

การพัฒนากรอบการพยากรณ์ด้วย LSTM สำหรับการทำนายทิศทางราคารายวัน
ของ EUR/USD ในตลาดแลกเปลี่ยนเงินตราต่างประเทศ
Developing an LSTM-Based Forecasting Framework for Predicting Daily
EUR/USD Price Direction in the Foreign Exchange Market

วรรณะ พงษ์เสนา^{1*}

Watthana Pongsena^{1*}

¹คณะศิลปศาสตร์และวิทยาศาสตร์ มหาวิทยาลัยราชภัฏศรีสะเกษ จังหวัดศรีสะเกษ

¹Faculty of Liberal Arts and Sciences, Sisaket Rajabhat University, Sisaket Province

*Corresponding Author: w.pongsena@sskru.ac.th

Received: February 3, 2026

Revised: March 10, 2026

Accepted: March 16, 2026

บทคัดย่อ

ตลาดแลกเปลี่ยนเงินตราต่างประเทศ (Foreign Exchange: Forex) เป็นหนึ่งในตลาดการเงินที่มีขนาดใหญ่และมีสภาพคล่องสูงที่สุดในโลก การพยากรณ์ทิศทางราคาส่งออกของอัตราแลกเปลี่ยนยังคงเป็นเรื่องที่ท้าทาย เนื่องจากข้อมูลอนุกรมเวลาทางการเงินมีลักษณะผันผวนสูง ไม่เป็นเชิงเส้น และไม่คงที่ตามเวลา งานวิจัยนี้พัฒนากรอบการพยากรณ์โดยใช้แบบจำลอง Long Short-Term Memory (LSTM) สำหรับการทำนายทิศทางราคารายวันของคู่เงิน EUR/USD โดยใช้ข้อมูลอัตราแลกเปลี่ยนรายวันที่ส่งออกจากแพลตฟอร์ม MetaTrader 4 ช่วง วันที่ 2 มกราคม ค.ศ. 2018 ถึง 30 เมษายน ค.ศ. 2019 จำนวน 345 ระเบียบ ข้อมูลถูกกำหนดให้อยู่ในรูปแบบปัญหาการจำแนกสองคลาสผ่านตัวแปร Trend และสร้างข้อมูลลำดับด้วยเทคนิค sliding window สำหรับการฝึกแบบจำลอง กรอบการวิจัยนี้ศึกษาผลของพารามิเตอร์สำคัญ ได้แก่ ความยาวหน้าต่างข้อมูลนำเข้า ปริมาณข้อมูลฝึก และความถี่ในการปรับปรุงแบบจำลอง นอกจากแบบจำลอง LSTM พื้นฐานแล้วยังพัฒนาแบบจำลองปรับเหมาะและแนวทางการแบ่งกลุ่มลำดับข้อมูลตามลักษณะการเคลื่อนไหวของราคา (LSTM-CS) เพื่อให้แบบจำลองสามารถเรียนรู้รูปแบบตลาดที่แตกต่างกันได้ ผลการทดลองพบว่าแบบจำลองพื้นฐานให้ความแม่นยำร้อยละ 52.5 และเพิ่มขึ้นเป็นร้อยละ 67.5 หลังการปรับแต่งพารามิเตอร์ ขณะที่แบบจำลอง LSTM-CS ให้ผลดีที่สุดที่ร้อยละ 69.7

คำสำคัญ: หน่วยความจำระยะสั้นและระยะยาว, การพยากรณ์อนุกรมเวลา, ตลาดแลกเปลี่ยนเงินตราต่างประเทศ, การพยากรณ์ทิศทางราคา

Abstract

The foreign exchange (Forex) market is one of the largest and most liquid financial markets in the world. Predicting exchange-rate direction remains challenging because financial time-series data are highly volatile, nonlinear, and non-stationary. This study develops an LSTM-based forecasting framework for predicting the daily price direction of the EUR/USD exchange rate. The dataset consists of daily exchange-rate data exported from the MetaTrader 4 platform from 2 January 2018 to 30 April 2019, comprising 345 observations. The prediction task is formulated as a binary classification problem through a Trend variable, and sequential samples are generated using a sliding window technique for model training. The proposed

framework investigates the effects of key parameters, including input window length, training data size, and model update frequency. In addition to the baseline LSTM model, optimized models and a classification-based sequence segmentation approach based on price movement characteristics (LSTM-CS) are developed to enable the model to learn different market patterns. The experimental results show that the baseline model achieved an accuracy of 52.5%, which increased to 67.5% after parameter optimization, while the LSTM-CS model produced the best result at 69.7%.

Keywords: Long Short-Term Memory, Time-Series Forecasting, Foreign Exchange Market, Price direction prediction

1. Introduction

The foreign exchange (Forex) market is the largest and most liquid financial market in the world and plays a critical role in global trade, international investment, and monetary systems (Lyons, 2001). According to the Bank for International Settlements (BIS, 2022), the global Forex market records a multi-trillion-dollar daily trading volume, making it one of the most actively traded financial markets worldwide. Despite its high liquidity and continuous trading activity, predicting exchange-rate movements remains challenging because financial time-series data are complex, nonlinear, and non-stationary (Poon & Granger, 2003; Lim & Zohren, 2021). These characteristics make it difficult for traditional forecasting methods to capture the dynamic relationships that govern exchange-rate fluctuations.

