

Enhancing Cryptocurrency Price Prediction by Applying Temporal Fusion Transformers on Subseries

Arash Peik[†], Mohammad Ali Zare Chahooki^{†,*}, Amin Milani Fard[‡], Mehdi Agha Sarram[†]

[†]Yazd University, Yazd, Iran

[‡]New York Institute of Technology, Vancouver, Canada

Abstract

In this work, we categorize a financial time series into a number of subseries with similar behavior to increase prediction accuracy by learning the subseries category. We create a deep learning model for each category based on the attention mechanism to predict its next step. Due to the limited amount of cryptocurrency data for training models, if the number of categories increases, the amount of training data for each model will decrease, and some complex models will not be trained well due to the large number of parameters. To overcome this challenge, we propose to combine the time series data of other cryptocurrencies to increase the amount of data for each category, thus increasing the precision of the models corresponding to each category.

Keywords: Cryptocurrency trading, Financial time series, Temporal Fusion Transformers, Categorization of time series

1. Introduction

A financial time series is generally a set of discrete sampled values for the price of a commodity or currency. Predicting financial time series, in particular the price of cryptocurrencies, is a critical task in financial markets, where machine learning techniques have been shown to be effective. Digital currencies fluctuations depend on factors such as volume, the difficulty of mining, popularity, alternative coin prices, transaction costs, social media sentiments, and countries' laws and regulations, and hence cryptocurrencies price forecasting is a challenging task. Akyildirim et al. [1] examined the performance of different machine learning algorithms for predicting the mid-price movement for Bitcoin. Chen et al. [2] used classic machine learning to predict Ethereum price. Fazlollahi and Ebrahimijam [3] presented a multi-input LSTM neural network with technical analysis indicator lags as input to forecast the price returns of Bitcoin, Ethereum, and Ripple. Bouteska et al. [4] provided a comparative analysis of ensemble learning and deep learning models for cryptocurrencies price prediction. The Temporal Fusion Transformer (TFT) [5] is a recent attention-based deep neural network architecture for forecasting time series as an extension of the traditional transformer architecture [6]. TFT incorporates temporal fusion mechanisms to capture temporal patterns and dependencies within sequential data. This model has shown state-of-the-art performance in various time series forecasting benchmarks and can effectively capture the intricate temporal patterns in financial markets [7].

Given the dynamic nature of financial data influenced by external factors, the TFT's adaptability and capability to model complex dependencies make it a promising tool for accurate and robust financial time series predictions. Considering the functionality of the TFT model and its accuracy in repeatable processes and events over time, instead of directly applying it on a given time series, in this work we investigate the effect of splitting the input time series into similar subseries (categories) and feeding these subseries as separate time series to TFT so that the model learns the repeated behavior of each category. Our

*chahooki@yazd.ac.ir




Figure 1. Overview of our proposed method.

hypothesis is that a more accurate prediction of subseries could result in higher accuracy of the series prediction.

Contributions. In this work we propose to categorize the financial time series into several subseries with similar behavior and train a TFT model for each category of the subseries separately. We train these models to predict subseries of other cryptocurrencies that behave similar to the corresponding category. By using different selectors for each cryptocurrency, which recognize the subseries category, it would be possible to predict each cryptocurrency’s bullish or bearish status. By training repeated events in the time series, we increased the prediction accuracy of the model compared to deep models such as LSTM and simple TFT, and increased the profitability of algorithmic trading. Our code and tests are publicly available for download¹.

2. The Proposed Approach

Our main idea is to increase the prediction accuracy by categorizing the financial time series into several similar subseries and more accurate learning of classes of similar subseries. We use a Markov chain to predict the next subseries category after finding the category of the last observed subseries. Finally, with the help of a transformer-based model related to that category, we predict the value of the cryptocurrency at the next moment.

