

Received 31 January 2024, accepted 17 February 2024, date of publication 20 February 2024, date of current version 29 February 2024.

Digital Object Identifier 10.1109/ACCESS.2024.3367719

## RESEARCH ARTICLE

# Enhanced Bitcoin Price Direction Forecasting With DQN

AZAMJON MUMINOV<sup>1</sup>, OTABEK SATTAROV<sup>2</sup>, AND DAEYOUNG NA<sup>3</sup>, (Member, IEEE)

<sup>1</sup>Department of Computer Engineering, Gachon University, Seongnam 13120, South Korea

<sup>2</sup>School of Computing, Gachon University, Seongnam 13120, South Korea

<sup>3</sup>Global Leadership School, Handong Global University, Pohang 37554, South Korea

Corresponding author: Daeyoung Na (soriru@handong.edu)

This work was supported in part by the Basic Science Research Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Education under Grant 2019R1F1A105775713, and in part by the Gachon University Research Fund of 2020 under Grant GCU-202008460007.

**ABSTRACT** In the Bitcoin trading landscape, predicting price movements is paramount. Our study focuses on identifying the key factors influencing these price fluctuations. Utilizing the Pearson correlation method, we extract essential data points from a comprehensive set of 14 data features. We consider historical Bitcoin prices, representing past market behavior; trading volumes, which highlight the level of trading activity; network metrics that provide insights into Bitcoin's blockchain operations; and social indicators: analyzed sentiments from Twitter, tracked Bitcoin-related search trends on Google and on Twitter. These social indicators give us a more nuanced understanding of the digital community's sentiment and interest levels. With this curated data, we forge ahead in developing a predictive model using Deep Q-Network (DQN). A defining aspect of our model is its innovative reward function, tailored for enhancing predicting Bitcoin price direction, distinguished by its multi-faceted reward function. This function is a blend of several critical factors: it rewards prediction accuracy, incorporates confidence scaling, applies an escalating penalty for consecutive incorrect predictions, and includes a time-based discounting to prioritize recent market trends. This composite approach ensures that the model's performance is not only precise in its immediate predictions but also adaptable and responsive to the evolving patterns of the cryptocurrency market. Notably, in our tests, our model achieved an impressive F1-score of 95%, offering substantial promise for traders and investors.

**INDEX TERMS** Bitcoin, reinforcement learning, deep Q-network, Pearson correlation, reward function.

## I. INTRODUCTION

Introduced to the world in 2009 by an anonymous person known as Satoshi Nakamoto [1], Bitcoin was a novel concept designed to operate as a decentralized digital currency without the need for a central bank or single administrator [2]. This revolutionary concept transformed from just an academic idea into a global financial phenomenon over the years. Initially, Bitcoin primarily intrigued cryptographic enthusiasts and those skeptical of centralized financial systems [3]. However, its decentralized nature, finite supply, and the potential for peer-to-peer transactions without intermediaries

began capturing a broader audience's attention [4]. By the end of the 2010s, Bitcoin had firmly established itself as the world's leading cryptocurrency by market capitalization, with significant volatility marking its ascent [5]. Understanding this volatility is essential as predicting Bitcoin's price direction is vital for traders and investors to make informed decisions [6]. In particular, accurately predicting its price direction can lead to substantial profits, while inaccurate predictions can lead to equally significant losses. Thus, the emphasis on predicting its movement is not merely academic but has financial implications for a wide group of stakeholders, from individual retail traders to institutional investors [7]. Given the multiple unique factors influencing Bitcoin's price, this task becomes notably challenging.

The associate editor coordinating the review of this manuscript and approving it for publication was Thomas Canhao Xu<sup>id</sup>.

Obtaining a precise prediction is important, and in achieving this precision, two primary considerations emerge.

**Feature selection.** It's crucial to select the features that most influence Bitcoin's price direction. Bitcoin's price is governed by a mixture of factors, each adding a layer of complexity to its prediction.

One of the primary determinants is its own price and volume data. The historical Bitcoin prices, which include opening, closing, highest, and lowest figures, give insights into its past behavior, potentially indicating future trends. Concurrently, the volume of trades provides a perspective on the currency's liquidity and the level of its trading activity, acting as a potential marker of market sentiment [8].

On-chain metrics provide a detailed view of the activities within the Bitcoin ecosystem. The hash rate, reflecting the computational power of the Bitcoin network, indicates its security and the overall mining effort [9]. Similarly, factors like the number of transactions, average transaction value, and the count of active addresses show the currency's adoption rate and its usage patterns. Data regarding exchange inflows and outflows, miner outflows, stablecoin inflows, and options market data are vital. These metrics highlight supply and demand dynamics, the decisions of miners, and how Bitcoin interacts with other cryptocurrencies and derivatives markets.

Social indicators like Twitter sentiments, as well as Google and Twitter trends data for Bitcoin, are also significant. The shifts in public sentiment, measured on these platforms, can provide clues about potential price movements. For instance, a rise in positive sentiment might hint at an upcoming increase in Bitcoin's price and vice-versa [10].

However, even with these rich data sources, it's essential to choose the right features for analysis. Not all data points have an equal impact on Bitcoin's price. The importance of feature selection lies in ensuring efficiency and accuracy. By using techniques like the Pearson correlation, researchers can identify which of these metrics most strongly correlates with Bitcoin's price direction, eliminating unnecessary data and focusing on the most relevant predictors [11].

**Algorithm selection.** It's important to select the algorithm used to build the model that generates the prediction. The rapidly growing interest in predicting Bitcoin's price direction has led to the exploration of various algorithms and methodologies. Classical approaches have typically involved time series analysis, linear regression models, and even some traditional machine learning techniques like support vector machines and decision trees [12]. In addition, recently, Transfer Learning has emerged as an addressing the one of the main limitations of these approaches, where the availability of large, annotated datasets may be limited, and the market conditions are rapidly evolving [13]. However, the inherent complexity and multifaceted nature of Bitcoin's price dynamics, influenced by multiple factors, have often stretched these traditional methods to their limits.

Among the vast algorithmic landscape, the Deep Q-Network (DQN) has emerged as a particularly promising candidate [15], [16]. Unlike other algorithms, DQN, a type of deep reinforcement learning, is adept at managing high-dimensional input spaces, making it ideal for handling the rich set of features associated with Bitcoin [17]. The DQN combines the strength of Q-learning with deep neural networks. Q-learning, a model-free reinforcement learning algorithm, aims to find an optimal action-selection policy for a given finite Markov decision process. When placed against the backdrop of Bitcoin's price prediction, this means determining actions (or predictions) that maximize the expected value of total rewards. With the comprehensive historical data of Bitcoin and its intricate inter-feature relationships, the state space becomes vast, presenting challenges to traditional Q-learning algorithms. This complexity is where deep neural networks excel. Deep neural networks, celebrated for recognizing patterns in substantial datasets when engaged with Q-learning in DQNs, allow the algorithm to generalize over the state space. This capability makes DQNs proficient at capturing nuanced patterns among features, translating to efficient predictions of Bitcoin's price direction. Furthermore, the adaptive nature of DQN is invaluable in the context of cryptocurrency markets, known for their volatility. Patterns evolve, and DQNs can adjust to new data, refining their predictions as newer information streams in, making them inherently appropriate for such dynamic ecosystems [18].

In this study, we leveraged the capabilities of DQN to offer a rigorous approach to predicting Bitcoin's price direction changes. Our main contributions are:

- (a) First, we assembled a dataset, collected on an hourly basis, that captures various factors of the Bitcoin ecosystem. This dataset encompasses traditional price and volume data, crucial on-chain metrics reflective of Bitcoin's underlying network activities, and social indicators drawn from platforms like Twitter and Google Trends. Each feature carries its own significance: while historical price points provide immediate context, on-chain metrics elucidate the Bitcoin network's health, and social indicators offer insights into the collective sentiment of the wider digital community.
- (b) Second, to ensure the efficacy of our model, we recognized the necessity to prioritize these features. Utilizing the Pearson correlation method, we discerned and ranked the most influential features in terms of their effect on Bitcoin's price direction changes. This step ensured that our model was equipped with the most pertinent information, streamlining its predictive capability.
- (c) Finally, a key contribution of our work is the development and integration of a novel reward function within the DQN model, specifically designed to enhance the prediction of Bitcoin price movements. This reward function, characterized by its multifaceted structure, not

**TABLE 1.** Taxonomy of related works.

Data Sources	Methods
Price and volume data	Grey prediction model [24], LSTM and ADAM optimization [28], ARIMA-GARCH model [29], DRL model [30], Hybrid LSTAR-GARCH model [31], Grey Model First Order One Variable- GM(1,1) [32], Moving averages [33], ARIMA model [25], <b>DQN [This paper]</b> .
On-chain metric data	Prophet model [34], Q-learning [35], Synthesizer Transformer model [36], <b>DQN [This paper]</b> .
Social indicator data	Random Forest based-model [37], hybrid model [38], Q-learning [39], 27 ML models [40], multimodal model [41], Granger causality-in-quantile approach with wavelet analysis [42], XGBoost classifier [43], <b>DQN [This paper]</b> .

only reinforces accurate predictions but also incorporates elements like confidence scaling, penalties for consecutive inaccuracies, and time-based relevance, addressing the three potential outcomes: an increase, a decrease, or no-change in Bitcoin's price. The synergy between the proposed reward mechanism and our model's advanced architecture, augmented by carefully selected features, has yielded a notable prediction F1-score of 95%, an indication of the efficacy and robustness of our proposed approach in financial market forecasting.

Speaking in short, we recognize the importance of Bitcoin and its unpredictable nature in the financial world. As the leading cryptocurrency, Bitcoin's price movements are a subject of significant interest for traders, investors, and researchers alike. However, the volatile and multifaceted nature of the cryptocurrency market presents unique challenges in price prediction. Our motivation is to harness and analyze various data sources – ranging from historical prices and trading volumes to social media sentiments – to develop an advanced predictive model using DQN. This endeavor is driven by the objective of providing more accurate and reliable price direction predictions, crucial for informed trading and investment decisions in the dynamic cryptocurrency market. By addressing this research gap, our study aims to contribute significantly to the field of financial analytics and cryptocurrency market prediction.

The remainder of this paper is organized as follows: In Section II, we delve into related work, setting the context for our study and highlighting the gaps our research aims to fill. Section III provides a detailed account of how we gathered data from various sources, emphasizing the importance and relevance of each type of data in predicting Bitcoin's price direction. In Section IV, we furnish the background of the DQN model algorithm, tracing its roots from reinforcement learning, elaborating on the specifics of DQN, and detailing its implementation for our prediction task. Our experimental results, which underline the efficacy of our approach, are presented in Section V. Section VI offers a discussion, where we dissect the findings, draw comparisons with existing methodologies, and muse on the implications of our results. Finally, Section VII concludes the paper, summarizing our

contributions and pointing toward potential avenues for future research.

## II. RELATED WORKS

The drastic rise of Bitcoin and its subsequent integration into mainstream financial markets has intensified the quest to accurately predict its price direction. As the world's premier cryptocurrency, understanding the factors that influence Bitcoin's price has become paramount for investors, traders, and researchers alike. Early studies primarily relied on historical price data, but with the evolution of the digital currency landscape, multiple data sources have been employed to enhance prediction accuracy. For instance, recent research has explored the use of off-chain data such as cryptocurrency price streams for Ethereum gas price prediction [19].

