

THE POTENTIAL OF HYBRID LSTM-GENERATIVE AI ECO-MODEL IN FORECASTING FINANCIAL AND ECONOMIC INDICATORS

Andrii Ivashchenko¹, Tetiana Ivashchenko^{2*}

Received 25.03.2025.

| Send to review 06.04.2025.

| Accepted 19.07.2025.

Original Article



¹ Oracle America, Inc., USA

² Academy of Labor, Social Relations and Tourism, Faculty of Economics, Social management and Tourism, Department of Economics and Management, Ukraine

*Corresponding Author:

Tetiana Ivashchenko

Email: tanita.ivash@gmail.com

JEL Classification:

C45, C53, E37, F37, F65, G17

Doi: 10.61432/CPNE0301171i

UDK: 519.816:005.5

Article Classification:

Original scientific paper

ABSTRACT

This study presents the development and evaluation of The Hybrid LSTM-Generative AI ECO-Model for forecasting financial and economic indicators such as EUR/USD exchange rate, utilizing a combination of Long Short-Term Memory (LSTM) networks and generative AI models (GPT-2 and Llama-3.2-1B). The primary objective was to achieve high prediction accuracy while minimizing computational resource consumption and ensuring ease of use of the model on various devices. The model was trained and tested on historical financial and economic data, including exchange rates, macroeconomic indicators, commodity prices, and sentiment analysis of financial news.

Our findings indicate that traditional LSTM models outperform generative AI models in time-series forecasting tasks due to their ability to capture temporal dependencies. However, integrating generative AI for dataset refinement and model optimization significantly improved forecasting performance. The hybrid ECO-model, leveraging generative AI-driven parameter selection and sentiment analysis, demonstrated superior accuracy for long-term predictions. The most influential parameters included historical exchange rate trends, gold and oil prices and news sentiment.

By implementing an ECO-approach that optimizes dataset size, minimizes training iterations, and employs a lightweight model architecture, our study highlights a path toward efficient and sustainable financial forecasting. Future research directions include enhancing anomaly detection mechanisms, incorporating additional weak predictors, and refining the role of generative AI in hybrid time-series forecasting models.

Keywords: Financial forecasting, economic indicators prediction, LSTM-Generative AI model, currency exchange rate, hybrid AI modelling, sentiment analysis, ECO-model

1. INTRODUCTION

Financial and economic forecasting plays a critical role in shaping policies, business strategies, and investment decisions. Accurately predicting key indicators such as GDP growth, inflation, stock market trends, or currency exchange rates remains a complex task due to the dynamic, non-linear, and interconnected nature of financial systems. Historically, various statistical and econometric models have been employed to forecast economic trends. These include linear regression, ARIMA (Auto-Regressive Integrated Moving Average), and VAR (Vector Auto-Regression) models, which rely on historical data patterns to predict future outcomes. While these models have been instrumental, their inherent assumptions, such as linearity and stationarity, often limit their applicability in capturing complex economic behaviours.

In recent years, advances in machine learning and artificial intelligence have revolutionized financial forecasting. Technologies such as support vector machines (SVM), random forests, and neural networks have enabled models to capture non-linear relationships and adapt to high-dimensional datasets. These methods have outperformed traditional approaches in various contexts but often face limitations in their ability to process sequential and time-dependent data effectively. The further advancements in artificial intelligence (AI) have significantly revolutionized this field, enabling the development of predictive tools like Long Short-Term Memory (LSTM) networks and Generative AI models that outperform many methods that have been used before. The benefits of these technologies are manifold. LSTMs provide accurate short- and medium-term predictions by effectively handling the temporal dependencies in financial data, while Generative AI enhances robustness by modelling complex interactions and creating synthetic data for underrepresented scenarios. These technologies not only improve prediction accuracy but also offer tools to navigate uncertainties, optimize strategies, and mitigate risks.

Despite their individual strengths, LSTMs and generative AI models have limitations when used in isolation. LSTMs, for instance, may struggle with highly sparse or noisy datasets, while generative AI models may lack precise temporal alignment in predictions. To overcome these challenges, scientists are increasingly adopting hybrid approaches that integrate LSTMs with other forecasting tools, such as generative AI models, statistical methods, or machine learning techniques. Hybrid models leverage the strengths of complementary techniques to address the limitations of individual methods. For instance, combining LSTMs with generative AI allows for the simulation of complex data distributions, enhancing robustness and enabling probabilistic forecasting. Integrating statistical models can improve interpretability, while machine learning techniques contribute to feature selection and dimensionality reduction. These synergies lead to improved predictive accuracy, adaptability, and resilience in uncertain financial environments.

This paper examines the potential of a Hybrid LSTM-Generative AI ECO-Model to address the limitations of existing methods, proposing an innovative approach to achieve more accurate and scalable predictions. By exploring this hybrid framework, we aim to advance the field of financial and economic forecasting and contribute to the development of solutions that are both robust and practical in real-world applications.

1.1. LITERATURE REVIEW

Many scientists conducting research in this area are trying to find the optimal mechanism for constructing forecasts of macroeconomic indicators and develop universal models capable of producing the most accurate results. The use of machine learning in macroeconomic forecasting has been gaining traction, with particular focus on Long Short-Term Memory (LSTM) models and Generative AI models. This section reviews key studies comparing these approaches, providing insights into their relative strengths and weaknesses for predicting macroeconomic indicators such as GDP, exchange rate, inflation, and stock market trends. It is commonly known that LSTM model has the ability to capture long-term dependencies. The memory units in the LSTM can transfer and update the information between different time steps, which makes the neural network more suitable for time-series analysis and modeling (Aumeboonsuke, 2025). That is why many analysts consider LSTM-based models as a basis for making forecasts. Among them Hall (Hall, 2018), who underscores the adaptability of LSTM models in macroeconomic forecasting, highlighting their ability to capture long-term dependencies and nonlinearities in time series data. However, he also notes that LSTM models can be computationally expensive and may struggle when dealing with large datasets. Medeiros et al. (Medeiros et al., 2021) similarly highlight LSTM's superiority in capturing complex interactions between macroeconomic variables in data-rich environments, particularly in inflation forecasting. Kamolthip (Kamolthip, 2021) supports this view, finding that LSTM models handle mixed-frequency data effectively, which is a common challenge in macroeconomic data due to the varying frequencies of economic indicators. Despite LSTM's strong performance, Kamolthip acknowledges that its effectiveness diminishes in cases of severe frequency mismatches. The utility of LSTM models in forecasting volatile financial data is echoed by Shahi et al. (Shahi et al., 2020), who compared LSTM and Gated Recurrent Unit (GRU) models in stock price forecasting. Their findings suggest

that LSTM outperforms GRU in handling non-linear data patterns, which is relevant to macroeconomic forecasting given the similarly unpredictable nature of economic indicators like inflation and GDP. Nakamura came to the same conclusion when she compared the results of using the LSTM model and traditional methods in the process of forecasting inflation (Nakamura, 2005). His study demonstrates that LSTM can capture temporal dynamics better than other approaches. Comparing classical statistical algorithms to LSTM and GRU models Chang (Chang et al., 2024) claims that the latter can identify and retain important data sequences, enabling more accurate predictions. But his experimental results show that the GRU model outperforms the LSTM model in terms of prediction accuracy and training time across multiple metrics such as RMSE and MAE.

While the effectiveness of LSTM is widely recognized, several authors highlight its limitations. For instance, Kamolthip (Kamolthip, 2021) and Jahn (Jahn, 2018) note that LSTM models often struggle when faced with highly volatile or multi-frequency datasets, a common scenario in macroeconomic forecasting. This thought is supported by Lim and Zohren (Lim & Zohren, 2021), who point out that deep neural networks typically require time series to be discretized at regular intervals, making it difficult to forecast datasets where observations can be missing or arrive at random intervals. This is where generative AI models, particularly those using GANs, offer potential advantages.

