the possible values of the hyperparameters are listed in Table 13.

Selected hyperparameters are presented in Table 14. It can be observed that mostly the parameters that require high model confidence to open/close the position were selected. This may be caused by the fact, that each position incurs a fee and each mistake is extremely costly. With transaction fee equal to 0.1% even if the asset price stays the same, already with  $\log_{0.999}(0.5) \approx 692$  incorrect position changes the portfolio value is halved - which is not that many considering the minute frequency data.

<sup>23</sup>The window sizes used for computing MACD were: 12 for fast window, 26 for slow window and 9 for signal window.

<sup>24</sup>E.g. an interval with close time at 2023-10-18 01:29:59 will have values of 1 (second hour of the day) for the `hour` variable and 2 (third day of the week - Wednesday) for the `weekday` variable

<sup>25</sup>The training was done using a single Nvidia GTX 1080TI card.

Table 12: Selected hyperparameters for Informer model.

Parameter	Quantile Informer			GMADL Informer		
	5min	15min	30min	5min	15min	30min
past window	22	28	72	28	30	26
batch size	64	64	128	256	128	128
model dimensionality ( $d$ )	256	256	1024	256	256	256
fully connected layer dim	512	256	1024	256	256	256
attention heads $h$	2	6	2	2	2	4
dropout	0.05	0.1	0.05	0.01	0.01	0.01
number of encoder layers	1	1	1	1	1	1
number of decoder layers	1	3	3	3	1	3
learning rate	0.0001	0.001	0.0001	0.0001	0.001	0.0001

Note: Tabpresentstes hyperparameter values selected for Informer model trained on 5min, 15min and 30min interval data with Quantile Loss and GMADL loss function.

Table 13: Hyperparameter space for Quantile Informer Strategy

Parameter	Values
<i>enter long</i>	[-, 0.75, 0.9, 0.95, 0.97, 0.98, 0.99]
<i>exit long</i>	[-, 0.75, 0.9, 0.95, 0.97, 0.98, 0.99]
<i>enter short</i>	[-, 0.75, 0.9, 0.95, 0.97, 0.98, 0.99]
<i>exit short</i>	[-, 0.75, 0.9, 0.95, 0.97, 0.98, 0.99]
<i>threshold</i>	[0.001, 0.002, 0.003]

Note: Considering values of hyperparameters for Quantile Informer strategy, all 12288 combinations were searched exhaustively.

**GMADL Informer strategy hyperparameters** The other strategy was created using the Informer trained with GMADL loss function. The hyperparameters of this strategy are *enter long*, *exit long*, *enter short* and *exit short*. They are described in Section 4.2.4. The tested values range from 0.0001 to 0.007. "-" indicates that the position will not be taken regardless of the model prediction. The possible values of the hyperparameters are listed in Table 15.

The selected hyperparameter values are presented in Table 16. Similarly as with the RSI strategy, the values for *exit long* and *exit short* were rarely selected, meaning that throughout the period, the **long position** or **short position** was always open. There seems to be no bias towards a larger return threshold in the lower frequency data, indicating that the sudden price changes happen and can be leveraged even in such frequent intervals as 5 min.

## 5.2 Evaluation Results

After the hyperparameters were selected, all strategies were evaluated on the corresponding test window parts. The metrics were computed for each window and the whole period by concatenating the positions selected by each strategy for every period. In this section, first, the evaluation results are analyzed separately for each frequency, and then the strategies that achieved the best overall results are compared and discussed.

### 5.2.1 Evaluation on 5 min data

Evaluation results for the 5 min interval throughout the testing period are presented in Figure 8, a GMADL Informer strategy significantly outperformed the other strategies. Throughout the testing period, it achieved annualized returns of 115%, while maintaining an annualized standard division at a level similar to the Buy&Hold benchmark. What is more, the maximum drawdown of this strategy was only 32.7% which is significantly lower than the 77.3% of Buy&Hold. The strategy kept the long position and *short position* for approximately a similar period. The next best strategy was an RSI-based strategy, which however for almost 40% of the time did not hold a position. These are the only two strategies that beat the benchmark. It is worth noting that they had much fewer position changes than the other strategies that performed worse, suggesting that they were better at detecting the trend changes and filtering out the insignificant signal.

Table 14: Selected hyperparameter values for Quantile Informer strategy.

Window	<i>enter long</i>	<i>exit long</i>	<i>enter short</i>	<i>exit short</i>	<i>threshold</i>
W1-5min	-	-	0.980	0.970	0.003
W2-5min	0.970	-	0.950	-	0.003
W3-5min	0.990	-	0.950	-	0.002
W4-5min	0.970	0.990	-	-	0.003
W5-5min	0.900	0.900	-	-	0.002
W6-5min	0.950	0.950	-	-	0.003
W1-15min	0.980	-	0.990	-	0.003
W2-15min	0.990	0.950	-	-	0.003
W3-15min	-	-	0.990	0.990	0.003
W4-15min	0.990	0.970	-	-	0.003
W5-15min	0.950	0.990	-	-	0.001
W6-15min	0.990	0.950	-	-	0.002
W1-30min	0.980	-	0.950	0.970	0.003
W2-30min	-	-	0.980	0.990	0.001
W3-30min	-	-	0.990	0.900	0.003
W4-30min	0.950	0.990	-	-	0.001
W5-30min	0.900	0.990	-	-	0.001
W6-30min	0.970	0.980	-	-	0.003

Note: Table presents hyperparameter values selected for the Quantile Informer strategy for each window of 5min, 15min, and 30min interval data.

Table 15: Hyperparameter space for GMADL Informer Strategy

Parameter	Values
<i>enter long</i>	[-, 0.001, 0.002, 0.003, 0.004, 0.005, 0.006, 0.007]
<i>exit long</i>	[-, -0.001, -0.002, -0.003, -0.004, -0.005, -0.006, -0.007]
<i>enter short</i>	[-, -0.001, -0.002, -0.003, -0.004, -0.005, -0.006, -0.007]
<i>exit short</i>	[-, 0.001, 0.002, 0.003, 0.004, 0.005, 0.006, 0.007]

Note: Considering values of hyperparameters for GMADL Informer strategy, all 4096 combinations were searched exhaustively.

Quantile Informer strategy performed poorly, as the second worst strategy, suggesting either poor quality of the predictions or inability to effectively use the predictions. Note that this strategy also had the most position changes, suggesting insufficient filtering of when the position should be changed, or incorrect predictions of large price changes. The worst-performing strategy was one based on the MACD indicator.

Looking at the evaluation results for each of the data windows separately, presented in Figure 9, confirms the superiority of the GMADL informer strategy. It was the top performing strategy on all the evaluation periods but the last one, on which the best performing strategy was Buy&Hold, suggesting this was the most difficult period to trade. Most of the GMADL Strategy trades were done during the second window period, while during the other periods, the positions were changed sparingly. In the third period, the number of transactions was the smallest, only 16, and the strategy did not hold any position for 80% of the time; however, it still achieved the best results compared to the other strategies, as it correctly predicted and leveraged the sudden price swings.


The RSI strategy performed quite well in the first three periods; however, in the next three, the quality suddenly dropped. A similar situation can be observed for the Quantile Informer strategy. This might suggest a more unpredictable market during the later periods. The poor overall quality of the MACD strategy in the whole period could be explained by the extremely poor performance in the first two periods (those were the periods with higher volatility, see Table 2), in which the value of the portfolio dropped to 0.375 and 0.498 respectively.

Table 16: Selected hyperparameter values for GMADL Informer Strategy

Window	<i>enter long</i>	<i>exit Long</i>	<i>enter Short</i>	<i>exit Short</i>
W1-5min	0.004	-	-0.005	-
W2-5min	0.002	-	-0.001	-
W3-5min	-	-	-0.006	0.003
W4-5min	0.002	-	-0.005	-
W5-5min	0.002	-	-0.003	-
W6-5min	0.001	-	-0.007	-
W1-15min	0.005	-	-0.002	-
W2-15min	0.006	-	-0.002	-
W3-15min	-	-	-0.001	0.005
W4-15min	0.007	-	-0.005	-
W5-15min	0.001	-	-0.004	-
W6-15min	0.002	-0.002	-	-
W1-30min	0.007	-	-0.001	-
W2-30min	0.003	-	-0.007	-
W3-30min	-	-	-0.001	0.005
W4-30min	0.002	-	-0.006	-
W5-30min	0.005	-0.001	-	-
W6-30min	0.003	-	-0.004	-

Note: Table presents hyperparameter values selected for the GMADL Informer strategy for each window of 5min, 15min, and 30min interval data.