Machine learning has emerged as a dominant force in this endeavor, with studies emphasizing its potential in forecasting Bitcoin's price movements [20]. The intricate dynamics of Bitcoin's price have also led researchers to explore advanced techniques such as the centralized decomposition approach in LSTM [21] and neural network-based forecasting [22]. These methodologies underscore the significance of employing sophisticated algorithms to navigate the volatile terrains of cryptocurrency markets.

However, the prediction landscape is not limited to Bitcoin alone. The cryptocurrency ecosystem is vast, and the interplay between different digital currencies can offer valuable insights. A study highlights the importance of considering multiple cryptocurrencies and their high-frequency trends for more holistic and accurate predictions [23]. As we delve deeper into the relevant literature, we will uncover the nuances, methodologies, and algorithms that have shaped the current state of Bitcoin price prediction. Table 1 provides a summary of these related studies, detailing the algorithms and methods used in each. This overview helps to place our research within the broader context of existing literature in this field.

### A. HISTORICAL EVOLUTION OF BITCOIN PRICE PREDICTION

The historical trajectory of Bitcoin price prediction has witnessed a fascinating evolution, marked by the integration

of diverse data sources and the application of increasingly sophisticated algorithms. Early efforts to predict Bitcoin's price direction were primarily rooted in analyzing historical price data, leveraging traditional statistical methods to discern patterns and potential future movements. A study exemplifies the utilization of historical prices to build a grey forecasting model to predict the price trend of volatile assets, such as Bitcoin [24]. As the field matured, researchers began to explore more nuanced and multifaceted approaches to predict Bitcoin prices. A notable shift was observed towards the application of machine learning algorithms to enhance prediction accuracy, as evidenced by a study that employed machine learning for Bitcoin price prediction [25]. Furthermore, the complex and volatile nature of Bitcoin prices during various historical crises, such as the COVID-19 pandemic, has been dissected using models that account for long-term memory or long-range dependence in time-series data [26]. The evolution of Bitcoin price prediction has not only been confined to the refinement of algorithms but also to the diversification of data sources and methodologies adopted. Some study underscores the application of machine learning algorithms in predicting cryptocurrency prices and improving trading strategies [27]. As we delve deeper into the major milestones and significant breakthroughs in the field, we will explore the various data sources, algorithms, and methodologies that have been pivotal in shaping the research landscape of Bitcoin price prediction.

## B. DATA SOURCES IN BITCOIN PRICE PREDICTION

The exploration into Bitcoin price prediction has witnessed a rich tapestry of methodologies, each employing various data sources to enhance the precision and reliability of their forecasts. Historically, researchers have leaned heavily on price and volume data, utilizing historical prices and trade volumes to discern patterns and potential future movements [28], [29], [30], [31]. A notable approach, as presented in a study [32], leverages the GM (1, 1) model to predict the closing prices of gold and Bitcoin, utilizing historical price data spanning from November 9, 2016, to October 9, 2021. The study underscores the potential of grey models in predicting cryptocurrency price fluctuations and informing trading decisions, thereby highlighting the important role of historical data in shaping predictive models. Complementing this, another research intertwines technical analysis with historical price data, exploring the correlation between historical price movements and future trends in the cryptocurrency market [33]. The study carefully investigates the efficacy of employing technical analysis strategies, providing a nuanced perspective on the symbiotic relationship between historical price movements and future cryptocurrency valuation, and thereby, extending the discourse on the multifaceted influences impacting Bitcoin prices. Further enriching this discourse, a study [25] delves into the application of a quantifier trading system, focusing on time series data to predict the prices of gold and Bitcoin.

The research, which also utilizes a five-year span of historical data, employs a time series prediction model, thereby providing valuable insights into the applicability of time series models in predicting cryptocurrency prices. This study, in conjunction with the aforementioned research endeavors, collectively illuminates the diverse methodologies and data sources employed in the realm of Bitcoin price prediction, each contributing uniquely to the overarching narrative and understanding of cryptocurrency valuation dynamics.

However, a recent study, underscores the significance of employing other data sources, including on-chain data, to predict cryptocurrency trends, thereby highlighting the evolution in data utilization for price prediction [34]. The study underscores the significance of on-chain metrics, such as hash rate and transaction numbers, in predicting cryptocurrency prices, particularly Bitcoin. The research provides valuable insights into the applicability of on-chain metrics in navigating the volatile cryptocurrency market, thereby contributing to the broader narrative of utilizing intrinsic blockchain data for predictive modeling in the financial domain. The utilization of on-chain data for predicting the direction of Bitcoin's price has demonstrated its efficacy in our previous research [35]. In that study, we employed the Q-learning algorithm to construct a predictive model, which considered four types of on-chain data: exchange inflows and outflows, miner outflows, stablecoin inflows, and options market data. The experimental results indicated that the Q-model successfully predicted one of three price changes: increase, no-change, or decrease, with an accuracy exceeding 85%. Similarly, research done by Herremans et.al. [36] also investigated the role of on-chain data in predicting Bitcoin's next-day volatility, with a focus on extreme volatility spikes. They proposed a deep learning Synthesizer Transformer model and reported the model's outperforming results over existing state-of-the-art models.

In the prediction of the Bitcoin price movements, it has become important to explore social indicators, particularly sentiments derived from platforms like Twitter and Google Trends. Therefore, there have been several studies that considered finding the relationship between Bitcoin price and public opinions [37], [38], [39]. For example, Naman et.al [40] utilized sentiment analysis and technical indicators to predict whether the price change would be bearish or bullish, demonstrating the potential of social media sentiments in forecasting cryptocurrency market movements. Another study [41] employed a multimodal model with Twitter FinBERT embeddings to predict extreme price movements of Bitcoin, highlighting the significance of utilizing social media data, especially Twitter, in predicting cryptocurrency price fluctuations. Furthermore, the researcher Remzi et.al. [42] explored the relationship between Twitter-based economic uncertainty and Bitcoin returns, providing insights into the impact of social media-based economic indicators on cryptocurrency valuation. Moreover, another study has utilized sentiment analysis in social media and considered the influence of celebrities, along with data mining techniques,

to predict Bitcoin price changes, thereby offering a unique perspective on the multifaceted influences impacting Bitcoin prices [43]. These studies collectively underscore the growing relevance and applicability of social indicators, especially sentiments derived from social media platforms, in predicting Bitcoin price movements, thereby contributing to the burgeoning field of cryptocurrency research and offering valuable insights for future studies.

### C. CHALLENGES AND LIMITATIONS OF DATA SOURCES IN BITCOIN PRICE PREDICTION

In the Bitcoin price prediction task, the choice of data sources is paramount, but these sources often come with inherent challenges and limitations. When considering price and volume data, one immediately confronts issues such as potential incompleteness in historical datasets. Missing data points can introduce inaccuracies or biases, skewing our predictive outcomes. Furthermore, the data is often rife with noise, as short-lived, non-recurring events temporarily distort prices and volumes. Such noise might drown out more significant, long-term trends. Another challenge lies in the granularity of the data. Depending on how frequently data points are recorded—be it hourly or daily—certain important trends might either be missed or be overemphasized.

On-chain metrics, while providing a direct lens into the Bitcoin ecosystem, pose their own set of challenges. For instance, quantitative data such as the hash rate or transaction volume might be straightforward to compute, but their interpretation vis-a-vis price prediction remains non-trivial. A notable lag can often exist between a change in an on-chain metric and its market effect, making real-time predictions tricky. Furthermore, these metrics sometimes risk giving undue importance to the behaviors of large entities, like major miners, thereby not accurately capturing the sentiment of the broader Bitcoin community.

Social indicators introduce a unique set of challenges. Extracting sentiments from platforms like Twitter is a complex endeavor. The often ambiguous nature of human language, combined with the presence of automated bots, can lead to misinterpretations. Furthermore, platforms like Twitter can amplify the effects of short-term news, leading to potential market overreactions that might not be indicative of long-term trends. An additional concern is the potential geographic and demographic biases. Relying solely on platforms popular in certain regions might not give a holistic view of the global Bitcoin sentiment.

Last but not least, a recurring challenge has been the design and implementation of an effective reward function. Traditional approaches often relied on overly simplistic reward structures, typically rewarding accuracy without a nuanced understanding of the market's volatile nature. These models frequently overlooked critical aspects such as the confidence level of predictions, the impact of consecutive errors, and the need to adapt to rapidly changing market conditions. As a result, while they could achieve a degree of predictive accuracy, these models often fell short in terms of

robustness and adaptability, leading to potential overfitting to historical data and a lack of responsiveness to new market dynamics. Furthermore, the absence of a time-based discounting mechanism meant that older predictions, possibly irrelevant in the fast-paced cryptocurrency market, were given equal weight to more recent, pertinent information. This oversight could lead to a misalignment between the model's outputs and the current market state, ultimately impacting the reliability and practical applicability of these predictive tools in real-world trading scenarios.

Recognizing these challenges, our study seeks to carve a path forward. We address the data source limitations through feature selection. Using the Pearson correlation method, we prioritize the most impactful metrics across these sources, thereby focusing our analysis on the most pertinent indicators. Moreover, by integrating these selected features with the DQN model, we aim to adeptly capture the nuanced relationships between these metrics and Bitcoin's price direction. Moreover, in tackling the challenge of designing a proper reward function for price direction change prediction, our approach introduces a sophisticated, multi-dimensional reward system. This system is meticulously engineered to not only reward accuracy but also to consider the confidence level of each prediction, apply a graduated penalty for consecutive errors, and incorporate a time-based discounting element. By doing so, we ensure that the model remains agile and aligned with the rapidly evolving market conditions, effectively addressing the limitations of previous studies. This innovative reward function is pivotal in enhancing the model's ability to make robust, well-informed predictions that are reflective of the complex dynamics of the cryptocurrency market, thereby substantially improving prediction reliability and practical applicability.


### III. DATA PREPARATION

The nature of Bitcoin's price prediction is intrinsically linked to the quality and relevance of the data sources utilized. In this section, we aim to offer a comprehensive understanding of the diverse datasets employed in our study, providing both a rationale for their inclusion and details on their acquisition.

Our research focuses on the period in the Bitcoin price market, specifically from April 1, 2014, to November 14, 2018. This selection was influenced by notable price fluctuations observed during this period, as can be seen in Figure 1, which visually showcases the Bitcoin price trajectory over these years. Such pronounced volatility not only presents a rich tapestry for analysis but also poses formidable challenges for any predictive model. Thus, this period serves as an ideal litmus test to gauge the robustness and effectiveness of our proposed method, particularly given the intricate dynamics at play in the Bitcoin market during these years.

#### A. PRICE AND VOLUME DATA

In the sphere of financial analytics, Price and Volume data stand out as foundational elements. For Bitcoin, the close



**FIGURE 1.** Bitcoin Price Trajectory from April 1, 2014, to November 14, 2018. The chart visually highlights the volatility inherent in the cryptocurrency landscape during this time period.

price typically represents the last traded price within a specified period, offering a snapshot of its market valuation at that moment. In our study, the focus on the close price, as opposed to other metrics like opening or average price, was primarily due to its definitive nature, encapsulating the consensus value of Bitcoin at the end of each hour. This choice offers a clear representation of Bitcoin's market value at consistent intervals, allowing for a granular view of its price movement over time.

Volume, on the other hand, gives us an insight into the activity level within the Bitcoin market. It signifies the total quantity of Bitcoin traded within a specific timeframe. A surge in volume often indicates heightened interest or activity, potentially due to external triggers like significant news events or regulatory changes. Conversely, reduced volume might suggest a period of stagnation or reduced market enthusiasm. For our research, the inclusion of volume data, in tandem with close price, aims to provide a holistic understanding of both Bitcoin's market valuation and the underlying market activity propelling those price shifts. To ensure the accuracy and relevance of our analysis, we collected close price and volume information of Bitcoin for our defined experimental time frame. The data was extracted on an hourly basis, providing us with a rich dataset comprising 40,824 hourly records that capture the intricacies and hourly shifts in the Bitcoin market. This high-resolution data serves as a critical foundation for our subsequent analytical endeavors.