Historically, exchange-rate forecasting has relied on two main analytical frameworks: fundamental analysis and technical analysis. Fundamental analysis evaluates macroeconomic variables such as interest rates, inflation, and monetary policy to estimate currency valuation (Rapach & Zhou, 2013). In contrast, technical analysis focuses on historical price patterns and trading indicators to identify potential future movements (Lo et al., 2000; Atsalakis & Valavanis, 2009). Although these approaches have been widely applied in financial markets, their predictive performance may be limited when market dynamics exhibit complex nonlinear behavior and rapidly changing statistical properties.

With the increasing availability of financial data and computational power, machine learning techniques have been widely explored as alternative approaches to financial time-series forecasting. Early studies demonstrated that artificial neural networks (ANNs) and support vector machines (SVMs) could improve forecasting performance under certain market conditions (Cao et al., 2003; Kamruzzaman & Sarker, 2003). These models can capture nonlinear relationships in financial data that are often difficult to model using conventional statistical techniques.

More recently, deep learning architectures have gained considerable attention in financial forecasting applications. In particular, recurrent neural networks (RNNs) are well suited to sequential data modeling because they incorporate temporal dependencies between observations. However, conventional RNN models often suffer from vanishing and exploding gradient problems when learning long-term dependencies (Hochreiter, 1998; Pascanu et al., 2013). To overcome this limitation, Hochreiter &

Schmidhuber (1997) introduced the Long Short-Term Memory (LSTM) architecture, which incorporates memory cells and gating mechanisms that regulate the flow of information across time steps. Because of these characteristics, LSTM networks have been widely applied to financial forecasting problems, including stock-price prediction, volatility forecasting, and exchange-rate modeling (Fischer & Krauss, 2018; Sezer et al., 2020; Lim & Zohren, 2021; García et al., 2024).

Recent studies have continued to explore the application of deep learning methods to exchange-rate prediction. For example, García et al. (2024) compared LSTM and bidirectional LSTM models for foreign exchange forecasting, while Nguyen et al. (2024) investigated deep learning-based predictive models for identifying Forex market trends. These studies demonstrate that deep learning approaches can capture complex nonlinear relationships in financial time-series data and may provide better predictive performance than conventional statistical methods. However, the effectiveness of such models often depends not only on network architecture but also on the configuration of important modeling parameters.

Despite the growing body of research on deep learning-based financial forecasting, many existing studies focus primarily on improving model architectures or developing hybrid neural networks. In practical forecasting environments, however, model performance can also be strongly influenced by operational parameters such as input window length, the amount of historical data used for training, and the frequency at which the predictive model is updated. These factors are particularly important in financial markets, where statistical relationships may evolve over time because of changing economic conditions, policy decisions, and market sentiment (Gu et al., 2020; Zhang et al., 2022). Systematically investigating the effects of these parameters can therefore provide valuable insights into the practical implementation of deep learning models for financial forecasting.

Another important challenge in financial time-series prediction is the heterogeneous nature of market movements. Exchange-rate sequences may contain different behavioral patterns such as persistent trends, short-term reversals, and irregular fluctuations. When a single predictive model is trained on all sequence types simultaneously, the presence of heterogeneous patterns may reduce the model's ability to learn consistent relationships between historical observations and future price movements. Segmenting sequences into more homogeneous groups may therefore improve the learning capability of forecasting models by allowing specialized models to focus on specific market behaviors.

To address these issues, this study develops an LSTM-based forecasting framework for predicting the daily directional movement of the EUR/USD exchange rate in the foreign exchange market. The EUR/USD currency pair is selected because it represents the most actively traded currency pair globally and accounts for the largest share of global Forex trading volume (Evans & Lyons, 2002; King et al., 2013; BIS, 2022). Due to its high liquidity and extensive market participation, EUR/USD is frequently used as a benchmark dataset in empirical studies of exchange-rate forecasting and algorithmic trading.

Rather than proposing a new neural network architecture, this research focuses on systematically examining how key model configuration parameters influence predictive performance in a realistic forecasting setting. The study first develops a baseline LSTM model for daily exchange-rate direction prediction using historical EUR/USD data. It then investigates the effects of three important parameters:

input window length, training dataset size, and model update frequency. In addition, a classification-based sequence segmentation approach, referred to as LSTM-CS, is introduced to partition price sequences into groups with similar movement patterns so that specialized LSTM models can learn more homogeneous market behaviors.

The main contributions of this study can be summarized as follows. First, the study develops an LSTM-based forecasting framework for predicting the daily direction of exchange rate movements using historical EUR/USD price data. Second, it systematically investigates several key model configuration parameters, including input window length, training dataset size, and model update frequency, in order to understand their influence on predictive performance. Third, the study introduces a classification-based sequence segmentation approach (LSTM-CS), which enables the model to specialize in learning different price movement patterns. Finally, comprehensive empirical evaluations are conducted to demonstrate how appropriate parameter configuration and sequence segmentation strategies can significantly improve predictive performance in Forex direction forecasting.

2. Materials and Methods

2.1 Long Short-Term Memory Networks

Long Short-Term Memory (LSTM) is a specialized recurrent neural network (RNN) architecture designed to capture temporal dependencies in sequential data (Hochreiter & Schmidhuber, 1997). RNN models process input sequences while maintaining hidden states that propagate information across time steps, allowing temporal relationships within time-series data to be modeled (Lim & Zohren, 2021). The basic structure of an RNN used for sequential processing is illustrated in Figure 1.




