

EUR/USD Exchange Rate Forecasting Based on Information Fusion with Large Language Models and Deep Learning Methods

Hongcheng Ding*, Xuanze Zhao, Ruiting Deng, Shamsul Nahar Abdullah,
Deshinta Arrova Dewi

INTI International University, Malaysia
i24025877@student.newinti.edu.my

June 30, 2025

Abstract

Accurate forecasting of the EUR/USD exchange rate is crucial for investors, businesses, and policymakers. This paper proposes a novel framework, IUS, that integrates unstructured textual data from news and analysis with structured data on exchange rates and financial indicators to enhance exchange rate prediction. The IUS framework employs large language models for sentiment polarity scoring and exchange rate movement classification of texts. These textual features are combined with quantitative features and input into a Causality-Driven Feature Generator. An Optuna-optimized Bi-LSTM model is then used to forecast the EUR/USD exchange rate. Experiments demonstrate that the proposed method outperforms benchmark models, reducing MAE by 10.69% and RMSE by 9.56% compared to the best performing baseline. Results also show the benefits of data fusion, with the combination of unstructured and structured data yielding higher accuracy than structured data alone. Furthermore, feature selection using the top 12 important quantitative features combined with the textual features proves most effective. The proposed IUS framework and Optuna-Bi-LSTM model provide a powerful new approach for exchange rate forecasting through multi-source data integration.

Keywords: exchange rate forecasting, EUR/USD, sentiment analysis, textual data, large language models, feature generation, Bi-LSTM, Optuna

1 Introduction

The exchange rate between the Euro and the US Dollar is a significant indicator in the global financial market, reflecting the economic dynamics between two of the world's largest economies. Precise prediction of the EUR/USD exchange rate is crucial for individual investors, businesses engaged in international trade, and policymakers responsible for economic stability and growth. Traditionally, econometric models have been utilized to forecast exchange rates, relying heavily on historical market data and macroeconomic indicators released by governments and financial organizations [1]. Although these datasets are comprehensive, their low publication frequency makes it difficult to capture real-time market volatility and nonlinear dynamics [2].

The integration of unstructured data from diverse sources, such as news articles, financial reports and social media platforms, has the potential to improve the accuracy of exchange rate forecasting. In recent years, the significant impact of political events, global economic crises, and unexpected international incidents on currency fluctuations has been recognized [3], suggesting that considering a wider range of information beyond traditional structured data may be beneficial. The great amount of textual data may contain valuable insights into market sentiment, economic trends, and key events that can influence exchange rates [4]. However, exchange rate forecasting presents two significant challenges. Firstly, while the relationship between news information and market trends is relatively straightforward in traditional financial markets, the complex semantics of market-driven news and analysis texts in the context of exchange rates pose difficulties for sentiment analysis. Secondly, traditional methods struggle to adequately capture the complex nonlinear patterns and unstructured

relationships hidden within the textual data, limiting their ability to provide accurate predictions in the dynamically changing foreign exchange market[5].

Recent attempts to address these challenges by integrating textual data and deep learning methods show promise, but limitations still exist. Singh et al. [6] propose using Word2Vec and LSTM to classify collected Weibo text data and assign sentiment weights to each word. The sentiment analysis results are then incorporated into a CNN-LSTM hybrid model for exchange rate prediction. However, the texts related to exchange rates often contain news or comments about two countries simultaneously, making it challenging to accurately attribute the content to a specific country. The existence of massive noise and the involvement of information related to two currencies in the text used for exchange rate forecasting create challenges in handling lengthy texts and complex semantics. Tadphale et al. [7] utilize news headlines for sentiment analysis and then combine with other market indicators as inputs to an LSTM model for exchange rate prediction. While this approach shows potential, news headlines have limited information content and lack complete context due to text length. Moreover, the news of each day may also have a potential impact on the exchange rate movement of the next day, which is not fully captured by the proposed approach.

In this paper, we present a novel framework that integrates unstructured and structured (IUS) data to forecast EUR/USD exchange rates. We use ChatGPT-4.0 to filter out noise from collected news and analysis texts, extracting segments related to the exchange rate, in order to obtain the initial dataset. The dataset is then annotated with sentiment polarity scores by ChatGPT-4.0 and next-day exchange rate movements, e.g. up or down. Subsequently, we fine-tune two large language models (LLMs) on this dataset to create textual features, textual sentiment and exchange rate movement features. Next, the collected exchange rate and financial market data are utilized to generate quantitative features, which are integrated with textual features and inputted into a Causality-Driven Feature Generator. Finally, all generated features are fed into an Optuna-Bi-LSTM model to predict the EUR/USD exchange rate.

Our proposed method effectively addresses the challenges of processing news articles containing information about both countries in an exchange rate pair, handling high noise, complex semantics, the lack of contextual information in brief texts, and the various extraction of textual features. We validate the effectiveness of our proposed exchange rate prediction approach on our EUR/USD exchange rate dataset. The results demonstrate that our model outperforms the strongest benchmark models by 10.69% in terms of Mean Absolute Error (MAE) and by 9.56% in terms of Root Mean Squared Error (RMSE). As for data fusion, by combining unstructured and structured data, the Optuna-Bi-LSTM model is able to enhance prediction accuracy beyond what is possible with structured data alone. Furthermore, using the top 12 important features selected by the RFE method combined with 31 textual features proves to be more effective compared to analyzing all textual features, as it more directly corresponds to the actual exchange rate response to market conditions.

This research makes several significant contributions to the field of exchange rate prediction:

- We propose a novel IUS framework that integrates unstructured and structured data to predict exchange rates, setting a new benchmark in the industry.
- By utilizing LLMs, we develop an annotated EUR/USD exchange rate text dataset through a multi-step processing of raw textual data, introducing a new approach to sentiment analysis for complex semantic texts.
- Through extensive literature review and the application of a Causality-Driven Feature Generator, we construct a comprehensive feature set that incorporates a wide range of economic, financial, and textual indicators.
- To capture the nonlinear dynamics of exchange rate time series, we employ a Bi-LSTM deep learning model which is further optimized by the Optuna hyperparameter optimization framework. This ensures optimal predictive performance and enhances the generalizability of our methods.

2 Motivation

Traditional sentiment analysis methods in exchange rate prediction have primarily focused on brief texts such as Twitter posts and news headlines. These methods neglect the rich semantics and diverse information contained within longer content, which are essential for a thorough understanding of market sentiment. Additionally, the effectiveness of dictionary-based sentiment analysis heavily




Figure 1: Comparative analysis of sentiment impact on traditional vs. exchange rate markets.

depends on the quality and comprehensiveness of predefined word lists. These methods face challenges in adapting to a changing market because sentiment dictionaries tend to be static and are not immediately updated, thus failing to reflect domain-specific vocabulary. More critically, traditional methods struggle to capture the differences in numerical data and its implications for market sentiment. However, in a financial context, the extent of data growth and the variations in percentages often reflect market fluctuations. For instance, a 0.005% increase in a specific data point might influence the market differently compared to a 50% increase.

The relationship between news information and market trends is relatively straightforward in traditional markets, as illustrated in Figure 1. However, the complex semantics of news and analysis texts in the context of exchange rates pose difficulties for sentiment analysis. The nature of financial texts makes sentiment analysis and annotation a particularly challenging task, as it often contains specialized jargon, implicit sentiments linked to market conditions, and subtle variations across different sectors. Moreover, exchange rates involve two countries, and the impact of primarily positive or negative news on financial markets can be significant. For instance, an abundance of positive news about the U.S. can lead to a strengthening of the U.S. dollar, which typically results in a decrease in the EUR/USD exchange rate, and vice versa. However, news texts often contain associated information related to both countries, making classification difficult, and the positive or negative news from both countries has a zero-sum relationship in the final sentiment analysis.

To tackle the challenges of sentiment analysis in extensive financial texts, we utilize the capabilities of LLMs. As depicted in Figure 2(A), LLMs are structured with deep and hierarchical Transformer blocks, each featuring a multi-head self-attention mechanism. This architecture allows LLMs to recognize complex patterns and capture long-range dependencies within the text. As texts progress through these Transformer blocks, illustrated in Figure 2(B), the multiple attention heads within a block enable simultaneous focus on various textual elements. The different lines represent the varying attention paid by the multi-head attention mechanism to elements at different positions, revealing complex interrelationships even among distant sentences. Following this, the layered structure shown in Figure 2(C), with its consistent application of self-attention, effectively addresses the complexities of lengthy financial texts. By recognizing long-distance dependencies and maintaining focus on relevant elements throughout the input, LLMs can analyze every detail of the text to generate an overall sentiment polarity score. Moreover, LLMs have the ability to capture the differences in numerical data, enabling them to provide more accurate sentiment polarity scores that reflect the actual impact of these figures on market sentiment.

We employ ChatGPT-4.0 and several traditional dictionary-based sentiment analysis tools, including Vader, TextBlob, and Afinn, to conduct sentiment analysis on collected news and analysis texts and compare these data with trends in the forex market. As demonstrated in Figure 3, the traditional tools exhibit poor performance when processing longer texts, exhibiting significant variety in the sentiment scores they produce, which correlate poorly with actual market movements in forex trading. In contrast, the results from ChatGPT-4.0 show a higher consistency with both the next-day movements and change percentages.

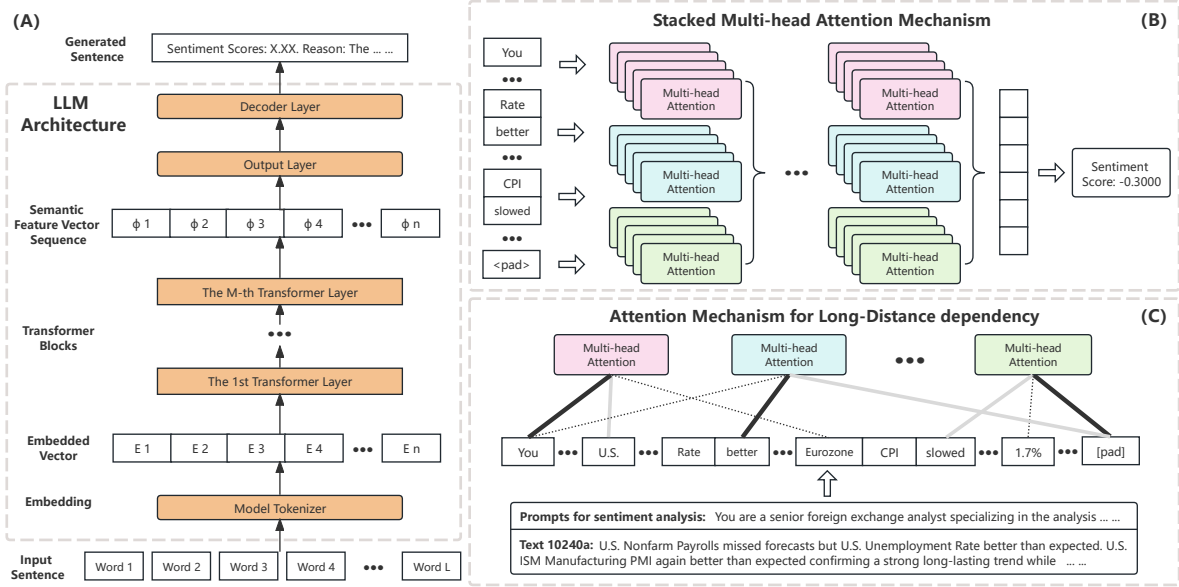


Figure 2: Capturing long-distance dependencies in large language models (LLMs): (A) LLM architecture with stacked transformer blocks and multi-head attention mechanisms; (B) Attention mechanism focusing on different positions within the input sequence; (C) Stacked multi-head attention mechanism maintaining focus on relevant elements throughout the lengthy text.