Figure 8: Evaluation results on 5 min data



Strategy	VAL	ARC	ASD	IR*	MD	IR**	N	LONG	SHORT
Buy&Hold	1.441	13.1%	57.7%	0.228	77.3%	0.039	2	100.0%	0.0%
MACD Strategy	0.516	-20.0%	54.1%	-0.370	85.8%	-0.087	2536	50.4%	32.4%
RSI Strategy	3.341	50.3%	50.4%	0.999	30.0%	1.676	846	28.3%	33.5%
Quantile Informer	0.956	1.5%	47.9%	-0.032	54.0%	-0.001	3396	40.2%	27.8%
GMADL Informer	9.747	115.9%	54.4%	2.129	32.7%	7.552	846	44.8%	41.5%

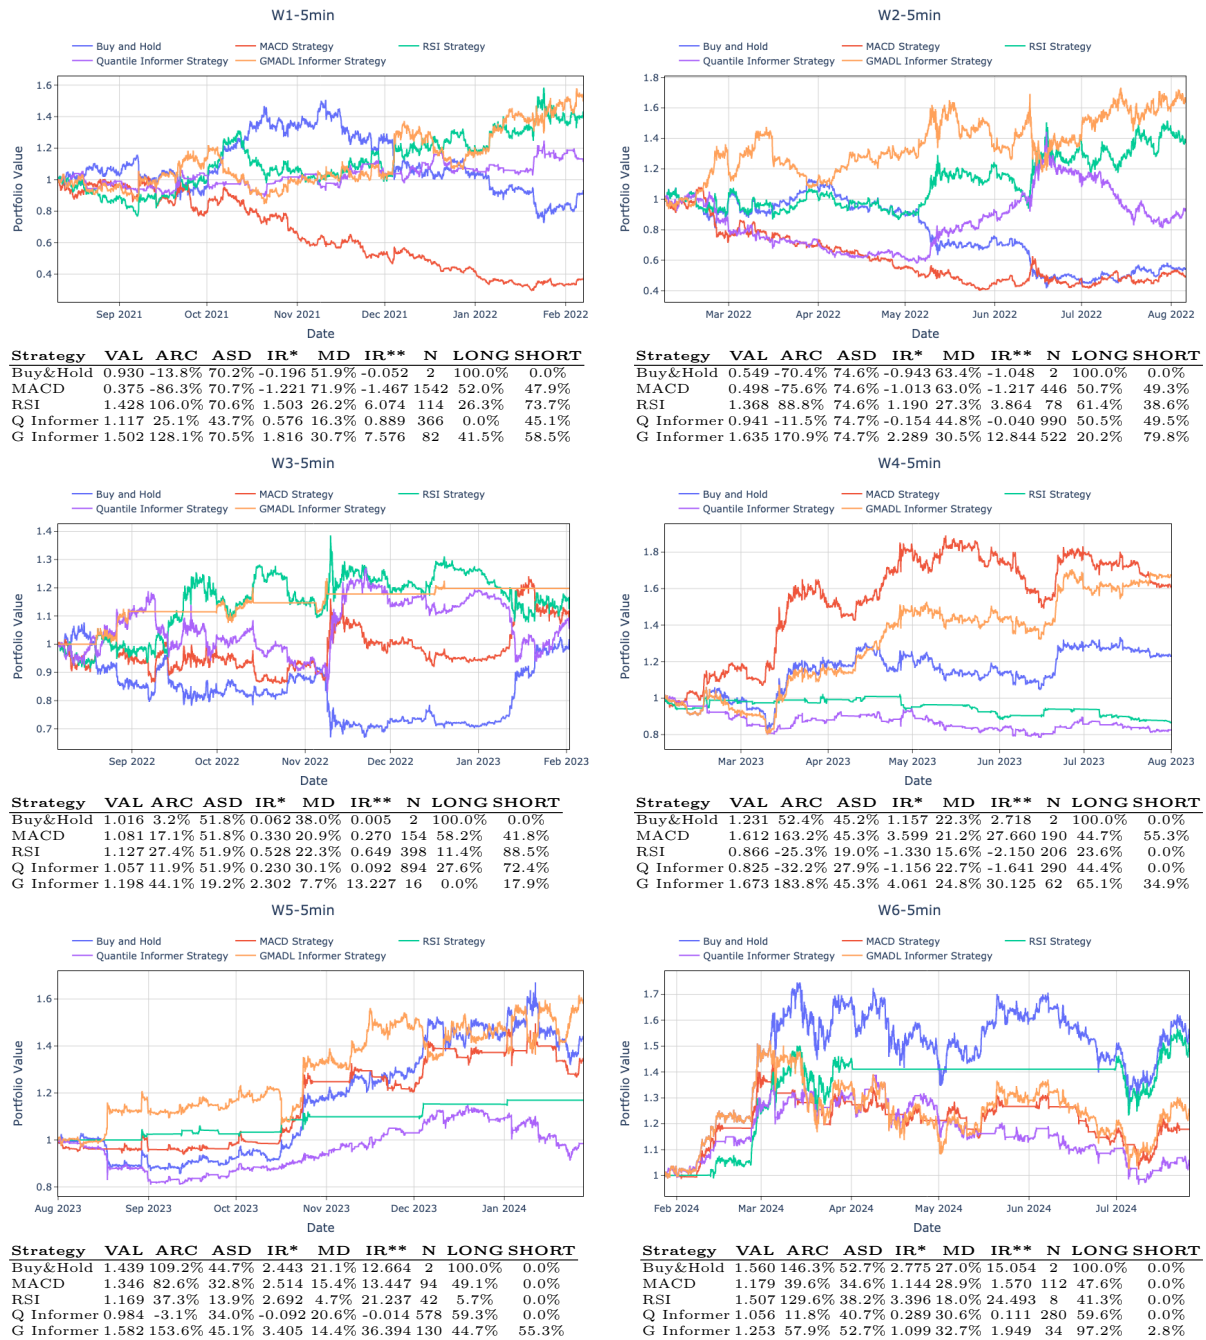
Note: Evaluation results on the whole testing period of 5-minute interval BTC/USDT data. The presented metrics are: portfolio value at the end of the evaluation period (VAL), Annualized Return Compound (ARC), Annualized Standard Deviation (ASD), Information Ratio (IR\*), Maximum Drawdown (MD), Modified Information Ratio (IR\*\*), number of trades (N) and percent of the long/short positions (LONG/SHORT).

### 5.2.2 Evaluation on 15 min data

The results for the 15 min data are presented in Figure 10. Here, the GMADL Informer strategy also yielded the best results. While the modified information ratio was smaller than the GMADL Informer strategy from the 5 min data it was the best one from the strategies evaluated on 15 min interval. The



Figure 9: Evaluation results for individual windows on 5 min data




Note: Evaluation results for each of the out-of-sample data windows of 5-minute interval BTC/USDT data. The presented metrics are: portfolio value at the end of the evaluation period (VAL), Annualized Return Compound (ARC), Annualized Standard Deviation (ASD), Information Ratio (IR\*), Maximum Drawdown (MD), Modified Information Ratio (IR\*\*), number of trades (N) and percent of the long/short positions (LONG/SHORT). MACD = MACD strategy. RSI = RSI strategy. Q Informer = Quantile Informer. G Informer = GMADL Informer.

next best strategy was the Buy&Hold benchmark. The other strategies performed poorly, achieving results worse than Buy&Hold. Quantile Informer Strategy had a position only for the 40% of time suggesting the selection of the hyperparameters that too aggressively limit the event when the position is entered/closed. MACD and RSI Strategies had more position changes, which might have negatively contributed to their performance due to high exchange fees.

The results of the 15-minute individual windows presented in Figure 11, show that the GMADL Informer strategy was the best only in two windows (first and third) and the majority of the value was accumulated during the first period. What is more, it was the worst-performing strategy in the fourth period.

Although the Quantile Informer strategy again presented poor overall results, it was the best-performing strategy in the fourth window. However, during the other periods, the positions either were changed too

Figure 10: Evaluation results on 15 min data



Strategy	VAL	ARC	ASD	IR*	MD	IR**	N	LONG	SHORT
Buy&Hold	1.440	13.1%	56.0%	0.234	77.2%	0.040	2	100.0%	0.0%
MACD Strategy	0.468	-22.6%	52.4%	-0.432	83.2%	-0.118	1312	51.8%	31.3%
RSI Strategy	0.800	-7.3%	55.7%	-0.131	66.7%	-0.014	1206	54.0%	43.1%
Quantile Informer	0.945	-1.9%	37.8%	-0.051	48.4%	-0.002	824	28.5%	16.8%
GMADL Informer	3.296	49.6%	52.7%	0.942	47.4%	0.987	362	49.4%	37.7%

Note: Evaluation results on the whole testing period of 15-minute interval BTC/USDT data. The presented metrics are: portfolio value at the end of the evaluation period (VAL), Annualized Return Compound (ARC), Annualized Standard Deviation (ASD), Information Ratio (IR\*), Maximum Drawdown (MD), Modified Information Ratio (IR\*\*), number of trades (N) and percent of the long/short positions (LONG/SHORT).

frequently or were missing on the high price swings. With data of 15-minute intervals, similarly as with 5 min, the last period proved to be the most difficult, and again no strategy beat the Buy&Hold benchmark during that period. Despite poor overall results, the MACD Strategy wasn't significantly worse than the other strategies except on the first window, on which a significant loss of portfolio value (value at the end of this period was 0.35) influenced the result of the whole testing period.