Figure 1 depicts the overview of our proposed approach. We extract the Open, High, Low, Close, and Volume (OHLCV) information from a target cryptocurrency time series (block 1). Then, because a cryptocurrency may fluctuate in completely different price ranges over time, and to use the entire time series in model training, we change the price time series to the volatility rate time series (block 2). For example, on October 13, 2023, the price of Bitcoin was around \$26,750. Several days before and after, the price has increased and decreased in the same range, and our goal is to predict this increase and decrease in prices. Five months later, the price of Bitcoin fluctuated around \$71,500. If we use the price time series, the decrease in the price of \$100 will be considered the same in the entire time series if

¹<https://github.com/arashitc2/tft-on-subseries-for-cryptocurrency-forecasting>

the decrease in the price of \$100 in October 2023 is much more significant than this decrease in March 2024. By changing the time series of the price to the time series of the volatility, in short time frames of a few minutes, we have somehow normalized the financial values in the range of -1 to 1 (percent). Because the rise or fall of the cryptocurrency price and its amount are important to us, this change makes the model more focused on this issue.

In the next step, we extract subseries with the same length from the created time series (block 3). These subseries are created by moving a sliding window on the time series. We will do all the steps of block 1 to 3 for the time series of other cryptocurrencies as well. We then combined similar series of other cryptocurrencies with the data (block 4) to increase the training data for better learning of the models. Next, we categorized them based on the upward or downward volatility of the price of each time frame compared to the previous one (block 5). We explain this categorization in more detail in section 2.2.

We then create a TFT model for each category of subseries (block 6). From the subseries of target cryptocurrency extracted from block 3 and their corresponding categories, we create a probabilistic model to predict the next category after observing a category of subseries (block 7). This model predicts upward/downward volatility of the price of the next time frame compared to the current time frame. With this prediction, we will not know the exact amount of upward/downward volatility, which will be predicted in the next step. We note that the volatility may be decreasing, but the price is increasing, or the volatility may be increasing, but the price is decreasing, which we will explain further. When a new subseries is observed until time t , we intend to predict the price and status in time $t+1$. After detecting the observed subseries category (block 8), we send it to a model selector to predict the next subseries category with probabilities (block 9). Finally, we send the subseries to the TFT model corresponding to the predicted category to predict for $t+1$ (block 10).

Algorithm 1 shows the steps of our proposed method. It gets as input the time series of all available cryptocurrencies for training the TFT models. The preprocessing steps (lines 1-5) include conversion of price time series to price volatility time series (line 2 corresponding to block 2 in Figure 1), converting the time series into subseries based on the movement of the sliding window to the length of the number of time frames n (line 3 according to block 3). We considered n to be 8 so that the trained model can predict the value of the 8th time frame after seeing a time series including 7 time frames. Next, all subseries of cryptocurrencies are merged together (line 4 corresponding to block 4). Lines 6-9 perform subseries categorization (block 5) and build *allCatSubseries* that is a key-value array, where the key is *catID* and the value is a list of subsets key is the *catID* and value is a list of subseries. TFT models will be trained in lines 10-12, and finally, in lines 13-17 according to the desired cryptocurrency for prediction, the selective model specific to the same cryptocurrency will be produced according to the categories after performing the aforementioned preprocessing steps.

2.1. Data Preprocessing

We collected a dataset of Binance transactions in 1-minute time frames for 2 years, which has not been published before at this granularity. The details of the dataset are described in section 3.1. The important features used are "volume", "closing price", "number of trades", "quote Asset Volume". Among the price features, we will change the price of the close moment to the feature of the percentage change of the price compared to the previous moment as Volatility rate. For example, consider a subseries of Litecoin price based on Tether as [83.97 , 83.92 , 83.79 , 83.67 , 83.67 , 83.64 , 83.61 , 83.61]. We change this subseries to Volatility rate subseries using $V_t = (\frac{P_t}{P_{t-1}} - 1) * 100$, that modifies time series to [-0.05 , -0.15 , -0.14 , 0 , -0.03 , -0.03 , 0]. Note that since P_{t-1} is used to produce the subseries of changes and the first value of the subseries does not have the previous moment, the number of members of the changed subseries is one less, i.e. 7 numbers compared to 8

Algorithm 1 TFT Model and Model Selector Generator.

