

Sequence-Based Target Coin Prediction for Cryptocurrency Pump-and-Dump

SIHAO HU, National University of Singapore, Singapore

ZHEN ZHANG, National University of Singapore, Singapore

SHENGLIANG LU, National University of Singapore, Singapore

BINGSHENG HE, National University of Singapore, Singapore

ZHAO LI, Link2Do Technology & Zhejiang University, China

With the proliferation of pump-and-dump schemes (P&Ds) in the cryptocurrency market, it becomes imperative to detect such fraudulent activities in advance to alert potentially susceptible investors. In this paper, we focus on predicting the pump probability of all coins listed in the target exchange before a scheduled pump time, which we refer to as the *target coin prediction* task. Firstly, we conduct a comprehensive study of the latest 709 P&D events organized in Telegram from Jan. 2019 to Jan. 2022. Our empirical analysis reveals some interesting patterns of P&Ds, such as that pumped coins exhibit intra-channel homogeneity and inter-channel heterogeneity. Here *channel* refers a form of group in Telegram that is frequently used to coordinate P&D events. This observation inspires us to develop a novel sequence-based neural network, dubbed SNN, which encodes a channel's P&D event history into a sequence representation via the positional attention mechanism to enhance the prediction accuracy. Positional attention helps to extract useful information and alleviates noise, especially when the sequence length is long. Extensive experiments verify the effectiveness and generalizability of proposed methods. Additionally, we release the code and P&D dataset on GitHub¹, and regularly update the dataset.

CCS Concepts: • **Security and privacy** → *Web application security*; • **Computing methodologies** → *Machine learning*; • **Information systems** → *Data mining*.

Additional Key Words and Phrases: blockchain security, time-series forecasting, datasets

ACM Reference Format:

Sihao Hu, Zhen Zhang, Shengliang Lu, Bingsheng He, and Zhao Li. 2023. Sequence-Based Target Coin Prediction for Cryptocurrency Pump-and-Dump. *Proc. ACM Manag. Data* 1, 1, Article 6 (May 2023), 19 pages. <https://doi.org/10.1145/3588686>

1 INTRODUCTION

Pump-and-dump (P&D) is a manipulative scheme that attempt to boost an asset price before selling the cheaply purchased assets at a higher price. While its origins can be traced back to the stock market [27], the burgeoning popularity of cryptocurrencies [5, 10, 28] has led to the proliferation of P&Ds within the cryptocurrency realm. Organizing cryptocurrency P&Ds typically took place in encrypted messaging platforms such as Telegram, where organizers create *pump channels* to attract speculators. For a planned pump, its organizer posts an announcement several days in advance,

¹<https://github.com/Bayi-Hu/Pump-and-Dump-Detection-on-Cryptocurrency>

Authors' addresses: Sihao Hu, National University of Singapore, Singapore, hushihao26@gmail.com; Zhen Zhang, National University of Singapore, Singapore, zhen@nus.edu.sg; Shengliang Lu, National University of Singapore, Singapore, lusl@nus.edu.sg; Bingsheng He, National University of Singapore, Singapore, dcsheb@nus.edu.sg; Zhao Li, Link2Do Technology & Zhejiang University, China, lzjoey@gmail.com.



This work is licensed under a Creative Commons Attribution International 4.0 License.




Fig. 1. The \$NAS pump-and-dump coordinated by the Telegram channel "Binance Trading Signals".

clearly stating the target date, time, and exchange. As the scheduled time approaches, organizers use a *countdown* strategy to intermittently remind members, without disclosing the target coin until the scheduled time. As illustrated in Figure 1, a P&D coordinated by the Telegram channel "Binance Trading Signals" (a channel with more than 300,000 subscribers) took place on Binance on August 22, 2021, at 17:00 UTC. After the coin NAS was disclosed, the price inflated immediately and reached 250% in one minute, before dropping shortly afterward. Although the entire process lasted only five minutes, a staggering \$27 million was transferred between participants.

As P&D schemes occur more frequently and attract greater public attention, it is imperative to detect such fraudulent activities [13, 14, 30] and inform susceptible investors before they fall victim. Existing efforts [16–18, 22] mainly focus on P&D post-detection task, which aims to detect whether a P&D has happened or is happening. However, we argue that this task fails to meet practical needs, as P&Ds typically occur rapidly, leaving no time to alert investors. In this paper, we concentrate on the *target coin prediction* task, *i.e.*, to predict the pump likelihood of all coins listed in the target exchange one hour before the pump time, given the exchange and pump time announced in advance. The target coin prediction task is an interesting and challenging data science problem. To achieve accurate predictions, we rely on heterogeneous data sources, including historical statistics of trading price and volume, as well as text data that collected from Telegram. Different data sources require different tools and algorithms to extract useful and timely information.

As shown in Figure 2, we propose an efficient data science pipeline consisting of two main stages: data collection and target coin prediction. The data collection stage targets on identifying as many historical P&Ds as possible. It explores pump channels based on verified seed channels, gathers tremendous text messages from Telegram, and identifies important pump messages through keyword filtering and machine learning algorithms. The data collection stage works offline, maintains a P&D dataset, and updates it regularly.

Given the pump messages identified in the data collection stage, the target coin prediction stage aims to make effective and efficient predictions. It collects multi-modal features from heterogeneous data sources, including historical trading prices and volumes from cryptocurrency market statistics, and coin embeddings pre-trained on the large text corpus collected from Telegram. We then generate features to feed into the prediction model and obtain the pump probability of all coins listed on the target exchange one hour before the scheduled pump time. The entire process of target coin prediction can achieve real-time efficiency to ensure the timeliness.




Fig. 2. The data science pipeline consists of two stages: data collection and target coin prediction.

Empirically, our data pipeline filters 709 P&Ds out of 4,674,822 Telegram messages posted from Jan. 2019 to Jan. 2022. By analyzing them we discover that: 1) Coins with middle-cap value and high discussion degree in social media and forums are more likely to be targeted, indicating that the choice of the coin is well thought out; 2) The buy-in behavior of organizers starts within 60 hours before the scheduled pump time, and the pre-pump phenomenon confirms the existence of the hierarchy of Telegram channels. Precursors like the boost in price and trading volume caused by insiders (organizers and VIP members) can be a powerful signal for detecting P&D coins in advance; 3) Coins pumped by the same channel exhibit homogeneity in both terms of statistics and semantics, whereas, coins selected by different channels exhibit distinct heterogeneity.