### 5.2.3 Evaluation on 30 min data

In the longest 30-minute interval dataset, the strategies based on the technical analysis indicator outperformed the ones based on the Informer model. The results presented in Figure 12, show that both the RSI strategy and the MACD strategy achieved a better-modified information ratio than Buy&Hold, while Informer-based strategies fell below the benchmark. This is also the first time the MACD-based strategy achieved better results than the Buy&Hold strategy. Interestingly, despite the best overall performance, the RSI strategy has visible long periods of holding no position.

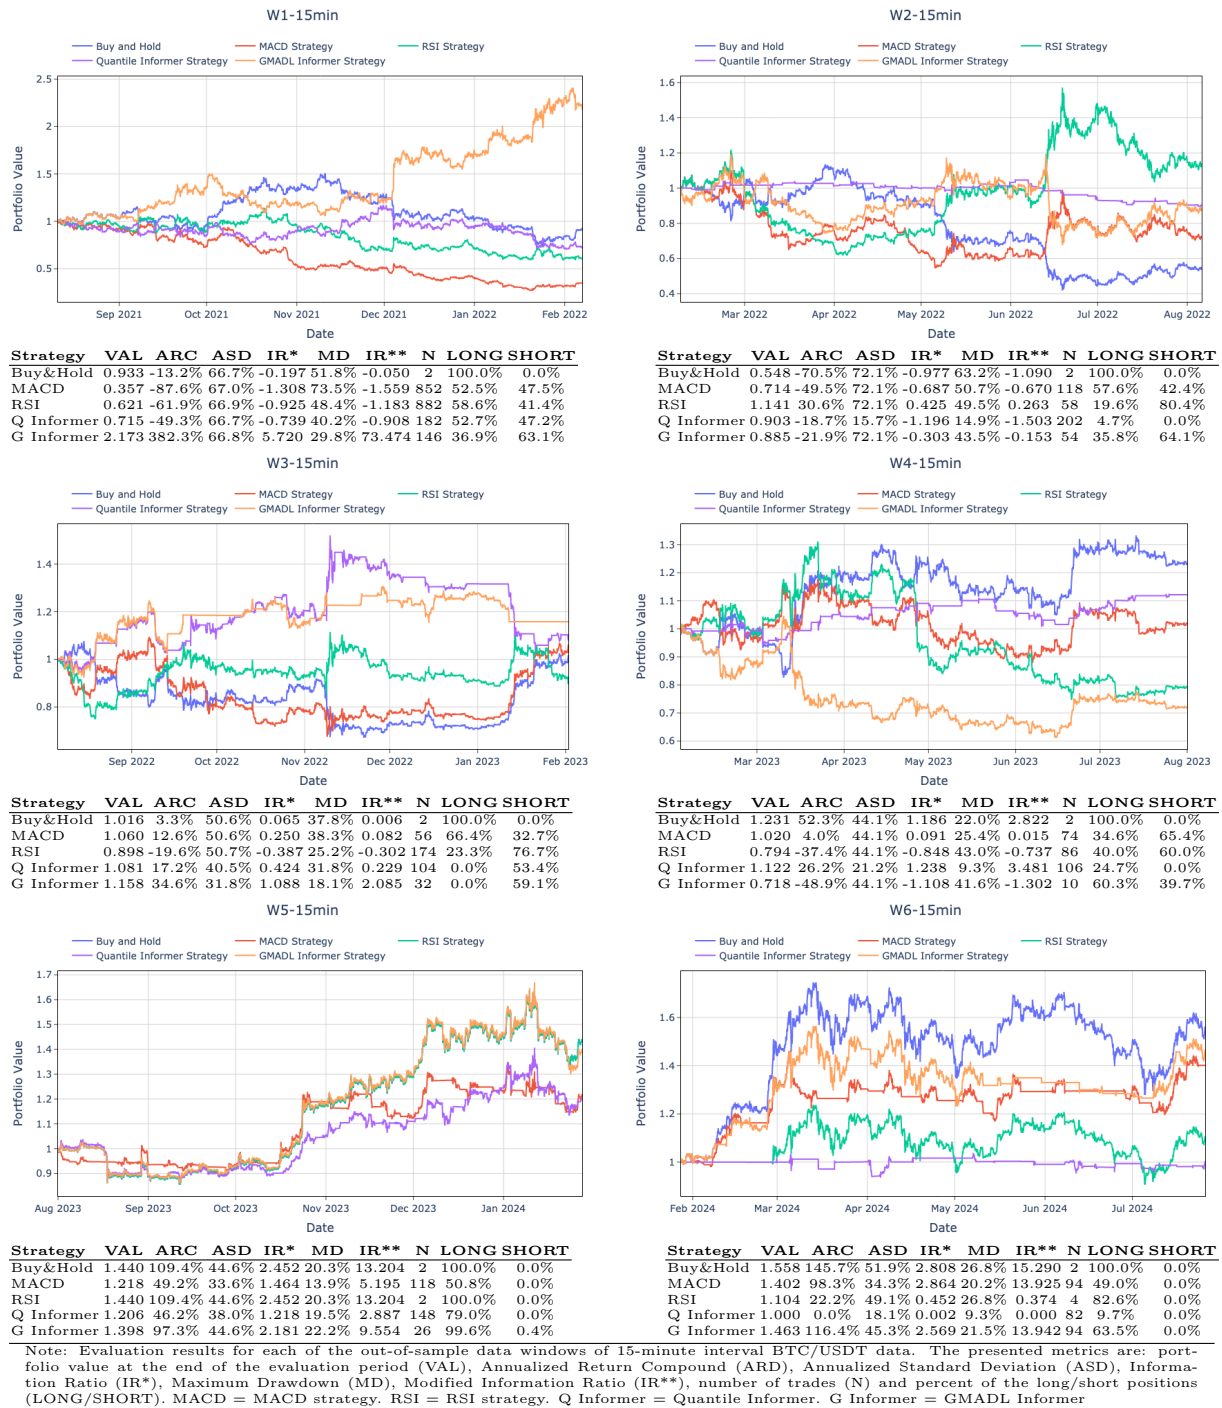
Those periods can be further analyzed by looking at individual windows in Figure 13. In fact, during the first of these periods, the RSI strategy took a position for only 0.02% of the time in the fourth window, and no position in the sixth window. This may suggest that the distribution of validation data on which the hyperparameters were selected differed from the testing data or that there was high uncertainty during those periods.

The GMADL strategy performed worse than the Buy&Hold benchmark in four of six windows. However, it performed well on the windows in which the RSI strategy was not taking position. Quantiles Informer strategy once more turned out to be poor, being the worst strategy in half of the evaluation windows and losing almost half the portfolio value during the fourth period. Bad results of the machine learning model-based strategies may suggest that the dataset with 30-minute windows had too few examples for the models to learn effectively.

### 5.2.4 Best strategies

The top five strategies that performed best overall during the testing period are again presented in Figure 14. All strategies presented on the chart performed better than the Buy&Hold benchmark.


Figure 11: Evaluation results for individual windows on 15 min data



GMADL Informer strategy seems to be benefiting from the higher frequency data, and while the result on 15-minute data is also in the top 5, the performance of the strategy using 5-minute data is significantly better. The worse performance on 15-minute data and 30-minute data can be explained by the fact that those datasets contain fewer examples, so it might have been too few for the model to train well.

The RSI-based strategy turned out to give quite good results, both for 5min and 30min data. This indicates that despite the long-standing presence of technical indicators and the availability of more advanced and intricate models for generating trading signals, these indicators remain valuable and can be employed to develop effective trading strategies. MACD strategy was performing poorly for the higher frequency data, however, for the 30-minute interval data the strategy was in the top five best. This might indicate that the MACD oscillator might not be applicable for the higher frequencies but is still useful

Figure 12: Evaluation results on 30 min data



Strategy	VAL	ARC	ASD	IR*	MD	IR**	N	LONG	SHORT
Buy&Hold	1.440	13.1%	55.9%	0.235	77.2%	0.040	2	100.0%	0.0%
MACD Strategy	1.952	25.4%	52.4%	0.485	59.2%	0.207	328	52.3%	28.3%
RSI Strategy	4.542	66.8%	46.2%	1.444	39.9%	2.415	378	30.7%	28.0%
Quantile Informer	0.629	-14.5%	36.9%	-0.393	55.1%	-0.104	1784	30.2%	15.2%
GMADL Informer	0.834	-6.0%	51.9%	-0.115	68.2%	-0.010	1298	42.9%	38.2%

Note: Evaluation results on the whole testing period of 30-minute interval BTC/USDT data. The presented metrics are: portfolio value at the end of the evaluation period (VAL), Annualized Return Compound (ARC), Annualized Standard Deviation (ASD), Information Ratio (IR\*), Maximum Drawdown (MD), Modified Information Ratio (IR\*\*), number of trades (N) and percent of the long/short positions (LONG/SHORT).

for the lower frequencies. This hypothesis seems to be confirmed by the hyperparameters selected for the strategies that had a bias toward longer sequences.