Figure 1. Structure of a recurrent neural network (RNN) for sequential data processing

However, conventional RNN models often suffer from vanishing and exploding gradient problems during backpropagation through time, making it difficult for the network to learn long-term dependencies (Hochreiter, 1998; Pascanu et al., 2013). To overcome this limitation, the LSTM architecture introduces a memory cell together with three gating mechanisms, namely the input gate, forget gate, and

output gate, which regulate the flow of information within the network. The internal structure of these gates is illustrated in Figure 2.




Figure 2. LSTM gating mechanisms, including the input gate, forget gate, and output gate

Because of these mechanisms, LSTM networks can retain relevant historical information over longer time horizons and have therefore been widely applied in financial time-series forecasting tasks such as stock-price prediction, volatility forecasting, and exchange-rate modeling (Fischer & Krauss, 2018; Sezer et al., 2020; Lim & Zohren, 2021).

For time-series prediction problems, LSTM models can be configured using different sequence structures. In this study, a many-to-one sequence prediction structure is employed, in which a sequence of historical observations is used to predict the direction of the next trading day. This structure is commonly used in financial forecasting applications, where past information influences future market movements, as illustrated in Figure 3.




Figure 3. Many-to-one sequence prediction structure used in the LSTM forecasting model

2.2 Research Framework

The overall research framework used in this study is illustrated in Figure 4. It consists of three main stages: (1) data collection and preprocessing, (2) model development and optimization, and (3) model evaluation.




Figure 4. Research framework of the proposed LSTM-based forecasting system

In the first stage, historical Forex price data are collected and transformed into a format suitable for machine learning analysis. This stage includes data cleaning, transformation of raw price data into directional labels, and generation of sequential datasets using a sliding window technique.

In the second stage, predictive models are developed and optimized. A baseline LSTM model (LSTM_Basic) is first constructed to establish a reference model for performance comparison. Several optimized model variants are then developed by systematically adjusting important parameters that may influence predictive performance.

The models examined in this study include:

- LSTM_OptIL – Optimization of input length (window size)
- LSTM_OptATD – Optimization of the amount of training data
- LSTM_OptUM – Optimization of model update frequency
- LSTM-CS – classification-based sequence segmentation model

In the final stage, the models are evaluated using historical exchange-rate data and their predictive performance is compared using classification-based evaluation metrics derived from the confusion matrix.

2.3 Data Processing: Data Collection, Preparation, and Transformation

This section describes the procedures used for data collection, preparation, and transformation.

2.3.1 Data Collection

The dataset used in this study consists of daily historical exchange-rate data from the Forex market. The data were exported from the MetaTrader 4 trading platform provided by FXPro, which is widely used in retail algorithmic trading environments (Aldridge, 2013).

The selected currency pair is EUR/USD, covering the period from 2 January 2018 to 30 April 2019, resulting in a dataset of 345 daily observations. Although the dataset contains a limited number

of daily observations, the sliding window technique generates a large number of sequential training samples, allowing the LSTM model to learn temporal patterns effectively.

The EUR/USD pair was selected because it represents the most actively traded currency pair in the global Forex market and is commonly used as a benchmark in empirical exchange-rate research (Evans & Lyons, 2002; King et al., 2013; BIS, 2022).

The dataset follows the standard OHLC format and includes the following attributes: (1) trading date, (2) opening price, (3) highest price, (4) lowest price, (5) closing price, and (6) trading volume. An example of the exported dataset is illustrated in Figure 5.

Date	Time	Open	High	Low	Close	Volume
2018.01.02	0:00	1.2087	1.20888	1.19967	1.20004	77692
2018.01.04	0:00	1.19114	1.1953	1.18672	1.1946	6936
2018.01.05	0:00	1.19495	1.19753	1.18864	1.19366	37558
2018.01.06	0:00	1.19361	1.19683	1.18486	1.18715	38120
2018.01.07	0:00	1.1872	1.18954	1.18013	1.18327	40248
2018.01.08	0:00	1.18321	1.18466	1.17537	1.17913	36991
2018.01.09	0:00	1.17905	1.18443	1.17621	1.18387	63687
2018.01.11	0:00	1.18523	1.18691	1.18357	1.18689	4385
2018.01.12	0:00	1.18676	1.18704	1.17856	1.18312	73332
2018.01.13	0:00	1.18309	1.1859	1.17531	1.17761	79404
2018.01.14	0:00	1.17759	1.18449	1.17267	1.17832	79148
2018.01.15	0:00	1.17828	1.17892	1.15675	1.16353	100901
2018.01.16	0:00	1.16357	1.16484	1.14593	1.15632	99827
2018.01.18	0:00	1.1549	1.15724	1.15457	1.15603	3663
2018.01.19	0:00	1.15602	1.16383	1.1551	1.15939	75745
2018.01.20	0:00	1.15951	1.16139	1.15397	1.15446	76298
2018.01.21	0:00	1.1544	1.16739	1.15406	1.16137	76752

Figure 5. Example of EUR/USD historical price data exported from MetaTrader 4

2.3.2 Data Preparation

To formulate the prediction problem as a classification task, an additional variable called Trend is constructed to represent the daily price movement direction. The Trend variable is defined as follows:

$$Trend = \begin{cases} 1, & \text{if } Close_t \geq Open_t \\ 0, & \text{if } Close_t < Open_t \end{cases} \quad (1)$$

where $Close_t$ and $Open_t$ represent the closing and opening prices of trading day t .

This transformation converts the forecasting problem into a binary classification problem in which the model predicts whether the exchange rate will move upward or downward within a trading day.