Economists can reap significant productivity gains by taking advantage of generative AI to automate micro- and later macro-tasks. Large language models (LLMs) such as ChatGPT can assist economists by describing dozens of use cases in six areas: ideation and feedback, writing, background research, data analysis, coding, and mathematical derivations. That's why recent studies have begun to explore the potential of generative AI models in economic development (Korinek, 2023). But particular attention is paid to the analysis of the capabilities of Large Language models in macroeconomic forecasting. Thus, Staffini et al. (Staffini, 2022) argue that generative models can outperform LSTM in generating future scenarios, making them more suitable for long-term macroeconomic predictions. This position is supported by Liu and Motani (Liu & Motani, 2022) who suggest a generative model, GenF, that uses synthetic data to improve long-range forecasts. Their work highlights how generative models can outperform LSTM models in long-term forecasting by reducing bias and variance, particularly in scenarios that require predictions over extended time horizons. In turn, Yoshimatsu (Yoshimatsu et al., 2024) points out that similar to large language models (like ChatGPT) time series transformers become experts at picking up on subtle patterns of relationships between the values in these series with demonstrated success across a variety of domains. Author considers four popular foundation time series transformer models: Chronos, TimesFM, Moirai, TimeGPT and their effective use in identifying complex economic patterns through large-scale time-series data. The research results of Che et al. (Che et al., 2024) show that the conditional generative adversarial networks (cGAN) model can effectively capture the complexity of financial market data, and the deviation between the prediction results and the actual market performance is minimal, showing a high degree of accuracy. Lee also points out the transformative impact of finance-specific large language models (LLMs), the innovative use of generative adversarial networks (GANs) in synthetic financial data generation, but at the same time authors warn that there is a pressing necessity of a new regulatory framework to govern the use of GAI in the finance sector (Lee et al., 2024).

However, concerns remain about the reliability and interpretability of generative models as well. Thus, Hassani and Silva argue that while generative AI models show promise in surpassing traditional methods such as ARIMA, they still face challenges related to model interpretability and the "black-box" nature of their predictions (Hassani & Silva, 2024). This sentiment is echoed by Brynjolfsson and Li, who highlight that while generative AI models hold promise for improving financial forecasting, their lack of transparency may limit their adoption in high-stakes macroeconomic contexts (Brynjolfsson & Li, 2024). Emphasizing on the ability of LLM to help us better understand human beliefs and navigate possible patterns of irrational expectations, Bybee draws attention to the fact that LLM Macroeconomic Expectations exhibits the underreaction typically found in SPF consensus forecasts and its return expectations are extrapolative, disconnected from objective measures of expected returns, and negatively correlated with future realized returns (Bybee, 2023). Additionally, financial professionals should always keep in mind that generative AI has some current limitations that have

not yet been overcome like bias amplification, reliability of output including those regarding forecast indicators, security of information and sovereignty of data (Schaefer et al., 2023).

The question of which approach is better - LSTM or generative AI for macroeconomic forecasting remains open, since there is no clear solution that would satisfy everyone yet. In order to improve the performance and prediction accuracy of the models, scientists and engineers have been actively working on developing hybrid approaches, such as: adding and changing parameters in existing models, using a chain of ANNs, and combining the capabilities of LSTM and GAI. For instance, Nguyen et al. (Nguyen et al., 2021) suggest a hybrid approach, integrating LSTM with normalizing flows to capture non-linear temporal dependencies in macroeconomic data more effectively than either model could on its own. Smyl (Smyl, 2020) also developed a hybrid model combining LSTM with exponential smoothing, which proves effective in managing non-linear patterns in macroeconomic time series. Hamiane et al. (Hamiane et al., 2024) found out that hybrid models combining LSTM and ARIMA outperform standalone models in predicting GDP, especially when non-linear trends are prominent. A more complex approach suggests Atif (Atif, 2025) by combining the strengths of linear autoregressive integrated moving average (ARIMA) models with the memory-retaining capabilities of long short-term memory (LSTM) networks and temporal convolutional networks (TCN), and creating hybrid architectures that capture diverse patterns in GDP time series. Sarveswararao (Sarveswararao & Ravi, 2021), in turn, by using LSTM enhanced the Lower Upper Bound Estimation (LUBE) method in order to generate prediction intervals for quantifying uncertainty in macroeconomic variables and proved that modified LUBE can be more effectively applied to generate prediction intervals for financial and non-financial time series. Sagheer and Kotb further contribute to this discussion by demonstrating that LSTM models, when pre-trained in an unsupervised manner, can enhance accuracy in predicting multivariate economic indicators (Sagheer & Kotb, 2019). In a study of forecasting directional movement of Forex data focusing on macroeconomic indicators influencing currency values Yildirim (Yildirim et al., 2021) has proved that macroeconomic LSTM model when combined with technical indicator LSTM model into a one hybrid model, can be quite successful and effectively applied to financial forecasting. At the same time, according to Zhao, the introduction of the unbiased and the improved grey model which is weighted and averaged into the LSTM model can make a better scientific prediction of macroeconomic indicators like GDP and CPI (Zhao, 2024).

An example of the successful application of the combined LSTM-GAN model can be the work of Ngo (Ngo, 2021), who using adversarial training has tuned an LSTM prediction model and developed a system that predicted closing prices of stocks based on the previous timesteps' prices and movement as a whole. The approach by Kim et al. (Kim et al., 2024) that implies the use of AI-based generative news sentiment analysis for the improvement of LSTM and VAR models in real estate price prediction has proven to be quite effective as well. Another promising solution called the scheme of generative LSTM (GLSTM) has been suggested by Liu et al. (Liu, 2019). It consists of a generative model, to generate synthetic data for the next few time steps, and a predictive model, to directly make long-range predictions based on observed and generated data (Lemaallem & Saadi, 2025). This approach is aimed at achieving a significant increase in the accuracy of long-term forecasting of indicators with high variability.

The literature reveals a divide between the strengths of LSTM and generative AI models in macroeconomic forecasting. LSTM models are praised for their ability to capture long-term dependencies and handle non-linearities in data, making them suitable for short to medium-term forecasts. However, for long-term forecasting, generative models, particularly those that utilize synthetic data generation, offer a distinct advantage by reducing forecast bias and improving accuracy over extended horizons. Based on the literature review the future of macroeconomic forecasting may lie in hybrid approaches, combining the strengths of both LSTM and generative AI models to enhance predictive power and overcome the limitations of each approach individually.

While many authors in the process of developing LSTM-GAI combined models focus their attention more on one aspect of using the advantages of generative models to improve LSTM. In our opinion, this limits the predictive capabilities of hybrid models and leads to less accurate results when applying it. Our approach is distinguished by greater involvement of generative models in the forecasting

process, but at the same time we consider it important that our model remains economical (lightweight), as simple as possible, in terms of resource consumption and easiness of use. That's why we called it Hybrid LSTM-Generative AI Prediction ECO-Model.

1.2. METHODOLOGY

Building on the insights gained from the literature review, we aimed to investigate, implement, and evaluate a model designed to forecast financial and economic indicators. For the financial parameter, we selected the EUR/USD exchange rate due to its significance as a key financial indicator, characterized by frequent fluctuations compared to other economic metrics such as GDP, inflation, or unemployment rates. Furthermore, the EUR/USD exchange rate is publicly accessible and offers a reliable historical dataset, making it particularly well-suited for predictive modeling. Other financial parameters can similarly be predicted using our model, provided that datasets of equivalent quality are prepared.

Our objective was to develop an Eco-model that enables accurate forecasting of the EUR/USD exchange rate using lightweight models and limited computational resources. Specifically, we sought to achieve this with a standard workstation (Intel(R) Core (TM) i7-1355U 1.70 GHz, 32 GB RAM) without specialized GPU modules. Additionally, we aimed to minimize training time and reduce overall carbon dioxide emissions associated with the computational process.

In this study, we utilized a combination of experiment-based methods, modeling, training techniques and optimization procedures to forecast the EUR/USD exchange rate.

- *Experiment Method:* We conducted an experimental comparison between traditional (LSTM) and generative AI models (GPT-2 and Llama-3.2-1B) for time-series forecasting of the EUR/USD exchange rate. The models were evaluated on their ability to predict exchange rates over multiple time horizons (7, 30, 180, and 365 days), allowing us to compare the performance of these distinct model types.
- *Modeling Method:* We employed a structured modeling methodology to develop the architecture of the Hybrid LSTM-Generative AI Eco-Model. This approach involved systematic design and integration of Long Short-Term Memory (LSTM) networks with Generative AI components to enhance the model's predictive accuracy and adaptability in economic forecasting.
- *Synthesis and Analysis:* Through a thorough literature review, we synthesized existing research on the use of LSTM networks and generative AI models for time-series prediction. This analysis formed the basis for hypothesizing the potential advantages of generative models in forecasting financial data, motivating our exploration of Llama and GPT-2 in this context.
- *Quantitative Analysis for Data Processing:* We used quantitative evaluation metrics, namely Root Mean Squared Error (RMSE) and Mean Absolute Error (MAE), to assess the models' prediction accuracy. These metrics provided a clear comparison between the performance of LSTM and generative AI models, with lower RMSE and MAE values indicating better predictive accuracy.