Input: Number of time frames n , cryptocurrencies time series T , cryptocurrency to predict c
Output: TFT model M for all cryptocurrencies, and TFT model selector S for cryptocurrency c

```

1: for each cryptocurrency time series  $t \in T$  do
2:    $v \leftarrow \text{PriceToVolatility}(t)$ 
3:    $\text{subseries} \leftarrow \text{SplitToSubseries}(v, n)$ 
4:    $\text{allSubseries} \leftarrow \text{allSubseries} \cup \text{subseries}$ 
5: end for
6: for each  $\text{sub} \in \text{allSubseries}$  do
7:    $\text{catID} \leftarrow \text{Categorize}(\text{sub})$ 
8:    $\text{allCatSubseries}[\text{catID}] \leftarrow \text{allCatSubseries}[\text{catID}] \cup \text{sub}$ 
9: end for
10: for each key-value  $\{\text{catID} \Rightarrow \text{catSubseries}\} \in \text{allCatSubseries}$  do
11:    $M[\text{catID}] \leftarrow \text{LearnTFTModel}(\text{catSubseries})$ 
12: end for
13: for each  $\text{cryptoTS} \in T[c]$  do
14:    $v \leftarrow \text{PriceToVolatility}(\text{cryptoTS})$ 
15:    $\text{subseries} \leftarrow \text{SplitToSubseries}(v, n)$ 
16:    $S \leftarrow \text{LearnTFTSelectorModel}(\text{allCatSubseries}, \text{subseries})$ 
17: end for
18: return  $\{M, S\}$ 

```

in this example. This will create a new feature based on the price, which does not depend on the type of cryptocurrency and its price change interval. In order to better train models that need large amounts of data for training, the data of different cryptocurrencies can be used collectively to improve the quality of training.

To normalize some other features, such as the "number of trades", we divide them by the maximum value of that feature during a period when the market is neither strongly bullish nor strongly bearish. If the market is bullish or bearish and the value of these features is greater than the selected maximum, we consider it as 1 in normalization. Considering that we have set the closing price as a prediction criterion, we replace the open, high and low features for each period with their ratio to the closing price of that period. In other words, we perform feature extraction by dividing the values of each attribute by the value of the "close" attribute of the previous record.

In a deep model based on transformers, if the mentioned features and indicators, such as Moving Average Convergence/Divergence (MACD), are used, it can effectively improve the accuracy, which is also mentioned in our results section. However, after checking the correlation of the features, we realized that the current close price and the features extracted from it do not have a significant correlation with other features. Hence, we only used the "close price volatility rate" feature for the sake of simplicity of the models. Moreover, we entered the position of the moment in the time series as covariates in the models.

2.2. Time Series Categorization

Clustering a financial time series does not generally divide the subseries into behavioral similar clusters. Figure 2a shows an example of clustering in 1-dimensional space, which shows the red points of one cluster and the blue points of another cluster. Blue cluster points are different from each other in terms of behavior. If the chart shows the volatility of the price, two points of this cluster show an increase in the price and the other two points show a decrease in the price. Figure 2b also shows the three subseries that are placed in one cluster after running the k -means algorithm to divide into 256 clusters. These three subseries have completely different behaviors compared to each other.

Clustering in the n -dimensional space, where n equals the number of moments of each subseries, results in placing subseries with similar vectors in the same category. Thus, two




Figure 2. Cluster with different behavior.



Figure 3. Time subseries of 7 consecutive moments of the Bitcoin-Tether currency pair

subseries that differ in behavior but have small distance in the n -dimensional space may be put in the same cluster. It is also possible that two similar subseries in behavior and logic may not be placed in the same cluster. On the other hand, the clustering operation requires its own processing time and resources. For these reasons, they can be labeled with little processing load without using clustering and considering the behavior of each subseries. Since in the labeling operation, the categories are already known and do not require special learning, we use the term "categorization" for this.

This Categorization is such that if we divide the initial time series into subseries with 8 consecutive moments and consider the behaviors to be only two states of decreasing or increasing compared to the previous moment, each subseries can be a member of 2^7 different states (categories). Figure 3 (upper part) shows a subseries consisting of 7 consecutive moments of the time series of the price of Bitcoin against Tether. The bottom part of figure 3 shows the price volatility of the top part.