Date	Time	Open	High	Low	Close	Volume	Trend
2018.01.02	0:00	1.2087	1.20888	1.19967	1.20004	77692	0
2018.01.04	0:00	1.19114	1.1953	1.18672	1.1946	6936	1
2018.01.05	0:00	1.19495	1.19753	1.18864	1.19366	37558	0
2018.01.06	0:00	1.19361	1.19683	1.18486	1.18715	38120	0
2018.01.07	0:00	1.1872	1.18954	1.18013	1.18327	40248	0
2018.01.08	0:00	1.18321	1.18466	1.17537	1.17913	36991	0
2018.01.09	0:00	1.17905	1.18443	1.17621	1.18387	63687	1
2018.01.11	0:00	1.18523	1.18691	1.18357	1.18689	4385	1
2018.01.12	0:00	1.18676	1.18704	1.17856	1.18312	73332	0
2018.01.13	0:00	1.18309	1.1859	1.17531	1.17761	79404	0
2018.01.14	0:00	1.17759	1.18449	1.17267	1.17832	79148	1
2018.01.15	0:00	1.17828	1.17892	1.15675	1.16353	100901	0
2018.01.16	0:00	1.16357	1.16484	1.14593	1.15632	99827	0
2018.01.18	0:00	1.1549	1.15724	1.15457	1.15603	3663	1
2018.01.19	0:00	1.15602	1.16383	1.1551	1.15939	75745	1
2018.01.20	0:00	1.15951	1.16139	1.15397	1.15446	76298	0
2018.01.21	0:00	1.1544	1.16739	1.15406	1.16137	76752	1

Figure 6. Example of dataset with the generated Trend variable

The distribution of upward and downward movements was examined to evaluate class balance. The dataset shows a relatively balanced distribution of both classes, and therefore no additional class-balancing techniques were applied.

2.3.3 Data Transformation

After constructing the Trend variable, the dataset is transformed into sequential samples using the sliding window technique, which is commonly used in deep learning-based time-series forecasting (Brownlee, 2018; Fischer & Krauss, 2018; Lim & Zohren, 2021).

The sliding window moves across the dataset and extracts consecutive sequences of historical observations to be used as inputs for the LSTM model, as illustrated in Figure 7.




Figure 7. Sliding window technique used to generate sequential samples

Each sequence represents a set of historical directional movements used to predict the direction of the next trading day. An example of the generated sequence dataset is shown in Figure 8.

1	0	0	1	0	0	0	0	1	0	0
2	0	1	0	0	0	0	1	0	0	1
3	1	0	0	0	0	1	0	0	1	1
4	0	0	0	0	1	0	0	1	1	0
5	0	0	0	1	0	0	1	1	0	1
6	0	0	1	0	0	1	1	0	1	0
7	0	1	0	0	1	1	0	1	0	0
8	1	0	0	1	1	0	1	0	0	1
9	0	0	1	1	0	1	0	0	1	1
10	0	1	1	0	1	0	0	1	1	1
11	1	1	0	1	0	0	1	1	1	0
12	1	0	1	0	0	1	1	1	0	1
13	0	1	0	0	1	1	1	0	1	1
14	1	0	0	1	1	1	0	1	1	1
15	0	0	1	1	1	0	1	1	1	0

Figure 8. Example of binary sequence data generated from daily price movements.

2.4 Model Development and Optimization

This section describes the construction of the LSTM-based predictive models and the strategies used to improve prediction performance.

2.4.1 Problem Formulation

The forecasting task in this study is formulated as a binary classification problem. Given a sequence of historical observations:

$$X = (x_{t-n+1}, x_{t-n+2}, \dots, x_t) \quad (2)$$

the objective of the model is to predict the directional movement of the exchange rate at time $t + 1$. The conceptual prediction structure is illustrated in Figure 9.




Figure 9. Conceptual illustration of daily exchange-rate direction prediction.

Three important modeling parameters are examined: (1) input window length (lookback period), (2) training dataset size, and (3) model update frequency. These parameters may influence the predictive capability of machine learning models in non-stationary financial markets (Gu et al., 2020; Sezer et al., 2020).

2.4.2 LSTM Model Architecture

The architecture consists of an input layer, an LSTM layer for temporal feature extraction, a fully connected dense layer, and a sigmoid output layer for binary classification.

The model is trained using the binary cross-entropy loss function and optimized using the Adam optimizer, which is widely applied in deep learning optimization tasks. The main hyperparameters used in the baseline model are summarized in Table 1.

Table 1. Baseline LSTM model hyperparameters

Parameter	Value
LSTM units	32
Dense units	1
Activation	Sigmoid
Optimizer	Adam
Loss function	Binary cross-entropy
Batch size	16
Epochs	50
Learning rate	0.001

2.4.3 Development of the baseline LSTM Model (LSTM_Basic)

The baseline LSTM model is first developed to establish a reference benchmark for evaluating the effectiveness of the proposed model configurations. The overall development procedure of the baseline model is illustrated in Figure 10.

The generated sequence dataset is divided into training and testing subsets following chronological order. In time-series forecasting problems, chronological data partitioning is necessary to prevent information leakage between training and testing datasets. Historical observations occurring earlier in time are therefore used for model training, while later observations are reserved for model evaluation (Lim & Zohren, 2021).

Although the sliding window technique produces overlapping sequences, the training and testing datasets are constructed independently after the chronological split. As a result, no overlapping observations exist between the two datasets, ensuring that the predictive model does not gain access to future information during training. This approach provides a more realistic evaluation framework for financial forecasting models.




