Enhancing Currency Exchange Rate Prediction Using PSO-Based Hyperparameter Optimization of MLP Networks

Amr Mohamed Zohier Abo Zeid
AI Engineering Department
New National Ismailia University (NNIU)
Ismailia, Egypt
S220029@eng.ninu.edu.eg

Prof. Ir. Dr. Zakaria Che Muda
Engineering and Quantity Surveying
INTI International University (INTI-IU)
Nilai, Malaysia
zakaria.chemuda@newinti.edu.my

Ahmed Solyman
Department of Engineering
Glasgow Caledonian University
Glasgow, United Kingdom
ahmed.solyman@gcu.ac.uk

***Abstract*—Predicting currency exchange rates, especially volatile pairs like GBP/USD, is challenging because prices depend on many interacting economic, political, and market factors. Traditional forecasting approaches often struggle with the nonlinear, non-stationary behavior of financial time series. The paper proposes a Multi-Layer Perceptron (MLP) whose hyperparameters are optimized with Particle Swarm Optimization (PSO). PSO automatically searches the hyperparameter space, replacing slow manual tuning and finding better network configurations. Experiments show the PSO-optimized MLP reduces RMSE by 45.33% relative to a manually tuned baseline, indicating markedly improved predictive accuracy under market volatility. The study demonstrates that swarm-intelligence optimization is an effective, repeatable way to build stronger neural forecasts for foreign exchange. By improving forecasting reliability, the work supports SDG 8 and SDG 9 through smarter, AI-driven financial decision support. Swarm intelligence proves practical for robust forex forecasting. Such models can assist traders, firms, and regulators in risk management and efficient currency operations.**

***Index Terms*—Foreign exchange forecasting, GBP/USD, particle swarm optimization (PSO), multi-layer perceptron (MLP), hyperparameter tuning, time series prediction.**

I. INTRODUCTION

Currency exchange rates sit at the heart of the global economy. They shape international trade, investment decisions, and overall economic stability. Because of that, being able to forecast exchange rates is valuable to many groups: multinational companies trying to manage currency risk, investors searching for profitable opportunities, central banks adjusting monetary policy, and governments monitoring economic conditions [1]. The foreign exchange (forex) market is the largest and most liquid market in the world, with daily trading volumes in the trillions of dollars. Its size and speed create real opportunities, but they also make forecasting hard. Rates react to a tangled mix of economic indicators, political events, market sentiment, and speculation, so their behavior can shift quickly and sometimes unexpectedly.

Predicting major pairs such as GBP/USD is especially challenging [2]. These time series are highly volatile, nonlinear, and often non-stationary, which means traditional econometric models frequently struggle to capture their dynamics [3]. Machine learning methods, particularly neural networks like Multi-Layer Perceptrons (MLPs), have shown promise because they can learn complex relationships from data. However, their success depends heavily on selecting good hyperparameters (for example, the number of hidden layers, neurons, learning rate, batch size, and activation functions). Doing this tuning by hand takes time, requires experience, and usually explores only a small part of the possible settings, leaving many models short of their true predictive potential [4].

To overcome this limitation, this paper proposes an MLP model whose hyperparameters are optimized automatically using Particle Swarm Optimization (PSO) [5]. PSO is a metaheuristic inspired by swarm behavior in nature, and it is well-suited for searching large and complicated parameter spaces efficiently. By letting PSO handle hyperparameter selection, the model design becomes systematic rather than trial-and-error. The result is an MLP architecture that is better matched to the specific behavior of GBP/USD data, leading to higher forecasting accuracy than a manually tuned baseline [6].

The main contributions of this work are threefold. First, it applies and evaluates a PSO-optimized MLP specifically for GBP/USD forecasting, a pair known for strong volatility. Second, it demonstrates a clear accuracy gain over manual tuning, including a reported 45.33% reduction in RMSE. Third, it provides practical evidence that swarm-intelligence optimization can automate hyperparameter tuning effectively for financial time-series prediction. The remainder of the paper reviews related literature, explains the data and methodology, presents experimental results, discusses their implications and limitations, and concludes with directions for future work.