1.2.1. DATASET FORMATION

The quality of data preparation plays a pivotal role in the effectiveness of prediction model training. The accuracy of the predicted model is directly influenced by both the quantity and the quality of the data used. For the EUR/USD rate prediction, we have compiled extensive datasets of historical data beginning from January 1, 2010.

We utilized **various parameters** to construct the training datasets, which include: historical EUR/USD Exchange Rates; Financial News; Inflation Rates; GDP Growth Rates; Commodity Prices (Gold and Oil); Unemployment Rates; Economic News; Government Stability Index; Monetary Policy Statements; Interest Rates and Bond Yields; Trade Balance (EU-US Imports/Exports).

For the collection of historical EUR/USD exchange rates, as well as prices for oil and gold, we utilized the API provided by Yahoo Finance. This service allowed us to efficiently gather large-scale time-series data, essential for model training and analysis. Macroeconomic indicators, such as GDP growth, inflation rates, unemployment statistics, and government stability indices, were sourced from au-

thoritative government databases, specifically the Federal Reserve Economic Data (FRED) and Eurostat. These platforms offer reliable and up-to-date economic data from the US and EU regions.

To capture the influence of economic sentiment, we conducted sentiment analysis on relevant economic news articles. The news data was gathered from a wide range of media outlets, including social media platforms, financial news agencies, and specialized economic websites. We used the News API (<https://newsapi.org>), which aggregates data from various sources such as Google News. In addition, historical data was extracted from renowned economic portals like Bloomberg and Reuters. News articles were scraped and processed using a combination of APIs and web scraping techniques to ensure comprehensive coverage of economic events.

All collected data underwent an extensive analysis phase, where generative AI models were employed to process the information and extract meaningful insights. This AI-driven analysis provided us with aggregated output, which was further refined to ensure consistency and quality.

Prior to model training, rigorous data cleaning and scaling methods were applied to ensure the accuracy and relevance of the datasets. Data consistency checks were performed to identify and correct any discrepancies, and normalization techniques were used to scale numerical values for uniformity.

To automate the process of dataset creation and preparation, we developed a Dataset Generation Agent, which streamlined the entire data collection and preprocessing workflow. A detailed description of this agent, including its architecture and operational principles, is provided in Section 3.4.2.

Finally, the entire dataset was partitioned into three subsets: 70% of the data was allocated for training purposes, 20% for validation, and 10% for test. This partitioning ensures an unbiased evaluation of model performance and allows for robust model generalization.

1.2.2. MODELS' SELECTION AND TRAINING

Traditionally, statistical models and neural network architectures have been the predominant methods for forecasting sequential data in time-series prediction tasks. Based on our literature review, Long Short-Term Memory (LSTM) networks - a class of recurrent neural networks (RNNs) - have demonstrated considerable effectiveness due to their capacity to capture long-term dependencies and complex temporal patterns within sequential datasets. However, recent advancements in generative artificial intelligence have opened new avenues for time-series forecasting. In this context, we considered it essential to explore the applicability of state-of-the-art generative models for predicting time-series data, with a specific focus on the EUR/USD currency exchange rate. Models such as LLaMA and ChatGPT, which have shown exceptional performance in natural language generation and comprehension, motivated our investigation into their potential to model and forecast continuous time-series sequences.

In our study, we focused on testing two lightweight generative models: Llama-3.2-1b and GPT-2 correspond to the selected ECO approach that is described in Section 3.3.

The primary motivation behind using Llama and GPT-2 lies in their ability to process large sequences of data, and their underlying transformer architecture, which has demonstrated strong capabilities in learning complex patterns. By applying these models to consecutive time-series data like EUR/USD exchange rates, we aim to assess whether such generative approaches can yield competitive or even superior results compared to traditional models like LSTM.

Long Short-Term Memory (LSTM) networks and Generative AI models, specifically Llama-3.2-1B and GPT-2 were initially tested on historical EUR/USD exchange rate data without any optimization or additional parameters, with the goal of forecasting EUR/USD rates over time periods of 7, 30, 180, and 365 days. The results showed that the LSTM model outperformed the Generative AI models in terms of prediction accuracy (see Table 1).

Table 1. Predictions of the EUR/USD rates by LSTM, Llama-3.2-1B and GPT-2 before optimization

Model	7 days		30 days		180 days		365 days	
	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE
LSTM	0.06278	0.05789	0.07292	0.07145	0.0715	0.07086	0.06346	0.06336
GPT-2	0.08308	0.08286	0.08894	0.08851	0.12113	0.11861	0.1204	0.11876
Llama-3.2-1B	0.08211	0.08003	0.08653	0.08479	0.11905	0.11634	0.1173	0.11467

Source: Authors' calculation

To evaluate the performance of the selected models, we employed several evaluation metrics, including **Root Mean Squared Error** (RMSE) and **Mean Absolute Error** (MAE).

RMSE is a widely used metric that measures the average magnitude of the errors between predicted and actual values. It is calculated as the square root of the average of the squared differences between predicted and actual values, as follows:

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

where:

- n is the total number of observations,
- y_i is the actual value,
- \hat{y}_i is the predicted value.

A lower RMSE value indicates better model performance, as it signifies that the predictions are closer to the actual values. RMSE is sensitive to large errors due to the squared term, making it suitable for cases where significant errors need to be penalized.

MAE is another commonly used metric that measures the average of the absolute differences between predicted and actual values. It is calculated as:

$$MAE = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i| \quad (2)$$

Unlike RMSE, MAE treats all errors equally, making it more robust to outliers. A lower MAE value suggests a model with higher accuracy.

For the loss function of the LSTM model was used MSE (Mean Squared Error) and for fine-tuning of the generative AI models was used Cross-Entropy Loss.

$$MSE = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2 \quad (3)$$

$$L = -\frac{1}{n} \sum_{i=1}^n \sum_{c=1}^C y_{i,c} \log \hat{y}_{i,c} \quad (4)$$

where:

- n is the number of samples,
- C is the number of classes,
- y_{ij} is the actual binary indicator (1 if the class label is correct, 0 otherwise) for sample i and class j,
- \hat{y}_{ij} is the predicted probability for class j for sample i,
- The logarithm is taken element-wise.

This loss function is commonly used for classification tasks, particularly in neural networks, where it measures the difference between the predicted probability distribution and the true distribution. Lower cross-entropy loss values indicate better model performance.

Generative AI models, such as Llama-3.2-1B and GPT-2, are primarily designed for natural language processing tasks, where they predict sequences of text tokens. As such, they struggled to accurately predict continuous numerical values, like the EUR/USD exchange rates. The models did not grasp

the context of the financial data and produced the same predicted value regardless of the forecast horizon. In contrast, the LSTM model, which is specifically designed to capture temporal dependencies in sequential data, showed natural suitability for time-series forecasting, leading to better prediction outcomes.

To improve the predictive performance of both the LSTM model and the Generative AI models, a series of targeted optimizations were implemented. For the LSTM model, we utilized a Training Agent (as detailed in Subsection 1.2.3) to explore a range of hyperparameter configurations. These included adjustments to the number of training epochs, the number and size of hidden layers, and the application of various dropout rates to mitigate overfitting.

In parallel, we applied fine-tuning techniques to the Generative AI models. Given the eco-conscious design of our approach and constraints on computational resources, the fine-tuning process was deliberately limited to a maximum of 10 training epochs per model. This strategy was intended to balance performance gains with resource efficiency.

Following these optimization steps, we evaluated the forecasting results of the optimized LSTM model and the fine-tuned generative models—LLaMA-3.2-1B and GPT-2—using standard performance metrics. The results, summarized in Table 2, indicate that the optimized LSTM model achieved a marked improvement in prediction accuracy compared to its baseline configuration, thereby confirming the effectiveness of the hyperparameter tuning process. Conversely, the fine-tuned generative models exhibited only marginal performance gains, suggesting a more limited sensitivity to fine-tuning under the given constraints.