Figure 10. Development process of the baseline LSTM model

Model evaluation in this study follows a rolling chronological prediction procedure that mimics a realistic forecasting environment. At each evaluation stage, the model is trained using historical observations available up to that point and then used to generate predictions for subsequent unseen data. This procedure preserves the temporal ordering of financial time-series data and prevents information from future observations from influencing the training process. Such chronological evaluation strategies are commonly recommended for financial forecasting studies in order to avoid data leakage and ensure realistic model assessment (Sezer et al., 2020; Lim & Zohren, 2021).

The baseline model configuration uses the following parameters:

- Input length: 30 trading days
- Training dataset size: 90 trading days
- Model update: no retraining during the evaluation period

These parameter settings serve as the initial configuration for the baseline model before performing further parameter optimization in subsequent experiments.

2.4.4 Determining the Suitable Input Length (LSTM_OptIL)

To determine the most suitable input window length for daily Forex prediction, multiple configurations ranging from 5 to 30 trading days are evaluated. Shorter windows capture recent market information, while longer windows may include additional historical context. Each configuration is evaluated using the same experimental protocol to identify the window length that provides the most stable prediction performance.

The tested input lengths were selected to cover a range from short-term recent memory to longer observation windows. This range allows the experiment to examine the trade-off

between temporal relevance and the amount of historical context used by the predictive model. Similar ranges of observation windows have also been explored in previous financial time-series forecasting studies.

2.4.5 Determining the Suitable Amount of Training Data (LSTM_OptATD)

The amount of historical data used for model training may influence model performance in financial forecasting. Excessively long training datasets may introduce outdated patterns that no longer reflect current market conditions (Sezer et al., 2020).

To investigate this effect, multiple training dataset sizes are evaluated, including: 5, 7, 14, 30, 60, and 90 trading days. These experiments analyze how the amount of training data affects the predictive capability of the LSTM model.

These values were selected to represent short-, medium-, and relatively longer historical horizons commonly used in practical financial forecasting experiments. Evaluating multiple training durations allows the study to examine how the amount of historical data influences the adaptability of the model in a non-stationary financial environment.

2.4.6 Determining the Suitable Model Update Frequency (LSTM_OptUM)

Financial markets are dynamic and non-stationary environments. Prediction models trained on historical data may gradually lose predictive power if they are not updated regularly (Gu et al., 2020; Zhang et al., 2022).

To evaluate this effect, the trained model is tested over multiple consecutive trading days without retraining. The change in prediction accuracy is monitored to determine the most appropriate model update frequency.

2.4.7 Classification-Based Sequence Segmentation (LSTM-CS)

Financial time-series sequences often contain heterogeneous movement patterns, including trends, reversals, and short-term fluctuations (Sezer et al., 2020). Training a single predictive model on heterogeneous data may reduce the ability of the model to learn consistent relationships between historical observations and future price direction.

To address this issue, this study introduces a classification-based sequence segmentation approach (LSTM-CS).

The approach consists of three steps:

- 1) Generated sequences are analyzed according to the proportion of upward and downward movements within each sliding window.
- 2) Sequences are grouped into pattern categories representing different movement behaviors.
- 3) Independent LSTM models are trained for each sequence group using optimized parameter configurations.

The LSTM-CS configuration uses the following parameters:

- Input length: 10 trading days
- Training dataset size: 14 trading days
- Model update frequency: daily retraining

By training specialized models on more homogeneous subsets of data, the LSTM-CS approach aims to improve the ability of the model to capture distinct price dynamics in financial time-series data.

2.5 Performance Comparison of the LSTM Models

The predictive performance of the proposed models is evaluated using **directional prediction accuracy**, which measures the proportion of correctly predicted price movement directions relative to the total number of observations.

In financial forecasting applications, correctly predicting the **direction of price movement** is often more important than predicting the exact magnitude of price changes. Directional prediction accuracy directly influences trading decisions such as entering buy or sell positions. Therefore, directional accuracy has been widely used as an evaluation metric in financial time-series forecasting studies (Blaskowitz & Herwartz, 2011; Moosa, 2015).

The prediction accuracy is computed using Equation (3):

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \quad (3)$$

where TP (True Positive) represents correctly predicted upward price movements, TN (True Negative) represents correctly predicted downward movements, FP (False Positive) represents incorrect upward predictions, and FN (False Negative) represents missed upward movements.

These values are obtained from the confusion matrix, which summarizes the classification results of the predictive model. The structure of the confusion matrix used in this study is presented in Table 2.

Table 2. Confusion matrix used for model evaluation

	Actual Positive	Actual Negative
Predicted Positive	TP	FP
Predicted Negative	FN	TN

The confusion matrix provides a structured representation of prediction outcomes by distinguishing between correct and incorrect directional predictions. Using these values, the overall directional accuracy of each predictive model can be computed and compared.

In financial forecasting problems, even prediction accuracies slightly above random chance may still provide practical value when incorporated into trading strategies. Traders often design trading systems with asymmetric risk–reward structures, where profitable trades generate larger gains than losing trades. Under such conditions, a trading strategy may remain profitable even when prediction accuracy is close to 50%, provided that the average payoff of winning trades exceeds that of losing trades (Tharp, 2008; Vince, 1992).

Previous studies have also emphasized that directional forecasting accuracy modestly above random benchmarks can still produce economically meaningful signals when integrated into systematic trading frameworks (Blaskowitz & Herwartz, 2011; Moosa, 2015). Therefore, the purpose of this evaluation is to determine whether the proposed LSTM-based models can achieve prediction accuracy above the random benchmark while maintaining stable performance across different parameter configurations.