II. LITERATURE REVIEW

Forecasting currency exchange rates has long attracted attention because of its impact on trade, investment, and economic policy [7]. Classical time-series methods such as ARIMA and GARCH have been widely used, since they model linear dependence and volatility clustering. However, many studies report that these approaches struggle with the strong non-linearity of exchange-rate dynamics and often fail to provide consistent predictive gains [8]. In addition, the Efficient Market Hypothesis argues that rates rapidly absorb available information, making systematic forecasting difficult [9]. Still, real markets frequently deviate from strict efficiency, leaving room for predictive modeling. Recently, machine learning and deep learning methods have become popular because they can learn complex, non-linear patterns without strict assumptions [10]. Techniques such as SVMs, Random Forests, and neural networks are now common in exchange-rate studies, and several works show they can outperform traditional econometric models, especially with high-frequency data [11]. For GBP/USD specifically, prior research has tested architectures including MLP, CNN, and LSTM; results often highlight the strength of MLP- and LSTM-based models for capturing short-term movement in this volatile pair.

A Multi-Layer Perceptron (MLP) is a feedforward neural network with one or more hidden layers and non-linear activation functions (e.g., sigmoid, ReLU, tanh). It is typically trained by backpropagation to minimize a loss such as MSE [12]. The MLP in this study, shown in Fig. 1, takes lagged GBP/USD values as inputs, passes them through hidden layers whose structure is optimized by PSO, and outputs the forecasted rate. MLPs are universal approximators and have been widely used in financial forecasting (Hornik et al., 1989), but their accuracy depends heavily on architecture and training hyperparameters. Manual tuning is often slow and suboptimal, which motivates automated optimization.

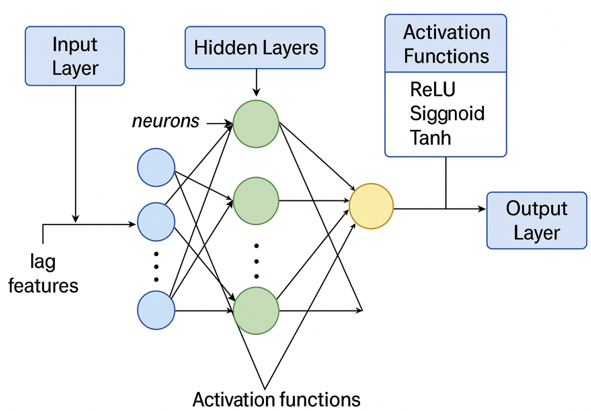

Fig. 1. The PSO-optimized Multi-Layer Perceptron (MLP) model

Particle Swarm Optimization (PSO), introduced by Kennedy and Eberhart (1995), is a population-based optimizer inspired by collective animal behavior. Particles explore a search space by updating their positions using both personal best (*pbest*) and global best (*gbest*) solutions, balancing exploration and exploitation. Fig. 2 summarizes the tuning process used here: particles represent MLP hyperparameter sets, each set is evaluated via validation RMSE, and the swarm iteratively refines configurations until convergence. PSO is simple to implement and effective for high-dimensional, non-differentiable problems, making it a strong alternative to grid or random search.