If we consider the fluctuation (increase-decrease rate) of the points of a graph as important for training a model, we should refer to the price acceleration graph (price fluctuation fluctuation) from the price fluctuation graph. In this case, the two subseries of price fluctuations, which show the amount of price fluctuation at each moment compared to the price of the previous moment, if they have a similar shape in terms of appearance, they will have relatively the same acceleration graph. Even if, in terms of appearance, one subseries of fluctuations is above the zero line of the chart, and another subseries is exactly the same, but below the chart, they will have the same slope chart. Figure 4 shows two subseries with the same slope, one completely above and one completely below the zero line. In fact, the two mentioned subseries show completely different price trends. In the first case, when the subseries is completely above the zero line, it indicates the upward trend of the price; in the




Figure 4. Two subseries with similar appearance, same acceleration, but different behavior.




Figure 5. Candlestick chart of two subseries with the same acceleration.

second case, when the subseries is completely below the chart, it indicates the downward trend of the price. These two completely different states will likely show the same behavior in the next moment. To show this issue, you should refer to the candlestick chart of the initial time series. Figure 5 shows the candlestick chart of two subseries of Figure 4, where part A shows a completely bullish price and part B shows a completely bearish price.

One of the basic technical analyses in the candlestick structure is that if the increase in price gradually decreases, we will probably see a downward price trend in the future, and it can be a good point to sell that cryptocurrency. In another analysis, if the downward trend of the moment-to-moment price increases, the price of the next moment will probably decrease as well. Therefore, in both cases, which are completely different, there is a possibility that the next price will be downward. In Figure 4 blue line, if we predict that the speed of price changes is downward, the degree of downward speed is very important. Because if the rate of change of the downward price and a number is higher than 0, the price will be upward. Also, if the rate of price change is downward and the number is smaller than 0, then the price will be bearish. Here, the rate of change is decreasing, which has a great impact on investors' decision, will be predicted by the transformer-based model corresponding to the same category. In general, the model would predict the decrease of the fluctuation, but its amount will be predicted according to the previous moments of the subseries.

2.3. Probabilistic Model Selector

We build a probability-based model with fast self-updating capability to predict the next subseries category and select its corresponding model. The model is made according to the history and possibilities of transfer from one category to another. If, when creating subseries from the main series, we move each time window stage by only one transfer step, and categorization according to the direction (up/down) of the price, in this model each

category can only be transferred to two other categories. This way, there will be a possibility of transfer from each category to only two other categories.

2.4. Temporal Fusion Transformer Construction

The transformer model based on the attention mechanism was introduced by Vaswani et al. [6]. Attention models attempt to address the potential limitations of the encoder-decoder model and enable the network to focus on a few specific aspects at a time and ignore the rest. These models perform well in predicting time series due to paying attention to previous effective situations. Each time series record can be seen as an entity with three representation vectors: query vector, key vector, and value vector. The attention of an entity on a set of entities is a linear combination of the value vectors of the set, where the weights are determined by how similar the query vector of the entity is to the key vectors of each entity in the set. Cosine multiplication is used to measure the similarity, and then a softmax layer normalizes the similarity scores between zero and one.

Recently, Lim et al. [5] introduced a method called Temporal Fusion Transformer (TFT), that in addition to addressing the complexities in recurrent neural networks, takes covariates of each time frame, typically time-dependent, as inputs. The model incorporates attention to other previous time steps using the transformer approach, as proposed by Vaswani et al. [6]. Another feature in almost all TFT stages and gates is the ability to jump and direct connection, which makes it possible to skip stages and gates in situations where our data is simpler than the model. In this way, the model can adjust its complexity during training. Also, this model can receive auxiliary time-dependent variables in addition to the main time series and adjust the time in the time series according to the auxiliary variables.

If our data is less complex than the model, with the help of this model, it is possible to skip steps and gates. In this way, the model can adjust its complexity during training. The specialty of this model lies in its ability to observe a time series and predict its future continuation of the series. This model can receive time-dependent covariates in addition to the main time series and adjust the time in the time series according to the covariates. This feature of the model helps us to easily and without creating additional noise, connect the subseries of the same category that are separate and related to different times and give it as training data to the model.