3 Related Work

1. Current studies on exchange rate forecasting

Haider et al. [8] examine whether commodity prices can forecast exchange rates in commodity-dependent economies using both in-sample and out-of-sample techniques. By modeling commodity prices to predict USD rates, their findings indicate this approach is more effective than a random walk model, providing valuable perspectives across various economies. Sarkar and Ali [9] analyze linear regression for predicting EUR/USD exchange rates using normalized daily and hourly data. Their research applies this approach to different time series, offering strategies to help traders mitigate issues and enhance profitability in the forex market. Ruan et al. [10] evaluate whether economic policy uncertainty (EPU) outperforms traditional macroeconomic indicators in predicting exchange rate volatility in both developed and emerging markets. Their results demonstrate the superior predictive capability of EPU, suggesting significant implications for risk management and policy-making, and recommending broader application to verify these findings' generalizability. Windsor and Cao [11] develop a comprehensive system using market indicators and investor sentiments to predict the USD/CNY exchange rate. This innovative system effectively captures complex interactions among various financial factors, providing a precise and robust forecasting tool. Salisu et al. [12] demonstrate that oil prices are a reliable indicator of exchange rate returns for both net oil exporters and importers. Their study emphasizes the importance of considering asymmetries in the data, which substantially enhances the predictability of an oil-based model. The results underscore the potential of oil prices as a crucial factor in financial forecasting models. Neghab et al. [13] employ machine learning techniques, including linear regression, tree-based models, and deep learning, to forecast exchange rates based on macroeconomic fundamentals. The study addresses challenges such as nonlinearity, multicollinearity, time variation, and noise in modeling.

2. The application of unstructured data in predictions

Ito and Takeda [14] improve the accuracy of exchange rate models by using sentiment indices constructed from Google search volumes of financial terms. This approach effectively captures timely market sentiments, although the generalizability of their findings requires further exploration due to the analysis's limited scope. Ben Omrane et al. [15] investigate the impact of US and EU macroeconomic news on the volatility and returns of the EUR/USD exchange rate using regime smooth transition regression. Their findings indicate that the effects of news vary between economic states, with US news generally having a larger impact than EU news on currency fluctuations. Li et al. [16] introduce

	ChatGPT-4.0	Vader	TextBlob	Afinn	Next-day Movement	Next-day Change %
Text 2562a: After starting the week with gains, the euro has ...	-0.80	-0.31	0.04	-0.52	Down	-0.68%
Text 7389c: EURUSD on Sharp Bearish Decline As Interest ...	-0.60	0.80	0.07	0.06	Down	-0.42%
Text 10877a: The European election risk has undoubtedly influenced ...	-0.30	0.73	0.09	-0.11	Down	-0.17%
Text 4230a: EUR/USD meets minor resistance at 2-month highs ...	0.00	0.04	0.18	-0.23	-	-0.01%
Text 5646a: The EUR/USD daily Forex chart has been sideways ...	0.00	0.66	0.05	-0.11	-	0.00%
Text 9959a: EUR/USD meets first support at 1.1370/60 and with ...	0.00	-0.61	0.06	-0.01	-	0.01%
Text 4445a: EUR/USD Mid-Session Technical Analysis for August ...	0.25	0.77	0.04	-0.11	Up	0.28%
Text 4763b: The Bull Case The euro is a buy for one big reason: ...	0.60	0.96	0.08	-0.01	Up	0.74%
Text 4221a: EUR/USD edged north on Friday, breaking above the ...	0.90	-0.40	0.07	-0.04	Up	1.19%

Figure 3: Comparative analysis of sentiment scores by ChatGPT-4.0 and traditional tools.

a novel approach for forecasting crude oil prices that incorporates online news text mining to extract sentiment features and group news by topic, effectively capturing the immediate market impacts of various events. This method enriches traditional forecasting models, although it is limited by its reliance on a single news source and potential noise in online news data. Bai et al. [17] propose a robust framework for forecasting crude oil prices that exploits advanced text mining techniques on news headlines to construct high-quality features. They introduce novel indicators for topic and sentiment analysis tailored for short texts, enhancing the model’s performance. However, the study does not deeply investigate the relative importance of textual features compared to non-textual factors. Swathi et al. [18] employ Twitter sentiment analysis to predict stock prices by analyzing emotions and opinions in stock-related tweets. This method provides deeper insights into market sentiment and enhances the predictive performance of their model, although it predominantly relies on Twitter, suggesting the potential benefits of incorporating more diverse data sources. Kalamara et al. [19] explore the use of newspaper text to extract economic signals, demonstrating that such information can substantially improve macroeconomic forecasts. By combining a large array of text-derived regressors with supervised machine learning, they achieve significant forecast improvements, especially during periods of economic stress. Naeem et al. [20] utilize a machine learning approach to forecast the USD/PKR exchange rate, employing sentiment analysis of finance-related tweets and various machine learning classifiers to process and optimize the dataset. This innovative method showcases the potential of leveraging social media data for financial predictions. Lv et al. [21] develop a hybrid model that combines sentiment analysis of Weibo text data with historical exchange rate information to predict market trends. This method improves prediction accuracy by integrating the perspectives of market participants, demonstrating the substantial impact of social media sentiment on forecasting exchange rates. Küçüklerli and Ulusoy [22] integrate Twitter sentiment analysis with economic indicators to predict the USD/TL exchange rate using machine learning techniques. They collect exchange rate data and finance-related tweets, preprocess the data, and employ various ML models. The LSTM model achieves the highest accuracy of 65% in forecasting the exchange rate. Semiromi et al. (2024) predict currency pairs using news story events from the economic calendar and machine learning techniques. They use text mining methods, sentiment analysis with a new sentiment dictionary, and machine learning algorithms, achieving over 60% accuracy in predicting exchange rate movements.

3. Traditional Predictive Methods

Traditional econometric models and machine learning techniques have been extensively applied to exchange rate forecasting. Li et al. [23] utilize ARIMA and GARCH models to predict USD/EUR exchange rate fluctuations, showing GARCH’s effectiveness in capturing financial data volatility while noting the limitations of ARIMA in exchange rate forecasting. Zhang [24] applies Simple Exponential

Smoothing and ARIMA models to forecast the RMB/USD exchange rate, exploring their advantages in capturing exchange rate dynamics and providing accurate predictions. Colombo and Pelagatti [25] use various statistical learning methods, including regularized regression splines, RF, and SVM, to assess exchange rate models' predictive power. The sticky price monetary model with error correction specification shows strong forecasting performance at different horizons, outperforming the random walk benchmark, providing insights into the non-linear relationship between exchange rates and fundamentals. Pfahler [26] employs artificial neural networks and gradient-boosted decision trees to forecast exchange rate movements using macroeconomic fundamentals from purchasing power parity, uncovered interest rate parity, and monetary model theories. These machine learning models outperform the random walk benchmark, especially when time dummies are included, highlighting the potential of complex interactions between time dummies and fundamentals in exchange rate forecasting. Khoa and Huynh [27] apply SVR under the uncovered interest rate parity framework to forecast the VND/USD exchange rate during the COVID-19 pandemic. Combining the framework with the robust SVR technique demonstrates superior performance compared to OLS regression and random walk models. However, these econometric models and traditional machine learning techniques face limitations in capturing the complex, non-linear dynamics and high-dimensional relationships inherent in exchange rate data.

4. Modern methods

Sun et al. [28] introduce a methodology that combines LSTM with bagging ensemble learning, significantly enhancing the accuracy and profitability of exchange rate forecasts. This method outperforms traditional benchmark models, although it is solely applied to univariate exchange rate series without considering additional influencing factors. Liu et al. [29] develop an innovative LASSO-Bi-LSTM ensemble learning strategy that merges the LASSO with Bi-LSTM networks for forecasting the USD/CNY exchange rate. This approach demonstrates superior accuracy over other deep learning models but is limited to short-term forecasts and is not tested in broader financial markets. Islam and Hossain [30] craft a hybrid model that integrates GRU and LSTM networks to predict currency prices in the forex market for key pairs. The hybrid model surpasses the performance of LSTM, GRU and a simple moving average approach in terms of accuracy and risk-adjusted returns. Despite its success, the model sometimes struggles with abrupt price fluctuations and requires specific adjustments for optimal performance with the GBP/USD pair. Dautel et al. [31] conduct an empirical analysis comparing various deep learning frameworks, including LSTM and GRU networks, for exchange rate prediction. The study provides valuable insights into the practical application of these models for financial forecasting, though it acknowledges challenges in model tuning and the disparity between statistical accuracy and economic relevancy. Wan et al. [32] introduce the Multivariate TCN, which uses a deep CNN with multichannel residual blocks and an asymmetric structure to enhance forecasting in aperiodic multivariate time series. This model shows marked improvements over other algorithms, and the study suggests that focusing on higher-order statistical features could simplify the model and boost performance. Zeng et al. [33] critically assess the efficacy of Transformer-based solutions for long-term time series forecasting, proposing a straightforward one-layer linear model, which outperforms more complex Transformer-based models across multiple datasets. However, the simplicity of this model limits its capacity. Optimization algorithms play a vital role in improving the performance of deep learning networks by systematically tuning their hyperparameters. Xu et al. [34] propose a Particle Swarm Optimization with LSTM (PSO-LSTM) model that leverages particle swarm optimization to optimize LSTM hyperparameters, mimicking the social behavior of bird flocking to discover optimal solutions. Hamdia et al. [35] employ a Genetic Algorithm, drawing inspiration from natural selection and evolution, to optimize deep neural network architectures. Victoria and Maragatham [36] utilize Bayesian Optimization, which constructs a probabilistic model of the objective function to guide the search for the best hyperparameters. Dong et al. [37] introduce an approach using Deep Reinforcement Learning, specifically a continuous Deep Q-learning algorithm with a heuristic strategy, for adaptive hyperparameter optimization in visual object tracking. García Amboage et al. [38] propose Swift-Hyperband, integrating performance prediction via SVR with early stopping methods to streamline the optimization process. Brodzicki et al. [39] apply the Whale Optimization Algorithm (WOA), inspired by the foraging behavior of humpback whales, for deep neural network hyperparameter optimization.

4 Methodology

4.1 The IUS Framework

In this work, we introduce the IUS Framework, as illustrated in Figure 4, which consists of five technical components. The first component is the Sentiment Polarity Scoring Module (SPSM), which employs an embedding generator based on a fine-tuned version of the RoBERTa-Large model, specifically adapted for sentiment analysis. This module generates the sentiment polarity feature tensor (S_{fx}) from news and analysis texts related to the target exchange rate over D trading days, capturing the sentiment representation of all texts during this period. The second component, the Movement Classification Module (MCM), utilizes the original RoBERTa-Large model. Configured as an embedding generator, it produces the feature tensor (M_{fx}) for analyzing the next-day exchange rate movement, e.g., up or down, capturing the movement representation of texts within the same period. Subsequently, we extract quantitative indicator subsequences for the target exchange rate and its related exchange rates, as well as financial market indicators over D trading days. E_{fx} and F_{fx} represent the quantitative features of these rates and financial indicators, respectively, over the same period. Finally, all these features are integrated and processed through a Causality-Driven Feature Generator, then input into the Bi-LSTM model to predict the closing price of the target exchange rate for the next trading day y_{t+1} .

4.1.1 The IUS Framework

For the textual dataset, we collect data from investing.com and forexempire.com, covering the period from February 6, 2016, to January 19, 2024. The dataset includes all accessible data on these platforms within the period, totaling 35,427 texts. However, due to the potential existence of noise and irrelevant information, we utilize ChatGPT-4.0 and prompt engineering techniques to filter the raw dataset. This data annotation approach aligns with prior research, such as using LLMs for automatic data annotation to detect hallucinations [40], keyword annotation and document content description generation [41], and math problem knowledge tagging with few-shot learning [42]. After filtering, we observe that the news and analysis texts inherently contain a higher level of noise, which may be attributed to the necessity of catering to the diverse needs of readers. In addition, we notice that often only individual paragraphs or multiple segments within an article are directly relevant to the EUR/USD exchange rate. To further refine the dataset and extract the relevant segments, we employ ChatGPT-4.0 again to process the text data, creating a final textual dataset consisting of 20,329 texts. We employ ChatGPT-4.0 to annotate the sentiment polarity scores for the textual training dataset by integrating prompt engineering techniques. To safeguard against potential and unidentified errors, the model is also required to provide explanations for the polarity scores it assigns. In our prompt engineering, we define the polarity score range as $[-1, 1]$, where scores approaching 1 indicate a strongly positive sentiment, and vice versa. Scores near zero represent a neutral sentiment. In terms of next-day exchange rate movement, e.g., up or down, M_d is used to annotate the text data. M_d is defined as:

$$M_d = \begin{cases} 0, & CP_{d+1} < CP_d, \\ 1, & CP_{d+1} \geq CP_d, \end{cases} \quad (1)$$

where CP_{d-1} is the closing price of the exchange rate on trading day $d + 1$ and CP_d is the closing price on day d . We do not introduce an additional label for $CP_{d+1} = CP_d$, as it is rare for the closing prices to be the same on two consecutive transaction days.