Prior to our work, P&Ds are treated as isolated, unrelated events [23, 29]. However, our analysis show that a channel's pump coin selection follows some specific patterns and can be used to predict the future coin. Therefore, we construct channels' pump history as sequences and use sequential model to encode these sequences into embedding vectors to represent their coin selection strategies/preferences. We propose an effective Sequence Neural Network (SNN) with the core design of a positional attention module, which can directly capture the skip-correlation in one layer without serial computation as RNNs, and decouple the constraint of CNNs in the depth of layers and length of sequences. Secondly, we identify and address the *coin-side cold start problem* that are largely overlooked by previous works but is essential in the practical settings.

Experiments on a real-world P&D dataset show that SNN outperforms its competitors with a large margin. To further verify the generalizability of the data science pipeline, we extend it to a different task, *i.e.*, cryptocurrency price forecasting. Experiments show that our methodology is extensible and well-suited for other tasks that involves multiple sources of sequential data.

Contributions: To summarize, the contributions are as follows:

- **Pipeline.** We build a data science pipeline that extracts useful and timely information from heterogeneous data sources, to support effective and efficient target coin prediction.
- **Analysis.** We present a comprehensive study on the latest P&Ds, and uncover some interesting characteristics of P&Ds such as intra-channel homogeneity and inter-channel heterogeneity.
- **Model.** We propose SNN with a positional attention module to capture the pump history of channels when predicting the target coin. We identify and address the coin-side cold-start problem in practical settings.

2 A TYPICAL PUMP-AND-DUMP PROCESS

Set up: A successful pump requires a large number of participants to buy the target cryptocurrency at the same time. Therefore, the organizer would establish a public channel, which is a specific form of group chat that only allows the organizer to post messages. Then the organizer recruits as many subscribers as possible by posting invitation links on social media like Twitter, or other channels for advertisement. Moreover, we notice that many pump organizers create another private VIP channels, which only accept paid members to subscribe. In return, VIP members will be informed of the target coin name minutes or even hours before the coin name is released on the public channel.

Pump Announcement: When the public channel gets enough members, the pump organizer releases a pump announcement several days before a pump activity. The announcement usually includes three important information: the precise time to initiate the pump, the target exchange, and the pairing coin. As the scheduled pump time approaches, organizers keep reminding members of the pump event by counting down. The counting frequency gradually increases from once a day to once a few hours, and then once a few minutes before the pump time. During this process, the organizers advise members to transfer sufficient pairing coin to the exchange and post some "pump rules" like "Do not sell immediately".

Pre-pump: VIP members in private channels know the target coin name in advance, so they can buy coins at a lower price and wait to sell them at a higher price at the pump time. If they buy a large volume of coins in a short time, it can lead to a price hike before the pump time, which we call "pre-pump."

Pump: Upon a few minutes before the pre-arranged pump time, the organizers will end the countdown with a typical message like: "The next message will be the coin name!" The coin name is then posted at the pump time, usually in the uppercase symbol name of the coin, such as "FIC," and occasionally in the format of an OCR-proof image to prevent automatic recognition. Then the coin's price immediately surges, often reaching several times the original price within one minute.

Dump: Although the organizer may urge members to hold onto the coin and not sell it immediately, participants typically know that the price will quickly drop and do not want to sell later than others. Thus, just a few moments after the price reaches its peak, it starts to drop, causing more participants to panic-sell in turn. Eventually, the coin price will bounce back to the original value or even lower after the P&D.

3 DATA COLLECTION

3.1 Channel Exploration

We explore the pump channels in Telegram, which is currently the largest platform for P&D schemes to thrive [24, 29]. Firstly, we collect a total of 1,142 channels from PumpOlymp [3], a website that provides verified P&D channels in Telegram. Then we use the Telethon API [4] to check their status. We find that nearly half of channels (564/1,142) have been deleted due to inactivity for over six months. Secondly, we adopt the *snowball* strategy [24] to explore more pump channels, based on an observation that pump channel organizers sometimes post invitation links across other channels to attract more participants. Specifically, we collect all historical messages posted in seed channels and extract the invitation links from these seed pump channels to explore new channels. To ensure high relatedness, we set the exploration hop to 2. After filtering out the expired links, we get 137 channels that are still in active and not duplicated with the seed channels. We retrieve a total of 4,674,822 messages that are posted on these 715 channels between Jan. 1, 2019, and Jan. 13, 2022.

Table 1. Model performance on pump message detection

Model	AUC	Precision	Recall	F1
LR	0.988	0.892	0.913	0.902
RF	0.994	0.901	0.939	0.920

Table 2. Statics of P&D dataset

# Sample	# Event	# Channel	# Coin	# Exchange
1,335	709	108	278	18

Table 3. Examples of P&D dataset

Channel	Coin	Exchange	Pair	Timestamp
1394568162	EVX	Binance	BTC	2021-10-24 17:01:59
1394568162	BRD	Binance	BTC	2021-10-31 15:01:30
1394568162	MTH	Binance	BTC	2022-01-13 17:00:19
...	...	...	...	...
1250076602	EVX	Binance	BTC	2021-10-24 17:00:07
1250076602	EZ	Binance	BTC	2022-01-05 19:00:05

3.2 Pump Message Detection

The amount of Telegram messages is so tremendous that manual annotation becomes prohibitively expensive and time-consuming. Fortunately, the majority of pump-relevant messages conform to specific patterns that can be readily discerned by humans, rendering them identifiable by machine learning algorithms.

The workflow for pump message detection is illustrated in Figure 2: Firstly, a keyword matching strategy is employed to reduce the proportion of non-pump messages in the text dataset, which reserves any message that mentions a coin or exchange name, or includes keywords such as "pump," "target," "hold," "sell," etc. This simple but useful strategy selects 2,193,549 messages out of 4,674,822 messages. Secondly, we randomly select and label 5,050 messages as ether pump and non-pump messages. Pump messages are labeled as those including pump announcements, pump countdowns, target coin releases, and post-pump reviews, while all other messages are labeled as non-pump messages. Subsequently, we remove punctuation marks, stop words, URLs, and emojis from labeled messages and represent each message as a vector using the TF-IDF technique. We separate 70% samples to train a Random Forest (RF) model and a Logistic Regression (LR) model, and test them on the 30% left samples.