In order to send the subseries of each category to the TFT model, we first join all the subseries of each category and generate a new time series. Then, we create a matrix with the length of the new time series, the value of each of which is equal to the index of the corresponding moment in the initial subseries. In this way, if each time subseries contains 7 moments, the constructed matrix contains the numbers 1, 2, 3, ..., 7 that are repeated alternately. This matrix will be sent as time-dependent covariates data to the TFT model. The mentioned matrix is one-dimensional, which can be converted into two-dimensional, and other time parameters, such as hour, day of the year, day of the week, etc., from the primary subseries can be placed in the other dimension of the matrix.

After creating the time series of each category and the corresponding covariates data matrix, we send them to the TFT models for training to have a TFT model corresponding to each category. After training the models, to predict the new data, we must first find the category of the last subseries of the new data. We then use a probabilistic selector model to predict the next subseries category. Therefore, to predict the price of the next moment, it is enough to send the next subseries to the TFT model of the predicted category.

3. Experiments and Results

3.1. Dataset

A large amount of cryptocurrency transfers happen through decentralized exchanges, which were used in this research. All transactions are entered in currency pairs, i.e., the price of the first cryptocurrency compared to the second. We store transaction information of Binance exchange as the largest decentralized cryptocurrency exchange cumulatively in 1-minute time intervals for 18 months since December 2021. Therefore, each database record represents the transactions of a currency pair in a specific minute. Due to the unavailability of one-minute transaction details for an extended period in the existing datasets and APIs, we had to gather the necessary information. All buying and selling orders are divided into two categories: taker and maker. If a cryptocurrency buyer buys instantaneously at the price listed by the sellers, it is called a *taker*. The Taker buyer chooses from the sale list and makes his purchase instantly. However, if the buyer wants to buy at a price lower than the selling list, he will register his order in the buying queue so that whenever a seller agrees to that price, the transaction will be done. In this case, the buyer is called a *maker*. The Maker buyer does not make his purchase instantly and registers his order in the shopping queue. In the same way, the sellers are also divided into two categories: Takers and Makers.

The information stored in each record includes the time of the start and end of that minute in the form of a timestamp, the title of the cryptocurrency pair, the value of the first cryptocurrency compared to the second cryptocurrency at the exact start time (the first transaction of that minute), the value of the first cryptocurrency compared to the second one at the end of the minute (the last transaction of that minute), the maximum and minimum value of the first cryptocurrency compared to the second one in that minute, the trading volume of that minute based on the first cryptocurrency (V), the trading volume based on the second cryptocurrency (QAV), the number of transactions, the volume purchased by the taker buyer, the payment amount of the second cryptocurrency to buy the first cryptocurrency by the taker buyer, and the name of the exchange. We got this information from Binance decentralized exchange, which is known as the largest and most effective decentralized cryptocurrency exchange. We examined 13 currency pairs for this research and collected in our database. The first cryptocurrencies of these currency pairs include Bitcoin, Ethereum, Bitcoin Cash, Litecoin, Dogecoin, PolkaDOT, and Ripple.

3.2. Implementation and Settings

We used Python *Darts* library to implement the main predictive models of TFT. We also implemented and tested the TFT Model Selector in different ways including Markov chain and LSTM model using *Keras* library.

Time frame. We initially conducted our experiment on 1-minute time frames; however, since the time frame is very small and at any moment the minimum price of the selling queue is different from the maximum price of the buying queue and the last transaction of the exchange may be done from one of these two queues, therefore, small fluctuations are created at any moment. Transaction information of these moments and fluctuations alone are not meaningful, and by aggregating a volume of these moments, we will get meaningful information. After conducting experiments, we empirically concluded that the minimum meaningful small time frame is 7 minutes.

TFT models. Temporal Fusion Transformers are powerful models for predicting multivariate time series [5]. In our proposed method, we use a TFT model to predict each category of subseries to predict the next moment price change of those subseries with higher accuracy. The TFT gets as input static metadata, time-varying past inputs, and time-varying a priori known future inputs. Given time subseries in 8 time frames, we want to predict the price