Building on the imperative to harness accurate and reliable data, our methodology for data collection was both rigorous and systematic. The primary source for extracting the close price and volume information was the CryptoCompare API [44]. Renowned for its comprehensive databases on cryptocurrency metrics, CryptoCompare is one of the most reputable platforms in the digital currency domain. Its ability to provide historical hourly data, especially over an extended

time frame like ours, made it an invaluable resource for this study. Utilizing the API, we systematically queried and retrieved the specified metrics for Bitcoin. The data extraction was executed with care to prevent any potential gaps or inconsistencies. To ensure the integrity of the data, a validation process was incorporated post-extraction. Any potential outliers or anomalies were cross-referenced with secondary databases, such as CoinGecko [45] and CoinMarketCap [46], ensuring that our dataset was both robust and reliable.

## B. ON-CHAIN METRICS

On-chain metrics grant an intrinsic, clear view of the ongoing activities and patterns within a blockchain. Distinct from external sources of information, these metrics originate directly from the inner operations of the Bitcoin network, serving as a transparent window into its core dynamics. While the blockchain encompasses multiple metrics, it is essential to curate those that hold the highest potential for influencing Bitcoin's future price directions. In this vein, we've selectively focused on specific on-chain data points, deeming them more impactful than others based on our preliminary assessments and their correlations with price movements:

- **Hash Rate:** Demonstrates the Bitcoin network's computational prowess, signifying security and miner confidence in anticipated rewards.
- **Number of Transactions:** Serves as an indicator of the Bitcoin network's activity level and rate of adoption.
- **Average Transaction Value:** Highlights the prevailing economic behavior on the network, hinting at dominant transaction sizes.
- **Active Addresses:** Reflects user activity and can indicate external interest in the Bitcoin network.
- **Exchange Inflows and Outflows:** Indicators of liquidity movements and possible market sentiment shifts.

- **Miner Outflows:** Illuminate miner behavior, particularly their strategies in offloading rewards.
- **Stablecoin Inflows:** As stablecoin dynamics often serve as precursors to market movements, significant inflows could hint at impending buying activity.
- **Options Market Data:** Observations here can provide foresight into potential Bitcoin price shifts, as derivatives often lead spot price actions.

Adhering to our experiment's timeframe, from April 1, 2014, to November 14, 2018, we sourced this on-chain data on an hourly basis. This granular approach ensures that our dataset captures the intricate hourly shifts within the network, offering richer insights. For this endeavor, we mostly relied on Blockchair [47]. This sophisticated analytics engine offers detailed insights, making it an ideal tool for our purposes. To support the reliability of our data, we periodically cross-check our extractions against Blockchain.com's explorer tool [48], reinforcing the integrity of our dataset.

After obtaining the extensive on-chain data for our experiment's duration, it is vital to ensure its integrity and readiness for the subsequent modeling phase. We undertook several preprocessing steps:

- Handling Missing Values:** The robustness of Blockchair notwithstanding, sporadic data gaps emerged. These missing values were especially prevalent in metrics that depend on external factors, like exchange inflows and outflows. We approached this by employing time series-specific interpolation methods, ensuring that filled values are consistent with the adjacent data, preserving the integrity of trends.
- Noise Reduction:** Cryptocurrency markets exhibit frequent price fluctuations. Especially metrics like exchange inflows and hash rates can exhibit transient spikes. To mitigate this, we applied a rolling window average, spanning several hours, to dampen short-term volatility and amplify significant trends.
- Feature Scaling:** Metrics like the number of transactions and the hash rate can differ drastically in their numerical ranges. To bring them onto a common scale and prevent any single metric from unduly influencing the model, we employed Min-Max normalization. This ensured that a unit change in any feature had a consistent impact.
- Temporal Alignment:** While metrics like active addresses were natively hourly, others, such as stablecoin inflows, might initially present at a daily granularity. We resampled such data, distributing daily values uniformly across the 24-hour span, ensuring that every hour was representatively populated.
- Feature Engineering:** Delving deeper, we computed secondary metrics. For instance, from the hash rate, we derived a "hash rate momentum" feature, capturing its rate of change. Similarly, for exchange inflows, we calculated a "net inflow" feature, representing the

difference between inflows and outflows, giving a direct insight into exchange liquidity dynamics.

With these in-depth preprocessing measures, our on-chain metrics were transformed from raw data points into a cohesive, detailed, and actionable dataset, primed for high-quality analytical processing.

### C. SOCIAL INDICATOR DATA

Social indicators, originating from platforms deeply embedded with public sentiment, are important in modern financial markets. Especially in the world of digital assets like Bitcoin, these indicators offer an intricate tapestry of insights into the collective perspective of market participants. Bitcoin, largely steered by public perception and acceptance, is intimately impacted by how it is perceived, discussed, and debated in the digital world. In this volatile cryptocurrency sphere, the resonance of public sentiment often transposes directly into market dynamics. Platforms like Twitter, with their real-time nature and widespread embrace among persons who are engaged with cryptocurrency, stand as reliable sentinels for immediate market sentiment. Every tweet, every hashtag, and every trend might hint at or amplify market movements. Complementarily, Google Trends paints a macro picture, chronicling the rise and fall of general public interest in Bitcoin. Together, these platforms don't just mirror public sentiment—they often presage market trends. To encapsulate these valuable insights, a systematic data extraction approach formulated as follows:

- **Twitter Sentiments:** Our endeavor began by harnessing the Twitter API [49]. Data collection involved aggregating over seven million tweets related to Bitcoin from April 1, 2014, to November 14, 2018. This extraction process leaned on Twitter's streaming API, zeroing in on targeted keywords such as #Bitcoin, #bitcoin, #BTC, and #btc. The raw tweets, inherently muddled with noise, demanded preprocessing. To distill quality from the quantity, the dataset was subjected to rigorous cleaning: removal of URLs, extraneous hashtags, redundant symbols, and other miscellaneous content. After this cleaning, sentiment analysis was performed using the VADER Python library [50]. The derived sentiment scores are being categorized into three classes. Scores between -1 and 0 were deemed as negative sentiments; a neutral sentiments was denoted by a score of 0, while scores between 0 and 1 flagged positive sentiments. For our study and to finesse the input state of our DQN model, these continuous sentiment scores were rounded off to two decimal places and treated as discrete values.
- **Google Trends Data** To measure the volume of public interest surrounding Bitcoin, we turn to Google Trends—a service offered by Google that reflects how often a specific search term, in this case, 'Bitcoin', is entered relative to the total search volume on Google over a specified time frame. Aligning with our experiment's temporal bounds, we capture data

from April 1, 2014, to November 14, 2018, using the PyTrends library [51], an API for Google Trends in Python. Given that Google Trends offers data in a normalized format, the output ranges from 0 to 100, where the peak of the search interest for the time and region specified represents 100. This normalized format ensures relative comparisons across varying time frames and regions, but can introduce volatility. To address this potential issue, we apply a moving average to smoothen out short-term fluctuations and capture the longer-term trend. After gathering the Google Trends data, preprocessing steps are necessary to ensure consistency. This involves handling missing values by applying interpolation methods and normalizing the entire dataset to ensure that the values range between 0 and 1. These transformed values provide a continuous scale representing the level of public interest in Bitcoin over the specified period.

- **Twitter Trends Data** Twitter trends represent a real-time snapshot of what topics or hashtags are currently popular or garnering the most attention on the Twitter platform. These trends can be globally aggregated or localized to specific regions or cities. Given the global reach and influence of Bitcoin, understanding the sentiments and discussions around it on a platform like Twitter can offer valuable insights. We harness the Twitter API, specifically focusing on the “Trends” endpoint, to access the most popular and trending discussions related to Bitcoin. The Python-based Tweepy library [52] facilitates our data extraction from this API. Aligning with our overarching time frame from April 1, 2014, to November 14, 2018, we extract trending data on an hourly basis. The raw data retrieved comprises trending topics, their tweet volume, and associated metadata. It’s essential to discern that these trends are not strictly numerical; instead, they encompass hashtags, topics, and associated keywords. To transform this into a quantifiable metric for our analysis, we apply a weighted scoring mechanism. Each trending topic related to Bitcoin receives a score based on its position in the trend list (with higher positions receiving higher scores) and its tweet volume. However, like any raw dataset, this data demands cleansing. Preprocessing involves filtering out irrelevant trends, addressing potential missing values using forward-fill or backward-fill methods, and then scaling the weighted scores to lie between 0 and 1. This streamlined dataset offers a representation of the popularity and discussions around Bitcoin on Twitter over the chosen period.

As we delve deeper into each data source, we shall elaborate on its significance in capturing the multifaceted nature of Bitcoin’s price movement and the methodology adopted for its collection. In order to provide clarity regarding the diverse range of data sources used in our research, we have tabulated a detailed summary that encompasses essential

details for each type of data. Table 2 offers a comprehensive overview of this.

In summary, the diversified data sources elaborated upon in this section serve as the raw ingredients for our ensuing analytical processes. Before diving directly into predictive modeling, it’s important to recognize the salient features among the vast array of data available. To this end, we employ the Pearson Correlation method, systematically identifying the metrics most crucially linked to Bitcoin’s price direction changes. Only with this nuanced understanding do we then harness the capabilities of the DQN predictive model. By focusing on the most influential data points, our approach aims to offer a precise, streamlined, and effective prediction of Bitcoin’s price direction. The subsequent sections provide a deeper exploration of this methodology, its implementation, and the resulting insights.

#### IV. METHODS AND IMPLEMENTATION

In this section, we delve deeper into the methodological underpinnings of our study, laying out the foundations that enabled our research. The output of the prediction model can be seen in two distinct yet interconnected phases: identifying the right features and constructing a prediction mechanism. Firstly, to ensure the model operates on relevant information, we prioritize the features that have a pronounced impact on Bitcoin’s price direction. For this, we employ the Pearson correlation method, which aids in discerning the relationship between individual features and price direction changes.

Subsequently, we harness the capabilities of the DQN algorithm for the predictive task. While DQN’s origins and foundational concepts form an essential backdrop, our primary focus is on its adaptation to our specific challenge—predicting Bitcoin’s price direction. In this adaptation, elements of the Markov Decision Process play a critical role, from defining the state using our selected features to crafting a reward function that aligns with our objectives.

##### A. PEARSON CORRELATION METHOD

The Pearson correlation coefficient, often represented as  $r$ , is a statistical measure that quantifies the degree of linear relationship between two variables. Its value ranges between -1 and 1, with -1 indicating a perfect negative linear correlation, 1 signifying a perfect positive linear correlation, and 0 meaning no linear correlation. The formula to compute the Pearson correlation coefficient between two variables,  $x$  and  $y$ , is given by:

$$r_{x,y} = \frac{\Sigma(x - \bar{x})(y - \bar{y})}{\sqrt{\Sigma(x - \bar{x})^2 \Sigma(y - \bar{y})^2}} \quad (1)$$

where,  $x_i$  and  $y_i$  are the values of the two variables,  $\bar{x}$  and  $\bar{y}$  are the mean values of  $x$  and  $y$ , respectively. The resultant coefficient  $r$  lies in the range of [-1, 1]. A value close to 1 implies a strong positive correlation: as one variable increases, the other also tends to increase. A value close to -1 implies a strong negative correlation: as one variable increases, the other tends to decrease.