Table 2. Predictions of the EUR/USD rates by LSTM, Llama-3.2-1B and GPT-2 after optimization

Model	7 days		30 days		180 days		365 days	
	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE
LSTM	0.01133	0.01104	0.01268	0.01239	0.0114	0.01107	0.01041	0.01019
GPT-2	0.07512	0.07387	0.08156	0.08042	0.10589	0.10231	0.09745	0.09401
Llama-3.2-1B	0.07344	0.07205	0.07871	0.07012	0.10003	0.09021	0.09531	0.08873

Source: Authors' calculation

Based on these results, we conclude that for forecasting time-series data, such as the EUR/USD exchange rates, an optimized LSTM model is more effective than Generative AI models like Llama-3.2-1B and GPT-2. The LSTM model's ability to handle sequential data and learn from temporal dependencies makes it better suited for this task.

1.2.3. HYBRID ECO-MODEL ELABORATION

In the development of our model, we adopted an eco-approach aimed at creating a lightweight, efficient system that minimizes computational resource requirements, which in turn reduces carbon emissions associated with model training and execution. The broader goal of this eco-friendly approach was not only to optimize for energy consumption but also to make the model more accessible, enabling it to run on a wide range of devices and operate at higher speeds.

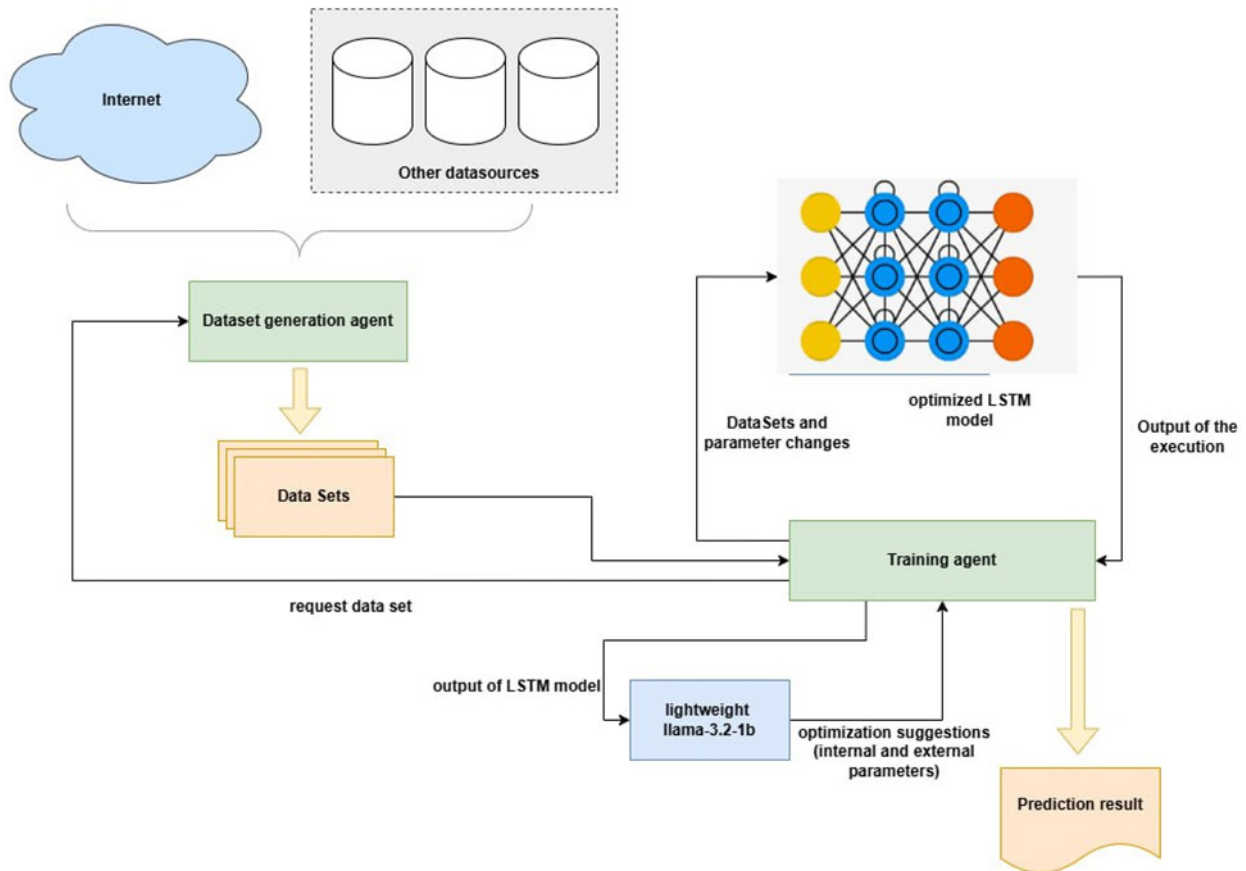
Our ECO-approach encompasses the following key strategies:

- **Use of Lightweight Models:** We prioritized the use of lightweight models for developing a hybrid prediction solution.
- **Data Set Optimization Based on Prediction Horizon:** The size and complexity of the datasets used for training were optimized based on the prediction horizon. For example, for short-term forecasting (7-day prediction), we utilized smaller datasets, which reduced both the time required for training and the computational power needed. For medium-term forecasts (30-day predictions), a medium-sized dataset was used, while for longer-term predictions (180 and 365 days), we employed larger datasets. This tiered approach ensures that only the necessary amount of data is processed at each stage, minimizing resource consumption without compromising the accuracy of the forecasts.

- **Use of Training Agent for Model Optimization:** A central component of our eco-approach is the Training Agent, which plays a crucial role in maintaining a minimal model size. The agent works by interacting with the generative AI model (Llama-3.2-1B) to select and optimize the hyperparameters. This allows us to focus on identifying and adjusting only those parameters that have the most significant impact on the model's performance.
- **Efficient Deployment Across Devices:** A key goal of our eco-approach is to ensure that the model is not only lightweight in terms of training but also optimized for inference. By focusing on computational efficiency, we aim to make the model deployable on a wide range of devices, from standard workstations to mobile devices.

To provide accurate and effective predictions of the EUR/USD exchange rate, we have developed a hybrid eco-model, combining various AI techniques to enhance the model's performance. The architecture of this hybrid model is illustrated in Figure 1.

Figure 1. Hybrid ECO-Model high-level architecture



Source: Authors' design

The system is composed of two primary agents and two AI models. The first agent is the Dataset Generation Agent (discussed in section 3.4.2), responsible for collecting and preparing the data needed for training. The second is the Training Agent (see section 3.4.1 for more details), which manages the overall training process. The AI models involved include an optimized Long Short-Term Memory (LSTM) model, designed to predict the EUR/USD rate, and a lightweight generative AI model (Llama-3.2-1B), which analyzes the outputs of the LSTM model and provides valuable feedback for refining both internal and external model parameters.

The **Training Agent** (TA) initiates the training process by requesting the LSTM model to predict exchange rates for a specific forecasting horizon. It then queries the Dataset Generation Agent for the necessary datasets, gathers this data, and trains the LSTM model with default parameters. The LSTM model's output is collected by the Training Agent and sent to the Llama model for further analysis.

The **Generative AI Model** (Llama-3.2-1B) processes the LSTM model's outputs, analyzing the results to identify potential improvements. These suggestions may include adjustments to internal

parameters such as the number of epochs, dropout rates, and the number of layers in the LSTM model. Additionally, the generative model may propose changes to the dataset, such as the inclusion or removal of specific training parameters.

TA plays a central role in the Hybrid ECO-Model, acting as the system's coordinator for model training and optimization. It operates iteratively, incorporating feedback from the generative AI component to refine the LSTM model's performance. Based on this feedback, the agent may modify the model's architecture or request new datasets as needed, re-running the training process to improve accuracy. This cycle repeats until the generative AI determines that the LSTM model is sufficiently optimized for the target prediction horizon.

Once the optimization process is complete, the Training Agent saves the finalized LSTM model for future deployment. It also documents and publishes the outcomes of both the training and testing phases, ensuring transparency and reproducibility of the results.

To better understand how this system functions, we now examine the key components of the Hybrid ECO-Model, starting with the Training Agent itself.

As a core component of the model TA is responsible not only for executing training scenarios but also for facilitating communication between different modules within the system. It ensures smooth data flow and model interaction through its integrated API interfaces. Specifically, the agent is equipped with a REST API that allows efficient communication with external systems and components.