or direction of the closing moment of the 8th time frame by observing 7 time frames of 7 minutes. Therefore, in the created TFT models, we set the value of input-chunk-length to 7 and the value of output-chunk-length to 1. We also set the Hidden state size of the TFT to 70, used 4 LSTM layers and 4 attention heads, and trained the model with only 7 epochs. According to the number of 8 time frames in the existing time subseries, if we consider the upward or downward conditions of the close price of each time frame compared to the previous time frame, we will have 2^7 different subseries. We created a TFT model for each state of the subseries. We further created a Markov chain-type model that has the status of the number of TFT models made (128 models). From each situation, there is a possibility of moving to only two other situations (upward or downward price at the close of the next time frame). This model is responsible for selecting the appropriate TFT model according to the probabilities to predict the 8th price direction. For testing purposes, instead of 128 TFT models produced and the Markov selector model, TFT models were created that did not use the last minute time frame in model making. These models predict the last moment situation (upward or downward) without using a selector. Therefore, in this experiment, 2^6 classes of subseries were created, equivalent to 64 different TFT models.

The TFT model can learn the pattern well by receiving covariates that show a pattern in the time series. Since the TFT model receives a continuous time series for training, we concatenate them to send the subseries of a category to the model. In order for the model to understand and focus on the repeatability of the subseries behavior, we assigned a sequence number to each moment of each subseries and sent them to the model as covariates. In these models, we used 4 layers for the LSTM Encoder and Decoder. Also, we used 4 attention heads in making the model.

TFT model selector. The selector decides which TFT model corresponding to the category can better predict the price volatility rate. We evaluate 3 different solutions for this selector. First, we create a Markov chain to predict the next category that will be observed according to the probabilities (HMM). Second, we use LSTM model to predict the next category. After predicting the next category, we will send the last sub-time series to the corresponding model of that category to predict the price situation of the next moment. Third, we delegate the prediction of the next category to the same TFT model. In this way, from the same stage of categorization, we do not enter the moment of the tag in the categorization and let our TFT model recognize that the price volatility rate of the final moment has decreased or increased compared to the previous moment. In this case, the number of moments based on which the categorization is done will be reduced by one moment, and as a result, the number of categories will be halved. Note that in the first and the second method, i.e., HMM and LSTM, we predict the the price volatility rate increase/decrease compared to the previous moment with the selector model. However, we use the TFT model to predict whether this increase/decrease will cause the volatility rate to become positive or negative, i.e., increase/decrease in the price of the cryptocurrency.

3.3. Comparisons and Results

While our methodology leverages data from multiple cryptocurrencies during the training phase to learn generalized subseries behaviors, for our evaluation in this early work we focus specifically on the LTC-USDT pair as a proof-of-concept to demonstrate the effectiveness of the proposed approach within a manageable experimental scope.

Performance. We assess our model’s prediction accuracy and precision compared to the baselines. For testing, we used LTC-USDT transaction data that were not present in the model training from 08-22-2023 for two weeks. We checked the criteria only for bullish or bearish prices. In trading, if a bearish time frame is wrongly recognized as bullish due to the purchase of cryptocurrency and entry into the trade, in addition to not achieving the

estimated profit, the previous asset will also be damaged. For this reason, the precision of the methods will be calculated and evaluated. On the other hand, if the bullish time frame is wrongly recognized as bearish, the previous asset will not suffer a loss. Table 1 shows the accuracy and precision of our proposed method compared to recent basic techniques.

In this comparison, we did not include financial covariates in any of the TFT models. The results do not show high values for accuracy and precision and the improvement does not seem to be considerable. However, note that in this work in order to avoid very large amount of transactions, we confined our time frame to 7 minutes which is very small compared to hourly or daily time frames. Cryptocurrency markets are chaotic in small time frames and due to the large number of transactions any small improvement in accuracy and precision values can have a great impact on long-term profits, e.g. on a monthly basis. In general, we are focused on short time frames and quick profitable trades. This strategy is designed to take advantage of short-term market movements by analyzing price action on a short time frames, and is known for its agility and ability to capture quick profits in volatile markets.

Table 1. Accuracy and Precision Comparison.

Model	Accuracy	Precision
Trading using LSTM model [8]	0.496	0.492
Trading using a TFT model [5] without financial covariates	0.535	0.529
Trading with subseries categorization (without selector model and 64 categories) ¹	0.539	0.534

Selector accuracy. In two of our experiments, we used a separate selector model including Markov chain and LSTM, and in another experiment, we combined the selector with the original TFT models. For LTC-USDT currency pair, the accuracy of TFT model selector was 0.79 for Markov chain selector and 0.74 for LSTM selector.