The results are presented in Table 1, where metrics are calculated using a relatively low threshold of 0.2 to identify as many pump messages as possible. Because a single pump message is insufficient to identify a P&D event, we aggregate messages from into *sessions* if the time interval between two adjacent messages is less than 24 hours. In this case, a session is the minimum time unit for a channel to hold a P&D, and we assume that a P&D occurs at most once in a session (which is verified by us manually). Finally, we identify 1,335 P&D samples from total 2,006 sessions based on the 37,525 pump messages identified by the RF model. Each P&D sample is represented by a

quintuple (channel_id, target, exchange, pairing coin, timestamp), as exemplified in Table 3. A summary of the dataset statistics is presented in Table 2.

4 PUMP-AND-DUMP SCHEMES ANALYSIS

This section presents analytical studies at the coin, event, and channel levels, aimed at answering the following research questions:

- **Q1:** What types of coins are susceptible to P&D schemes?
- **Q2:** Is it possible to detect P&D schemes in advance?
- **Q3:** Do the target coin selection strategies differ across channels?

4.1 Coin-Level Analysis (Q1)

To investigate Q1, we utilize the CoinGecko [2] API to collect historical daily statistics on coins, such as market capitalization, volume data, and social media indices, etc. For pumped coins, we collect the data from three days preceding the pump time, as this data is generally more stable before P&Ds. For a fair comparison, we randomly retrieve historical data from Jan. 1, 2019, to Jan. 13, 2021 for the top 4,000 coins ranked on the CoinGecko website (almost all pumped coins are ranked above 4000). Figure 3 demonstrates the distribution of pumped coins and the top 4,000 coins based on several features. Due to space limitations, we only report on four features: market capitalization, Alexa rank (an index representing global web popularity), the number of Reddit subscribers, and the number of Twitter followers.




Fig. 3. Analysis of four features in distributions of pumped coins and the top 4000 coins.

From Figure 3(a) and (b), we observe that the market capitalization and Alexa rank distributions of pumped coins are most similar to those of coins ranked in the top1001-2000 range, suggesting that P&D organizers tend to target middle-cap coins. This is because it is much more challenging to manipulate the prices of large-cap coins [19, 29], especially for pump channels with a small number of participants. Additionally, organizers rarely select coins with low market capitalization and web popularity (Alexa rank), which typically reflect low trading liquidity. For example, if the market is

frozen, there will be no other traders being attracted to get involved and transactions will be limited to the group, making common members hesitant to engage in the pump due to the likelihood of losing money. Another observation, illustrated in Figure 3(c) and (d), is that the distributions of numbers of Reddit subscribers and Twitter follower are similar to those of coins ranked in the top-1-1000 range with minor variances, indicating that pumped coins have a significant presence on social media platforms. According to our statistics, the likelihood of a coin being pumped again is 60.1%, indicating that even with data collect three days prior to the pump time, the coin might have already undergone a previous pump. Once a coin is pumped, it is widely discussed on social media, and pump participants spread misinformation about the coin to lure potential victims.

4.2 Event-Level Analysis (Q2)

Event Distribution Drift across Exchanges: Out of the 709 P&D events in our dataset, **62.8%** (445) took place in Binance, 20.6% (146) in Yobit, 8.7%(62) in Hotbit and 3.0%(21) in Kucoin, which differs significantly from previous literature: Xu et al. [29] report 51% of P&Ds took place in Cryptopia, 27% in Yobit, and only **17%** in Binance from October 2018 to February 2019; Li et al. [19] report 52.6% of P&Ds happened in Bittrex, 32.2% in Yobit, and only **15.2%** in Binance from May 2017 to August 2018; It is noteworthy that the proportion of P&Ds occurring on Binance has drastically increased (from 15.2% to 62.8%) in the past few years. We identify two key factors that lead to this drift: firstly, Bittrex's ban on P&Ds in Nov. 2017 [7], and secondly, Cryptopia's liquidation following a severe hack in Jan. 2019 [8]. Additionally, Binance's advantages in terms of coin number, user volume, and commissions have made it the favored exchange for pump organizers.

The drift has impacted on the pattern of P&D attacks. Previous studies show that attacks on Binance typically have the lowest average return rate, approximately 29% of that on Yobit [19, 29], and involved an average of 1.42 channels per attack. This is likely due to the fact that Binance's trading volume is much larger than that of other exchanges, necessitating a greater number of participants for a successful P&D. However, our analysis shows that each P&Ds on Binance involves, on average, 2.25 channels, suggesting that pump organizers improve return rates by coordinating multiple channels for a single pump event.

Precursors to P&Ds: After the organizer decides which coin to pump, they gradually purchase the coin and release the name of the coin in the private channel, allowing VIP members to buy in before the scheduled pump time. The purchase behavior of organizers and VIP members can cause significant market movements for middle-cap coins. To study the precursors to P&Ds, we plot the time series of prices and trading volumes 72 hours before and 24 hours after the pump time in Figure 4(a) and (b). The time series are averaged from 445 pump events that occurred on Binance and were paired with BTC. For each P&D, we select the earliest announce time across multiple channels as the pump time if the event is organized by multiple channels, and retrieve historical OHLCV (open, high, low, close, volume) data for each minute using the Binance API [1].

One prominent finding is the gradual increase in coin price tens of hours prior to the pump time, as observed in Figure 4(a). In Figure 4(b), we note that the start point of frequent trading is about 57 hours before the pump time, with trading volume remaining relatively low before this point. Correspondingly, the average coin price in Figure 4(a) gradually starts to rise. We also notice that several hikes appear from 48 hours to 1 hour before the pump time, very similar to the largest trading volume hike around the pump time. We posit that these hikes are caused by the pre-pump events in which VIP members purchase a large number of coins in a short time. Figure 4(d) gives a verified pre-pump in a VIP channel 5 hours before the pump time. The observed boost in price and trading volume can be utilized as a powerful signal to detect P&Ds in advance. To demonstrate this, we calculate the average returns within a time window from $x + 1$ hours to 1 hour before the pump time ($x = 1, 3, 6, 12, 24, 36, 48, 60, 72$). We compare the results with randomly selected coins and




Fig. 4. The observational study of P&D events.

timestamps, calculating the statistics in the same manner. As presented in Figure 4(c), the highest average return reaches 9.5% when $x = 60$, and the average returns of randomly selected coins are close to zero.