The Quantile Informer strategy did not beat the Buy&Hold benchmark. In fact, it was one of the worst-performing strategies. This suggests that the Quantile loss function is not the best choice for training a model for trading strategy or that the proposed strategy does not use the model predictions effectively. However, more in-depth analysis should be carried out, testing the training with various quantile combinations to confirm this hypothesis. Specifically, in this study, a decision has been made to use a Quantile Loss with many different quantiles and then select the quantile value for entering/exiting the position as a part of the strategy hyperparameter search. This method provided the advantage of eliminating the requirement to retrain the model every time a new quantile pair is selected. However, training the model once with many quantiles could have an impact on the quality of the predictions.

To verify the statistical significance of the results, a probabilistic t-test that compares the Information Ratios of the best strategies against the Buy&Hold benchmark was conducted. The test uses a simplified formula for the standard error of Sharpie ratio, expressed as  $\frac{\sigma}{\sqrt{N}}$ , where  $N$  is the sample size of the returns and  $\sigma$  is the standard deviation of the excess of the strategy return over the benchmark return. The full test statistic is then as follows:

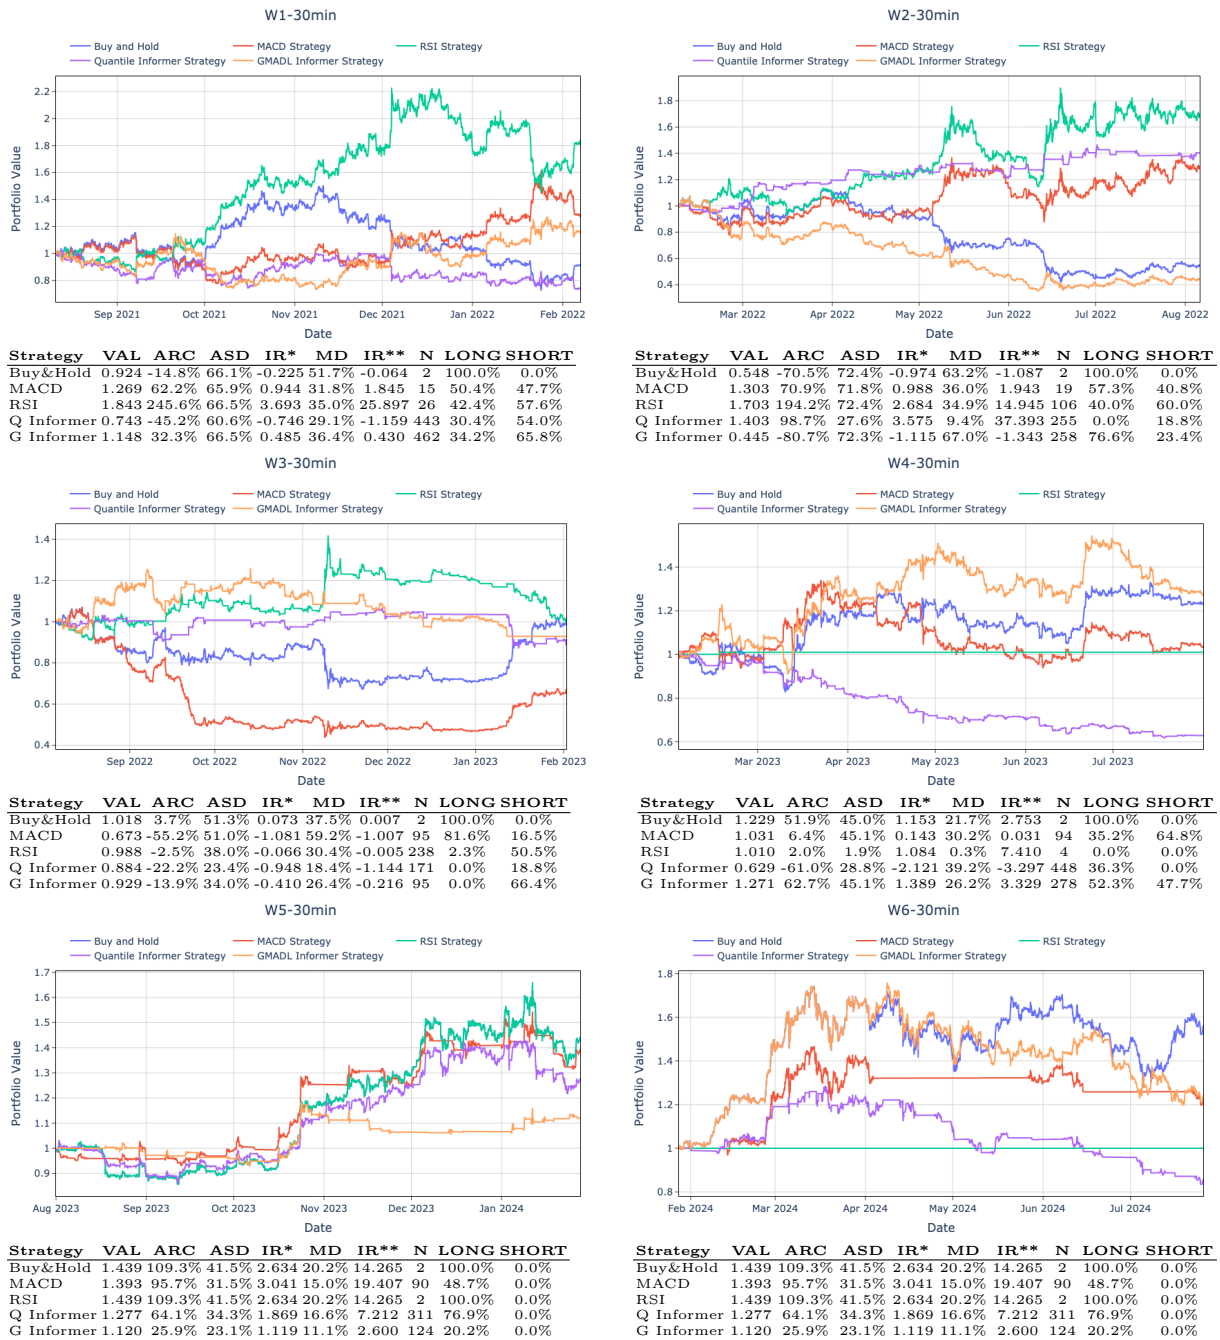
$$t = \frac{IR_{\text{strategy}}^* - IR_{\text{benchmark}}^*}{\frac{\sigma}{\sqrt{N}}}$$

The test assumes the differences in returns follow a normal distribution. The results are presented in Table 17. The null hypothesis that the Information Ratio of the top-performing strategies is not greater than the Buy&Hold Modified Information Ratio was rejected for all the tested strategies.

## 6 Sensitivity Analysis

The sensitivity analysis of the parameters of the Informer model requires retraining of the Informer model for each modified parameter and is outside the scope of this study due to the high computational cost. However, this section presents a brief analysis of how factors that impact the selection of strategy hyperparameters influence the strategy performance.

Figure 13: Evaluation results for individual windows on 30 min data



Note: Evaluation results for each of the out-of-sample data windows of 30-minute interval BTC/USDT data. The presented metrics are: portfolio value at the end of the evaluation period (VAL), Annualized Return Compound (ARC), Annualized Standard Deviation (ASD), Information Ratio (IR\*), Maximum Drawdown (MD), Modified Information Ratio (IR\*\*), number of trades (N) and percent of the long/short positions (LONG/SHORT). MACD = MACD strategy. RSI = RSI strategy. Q Informer = Quantile Informer. G Informer = GMADL Informer

## 6.1 Validation part size

The hyperparameter selection for the strategies (described in Section 5.1) was carried out using the validation part of the data windows. The size of this part was arbitrarily selected to be 0.2 of the out-of-sample data size, which, in terms of the length of the period, is roughly equal to 6 months. The length of this period influences what strategy hyperparameters are selected. If the period is too short, the data sample will not be representative enough to select meaningful values for the hyperparameters. On the other hand, if the period is too long, the strategy might not be relevant to the current market conditions. Thus, it is important to perform an analysis of how this parameter influences the strategy performance.