This evaluation framework enables a consistent comparison between the baseline LSTM model and the optimized model variants developed in this study.

3. Results

3.1 Baseline LSTM Model Performance

The baseline LSTM model was first evaluated to establish a reference benchmark for subsequent model optimization experiments. The baseline configuration employed an input window length of 30 trading days and a training dataset size of 90 trading days, without retraining during the evaluation period.

The prediction outcomes of the baseline model are summarized using a confusion matrix, as illustrated in Figure 11.

Confusion Matrix

Output Class	Up (1)	685	584	54.0%
	Down (0)	669	702	51.2%
		50.6%	54.6%	52.5%
		Up (1)	Down (0)	
		Target Class		

Figure 11. Confusion matrix of the baseline LSTM model for daily Forex price direction prediction

During the evaluation period, the model generated predictions for 2,640 testing samples obtained from sequential rolling prediction windows. Among these predictions, 1,387 observations were correctly classified, resulting in an overall directional prediction accuracy of 52.5%.

The confusion matrix values obtained from the baseline model are presented in Table 3.

Table 3. Confusion matrix of the baseline LSTM model

	Actual Up	Actual Down
Predicted Up	685	634
Predicted Down	619	702

Based on the confusion matrix, additional classification metrics were computed to provide a more comprehensive evaluation of the baseline model performance. These metrics are summarized in Table 4.

Table 4. Classification performance of the baseline LSTM model

Metric	Value (%)
Accuracy	52.5
Precision	51.9
Recall	52.5
F1-score	52.2

The results indicate that the baseline LSTM model performs only slightly better than random guessing. This outcome is consistent with previous studies showing that short-term financial direction prediction remains challenging due to the stochastic and highly dynamic nature of financial markets (Fischer & Krauss, 2018; Sezer et al., 2020)

3.2 Effect of Input Length (Lookback Window)

To determine the most appropriate observation window for daily Forex direction prediction, multiple input lengths ranging from 5 to 30 trading days were evaluated.

The prediction accuracy obtained for each configuration during the testing period from January to April 2019 is summarized in Table 5.

Table 5. Prediction accuracy of the LSTM model with different input window lengths

Testing Period	5	7	9	10	12	15	20	25	30
JAN 2019	46.6	49.9	55.4	56.1	52.3	54.2	54.0	55.1	47.8
FEB 2019	48.4	50.0	54.1	55.2	54.0	52.0	53.8	51.8	45.9
MAR 2019	50.1	50.7	52.2	54.8	56.2	54.9	54.5	49.7	52.6
APR 2019	47.7	50.5	49.7	52.0	51.5	48.0	47.2	45.0	49.7
AVG	48.2	50.3	52.9	54.5	53.5	52.3	52.4	50.4	49.0

The results show that the 10-day input window produced the highest average prediction accuracy (54.5%). Although slightly higher accuracy was observed for the 12-day window in March 2019, the 10-day window produced the most stable performance across different testing periods. Therefore, the 10-day input window was selected for the subsequent optimization experiments.

3.3 Effect of Model Update Frequency

To evaluate the impact of model updating, prediction accuracy was monitored across consecutive trading days without retraining. The experimental results are illustrated in Figure 12.




Figure 12. Prediction accuracy of the LSTM model across consecutive trading days without model retraining

The results show that prediction accuracy tends to be highest immediately after model training. In several testing periods, prediction accuracy exceeded 65% on the first evaluation day. However, prediction performance gradually declined as the model continued to generate predictions without retraining.

This decline suggests that the predictive relationships learned by the model may become less relevant as market conditions evolve over time. Therefore, frequent model updating may be necessary to maintain predictive performance in non-stationary financial environments.

3.4 Summary of Model Improvements

The overall performance of the baseline and optimized models is summarized in Table 6.

Table 6. Performance comparison of different LSTM model configurations

Model	Input Length	Training Data	Frequency Update	Accuracy (%)	Precision (%)	Recall (%)	F1-score (%)
LSTM_Basic	30	90	–	52.5	51.9	52.5	52.2
LSTM_OptIL	10	90	–	54.5	53.8	54.5	54.1
LSTM_OptATD	10	14	–	61.2	60.4	61.2	60.8
LSTM_OptUM	10	14	Daily	67.5	66.9	67.5	67.2
LSTM-CS	10	14	Daily	69.7	69.1	69.7	69.4

The results demonstrate that systematic parameter optimization significantly improves the predictive performance of the LSTM model.

First, adjusting the input window length from 30 days to 10 days improved prediction accuracy from 52.5% to 54.5%.

Second, reducing the training dataset size from 90 days to 14 days further increased prediction accuracy to 61.2%, indicating that more recent data may provide more relevant predictive signals.

Third, introducing daily model retraining increased prediction accuracy to 67.5%, suggesting that frequent model updates help maintain predictive performance under evolving market conditions.

Finally, the classification-based sequence segmentation approach (LSTM-CS) achieved the highest performance, with an accuracy of 69.7% and improved precision, recall, and F1-score values compared with the other models.

These results indicate that both model configuration and sequence segmentation strategies can contribute to improving prediction performance in daily Forex direction forecasting.