4.1.2 RoBERTa-Large

RoBERTa-Large, developed by Facebook AI, is a SOTA pre-trained LLM incorporating several key enhancements to improve its training process and architecture [43]. Utilizing increased training data, larger batch sizes, and extended training periods, RoBERTa-Large has demonstrated outstanding performance across a variety of benchmark tasks, showcasing its robust capabilities. Its architecture, as illustrated in Figure 5(A), is based on the transformer model and primarily consists of a tokenization module and 24 encoder blocks. Each encoder block, detailed in Figure 5(B), includes a multi-head self-attention mechanism followed by a feed-forward neural network. The self-attention mechanism allows the model to focus on different positions within the input sequence, capturing the relationships

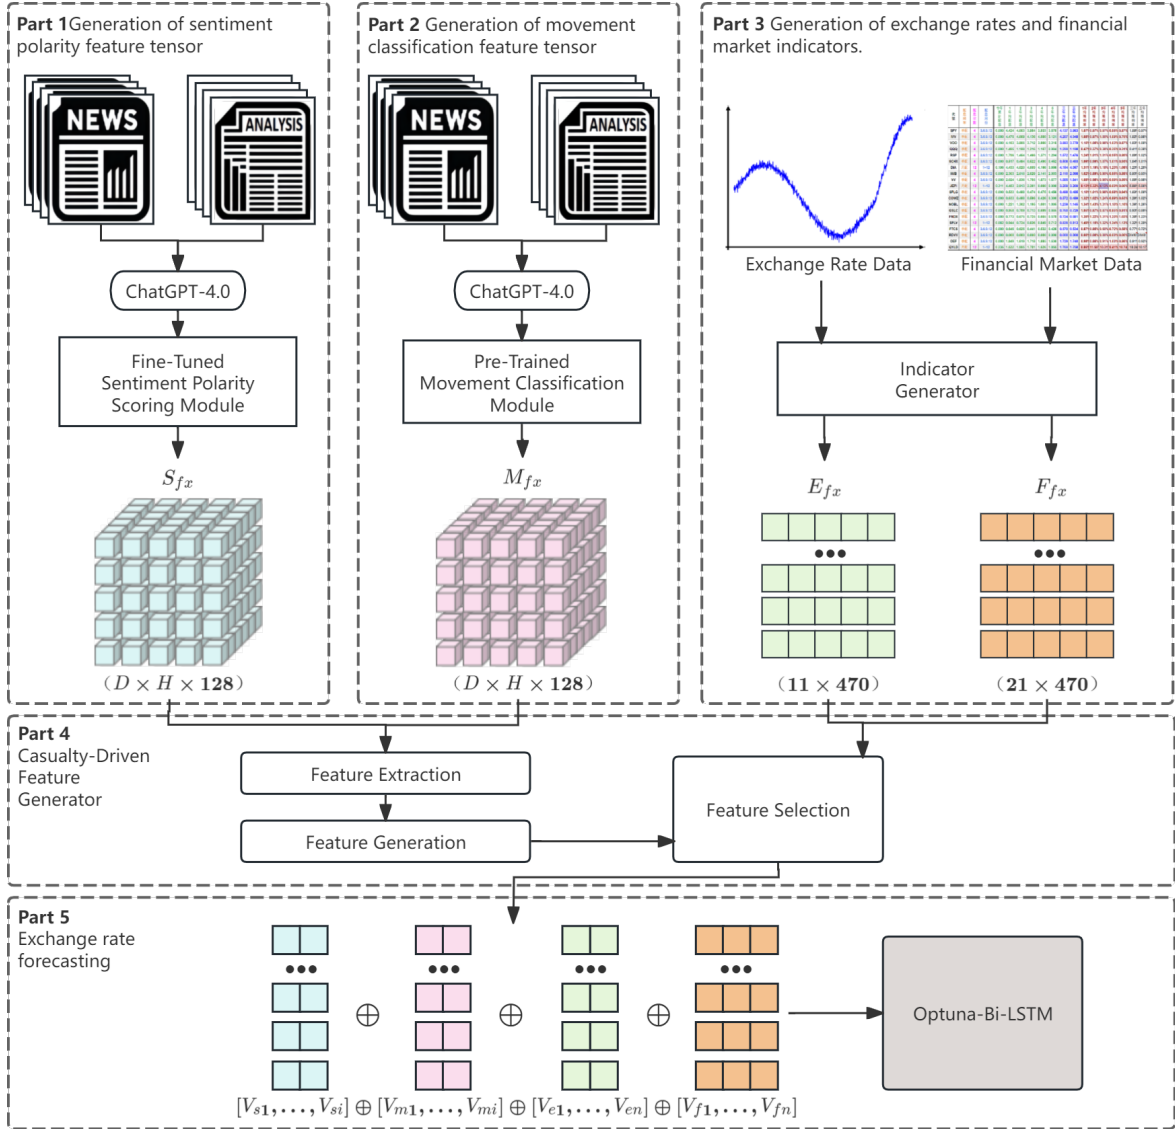


Figure 4: The framework of IUS to forecast EUR/USD exchange rate.

and dependencies among tokens. One of the primary advantages of RoBERTa-Large is its ability to learn robust and transferable language representations. By pre-training on extensive unlabelled text data, RoBERTa-Large develops a deep comprehension of language structure and semantics. As the input flows through each encoding layer, data move from the bottom to the top of the model, with representations becoming increasingly improved and enriched, thereby constructing hierarchical and contextualized embeddings of the text. Therefore, in our SPSM and MCM, we employ two RoBERTa-Large models as the embedding generators. Each text is processed by the RoBERTa-based tokenizer and then fed into the Encoder layer ($encoder(\cdot)$), generating the model's hidden states (T):

$$T = encoder([CLS], W_1, W_2, \dots, W_t, \dots, [PAD], [SEP]). \quad (2)$$

Here, $T \in \mathbb{R}^{1 \times 1024}$, W_i represents the i -th word in the text, and t denotes the length of the text. Positions not occupied by input words are filled with the [PAD] token to maintain a uniform sequence length of 512. The embeddings (e) extracted from the CLS to SEP positions of T as the text embeddings:

$$e = T([CLS] \text{ to } [SEP]), \quad (3)$$

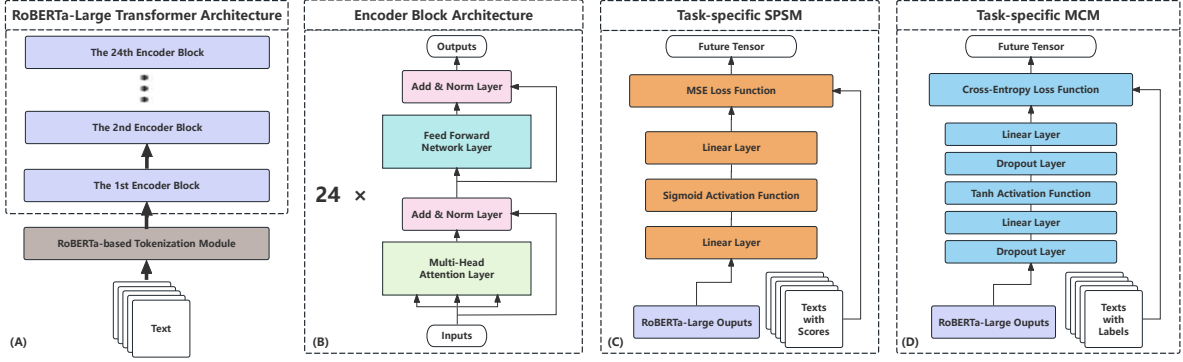


Figure 5: Components of the RoBERTa-Large Model, including (A) RoBERTa-Large Transformer Architecture, (B) The Encoder Block Architecture, (C) The Sentiment Polarity Scoring Module, and (D) The Movement Classification Module.

where e is a 1×1024 vector that passes through all transformer encoder blocks of RoBERTa-Large ($RL(\cdot)$), resulting in the final output hidden state F :

$$F = RL(e). \quad (4)$$

Here, $F \in \mathbb{R}^{1 \times 1024}$, the final hidden state will enter different modules to generate various feature tensors.

4.1.3 Sentiment Polarity Scoring Module

In SPSM, our RoBERTa-Large model utilizes weights from the Twitter-RoBERTa-Large-2022-154m model, which is fine-tuned by the CardiffNLP team on a large dataset containing 154 million tweets [44]. As shown in Figure 5(C), when using the RoBERTa-Large model for sentiment analysis, we expanded the model architecture with a module designed for regression. This module includes a Sigmoid activation function, two linear layers, and a mean squared error (MSE) loss function. The output of the RoBERTa-Large model initially passes through the first linear layer ($LL1_P(\cdot)$), which reduces the high-dimensional text representation from the final hidden state F of 1024 dimensions to a lower-dimensional space. A Sigmoid activation function ($Sig(\cdot)$) then compresses this output to a range between 0 and 1, and this output is further transformed by the second linear layer ($LL2_P(\cdot)$) to produce the final feature tensor S_{fx} :

$$S_{fx} = LL2_P(Sig(LL1_P(F))). \quad (5)$$

Here, $S_{fx} \in \mathbb{R}^{d \times h \times 128}$, d represents the total number of trading days and h indicates the maximum number of texts on all trading days. During training, the predicted sentiment score and the annotated polarity scores are fed into the MSE loss function [45–47]. The MSE loss function is defined as:

$$L_{MSE} = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2, \quad (6)$$

where y_i is the annotated polarity score for the i -th text in the dataset, and \hat{y}_i is the polarity score predicted by the model for the i -th text. n is the total number of texts over the training period. By minimizing the MSE loss function, the model learns the mapping relationship between text and sentiment scores. The backpropagation algorithm is used to compute gradients and update the parameters of RoBERTa-Large and SPSM, continually refining the predicted sentiment scores to approach the annotated scores. In addition, we discover that when using the Twitter-RoBERTa-Large-2022-154m model, convergence is faster, and performance on evaluation metrics is superior compared to the base model without fine-tuning, given the same number of training epochs.

4.1.4 Movement Classification Module

In MCM, we employ the RoBERTa-Large-Base model to uncover hidden patterns between textual information and the EUR/USD exchange rate movement on the following day [43]. As illustrated in

Figure 5(D), we augment the model architecture with a module specifically designed for classification, consisting of five layers. The final hidden state first goes through a dropout layer ($DP1(\cdot)$) which serves for regularization by randomly masking some neurons during training, thus reducing the risk of overfitting. After this, the output is processed by a linear layer ($LL1_M(\cdot)$), which linearly transforms the hidden state F from 1024 dimensions to a smaller dimension of 128. This reduction in dimensions decreases the number of parameters and enhances computational efficiency. We then obtain an intermediate representation F' :

$$F' = LL1_M(DP1(F)), \quad (7)$$

where F' refers to a vector of size 1×128 . A $Tanh$ activation function ($Tanh(\cdot)$) introduces non-linearity, enhancing the model's expressive power and mapping F' to the range of $[-1, 1]$. Another dropout layer ($DP2(\cdot)$) is added to further regularize the model. The final linear layer ($LL2_M(\cdot)$) maps the intermediate representations to feature tensor M_{fx} , which is defined as:

$$M_{fx} = LL2_M(DP2(Tanh(F'))), \quad (8)$$

where $M_{fx} \in \mathbb{R}^{d \times h \times 128}$, d represents the total number of trading days and h indicates the maximum number of texts on all transaction days.

In the model training, the Cross-Entropy (CE) loss function measures the difference between the model's predicted probability distribution and the true labels [48–50], which can be represented as:

$$L_{CE} = -\frac{1}{n} \sum_{i=1}^n [y_i \log(p_i) + (1 - y_i) \log(1 - p_i)], \quad (9)$$

where y_i is the true exchange rate movement label for the i -th text in the dataset, and p_i is the predicted probability of this movement by the model for the i -th text. n is the total number of texts over the training period. During the training process, the model learns the relationship between texts and exchange rate movements by minimizing the cross-entropy loss function, using the backpropagation algorithm to update the parameters of RoBERTa-Large-Base and MCM.