4.3 Channel-Level Analysis (Q3)

To answer Q3, we investigate coin selection strategies of different channels in terms of several characteristics. Figure 5(a) presents a scatter plot of the market capitalization of the pumped coins in different channels. Each dot in the figure represents a pumped coin, with the channel index on the X-axis and the market capitalization value on the Y-axis. The first observation is that the market capitalization of coins pumped by a specific channel exhibits homogeneity (similarity) within a specific range, while exhibits heterogeneity across different channels. The underlying reason for this phenomenon is straightforward: the success of a pump event heavily relies on the number of participants. A channel with larger population of members generally has access to a wider selection of coins to pump and tends to target coins with larger market capitalization to generate higher trading volume, which is difficult feat for smaller channels to achieve. Figure 5(b) and (c) show the Alexa rank and number of Reddit subscribers of the pumped coins respectively, which exhibit similar effects as Figure 5(a).

Semantic Similarity: We further investigate whether coins pumped by the same channel tend to share similar semantics. Specifically, we pre-train SkipGram embeddings [21] on a huge cryptocurrency-related corpus and extract the word embedding of the coin symbol to represent its semantics. Figure 6




Fig. 5. Scatter plots of three pumped coin features across different channels.




Fig. 6. Distributions of semantic similarity for coin pairs selected using three strategies.

illustrates the cosine similarity distributions of pairs selected under three distinct strategies: 1) selection from the same channel's pump history; 2) selection from all the pumped coins; 3) random selection from all available coins. Evidently, the semantic similarity distribution of coins pumped by the same channel exhibit a significant high average value (0.92) with minor variance compared to pairs selected with strategy 2 (with average value of 0.80) and strategy 3 (with average value of 0.72). This empirical study further proves that a channel has its distinct coin selection pattern.

Summary: We summarize the key findings:

- **A1:** Coins with middle capitalization and high levels of social media discussion are more likely to be targeted.

- **A2:** P&Ds are predictable as the purchase behavior of organizers and VIP members can lead to significant market movement before the scheduled pump time.
- **A3:** Each pump channel has its distinct coin selection pattern that varies from other channels. The coins pumped by a given channel exhibit homogeneity in both statistics and semantics.

5 TARGET COIN PREDICTION

Inspired by our analysis, we in this section present a sequence-based neural network to leverage channels' historical coin selection information to predict the target coin. We also identify and address the coin-side cold-start problem to further improve the performance.

5.1 Feature Generation

The features can be categorized into three groups: channel, target coin, and sequence. Regarding channel features, we collect *channel_id* and *subscriber_num*; As for the target coin, we collect *coin_id* and stable statistics (*market_cap*, *price*, *volume*, *Alex_rank*, *Twitter_follower*, *Reddit_subscriber*, etc.) three days prior to the pump event. We further calculate market movement signals (*price*, *return*, *volume*, *trade_cnt*, etc.) within $(x+1)$ hour to 1 hour prior to the pump time with varying time window ranges ($x = [1, 3, 6, 12, 24, 48, 60, 72]$). For sequence features, we group pumped coins by *channel_id* and sort them chronologically. Features like *coin_id* and stable statistics are incorporated for pumped coins in the sequence.




Fig. 7. Model architecture of SNN

5.2 Sequence Neural Network (SNN)

In SNN, we develop a sequence encoder that can extract a channel's historical pump sequence as a embedding vector to represent its coin selection pattern. As shown in Figure 7, SNN mainly consists of several parts:

Embedding Layer: The input feature contains categorical features, e.g., channel id and coin id, which cannot be directly used due to their high dimensionality. We adopt embedding techniques to embed the sparse features into low-dimensional dense vectors, which significantly eases computing. To reduce the redundancy of parameters, we make the sequence *coin_id* and target *coin_id* share

the same latent space. The generated embeddings are concatenated with other numeric features together to generate the overall features for the channel, target coin as follows:

$$\mathbf{h}_c = h_c^1 \oplus h_c^2 \oplus \dots \oplus h_c^X \quad (1)$$

$$\mathbf{h}_t = h_t^1 \oplus h_t^2 \oplus \dots \oplus h_t^Y \quad (2)$$

where h_c^i and h_t^i are the i -th feature attached to $\mathbf{h}_c$ and $\mathbf{h}_t$, respectively. $\oplus$ is the concatenation operator, X and Y are the number of features of channel and target coin. For coins in the sequence, we generate the whole sequence representation $\mathbf{h}_s$ by the positional attention module.

Positional Attention: A coin in a sequence has several features that may exhibit different sequential patterns. For example, we find that a channel might pump a specific coin periodically but never pump the coin continuously because otherwise it will be easily guessed by others. However, the organizer is likely to choose another coin with similar market capitalization for the next pump event, because the number of participants in the group remains stable over a period of time. This suggests the existence of two distinct patterns within the sequence: 1) temporal proximity pattern, suggesting that feature values that are closer in time are more likely to be correlated; 2) skip-correlated pattern, in which the most related timestamp is not the closest one but the previous ones, representing the periodicity or temporal delay.

Modeling skip-correlation is nontrivial in our scenario as this pattern widely exists in features collected from multiple data sources. Nevertheless, we find existing RNN-based methods are not well suited because the serial computation of RNNs makes them hard to retrieve previous information, especially when the sequence length and span of skip correlation are both large; Moreover, CNN-based models (TCN [6], WaveNet [25]) require relative deep layers to cover the whole sequence as the sequence length goes longer. During this process, convolution operations across different features over and over again may hurt the expression of the pattern of a specific feature.