4. Discussion

4.1 Interpretation of Experimental Results

The experimental results demonstrate that predictive model configuration plays a critical role in the forecasting performance of Long Short-Term Memory (LSTM) networks for daily exchange-rate direction prediction. The baseline LSTM model achieved a directional prediction accuracy of 52.5%, which is only slightly higher than random guessing. This outcome is consistent with previous studies indicating that short-term financial direction prediction remains inherently difficult because of the stochastic and highly dynamic nature of financial markets (Fischer & Krauss, 2018; Sezer et al., 2020; García et al., 2024).

One of the key findings of this study is that reducing the input window length significantly improved predictive performance. The experimental results showed that the 10-day input window produced the most stable prediction accuracy across different testing periods. This observation suggests that excessively long observation windows may introduce outdated or irrelevant historical information into the learning process. Because financial markets evolve rapidly, older observations may no longer reflect

current market conditions. Therefore, shorter observation windows may help the predictive model focus on more recent price dynamics that contain stronger predictive signals. Similar findings have been reported in previous studies of financial time-series forecasting, in which recent data often provide more informative signals for short-horizon prediction tasks (Lim & Zohren, 2021).

Another important observation concerns the effect of training dataset size on prediction accuracy. The experimental results show that reducing the training dataset size from 90 trading days to 14 trading days improved prediction accuracy from 54.5% to 61.2%. This result suggests that incorporating large amounts of historical data does not necessarily improve predictive performance in non-stationary financial environments. When historical data span longer periods, the statistical relationships within the dataset may change because of evolving economic conditions, market sentiment, or policy changes. As a result, including excessive historical data may introduce outdated patterns that reduce the predictive capability of machine learning models. Similar conclusions have been reported in previous research on financial forecasting, where adaptive training strategies are often required to address the non-stationarity of financial time-series data (Gu et al., 2020; Sezer et al., 2020).

The experiment on model update frequency further supports this interpretation. The results show that prediction accuracy gradually declines when the trained model is applied across multiple consecutive trading days without retraining. This performance degradation indicates that predictive relationships learned from historical data may become less relevant as market conditions evolve over time. Frequent model updating therefore appears necessary to maintain forecasting performance in dynamic financial markets. Previous studies on machine learning-based financial forecasting have similarly emphasized the importance of periodic model retraining to adapt to changing market conditions (Gu et al., 2020; Zhang et al., 2022).

The highest predictive performance in this study was achieved when the classification-based sequence segmentation approach (LSTM-CS) was applied, resulting in a prediction accuracy of 69.7%. This approach partitions price sequences into groups with similar movement characteristics, allowing specialized LSTM models to learn more homogeneous patterns. Financial time-series data often exhibit heterogeneous structures, including trending periods, reversals, and irregular fluctuations. Training a single predictive model on heterogeneous sequences may reduce the model's ability to capture consistent patterns. By grouping sequences according to their directional characteristics, the LSTM-CS approach reduces pattern heterogeneity and enables each model to specialize in learning specific price movement behaviors.

In the context of financial forecasting research, achieving prediction accuracies substantially above random chance remains challenging. Previous studies commonly report directional prediction accuracies in the range of approximately 55% to 65% for short-term financial forecasting tasks (Fischer & Krauss, 2018; Sezer et al., 2020; Garcia et al., 2024). In this context, the maximum prediction accuracy of 69.7% obtained in this study suggests that combining parameter optimization with sequence segmentation strategies may be a promising direction for improving deep learning-based forecasting models. Recent research focusing specifically on EUR/USD direction prediction using machine learning methods further confirms the continued importance of this problem in quantitative finance (Guyard & Deriaz, 2024).

4.2 Limitations of the Study

Despite the promising results, several limitations should be acknowledged. First, the experiments were conducted using only the EUR/USD currency pair, and the predictive performance of the proposed framework may vary when applied to other currency pairs with different market characteristics. Second, the dataset used in this study covers a relatively limited historical period, which may restrict the generalizability of the results. Third, this study focuses on performance differences among LSTM-based configurations rather than comparisons across different model families. Therefore, the results should be interpreted as improvements within the LSTM framework rather than a claim of superiority over alternative forecasting approaches. Finally, the evaluation focused primarily on directional prediction accuracy, and the study did not examine the economic profitability of the model under realistic trading conditions such as transaction costs and market frictions.

5. Conclusion

This study developed and evaluated an LSTM-based framework for predicting the daily directional movement of exchange rates in the foreign exchange market using historical EUR/USD data. The results demonstrate that model configuration plays an important role in the predictive performance of LSTM models for financial time-series forecasting. The baseline LSTM model achieved a prediction accuracy of 52.5%, which is only slightly higher than random chance. However, systematic parameter optimization significantly improved forecasting performance.

The experimental results indicate that reducing the input window length to 10 trading days and limiting the training dataset to 14 recent trading days improved prediction accuracy, suggesting that recent market information provides more relevant predictive signals in dynamic financial environments. In addition, frequent model updating helped maintain prediction performance over time. The proposed classification-based sequence segmentation approach (LSTM-CS) further improved prediction accuracy to 69.7%, indicating that training specialized models on more homogeneous sequence groups can enhance forecasting capability.

Future research may extend this framework by evaluating multiple currency pairs, incorporating additional market features, and examining the economic profitability of the proposed forecasting approach under realistic trading conditions.