4.1.5 Experience Rule

We also collect a financial indicator dataset, which primarily includes data related to EUR/USD exchange rates and financial markets, sourced mainly from financial platforms such as investing.com and finance.yahoo.com, among others. We utilize a comprehensive approach by extensively collecting and analyzing relevant literature to identify the indicators that potentially predict fluctuations in the target exchange rate. The construction of this financial indicator system, displayed in Table 1, is based on the evaluation of the relationships between these indicators and the target exchange rate, considering their leading, lagging, and potential non-linear relationships.

All collected raw financial data are fed into an indicator generator, which aligns the data and fills in missing values, producing the final quantitative features. We use linear interpolation to fill in these gaps:

$$v_i = v_a + \frac{(v_b - v_a)(t_i - t_a)}{t_b - t_a}, \quad (10)$$

here, v_i represents the interpolated value at the specific time t_i . v_a and v_b are the known values at the time points t_a and t_b , respectively. These known values are used to estimate v_i , assuming a linear change between t_a and t_b .

4.2 Causality-Driven Feature Generator

We employ a Causality-Driven Feature Generator to extract text, exchange rate, and financial market features. Specifically, Figure 6 shows the stage of textual feature extraction, the feature tensor is fed into a feature extractor and produces various types of feature tensors. These produced tensors have the same dimensions as the original feature tensor, which are then processed through a task-specific linear layer, mapping the three-dimensional feature tensors to feature matrices. As for the stage of feature generation, the feature matrices are inputted into an average pooling layer to yield a diverse set of textual features. Subsequently, all textual features combined with other features undergo feature selection to obtain the final set of features inputted into the forecasting model.

Classification	No.	Indicator Name	Reference
Target Series	1	EUR/USD Exchange Rate	
US Exchange Rate (Top 5 Trading Partners)	2	USD/CAD Exchange Rate	[51–54]
	3	USD/MXN Exchange Rate	
	4	USD/CNY Exchange Rate	
	5	USD/JPY Exchange Rate	
	6	USD/KRW Exchange Rate	
EU Exchange Rate (Top 5 Trading Partners)	7	EUR/CNY Exchange Rate	[51–54]
	8	EUR/GBP Exchange Rate	
	9	EUR/CHF Exchange Rate	
	10	EUR/RUB Exchange Rate	
	11	EUR/TRY Exchange Rate	
Currency Index	12	US Dollar Index	[55–58]
	13	Euro Index	[55–58]
Currency Futures	14	US Dollar Futures-Jun	[59–61]
	15	Euro Futures-Jun	[62]
Commodities	16	Crude Oil WTI Futures	[63, 64]
	17	Natural Gas Futures	[65, 66]
	18	Gold Futures	[67, 68]
	19	Copper Futures	[60, 69]
	20	Corn Futures	[70–73]
	21	Soybeans Futures	[70–73]
Bond Yield	22	US 10-Year Bond Yield	[74–76]
	23	Eurozone 10-Year Bond Yield	[77, 78]
Interbank Offered Rate	24	SOFR - 1 month	[79–83]
	25	EURIBOR - 1 month	[79–81, 84]
US Stock Index	26	Dow Jones Industrial Average	[85–93]
	27	S&P 500	
EU Stock Index	28	Euro Stoxx 50	[85–93]
	29	STOXX 600	
US Stock Index Futures	30	Dow Jones Futures - Jun	[94–96]
EU Stock Index Futures	31	EURO STOXX 50 Futures - Jun	[94–96]
Chicago Board Options Exchange	32	VIX	[97–99]

Table 1: Comprehensive list of financial indicators used for analysis.

4.2.1 Causality-Driven Feature Generator

In terms of feature extraction, we employ two types of feature extractors. The first one is a classifier based on information sources, inspired by the research of Angeletos et al. [100] and Ke et al. [101]. We construct a classifier for news and analysis texts. The news texts encompass a broad spectrum of content related to exchange rates, including central bank monetary policies, economic data releases, geopolitical events, and other significant factors influencing exchange rate trends. In contrast, the analysis texts primarily focus on interpreting, forecasting, and providing recommendations concerning exchange rate movements, typically authored by professional analysts or traders, characterized by their expertise and forward-looking nature. This feature classifier, as depicted in Figure 7(A), generates news text classification matrix A and analysis text classification matrix B . An element of 1 in these matrices indicates that the corresponding vector in the feature tensor is selected by the classifier, and an element of 0 indicates it is not part of that category. The feature tensor is then multiplied by the corresponding matrices to yield the classified feature tensors.

The second extractor is a text-based topic cluster. We utilize an LDA model to cluster the feature tensor, as shown in Figure 7(B), assigning each text’s feature vector to a specific topic [102, 103]. Furthermore, the optimal number of topics are determined by locating the peak in the topic coherence curve. To cautiously assess the impact of the number of topics on our research, we visualize the outputs of the LDA topic model using the pyLDAvis tool [104, 105]. This tool displays the distance between

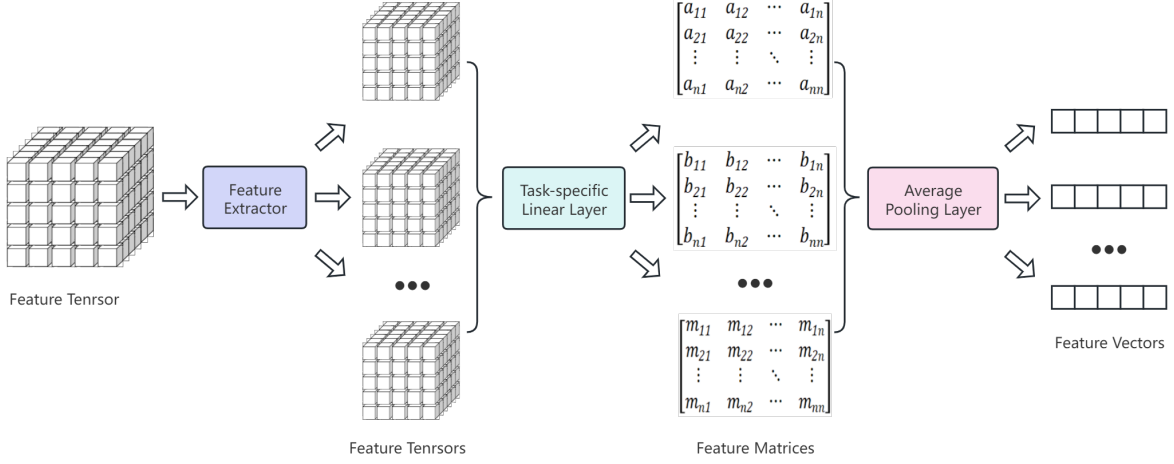


Figure 6: The overview of extracting textual features.

topics through multidimensional scaling and lists the top significant terms for each topic. To validate the stability of the extracted features, we also conduct dynamic topic modeling (DTM) to examine if the clustered features vary over time [16, 106]. During the feature extraction process, we obtain a series of topics, Topic 1, Topic 2, ..., Topic n , and the corresponding text cluster matrices A , B , ..., K . An element of 1 in these matrices indicates that the corresponding vector in the feature tensor is selected by the cluster; an element of 0 indicates it is not part of that category. The feature tensor is then multiplied by the corresponding matrices to yield the clustered feature tensors.

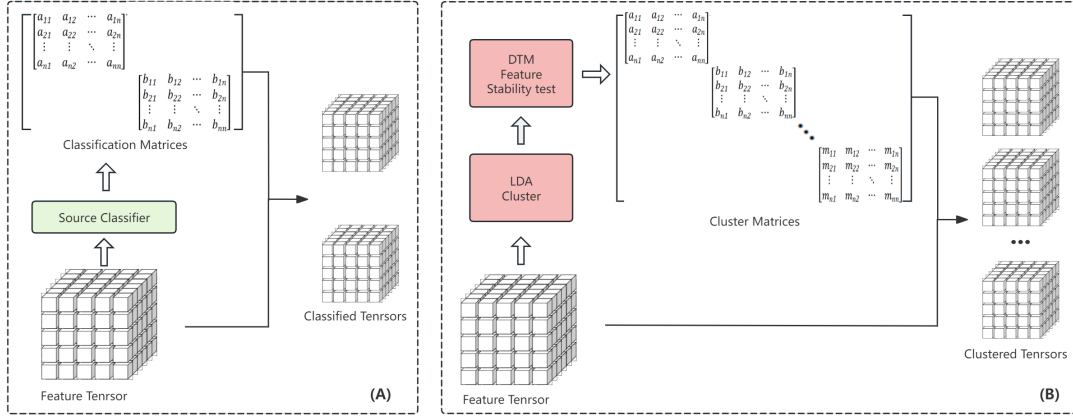


Figure 7: The feature extractors include (A) The source classifier, and (B) The LDA cluster and DTM feature stability test.

4.2.2 Feature Generation

When generating features, we need to handle two types of features: S_{fx} and M_{fx} . In terms of S_{fx} , after obtaining the sub-tensors from the feature extractor $FE(\cdot)$, we input these sub-tensors into a linear layer $LLS(\cdot)$ specifically designed for processing S_{fx} , mapping them to the corresponding sentiment feature sub-matrix, M_s , each element of which is compressed and normalized to range between $[-1, 1]$:

$$[M_{s1}, M_{s2}, \dots, M_{sn}] = LLS(FE(S_{fx})). \quad (11)$$

Here, $M_{si} \in \mathbb{R}^{d \times h}$ is the i -th feature sub-tensor obtained from S_{fx} through the feature extractor and subsequently compressed by a linear layer into the corresponding feature matrix, and n denotes the number of feature sub-tensors of the feature S_{fx} .

Each sentiment feature sub-matrix serves as the input for an average pooling layer, resulting in the corresponding feature vector V_s :

$$V_{si} = \frac{1}{h} \sum_{j=1}^h M_{sij}. \quad (12)$$

Here, M_{si} denotes the feature sub-matrix, h is the number of dimensions in M_{si} . M_{sij} represents the j -th column of M_{si} , and $V_{si} \in \mathbb{R}^{1 \times d}$ is the resultant vector obtained by averaging each dimension of M_{si} .

As for M_{fx} , after obtaining the sub-tensors from the feature extractor $FE(\cdot)$, we input these sub-tensors into a linear layer $LLM(\cdot)$ specifically designed for processing M_{fx} . This process maps each high-dimensional feature vector into a one-dimensional element, which corresponds to two categories in a binary classification task: one category for 'increase and no change', and another for 'decrease'. This mapping forms the corresponding movement feature sub-matrix, M_w :

$$[M_{w1}, M_{w2}, \dots, M_{wn}] = LLM(FE(M_{fx})). \quad (13)$$

Here, $M_{wi} \in \mathbb{R}^{d \times h}$ is the i -th feature sub-tensor obtained from M_{fx} through the feature extractor and subsequently compressed by a linear layer into the corresponding feature matrix, and n denotes the number of feature sub-tensors of the feature.

Each exchange movement feature sub-matrix serves as the input for an average pooling layer, resulting in the corresponding feature vector V_m :

$$V_{mi} = \frac{1}{k} \sum_{j=1}^k M_{wij}. \quad (14)$$

Here, M_{wi} denotes the feature sub-matrix, k is the number of dimensions in M_{wi} . M_{wij} represents the j -th column of M_{wi} , and $V_{mi} \in \mathbb{R}^{1 \times d}$ is the vector obtained by averaging each dimension of M_{wi} .

Overall, S_{fx} and M_{fx} are fed into a Causality-Driven Feature Generator and finally generates 8 textual features as shown in Figure 8.

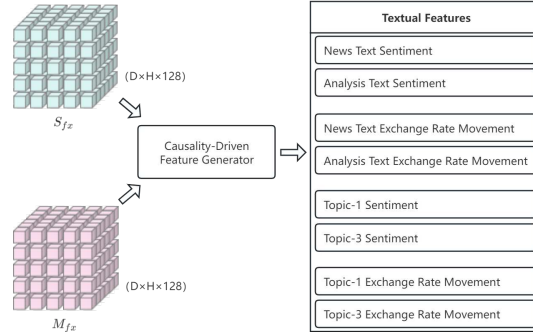


Figure 8: The final textual features generated by the Causality-Driven Feature Generator.