**TABLE 2.** Detailed summary of data sources for bitcoin price direction prediction from april 1, 2014, to november 14, 2018, collected on an hourly basis.

Data Category	Data Type	Description	Source
Price and Volume Data	Close Price	Closing price of Bitcoin for each hour.	CryptoCompare
	Trading Volume	Volume of Bitcoin trades occurring in each hour.	
On-Chain Metrics	Hash Rate	Reflects the computational power of the Bitcoin network.	Blockchair
	Number of Transactions	Total transactions on the Bitcoin network in each hour.	
	Average Transaction Value	Average value of all Bitcoin transactions in each hour.	
	Active Addresses	Total unique active addresses transacting on the Bitcoin network in each hour.	
	Exchange Inflows	Total inflow volume to exchanges.	
	Exchange Outflows	Total outflow volume from exchanges.	
	Miner Outflows	Volume of Bitcoin sent from miners to other addresses.	
	Stablecoin Inflows	Inflows of stablecoins to the Bitcoin network.	
	Options Market Data	Data related to Bitcoin options trading, including volume and open interest.	
Social Indicator Data	Twitter Sentiments	Sentiments extracted from over seven million tweets related to Bitcoin, categorized by VADER scores.	Twitter API
	Google Trends Data	Relative search interest for the term 'Bitcoin'.	Google Trends via PyTrends
	Twitter Trends Data	Data showing trending topics and hashtags related to Bitcoin with associated tweet volumes.	Twitter API

A value close to 0 suggests a weak or no linear correlation between variables [53]. The Pearson correlation coefficient is beneficial in various analytical scenarios due to its simplicity and ability to encapsulate the linear relationship between two variables into a solitary value. It's particularly valuable in assessing how a particular feature might influence the outcome in predictive modeling tasks.

In contexts with multiple features, especially when focusing on feature selection, Pearson correlation extends its utility. The aim is to evaluate the linear relation between each feature and the desired outcome, aiding in prioritizing those with significant relationships. For a dataset with  $n$  features, denoted as  $x_1, x_2, \dots, x_n$ , and a target variable  $y$ , the Pearson correlation coefficient for each  $x_i$  with respect to  $y$  is determined as:

$$r_{x_i, y} = \frac{\sum(x_i - \bar{x}_i)(y - \bar{y})}{\sqrt{\sum(x_i - \bar{x}_i)^2 \sum(y - \bar{y})^2}} \quad (2)$$

After determining the correlation coefficients for all features with respect to the target, they can be ranked. Features with insignificant correlation might be considered less relevant, leading to their exclusion and potentially enhancing the model's efficiency. This approach becomes particularly salient in the case of Bitcoin price direction prediction. Using this technique, after ranking all features by their correlation values, those with negligible associations can be de-emphasized or omitted outright. Modern computational tools, notably Python libraries like *scipy* [54] and *pandas* [55], offer effective methods to compute Pearson correlation

matrices for datasets. This ensures computational feasibility even with voluminous datasets, facilitating a seamless analytical journey.

## B. DEEP Q-NETWORK AND ITS IMPLEMENTATION

The DQN represents a confluence of two significant domains: deep learning and reinforcement learning. Its origins trace back to a groundbreaking paper by DeepMind in 2015, wherein they successfully trained agents to play a range of Atari 2600 games using visual input [17]. This achievement underscored the potency of combining the generalization capabilities of deep neural networks with the decision-making acumen of Q-learning, a classical reinforcement learning algorithm.

At its core, Q-learning seeks to learn an action-value function that assigns a value to each possible action in every possible state. This function is typically represented as  $Q(s, a)$ , which signifies the expected cumulative reward from taking action  $a$  in state  $s$  and following an optimal policy thereafter. Mathematically, the Q-value update rule in Q-learning is:

$$Q(s, a) \leftarrow Q(s, a) + \alpha (r + \gamma \max_{a'} Q(s', a') - Q(s, a)) \quad (3)$$

where,  $\alpha$ - is the learning rate,  $\gamma$ - is the discount factor,  $r$ - is the immediate reward, and  $s'$ - is the next state. In this equation above, the term  $r + \gamma \max_{a'} Q(s', a')$  represents the learned value, and  $Q(s, a)$  is the old value. The difference

between these two is the Temporal Difference (TD) error. This error is pivotal for updating the Q-values.

In simplistic scenarios, the action-value function can be tabulated. However, as the state or action space grows or becomes continuous, a tabular approach becomes infeasible. This is where deep learning intervenes. Instead of a table, a neural network, known as the Q-network, approximates the action-value function. Given a state  $s$  as input, this network produces Q-values for all possible actions. The action with the maximum Q-value is typically chosen as the next action to execute. The training of this Q-network is anchored in the Bellman equation. By minimizing the difference between the predicted Q-value and the actual Q-value, known as the TD error, the network's weights are iteratively updated. The DQN introduces several innovations to this paradigm, including Experience Replay and Target Networks, to stabilize training and enhance convergence.

Diving deeper into the fabric of reinforcement learning reveals the pivotal role of the Markov Decision Process (MDP). MDP provides a structured framework within which agents interact with their environment, make decisions, and receive feedback. It comprises three main elements: the state space, which enumerates all possible situations or conditions the agent might encounter; the action space, detailing the set of feasible actions the agent can undertake in each state; and the reward function, offering feedback on the quality of actions, guiding the agent's learning process. As we transition to the specifics of our study, a closer examination of these MDP elements and their intricacies in the context of predicting Bitcoin price directions becomes paramount.

Continuing from the aforementioned foundations of DQN, our research facilitates the predictions of Bitcoin's price directions by methodically establishing the MDP elements. An MDP provides a framework that encapsulates the interaction of an agent with its environment, and for our context, these MDP elements are pivotal.

- **State Space  $S$**  - In the context of our research state space is a critical construct that encapsulates a comprehensive representation of the environment within which the predictive model operates. This environment is characterized by distinct features that, collectively, have been shown to influence Bitcoin's price direction. To build a precise and meaningful state space, we first employ the Pearson correlation method. This method assists us in sieving through an extensive range of features and spotlighting only those that have a significant bearing on Bitcoin's price trajectory. However, given the multifaceted nature of the data, the resultant state isn't a simple scalar. Instead, it manifests as a matrix. The dimension of this matrix is  $n$ , where  $n$  denotes the number of salient features we've shortlisted using the Pearson correlation. It's imperative to note that each entry within this state matrix corresponds to the data of a particular feature, effectively encapsulating the entire feature set in one coherent structure. To further

the granularity and align with the overarching research methodology, each state corresponds to an hourly snapshot of these features, ensuring the state space is both detailed and time-relevant.

- **Action Space  $A$**  - In the current research, the potential decisions or predictions the DQN model can make about the future price direction of Bitcoin represent the action space. Given the inherent variability and the unpredictable nature of cryptocurrency markets, we've simplified the action space to embody three quintessential movements that are paramount for traders and investors:

- 1) **Increase:** This action corresponds to a prediction that Bitcoin's price will ascend in the subsequent hour. An increase suggests that market conditions or other influential factors are pushing the price upwards. Traders often look for such signals to make buying decisions or to hold their existing assets in anticipation of higher returns.
- 2) **Decrease:** Representing the converse scenario, a decrease indicates the DQN model's anticipation that Bitcoin's price will diminish in the upcoming hour. Such predictions can be vital for traders aiming to sell their assets to avoid potential losses or for those looking to buy assets at a lower price in the near future.
- 3) **No-Change:** This action is an assertion that the price of Bitcoin will remain relatively stable, neither increasing nor decreasing appreciably. In a market known for its volatility, periods of stability can be both a respite and a strategy point. Traders might interpret this as a period of consolidation, or they might use it as a brief window to analyze other external factors before making their next move.

It's important to understand that these actions aren't mere speculations but are the outputs of the DQN model after it analyzes the current state (a matrix of influential features). By categorizing the action space into these three distinct directions, we ensure that the model's predictions are actionable, clear, and straightforward, catering to the pragmatic needs of Bitcoin market participants.

- **Reward  $R_t$**  The reward function is the cornerstone of any reinforcement learning system, providing the necessary signals to the model to facilitate its learning journey. In the context of Bitcoin price direction prediction, our reward function is designed to reflect the intricacies of the cryptocurrency market. It's constructed to appreciate the subtleties of accurate forecasting and to penalize mistakes, embodying the delicate balance traders must strike in real-world conditions. The reward function is built upon several core components, each tailored to reinforce the model's proficiency in predicting market dynamics. Below, we elaborate on the components that comprise our sophisticated reward function:

- 1) **Predictive Accuracy Incentive:**  $S(\hat{a}_t, a_t)$ . This aspect of the reward function encapsulates the dual nature of prediction outcomes—rewarding accuracy and penalizing errors. It is a direct reflection of the principle that correct predictions should be rewarded, while incorrect ones should be penalized, fostering a model that strives for high accuracy in its forecasts. The model is incentivized for each precise direction prediction—whether it's an uptick or a downtick in price—through a reward function  $S$ , which is defined as:

$$S(\hat{a}_t, a_t) \begin{cases} 1 & \text{if } \hat{a}_t = a_t, \\ -1 & \text{otherwise.} \end{cases} \quad (4)$$

In this function,  $\hat{a}$  represents the model's predicted direction of price movement at time  $t$  (e.g., increase, decrease, or no-change), and  $a_t$  denotes the actual direction that the price moved at time  $t$ . The function grants a value of +1 for correct predictions. This part of the reward function emphasizes the model's alignment with the primary goal of prediction: accuracy. Reflecting the high stakes of cryptocurrency speculation, our model faces a penalty for every misprediction, mirroring the adverse consequences of erroneous decisions in the trading domain. This penalty, also determined by  $S$ , deducts a value of -1, reinforcing the critical nature of precision.

- 2) **Confidence-Adjusted Scaling:**  $C(c_t)$ . Confidence in prediction is a measure of the model's certainty in its own forecasts. This confidence level, represented by  $c_t$ , plays a critical role in our reward function. The confidence score at time  $t$ , is typically derived from the predictive model itself. For instance, in a neural network-based model, it could be the output of a softmax layer for classification tasks, representing the probability assigned to the chosen action. This score ranges from 0 (no confidence) to 1 (absolute confidence). The reward function scales the accuracy reward by a factor of the model's confidence, denoted by  $C$ , which is defined as:

$$C(c_t) = \gamma * c_t \quad (5)$$

The scaling factor  $\gamma$  is a hyperparameter that modulates the influence of confidence on the overall reward. A suitable value for  $\gamma$  can be chosen based on empirical testing and model performance. A common starting point could be  $\gamma = 1$ , giving direct weight to the confidence level. However, this value might be adjusted during the model's training phase to optimize performance. A higher value emphasizes the role of confidence in the reward, potentially making the model more conservative in its predictions, as high confidence is needed to achieve significant rewards. Conversely, a lower

value reduces the impact of confidence, allowing the model to be rewarded more for accuracy regardless of its confidence level. By integrating this confidence-adjusted scaling, we ensure that the model not only seeks to predict accurately but also develops a refined sense of the reliability of its predictions. This approach encourages the model to be judicious in its forecasts, aligning rewards more closely with the certainty of its predictions.