The analysis has been carried out only for the two best strategies, that is GMADL Informer with 5 min data and RSI Strategy with 30 min data.



Figure 14: Top 5 performing strategies



Strategy	VAL	ARC	ASD	IR*	MD	IR**	N	LONG	SHORT
Buy&Hold	1.441	13.1%	57.7%	0.228	77.3%	0.039	2	100.0%	0.0%
GMADL Informer (5min)	9.747	115.9%	54.4%	2.129	32.7%	7.552	846	44.8%	41.5%
RSI Strategy (5min)	3.341	50.3%	50.4%	0.999	30.0%	1.676	846	28.3%	33.5%
GMADL Informer (15min)	3.296	49.6%	52.7%	0.942	47.4%	0.987	362	49.4%	37.7%
RSI Strategy (30min)	4.542	66.8%	46.2%	1.444	39.9%	2.415	378	30.8%	28.0%
MACD Strategy (30min)	1.952	25.4%	52.4%	0.485	59.2%	0.207	328	52.3%	28.3%

Note: Top 5 strategies evaluated on testing period with BTC/USDT data. The presented metrics are: portfolio value at the end of the evaluation period (VAL), Annualized Return Compound (ARC), Annualized Standard Deviation (ASD), Information Ratio (IR\*), Maximum Drawdown (MD), Modified Information Ratio (IR\*\*), number of trades (N) and percent of the long/short positions (LONG/SHORT).

Table 17: Statistical t-test for comparing the performance of strategies over Buy&amp;Hold.

Strategy	N	$\sigma$	t-statistic	p-value
GMADL (5min)	311040	2.820834	375.84	0.000000***
GMADL (15 min)	103680	0.558743	408.13	0.000000***
RSI (5 min)	311040	0.648741	662.77	0.000000***
RSI (30 min)	51840	1.065079	258.49	0.000000***
MACD (30 min)	51840	0.326504	174.34	0.000000***


Note: T-test H0: The information Ratio of the strategy is not greater than the Buy&Hold information ratio. The values marked with \*\*\* indicate the p-value lower than the critical value  $\alpha = 0.01$ , rejecting the null hypothesis.

**GMADL Informer (5min):** The results for evaluating GMADL Informer on a 5-minute interval dataset when selecting hyperparameters using various validation part sizes are presented in Figure 15. For this strategy, selecting a longer or shorter validation period than 6 months resulted in decreased strategy performance. However, regardless of the length of the validation part, the strategy outperformed the benchmark, and the Modified Information Ratio stayed at a relatively high level. It can be seen that the number of trades decreases monotonically when the validation window is elongated. This suggests that the longer validation periods require the strategy to adapt to varying market conditions and more aggressively reduce the signals to change positions. The worst-performing variant was with a validation window length of three months; in this case, the strategy recorded significant losses during the last two windows of the testing period.

**RSI Strategy (30min):** Figure 16 presents the analogous analysis for the RSI strategy on 30-minute interval data. The RSI strategy seems to be more sensitive to the change in the validation part length than the GMADL strategy. Strategies that had hyperparameters selected on validation windows of length 3 and 12 months fell below the benchmark. Also, it seems that with hyperparameters selected on longer validation windows, the RSI strategy cut out too much signal and did not selecting any position during the last three windows of the testing period.




Figure 15: GMADL Informer strategies with parameters selected using different lengths of validation windows



Note: Best GMADL Informer strategy with hyperparameters selected on various lengths of validation windows. The presented metrics are: portfolio value at the end of the evaluation period (VAL), Annualized Return Compound (ARC), Annualized Standard Deviation (ASD), Information Ratio (IR\*), Maximum Drawdown (MD), Modified Information Ratio (IR\*\*), number of trades (N) and percent of the long/short positions (LONG/SHORT).

Figure 16: RSI strategies with hyperparameters selected using different lengths of validation windows




Note: Best RSI strategy with hyperparameters selected on various lengths of validation windows. The presented metrics are: portfolio value at the end of the evaluation period (VAL), Annualized Return Compound (ARC), Annualized Standard Deviation (ASD), Information Ratio (IR\*), Maximum Drawdown (MD), Modified Information Ratio (IR\*\*), number of trades (N) and percent of the long/short positions (LONG/SHORT).

## 6.2 Number of data windows

Another implicit parameter that may affect the evaluation results is the number of data windows in the period in which the evaluation is carried out. In the original experiment, six data windows, with

6-month out-of-sample testing periods were considered. This section explores how the results of the two best strategies: GMADL Informer with 5 min data and RSI Strategy with 30 min data are affected if the testing period is split differently: into three or twelve windows. Note that while selecting a different number of windows affects the length of the testing period of each window, the total testing period does not change and the in-sample (training and validation parts) length was kept constant.

Figure 17: GMADL Informer strategies evaluated on different number of windows



Strategy	VAL	ARC	ASD	IR*	MD	IR**	N	LONG	SHORT
Buy&Hold	1.44	13.1%	57.7%	0.23	77.3%	0.04	2	100.0%	0.0%
12 windows	2.06	27.6%	53.1%	0.52	53.3%	0.27	568	43.0%	37.9%
6 windows	9.75	115.9%	54.4%	2.13	32.7%	7.55	846	44.8%	41.5%
3 windows	3.08	46.2%	51.6%	0.90	45.9%	0.90	426	27.1%	45.9%

Note: Best GMADL Informer strategy evaluated on different number of windows. The presented metrics are: portfolio value at the end of the evaluation period (VAL), Annualized Return Compound (ARC), Annualized Standard Deviation (ASD), Information Ratio (IR\*), Maximum Drawdown (MD), Modified Information Ratio (IR\*\*), number of trades (N) and percent of the long/short positions (LONG/SHORT).

**GMADL Informer (5min):** Evaluation of the GMADL Informer strategy with a different number of windows required retraining of the underlying Informer model with the new data. For each of the cases considered: three and twelve windows, a new model was trained for each window, keeping the values of all the other hyperparameters the same as in the base scenario. The evaluation results are presented in Figure 18. It can be seen that the performance of the strategy was worse both when the number of windows increased and decreased. Although the Modified Information Ratio was still higher than the benchmark, it was significantly lower than the one in the six-window evaluation. This may suggest that such an excellent performance of the GMADL Informer strategy might have been an isolated event caused by the selection of a concrete number of evaluation windows. However, this conclusion should not belittle the fact that the GMADL Informer strategy, evaluated on different numbers of windows, still achieves impressive results compared to the benchmark.


**RSI Strategy (30min):** The same approach was repeated for the RSI strategy with 30-minute interval data and the results are presented in Figure 18. The RSI also outperformed the Buy&Hold benchmark in all the tested cases. However, in the case of longer testing windows, the strategy seems to be sensitive to varying market conditions and it noted a Maximum Drawdown of over 73% resulting in performance very similar to the benchmark. Shortening the testing windows seems to have less of an effect on the overall performance than with the GMADL strategy.

### 6.3 Top n-th strategy

During the selection of hyperparameters, it was implicitly assumed that the best hyperparameters for a strategy come from the strategy that performs best in the validation period. This might not be true, as the best-performing strategy might be overfitted to said period.

Figure 19, Figure 20, and Figure 21 present the portfolio values after evaluating strategies with the top 10 hyperparameter sets from validation windows. In the case of GMADL Strategy with 5 min data,

Figure 18: RSI strategies with hyperparameters selected using different lengths of validation windows



	Strategy	VAL	ARC	ASD	IR*	MD	IR**	N	LONG	SHORT
6	Buy&Hold	1.44	13.1%	55.9%	0.23	77.2%	0.04	2	100.0%	0.0%
	12 windows	3.14	47.2%	44.2%	1.07	42.6%	1.18	548	28.7%	29.7%
	6 windows	4.54	66.8%	46.2%	1.44	39.9%	2.42	378	30.8%	28.0%
	3 windows	1.73	20.3%	52.7%	0.38	73.2%	0.11	596	50.9%	33.5%

Note: Best RSI strategy with hyperparameters selected on various lengths of validation windows. The presented metrics are: portfolio value at the end of the evaluation period (VAL), Annualized Return Compound (ARC), Annualized Standard Deviation (ASD), Information Ratio (IR\*), Maximum Drawdown (MD), Modified Information Ratio (IR\*\*), number of trades (N) and percent of the long/short positions (LONG/SHORT).

the best overall performing strategy is indeed the one with the first set of hyperparameters, and the consecutive sets gradually decrease the results of the strategy. However, in the case of 15min and 30min data, the second best would turn out to be better. A similar situation can be observed with the best RSI strategy, the overall best RSI strategy with 30-minute windows was the first one from the validation, but in the case of 15-minute data, the second best was significantly better. In the case of strategies that had overall poor performance, the selection of a top n-th set of hyperparameters wouldn't have a significant impact.