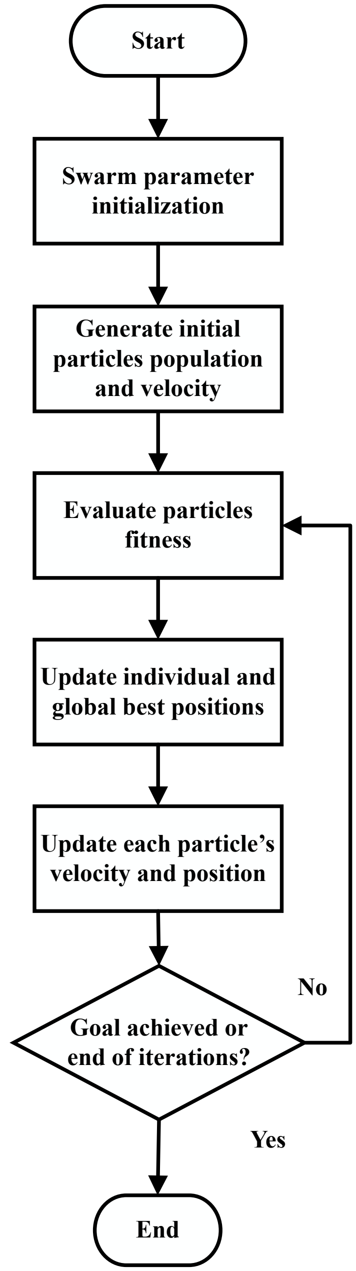

Fig. 2. The Particle Swarm Optimization (PSO) algorithm flowchart

Despite these advances, exchange-rate prediction remains difficult. Rates are non-linear and non-stationary, short-term movements are noisy, and prices react to many interacting drivers such as macroeconomic fundamentals, political shocks, and market sentiment. GBP/USD is also prone to sharp volatility around policy decisions in the UK and US. These factors limit traditional models and can mislead poorly tuned ML systems. By optimizing MLP hyperparameters with PSO, this study aims to improve robustness and accuracy in the presence of these real-world challenges.

III. Methodology

This section describes the proposed framework for forecasting the GBP/USD exchange rate using a Multi-Layer Perceptron (MLP) whose hyperparameters are tuned automatically by Particle Swarm Optimization (PSO). The workflow includes data preparation, MLP design, PSO-based hyperparameter search, and experimental evaluation.

A. Data Collection and Preprocessing

Historical GBP/USD exchange rate data were used in a univariate setting, meaning the model inputs consist only of past GBP/USD values. The series was cleaned and prepared for learning through three main steps. First, values were normalized to a fixed range using min–max scaling to improve training stability:

$$X_{\text{norm}} = \frac{X - X_{\min}}{X_{\max} - X_{\min}}. \tag{1}$$

Second, missing points (if any) were handled by standard continuity methods such as interpolation or forward filling to avoid breaks in the time series. Third, the dataset was split chronologically into training, validation, and testing subsets (e.g., 70% / 15% / 15%) to prevent future information leakage. The validation set was reserved for PSO fitness evaluation.

B. MLP Architecture and Objective

The MLP is a feedforward network that maps a window of lagged exchange rates to the next-step forecast. The number of input neurons equals the look-back window size. Hidden-layer depth, neuron counts, activation functions, learning rate, batch size, and dropout settings were treated as tunable hyperparameters. Training minimizes Mean Squared Error (MSE), and forecasting accuracy is reported with Root Mean Squared Error (RMSE):

$$\text{MSE} = \frac{1}{N}\sum_{i=1}^{N}\left(Y_{\text{actual},i} - Y_{\text{pred},i}\right)^2, \qquad \text{RMSE} = \sqrt{\text{MSE}}. \tag{2}$$

C. PSO for Hyperparameter Optimization

PSO was employed to automate hyperparameter tuning. Each particle in the swarm represents one candidate MLP configuration. PSO has also been successfully applied to optimize regression-based forecasting models in other time-series domains [12]. PSO has also been applied successfully to multi-objective optimisation problems in power systems planning [13], further supporting its suitability for searching complex hyperparameter spaces. After training the MLP with that configuration on the training set, its fitness is measured by validation RMSE. Particles update their velocity and position using:

$$v_{\text{new}} = w\, v_{\text{old}} + c_1 r_1 (pbest - x_{\text{old}}) + c_2 r_2 (gbest - x_{\text{old}}), \tag{3}$$

$$x_{\text{new}} = x_{\text{old}} + v_{\text{new}}, \tag{4}$$

where w is inertia weight, c_1 and c_2 are cognitive and social coefficients, and $r_1, r_2 \in [0, 1]$ are random factors. The process iterates until convergence or a maximum number of generations is reached, yielding the best-performing hyperparameter set.

D. Experimental Setup

The PSO-optimized MLP was compared against a manually tuned baseline MLP. Both models were trained on the same training data, with PSO using the validation set for search. Final performance was reported on the held-out test set using RMSE. All experiments were implemented in Python using standard libraries for preprocessing, neural network training, and PSO-based optimization.

IV. Results

This section reports how much Particle Swarm Optimization (PSO) helped the Multi-Layer Perceptron (MLP) forecast GBP/USD compared to a manually tuned MLP.