- 3) **Consecutive Error Penalty:**  $E(n)$ . Consistent incorrect predictions are indicative of a potential flaw in the model's learning process. To address this, we introduce a penalty for consecutive errors, which serves as a critical mechanism to incentivize the model to learn from its mistakes, adapt, and avoid falling into repetitive erroneous patterns. This penalty is formulated to grow exponentially with each successive error, providing a substantial disincentive for the model to persist in making the same mistakes. The penalty function, denoted by  $E$ , is defined as:

$$E(n) = \alpha^n \quad (6)$$

Here,  $n$  represents the number of consecutive incorrect predictions made by the model. The base of the exponent,  $\alpha$ , is a hyperparameter greater than 1, which determines the rate at which the penalty increases for each successive error. A typical starting value for  $\alpha$  might be in the range of 1.05 to 1.10, providing a moderate increase in the penalty for each consecutive error. The selection of  $\alpha$  should be carefully tuned based on empirical results to ensure it effectively deters the model from making repetitive errors without being overly punitive for isolated mistakes. The exponential nature of  $E$  means that the penalty becomes significantly larger with an increasing streak of incorrect predictions. This approach encourages the model to break out of any potential error loops and stimulates a more diverse exploration of strategies. By implementing this penalty, we aim to enhance the model's adaptability and robustness, promoting a learning process that is more reflective of the complexities and uncertainties inherent in Bitcoin price direction prediction.

- 4) **Time-Based Discounting:**  $D(t)$ . The Bitcoin market is inherently volatile, characterized by rapidly shifting trends and patterns. To ensure that the model remains agile and responsive to these changes, we incorporate a time-based discounting factor,  $D$ , into our reward function. This factor is crucial for emphasizing the relevance of more recent predictions and diminishing the weight of older ones, thereby aligning the model's focus with the latest market information. The time-based

discounting factor is defined as:

$$D(t) \delta^t \quad (7)$$

In this formulation,  $t$  represents the time step or the age of the prediction, with more recent predictions having a lower  $t$  value. The base of the exponent,  $\delta$ , is a hyperparameter between 0 and 1, which dictates the rate at which the value of past predictions decays over time. A common starting point for  $\delta$  might be around 0.95 to 0.99, indicating a gradual decrease in the importance of older predictions. This parameter should be tuned based on the model's performance and the specific dynamics of the Bitcoin market it's being trained on. By applying this time-based discount, predictions that are more recent and therefore more likely to be relevant in the rapidly changing market of Bitcoin are given more weight. This approach ensures that the model does not rely excessively on outdated information and remains adaptable to the latest market dynamics. It encourages the model to continually update its understanding of the market, staying in tune with the most current trends and shifts. This aspect of the reward function is particularly important in the context of cryptocurrency trading, where market conditions can change dramatically in a short period.

In summary, we propose to compute the total reward as the multiplication of the above-given factors—accuracy, confidence, consecutive errors, and time-based discounting. The composite reward function  $R_t$  at time  $t$  can be expressed succinctly as:

$$R_t S(\hat{a}_t, a_t) \times C(c_t) \times E(n) \times D(t) \quad (8)$$

This multiplicative approach is chosen to ensure a holistic and balanced influence of each component on the model's learning process. By multiplying these factors, the model is encouraged to maintain a consistent performance across all dimensions. Crucially, this method amplifies the impact of each component, where a deficiency in any one factor, such as low confidence or a high consecutive error count, significantly reduces the total reward. This is especially pertinent in Bitcoin price direction prediction, where the accuracy of a prediction is paramount, and any incorrect prediction leads to a substantial penalty, reflecting the high-risk nature of cryptocurrency trading. Furthermore, this multiplicative structure promotes a comprehensive learning strategy, ensuring the model develops a well-rounded understanding and responsiveness to the dynamics of the market. It aligns with the objectives of reinforcement learning in complex, multifaceted environments, fostering an integrated approach to decision-making.

In synthesizing the above MDP elements, we tailor a structured framework wherein our DQN model, equipped

with essential features, can autonomously predict the price directions of Bitcoin. This confluence of carefully chosen states, delineated actions, and a potent reward function ensures that our predictive model is both robust and high-performing.

## V. EXPERIMENT RESULTS


In this section, a comprehensive examination of the experimental results and findings is presented, showcasing the applicability and efficacy of the deployed DQN model in predicting Bitcoin price direction changes. Emphasis is placed on ensuring model robustness and validating its predictive prowess through experimental results. A notable aspect of the model validation process involves leveraging  $k$ -fold cross-validation to maximize the utilization of the available data, especially given the inherently limited quantity of financial data. Specifically, the dataset is partitioned into  $k(k - 10)$  distinct subsets. In each iteration of the validation, one subset is designated as the validation set, while the remaining  $k - 1$  subsets collectively form the training set. This approach guarantees that each data point is validated exactly once and is included in the training phase  $k - 1$  times. The error estimations obtained across all  $k$  trials are then averaged to furnish a more reliable and generalized model performance metric.

The experimental framework was executed on a system equipped with a Core i7 microprocessor, and 16 GB RAM, operating on a Windows 11 platform. The experiments were conducted using Google Colab, a platform renowned for its robust support for machine learning applications and complimentary access to GPU computational resources. The data analysis and model development heavily utilized Python's statistical and numerical libraries, NumPy and Pandas. Additionally, the DQN model was constructed and trained to utilize the TensorFlow and Keras libraries, which are integrated into the methodology to facilitate the development of a predictive model capable of discerning future price trends of Bitcoin.

### A. FEATURE SELECTION

Given the multifaceted nature of Bitcoin price movements, a total of 14 distinct features initially considered for the model, each offering a unique perspective into potential predictive attributes: Historical Prices, Volume of Trades, Hash Rate, Number of Transactions, Average Transaction Value, Active Addresses, Exchange Inflows, Exchange Outflows, Miner Outflows, Stablecoin Inflows, Options Market Data, Twitter Sentiments, Google Trends Data, and Twitter Trends Data.

The Pearson correlation method stands as the chosen technique to scrutinize the linear relationship between these features, providing a correlation coefficient that signifies the strength and direction of their linear relationship. This statistical method quantifies the degree to which variables in a dataset are linearly related, producing a correlation coefficient, which ranges from  $-1$  to  $+1$ . A coefficient



**FIGURE 2. Correlation Matrix of Features.**

close to +1 indicates a strong positive linear relationship, -1 indicates a strong negative linear relationship, and coefficients close to 0 suggest no linear correlation.

Figure 2 shows the correlation matrix heatmap, providing a visually intuitive representation of the Pearson correlation coefficients among the 14 features. Each cell in the heatmap signals the strength and direction of the correlation between two features, wherein warmer colors indicate stronger positive correlations and cooler colors denote stronger negative correlations.

When examining the correlation matrix, we focused on features that had a strong relationship with the target variable: Bitcoin’s price direction change. At the same time, we had to be careful about including features that were too closely related to each other, to avoid redundancy in the model. Subsequent steps in the selection process involve analyzing these chosen features against each other to ascertain any potential multicollinearity, which might compromise the robustness and interpretability of the model. For instance,

the “Historical Prices” align prominently with public sentiment and interest metrics - specifically, “Google Trends Data” (0.8) and “Twitter Sentiments” (0.7). These figures underscore a robust relationship between public engagement, online sentiment, and Bitcoin’s price, suggesting that public interest and positive sentiment on social platforms potentially drive Bitcoin’s price movement.

In contrast, the “Hash Rate,” which sits at a correlation value of 0.2, exhibits a relatively mild correlation with Bitcoin’s price direction. This indicates that, despite its importance for network security and transaction verification, the hash rate does not robustly correlate with price changes. Turning to inter-feature correlations, the tight bond between “Google Trends Data” and “Twitter Sentiments” (0.9) signals a parallel movement between public interest and the prevailing sentiment on social media. Conversely, “Exchange Inflows” and “Exchange Outflows” showcase an inverse relationship of -0.9, indicating that an influx in one often coincides with a reduction in the other, hinting at a

compensatory mechanism in the exchange reserve behaviors among market participants.

The initial selection features that have a correlation coefficient magnitude higher than the threshold with the Historical Prices, to ensure a moderate to strong linear relationship with the predicted variable. In essence, a 0.6 threshold is chosen to ensure the model is robust, computationally feasible, and avoids overfitting, while still providing valuable, interpretable insights through a meaningful, empirically-supported linear relationship with the target variable. This aids in achieving accurate, reliable predictive outcomes in forecasting Bitcoin's future price directions. Consequently, the following features show a correlation coefficient above the established threshold: Volume of Trades ( $r$  0.7), Number of Transactions ( $r$  0.6), Active Addresses ( $r$  0.6), Options Market Data ( $r$  0.6), Twitter Sentiments ( $r$  0.7), Google Trends Data ( $r$  0.8), Twitter Trends Data ( $r$  0.7).

Through evaluation of the Pearson correlation results, features that exhibited substantial correlations with Bitcoin's future price direction were prioritized, while also maintaining a judicious consideration to minimize inter-feature correlations. Subsequently, 7 features were ultimately selected to be integrated into the model for training, ensuring a pragmatic balance between robust predictive capability and model simplicity. These selected features encompass, each providing a unique lens through which the model interprets historical data to forecast future price directions of Bitcoin. In the model training phase, selected features serve as key inputs, utilized to predict Bitcoin's future price trends, ensuring a model that embodies both stability and predictive accuracy. Future sections will delve deeper into the model training, validation, and evaluation, exploring the implications and insights drawn from these selected features.

## B. DQN MODEL STRUCTURE AND PERFORMANCE EVALUATION

The DQN model, employed in this research, assimilates information from the seven selected features to facilitate accurate predictions of Bitcoin's future price trends. Given the complexity and high volatility inherent in cryptocurrency markets, the model structure is crafted to encapsulate potential non-linear relationships among the input features. The model initiates with an input layer comprising seven nodes, each corresponding to one of the chosen features: 'Historical Prices', 'Volume of Trades', 'Number of Transactions', 'Average Transaction Value', 'Active Addresses', 'Google Trends Data', and 'Twitter Sentiments'. This specific architecture ensures that the model comprehensively integrates diverse market dimensions to form its predictive output.

For the hidden layers, we implement a three-layer architecture. The first hidden layer incorporates 14 nodes, doubling the input size to facilitate the network's learning from intricate relationships within the data. The second hidden layer, designed for further abstraction, consists of 7 nodes, ensuring that the model sustains its focus on extracting nuanced patterns without leaning towards overfitting. The

final hidden layer, before the output, reintroduces a larger space of 14 nodes, aiming to synthesize the abstracted information into a cohesive understanding.

On arriving at the output layer, the DQN model culminates in three nodes, each representing a distinct potential action: 'Increase', 'Decrease', or 'No-change' in Bitcoin investment. This division allows the model to forecast the future directionality of Bitcoin prices by picking the action with the highest Q-value, which is derived from the Q-learning algorithm, in accordance with the state presented by the input layer.

Hyperparameters, such as the learning rate, discount factor, and exploration-exploitation trade-off, are carefully selected to balance efficient and explorative learning. A learning rate of 0.001 ensures gradual adaptation during training, mitigating the risk of overshooting optimal weight configurations. We set the discount factor ( $\gamma$ ) at 0.95 to provide a favorable balance between prioritizing immediate and future rewards. Furthermore, an epsilon-greedy strategy with an initial  $\epsilon$  of 1, decaying at a rate of 0.995 per episode, is employed to balance exploration and exploitation throughout the learning process. Moving towards the structural aspects of the model, the DQN integrates an activation function, specifically the Rectified Linear Unit (ReLU), renowned for mitigating the vanishing gradient problem and facilitating the model in capturing non-linear dependencies within the data. The Mean Squared Error (MSE) loss function is employed, honing in on minimizing the discrepancies between the predicted Q-values and the target Q-values, ensuring a focused and consistent training objective. Optimizing this loss function, the Adam optimizer is selected for its aptitude in adaptively tuning learning rates, thereby accommodating a more nuanced and adaptive optimization process. Further, synchronization in model training is achieved through updating the target network every 1000 steps, ensuring a stable and consistent target for the iterative updates, while concurrently averting excessively rapid adaptations that could destabilize the learning process. Additionally, a batch size of 64 is set for experience replay, ensuring the model is exposed to a diverse and representative subset of experiences during training, fostering a more generalized and robust policy development. Table 3 provides a concise overview of the architectures and hyperparameters used in the proposed DQN model of this study.

