



Multi-Horizon Gold Price Forecasting and Its Implications for Financial Markets

Asifa Iqbal^{1,*} Marwa M. Eid^{2,3}

¹ School of International Languages, Zhengzhou University, Henan, China

² Faculty of Artificial Intelligence, Delta University for Science and Technology, Mansoura 11152, Egypt

³ Jadara Research Center, Jadara University, Irbid 21110, Jordan

Emails: asifaqbal615@gmail.com · mmm@ieee.org

Received: May 31, 2025 Revised: June 28, 2025 Accepted: August 14, 2025 ★ Corresponding author

ABSTRACT

Accurate forecasting of gold prices remains a critical challenge in financial markets due to the nonlinear, nonstationary, and regime-dependent nature of commodity price dynamics, particularly for gold quoted against the US dollar (XAU/USD), which plays a central role as a safe-haven asset, inflation hedge, and portfolio diversifier. Motivated by the growing limitations of traditional econometric and manually tuned machine learning approaches in handling long-horizon, multi-timeframe financial data, this study proposes a robust forecasting framework that integrates deep learning with metaheuristic optimization. The main contribution of this work lies in the systematic combination of a Deep Pyramid Recurrent Neural Network (DPRNN) with advanced metaheuristic algorithms for automated hyperparameter optimization, with particular emphasis on Greylag Goose Optimization (GGO), alongside other state-of-the-art optimizers. Using historical XAU/USD data spanning from 2004 to February 2025 across multiple temporal resolutions, baseline model evaluation demonstrates that DPRNN outperforms other deep learning architectures prior to optimization, achieving a Mean Squared Error (MSE) of 0.0589, Root Mean Squared Error (RMSE) of 0.2426, and coefficient of determination (R^2) of 0.79. Following optimization, the proposed GGO-optimized DPRNN framework yields a substantial performance enhancement, reducing the MSE to 2.05×10^{-5} and RMSE to 4.52×10^{-3} , while simultaneously increasing the correlation coefficient to 0.987 and R^2 to 0.983, with near-perfect agreement metrics reflected by a Nash–Sutcliffe Efficiency of 0.986 and Willmott Index of 0.988. These results confirm the effectiveness of GGO in navigating complex hyperparameter search spaces and significantly improving predictive accuracy and stability. From an economic and financial perspective, the findings underscore the practical value of metaheuristic-optimized deep learning models for enhancing gold price forecasting, supporting more informed investment decisions, improved risk management, and greater market efficiency in volatile and uncertain financial environments.

Keywords: Gold price forecasting ▪ Financial time-series modeling ▪ Safe-haven assets ▪ Metaheuristic optimization ▪ Deep learning in financial markets

1. INTRODUCTION

Gold, commonly represented in international financial markets through the XAU/USD exchange rate, has long occupied

a pivotal position within the global economic and financial system [1, 2, 3]. Unlike conventional financial assets, gold is not directly linked to the creditworthiness of any issuing authority, which grants it a distinctive status as a store of

value and a hedge against systemic risk [4, 5]. Historically, gold has served as a monetary anchor, a reserve asset for central banks, and a strategic instrument for wealth preservation [6, 7]. Its unique economic characteristics have made it particularly attractive during periods of macroeconomic instability, financial crises, geopolitical uncertainty, and heightened inflationary pressures [8, 9]. As a result, gold prices reflect a complex interplay of monetary, financial, and behavioral factors, rendering their accurate prediction a subject of substantial academic and practical importance.

From a monetary and macroeconomic perspective, gold prices are closely associated with inflation expectations, real interest rates, and currency valuation, especially with respect to the United States dollar [10]. In environments characterized by expansionary monetary policy or declining real yields, gold is frequently perceived as a hedge against the erosion of purchasing power. Central banks and institutional investors often interpret gold price movements as an implicit signal of market confidence in fiat currencies and long-term price stability. Furthermore, gold plays a critical role in portfolio diversification strategies due to its historically low correlation with traditional asset classes such as equities and fixed-income securities, particularly during periods of market stress. Consequently, fluctuations in XAU/USD prices have direct implications for asset allocation, risk management, and long-term investment decision-making [11].