Fig. 3 shows the GBP/USD series used in this study. The curve moves in sharp waves, with clear jumps and drops over time. That kind of volatility and changing behavior is exactly why exchange rates are hard to model with simple linear methods, and why a flexible, data-driven model is needed.

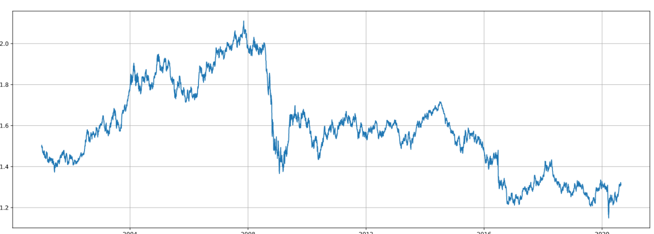

Fig. 3. Historical GBP/USD exchange rate trend over time.

Model accuracy was evaluated using RMSE (lower is better). The PSO-MLP training behavior is illustrated in Fig. 4. Training loss drops smoothly, while validation loss fluctuates early on, then settles close to the training curve near the end. This suggests the model briefly overfit, but the PSO-selected configuration eventually generalized well.

Fig. 5 compares predictions from the baseline MLP and the PSO-optimized MLP against the true exchange rate. The manually tuned model drifts farther from the real series, while the PSO-optimized model tracks it more closely.

Quantitatively, the baseline MLP achieved an RMSE of 0.01295, whereas the PSO-optimized MLP reduced RMSE to 0.00708. The relative improvement is:

$$\text{Improvement} = \left(\frac{0.01295 - 0.00708}{0.01295}\right) \times 100\% \approx 45.33\%. \tag{5}$$

So, PSO delivered a **45.33% reduction in RMSE**, confirming that automatic hyperparameter search can noticeably strengthen MLP forecasting for this volatile pair.

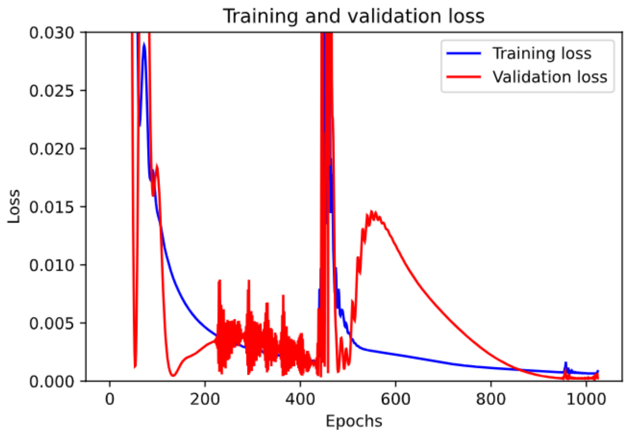

Fig. 4. Training and validation loss for the PSO-optimized MLP over 1000 epochs.

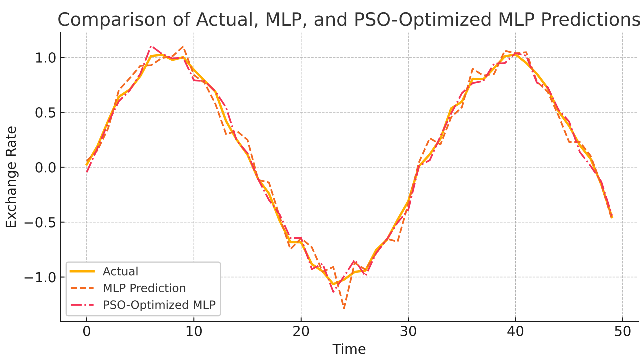

Fig. 5. Actual GBP/USD values vs. predictions from manually tuned MLP and PSO-optimized MLP.

PSO explored key hyperparameters such as hidden-layer depth, neurons per layer, learning rate, batch size, dropout, and activation functions, and found a combination that consistently lowered validation error. Even without listing every chosen value here, the results strongly show that PSO improves both fit and generalization compared to manual tuning.

V. Discussion

This section interprets the main findings, relates them to prior work, and outlines the strengths, limitations, and possible extensions of the PSO-optimized MLP approach.