**Performance Evaluation.** In consideration of a predicted output manifesting as one of three possible outcomes (namely, an increase in price, a decrease, or no alteration), this task is identified as a multi-classification problem. Our strategy facilitates the assessment of our model's efficacy through the calculation of its F1 score. This score, a prevalent metric for gauging the aptitude of classification models [56], embodies the harmonic mean of precision and recall. Consequently, it yields a consolidated metric for evaluating the model's aptitude to accurately classify both positive and negative instances. To derive the F1 score for a prediction, several values were initially computed:

TABLE 3. DQN model architecture and hyperparameters.

Input Nodes	Output Nodes	Hidden Layers	Learning Rate ( $\alpha$ )	Discount Factor ( $\gamma$ )	Initial Exploration ( $\epsilon_{init}$ )
7	3	3 layers (14, 7, 14 nodes each)	0.001	0.95	1
Loss Function	Activation Function	Exploration Decay Rate ( $\epsilon_{decay}$ )	Optimizer	Target Network Update	Batch Size
MSE	ReLU	0.995	Adam	1000	64

- **True Positives (TP):** TP represents instances where the model's predictions and the actual outcomes both align in denoting a positive result. In the context of our model, a true positive would occur when both the predicted and actual outcomes suggest a price increase on a given time (in our case on a given hour). Ensuring high TP values is crucial as it reflects the model's capacity to correctly identify positive outcomes.
- **False Positives (FP):** FP indicates instances where the model erroneously identifies an outcome as positive when it is, in fact, negative. In the scenario of predicting Bitcoin prices, a false positive means that the model anticipates a price rise when the price actually falls. Managing FP is vital to prevent falsely optimistic predictions and ensuing erroneous decision-making.
- **False Negatives (FN):** FN accounts for instances where the model inaccurately predicts a negative outcome while the actual result is positive. In our model, a false negative occurs when a price increase occurs but is incorrectly predicted as a price drop or stable price by the model. Mitigating FN is critical to avert overlooking potential positive opportunities.
- **Precision:** Precision represents the accuracy of positive predictions, quantifying the ratio of correctly predicted positive observations to the total predicted positives. High precision indicates an efficacy in predicting positive instances and a reduction in the false-positive rate, thereby enhancing the reliability of positive predictions.

$$Precision = \frac{TP}{TP + FP} \quad (9)$$

- **Recall:** Recall quantifies the model's ability to identify all relevant instances, by measuring the ratio of correctly predicted positive observations to all observations in the actual class. Ensuring a high recall is crucial to guarantee that most positive instances are not being misclassified, thus preventing opportunity costs associated with missed positive instances.

$$Recall = \frac{TP}{TP + FN} \quad (10)$$

Subsequent to determining the values for precision and recall, the F1 score can be calculated as follows:

$$F1 \text{ score} = \frac{2 \times Precision \times Recall}{Precision + Recall} \quad (11)$$

The F1 score, which ranges between 0 and 1, serves as an indicator of model performance, with higher values signifying superior performance. The F1 score measures the accuracy of a classification model by comparing its predictions with the actual labels in the dataset. In this research, the absence of a pre-labeled dataset was mitigated by employing a Python code. This code observes the price of Bitcoin on the  $n$ -th hour (with  $n \in (1, \text{number of experiment hours})$ ) and compares it with the preceding hour's ( $n - 1$ ) Bitcoin price. This comparison aids in defining whether the price increased, decreased, or kept its value, thereby enabling the labeling of the  $n$ -th hour accordingly. By replicating these steps for each hour under review, a labeled dataset was constructed. Consequently, this dataset enabled the calculation of the model's F1 score.

### C. PREDICTION OF PRICE DIRECTION CHANGES

This experiment involves a dataset that combines various types of financial data, all related to the Bitcoin market. We choose our data sources carefully to span a range of financial indicators and select seven based on the Pearson correlation method for use in our model. To execute this experiment, we download Bitcoin data from various financial databases and use a DQN model for our predictions.

To ensure detailed and effective data analysis, we have implemented a dataset split process to validate our model's prediction accuracy more robustly. We utilized the k-fold cross-validation technique, which is particularly beneficial for avoiding overfitting and ensuring that our model is tested comprehensively on different subsets of the data.

For a more detailed breakdown: our dataset, providing hourly Bitcoin data from April 1, 2014, to November 14, 2018, was divided into training, validation, and testing sets. We allocated 70% of the data for training, allowing the model to learn and adapt to various patterns and trends in the Bitcoin market. To fine-tune the model and adjust hyperparameters effectively, 15% of the data was dedicated to validation. This validation set plays a crucial role in preventing overfitting, as it helps in evaluating the model's performance on data that it has not encountered during the training phase. The remaining 15% of the data was reserved for testing the model's performance, providing an unbiased evaluation of its predictive power.

Employing the k-fold cross-validation technique, where 'k' was set to 10, ensured that every data point had the

**TABLE 4.** Evaluation of DQN model's prediction.

Metrics	Increase	No-change	Decrease
True Positives (TP)	56451	5706	5463
False Positives (FP)	558	238	349
False Negatives (FN)	425	364	607
Precision	0.91	0.96	0.94
Recall	0.93	0.94	0.90
F1-score	0.92	<b>0.95</b>	0.92

opportunity to be in the test set exactly once and in the training set  $k - 1$  times. This method significantly enhances the reliability of our model's performance metrics, as it allows us to affirm that the accuracy figure is not merely a result of the model's overfitting to a specific subset of the data but is indeed indicative of its true predictive capability.

In line with the basic principles of the DQN model, an agent learns from its experiences and continuously improves its decision-making strategy over subsequent iterations. Therefore, we expect the model's performance to improve when the remaining dataset is used. Through the conduction of this experiment, our aim is to gain a deeper understanding of the relationship between various financial indicators and the hourly fluctuations in the Bitcoin market.

As mentioned earlier, we constituted 15% of our total data as a validation dataset. It is important to strike a balance between learning complexity and model generalizability, ensuring that our model neither overfits nor underperforms. After each training epoch, the model is evaluated on this validation set. This evaluation is to check if the model is learning general patterns rather than memorizing the training data. This allows us to fine-tune hyperparameters such as the learning rate, the number of layers, or the number of neurons in each layer. Such tuning is important to optimize the model's performance, ensuring it can accurately predict Bitcoin prices under varying market conditions. Furthermore, this approach is instrumental in preventing overfitting - where the model performs exceptionally well on training data but poorly on unseen data. By employing this method we stopped the training at the right moment, preserving the model's ability to generalize to new, unseen data, thereby enhancing its practical applicability.

When evaluating the efficacy of the DQN model, our primary focus extends beyond quantitative results to a comprehensive error analysis, as presented in Table 4. This table includes an in-depth evaluation of the model's performance in predicting Bitcoin price direction changes for each trend: 'Increase', 'No-change', and 'Decrease'.

The inclusion of True Positives (TP), False Positives (FP), and False Negatives (FN) provides a deeper insight into the model's predictive accuracy. For 'Increase' predictions, the model correctly identifies 56,451 instances as TP, with 558 instances incorrectly predicted (FP) and 425 actual

'Increase' instances missed (FN). In the case of 'No-change', the model achieves 5,706 TP, with a lower FP count of 238 and 364 FN, showcasing its strong ability to identify stable market conditions. For 'Decrease' predictions, the TP count is 5,463, accompanied by 349 FP and 607 FN, reflecting the model's capability to predict downward trends with reasonable accuracy.

The 'Precision' row in Table 4 measures the model's ability to correctly identify each trend, implying how many instances are accurately predicted relative to the number of total predictions for each trend. Our model showcases a commendable precision of 0.91 for 'Increase', indicating that 91% of all predictions for 'Increase' are indeed accurate. Similarly, for 'No-change' and 'Decrease' predictions, the model attains notable precisions of 0.96 and 0.94, respectively.

The 'Recall' illustrates the ratio of correctly identified positive predictions to the actual positives. Our model successfully identifies 93% of all actual 'Increase' instances, achieving a recall of 0.93. Meanwhile, it exhibits a recall of 0.94 for 'No-change' and 0.9 for 'Decrease', reflecting its proficiency to capture the majority of actual instances for these trends.

The 'F1-score', represented in the last row, serves as a harmonizing metric, providing a singular score that balances precision and recall. Our model attains F1-scores of 0.92, 0.95, and 0.92 for 'Increase', 'No-change', and 'Decrease', respectively, indicating a balanced performance in terms of both false positives and false negatives.

To offer a more intuitive understanding and facilitate the analysis of the results, our model's performance is visualized in a heatmap, denoted as Figure 3. This heatmap portrays the model's predictions across three potential trends in the Bitcoin market: 'Increase', 'No-change', and 'Decrease'. The gradation of color within the heatmap, transitioning from darker blue to lighter red, symbolizes the model's predictive accuracy, with lighter colors indicating a higher level of accuracy. Examining this figure, our model demonstrates robustness, especially in predicting the 'No-change' trend, where it achieves a compelling F1-score of approximately 95%. This underscores its adept ability to reliably forecast periods of stability within the Bitcoin market. Furthermore, in forecasting 'Increase' and 'Decrease' trends, the model maintains notable performance, with accuracies of 92% and 92%, respectively.

However, it is important to note that while the model reveals notable predictive capabilities, instances of misclassification are observable. Misclassifications among 'Increase', 'No-change', and 'Decrease' trends become evident, pinpointing avenues for future model refinement and exploration.

## VI. DISCUSSION

In this section, we explore a comparative analysis, comparing the performance of our DQN model with other established methods for Bitcoin price trend prediction. The primary goal is to assess how our model stands in the broader context






FIGURE 3. Heatmap Visualization of Predictive Metrics for Bitcoin Price Trend Classification using DQN Model.

of prediction techniques. This comparative view not only provides a perspective to measure the relative efficiency and robustness of our DQN model but also highlights potential areas for further improvement and development. The forthcoming table offers a systematic comparison, detailing the metrics of each method. By extracting insights from this comparison, we aim to emphasize the unique advantages of our approach and outline the direction for subsequent improvements in this field.

A crucial element of our discussion is Table 5, which presents a comprehensive performance comparison of our research against previous studies tackling the same issue. By showcasing the accuracy of each method over different years, the table gives a clear picture of the progression in this area of research. This comparison not only highlights the effectiveness of our proposed technique but also places our work within the broader academic discourse.

Bitcoin price prediction research is a diverse field, with many researchers using a variety of methodologies and datasets to predict future price movements. The choice of the dataset, combined with the prediction algorithm, determines the effectiveness of the research, often represented by its accuracy rate. Our previous research [35] utilized a Q-learning-based model that used Bitcoin’s on-chain data and whale transaction tweet data, achieving an accuracy rate of 90.02%. However, the progression in research methods highlights our dedication to improving results. Our current study builds upon our previous findings and incorporates a wide-ranging dataset that includes price and volume data, on-chain data, Twitter sentiments, Google trends, and Twitter trends data. This comprehensive approach has elevated our model’s F1-score to an impressive 95%. In comparison to the findings of Basher et al. [57], our research stands out even more. They employed tree-based classifiers on technical indicators of Bitcoin price data and reported

TABLE 5. Performance comparison with other results.