Figure 19: Strategies with top 10 hyperparameter combinations for 5min data

	1	2	3	4	5	6	7	8	9	10
Buy and Hold	0.039	0.039	0.039	0.039	0.039	0.039	0.039	0.039	0.039	0.039
MACD Strategy	-0.087	-0.066	-0.02	-0.002	-0.013	-0.1	-0.201	-0.137	-0.001	-0.005
RSI Strategy	1.676	0.079	1.948	0.145	0.972	-0.019	0.022	0.549	0.878	2.125
Quantile Informer	-0.001	-0.074	-0.129	-0.016	-0.237	0.296	-0.582	0.007	-0.165	-0.064
Gmadi Informer	7.552	2.758	2.331	2.179	2.15	2.498	2.028	0.026	1.874	0.677

Note: Modified Information Ratio (IR\*\*) of the strategies evaluated on the whole testing period of 5 min data, with the strategies using the top 10 hyperparameter sets according to IR\*\* from the evaluation on the validation set.

## 7 Conclusion

The research aimed to explore different methods to create automated investment strategies for trading Bitcoin. Four strategies were proposed, two using technical indicators signal and two using predictions of the Informer machine learning model. The strategies were evaluated independently over six periods on BTC/USDT data of various frequencies: 5 minutes, 15 minutes, and 30 minutes. For each data window

Figure 20: Strategies with top 10 hyperparameter combinations for 15min data

	1	2	3	4	5	6	7	8	9	10
Buy and Hold	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04
MACD Strategy	-0.118	-0.045	-0.024	-0.1	-0.008	-0.006	0.059	0.017	0.223	-0.028
RSI Strategy	-0.014	3.996	3.524	0.002	0.839	0.97	0.568	2.489	1.888	0.093
Quantile Informer	-0.002	-0.0	-0.155	-0.188	-0.108	0.002	-0.011	-0.099	-0.098	-0.097
Gmadl Informer	0.987	2.341	0.846	3.21	1.662	0.834	4.286	2.658	1.25	3.682

Note: Modified Information Ratio(IR\*\*) of the strategies evaluated on the whole testing period of 15 min data, with the strategies using the top 10 hyperparameter sets according to IR\*\* from the evaluation on the validation set.

Figure 21: Strategies with top 10 hyperparameter combinations for 30min data

	1	2	3	4	5	6	7	8	9	10
Buy and Hold	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04
MACD Strategy	0.207	1.035	0.275	0.5	0.271	0.281	0.093	0.197	0.873	0.309
RSI Strategy	2.415	2.058	0.436	0.251	0.159	0.44	0.199	0.412	0.826	1.524
Quantile Informer	-0.104	-0.067	-0.11	-0.313	0.001	0.197	-0.022	-0.037	-0.108	-0.246
Gmadl Informer	-0.01	1.448	3.861	3.602	0.13	0.1	0.074	0.11	0.797	0.394

Note: Modified Information Ratio(IR\*\*) of the strategies evaluated on the whole testing period of 30 min data, with the strategies using the top 10 hyperparameter sets according to IR\*\* from the evaluation on the validation set.

and every frequency, the strategies that use a machine learning model were first trained using the training part of the data window. The hyperparameters were then selected using the validation part of the data window. Finally, a strategy with the best set of hyperparameters according to the Modified Information Ratio was evaluated on the test part of the data window. Five strategies beat the Buy&Hold benchmark, when compared throughout the testing period, i.e., combined results from all data windows. The best-performing strategy was the one based on predictions of the Informer model trained on 5-minute data with the GMADL loss function. In addition, it significantly outperformed all the other strategies, proving to be best in almost all testing periods. A sensitivity analysis was performed to measure the impact of the selection of hyperparameters on strategy performance. It showed that the size of the validation window proved to have an impact on the selection of hyperparameters and overall strategy performance. This was also the case when modifying the number of testing periods, changing that parameter affected results and decreased the performance of the strategy, although the GMADL Informer strategy in all cases still managed to beat the benchmark.

Based on the results of the research an attempt to answer the research questions can be made:

**Q:** *Is it possible to create an algorithmic strategy for trading Bitcoin, that is more efficient than the Buy&Hold approach?*

The research showed that it is possible to create an algorithmic strategy that outperforms the Buy&Hold benchmark on the Bitcoin data. Five strategies achieved better results during the testing period than the benchmark. The strategies were evaluated in six different test windows, through an overall testing period of almost three years.

**Q:** *Does signal from Informer model allow to create strategies that are more efficient on trading Bitcoin than strategies based on technical indicators?*

The strategy based on the Informer model trained with GMADL loss function turned out to be more efficient than the strategies based on technical indicators, however only when the higher frequency data i.e. 5 minutes and 15 minutes was used. While this may suggest that the strategies using such a model are generally more efficient, more research should be conducted to unambiguously assert the superiority of such a strategy.

**Q:** *How does selection of the machine learning model loss function influence the strategy performance?*

The strategy involving the Informer model with the GMADL loss function exhibited superior performance, whereas the approach utilizing a model with the Quantile loss function underperformed relative to a benchmark. This suggests that the GMADL loss function is better fitted for creating models that are to be used in automated trading strategies. What is more selection of the loss function seems to be one of the deciding factors when training the model that generates the signal for algorithmic trading systems.

**Q:** *Does usage of higher frequency data allow to create more efficient strategies ?*

Usage of high-frequency data seemed to improve the performance of the GMADL Informer strategy. Strategies based on technical indicators did not seem to benefit much from the higher frequencies. The RSI indicator-based strategy performed similarly on the highest and lowest of the tested intervals, the MACD indicator-based strategy was underperforming with frequencies higher than 30 minutes.

The main contribution of the study was the analysis of the novel application of loss functions: Quantile and GMADL loss, to train the Informer model to predict returns of Bitcoin. Predictions were then used to construct automated trading strategies which were compared against the Buy&Hold benchmark and the two benchmark strategies based on MACD and RSI technical indicators. The exhaustive analysis included the usage of different testing periods, different data intervals, and sensitivity analysis. It showed that the GMADL loss function allows to effectively train the machine learning model, which is able to provide meaningful signals for the trading strategy. Such a strategy can be deployed to provide higher yields than the Buy&Hold approach.

There are various directions in which the research could be extended. A more extensive sensitivity analysis could be carried out to better understand the impact of various hyperparameters on the strategies, especially in the case of Informer-based strategies and loss functions used for training the model. Another direction would be to test even higher data frequencies and explore whether the performance of the Informer-based model would be further improved. Finally, datasets of other financial instruments could be used, to see whether the same strategies perform well on other data than Bitcoin.

## References

- Adebiyi, Ayodele, Aderemi Adewumi, and Charles Ayo (Mar. 2014). "Stock price prediction using the ARIMA model". In: DOI: 10.1109/UKSim.2014.67.
- Agarap, Abien Fred (2019). *Deep Learning using Rectified Linear Units (ReLU)*. arXiv: 1803.08375 [cs.NE]. URL: <https://arxiv.org/abs/1803.08375>.
- Appel, Gerald (2005). *Technical analysis: power tools for active investors*. First. FT Press. ISBN: 0131479024.
- Azari, Amin (Oct. 2018). *Bitcoin Price Prediction: An ARIMA Approach*. DOI: 10.48550/arXiv.1904.05315.
- Ba, Jimmy Lei, Jamie Ryan Kiros, and Geoffrey E. Hinton (2016). *Layer Normalization*. arXiv: 1607.06450 [stat.ML]. URL: <https://arxiv.org/abs/1607.06450>.
- Bahdanau, Dzmitry, Kyunghyun Cho, and Yoshua Bengio (2014). *Neural Machine Translation by Jointly Learning to Align and Translate*. cite arxiv:1409.0473Comment: Accepted at ICLR 2015 as oral presentation. URL: <http://arxiv.org/abs/1409.0473>.
- Barber, Brad and Terrance Odean (2001). "Boys will be Boys: Gender, Overconfidence, and Common Stock Investment". In: *The Quarterly Journal of Economics* 116.1, pp. 261–292. URL: <https://EconPapers.repec.org/RePEc:oup:qjecon:v:116:y:2001:i:1:p:261-292..>
- Barez, Fazl et al. (2023). "Exploring the Advantages of Transformers for High-Frequency Trading". In: *Social Science Research Network*. DOI: 10.2139/ssrn.4364833.
- Basodi, Sunitha et al. (2020). "Gradient amplification: An efficient way to train deep neural networks". In: *Big Data Mining and Analytics* 3.3, pp. 196–207. DOI: 10.26599/BDMA.2020.9020004.
- Benidis, Konstantinos et al. (July 2023). "Deep Learning for Time Series Forecasting: Tutorial and Literature Survey". en. In: *ACM Computing Surveys* 55.6, pp. 1–36. ISSN: 0360-0300, 1557-7341. DOI: 10.1145/3533382. URL: <https://dl.acm.org/doi/10.1145/3533382> (visited on 03/29/2023).