4.2.3 Feature Selection

We employ the VAR model to determine the optimal lag orders for all features within the final prediction feature set. The VAR model is an econometric model used to learn the dynamic relationships among multiple time-series variables, expressing each endogenous variable's current value as a linear combination of its own and all other endogenous variables' lagged values [107, 108]. The VAR(p) model can be represented as follows:

$$Y_t = c + A_1 Y_{t-1} + A_2 Y_{t-2} + \dots + A_p Y_{t-p} + \epsilon_t. \quad (15)$$

Here, Y_t is an $n \times 1$ vector of endogenous variables, c is an $n \times 1$ vector of constants, A_i are $n \times n$ coefficient matrices, and ϵ_t is an $n \times 1$ vector of error terms satisfying the white noise condition. To determine the optimal lag order p , we use the *Akaike Information Criterion (AIC)*:

$$AIC(p) = \ln(\det(\hat{\Sigma}_p)) + \frac{2pn^2}{T}. \quad (16)$$

Where, $\hat{\Sigma}_p$ is the estimated covariance matrix of the residuals for p lags, indicating the overall variance that the model fails to explain with smaller values being preferable. The natural logarithm of the determinant of this matrix, $\ln(\det(\hat{\Sigma}_p))$, quantifies the total unexplained variance by the model, serving as a measure of model fit. The penalty term, $\frac{2pn^2}{T}$, increases with the number of lags (p) and the number of endogenous variables (n), normalized by the sample size (T), to penalize model complexity and prevent overfitting, promoting a balance between fitting accuracy and model simplicity [109–111].

In order to select the optimal lag order, *VAR* models are estimated at different lag orders for each feature, ranging from 0 to 10 to encompass potential fluctuations. Corresponding *AIC* values are then calculated to determine the optimal lag for each feature. The lag order associated with the lowest *AIC* value is considered optimal, as a lower *AIC* value indicates a better balance between model fit and complexity[16, 17]. Table 2 displays the optimal lag periods for all features. After adjusting all

Feature Name	Lag	Feature Name	Lag	Feature Name	Lag
EUR/USD ER	2	Euro Futures-Jun	3	STOXX 600	1
USD/CAD ER	2	Crude Oil WTI Futures	1	EURO STOXX 50 Futures - Jun	1
USD/MXN ER	1	Natural Gas Futures	1	Dow Jones Futures - Jun	1
USD/CNY ER	1	Gold Futures	1	VIX	1
USD/JPY ER	1	Copper Futures	1	News Text Sentiment	4
USD/KRW ER	1	Corn Futures	1	Analysis Text Sentiment	1
EUR/CNY ER	1	Soybeans Futures	1	News Text ER Movement	4
EUR/GBP ER	1	US 10-Year Bond Yield	3	Analysis Text ER Movement	1
EUR/CHF ER	1	Eurozone 10-Year Bond Yield	2	Topic-1 Sentiment	4
EUR/RUB ER	1	SOFR - 1 Month	1	Topic-3 Sentiment	1
EUR/TRY ER	1	EURIBOR - 1 Month	1	Topic-1 ER Movement	5
US Dollar Index	3	Dow Jones Industrial Average	1	Topic-3 ER Movement	3
Euro Index	7	S&P 500	1		
US Dollar Futures-Jun	3	Euro Stoxx 50	1		

Note: ER = Exchange Rate

Table 2: The optimal lag periods for all features

features to their optimal lag orders, we employ the Recursive Feature Elimination (RFE) method with random forest regression to rank the importance of all indicators. RFE is a backward elimination algorithm that recursively removes the least important features until the desired number of features is determined [112–114]. Random forest regressor, an ensemble of decision trees each trained on a bootstrap sample of the training data, determines the importance of each feature by averaging the decrease in impurity this feature causes across all trees in the forest. This decrease in impurity, also known as Gini importance or mean decrease impurity, enhances prediction accuracy and minimizes overfitting. The formula for calculating the importance of a feature f , denoted as $I(f)$, is defined as:

$$I(f) = \sum_{t=1}^T \sum_{n=1}^{N_t} \Delta i(n, f). \quad (17)$$

Where $I(f)$ is the importance of feature f , T is the number of trees in the forest, N_t is the number of nodes in tree t , and $\Delta i(n, f)$ is the impurity decrease caused by feature f at node n . The RFE algorithm initiates by training a random forest regressor on the initial set of features, calculating feature importances, removing the least important features, and repeating these steps until the desired number of features is retained. The RFE process can be expressed as:

$$F_i = F_{i-1} \setminus \{f_j\}, \quad (18)$$

where F_i is the feature set at iteration i , F_{i-1} is the feature set from the previous iteration, and f_j is the least important feature removed at iteration i . This recursive removal of less important features and evaluation of model performance ultimately helps identify an optimal feature subset containing the most important features.

We use $FS(\cdot)$ to denote feature selection and all features obtained through the Causality-Driven Feature Generator are concatenated to form the input feature set I :

$$I = FS([V_{s1}, \dots, V_{sn}] \oplus [V_{m1}, \dots, V_{mn}] \oplus [V_{e1}, \dots, V_{en}] \oplus [V_{f1}, \dots, V_{fn}]). \quad (19)$$

Where, $I \in \mathbb{R}^{112 \times 470}$ represents the feature set, including 111 features over 470 trading days, comprising 80 financial features and 31 textual features. The symbol \oplus denotes the concatenation operation. Finally, before inputting the selected features into the predictive model, we perform min-max normalization to scale all features to the range of $[0, 1]$. The normalization is applied as follows:

$$I'_i = \frac{I_i - \min(I_i)}{\max(I_i) - \min(I_i)}, \quad (20)$$

where I'_i represents the i -th row of the matrix I , corresponding to the i -th feature across all dates. I'_i is the normalized i -th row of the matrix, $\min(I_i)$ and $\max(I_i)$ are the minimum and maximum values found in the i -th row of I . After normalization, rows I'_1 to I'_{112} are concatenated to form the final feature matrix $I' \in \mathbb{R}^{112 \times 470}$, which is then used as the input for the forecasting model.

4.3 Optuna-Bi-LSTM

4.3.1 Bi-LSTM

This study employs a Bi-LSTM model to analyze financial features and forecast EUR/USD foreign exchange rates. The Bi-LSTM is an efficient sequential learning model that enhances performance by integrating past and future feature information, demonstrating strong capabilities in time series forecasting [115–117]. Figure 9(A) illustrates the structure of our prediction model, where I' is inputted into two Bi-LSTM layers, $Bi-LSTM1(\cdot)$ and $Bi-LSTM2(\cdot)$, to identify temporal patterns for predicting EUR/USD exchange rate movements. The key to the model lies in the Bi-LSTM layers, which capture both forward and backward dependencies in the input time sequence. Figure 9(B) displays the structure of each Bi-LSTM layer, composed of a forward LSTM and a backward LSTM, processing the sequential information in both directions respectively. After each Bi-LSTM layer, a dropout layer, $DP1(\cdot)$ and $DP2(\cdot)$, is added to utilize regularization to reduce the risk of overfitting. The final hidden state from I' at the last time step, h_{t+1} , can be expressed as:

$$h_{t+1} = DP2(Bi-LSTM2(DP1(Bi-LSTM1(I')))), \quad (21)$$

where $h_{t+1} \in \mathbb{C}^{1 \times w}$, $c1$ is the number of feature combinations in the current model, and w is the sliding window length. h_{t+1} carries all the features necessary for predicting the EUR/USD exchange rate movement on day $t + 1$. Finally, h_{t+1} passes through a fully connected layer ($FC(\cdot)$) and an output layer ($Output(\cdot)$), generating the predicted EUR/USD exchange rate \hat{y}_{t+1} for day $t + 1$:

$$\hat{y}_{t+1} = Output(FC(h_{t+1})). \quad (22)$$

We utilize the MSE loss function to measure the discrepancy between predicted values and actual values, updating the model parameters accordingly:

$$L_{MSE} = \frac{1}{n} \sum_{i=1}^n (y_{t+i} - \hat{y}_{t+i})^2, \quad (23)$$

where y_{t+1} represents the actual price range of the EUR/USD exchange rate on day $t+1$. By minimizing the loss function, the model learns the relationship between the features within the rolling window and the subsequent trading day's EUR/USD exchange rate.

4.3.2 Optuna optimization framework

Optuna is an automatic hyperparameter optimization framework that efficiently searches for the optimal set of hyperparameters [118, 119]. Its core principle is to generate trials, each testing a different combination of hyperparameters, while using the results of previous trials to guide the sampling of parameters in subsequent trials. This sampling, combined with the ability to prune underperforming trials, allows Optuna to quickly converge on the best hyperparameter configuration.

The Figure 10 illustrates the process of hyperparameter optimization using Optuna. The process begins by defining an objective function for constructing, training, and evaluating the model. Subsequently, an Optuna Study object is created to manage the optimization process, and the optimization loop then begins, generating new hyperparameters based on the results of previous trials. Each trial

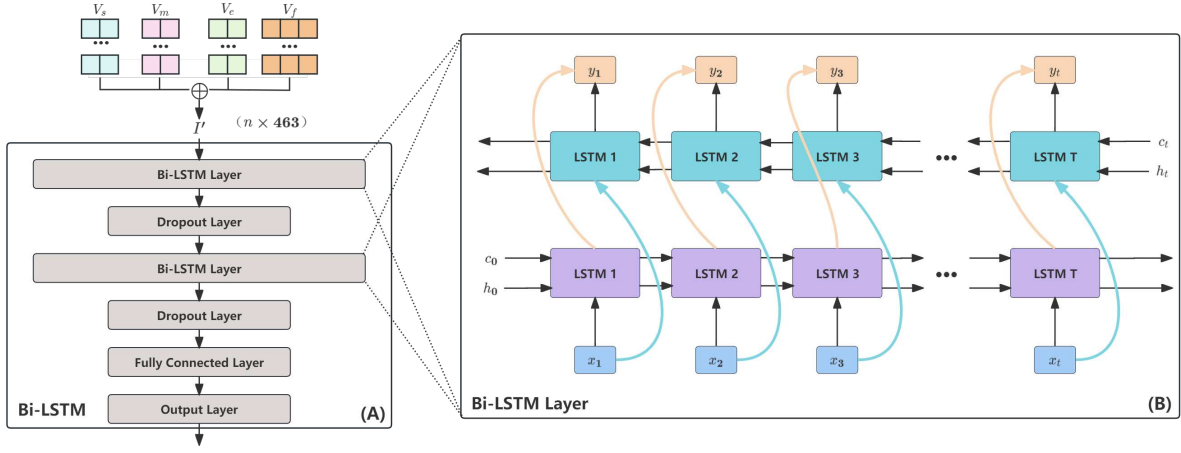


Figure 9: The Bi-LSTM model includes: (A) Bi-LSTM architecture and (B) the workflow of Bi-LSTM Layer.

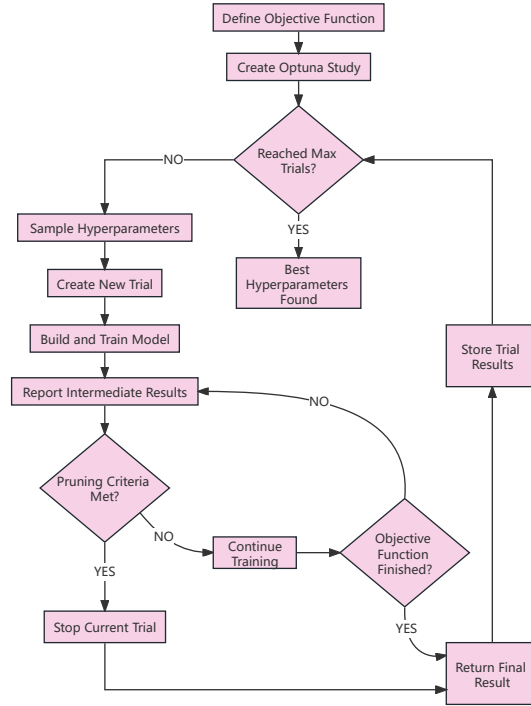


Figure 10: The workflow of the Optuna hyperparameter optimization process.

is represented by a Trial object and passed to the objective function. Inside the objective function, the suggest method of the Trial object is utilized to obtain hyperparameter values. The model is built and trained using these hyperparameters, and intermediate results are reported by the Trial's report method. If the pruning criteria are met, the trial is terminated; otherwise, training continues until the objective function is completed. The final evaluation result is returned, and the trial results are stored. This process is repeated until the maximum number of trials is reached. Finally, the best hyperparameter combination is obtained from the Study object.

The parameter space in this study consists of six dimensions, shown in Table 3. Among these, hidden layer size, fully connected layer size, batch size, and window size are discrete variables, but window size also includes continuous ranges within specific numbers. Dropout rate and learning rate are continuous variables, with the learning rate being sampled on a logarithmic scale. Additionally, the size of the fully connected layer is constrained by the size of the hidden layer and must be less than or equal to the hidden layer size.