Fig. 8. Positional attention mechanism

To capture skip-correlation and distinguish sequential patterns for different features to prevent interference, we propose a simple but effective encoder named *positional attention*. As presented in Figure 8, we utilize $P_1, P_2, \dots, P_N$ to designate the 1, 2, ..., N -th position of entities (coins) in the sequence, with each entity having a total of K features denoted as $F_1, F_2, \dots, F_K$. For ease of explanation, we use $F_{i,j}$ to indicate the j -th feature of P_i . For each feature F_j , we generate its

positional attention vector α_j as follows:

$$\mathbf{a}_j = (f(a_{1,j}, a_{2,j}, \dots, a_{N,j})) \quad (3)$$

$$\alpha_j = \text{Softmax}(\mathbf{a}_j) \quad (4)$$

where $a_{1,j}, a_{2,j}, \dots, a_{N,j}$ are zero-initialized learnable parameters, and $f(\cdot)$ is an adjustable mapping function ($\mathbb{R}^N \xrightarrow{f} \mathbb{R}^N$) such as an MLP layer. Since the $\text{Softmax}(\cdot)$ is computed across different positions, the attention vector can take positional information into account. Subsequently, we multiply each element $\alpha_{i,j}$ in the attention vector with its corresponding P_i 's j -th feature $F_{i,j}$ as shown below:

$$\begin{aligned} h_j &= \alpha_j \cdot (F_{1,j}, F_{2,j}, \dots, F_{N,j})^T \\ &= \sum_{i=1}^N \alpha_{i,j} \cdot F_{i,j} \end{aligned} \quad (5)$$

As illustrated in Figure 8, for each feature F_j , $j \in \{1, 2, \dots, K\}$, we initialize multiple attention vectors and repeat the attention computation multiple times to generate C_j *multi-channel* of h_j , denoted as $\mathbf{h}_j = (h_j^1, h_j^2, \dots, h_j^{C_j})$. This approach can increase the capacity of the positional attention, as some of the features may exhibit both of the two aforementioned patterns. Unlike the number of channels for filters in CNN-based models, the choice of channel number in our positional attention is independent for different features. In practice, we find that this hyper-parameter can be set larger for those non-skip-correlated features, so that besides capturing its major pattern, some of the channels can also learn the minor, skip-correlated patterns. Finally, we flatten and concatenate the multi-channel vectors of all the features to generate the final representation vector $\mathbf{h}_s$ for the entire sequence:

$$\mathbf{h}_s = \mathbf{h}_1 \oplus \dots \oplus \mathbf{h}_j \oplus \dots \oplus \mathbf{h}_K \quad (6)$$

where $\mathbf{h}_j$ is the multi-channel representation vector of F_j .

MLP Layer: The embedding and attention layers are primarily based on linear projections. We utilize several fully connected layers and activation functions to endow the model with non-linearity. The output $\hat{y}$ represents the predicted pump probability:

$$\hat{y} = \text{Sigmoid}(\text{MLP}(\mathbf{h}_c \oplus \mathbf{h}_t \oplus \mathbf{h}_s)) \quad (7)$$

Loss Function: The objective function used in SNN is the negative log-likelihood function defined as:

$$L = -\frac{1}{|\mathcal{D}|} \sum_{(y, \hat{y}) \in \mathcal{D}} (y \log \hat{y} + (1 - y) \log(1 - \hat{y})) \quad (8)$$

where $\mathcal{D}$ is the training set, and $y \in \{1, 0\}$ is the ground-truth label that denotes whether it is a P&D or not.

Summary: we summarize the key advantages of SNN:

- **D1:** SNN can capture skip-correlation directly in one layer without the need for serial computation as RNNs and decouples the depth of layers from the length of sequences covered by CNNs.
- **D2:** SNN distinguishes the sequence patterns at the feature-level, reducing interference between different features.
- **D3:** The time complexity of SNN is $\mathcal{O}(N \cdot K \cdot C)$, ($N \gg K, C$), and the attentions for K features and C channels can be calculated in parallel, making it computationally efficient.




Fig. 9. Comparing the ℓ_1 norm distributions of two types of coin id embeddings in training and testing sets.

5.3 Coin-Side Cold-Start Problem

To ensure consistency with real-world applications, we evaluate our method on a testing set consisting of samples all collected after the training set. This strategy also prevents any future information leakage, since the sequences of some latter samples may include label information of former samples. Under such practical settings, we observe the occurrence of the coin-side cold-start problem during end-to-end training. Similar to the item-side cold start problem in the recommendation field [12], the coin-side cold-start problem occurs in two cases: 1) coins that are pumped in the testing set are never be pumped in the training set; 2) coins in the testing set are never exist in the training set. Coins falling into either of these categories can cause the model to struggle in classifying them correctly due to the lack of robust representations to represent the coin appropriately.

Figure 9(a) shows the ℓ_1 norm of coin_id embedding of positive (pumped) and negative coins in the training set, revealing a pronounced difference between two distributions; Figure 9(b) illustrates the ℓ_1 norm of coin_id embeddings of four types of coins in the testing set. *Untrain* denotes the coins never exist in the training set, thus their coin_id embeddings remaining the same as random initialization. *Positive1* represents positive coins that are pumped in the training set, while *Positive2* represents the distribution of positive coins that are never pumped in the training set, which exhibit a similar distribution to that of the negative coins. The distributions of the *untrain* and *positive2* coins correspond to the two cases mentioned before and are the primary factors of the cold-start problem.

To address the cold-start problem, we propose to use the word embedding of a coin symbol to replace its coin_id embedding learned during end-to-end (E2E) training. Word embeddings learned on a large corpus can cover almost all of the coin symbols, making embedding vectors sufficiently trained and containing semantic information. Specifically, we utilize well-known word

embedding techniques, *e.g.*, SkipGram [21], CBoW [20], to pre-train the embeddings on a large corpus of Telegram messages we collected from the cryptocurrency-related channels and groups, not limited to the pump channels. Figure 9(c) and Figure 9(d) shows the ℓ_1 norm distributions of SkipGram embeddings of positive and negative coins in the training set and testing set, respectively. Apparently, two distributions are consistent in the two sets, indicating that using word embedding can alleviate the coin-side cold-start problem.

6 EXPERIMENTS

6.1 Experiment Settings

Dataset: We limit our experiments on predicting coins that were pumped on Binance and paired with BTC. Following this setting, 948 samples are selected (71%) out of the total 1,335 samples as the positive samples. For each positive sample, we label other eligible coins listed on Binance at the time of pump events as negative coins and generate features for them to construct the negative samples. We split the training, validation, and testing sets by the timestamp "2021-01-19 00:00:00" and "2021-05-10 00:00:00". Table 4 presents the statistics of the dataset used in the experiment. The difference between positive and negative ratios across three sets is due to the varying numbers of coins listed on Binance during different periods.