Profitability. We calculate the profitability through a simulation. We use the dataset from the previous section for testing. In this experiment, we simulate a space in which we have 100 USDT at the beginning of the asset period, and with the help of the proposed methods, based on the models’ output buy and sell signals, we trade at that moment’s price. In these transactions, we considered the 0 exchange fee. We entered the entire asset into the transaction with each signal, regardless of the transaction’s risk level and the stock portfolio’s management. The end-of-period asset represents a measure of the profitability of the proposed method. Table 2 shows the profitability of the proposed methods in this simulation. Note that a 1 USDT increase in profit in 2 weeks based on the initial balance of 100 USDT is equivalent to an increase of 2% in monthly profit compared to other methods, and this profit is achieved in a bear market without using futures and leverages. In our analysis time frame, without trading the 100 USDT value decreases to 98 and using LSTM model for training it gets even less to 97.55. Using other baselines will almost retain 100 USDT as shown in Table 2, except for trading using a TFT model using financial covariates that is comparable to ours, while in our experiments we did not consider financial covariates and left it for a future work. The 4 USDT increase means 4% profit in 2 weeks that is about 8% increase per month that is a considerable profit.

4. Discussion

Profitability and accuracy. We observed in our experiments with the dataset for testing that the price of the selected cryptocurrency is generally down 2% in the selected 14 days. Using our proposed model we achieved more than 4% profit, which compared to the profitability of holding cryptocurrency in this period, our temporal fusion transformer models with time series categorization have created more than 6% more profit in 2 weeks. In the

Table 2. Profitability Comparison.

Model	End of term asset amount
Without trading (buying at the beginning and selling at the end of the period)	98.01 USDT
Trading using LSTM model	97.55 USDT
Forecasting and trading using a TFT model without financial covariates ¹	100.81 USDT
Forecasting and trading using a TFT model using financial covariates	103.19 USDT
Trading by categorization of subseries and using LSTM selector model	100.35 USDT
Trading by categorization of subseries and using Markov chain based selector	103.16 USDT
Trading with subseries categorization (without selector model and 64 categories)	104.08 USDT

same conditions, the proposed model performs better than LSTM models and transformers (without classification), while only using price data and not any other covariates.

We also evaluated different selectors by using TFT models corresponding to each category. First, in order to evaluate the accuracy of the category models, we assumed a selector that selects the category (prediction model) of each subseries with 100% accuracy. In this case, the category models turn 100 USDT into 248 USDT in two weeks, and actually make a 148% profit in two weeks. Then we created a selector with LSTM, but contrary to expectations it reduced the amount of assets and after two weeks \$2.50 of the initial assets were gone. Another selector that was tested was based on probabilities and Markov chain. This model had more than 3% profitability in the mentioned two weeks. The important feature of this selector is that it can be used in predicting other cryptocurrencies without changing the transformer-based models and only by changing the probability values of the Markov chain. We also tested a method without using a selector so that our trained TFT models also predict the bullishness or the bearishness of the price velocity. When categorizing the subseries, we do not use the price fluctuation of the last time frame in the categorization, and the model should predict the entire situation of the last moment. In this method, due to the reduction of a time frame in categorization, we have 64 categories, and the training data of each category increases. We obtained more than 4% profit within two weeks of testing using this approach. Here the disadvantage is that the general models may show different qualities and accuracy in other cryptocurrencies.

While the observed improvements in accuracy and simulated profitability might appear modest in absolute terms, their significance should be evaluated within the challenging context of high-frequency cryptocurrency market forecasting. These markets are characterized by high volatility and near-chaotic dynamics, particularly on short time scales (7-minute intervals in our study), where achieving consistent predictive accuracy marginally above baseline levels (e.g., 50% accuracy) is notoriously difficult yet potentially impactful. In algorithmic trading operating at such frequencies, even small, consistent edges in predictive power can compound into substantial gains over a large number of trades.