- Billah, Baki et al. (Feb. 2006). "Exponential smoothing model selection for forecasting". In: *International Journal of Forecasting* 22, pp. 239–247. DOI: 10.1016/j.ijforecast.2005.08.002.
- Bollinger, J. (2002). *Bollinger on Bollinger Bands*. McGraw-Hill Education. ISBN: 9780071373685. URL: <https://books.google.pl/books?id=FLlxAz85iysC>.
- Bondt, Werner F. M. De and Richard Thaler (1985). "Does the Stock Market Overreact?" In: *The Journal of Finance* 40.3, pp. 793–805. ISSN: 00221082, 15406261. URL: <http://www.jstor.org/stable/2327804> (visited on 10/05/2024).
- Box, George.E.P. and Gwilym M. Jenkins (1976). *Time Series Analysis: Forecasting and Control*. Holden-Day.
- Chen, Kai et al. (2015). "A LSTM-based method for stock returns prediction: A case study of China stock market". In: *null*. DOI: 10.1109/bigdata.2015.7364089.
- Cho, Kyunghyun et al. (2014). "Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation". In: *CoRR* abs/1406.1078. arXiv: 1406.1078. URL: <http://arxiv.org/abs/1406.1078>.
- Chollet, Francois (July 2017). "Xception: Deep Learning with Depthwise Separable Convolutions". In: pp. 1800–1807. DOI: 10.1109/CVPR.2017.195.
- Clevert, Djork-Arné, Thomas Unterthiner, and Sepp Hochreiter (2016). *Fast and Accurate Deep Network Learning by Exponential Linear Units (ELUs)*. arXiv: 1511.07289 [cs.LG]. URL: <https://arxiv.org/abs/1511.07289>.
- Devlin, Jacob et al. (2018). "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding". In: *CoRR* abs/1810.04805. arXiv: 1810.04805. URL: <http://arxiv.org/abs/1810.04805>.
- Ding, Ying, Changsheng Zhang, and Chen Zhang (2023). "An Informer Based Method for Stock Intraday Price Prediction". In: *2023 19th International Conference on Natural Computation, Fuzzy Systems and Knowledge Discovery (ICNC-FSKD)*, pp. 1–6. DOI: 10.1109/ICNC-FSKD59587.2023.10280785.
- Duan, Chuyang and Wenjun Ke (2024). "Advanced Stock Price Prediction Using LSTM and Informer Models". In: *Journal of Artificial Intelligence General science (JAIGS) ISSN:3006-4023*. DOI: 10.60087/jaigs.v5i1.183.
- Eisenach, Carson, Yagna Patel, and Dhruv Madeka (2020). "MQTransformer: Multi-Horizon Forecasts with Context Dependent and Feedback-Aware Attention". In: *CoRR* abs/2009.14799. arXiv: 2009.14799. URL: <https://arxiv.org/abs/2009.14799>.
- Fama, Eugene (1970). "Efficient Capital Markets: A Review of Theory and Empirical Work". In: *Journal of Finance* 25, pp. 383–417.
- (Dec. 1991). "Efficient Capital Markets: II". In: *Journal of Finance* 46.5, pp. 1575–1617. URL: <https://ideas.repec.org/a/bla/jfinan/v46y1991i5p1575-617.html>.
- Fischer, Thomas et al. (2017). "Deep learning with long short-term memory networks for financial market predictions". In: *European Journal of Operational Research*. DOI: 10.1016/j.ejor.2017.11.054.
- Fischhoff, Baruch and Paul Slovic (June 1978). "A Little Learning...: Confidence in Multicue Judgment Tasks". In: p. 58.
- Gervais, Simon and Terrance Odean (Feb. 2001). "Learning To Be Overconfident". In: *Review of Financial Studies* 14, pp. 1–27. DOI: 10.2139/ssrn.36313.
- Grudniewicz, Jan and Robert Ślepaczuk (2023). "Application of machine learning in algorithmic investment strategies on global stock markets". In: *Research in International Business and Finance* 66, p. 102052. ISSN: 0275-5319. DOI: <https://doi.org/10.1016/j.ribaf.2023.102052>. URL: <https://www.sciencedirect.com/science/article/pii/S0275531923001782>.
- Hájek, Petr and Josef Novotný (2024). "Beyond Sentiment in Stock Price Prediction: Integrating News Sentiment and Investor Attention with Temporal Fusion Transformer". In: *Artificial Intelligence Applications and Innovations*. DOI: 10.1007/978-3-031-63219-8\_3.
- Hochreiter, Sepp and Jürgen Schmidhuber (Nov. 1997). "Long Short-Term Memory". In: *Neural Comput.* 9.8, pp. 1735–1780. ISSN: 0899-7667. DOI: 10.1162/neco.1997.9.8.1735. URL: <https://doi.org/10.1162/neco.1997.9.8.1735>.
- Honchar, Oleksandr and Luca Di Persio (Jan. 2016). "Artificial neural networks approach to the forecast of stock market price movements". In: Vol.1, pp. 158–162.
- Hossain, Mohammad et al. (Nov. 2018). "Hybrid Deep Learning Model for Stock Price Prediction". In: pp. 1837–1844. DOI: 10.1109/SSCI.2018.8628641.
- Hu, Xiaokang and Xiaokang Hu (2021). "Stock Price Prediction Based on Temporal Fusion Transformer". In: *2021 3rd International Conference on Machine Learning, Big Data and Business Intelligence (MLDBI)*. DOI: 10.1109/mlbdi54094.2021.00019.