One of its key responsibilities is triggering the Dataset Generation Agent, which handles the retrieval and preparation of training datasets. Once the dataset is available, the Training Agent initiates the execution of the Long Short-Term Memory (LSTM) model - an advanced type of recurrent neural network - by adjusting a range of hyperparameters to enhance model performance. These include:

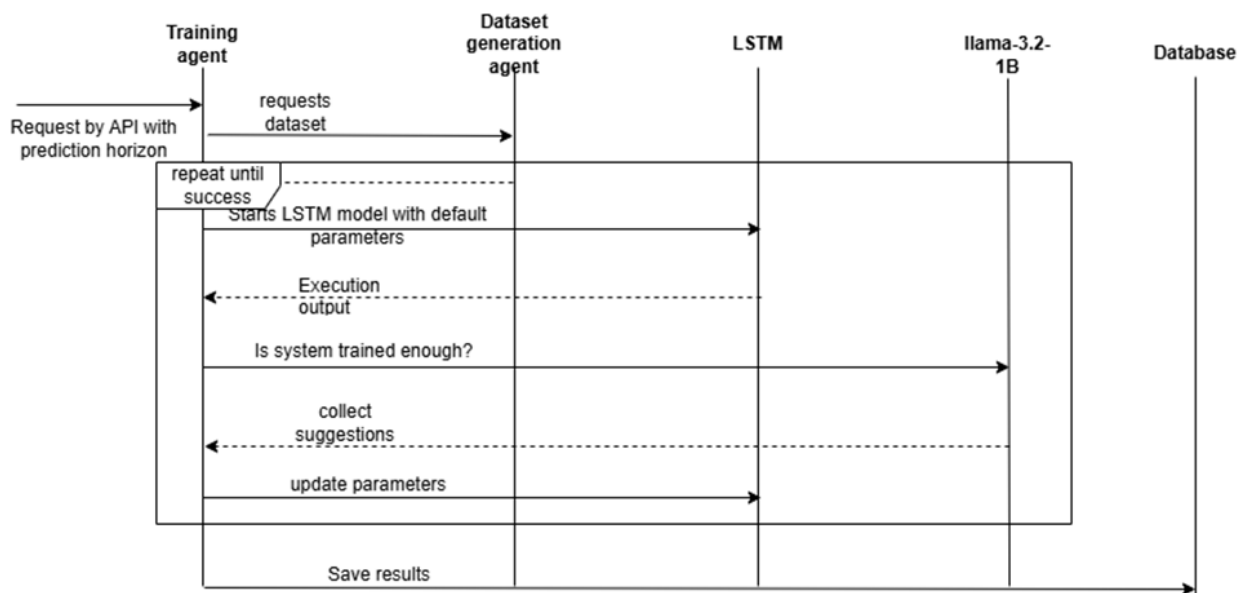
- Number of epochs: Defines how many times the training process iterates over the dataset.
- Dropout rate: Regulates the dropout layers in the LSTM to prevent overfitting.
- Loss function: Determines the metric used to optimize the model.
- Hidden layers and size: Specifies the architecture of the LSTM, including the number of hidden layers and their respective sizes.

TA executes the LSTM model using the specified parameters and collects the resulting outputs. To enhance the model's performance, it then engages in a feedback loop with the Generative AI model (LLaMA-3.2-1B), leveraging prompt engineering techniques. The generative model analyzes both the inputs and outputs of the LSTM and returns structured recommendations for improvement. These may involve changes to the model architecture, training configurations, or other hyperparameters that could boost predictive accuracy.

TA interprets these suggestions and applies the proposed adjustments to the LSTM model. Once the updates are made, it re-initiates the dataset retrieval and training process. This iterative cycle continues until the LLaMA-3.2-1B model confirms that the LSTM has been sufficiently optimized for the desired forecasting horizon.

After the optimization is complete, the final version of the LSTM model, along with its training results, is stored in a centralized database. This allows the model to be reused for future predictions or further analysis without requiring retraining—ensuring both efficiency and reproducibility. Figure 2 illustrates a sequence diagram depicting one of the training agent scenarios, showing the sequence of interactions between the agent and other system components during the model training process.

Figure 2. Sequence diagram of the training agent scenario



Source: Authors' design

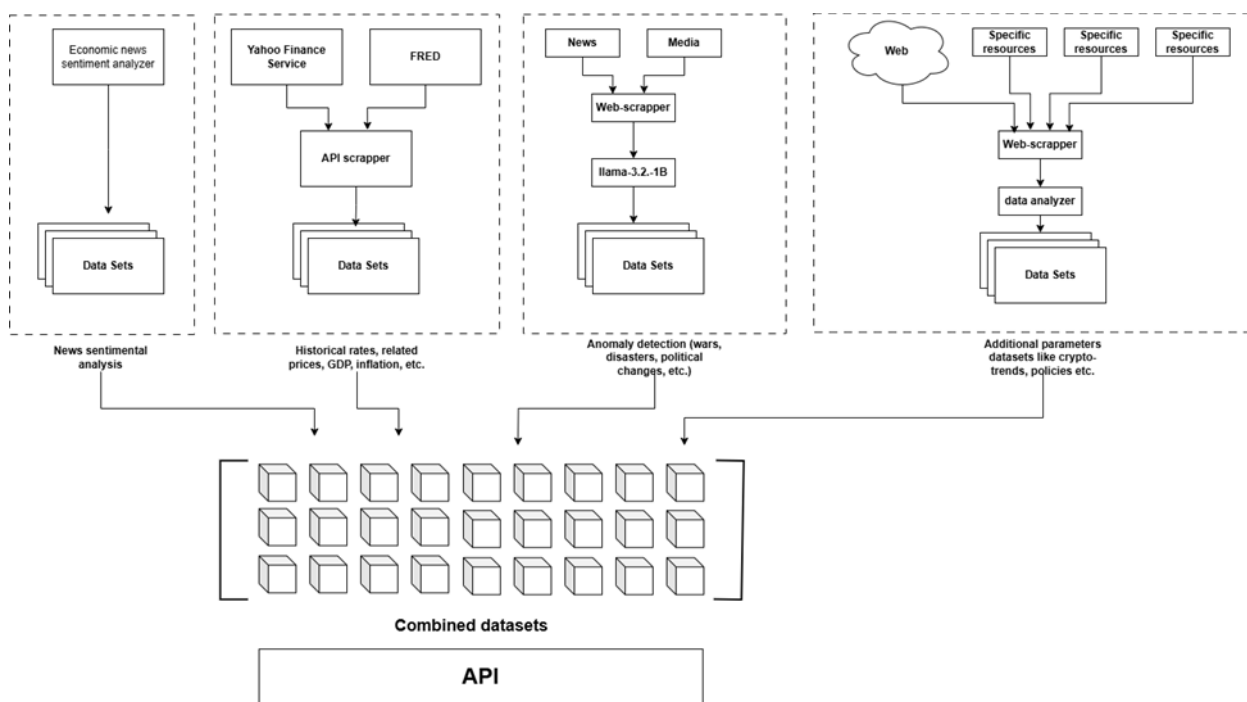
TA is also responsible for interaction with generative AI model to select the most optimal external input parameters that should be included in the final configuration (like price oil and gold).

Dataset Generation Agent is a critical component of the hybrid eco-model, responsible for collecting data and generating datasets. It consists of several micro-services, each dedicated to a specific task:

- News Sentiment Analyzer
- Financial and Economic Indicators Collector
- Anomaly Detection Analyzer
- Additional Parameters Collector

Figure 3 illustrates the architecture of the Dataset Generation agent.