Table 4. P&D dataset description

Pumped?	Train	Validate	Test	Total	Ratio
TRUE	648	100	200	948	0.48%
FALSE	106,900	24,666	64,099	195,665	99.52%
Total	107,548	24,766	64,299	196,613	100.0%

Competitors: We compare the proposed SNN with two types of baselines. The first type is machine-learning methods using handcrafted features, such as Logistic Regression (LR) and Random Forest (RF). We use the mean encoding technique to compensate for the lack of embedding layers in these models. The second type of methods are deep-learning algorithms including a DNN using the same features as SNN except the sequence, RNN-based models such as LSTM, BiLSTM, GRU and BiGRU, and CNN-based method TCN [6]. For RNN models, the hidden dimension of cells is set to 32; For TCN, the depth of convolution layer is set to 3 with 16 channels per layer, the kernel size is set to 4 to cover a 20-length sequence. For SNN, the number of channel is set to 8.

Evaluation Metrics: We employ the Hit Ratio (HR@ k) to measure the model performance. Specifically, we rank a positive sample and its corresponding negative samples as a list according to pump probabilities predicted by the model and use HR@ k to indicate whether the ground truth coin is included in the top k samples. The reported HR@ k is averaged across all the lists with $k = 1, 5, 10, 20, 30$. HR@ k is well suited to evaluate a model in a real-world setting.

6.2 Performance Comparison

Table 5 summarizes the results of all the competitors on the testing set. From the table, we can observe that: 1) Among sequential models, RNN-based models (LSTM, BiLSTM, GRU and BiGRU) perform worse than TCN and SNN, since RNN-based models fail to capture the skip-correlated pattern due to their serial computation paradigm; 2) SNN works better than TCN and other competitors by a margin, which can be attributed to its ability of reducing interference between different features. Although TCN can capture skip-correlated patterns to some extent, it requires relatively

deep layers to conduct hierarchical convolutions to cover the entire sequence and convolution operations across different features over and over again may hurt the expression of the pattern of a particular feature; 3) Comparing SNN with DNN, the improvements of **10.5** Absolute Percentage (AP) on HR@3, and **9.8** AP on HR@10 suggest that modeling channel's pump history can bring a huge benefit to this task. This performance boost can be easily extended to any other non-sequential methods, *e.g.*, traditional ML models, by incorporating sequence representations extracted by a trained SNN.

Table 5. Performance comparison for target coin prediction

Metric	LR	RF	DNN	LSTM	BiLSTM	GRU	BiGRU	TCN	SNN
HR@1	0.156	0.189	0.225	0.207	0.203	0.229	0.163	0.256	0.260
HR@3	0.269	0.348	0.278	0.339	0.344	0.339	0.335	0.348	0.383
HR@5	0.322	0.417	0.383	0.423	0.396	0.414	0.401	0.427	0.465
HR@10	0.449	0.537	0.498	0.551	0.546	0.529	0.555	0.573	0.596
HR@20	0.608	0.687	0.626	0.648	0.630	0.626	0.678	0.692	0.727
HR@30	0.662	0.731	0.727	0.696	0.696	0.714	0.709	0.770	0.797

6.3 Coin Embedding Test

Table 6 summarizes the results of six approaches on the testing set. E2E means a DNN model that only takes coin_id as input, while CW and SG are the DNN models trained on CBoW and SkipGram embedding, respectively. SNN_C and SNN_S share the same architecture as SNN but the original coin_id embeddings are replaced with CBoW and SkipGram embedding, respectively.

It is not surprising to see E2E gives the worst performance since it suffers from the cold-start problem. In comparison, both CW and SG show much superior performance, especially with SG achieving the HR@3 of 0.115, suggesting that well-trained word embedding can alleviate the cold-start problem. This conclusion is also confirmed by the observation that SNN_C and SNN_S achieve significant positive lifts compared to the original SNN.

Table 6. Evaluation of coin embeddings

Metric	E2E	CBoW	SG	SNN	SNN _C	SNN _S
HR@1	0.0	0.035	0.043	<u>0.260</u>	0.256	0.277
HR@3	0.0	0.090	0.115	<u>0.383</u>	0.391	0.414
HR@5	0.013	0.133	0.176	<u>0.465</u>	0.499	0.513
HR@10	0.057	0.253	0.286	<u>0.596</u>	0.617	0.623
HR@20	0.101	0.362	0.376	<u>0.727</u>	0.731	0.739
HR@30	0.242	0.472	0.487	<u>0.797</u>	0.806	0.823

6.4 Sequence Pattern Visualization

We try to reveal some meaningful sequence patterns by visualizing the positional attention vectors of several features learned during training. Figure 10(a) presents the results of attention vectors

for the target coin prediction task, where $F_1 \sim F_6$ are *coin_id*, *volume*, *price*, *Twitter_follower*, *market_cap*, and *Alexa_rank*, respectively. P_1 is the temporal closest position and P_{30} is the farthest position. The darker the color, the higher the attention score. We observe that F_1 to F_4 exhibit skip-correlated patterns, which supports our design of positional attention module to capture this type of feature pattern. F_5 and F_6 exhibit the temporal proximity pattern, indicating that organizers tend to pump coins with similar market cap and web popularity that close to the coins pumped last few times.




Fig. 10. Positional attention visualization w.r.t. different features

7 GENERALIZABILITY OF METHODOLOGY

To assess the generalizability of our methodology, we extend the data science pipeline to a different task, namely sentiment-enhanced cryptocurrency price forecasting. In Telegram trading groups, we observe that investors share news and predict market movements. Their sentiment can be used to predict the market movement, or the market movement might be affected by their sentiment because their subsequent trading operations.