- Huberman, Gur and Tomer Regev (2001). “Contagious Speculation and a Cure for Cancer: A Nonevent that Made Stock Prices Soar”. In: *Journal of Finance* 56.1, pp. 387–396. URL: <https://EconPapers.repec.org/RePEc:bla:jfinan:v:56:y:2001:i:1:p:387-396>.
- Kahneman, Daniel and Amos Tversky (Mar. 1979). “Prospect Theory: An Analysis of Decision under Risk”. In: *Econometrica* 47.2, pp. 263–291. URL: <https://ideas.repec.org/a/ecm/emetrp/v47y1979i2p263-91.html>.
- Khashei, Mehdi and Mehdi Bijari (2011). “A novel hybridization of artificial neural networks and ARIMA models for time series forecasting”. In: *Applied Soft Computing* 11.2. The Impact of Soft Computing for the Progress of Artificial Intelligence, pp. 2664–2675. ISSN: 1568-4946. DOI: <https://doi.org/10.1016/j.asoc.2010.10.015>. URL: <https://www.sciencedirect.com/science/article/pii/S1568494610002759>.
- Kijewski, Mateusz, Robert Slepaczuk, and Maciej Wysocki (Sept. 2024). “Predicting Prices Of S&P 500 Index Using Classical Methods and Recurrent Neural Networks”. In: *Proceedings of the 32nd International Conference on Information Systems Development (ISD 2024)*. DOI: 10.62036/ISD.2024.89.
- Kim, Jae H., Abul Shamsuddin, and Kian-Ping Lim (2011). “Stock return predictability and the adaptive markets hypothesis: Evidence from century-long U.S. data”. In: *Journal of Empirical Finance* 18.5, pp. 868–879. DOI: 10.1016/j.jempfin.2011.08. URL: <https://ideas.repec.org/a/eee/empfin/v18y2011i5p868-879.html>.
- Kiranyaz, Serkan et al. (2019). *1D Convolutional Neural Networks and Applications: A Survey*. arXiv: 1905.03554 [eess.SP]. URL: <https://arxiv.org/abs/1905.03554>.
- Kolen, John F. and Stefan C. Kremer (2001). “Gradient Flow in Recurrent Nets: The Difficulty of Learning LongTerm Dependencies”. In: *A Field Guide to Dynamical Recurrent Networks*, pp. 237–243. DOI: 10.1109/9780470544037.ch14.
- Kosc, Krzysztof, Paweł Sakowski, and Robert Ślepaczuk (2019). “Momentum and contrarian effects on the cryptocurrency market”. In: *Physica A: Statistical Mechanics and its Applications* 523, pp. 691–701. ISSN: 0378-4371. DOI: <https://doi.org/10.1016/j.physa.2019.02.057>. URL: <https://www.sciencedirect.com/science/article/pii/S037843711930216X>.
- LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton (May 2015). “Deep learning”. en. In: *Nature* 521.7553, pp. 436–444. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/nature14539. URL: <https://www.nature.com/articles/nature14539> (visited on 10/06/2024).
- Lim, Bryan, Sercan O. Arik, et al. (2020). *Temporal Fusion Transformers for Interpretable Multi-horizon Time Series Forecasting*. arXiv: 1912.09363 [stat.ML]. URL: <https://arxiv.org/abs/1912.09363>.
- Lim, Kian-Ping and Robert Brooks (Feb. 2011). “The Evolution Of Stock Market Efficiency Over Time: A Survey Of The Empirical Literature”. In: *Journal of Economic Surveys* 25.1, pp. 69–108. URL: <https://ideas.repec.org/a/bla/jecsur/v25y2011i1p69-108.html>.
- Lu, Yuze, Hailong Zhang, and Qiwen Guo (2023). *Stock and market index prediction using Informer network*. arXiv: 2305.14382 [q-fin.ST]. URL: <https://arxiv.org/abs/2305.14382>.
- Luong, Thang, Hieu Pham, and Christopher D. Manning (Sept. 2015). “Effective Approaches to Attention-based Neural Machine Translation”. In: *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. Ed. by Lluís Màrquez, Chris Callison-Burch, and Jian Su. Lisbon, Portugal: Association for Computational Linguistics, pp. 1412–1421. DOI: 10.18653/v1/D15-1166. URL: <https://aclanthology.org/D15-1166>.
- Malkiel, Burton (Feb. 2005). “Reflections on the Efficient Market Hypothesis: 30 Years Later”. In: *The Financial Review* 40, pp. 1–9. DOI: 10.1111/j.0732-8516.2005.00090.x.
- Malkiel, Burton. G. (1973). *A Random Walk Down Wall Street*. Norton, New York.
- Michańków, Jakub, Paweł Sakowski, and Robert Ślepaczuk (2022). “LSTM in Algorithmic Investment Strategies on BTC and S&P500 Index”. In: *Sensors* 22.3. ISSN: 1424-8220. DOI: 10.3390/s22030917. URL: <https://www.mdpi.com/1424-8220/22/3/917>.
- (2024). “Mean Absolute Directional Loss as a new loss function for machine learning problems in algorithmic investment strategies”. In: *Journal of Computational Science* 81, p. 102375. ISSN: 1877-7503. DOI: <https://doi.org/10.1016/j.jocs.2024.102375>. URL: <https://www.sciencedirect.com/science/article/pii/S1877750324001686>.
- Nagi, Jawad et al. (2011). “Max-pooling convolutional neural networks for vision-based hand gesture recognition”. In: *2011 IEEE International Conference on Signal and Image Processing Applications (ICSIPA)*, pp. 342–347. DOI: 10.1109/ICSIPA.2011.6144164.
- Nakamoto, Satoshi (Mar. 2009). “Bitcoin: A Peer-to-Peer Electronic Cash System”. In: *Cryptography Mailing list at https://metzdowd.com*.



- Nelson, David M. et al. (2017). "Stock market's price movement prediction with LSTM neural networks". In: DOI: 10.1109/ijcnn.2017.7966019.
- Odean., Terrance (Nov. 1996). *Are Investors Reluctant to Realize Their Losses?* Research Program in Finance Working Papers RPF-269. University of California at Berkeley. URL: <https://ideas.repec.org/p/ucb/calbrf/rpf-269.html>.
- OpenAI et al. (2024). *GPT-4 Technical Report*. arXiv: 2303.08774 [cs.CL]. URL: <https://arxiv.org/abs/2303.08774>.
- Pai, Ping-Feng and Chih-Sheng Lin (2005). "A hybrid ARIMA and support vector machines model in stock price forecasting". In: *Omega* 33.6, pp. 497–505. ISSN: 0305-0483. DOI: <https://doi.org/10.1016/j.omega.2004.07.024>. URL: <https://www.sciencedirect.com/science/article/pii/S0305048304001082>.
- Penmetsa, Siddharth and Maruthi Vemula (2023). "Cryptocurrency Price Prediction with LSTM and Transformer Models Leveraging Momentum and Volatility Technical Indicators". In: *2023 IEEE 3rd International Conference on Data Science and Computer Application (ICDSCA)*. DOI: 10.1109/icdsca59871.2023.10393319.
- Persio, Luca Di et al. (2017). "Recurrent Neural Networks Approach to the Financial Forecast of Google Assets". In: *null*. DOI: null.
- Salinas, David, Valentin Flunkert, and Jan Gasthaus (2019). *DeepAR: Probabilistic Forecasting with Autoregressive Recurrent Networks*. arXiv: 1704.04110 [cs.AI]. URL: <https://arxiv.org/abs/1704.04110>.
- Selvin, Sreelekshmy et al. (2017). "Stock price prediction using LSTM, RNN and CNN-sliding window model". In: *null*. DOI: 10.1109/icacci.2017.8126078.
- Sharpe, William F. (1998). "The Sharpe Ratio (Fall 1994)". In: *The Best of The Journal of Portfolio Management*. Ed. by Peter L. Bernstein and Frank J. Fabozzi. Princeton: Princeton University Press, pp. 169–178. ISBN: 9781400829408. DOI: doi:10.1515/9781400829408-022. URL: <https://doi.org/10.1515/9781400829408-022>.
- Siarni-Namini, Sima et al. (2018). "Forecasting Economics and Financial Time Series: ARIMA vs. LSTM." In: *arXiv: Learning*. DOI: null.
- Team, Gemini et al. (2024). *Gemini: A Family of Highly Capable Multimodal Models*. arXiv: 2312.11805 [cs.CL]. URL: <https://arxiv.org/abs/2312.11805>.
- Tversky, Amos and Daniel Kahneman (1981). "The Framing of Decisions and the Psychology of Choice". In: *Science* 211.4481, pp. 453–458. DOI: 10.1126/science.7455683. eprint: <https://www.science.org/doi/pdf/10.1126/science.7455683>. URL: <https://www.science.org/doi/abs/10.1126/science.7455683>.
- Vaserstein, L.N. (1969). "Markov processes over denumerable products of spaces, describing large systems of automata". In: *Probl. Peredachi Inf.* 5.
- Vaswani, Ashish et al. (2017). "Attention Is All You Need". In: *CoRR* abs/1706.03762. arXiv: 1706.03762. URL: <http://arxiv.org/abs/1706.03762>.
- Võ, Thi Ái Nguyễn and Robert Ślepaczuk (Jan. 2022). "Applying Hybrid ARIMA-SGARCH in Algorithmic Investment Strategies on S&P500 Index". In: *Entropy* 24, p. 158. DOI: 10.3390/e24020158.
- Wang, Ju-Jie et al. (2012). "Stock index forecasting based on a hybrid model". In: *Omega* 40.6. Special Issue on Forecasting in Management Science, pp. 758–766. ISSN: 0305-0483. DOI: <https://doi.org/10.1016/j.omega.2011.07.008>. URL: <https://www.sciencedirect.com/science/article/pii/S0305048311001435>.
- Wang, S. (2023). "A Stock Price Prediction Method Based on BiLSTM and Improved Transformer". In: *IEEE Access*. DOI: 10.1109/access.2023.3296308.
- Wilder, J.W. (1978). *New Concepts in Technical Trading Systems*. Trend Research. ISBN: 9780894590276. URL: <https://books.google.pl/books?id=WesJAQAAMAAJ>.
- Yadav, Y. (Nov. 2015). "How algorithmic trading undermines efficiency in capital markets". In: 68, pp. 1607–1671.
- Zhang, G.Peter (2003). "Time series forecasting using a hybrid ARIMA and neural network model". In: *Neurocomputing* 50, pp. 159–175. ISSN: 0925-2312. DOI: [https://doi.org/10.1016/S0925-2312\(01\)00702-0](https://doi.org/10.1016/S0925-2312(01)00702-0). URL: <https://www.sciencedirect.com/science/article/pii/S0925231201007020>.
- Zhao, Huali, Martin Crane, and Marija Bezbradica (2022). "Attention! Transformer with Sentiment on Cryptocurrencies Price Prediction". In: *International Conference on Complex Information Systems*. DOI: 10.5220/0011103400003197.
- Zhou, Haoyi et al. (2021). *Informer: Beyond Efficient Transformer for Long Sequence Time-Series Forecasting*. arXiv: 2012.07436 [cs.LG]. URL: <https://arxiv.org/abs/2012.07436>.



UNIVERSITY OF WARSAW  
FACULTY OF ECONOMIC SCIENCES  
44/50 DŁUGA ST.  
00-241 WARSAW  
[WWW.WNE.UW.EDU.PL](http://WWW.WNE.UW.EDU.PL)  
ISSN 2957-0506