Figure 3. Dataset Generation Agent architecture



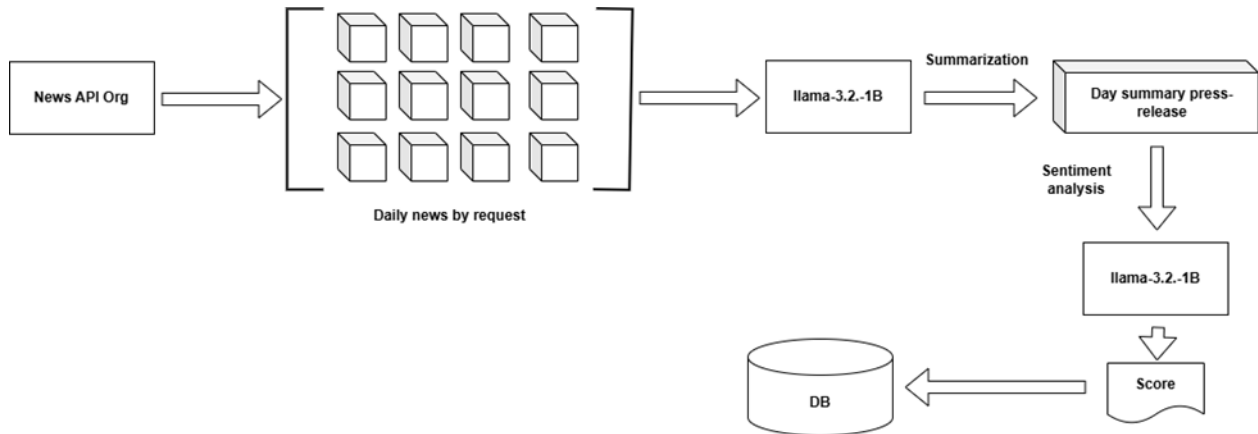
Source: Authors' design

The Dataset Generation agent provides a REST API to request the necessary datasets. It can trigger the execution of various analyzers either on-demand via API requests or based on a pre-defined schedule. The generated datasets are then stored in CSV format, either in file storage or a relational database.

Below is a more detailed description of each submodule:

The News Sentiment Analyzer is responsible for gathering daily financial news, generating daily press releases using the generative AI model (Llama-3.2-1B), and performing sentiment analysis on these press releases using the same model (see Figure 4).

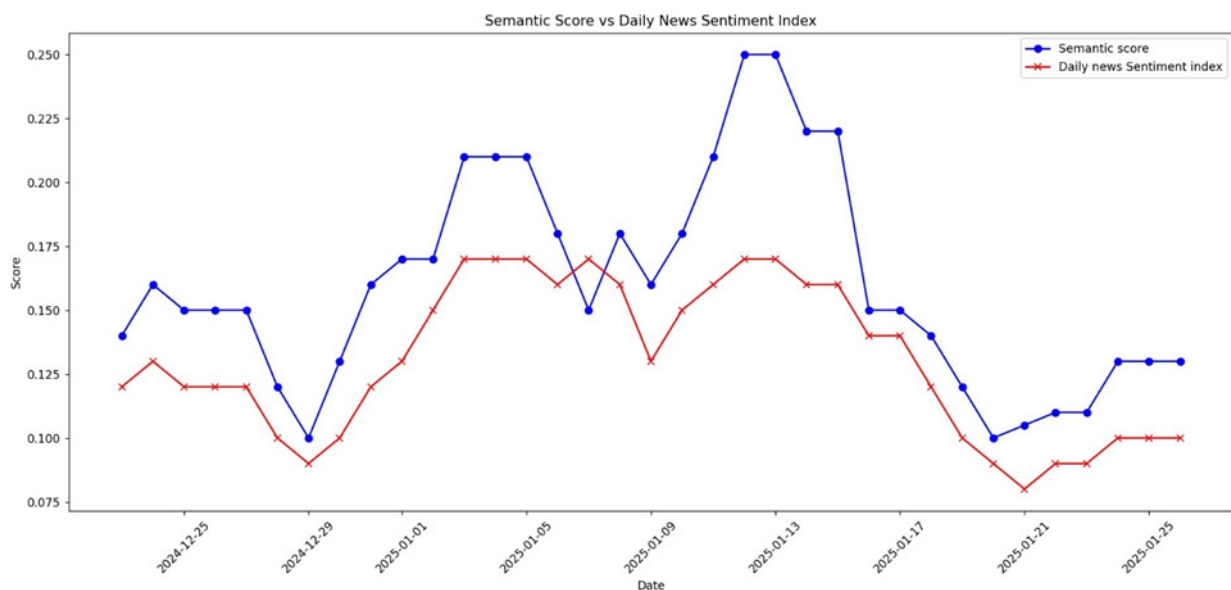
Figure 4. News Sentimental Analyzer architecture



Source: Authors' design

The results of the sentiment analysis are stored in the database. The analyzer collects historical news for a specified time period. We utilize the News API to gather data; however, this API, via a developer account, only allows access to 30 days of historical data. As a result, we were able to create datasets limited to a 30-day period. In the context of the hybrid eco-model, we require data spanning multiple years, ideally starting from 01/01/2010. To address this gap, we decided to use the Daily News Sentiment Index, as it closely correlates with the sentiment scores generated by the News Sentiment Analyzer (see Figure 5).

Figure 5. Correlation between Financial News Sentimental Score and 'Daily news sentiment index'



Source: Authors' calculations and graphing

Financial and economic indicators collector – The Financial and Economic Indicators Collector gathers relevant data from various sources, such as Yahoo Finance and the Federal Reserve Economic Data (FRED). The collector interacts with the REST APIs of these sources and saves the gathered data into CSV files for further use in the model.

The Anomaly Detection Analyzer is a promising yet currently unimplemented submodule. Its objective is to analyze various online sources to detect anomalies such as wars, natural disasters, pandemics, political changes, or other significant events that could influence financial and economic indicators. This component aims to incorporate such real-world events into the model for better accuracy and robustness.

The Additional Parameters Collector gathers supplementary minor economic parameters from open sources. It leverages both APIs and web scraping techniques to collect information, generating datasets that include these additional parameters to enhance the hybrid eco-model.

In summary, the Dataset Generation agent orchestrates the collection and preparation of diverse data sources, enabling the hybrid eco-model to train, predict, and simulate based on a wide range of financial and economic indicators.

Generative AI models play a crucial role in our hybrid ECO-model, driving decision-making processes based on the training outcomes of the LSTM model. Specifically, these models are responsible for adjusting the structure of the LSTM network, incorporating additional parameters into the dataset, modifying model coefficients, and other tasks that enhance the model's performance. Furthermore, they contribute significantly to the dataset generation process by providing capabilities such as text summarization and sentiment analysis.

Given that our objective is to minimize computational resource consumption in the hybrid ECO-model, we opted against fine-tuning the generative AI model. Instead, we employed prompt engineering to guide the model's behavior. While Retrieval-Augmented Generation (RAG) could be considered as an alternative, we found that prompt engineering alone was sufficient for the scope of our experiment.

In the overall system architecture, the generative AI model is tasked with several critical responsibilities:

- **Text Analysis and Suggestion Generation:** The model analyzes input data and generates recommendations for improving dataset quality or adjusting model parameters.
- Model also responsible for selecting the strongest parameters to the final dataset.
- **Input-Output Suggestion Based on Data:** The model offers insights and recommendations based on the relationships between the input data and output predictions.
- **Text Summarization:** It summarizes relevant news articles and reports to provide concise and relevant information for the model.
- **Sentiment Analysis of News:** It evaluates the sentiment of economic news to assess the market's potential impact on the EUR/USD exchange rate.

The operations of the generative AI model are orchestrated by the Training Agent and Dataset Generation Agent, ensuring that data preprocessing, model training, and decision-making are seamlessly integrated. Throughout our experiment, the generative AI module demonstrated high performance, effectively contributing to the accuracy and efficiency of the overall system.

2.RESULTS AND DISCUSSIONS

We implemented the Hybrid LSTM-Generative AI ECO-model and conducted an experiment to predict the EUR/USD exchange rates based on historical data and a specified prediction horizon. The model demonstrated the ability to effectively handle this task, providing fairly accurate predictions, making it suitable for use by a wide range of organizations in currency forecasting applications. Notably, the system's efficiency offers significant financial gains, particularly for financial institutions, investment firms, and currency traders who rely on precise exchange rate forecasts to inform their decision-making processes.

The system was trained and operated according to recommendations from the Generative AI module, which provided real-time guidance on LSTM model improvements. The results of the training process are summarized in Table 3, which presents the prediction accuracy achieved using various model configurations. Our analysis demonstrated that incorporating the "Inflation" parameter into the historical data resulted in reduced prediction accuracy compared to using the "GDP" parameter

alone. Furthermore, the inclusion of any of the “Inflation” and “GDP” parameters led to lower accuracy than adding the “Oil/Gold” parameters. This suggests that macroeconomic indicators may not have an immediate direct impact on the exchange rate, but they could exert a more significant influence over the long term.