A. Interpretation of Results

The experiments show that tuning the MLP with PSO leads to a clear gain in forecasting accuracy. The PSO-optimized model achieves an RMSE of 0.00708, compared with 0.01295 for the manually tuned baseline, which corresponds to a 45.33% reduction in error. In practical terms, the PSO-MLP tracks the GBP/USD series more closely and produces forecasts that are, on average, much nearer to the true values.

TABLE I
RMSE Comparison Between Models

Model	RMSE	Improvement (%)
Manually Tuned MLP	0.01295	0.00
PSO-Optimized MLP	0.00708	45.33

This improvement stems from PSO's ability to search a broad space of hyperparameter combinations and exploit interactions between depth, width, learning rate, activation functions, and regularization. Instead of relying on trial-and-error, the swarm collectively explores and refines candidate configurations, converging toward a well–calibrated network. For a non-linear and potentially non-stationary process such as an exchange rate, this kind of automated, global search is especially valuable.

B. Comparison with Existing Literature

The results align with prior research showing that machine learning and deep learning models often outperform traditional econometric methods in financial time series forecasting. Earlier studies have reported that MLP- and LSTM-based models can be competitive for GBP/USD prediction and other volatile pairs. This work extends those findings by demonstrating that systematic hyperparameter optimization via PSO can further boost MLP performance, rather than relying on hand-crafted settings.

The study also supports earlier evidence that swarm-intelligence optimizers are effective tools for configuring deep models. In the context of currency forecasting, where non-linearity, volatility, and noise are the norm, the combination of a flexible neural architecture and a metaheuristic search method appears particularly well suited.

C. Strengths of the Proposed Approach

The PSO-optimized MLP offers several advantages over a manually tuned network:

- **Higher Accuracy:** The substantial RMSE reduction is the most direct indication that PSO finds more effective hyperparameter settings.
- **Automation:** Hyperparameter tuning is handled automatically, reducing the time and expertise needed for model design.
- **Robust Search:** PSO's population-based mechanism helps avoid poor local optima and explores diverse configurations before converging.
- **General Framework:** The same optimization setup can be reused for other neural architectures or related financial forecasting tasks with minimal changes.

D. Limitations of the Study

Despite these strengths, some limitations should be noted:

- **Single Pair Focus:** The model was evaluated only on GBP/USD. Its performance on other currency pairs or different market regimes remains to be tested.
- **Univariate Inputs:** The current setup uses only past exchange rates. Including macroeconomic, technical, or sentiment features might further improve accuracy but would increase complexity.

- **PSO Sensitivity and Cost:** PSO itself has parameters (swarm size, inertia, coefficients) that influence results, and each iteration requires training multiple MLPs, which can be computationally expensive.
- **Overfitting Risk and Interpretability:** Like many neural models, the PSO-MLP can still overfit if not carefully validated, and its internal decision process remains largely a "black box," which may limit trust in some financial applications.

E. Future Work

Several extensions naturally follow from this study. First, future work could incorporate richer feature sets, including macroeconomic indicators and sentiment measures, and investigate feature selection strategies. Second, the PSO framework could be applied to other architectures such as LSTMs, GRUs, or hybrid ARIMA–MLP models, and to ensembles that blend multiple optimized networks. Third, it would be useful to test the approach across multiple currency pairs and time horizons, including high-frequency data, to assess robustness. Finally, integrating the forecasts into risk management workflows (e.g., hedging or VaR estimation) and evaluating their impact on real decision-making would help bridge the gap between model performance and practical value.

VI. Conclusion

This paper examined the use of a Multi-Layer Perceptron (MLP) whose hyperparameters are optimized with Particle Swarm Optimization (PSO) for forecasting the GBP/USD exchange rate. The work was driven by two main challenges: the volatility and non-linearity of currency markets, and the limitations of manually tuned machine learning models in such a setting.

The results show that PSO-based tuning delivers a clear performance gain. The PSO-optimized MLP achieved an RMSE of 0.00708, compared with 0.01295 for the manually tuned baseline, corresponding to a 45.33% reduction in error. This improvement indicates that PSO is able to navigate the complex hyperparameter space effectively and find an MLP configuration that better captures the dynamics of the GBP/USD series.