For this task, we use Bitcoin as an example, and the methodology can be easily extended to other coins. Following the same pipeline for target coin prediction, we explore several Telegram chat groups and collect user messages posted from Feb. 7, 2020, to May 31, 2022. Secondly, we apply a data filtering strategy to select Bitcoin-related messages. We extract sentiment information by VADER [15] and calculate sentiment statistical features within each hour (*avg_score*, *neg_avg_score*, *neg_num*, *pos_avg_score*, *pos_num*, *message_num*, etc). Finally, we construct a 200-hour-length sequence for each sample and use hourly price and sentiment features as the sequence feature. The label is Bitcoin's average price in the future 48 or 96 hours. The dataset description for this task is presented in Table 7. We adopt Mean Absolute Error (MAE) as the object loss during training, which is defined as:

$$L = -\frac{1}{|\mathcal{D}|} \sum_{(\hat{y}, y) \in \mathcal{D}} |y - \hat{y}| \quad (9)$$

where $\mathcal{D}$ is the training set, y is the ground-truth label and $\hat{y}$ is the prediction value. In this task, SNN solely takes the sequence features as input for price forecasting.

7.1 Experiment Settings

Dataset: We collect Bitcoin's hourly price from Feb. 7, 2020, to May 31, 2022, which is the same period as the Telegram message data. We split the training and testing sets by the timestamp "2021-12-19 00:00:00". For each sample, we construct a 200-hour sequence with extracted sentiment features and hourly price. The label is the average price of Bitcoin in the future 48 or 96 hours, as predicting the price in the future 1 hour is considered too easy.

Competitors: We use the same sequential competitors as the target coin prediction task. For TCN, we set the depth of convolution layer to 5 with 16 channels per layer, the kernel size is set to 8 to cover a 200-length sequence. For SNN, we set the channel number to 16 for *hour_price*, and

Table 7. Statics of BTC price forecasting dataset

# Message	# BTC msg	# Pos msg	# Neg msg	# Train	# Test
2,799,669	229,595	88,512	54,175	15,856	3,964

Table 8. Performance comparison for BTC price forecasting

Span	Metric	LSTM	BiLSTM	GRU	BiGRU	TCN	SNN
48h	MAE(P)	871.21	810.87	851.30	812.45	820.32	805.49
	MAE(P+T)	848.29	785.66	814.68	791.89	860.75	756.90
	Impr.	22.92	25.21	36.62	20.56	-40.43	48.59
96h	MAE(P)	1144.23	1078.13	1126.37	1049.85	1059.36	1051.57
	MAE(P+T)	1118.84	1043.70	1088.25	1027.45	1048.53	964.27
	Impr.	25.39	34.43	38.12	28.40	10.83	87.30
-	Cost	4.68s	5.41s	4.11s	4.61s	2.66s	0.36s

for other features, the channel numbers are set to 2. For RNN-based methods, we set the hidden dimensions to 32 based on empirically hyper-parameter tuning.

Evaluation Metrics: We use MAE as the metric, which is the same as the objective loss optimized during training.

7.2 Performance Comparison

Performance Comparison: Table 8 summarizes the results on the testing set, where MAE(P) and MAE(P+T) indicate the MAE loss when taking only the hour price or with Telegram sentiment features as input, Impr. denotes the relative improvement, and Cost is the average training time for 50 batches. It can be observed that: 1) SNN outperforms all the other methods when incorporating Telegram features, suggesting that it maximally exploits the sentiment information; 2) SNN runs significantly faster as discussed in Section 5.2; 3) The improvement brought by sentiment features for TCN is worse than other competitors, indicating that multiple times (5 layers to cover a 200-length sequence) of convolution operations might hurt the expressiveness of pattern of sentiment features.

Sequence Pattern Visualization: Figure 10(b) presents the visualization of positional attention for the price forecasting task, where $F_1 \sim F_6$ correspond to *hour_price*, *neg_num*, *pos_num*, *message_num*, *neg_avg_score*, *pos_avg_score*. We observe that F_1 strictly follows the temporal proximity pattern, and $F_2 \sim F_4$ show similar major patterns with slightly skip-correlated factors. On the other hand, F_5, F_6 demonstrate strong skip-correlation patterns, suggesting that the intensity of investor sentiment has a delayed impact on future price movement. Moreover, we find some of the attention channels of non-skip-correlated feature, e.g., *hour_price*, also demonstrates skip-correlated characteristics as C_3, C_4, C_5 presented in Figure 10(c). After manual review, it is find that some attention channels assign higher scores at the 24-th, 48-th positions, suggesting that through this way SNN captures the periodicity, a typical skip-correlation.

8 RELATED WORK

Anatomy: Kamps et al. [16] are the first to introduce cryptocurrency P&Ds and define their life cycle as accumulation, pump, and dump, and distinguish it from traditional P&Ds in the stock market. Following their setting, Xu et al. [29] present a comprehensive analysis of 412 P&Ds that occurred during Jun. 2018 and Feb. 2019, focusing on describing how a P&D activity is coordinated on Telegram. Li et al. [19] study P&Ds from a statistical perspective and point out the detrimental influence of P&Ds on the development of the cryptocurrency market. Some characteristics of P&Ds are also revealed and confirmed by the subsequent studies [9, 11, 18]. In addition, some works study P&Ds from the perspective of social media. Nizzoli et al. [24] highlight the vast presence of Twitter bots and their key part in spreading the invitation links of pump channels in Telegram. Similar results revealed by [26] show the widespread misinformation spread by Telegram channels to mainstream social media platforms.

P&D Post-detection: This task aims to detect whether a P&D has occurred or is ongoing, based on the statistics like price, trading volume, or signals from external sources, like social media. Kamps et al. [16] perform anomaly detection on the moving averages of price and trading volume data. La et al. [17, 18] improve this method by introducing much more fine-grained features, like *market buy orders*. In addition to using market statistical data, Mirtaheiri et al. [22] take into account the statistics of tweets and users that mentioned cryptocurrencies to make the detection. Even though some of the approaches [17, 18] argue they can achieve real-time response, the post-detection paradigm fails to meet the practical needs since the coin price can easily reach its peak in less than one minute after a P&D starts.

Target Coin Prediction: Xu et al. [29] are the first to introduce this task and they use a RF model to predict the target coin while treating P&D activities as isolated, unrelated events. Nghiem et al. [23] focus on modeling the time-series market movement of target coin by BiLSTM and CNN algorithms, which differs from our fundamental setting.

9 CONCLUSION

In this paper, we present a novel data science pipeline to tackle the target coin prediction task. We identify 709 P&Ds from millions of Telegram messages and extracts valuable and timely information from heterogeneous data sources to support efficient and effective target coin prediction. In our analysis, we discover some interesting phenomena, such as pumped coins exhibiting intra-channel homogeneity and inter-channel heterogeneity, which motivates us to propose SNN to encode a channel's pump history to accurately predict the target coin. Consequently, our method significantly improves over the base model, suggesting that the proposed data science pipeline is well-suited for real-world applications.