Generalizability. With regards to the generalizability of our proposed approach on other financial markets, such as Foreign Exchange (Forex) or commodities such as gold, or stock market, we can similarly utilize trading data of such markets. Our method does not depend on only one specific cryptocurrency due to the combination of sub-series of different cryptocurrencies in the construction of TFT models, and the models learn the behavior of a category of sub-series. Therefore, the learned model can be used in other financial markets by creating the TFT Model selector for that financial market. Also, according to the previous work done in fractal dimensions [9], with some changes, the data of other time frames can be predicted with this built model and we leave this as a future work. The produced TFT models can be affected by the length of time frames and in financial markets with extreme fluctuations, if the length of the time frame increases, the behaviors may intensify and the TFT models may not work well.

5. Conclusion and Future Work

We presented a method for predicting financial markets without extensive historical data. By using statistical models such as the hidden Markov model for the selector, we can update the model in real-time. This allows us to observe the varying behaviors of time series dynamically and enhance the training. Our results indicate that if the category selector model achieved 100% accuracy, a profit of 148% could be obtained. However, with a selector accuracy of 79%, we achieved a profit of 3.16% over a two-week period. Therefore, even minor improvements in the selector's accuracy can significantly increase profitability. Enhancing the accuracy of the selector model will be a primary focus of our future work. Our experiments show that the profitability of the model built using our proposed method, without utilizing any financial covariates, is higher than that of a simple TFT model that does use covariates. We will consider incorporating covariates into the construction of TFT models or the selector model as a future research. This work contributes to the state-of-the-art by offering an adaptive method for financial market prediction. The importance of our work lies in its potential application across various financial markets, providing a flexible and continually improving predictive model. Another potential direction for future improvement is to combine the subseries categorization approach with sentiment analysis derived from external sources such as news and social media, similar to the strategy used in the MFB model [10] to incorporate time-lagged sentiment.

Acknowledgments

This work was supported by the New York Institute of Technology - Vancouver Campus.

References

- [1] E. Akyildirim, O. Cepni, S. Corbet, and G. S. Uddin. "Forecasting mid-price movement of Bitcoin futures using machine learning". In: *Annals of Operations Research* 330.1 (2023), pp. 553–584.
- [2] M. Chen, N. Narwal, and M. Schultz. "Predicting price changes in ethereum". In: *International Journal on Computer Science and Engineering (IJCSSE) ISSN* (2019), pp. 0975–3397.
- [3] N. Fazlollahi and S. Ebrahimijam. "Predicting Cryptocurrency Price Returns by Using Deep Learning Model of Technical Analysis Indicators". In: *Global Economic Challenges: 6th International Conference on Banking and Finance Perspectives*. Springer. 2023, pp. 175–186.
- [4] A. Bouteska, M. Z. Abedin, P. Hajek, and K. Yuan. "Cryptocurrency price forecasting—A comparative analysis of ensemble learning and deep learning methods". In: *International Review of Financial Analysis* 92 (2024), p. 103055.
- [5] B. Lim, S. Ö. Arık, N. Loeff, and T. Pfister. "Temporal fusion transformers for interpretable multi-horizon time series forecasting". In: *International Journal of Forecasting* 37.4 (2021), pp. 1748–1764.
- [6] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. "Attention is all you need". In: *Advances in neural information processing systems* 30 (2017).
- [7] X. Hu. "Stock price prediction based on temporal fusion transformer". In: *International Conference on Machine Learning, Big Data and Business Intelligence*. IEEE. 2021, pp. 60–66.
- [8] X. Song, Y. Liu, L. Xue, J. Wang, J. Zhang, J. Wang, L. Jiang, and Z. Cheng. "Time-series well performance prediction based on Long Short-Term Memory (LSTM) neural network model". In: *Journal of Petroleum Science and Engineering* 186 (2020), p. 106682.
- [9] S. Bhatt, H. Dedania, and V. R. Shah. "Fractal dimensional analysis in financial time series". In: *International Journal of Financial Management* 5.2 (2015), pp. 57–62.
- [10] P. Han, H. Chen, A. Rasool, Q. Jiang, and M. Yang. "MFB: A Generalized Multimodal Fusion Approach for Bitcoin Price Prediction Using Time-Lagged Sentiment and Indicator Features". In: *Expert Systems with Applications* 261 (2025), p. 125515. issn: 0957-4174.