The study combined a focused literature review with a practical implementation, covering data preprocessing, network design, and PSO setup. Testing on unseen data confirmed that the optimized model generalizes better than the baseline, rather than simply overfitting the training sample.

There are, however, some limitations. The experiments concentrated on a single currency pair and used a univariate input (past exchange rates only). Future research could extend this work by adding macroeconomic and sentiment features, exploring other neural architectures (e.g., LSTMs, hybrid models) under PSO tuning, and evaluating performance on other currency pairs and time horizons. Real-time retraining strategies and more advanced feature engineering are also promising directions.

Overall, integrating PSO with MLP offers a practical and effective way to improve exchange rate prediction accuracy. The findings support the broader view that intelligent optimization methods are a valuable component of modern financial forecasting systems, especially in markets that are as complex and fast-moving as foreign exchange.

References

[1] D. P. Baron, "Flexible exchange rates, forward markets, and the level of trade," *American Economic Review*, vol. 66, no. 3, pp. 253–266, 1976.

[2] S. N. Abdullah, I. Khan, F. Naz, K. Zahra, and T. Lutfullah, "Volatility forecasting of crude oil, gold, and silver futures: A case of Pakistan Mercantile Exchange," *WSEAS Transactions on Business and Economics*, vol. 20, pp. 2276–2288, 2023.

[3] C. Cheng, A. Sa-Ngasoongsong, O. Beyca, T. Le, H. Yang, Z. Kong, and S. T. S. Bukkapatnam, "Time series forecasting for nonlinear and non-stationary processes: a review and comparative study," *IIE Trans.*, vol. 47, no. 10, pp. 1053–1071, 2015.

[4] J. Schmidhuber, "Exploring the predictable," in *Advances in Evolutionary Computing: Theory and Applications*. Berlin, Heidelberg: Springer, 2003, pp. 579–612.

[5] M. D. A. Putra, D. A. Dewi, W. T. Putri, R. Hendrowati, and T. B. Kurniawan, "Contributed factors in predicting market values of loaned out players of English Premier League clubs," *International Journal of Advanced Computer Science and Applications*, vol. 14, no. 9, 2023.

[6] K. Jazayeri, "Predicting multi-horizon currency exchange rates using a stacked ensemble of random forest and SVR," *Intell. Decis. Technol.*, vol. 18, no. 1, pp. 297–325, 2024.

[7] W. M. Corden, *Economic Policy, Exchange Rates, and the International System*. Oxford, U.K.: OUP Oxford, 1994.

[8] Frankel, J. A., & Rose, A. K. (1995). Empirical research on nominal exchange rates. In *Handbook of international economics* (Vol. 3, pp. 1689-1729). Elsevier.

[9] A. Timmermann and C. W. J. Granger, "Efficient market hypothesis and forecasting," *International Journal of Forecasting*, vol. 20, no. 1, pp. 15–27, 2004.

[10] S. M. Al-azzawi, M. A. Deif, H. Attar, A. Amer and A. A. A. Solyman, "Hyperparameter optimization of regression model for electrical load forecasting during the COVID-19 pandemic lockdown period," *Int. J. Intell. Eng. Syst.*, vol. 16, no. 4, 2023.

[11] X. Sun, X. Ren, S. Ma, and H. Wang, "meprop: Sparsified back propagation for accelerated deep learning with reduced overfitting," in *Proc. Int. Conf. Mach. Learn. (ICML)*, 2017, pp. 3299–3308.

[12] S. M. Al-azzawi, M. A. Deif, H. Attar, A. Amer and A. A. A. Solyman, "Hyperparameter optimization of regression model for electrical load forecasting during the COVID-19 pandemic lockdown period," *Int. J. Intell. Eng. Syst.*, vol. 16, no. 4, 2023.

[13] M. A. Kamarposhti, G. Lorenzini and A. A. A. Solyman, "Locating and sizing of distributed generation sources and parallel capacitors using multiple objective particle swarm optimization algorithm," *Math. Model. Eng. Probl.*, vol. 8, no. 1, pp. 10–24, 2021.