10 ACKNOWLEDGEMENT

The authors would like to thank Lingling Lu, Jianting He and Yuan Chen from Zhejiang University for their valuable suggestions regarding this work. This research was supported by the National Research Foundation, Singapore under its Industry Alignment Fund – Pre-positioning (IAF-PP) Funding Initiative. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore.

REFERENCES

- [1] Binance kline api. <https://data.binance.vision/?prefix=data/spot/monthly/klines>.
- [2] Coingecko api. <https://www.coingecko.com/en/api>.
- [3] Pumpolym. <https://pumpolym.com>.

- [4] Telethon api. <https://docs.telethon.dev>.
- [5] G. Attanasio, L. Cagliero, P. Garza, and E. Baralis. Quantitative cryptocurrency trading: exploring the use of machine learning techniques. In *Proceedings of the 5th Workshop on Data Science for Macro-modeling with Financial and Economic Datasets*, pages 1–6, 2019.
- [6] S. Bai, J. Z. Kolter, and V. Koltun. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv preprint arXiv:1803.01271*, 2018.
- [7] Bittrex. Bittrex global trading rules. <https://bittrexglobal.zendesk.com/hc/en-us/articles/360009687839-Bittrex-Global-trading-rules>.
- [8] CoinDesk. Hacked crypto exchange cryptopia files for us bankruptcy protection, 2019. <https://www.coindesk.com/markets/2019/05/27/hacked-crypto-exchange-cryptopia-files-for-us-bankruptcy-protection>.
- [9] A. Dhawan and T. J. Putniņš. A new wolf in town? pump-and-dump manipulation in cryptocurrency markets. *Pump-and-Dump Manipulation in Cryptocurrency Markets (November 17, 2020)*, 2020.
- [10] S. Gupta, J. Hellings, S. Rahnema, and M. Sadoghi. Building high throughput permissioned blockchain fabrics: challenges and opportunities. *Proceedings of the VLDB Endowment*, 13(12):3441–3444, 2020.
- [11] J. Hamrick, F. Rouhi, A. Mukherjee, A. Feder, N. Gandal, T. Moore, and M. Vasek. The economics of cryptocurrency pump and dump schemes. *Available at SSRN 3310307*, 2018.
- [12] S. Hu, Y. Cao, Y. Gong, Z. Li, Y. Yang, Q. Liu, W. Ou, and S. Ji. Gift: Graph-guided feature transfer for cold-start video click-through rate prediction. *arXiv preprint arXiv:2202.11525*, 2022.
- [13] S. Hu, X. Zhang, J. Zhou, S. Ji, et al. Turbo: Fraud detection in deposit-free leasing service via real-time behavior network mining. In *ICDE*, 2021.
- [14] S. Hu, Z. Zhang, B. Luo, B. H. Shengliang Lu, and L. Liu. Bert4eth: a pre-trained transformer for ethereum fraud detection. *the ACM Web conference (WWW)*, 2023.
- [15] C. Hutto and E. Gilbert. Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Proceedings of the international AAAI conference on web and social media*, volume 8, pages 216–225, 2014.
- [16] J. Kamps and B. Kleinberg. To the moon: defining and detecting cryptocurrency pump-and-dumps. *Crime Science*, 7(1):1–18, 2018.
- [17] M. La Morgia, A. Mei, F. Sassi, and J. Stefa. Pump and dumps in the bitcoin era: Real time detection of cryptocurrency market manipulations. In *2020 29th International Conference on Computer Communications and Networks (ICCCN)*, pages 1–9. IEEE, 2020.
- [18] M. La Morgia, A. Mei, F. Sassi, and J. Stefa. The doge of wall street: Analysis and detection of pump and dump cryptocurrency manipulations. *arXiv preprint arXiv:2105.00733*, 2021.
- [19] T. Li, D. Shin, and B. Wang. Cryptocurrency pump-and-dump schemes. *Available at SSRN 3267041*, 2021.
- [20] T. Mikolov, K. Chen, G. Corrado, and J. Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013.
- [21] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean. Distributed representations of words and phrases and their compositionality. *Advances in neural information processing systems*, 26, 2013.
- [22] M. Mirtaheri, S. Abu-El-Hajia, F. Morstatter, G. Ver Steeg, and A. Galstyan. Identifying and analyzing cryptocurrency manipulations in social media. *IEEE Transactions on Computational Social Systems*, 8(3):607–617, 2021.
- [23] H. Nghiem, G. Muric, F. Morstatter, and E. Ferrara. Detecting cryptocurrency pump-and-dump frauds using market and social signals. *Expert Systems with Applications*, 182:115284, 2021.
- [24] L. Nizzoli, S. Tardelli, M. Avvenuti, S. Cresci, M. Tesconi, and E. Ferrara. Charting the landscape of online cryptocurrency manipulation. *IEEE Access*, 8:113230–113245, 2020.
- [25] A. v. d. Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu. Wavenet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499*, 2016.
- [26] D. Pacheco, P.-M. Hui, C. Torres-Lugo, B. T. Truong, A. Flammini, and F. Menczer. Uncovering coordinated networks on social media: methods and case studies. In *Proceedings of the AAAI international conference on web and social media (ICWSM)*, 2021.
- [27] T. Renault. Pump-and-dump or news? stock market manipulation on social media. 2014.
- [28] P. Ruan, T. T. Anh Dinh, Q. Lin, M. Zhang, G. Chen, and B. Chin Ooi. Revealing every story of data in blockchain systems. *ACM Sigmod Record*, 49(1):70–77, 2020.
- [29] J. Xu and B. Livshits. The anatomy of a cryptocurrency {Pump-and-Dump} scheme. In *28th USENIX Security Symposium (USENIX Security 19)*, pages 1609–1625, 2019.
- [30] C. Ye, Y. Li, B. He, Z. Li, and J. Sun. Gpu-accelerated graph label propagation for real-time fraud detection. In *SIGMOD*, pages 2348–2356, 2021.

Received April 2022; revised July 2022; accepted August 2022