Parameter	Range
Hidden layer size	[8, 16, 32, 64, 128]
Fully connected layer size	[8, 16, 32, 64]
Dropout rate	[0.0, 0.5]
Learning rate	[1e-5, 1e-2]
Batch size	[8, 16, 32, 48, 64, 80, 96, 112, 128]
Window size (continuous)	[1, 24]
Window size (discrete)	[30, 40, 50, 60]

Table 3: The hyperparameter space.

5 Experiment

5.1 Dataset

During the training phase of two LLMs, the datasets span from February 6, 2017, to April 4, 2022. Subsequently, the two fine-tuned LLMs are utilized to score text sentiment polarity and classify exchange rate movements in the predictive dataset ranging from April 4, 2022, to January 19, 2024. As for the exchange rate prediction models, their training set consist of the first 315 transaction days from April 4, 2022, to January 19, 2024, and their prediction set comprises the following 155 transaction days, as shown in Figure 11. To prevent information leakage, we implement strict data processing methods, ensuring that the predictive data remains isolated from the training process.




Figure 11: Training and testing periods of the two LLMs and the exchange rate forecasting models.

5.2 Evaluation metrics

We evaluate the models' forecasting performances over the test period using two different criteria: the mean absolute error (MAE) and root mean squared error (RMSE):

$$\text{MAE} = \frac{1}{n} \sum_{i=1}^n |\hat{y}_i - y_i|, \quad (24)$$

$$\text{RMSE} = \sqrt{\frac{1}{n} \sum_{i=1}^n (\hat{y}_i - y_i)^2}. \quad (25)$$

Here, \hat{y}_i represents the model's forecasted exchange rate for day i , y_i represents the actual exchange rate for day i , and n is the total number of days in the testset [120].

To quantify the additional explanatory power of the textual features, we compute a metric named the percentage improvement (PI). This metric represents the percentage enhancement in predictive performance achieved by combining the two feature sets compared to using the quantitative financial features alone [16]. The two PIs are defined as follows:

$$PI_{\text{MAE}} = \left(\frac{\text{MAE of Quantitative Features} - \text{MAE of Combined Features}}{\text{MAE of Quantitative Features}} \right) \times 100\% \quad (26)$$

$$PI_{\text{RMSE}} = \left(\frac{\text{RMSE of Quantitative Features} - \text{RMSE of Combined Features}}{\text{RMSE of Quantitative Features}} \right) \times 100\% \quad (27)$$

6 Experiment Results

6.1 Main Results

6.1.1 Time series forecasting

We compare the performance of the Optuna-Bi-LSTM model with other benchmark models. The prediction results of different models for the EUR/USD exchange rate are shown in Table 4. It can be observed that our proposed method consistently outperforms other models in terms of both MAE and RMSE metrics, achieving an improvement of at least 10.69% in MAE and 9.56% in RMSE compared with the best other model.

Moreover, as illustrated in Figure 12, the prediction curve obtained by Optuna-Bi-LSTM aligns more closely with the raw data curve and exhibits a higher degree of trend similarity. This demonstrates the superior predictive performance of our proposed model.

Series	Category	Model	MAE	RMSE	Rank
Multivariate	Proposed method	Optuna-Bi-LSTM	0.003746	0.004982	1
	Deep learning methods	Bi-LSTM	0.004511	0.005814	4
		LSTM	0.004768	0.006212	5
		GRU	0.004958	0.006457	6
	Machine learning methods	Random Forest	0.005471	0.007672	7
		XGBoost	0.006809	0.009012	8
Univariate	Statistical methods	GARCH	0.004282	0.005695	2
		ARIMA	0.004456	0.005718	3

Table 4: Comparison of Forecasting Models based on MAE and RMSE



Figure 12: Training and testing periods of the two LLMs and the exchange rate forecasting models.

The results indicate that by combining structured and unstructured data for exchange rate prediction and feeding the features generated by the Causality-Driven Feature Generator into the model,

our proposed method achieves significant improvements over the benchmarks.

6.1.2 DM Test

To evaluate and compare the predictive performance of the eight models on the time series, we conducted the Diebold-Mariano (DM) test. The purpose of the DM test is to determine which models are statistically significantly superior to others in terms of prediction accuracy [1, 121, 122]. We first calculated the forecast errors for each model and constructed error difference series from them for pairwise comparisons between models. Using this approach, we performed a total of 28 DM tests, covering all possible combinations of model pairs. The test results are summarized in Table 5, where the LSTM model demonstrates the best performance, while the GRU model exhibits the poorest predictive ability.

Rank	Model
1	LSTM
2	Optuna-Bi-LSTM
3	Bi-LSTM
4	Random Forest
5	XGBoost
6	GARCH
7	ARIMA
8	GRU

Table 5: Ranking of forecasting models Based on DM Test Results.

6.1.3 Window Size Analysis

In this study, we analyze the impact of different window sizes on the prediction accuracy of the models. The tested window sizes range from 1 to 24, with additional extended sizes of 30, 40, 50, and 60, to assess their influence on model performance. This analysis helps determine the optimal window size that maximizes the accuracy of time series predictions. The results of this analysis are presented in Figure 13. By observing the changes in MAE and RMSE across different models, each window size is evaluated. When the window size is 3, the models exhibit the best performance. As the window size increases, the model performance generally deteriorates.




Figure 13: Performance metrics varying with window size.

6.2 Ablation Experiment

6.2.1 Textual Feature Breakdown

We investigate the relative predictive power of textual and other features by comparing the performance of using (1) only textual features (31 features), (2) only exchange rate and financial market features

(the top 12 important features selected by the RFE method), and all selected features (43 features), which is the combination of (1) and (2). The results show that in the scenario where only textual features are inputted, the model’s performance is far inferior to the scenarios using only financial features and all selected features. This indicates that textual features alone may not provide sufficient predictive information for exchange rate forecasting. Moreover, when using Optuna-Bi-LSTM, the combination of textual and financial features achieves the best performance, outperforming the case of using only financial features.

	(1)	(2)	Combination: (1) + (2)	PI from (2) to (1) + (2)
MAE				
Optuna-Bi-LSTM	0.025406	0.004270	0.003746	12.27%
Bi-LSTM	0.020712	0.004736	0.004511	4.75%
Random Forest	0.031247	0.005747	0.005471	4.80%
RMSE				
Optuna-Bi-LSTM	0.029987	0.005502	0.004982	9.45%
Bi-LSTM	0.024674	0.005816	0.005814	0.03%
Random Forest	0.032789	0.007813	0.007672	1.8%

Table 6: Forecasting performances with different features.

The last column of Table 6 presents the percentage improvement results defined by Formulas 26 and 27. Whether in terms of MAE or RMSE, the improvement of combined features over financial features alone in the Optuna-Bi-LSTM model significantly outperforms that of the Bi-LSTM model. The overall prediction results suggest that textual and financial features are complementary to each other and are suitable for exchange rate prediction. When unstructured data and structured data are combined in a predictive model, significant accuracy improvements can be achieved.

6.2.2 Breakdown Study

We conduct several ablation experiments to analyze the effectiveness of each textual feature generated by the IUS framework. In each experiment, we include the exchange rate and financial market features (the top 12 important features selected by the RFE method) as a fixed set of predictors. The textual feature sets we evaluate include: (A) features obtained through SSM and Source Classification, (B) features obtained through SSM and LDA Cluster, (C) features obtained through MCM and Source Classification, and (D) features obtained through MCM and LDA Cluster. We use (A + B + C + D) to represent the complete set of predictive features, including both textual and financial features, and (0) to represent no textual features, only the fixed set of financial features.

In Table 7, we calculate the weighted rank as follows: $\text{Weighted Score} = 0.5 \times \text{MAE Rank} + 0.5 \times \text{RMSE Rank}$. When the ranks are tied, we prioritize the MAE rank for ascending order. As shown in the table, using only (A) and (A)+(B) results in lower MAE and RMSE compared to our proposed method.




Figure 14: Prediction curves using textual features (A) and (A)+(B) for predicting exchange rates.

Textual features	MAE	Rank (MAE)	RMSE	Rank (RMSE)	Weighted Rank
(0)	0.004270	16	0.005502	16	16
(A)	0.003495	1	0.004716	1	1
(B)	0.003975	12	0.005164	10	10
(C)	0.003945	10	0.005115	6	8
(D)	0.004095	15	0.005342	15	15
(A)+(B)	0.003634	2	0.004879	2	2
(A)+(C)	0.003843	6	0.005015	4	5
(A)+(D)	0.003773	4	0.005062	5	4
(B)+(C)	0.003889	8	0.005144	9	9
(B)+(D)	0.004021	13	0.005270	13	14
(C)+(D)	0.003861	7	0.005121	7	7
(A)+(B)+(C)	0.004024	14	0.005207	11	13
(A)+(B)+(D)	0.003809	5	0.005131	8	6
(A)+(C)+(D)	0.003917	9	0.005287	14	11
(B)+(C)+(D)	0.003969	11	0.005268	12	12
(A)+(B)+(C)+(D)	0.003746	3	0.004982	3	3

Table 7: Performance metrics for different textual feature combinations

Figure 14 presents the prediction curves when using only (A) and (A)+(B) to predict the exchange rates. We observe that the two prediction curves are closer to the original data curve compared to our proposed method. Therefore, when using more features for prediction, the additional textual features (C)+(D) may introduce slight noise.

6.2.3 Recursive Factor Importance Feature Selection

We select different numbers of exchange rate and financial related features based on their importance, along with all textual features, to input into the predictive model. This approach allows us to observe the impact of the number of features on model predictive performance.




Figure 15: The workflow of the Optuna hyperparameter optimization process.

The Figure 15 illustrates the curve of MAE and RMSE as the number of features varies. As the financial features increase from 1 to 12, both curves reach their minimum values. Subsequently, both curves rise steeply. This indicates that as the number of features increases, noise is introduced into the predictive model, thereby gradually reducing its predictive capability. Overall, the top 12 important features selected by the RFE method, combined with 31 textual features, constitute a robust feature set for predicting exchange rates.

7 Conclusions and Future Work

To accurately forecast the EUR/USD exchange rate, we introduced the IUS framework and an Optuna-Bi-LSTM model. This framework integrates multi-source information, including news and analytical texts, other relevant exchange rates, and financial market indicators. We employ two large language models for sentiment polarity scoring and exchange rate movement classification, which are then combined with other quantitative indicators input into a Causality-Driven Feature Generator. All generated features are fed into the predictive model for exchange rate forecasting.

Experimental results demonstrate that compared to the strongest benchmarks, our method achieved the highest MAE and RMSE, improving by at least 10.69% and 9.56%, respectively. In terms of data fusion, by combining unstructured and structured data, the model is able to enhance prediction accuracy beyond what is possible with structured data alone. Furthermore, using the top 12 important features selected by the RFE method, combined with 31 textual features proves to be more effective compared to analyzing all textual features, as it more directly corresponds to the actual exchange rate response to market conditions. In summary, the proposed method achieves better performance in exchange rate forecasting and provides a comprehensive approach by integrating multi-source data for enhanced prediction accuracy.