Studies	Year	Dataset	Accuracy
Q-learning based model (Muminov et.al. [35])	2023	Bitcoin on-chain data & whale transactions tweet data	90.02%
Tree-based classifiers (Basher et.al. [57])	2022	Technical indicators of Bitcoin price data	>85%
Voting Classifier (Critien et.al. [58])	2022	Tweet’s sentiment & volume data	77.2%
CNN-LSTM (Livieris et.al. [59])	2021	Historical Bitcoin price data	55.03%
FastText model (Kilimchi et.al. [60])	2020	Tweet sentiment data	89.13%
Ensemble A:5-10-500 (Mallqui et.al. [61])	2019	Internal (Bitcoin on-chain data) & external (economic factors and external demands) data	62.91%
<b>Our paper</b>	2023	<b>Price and volume data, on-chain data, Twitter sentiments, Google trends, and Twitter trends data</b>	<b>95%</b>

an accuracy slightly above 85%. Their work shows the importance of data selection, which is further reinforced by Critien et al. [58] research. Using a voting classifier on tweet sentiments and volume data, they reported an accuracy of 77.2%. Other studies, like those by Livieris et al. [59] and Kilimchi et al. [60], achieved accuracies of 55.03% and 89.13% respectively.

However, despite our outcome result showing higher value than other existing studies’ results, it is important to recognize some facts that resist obtaining a fair comparison. A primary consideration is the disparity in datasets. While our research leverages a unique combination of data sources, including Bitcoin on-chain data, Twitter sentiments, and Google trends, other studies might rely on entirely different types or sets of data. This variation can substantially influence the outcomes, as the nature of the data inherently affects the patterns and insights derived from it. Furthermore, the timeline or period during which experiments are conducted can play a pivotal role. Market dynamics, external world events, and various other temporal factors can influence the behavior of datasets, leading to varied results even if other conditions remain the same. The choice of algorithms and network structures further complicates a direct comparison. Different algorithms have their strengths, weaknesses, and assumptions, which can lead to variations in performance, even when applied to similar data. Similarly, the network structure, especially in deep learning models, can be a determinant of how effectively patterns are recognized and generalized. Considering these challenges, while we provide comparison results to offer a broader perspective of the research landscape, it’s crucial to

approach them with caution. Direct comparisons, although insightful, should be made cautiously, keeping in mind the multiple underlying variations.

While our method excels in the model's performance and the ability to integrate a wide range of input data, it also has specific disadvantages when compared to existing methods. One such disadvantage is the computational complexity inherent in our DQN model, which requires significant processing power and can be more time-consuming than some of the simpler models presented in Table 5. Additionally, our approach, due to its reliance on a broad dataset, may be more sensitive to the quality and consistency of input data, making it potentially less robust in scenarios of data sparsity or inconsistency compared to more traditional methods. These aspects, while they do not overshadow the overall efficacy of our model, are important considerations for future improvements and adaptations of our approach in the field of cryptocurrency price prediction.

In addition to this, despite the promising results achieved in our study, several limitations and corresponding future works should be acknowledged.

- (i) **Integrating data diversity:** While we selected 14 different types of data that we believed would impact the price direction change of Bitcoin, there is a possibility that some influential data sources were overlooked. These could include factors such as the influence of regulatory changes in major Bitcoin markets, macroeconomic indicators like inflation rates in dominant fiat currencies, and the impact of major technological updates or forks within the Bitcoin protocol. To enhance the sophistication and accuracy of our model, expanding the variety of data sources to incorporate these additional potential influencers is crucial. This expansion may involve integrating a broader range of macroeconomic indicators, closely tracking regulatory shifts in key Bitcoin markets, or accounting for significant technological alterations in the Bitcoin ecosystem. Such enhancements are expected to add depth to our model, enabling it to forecast price changes with increased precision and reliability.
- (ii) **Multifaceted ML approaches:** Our predictive model employs the DQN algorithm, a choice that, while robust in numerous scenarios, is not without its challenges. One well-documented issue with the DQN algorithm is its tendency towards an overestimation bias, which can affect the accuracy of its predictions. To address this critical issue, prioritizing the mitigation of the overestimation bias in our model is essential. Exploring alternative reinforcement learning algorithms, such as Double DQN or Dueling DQN, could provide solutions to this problem. Additionally, augmenting our current algorithm with techniques specifically designed to counter this bias may lead to more reliable and accurate predictions. Furthermore, the development of hybrid models that amalgamate the strengths of various algorithms is another promising avenue. Such an approach could leverage the unique advantages of each algorithm, potentially creating a more robust and effective predictive tool.
- (iii) **Predictive depth:** Our study primarily focuses on predicting the directional change of Bitcoin prices, which, while critical, potentially overlooks the magnitude of these changes. Predicting the exact magnitude or percentage change is a nuanced and complex challenge. Our current model configuration might not be optimally designed to fully capture these magnitude variations, a significant oversight in volatile markets where the distinction between a 1% and a 10% change can have substantial implications for traders and investors. To enhance the utility and comprehensiveness of our tool for market participants, it would be beneficial to expand our model's capabilities to not only predict the direction of Bitcoin price changes but also the magnitude of these shifts. Developing a model that can accurately forecast both aspects requires a dual-focus approach: refining the model's architecture and fine-tuning its parameters to better adapt to the intricacies of volatile price movements. Beyond these technical enhancements, collaborating with experts in finance and economics can provide invaluable insights. Such interdisciplinary efforts can help in selecting more relevant data sources and refining modeling techniques, bridging the gap between raw data and real-world market dynamics. This synergy aims to make the model's predictions more actionable and relevant for end-users, offering a more comprehensive tool in the fast-paced world of cryptocurrency trading.
- (iv) **Broadening platforms:** While offering detailed insights into Bitcoin, this research encounters a limitation by focusing solely on this single cryptocurrency. The reliance on Bitcoin as the primary subject restricts our understanding to just one facet of the multifarious and ever-evolving cryptocurrency market. Additionally, the sentiment data utilized in this study, sourced from a limited set of platforms, may not comprehensively represent the complex sentiment dynamics that drive the broader market. This is a critical consideration, especially in the fast-paced and varied world of cryptocurrencies, where different platforms may be influenced by distinct factors. Therefore, it becomes imperative to expand our research horizons beyond Bitcoin. By incorporating data from a range of cryptocurrency platforms and considering other major cryptocurrencies like Ethereum and Ripple, we can significantly enhance the depth and relevance of our findings.
- (v) **Practical Application in Trading Scenarios:** While our model has demonstrated proficiency in predicting the direction of Bitcoin price changes using historical data, we acknowledge that real-time trading environments present additional complexities. The current study does not simulate the conditions of live trading, such as

execution delay, transaction costs, and market liquidity. Therefore, a critical area for future work will involve deploying our model in a simulated trading environment to rigorously evaluate its real-time effectiveness. This will allow us to observe the model's performance in the presence of factors intrinsic to live markets, providing a more realistic assessment of its practical utility for traders and investors. Collaborations with finance and economic experts will be sought to further refine our approach, ensuring that our model remains adaptive and robust amidst the rapidly evolving cryptocurrency market. This progression will be instrumental in transitioning our predictive model from a theoretical construct to a viable tool for real-world trading applications.

## VII. CONCLUSION

With the aim of predicting Bitcoin price direction changes, this study aimed to bring forth a more effective and accurate predictive model. Our main objective was to enhance the accuracy of forecasting Bitcoin price direction changes, given its significant implications for traders, investors, and the financial community at large. To achieve this, we initially considered 14 different types of data, spanning a broad spectrum from price and volume data to social media sentiments and trending topics. Such a dataset was considered to capture the nuances of the Bitcoin market. Using the Pearson correlation method, we narrowed down our selection to the seven most price-direction-effective features, namely: Volume of trades, number of transactions, active addresses, options market data, Twitter sentiments, Google trends data, and Twitter trends data. This methodological choice ensured that we focused on the most relevant data, enhancing the reliability and robustness of our predictions.

We employed the DQN algorithm to develop our predictive model. It is chosen for its capacity to manage complex environments and vast datasets. Moreover, we proposed a novel reward function tailored to improve the accuracy of the model and align it more closely with the real-world cases and objectives of our study. While the DQN algorithm has its inherent limitations, it played an important role in our study's success. With an accuracy rate of 95%, our study stands out, surpassing many other studies in the same area. This achievement is a clear reflection of the combined power of targeted data selection, advanced algorithmic techniques, and innovative reward functions.

In wrapping up, this research has set a significant benchmark in the field of cryptocurrency predictions by seamlessly combining diverse data sources with advanced machine learning techniques. Our research underscores the potential benefits of targeted data selection, efficient algorithmic implementation, and thoughtful innovation. We hope our findings provide a foundation upon which future studies can build and critically evaluate this intricate and ever-changing domain.