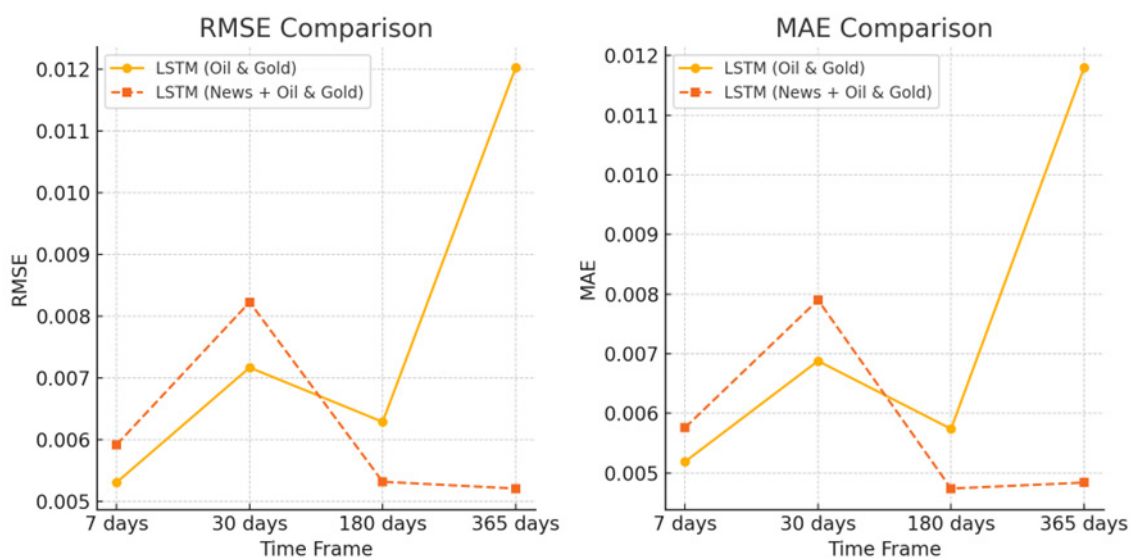
Table 3. Predictions of the EUR/USD rates by LSTM model with different parameters

Feature	7 days		30 days		180 days		365 days	
	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE
Inflation	0.05206	0.04842	0.0617	0.06054	0.0574	0.05711	0.05689	0.05677
GDP	0.05042	0.0463	0.05389	0.05283	0.05557	0.05535	0.05335	0.05323
Oil and gold	0.00531	0.00519	0.00717	0.00688	0.00629	0.00574	0.01203	0.0118
News semantic analysis	0.00959	0.00925	0.01363	0.013263	0.01024	0.00987	0.0118	0.0116
Oil/gold + semantic analysis	0.00592	0.00576	0.00823	0.007902	0.00532	0.00474	0.00521	0.00484

Source: Authors’ calculation

One of the key suggestions from the Generative AI model was to incorporate a combined oil/gold parameter. The inclusion of this combined parameter led to a significant improvement in prediction accuracy (from 5 to 12 times depending on the prediction horizon). As part of our goal to develop an optimal model with minimal resource usage, the Generative AI module further recommended adding sentiment analysis of economic news to enhance model performance. By integrating this additional feature, the final model contained three extra parameters: oil/gold prices, historical EUR/USD exchange rates, and news sentiment analysis.

Figure 6. Comparison of prediction accuracy for Oil&Gold and combined Oil&Gold+News parameters: a) Description of RMSE comparison is contained in the left part; b) Description of MAE Comparison is contained in the right part



Source: Authors’ calculation and graphing

For short-term prediction horizons, the final model with the added sentiment analysis parameter showed slightly lower accuracy compared to the LSTM model (in about 11% for 7 days prediction) that only included the oil/gold parameter. However, for longer prediction horizons, the inclusion of the news sentiment analysis parameter resulted in a dramatic improvement, with prediction accuracy more than doubling the performance of the model that used only oil/gold data. This underscores the importance of incorporating sentiment analysis for long-term forecasting, where external factors such as geopolitical events or economic shifts have a more pronounced impact.

The study conducted advances the field by proposing a more integrated and balanced hybrid architecture—the Hybrid LSTM-Generative AI Prediction ECO-Model—which ensures that both model types contribute meaningfully to the forecasting task. By designing a system where the generative

model plays a more active role in both pre-processing and direct forecasting tasks, this research moves beyond the auxiliary use of generative AI and demonstrates how deeper synergy can lead to improved forecast accuracy, especially in long-term horizons.

Moreover, the emphasis on efficiency and practicality—a lightweight, resource-conscious design—directly addresses the increasing demand for scalable and deployable forecasting tools in real-world economic policy and industry settings. In doing so, this study not only refines the theoretical understanding of hybrid forecasting models but also offers a viable path toward more accessible, interpretable, and cost-effective forecasting solutions.

Thus, the findings contribute meaningfully to the macroeconomic forecasting literature by:

- Challenging the traditional asymmetric hybrid model design;
- Demonstrating the performance benefits of a more balanced integration of LSTM and generative AI;
- Introducing an eco-efficient framework that aligns with practical deployment needs;
- Setting the stage for future work on generalizable, modular hybrid forecasting systems.

While the system demonstrated exceptional performance in predicting the EUR/USD exchange rate, several limitations were noted, which should be considered for future improvements and real-world application:

- **Limited Fine-Tuning of the Generative AI Model:** Due to the constraints of our eco-approach, we were unable to significantly fine-tune the Generative AI model. Instead, we relied on prompt engineering and, as an optional strategy, Retrieval-Augmented Generation (RAG). The inability to fine-tune the model restricts its flexibility in adapting to highly specific market conditions.
- **Limited Access to Historical Financial News:** The sentiment analysis module relied on available financial news data, but due to limited access to comprehensive historical and more specific (for example, exchange rates) financial news, the model's sentiment analysis was not as robust as it could be. This limitation affected the model's ability to accurately predict sentiment scores for specific days, which, in turn, impacted its performance.
- **Restricted Access to Broader Media and Social Media:** The model lacked access to a wide range of historical media and social media resources that could be used for anomaly detection. The absence of these data sources may have limited the model's ability to detect sudden shifts in market sentiment driven by social or media events.
- **Minimalist Approach to Parameter Inclusion:** In keeping with our goal of developing a minimalist LSTM model with the fewest parameters necessary, we did not include all potential weak parameters in the model. While this approach optimized resource consumption, it may have excluded some parameters that could have further enhanced the model's predictive capability.

3. CONCLUSION

The Hybrid LSTM-Generative AI ECO-Model presents a transformative approach to forecasting and economic indicator analysis, offering significant benefits to financial organizations and the broader economy. By integrating the predictive capabilities of Long Short-Term Memory (LSTM) networks with the creative adaptability of Generative AI, this model enhances the accuracy, adaptability, and efficiency of economic forecasts. Financial institutions, including banks, investment firms, and policy-making bodies, can leverage this model for more precise risk assessment, portfolio optimization, and macroeconomic planning.

Furthermore, the model's ability to generate synthetic economic scenarios enables stress testing and strategic decision-making, allowing businesses and governments to anticipate and mitigate potential financial disruptions. The increased forecasting accuracy enhances market stability, improves investment strategies, and fosters data-driven policymaking, ultimately leading to more resilient economic systems.

In this study, we demonstrated the feasibility of developing an efficient predictive model using an eco-approach, leveraging lightweight models and carefully selected parameters. The results indicate that, despite the advancements in generative AI models, they are not yet optimal for direct time-series forecasting. However, when integrated into hybrid systems alongside models such as LSTM, generative models can provide valuable insights and optimization strategies, contributing to improved overall predictive performance.

Our findings confirm that LSTM models, when optimized, outperform generative AI models in forecasting financial time-series data, such as the EUR/USD exchange rate. The results highlight the significance of incorporating domain-specific features, such as oil/gold prices and economic sentiment analysis, to enhance forecasting accuracy. The hybrid ECO-model, which integrates LSTM with Generative AI model proved effective in refining model predictions while maintaining minimal computational resource consumption.

Despite the success of our approach, certain limitations remain. The reliance on predefined hyperparameters, the limited fine-tuning of generative AI models, and the restricted access to historical financial news and broader social media data suggest avenues for further exploration.

Several potential directions for **future research and development** include:

- Anomaly Detection Module – Implementing an anomaly detection system to identify unusual patterns and external market shocks (e.g., political events, economic crises, or natural disasters) that could impact exchange rate predictions.
- Expansion of Data Sources – Enhancing dataset diversity by incorporating social media sentiment, broader financial news sources, and alternative economic indicators to improve model accuracy and robustness.
- Hybrid Model Enhancements – Exploring the potential of combining multiple weak parameters, which individually contribute little predictive power but may collectively enhance model performance.
- Adaptive Learning Strategies – Developing self-improving models that can dynamically adjust hyperparameters, input features, and prediction strategies based on real-time market changes.