References

- [1] Barbara Rossi. Exchange rate predictability. *Journal of economic literature*, 51(4):1063–1119, 2013.
- [2] Yin-Wong Cheung, Menzie D Chinn, Antonio Garcia Pascual, and Yi Zhang. Exchange rate prediction redux: New models, new data, new currencies. *Journal of International Money and Finance*, 95:332–362, 2019.
- [3] Omer Berat Sezer, Mehmet Ugur Gudelek, and Ahmet Murat Ozbayoglu. Financial time series forecasting with deep learning: A systematic literature review: 2005–2019. *Applied soft computing*, 90:106181, 2020.
- [4] Zexin Hu, Yiqi Zhao, and Matloob Khushi. A survey of forex and stock price prediction using deep learning. *Applied System Innovation*, 4(1):9, 2021.
- [5] Hamed Naderi Semiromi, Stefan Lessmann, and Wiebke Peters. News will tell: Forecasting foreign exchange rates based on news story events in the economy calendar. *The North American Journal of Economics and Finance*, 52:101181, 2020.
- [6] Harmanjeet Singh, Ashima Rani, Aarti Chugh, and Supreet Singh. Utilising sentiment analysis to predict market movements in currency exchange rates. In *2024 3rd International Conference for Innovation in Technology (INOCON)*, pages 1–7. IEEE, 2024.
- [7] Anushka Tadphale, Haripriya Saraswat, Omkar Sonawane, and PR Deshmukh. Impact of news sentiment on foreign exchange rate prediction. In *2023 3rd International Conference on Intelligent Technologies (CONIT)*, pages 1–8. IEEE, 2023.
- [8] Saba Haider, Mian Sajid Nazir, Alfredo Jiménez, and Muhammad Ali Jibran Qamar. Commodity prices and exchange rates: evidence from commodity-dependent developed and emerging economies. *International Journal of Emerging Markets*, 18(1):241–271, 2023.
- [9] MS Aziz Sarkar and UAM Ehsan Ali. Eur/usd exchange rate prediction using machine learning. *Int. J. Math. Sci. Comput*, 8(1):44–48, 2022.
- [10] Qingsong Ruan, Jiarui Zhang, and Dayong Lv. Forecasting exchange rate volatility: is economic policy uncertainty better? *Applied Economics*, 56(13):1526–1544, 2024.
- [11] Edmure Windsor and Wei Cao. Improving exchange rate forecasting via a new deep multimodal fusion model. *Applied Intelligence*, 52(14):16701–16717, 2022.
- [12] Afees A Salisu, Juncal Cuñado, Kazeem Isah, and Rangan Gupta. Oil price and exchange rate behaviour of the brics. *Emerging Markets Finance and Trade*, 57(7):2042–2051, 2021.

- [13] Davood Pirayesh Neghab, Mucahit Cevik, MIM Wahab, and Ayse Basar. Explaining exchange rate forecasts with macroeconomic fundamentals using interpretive machine learning. *Computational Economics*, pages 1–43, 2024.
- [14] Takumi Ito and Fumiko Takeda. Do sentiment indices always improve the prediction accuracy of exchange rates? *Journal of Forecasting*, 41(4):840–852, 2022.
- [15] Walid Ben Omrane, Robert Welch, and Xinyao Zhou. The dynamic effect of macroeconomic news on the euro/us dollar exchange rate. *Journal of Forecasting*, 39(1):84–103, 2020.
- [16] Xuerong Li, Wei Shang, and Shouyang Wang. Text-based crude oil price forecasting: A deep learning approach. *International Journal of Forecasting*, 35(4):1548–1560, 2019.
- [17] Yun Bai, Xixi Li, Hao Yu, and Suling Jia. Crude oil price forecasting incorporating news text. *International Journal of Forecasting*, 38(1):367–383, 2022.
- [18] T Swathi, N Kasiviswanath, and A Ananda Rao. An optimal deep learning-based lstm for stock price prediction using twitter sentiment analysis. *Applied Intelligence*, 52(12):13675–13688, 2022.
- [19] Eleni Kalamara, Arthur Turrell, Chris Redl, George Kapetanios, and Sujit Kapadia. Making text count: economic forecasting using newspaper text. *Journal of Applied Econometrics*, 37(5): 896–919, 2022.
- [20] Samreen Naeem, Wali Khan Mashwani, Aqib Ali, M Irfan Uddin, Marwan Mahmoud, Farrukh Jamal, and Christophe Chesneau. Machine learning-based usd/pkr exchange rate forecasting using sentiment analysis of twitter data. *Computers, Materials & Continua*, 67(3):3451–3461, 2021.
- [21] Lv Xueling, Xiong Xiong, and Shen Yucong. Exchange rate market trend prediction based on sentiment analysis. *Computers and Electrical Engineering*, 111:108901, 2023.
- [22] Kazım Berk Küçüklerli and Veysel Ulusoy. Sentiment-driven exchange rate forecasting: Integrating twitter analysis with economic indicators. *Journal of Applied Finance & Banking*, 14(3): 75–96, 2024.
- [23] Jiatong Li, Jiawen Yin, and Rui Zhang. Analysis and forecast of usd/eur exchange rate based on arima and garch models. In *International Conference on Economic Management and Green Development*, pages 566–575. Springer, 2023.
- [24] Ruimin Zhang. Rmb/usd exchange rate forecasting and analysis by ses and arima model. *Highlights in Business, Economics and Management*, 24:112–121, 2024.
- [25] Emilio Colombo and Matteo Pelagatti. Statistical learning and exchange rate forecasting. *International Journal of Forecasting*, 36(4):1260–1289, 2020.
- [26] Jonathan Felix Pfahler. Exchange rate forecasting with advanced machine learning methods. *Journal of Risk and Financial Management*, 15(1):2, 2021.
- [27] Bui Thanh Khoa and Tran Trong Huynh. Predicting exchange rate under uirp framework with support vector regression. *assessment*, 12:13, 2022.
- [28] Shaolong Sun, Shouyang Wang, and Yunjie Wei. A new ensemble deep learning approach for exchange rates forecasting and trading. *Advanced Engineering Informatics*, 46:101160, 2020.
- [29] Siyuan Liu, Qiqian Huang, Mingchen Li, and Yunjie Wei. A new lasso-bilstm-based ensemble learning approach for exchange rate forecasting. *Engineering Applications of Artificial Intelligence*, 127:107305, 2024.
- [30] Md Saiful Islam and Emam Hossain. Foreign exchange currency rate prediction using a gru-lstm hybrid network. *Soft Computing Letters*, 3:100009, 2021.
- [31] Alexander Jakob Dautel, Wolfgang Karl Härdle, Stefan Lessmann, and Hsin-Vonn Seow. Forex exchange rate forecasting using deep recurrent neural networks. *Digital Finance*, 2:69–96, 2020.

- [32] Renzhuo Wan, Shuping Mei, Jun Wang, Min Liu, and Fan Yang. Multivariate temporal convolutional network: A deep neural networks approach for multivariate time series forecasting. *Electronics*, 8(8):876, 2019.
- [33] Ailing Zeng, Muxi Chen, Lei Zhang, and Qiang Xu. Are transformers effective for time series forecasting? In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pages 11121–11128, 2023.
- [34] Yuanhao Xu, Caihong Hu, Qiang Wu, Shengqi Jian, Zhichao Li, Youqian Chen, Guodong Zhang, Zhaoxi Zhang, and Shuli Wang. Research on particle swarm optimization in lstm neural networks for rainfall-runoff simulation. *Journal of hydrology*, 608:127553, 2022.
- [35] Khader M Hamdia, Xiaoying Zhuang, and Timon Rabczuk. An efficient optimization approach for designing machine learning models based on genetic algorithm. *Neural Computing and Applications*, 33(6):1923–1933, 2021.
- [36] A Helen Victoria and Ganesh Maragatham. Automatic tuning of hyperparameters using bayesian optimization. *Evolving Systems*, 12(1):217–223, 2021.
- [37] Xingping Dong, Jianbing Shen, Wenguan Wang, Ling Shao, Haibin Ling, and Fatih Porikli. Dynamical hyperparameter optimization via deep reinforcement learning in tracking. *IEEE transactions on pattern analysis and machine intelligence*, 43(5):1515–1529, 2019.
- [38] Juan Pablo García Amboage, Eric Wulff, Maria Girone, and Tomás F Pena. Model performance prediction for hyperparameter optimization of deep learning models using high performance computing and quantum annealing. In *EPJ Web of Conferences*, volume 295, page 12005. EDP Sciences, 2024.
- [39] Andrzej Brodzicki, Michał Piekarski, and Joanna Jaworek-Korjakowska. The whale optimization algorithm approach for deep neural networks. *Sensors*, 21(23):8003, 2021.
- [40] Souvik Das and Rohini K Srihari. Compos mentis at semeval2024 task6: A multi-faceted role-based large language model ensemble to detect hallucination. In *Proceedings of the 18th International Workshop on Semantic Evaluation (SemEval-2024)*, pages 1449–1454, 2024.
- [41] Giuseppe Gallipoli, Simone Papicchio, Lorenzo Vaiani, Luca Cagliero, Arianna Miola, and Daniele Borghi. Keyword-based annotation of visually-rich document content for trend and risk analysis using large language models. In *Proceedings of the Joint Workshop of the 7th Financial Technology and Natural Language Processing, the 5th Knowledge Discovery from Unstructured Data in Financial Services, and the 4th Workshop on Economics and Natural Language Processing@ LREC-COLING 2024*, pages 130–136, 2024.
- [42] Hang Li, Tianlong Xu, Jiliang Tang, and Qingsong Wen. Knowledge tagging system on math questions via llms with flexible demonstration retriever. *arXiv preprint arXiv:2406.13885*, 2024.
- [43] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized BERT pretraining approach. *CoRR*, abs/1907.11692, 2019. URL <http://arxiv.org/abs/1907.11692>.
- [44] Daniel Loureiro, Kiamehr Rezaee, Talayeh Riahi, Francesco Barbieri, Leonardo Neves, Luis Espinosa Anke, and Jose Camacho-Collados. Tweet insights: A visualization platform to extract temporal insights from twitter. *arXiv preprint arXiv:2308.02142*, 2023.
- [45] Zhanghan Ke, Daoye Wang, Qiong Yan, Jimmy Ren, and Rynson WH Lau. Dual student: Breaking the limits of the teacher in semi-supervised learning. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6728–6736, 2019.
- [46] Jun Qi, Jun Du, Sabato Marco Siniscalchi, Xiaoli Ma, and Chin-Hui Lee. On mean absolute error for deep neural network based vector-to-vector regression. *IEEE Signal Processing Letters*, 27:1485–1489, 2020.

- [47] Vidya Muthukumar, Adhyayan Narang, Vignesh Subramanian, Mikhail Belkin, Daniel Hsu, and Anant Sahai. Classification vs regression in overparameterized regimes: Does the loss function matter? *Journal of Machine Learning Research*, 22(222):1–69, 2021.
- [48] Like Hui and Mikhail Belkin. Evaluation of neural architectures trained with square loss vs cross-entropy in classification tasks. *arXiv preprint arXiv:2006.07322*, 2020.
- [49] Anqi Mao, Mehryar Mohri, and Yutao Zhong. Cross-entropy loss functions: Theoretical analysis and applications. In *International conference on Machine learning*, pages 23803–23828. PMLR, 2023.
- [50] Zhaoqi Leng, Mingxing Tan, Chenxi Liu, Ekin Dogus Cubuk, Xiaojie Shi, Shuyang Cheng, and Dragomir Anguelov. Polyloss: A polynomial expansion perspective of classification loss functions. *arXiv preprint arXiv:2204.12511*, 2022.
- [51] Leixin Yang, Haiying Wang, Changgui Gu, and Huijie Yang. Weighted signed networks reveal interactions between us foreign exchange rates. *Entropy*, 26(2):161, 2024.
- [52] Jozef Baruník, Evžen Kočenda, and Lukáš Vácha. Asymmetric volatility connectedness on the forex market. *Journal of International Money and Finance*, 77:39–56, 2017.
- [53] Matthew Greenwood-Nimmo, Viet Hoang Nguyen, and Yongcheol Shin. Measuring the connectedness of the global economy. *International Journal of Forecasting*, 37(2):899–919, 2021.
- [54] Erdem Kilic. Contagion effects of us dollar and chinese yuan in forward and spot foreign exchange markets. *Economic Modelling*, 62:51–67, 2017.
- [55] Akner Alkan, Ali Fuat Alkaya, and Peter Schüller. Currency exchange rate forecasting with social media sentiment analysis. In *Intelligent and Fuzzy Techniques for Emerging Conditions and Digital Transformation: Proceedings of the INFUS 2021 Conference, held August 24-26, 2021. Volume 2*, pages 490–497. Springer, 2022.
- [56] Liang Ding. Conditional correlation between exchange rates and stock prices. *The Quarterly Review of Economics and Finance*, 80:452–463, 2021.
- [57] Orawan Chantarakasemchit, Siranee Nuchitprasitchai, and Yuenyong Nilsiam. Forex rates prediction on eur/usd with simple moving average technique and financial factors. In *2020 17th International Conference on Electrical Engineering/Electronics, Computer, Telecommunications and Information Technology (ECTI-CON)*, pages 771–774. IEEE, 2020.
- [58] Ikhlās Gurrib and Elgilani Elshareif. Optimizing the performance of the fractal adaptive moving average strategy: The case of eur/usd. *International Journal of Economics and Finance*, 8(2):171–179, 2016.
- [59] Aaron Tornell and Chunming Yuan. Speculation and hedging in the currency futures markets: Are they informative to the spot exchange rates. *Journal of Futures Markets*, 32(2):122–151, 2012.
- [60] Domenico Ferraro, Kenneth Rogoff, and Barbara Rossi. Can oil prices forecast exchange rates? an empirical analysis of the relationship between commodity prices and exchange rates. *Journal of International Money and Finance*, 54:116–141, 2015.
- [61] Yu-chin Chen and Kenneth Rogoff. Commodity currencies. *Journal of international Economics*, 60(1):133–160, 2003.
- [62] Richard K Lyons. A simultaneous trade model of the foreign exchange hot potato. *Journal of international Economics*, 42(3-4):275–298, 1997.
- [63] Paul Cashin, Luis F Céspedes, and Ratna Sahay. Commodity currencies and the real exchange rate. *Journal of Development Economics*, 75(1):239–268, 2004.
- [64] Ali Jadidzadeh and Apostolos Serletis. How does the us natural gas market react to demand and supply shocks in the crude oil market? *Energy Economics*, 63:66–74, 2017.