## REFERENCES

- [1] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," 2008. [Online]. Available: <https://bitcoin.org/bitcoin.pdf>
- [2] A. Narayanan, J. Bonneau, E. Felten, A. Miller, and S. Goldfeder, *Bitcoin and Cryptocurrency Technologies*. Princeton, NJ, USA: Princeton Univ. Press, 2016.
- [3] S. Barber, X. Boyen, E. Shi, and E. Uzun, "Bitter to better-how to make Bitcoin a better currency," in *Financial Cryptography and Data Security*. Berlin, Germany: Springer, 2012, pp. 399–414.
- [4] D. Yermack, "Is Bitcoin a real currency? An economic appraisal," Nat. Bureau Economic Res., Cambridge, MA, USA, 2013.
- [5] CoinMarketCap, *Cryptocurrency Market Capitalizations*, Green, Dover, DE, USA, 2021.
- [6] F. Glaser, K. Zimmermann, M. Haferkorn, M. C. Weber, and M. Siering, "Bitcoin-asset or currency? Revealing users' hidden intentions," in *Proc. ECIS*. 2014, pp. 1–14.
- [7] N. Gandal, J. Hamrick, T. Moore, and T. Oberman, "Price manipulation in the Bitcoin ecosystem," *J. Monetary Econ.*, vol. 95, pp. 86–96, May 2018.
- [8] J. Donier and J. P. Bouchaud, "Why do markets crash? Bitcoin data offers unique insights," *PLoS ONE*, vol. 10, no. 10, 2015, Art. no. e0139356.
- [9] A. Gervais, S. Capkun, G. O. Karame, and D. Gruber, "On the privacy provisions of Bloom filters in lightweight Bitcoin clients," in *Proc. 30th Annu. Comput. Secur. Appl. Conf.*, 2014, pp. 326–335.
- [10] R. C. Phillips and D. Gorse, "Predicting cryptocurrency price bubbles using social media data and epidemic modelling," in *Proc. IEEE Symp. Ser. Comput. Intell. (SSCI)*, Nov. 2017, pp. 1–7.
- [11] I. Guyon and A. Elisseeff, "An introduction to variable and feature selection," *J. Mach. Learn. Res.*, vol. 3, pp. 1157–1182, Mar. 2003.
- [12] P. J. Brockwell and R. A. Davis, *Introduction to Time Series and Forecasting*, 2nd ed. Berlin, Germany: Springer, 2002.
- [13] A. H. Ali, M. G. Yaseen, M. Aljanabi, and S. A. Abed, "Transfer learning: A new promising techniques," *Mesopotamian J. Big Data.*, vol. 2023, pp. 31–32, Feb. 2023.
- [14] M. Moradi, Y. Weng, and Y.-C. Lai, "Defending smart electrical power grids against cyberattacks with deep Q-learning," *PRX Energy*, vol. 1, no. 3, Nov. 2022, Art. no. 033005.
- [15] M. Moradi, Y. Weng, J. Dirkmann, and Y.-C. Lai, "Preferential cyber defense for power grids," *PRX Energy*, vol. 2, no. 4, Oct. 2023, Art. no. 043007.
- [16] S. Carta, A. Ferreira, A. S. Podda, D. Reforgiato Recupero, and A. Sanna, "Multi-DQN: An ensemble of deep Q-learning agents for stock market forecasting," *Exp. Syst. Appl.*, vol. 164, Feb. 2021, Art. no. 113820.
- [17] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, and S. Petersen, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [18] H. Van Hasselt, A. Guez, and D. Silver, "Deep reinforcement learning with double Q-learning," in *Proc. AAAI Conf. Artif. Intell.*, 2016, vol. 30, no. 1, pp. 1–7.
- [19] A. Ferenczi and C. Badica, "Prediction of Ethereum gas prices using DeepAR and probabilistic forecasting," *J. Inf. Telecommun.*, vol. 65, pp. 1–15, Sep. 2023.
- [20] S. Nithiya and G. Parimala, "Bitcoin price prediction using machine learning," in *Proc. Int. Conf. Comput. Commun. Informat. (ICCCI)*, Jan. 2023, pp. 1–6.
- [21] E. Koo and G. Kim, "Centralized decomposition approach in LSTM for Bitcoin price prediction," *Exp. Syst. Appl.*, vol. 237, Mar. 2024, Art. no. 121401.
- [22] K. Zela and L. Saliq, "Forecasting through neural networks: Bitcoin price prediction," *Eng. Appl.*, vol. 2, no. 3, pp. 218–224, 2023.
- [23] P. Peng, Y. Chen, W. Lin, and J. Z. Wang, "Attention-based CNN-LSTM for high-frequency multiple cryptocurrency trend prediction," *Exp. Syst. Appl.*, vol. 237, Mar. 2024, Art. no. 121520.
- [24] Z. Fu, S. Jia, and Z. Guo, "Volatile asset price prediction model based on grey prediction," in *Proc. 4th Manag. Sci. Informatization Econ. Innov. Develop. Conf.*, Chongqing, China, 2023, pp. 1–6.
- [25] Y. Li, "Research on quantifier trading system based on time series," *Academic J. Comput. Inf. Sci.*, vol. 6, no. 3, pp. 102–110, 2023.
- [26] T. Bouazizi, "Unpacking the complexities of cryptocurrency prices volatility in times of crisis: A time series data with long-term memory or long-range dependence," *Qeios*, vol. 322, Mar. 2023.
- [27] M. Puri, A. Garg, and L. Rani, "Cryptocurrency trading using machine learning," in *Proc. 3rd Int. Conf. Advance Comput. Innov. Technol. Eng. (ICACITE)*, Greater Noida, India, May 2023, pp. 2620–2624, doi: 10.1109/icacite57410.2023.10182501.

- [28] W. F. Hidayat, M. F. Julianto, Y. Malau, and A. Setiadi, "Implementation of LSTM and Adam optimization as a cryptocurrency polygon price predictor," in *Proc. Int. Conf. Inf. Technol. Res. Innov. (ICITRI)*, Aug. 2023, pp. 123–127.
- [29] J. Wang, "Application of portfolio price forecasting based on ARIMA-GARCH model," in *Proc. 2nd Int. Conf. Artif. Intell.*, 2023, pp. 296–303.
- [30] O. Sattarov, A. Muminov, C. W. Lee, H. K. Kang, R. Oh, J. Ahn, H. J. Oh, and H. S. Jeon, "Recommending cryptocurrency trading points with deep reinforcement learning approach," *Appl. Sci.*, vol. 10, no. 4, p. 1506, Feb. 2020.
- [31] Y. R. Fadillah and R. Widayarsi, "Application of hybrid LSTAR-GARCH model with expected tail loss in predicting the price movement of Bitcoin cryptocurrency against rupiah currency," *ZERO, Jurnal Sains, Matematika dan Terapan*, vol. 7, no. 1, pp. 33–42, 2023.
- [32] Y. Fu, "GM (1, 1) and quantitative trading decision model," in *Financial Engineering and Risk Management*. Ottawa, ON, Canada: Clausius Scientific Press, 2022.
- [33] D. Gazizova and K. Ageaoli, *Technical Analysis in Cryptocurrency Market*. Pittsburgh, PA, USA: Carnegie Mellon Univ., 2022, doi: [10.48762/2022.554](https://doi.org/10.48762/2022.554).
- [34] A. Sharma, "On-chain analysis and cryptocurrency price forecasting using on-chain metrics," Francis-Press, London, U.K., Tech. Rep., 2023.
- [35] M. Azamjon, O. Sattarov, and J. Cho, "Forecasting Bitcoin volatility through on-chain and whale-alert tweet analysis using the Q-learning algorithm," *IEEE Access*, vol. 11, pp. 108092–108103, 2023.
- [36] D. Herremans and K. Wee Low, "Forecasting Bitcoin volatility spikes from whale transactions and CryptoQuant data using synthesizer transformer models," 2022, [arXiv:2211.08281](https://arxiv.org/abs/2211.08281).
- [37] O. Sattarov, H. S. Jeon, R. Oh, and J. D. Lee, "Forecasting Bitcoin price fluctuation by Twitter sentiment analysis," in *Proc. Int. Conf. Inf. Sci. Commun. Technol. (ICISCT)*, 2020, pp. 1–4.
- [38] B. Amirshahi and S. Lahmiri, "Investigating the effectiveness of Twitter sentiment in cryptocurrency close price prediction by using deep learning," *Exp. Syst.*, vol. 43, Aug. 2023, Art. no. e13428.
- [39] S. Otabeek and J. Choi, "Twitter attribute classification with Q-learning on Bitcoin price prediction," *IEEE Access*, vol. 10, pp. 96136–96148, 2022.
- [40] N. Srivastava, O. Gowda, S. Bulbule, S. Bhandari, and A. Chaturvedi, "Bitcoin evolution analytics: Twitter sentiments to predict price change as bearish or bullish," in *Proc. IEEE Int. Conf. Big Data (Big Data)*, Dec. 2022, pp. 3547–3556.
- [41] Y. Zou and D. Herremans, "PreBit—A multimodal model with Twitter FinBERT embeddings for extreme price movement prediction of Bitcoin," 2022, [arXiv:2206.00648](https://arxiv.org/abs/2206.00648).
- [42] R. Gök, E. Bouri, and E. Gemici, "Can Twitter-based economic uncertainty predict safe-haven assets under all market conditions and investment horizons?" *Technol. Forecasting Social Change*, vol. 185, Dec. 2022, Art. no. 122091.
- [43] A. Qorbani, Y. Rabbani, R. K. Rad, and P. Falsafi, "Predicting Bitcoin price changes using sentiment analysis in social media and celebrities along with data mining," *J. Econ. Model.*, vol. 7, no. 3, pp. 163–182, 2022.
- [44] *CryptoCompare*. Accessed: Sep. 10, 2023. [Online]. Available: <https://min-api.cryptocompare.com/>
- [45] *CoinGecko*. Accessed: Sep. 10, 2023. [Online]. Available: <https://www.coingecko.com/>
- [46] *CoinMarketCap*. Accessed: Sep. 10, 2023. [Online]. Available: <https://coinmarketcap.com/>
- [47] *Blockchair*. Accessed: Sep. 10, 2023. [Online]. Available: <https://blockchair.com/>
- [48] *Blockchain*. Accessed: Sep. 10, 2023. [Online]. Available: <https://www.blockchain.com/>
- [49] *Twitter API*. Accessed: Sep. 10, 2023. [Online]. Available: <https://developer.twitter.com/en/docs/twitter-api>
- [50] *Vader Sentiment*. Accessed: Sep. 10, 2023. [Online]. Available: <https://pypi.org/project/vaderSentiment/>
- [51] *Pytrends API*. Accessed: Sep. 10, 2023. [Online]. Available: <https://pypi.org/project/pytrends/>
- [52] *Tweepy API*. Accessed: Sep. 10, 2023. [Online]. Available: <https://www.tweepy.org/>
- [53] *WallStreetMojo*. Accessed: Sep. 10, 2023. [Online]. Available: <https://www.wallstreetmojo.com/pearson-correlation-coefficient/>
- [54] *SciPy*. Accessed: Sep. 10, 2023. [Online]. Available: <https://scipy.org/>
- [55] *Pandas*. Accessed: Sep. 10, 2023. [Online]. Available: <https://pandas.pydata.org/>
- [56] D. M. Powers, "Evaluation: From precision, recall and F-measure to ROC, informedness, markedness and correlation," *J. Mach. Learn. Technol.*, vol. 2, no. 1, pp. 37–63, Oct. 2011.
- [57] S. A. Basher and P. Sadorsky, "Forecasting Bitcoin price direction with random forests: How important are interest rates, inflation, and market volatility?" *Mach. Learn. Appl.*, vol. 9, Sep. 2022, Art. no. 100355.
- [58] J. V. Critien, A. Gatt, and J. Ellul, "Bitcoin price change and trend prediction through Twitter sentiment and data volume," *Financial Innov.*, vol. 8, no. 1, pp. 1–20, Dec. 2022.
- [59] I. E. Livieris, N. Kiriakidou, S. Stavroyiannis, and P. Pintelas, "An advanced CNN-LSTM model for cryptocurrency forecasting," *Electronics*, vol. 10, no. 3, p. 287, Jan. 2021.
- [60] Z. H. Kilimci, "Sentiment analysis based direction prediction in Bitcoin using deep learning algorithms and word embedding models," *Int. J. Intell. Syst. Appl. Eng.*, vol. 8, no. 2, pp. 60–65, Jun. 2020.
- [61] D. C. A. Mallqui and R. A. S. Fernandes, "Predicting the direction, maximum, minimum and closing prices of daily Bitcoin exchange rate using machine learning techniques," *Appl. Soft Comput.*, vol. 75, pp. 596–606, Feb. 2019.



**AZAMJON MUMINOV** received the B.S. degree in informatics and information technologies from Tashkent University of Information Technologies, Tashkent, Uzbekistan, in June 2015, and the M.S. and Ph.D. degrees from the Computer Engineering Department, Konkuk University, South Korea, in 2018 and 2022, respectively. Since 2022, he has been an Assistant Professor with the Computer Engineering Department, Gachon University, South Korea. His research interests include artificial intelligence, machine learning algorithms, and especially the intersection of reinforcement learning techniques with the deep learning algorithms.



**OTABEEK SATTAROV** received the B.S. degree in information teaching methods from Gulistan State University, Syrdarya, Uzbekistan, in 2018, and the M.S. degree in computer engineering from Konkuk University, Chungju-si, South Korea, in 2021. He is currently pursuing the Ph.D. degree with the AI-Software Department, Gachon University, Seongnam, South Korea. His research interests include machine learning algorithms, deep reinforcement learning, and natural language processing.



**DAEYOUNG NA** (Member, IEEE) received the bachelor's and master's degrees in computer system engineering and the Ph.D. degree in mobile computing and computer engineering from Konkuk University, South Korea, in February 2006, in August 2009, and February 2019, respectively. Since September 2018, he has been an Assistant Professor with the Global Leadership School, Handong Global University, Pohang, South Korea.