6. References

- Aldridge, I. (2013). *High-frequency trading: a practical guide to algorithmic strategies and trading systems*. John Wiley & Sons.
- Atsalakis, G. S., & Valavanis, K. P. (2009). Surveying stock market forecasting techniques—Part II: Soft computing methods. *Expert Systems with applications*, 36(3), 5932–5941. <https://doi.org/10.1016/j.eswa.2008.07.006>
- Bank for International Settlements. (2022). OTC foreign exchange turnover in April 2022. Triennial Central Bank Survey. Retrieved December 12, 2025, from https://www.bis.org/statistics/rpfx22_fx.htm
- Blaskowitz, O., & Herwartz, H. (2011). On economic evaluation of directional forecasts. *International Journal of Forecasting*, 27(4), 1058–1065. <https://doi.org/10.1016/j.ijforecast.2010.09.003>

- Brownlee, J. (2018). Deep learning for time series forecasting: Predict the future with MLPs, CNNs and LSTMs in Python. Machine Learning Mastery. Retrieved December 12, 2025, from <https://machinelearningmastery.com/deep-learning-for-time-series-forecasting/>
- Cao, L. J., Tay, F. E. H., & Liu, J. (2003). Support vector machines with adaptive parameters in financial time series forecasting. *IEEE Transactions on Neural Networks*, 14(6), 1506–1518. <https://doi.org/10.1109/TNN.2003.820556>
- Evans, M. D. D., & Lyons, R. K. (2002). Order flow and exchange rate dynamics. *Journal of Political Economy*, 110(1), 170–180. <https://doi.org/10.1086/324391>
- Fischer, T., & Krauss, C. (2018). Deep learning with long short-term memory networks for financial market predictions. *European Journal of Operational Research*, 270(2), 654–669. <https://doi.org/10.1016/j.ejor.2017.11.054>
- García, F., Guíjarro, F., Oliver, J., & Tamošiūnienė, R. (2024). Foreign exchange forecasting models: LSTM and BiLSTM comparison. *Engineering Proceedings*, 68(1), Article 19. <https://doi.org/10.3390/engproc2024068019>
- Gu, S., Kelly, B., & Xiu, D. (2020). Empirical asset pricing via machine learning. *The Review of Financial Studies*, 33(5), 2223–2273. <https://doi.org/10.1093/rfs/hhaa009>
- Guyard, K. C., & Deriaz, M. (2024, August). Predicting foreign exchange EUR/USD direction using machine learning. In *Proceedings of the 2024 7th International Conference on Machine Learning and Machine Intelligence (MLMI)* (pp. 1-9). <https://doi.org/10.1145/3696271.3696272>
- Hochreiter, S. (1998). The vanishing gradient problem during learning recurrent neural networks. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 6(2), 107–116. <https://doi.org/10.1142/S0218488598000094>
- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>
- Kamruzzaman, J., & Sarker, R. A. (2003). Comparing ANN based models with ARIMA for prediction of forex rates. *ASOR Bulletin*, 22(2), 2–11.
- King, M. R., Osler, C., & Rime, D. (2013). The market microstructure approach to foreign exchange: Looking back and looking forward. *Journal of Economic Literature*, 51(4), 1090–1134. <https://doi.org/10.1257/jel.51.4.1090>
- Lim, B., & Zohren, S. (2021). Time-series forecasting with deep learning: A survey. *Philosophical Transactions of the Royal Society A*, 379(2194), Article 20200209. <https://doi.org/10.1098/rsta.2020.0209>
- Lo, A. W., Mamaysky, H., & Wang, J. (2000). Foundations of technical analysis: Computational algorithms, statistical inference, and empirical implementation. *The Journal of Finance*, 55(4), 1705–1765. <https://doi.org/10.1111/0022-1082.00265>
- Lyons, R. K. (2001). *The microstructure approach to exchange rates*. the MIT Press.
- Moosa, I. A. (2015). On the evaluation of directional forecasts. *Applied Economics*, 47(57), 6191–6199. <https://doi.org/10.1080/00036846.2015.1064086>
- Nguyen, P. D., Thao, N. N., Chi, D. T. K., Nguyen, H. C., Mach, B. N., & Nguyen, T. Q. (2024). Deep learning-based predictive models for forex market trends: Practical implementation and performance evaluation. *Science Progress*, 107(3), 00368504241275370. <https://doi.org/10.1177/00368504241275370>
- Pascanu, R., Mikolov, T., & Bengio, Y. (2013). On the difficulty of training recurrent neural networks. *Neural Networks*, 37, 1–11. <https://doi.org/10.1016/j.neunet.2012.02.010>
- Poon, S. H., & Granger, C. W. J. (2003). Forecasting volatility in financial markets: A review. *Journal of Economic Literature*, 41(2), 478–539. <https://doi.org/10.1257/002205103765762743>
- Rapach, D., & Zhou, G. (2013). Forecasting stock returns. In *Handbook of economic forecasting* (Vol. 2, pp. 328-383). Elsevier.
- Sezer, O. B., Gudelek, M. U., & Ozbayoglu, A. M. (2020). Financial time series forecasting with deep learning: A systematic literature review: 2005–2019. *Applied Soft Computing*, 90, 106181. <https://doi.org/10.1016/j.asoc.2020.106181>
- Tharp, V. K., Chabot, C., & Tharp, K. (2007). *Trade your way to financial freedom* (p. 343). New York, NY, USA: McGraw-Hill.

- Vince, R. (1992). *The mathematics of money management: risk analysis techniques for traders*. John Wiley & Sons.
- Zhang, Z., Zohren, S., & Roberts, S. (2022). Deep learning for volatility forecasting. *Quantitative Finance*, 22(2), 357–372.
<https://doi.org/10.1080/14697688.2021.1989139>