- [65] Radhamés A Lizardo and André V Mollick. Oil price fluctuations and us dollar exchange rates. *Energy economics*, 32(2):399–408, 2010.
- [66] Syed Abul Basher, Alfred A Haug, and Perry Sadorsky. The impact of oil shocks on exchange rates: A markov-switching approach. *Energy Economics*, 54:11–23, 2016.
- [67] Kuntara Pukthuanthong and Richard Roll. Gold and the dollar (and the euro, pound, and yen). *Journal of Banking & Finance*, 35(8):2070–2083, 2011.
- [68] Larry A Sjaastad. The price of gold and the exchange rates: Once again. *Resources Policy*, 33(2):118–124, 2008.
- [69] Hui Jun Zhang, Jean-Marie Dufour, and John W Galbraith. Exchange rates and commodity prices: Measuring causality at multiple horizons. *Journal of Empirical Finance*, 36:100–120, 2016.
- [70] Saban Nazlioglu and Ugur Soytas. Oil price, agricultural commodity prices, and the dollar: A panel cointegration and causality analysis. *Energy Economics*, 34(4):1098–1104, 2012.
- [71] Lateef Olawale Akanni. Returns and volatility spillover between food prices and exchange rate in nigeria. *Journal of Agribusiness in Developing and Emerging Economies*, 10(3):307–325, 2020.
- [72] Saban Nazlioglu, Cumhur Erdem, and Ugur Soytas. Volatility spillover between oil and agricultural commodity markets. *Energy Economics*, 36:658–665, 2013.
- [73] Anthony N Rezitis. The relationship between agricultural commodity prices, crude oil prices and us dollar exchange rates: A panel var approach and causality analysis. *International Review of Applied Economics*, 29(3):403–434, 2015.
- [74] Charles Engel and Steve Pak Yeung Wu. Forecasting the us dollar in the 21st century. *Journal of International Economics*, 141:103715, 2023.
- [75] Menzie D Chinn and Guy Meredith. Monetary policy and long-horizon uncovered interest parity. *IMF staff papers*, 51(3):409–430, 2004.
- [76] Natalja Lace, Irena Mačerinskienė, and Andrius Balčiūnas. Determining the eur/usd exchange rate with us and german government bond yields in the post-crisis period. *Intellectual Economics*, 9(2):150–155, 2015.
- [77] Antonio Afonso and Mina Kazemi. Euro area sovereign yields and the power of unconventional monetary policy. *Finance a úvěr-Czech Journal of Economics and Finance*, 68(2):100–119, 2018.
- [78] Martina Cecioni. Ecb monetary policy and the euro exchange rate. *Bank of Italy Temi di Discussione (Working Paper) No*, 1172, 2018.
- [79] Lena Tonzer. Cross-border interbank networks, banking risk and contagion. *Journal of Financial Stability*, 18:19–32, 2015.
- [80] Victoria Ivashina, David S Scharfstein, and Jeremy C Stein. Dollar funding and the lending behavior of global banks. *The Quarterly Journal of Economics*, 130(3):1241–1281, 2015.
- [81] Marcos Dal Bianco, Maximo Camacho, and Gabriel Perez Quiros. Short-run forecasting of the euro-dollar exchange rate with economic fundamentals. *Journal of International Money and Finance*, 31(2):377–396, 2012.
- [82] Darrell Duffie and Jeremy C Stein. Reforming libor and other financial market benchmarks. *Journal of Economic Perspectives*, 29(2):191–212, 2015.
- [83] Wenxin Du, Alexander Tepper, and Adrien Verdelhan. Deviations from covered interest rate parity. *The Journal of Finance*, 73(3):915–957, 2018.
- [84] Jens Eisenschmidt, Danielle Kedan, and Robin Darius Tietz. Measuring fragmentation in the euro area unsecured overnight interbank money market. *Economic Bulletin Articles*, 5, 2018.

- [85] Ming-Shiun Pan, Robert Chi-Wing Fok, and Y Angela Liu. Dynamic linkages between exchange rates and stock prices: Evidence from east asian markets. *International Review of Economics & Finance*, 16(4):503–520, 2007.
- [86] Yun-Cheng Tsai, Jun-Hao Chen, and Jun-Jie Wang. Predict forex trend via convolutional neural networks. *Journal of Intelligent Systems*, 29(1):941–958, 2019.
- [87] Trilok Nath Pandey, Alok Kumar Jagadev, Satchidananda Dehuri, and Sung-Bae Cho. A review and empirical analysis of neural networks based exchange rate prediction. *Intelligent Decision Technologies*, 12(4):423–439, 2018.
- [88] Chien-Chung Nieh and Cheng-Few Lee. Dynamic relationship between stock prices and exchange rates for g-7 countries. *The Quarterly Review of Economics and Finance*, 41(4):477–490, 2001.
- [89] Chien-Hsiu Lin. The comovement between exchange rates and stock prices in the asian emerging markets. *International Review of Economics & Finance*, 22(1):161–172, 2012.
- [90] Kate Phylaktis and Fabiola Ravazzolo. Stock prices and exchange rate dynamics. *Journal of international Money and Finance*, 24(7):1031–1053, 2005.
- [91] A Can Inci and Bong Soo Lee. Dynamic relations between stock returns and exchange rate changes. *European Financial Management*, 20(1):71–106, 2014.
- [92] Tomoe Moore and Ping Wang. Dynamic linkage between real exchange rates and stock prices: Evidence from developed and emerging asian markets. *International Review of Economics & Finance*, 29:1–11, 2014.
- [93] I-Chun Tsai. The relationship between stock price index and exchange rate in asian markets: A quantile regression approach. *Journal of International Financial Markets, Institutions and Money*, 22(3):609–621, 2012.
- [94] Kenneth A Tah and Geoffrey Ngene. Dynamic linkages between us and eurodollar interest rates: new evidence from causality in quantiles. *Journal of Economics and Finance*, 45:200–210, 2021.
- [95] Gaurav Agrawal, Aniruddh Kumar Srivastav, and Ankita Srivastava. A study of exchange rates movement and stock market volatility. *International Journal of business and management*, 5(12):62, 2010.
- [96] Elena Andreou, Maria Matsi, and Andreas Savvides. Stock and foreign exchange market linkages in emerging economies. *Journal of International Financial Markets, Institutions and Money*, 27: 248–268, 2013.
- [97] Markus K Brunnermeier, Stefan Nagel, and Lasse H Pedersen. Carry trades and currency crashes. *NBER macroeconomics annual*, 23(1):313–348, 2008.
- [98] John Cairns, Corrinne Ho, and Robert N McCauley. Exchange rates and global volatility: implications for asia-pacific currencies. *BIS Quarterly Review*, March, 2007.
- [99] Zhiyuan Pan, Yudong Wang, Li Liu, and Qing Wang. Improving volatility prediction and option valuation using vix information: A volatility spillover garch model. *Journal of Futures Markets*, 39(6):744–776, 2019.
- [100] George-Marios Angeletos, Fabrice Collard, and Harris Dellas. Business-cycle anatomy. *American Economic Review*, 110(10):3030–3070, 2020.
- [101] Zheng Tracy Ke, Bryan T Kelly, and Dacheng Xiu. Predicting returns with text data. Technical report, National Bureau of Economic Research, 2019.
- [102] David M Blei, Andrew Y Ng, and Michael I Jordan. Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022, 2003.
- [103] Saqib Aziz, Michael Dowling, Helmi Hammami, and Anke Piepenbrink. Machine learning in finance: A topic modeling approach. *European Financial Management*, 28(3):744–770, 2022.

- [104] Arş Gör Abdullah Kürşat MERTER, Yavuz Selim BALCIOĞLU, Arş Gör Sedat ÇEREZ, and ÖZER Gökhan. Evaluation of annual reports of companies with sentiment analysis: An application in bist100 index. In *8th International New York Academic Research Congress On Humanities and Social Sciences*, pages 197–204, 2023.
- [105] Sulong Zhou, Pengyu Kan, Qunying Huang, and Janet Silbernagel. A guided latent dirichlet allocation approach to investigate real-time latent topics of twitter data during hurricane laura. *Journal of Information Science*, 49(2):465–479, 2023.
- [106] David M Blei and John D Lafferty. Dynamic topic models. In *Proceedings of the 23rd international conference on Machine learning*, pages 113–120, 2006.
- [107] Ventzislav Ivanov and Lutz Kilian. A practitioner’s guide to lag order selection for var impulse response analysis. *Studies in Nonlinear Dynamics & Econometrics*, 9(1), 2005.
- [108] James H Stock and Mark W Watson. Vector autoregressions. *Journal of Economic perspectives*, 15(4):101–115, 2001.
- [109] Yahaya Musa, Ibrahim Ahmad, and Bilkisu Maijamaa. Application of vector autoregression (var) on modelling and forecasting average monthly rainfall and temperature. *Lloyd Business Review*, pages 1–23, 2024.
- [110] Semanto Mondal, Prakash Srinivasan, and Rajib Chandra Ghosh. Multivariate time series forecasting to forecast weight dynamics. In *2023 Second International Conference on Advances in Computational Intelligence and Communication (ICACIC)*, pages 1–6. IEEE, 2023.
- [111] Frank Schorfheide, Dongho Song, and Amir Yaron. Identifying long-run risks: A bayesian mixed-frequency approach. *Econometrica*, 86(2):617–654, 2018.
- [112] Burcu F Darst, Kristen C Malecki, and Corinne D Engelman. Using recursive feature elimination in random forest to account for correlated variables in high dimensional data. *BMC genetics*, 19: 1–6, 2018.
- [113] Qifeng Zhou, Hao Zhou, and Tao Li. Cost-sensitive feature selection using random forest: Selecting low-cost subsets of informative features. *Knowledge-based systems*, 95:1–11, 2016.
- [114] Baptiste Gregorutti, Bertrand Michel, and Philippe Saint-Pierre. Correlation and variable importance in random forests. *Statistics and Computing*, 27:659–678, 2017.
- [115] Mo Yang and Jing Wang. Adaptability of financial time series prediction based on bilstm. *Procedia Computer Science*, 199:18–25, 2022.
- [116] Yu Ma, Rui Mao, Qika Lin, Peng Wu, and Erik Cambria. Multi-source aggregated classification for stock price movement prediction. *Information Fusion*, 91:515–528, 2023.
- [117] Sima Siami-Namini, Neda Tavakoli, and Akbar Siami Namin. The performance of lstm and bilstm in forecasting time series. In *2019 IEEE International conference on big data (Big Data)*, pages 3285–3292. IEEE, 2019.
- [118] Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 2623–2631, 2019.
- [119] Yipeng Zhou, Ze Dong, and Xiongguan Bao. A ship trajectory prediction method based on an optuna–bilstm model. *Applied Sciences*, 14(9):3719, 2024.
- [120] Bolin Lei, Boyu Zhang, and Yuping Song. Volatility forecasting for high-frequency financial data based on web search index and deep learning model. *Mathematics*, 9(4):320, 2021.
- [121] Francis X Diebold and Robert S Mariano. Comparing predictive accuracy. *Journal of Business & economic statistics*, 20(1):134–144, 2002.
- [122] Tanya Molodtsova and David H Papell. Out-of-sample exchange rate predictability with taylor rule fundamentals. *Journal of international economics*, 77(2):167–180, 2009.