By addressing these research directions, future iterations of the hybrid ECO-model can achieve even greater predictive accuracy, broader applicability, and enhanced adaptability to real-world financial forecasting challenges.

REFERENCES

- Atif, D. (2025). Enhancing Long-Term GDP Forecasting with Advanced Hybrid Models: A Comparative Study of ARIMA-LSTM and ARIMA-TCN with Dense Regression. *Comput Econ*. <https://doi.org/10.1007/s10614-024-10683-5>
- Aumeboonsuke, V. (2024). Digital Convergence and Divergence in EU and ASEAN Economies. *Economics - Innovative and economics research journal*, 13(1), 153–177. <https://doi.org/10.2478/eoik-2025-0002>
- Brynjolfsson, E., Li D. (2024). The Economics of Generative AI(NBER). *The Reporter*, 1, 1-6. <https://www.nber.org/reporter/2024number1/economics-generative-ai>
- Bybee, L. (2023). Surveying Generative AI's Economic Expectations, *Papers* 2305.02823, arXiv.org, revised May. <https://doi.org/10.48550/arXiv.2305.02823>
- Chang, V., Chidozie, Q. A. Xu, A., Wang, H. (2024). Predicting Economic Trends and Stock Market Prices with Deep Learning and Advanced Machine Learning Techniques. *Electronics*, 13(17), 3396. <https://doi.org/10.3390/electronics13173396>
- Che, C., Huang, Z., Li, C., Zheng, H., Tian, X. (2024). Integrating Generative AI into Financial Market Prediction for Improved Decision Making. *Applied and Computational Engineering*, 64. 155-161. <https://doi.org/10.54254/2755-2721/64/20241376>
- Hall, A. S. (2018). Machine Learning Approaches to Macroeconomic Forecasting, Economic Review. *Federal Reserve Bank of Kansas City*, Q IV 63-81. <https://doi.org/10.18651/ER/4q18SmalterHall>
- Hamiane, S., Ghanou, Y., Khalifi, H., Telmem, M. (2024). Comparative analysis of LSTM, ARIMA, and hybrid models for forecasting future GDP. *Ingénierie des Systèmes d'Information*, 29(3) 853-861. <https://doi.org/10.18280/isi.290306>

- Hassani, H., Silva, E. S. (2024). Predictions from Generative Artificial Intelligence Models: Towards a *New Benchmark in Forecasting Practice*. *Information*, 15, 6: 291. <https://doi.org/10.3390/info15060291>
- Jahn, M. (2018). Artificial neural network regression models: Predicting GDP growth. *HWWI Research Papers 185*, Hamburg Institute of International Economics (HWWI), Hamburg. <https://ideas.repec.org/p/zbw/hwwirp/185.html>
- Kamolthip, S. (2021). Macroeconomic Forecasting with LSTM and Mixed Frequency Time Series Data. PIER Discussion Papers 165. *Puey Ungphakorn Institute for Economic Research*. <https://ideas.repec.org/p/pui/dpaper/165.html>
- Kim, S., Kwon, M. J., Kim, H.H. (2024). Sentiment Analysis of News Based on Generative AI and Real Estate Price Prediction: Application of LSTM and VAR Models. *The Transactions of the Korea Information Processing Society*, 13(5), 209–216. <https://doi.org/10.3745/TKIPS.2024.13.5.209>
- Korinek, A. (2023). Generative AI for Economic Research: Use Cases and Implications for Economists, *Journal of Economic Literature*, 61 (4), 1281–1317. <https://doi.org/10.1257/jel.20231736>
- Lemaallem, H., Saadi, A. (2025). Productivity and Trade Openness: A Sectoral Analysis of Morocco's Economy. *Economics - Innovative and economics research journal*, 13(2), 529–551. <https://doi.org/10.2478/eoik-2025-0052>
- Lee, D. K. C., Guan, C., Yu, Y., Ding, Q. (2024). A Comprehensive Review of Generative AI in Finance. *Fin-Tech*, 3(3), 460-478. <https://doi.org/10.3390/fintech3030025>
- Lim, B., Zohren, S. (2021). Time-series forecasting with deep learning: a survey. *Phil. Trans. R. Soc. A*.379 20200209. <http://doi.org/10.1098/rsta.2020.0209>
- Liu, S., Motani, M. (2022). Towards Better Long-range Time Series Forecasting using Generative Forecasting, *Preprint. National University of Singapore*. <https://doi.org/10.48550/arXiv.2212.06142>
- Liu, S., Yao, J., Motani, M. (2019). Early Prediction of Vital Signs Using Generative Boosting via LSTM Networks, 2019 *IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*. San Diego, CA, USA. 437-444. <https://doi.org/10.1109/BIBM47256.2019.8983313>
- Medeiros, M. C., Vasconcelos, G. F. R., Veiga, Á., Zilberman, E. (2021). Forecasting Inflation in a Data-Rich Environment: The Benefits of Machine Learning Methods. *Journal of Business & Economic Statistics, Taylor & Francis Journals*, 39(1), 98-119. <https://doi.org/10.1080/07350015.2019.1637745>
- Nakamura, E. (2005). Inflation forecasting using a neural network. *Economics Letters*, 86.3, 373-378. <https://doi.org/10.1016/j.econlet.2004.09.003>
- Ngo, D. (2021). Predicting Stocks with LSTM-based DRNN and GAN. *Master's Projects*. 1046. 1-32. <https://doi.org/10.31979/etd.fjsz-y926>
- Nguyen H.T. et al. (2021). Macroeconomic Forecasting Based on LSTM-Conditioned Normalizing Flows. In: Ngoc Thach, N., Ha, D.T., Trung, N.D., Kreinovich, V. (eds) Prediction and Causality in Econometrics and Related Topics, ECONVN. *Studies in Computational Intelligence. Springer*. 983. 658-669. https://doi.org/10.1007/978-3-030-77094-5_48
- Sagheer, A., Kotb, M. (2019). Unsupervised Pre-training of a Deep LSTM-based Stacked Autoencoder for Multivariate Time Series Forecasting Problems. *Sci Rep* 9, 19038. <https://doi.org/10.1038/s41598-019-55320-6>
- Sarveswararao, V., Ravi, V. (2021). Prediction Intervals for Macroeconomic Variables Using LSTM Based LUBE Method. In: Gunjan, V.K., Zurada, J.M. (eds) Modern Approaches in Machine Learning and Cognitive Science: A Walkthrough. *Studies in Computational Intelligence. Springer. Cham.*, vol 956, pp. 267-275. https://doi.org/10.1007/978-3-030-68291-0_20
- Schaefer G. et al., (2023). The implications of generative AI in Finance. *Deloitte*. <https://www.deloitte.com/us/en/services/consulting/articles/generative-ai-in-finance.html>
- Shahi, T.B., Shrestha, A., Neupane, A., Guo, W. (2020). Stock Price Forecasting with Deep Learning: A Comparative Study. *Mathematics*, 8, 1441. <https://doi.org/10.3390/math8091441>
- Smyl, S. (2020). A hybrid method of exponential smoothing and recurrent neural networks for time series forecasting. *International Journal of Forecasting*, 36/1 75-85. <https://doi.org/10.1016/j.ijforecast.2019.03.017>
- Staffini, A. (2022). Stock Price Forecasting by a Deep Convolutional Generative Adversarial Network. *Front Artif Intell*, 5:837596. <https://doi.org/10.3389/frai.2022.837596>
- Yıldırım, D.C., Toroslu, I.H., Fiore, U. (2021). Forecasting directional movement of Forex data using LSTM with technical and macroeconomic indicators. *Financ Innov.*, 7/1. <https://doi.org/10.1186/s40854-020-00220-2>
- Yoshimatsu, R., Jain, P., Smith, B. (2024). An Introduction to Time Series Forecasting with Generative AI. *Databricks*. <https://www.databricks.com/blog/introduction-time-series-forecasting-generative-ai>
- Zhao, J. (2024). The Macroeconomic Prediction Model based on LSTM Improved Algorithm, International Conference on Integrated Circuits and Communication Systems (ICICACS). *Raichur. Indi*, 1-7. <https://doi.org/10.1109/ICICACS60521.2024.10498821>