The increasing volatility and complexity of global financial markets have further intensified the need for accurate gold price forecasting [12]. Over the past two decades, financial markets have been shaped by repeated episodes of systemic disruption, including global financial crises, sovereign debt instability, pandemics, and abrupt shifts in monetary policy regimes [13]. These events have introduced nonlinear dynamics and regime-dependent behaviors that challenge conventional analytical frameworks. In such settings, price movements often exhibit heterogeneous patterns across different temporal horizons, underscoring the importance of multi-timeframe analysis [14]. Forecasting gold prices across intraday, daily, weekly, and monthly horizons is therefore essential for supporting diverse financial objectives, ranging from short-term trading and hedging to long-term strategic planning and policy analysis [15].

Despite extensive research on gold price modeling, traditional econometric and technical analysis methods exhibit notable limitations when applied to modern financial data. Linear econometric models often rely on restrictive assumptions regarding stationarity and linearity, which may not hold in the presence of structural breaks and evolving market conditions. Similarly, rule-based technical indicators are typically sensitive to noise and parameter selection, limiting their robustness and adaptability. These shortcomings become particularly pronounced when large-scale, high-frequency, and long-horizon datasets are considered simultaneously, as complex interactions across timeframes are difficult to capture using conventional techniques.

In response to these challenges, machine learning has emerged as a powerful paradigm for financial time-series forecasting. By leveraging advanced computational architectures and data-driven learning mechanisms, machine learning models are capable of capturing nonlinear relationships, tem-

poral dependencies, and hidden structures that are difficult to model explicitly. Their ability to process high-dimensional inputs and adapt to evolving patterns has positioned them as promising tools for improving the reliability and efficiency of gold price prediction across multiple temporal scales.

Notwithstanding the potential advantages of machine learning methodologies, forecasting gold prices remains a highly complex task due to several intrinsic characteristics of financial time series. One of the most prominent challenges arises from the high dimensionality associated with multi-timeframe financial data. When Open, High, Low, Close, and Volume information is collected across multiple temporal resolutions, the resulting feature space expands rapidly, increasing computational burden and exacerbating the risk of overfitting. This issue is commonly referred to as the curse of dimensionality and can significantly impair model generalization if not properly addressed.

Feature redundancy constitutes an additional challenge in gold price forecasting. Financial indicators derived from price and volume data often exhibit strong correlations, both within and across timeframes. Such multicollinearity can obscure the true informational contribution of individual features, reduce learning efficiency, and complicate model interpretation. Without effective feature selection mechanisms, machine learning models may allocate excessive capacity to redundant or irrelevant information, thereby limiting their predictive effectiveness and increasing computational cost.

Another fundamental difficulty stems from the nonlinear and nonstationary nature of gold price dynamics. Financial markets evolve continuously in response to macroeconomic developments, monetary policy decisions, and shifts in investor sentiment. Sudden structural changes, such as those induced by financial crises or major policy interventions, can alter the underlying data-generating process. These characteristics challenge learning algorithms, which must balance adaptability to new patterns with stability across long historical horizons.

Deep learning models, while capable of modeling complex dependencies, introduce further challenges related to hyperparameter sensitivity. Model performance is often highly dependent on architectural and training-related parameters, including learning rates, network depth, and regularization strategies. Inappropriate hyperparameter configurations may lead to poor convergence, excessive variance, or biased forecasts. Given the complexity and scale of financial datasets, manual tuning becomes impractical, underscoring the need for automated and systematic optimization strategies.

Overfitting represents a persistent concern, particularly when long-span historical datasets are employed. While extensive historical coverage enables the inclusion of diverse market regimes, it also increases the likelihood that models will learn regime-specific noise rather than generalizable patterns. Ensuring robust performance across heterogeneous economic conditions, including periods of crisis, expansion, and policy transition, remains one of the most critical challenges in XAU/USD price prediction.

The primary objective of this study is to develop and rigorously evaluate machine learning-based forecasting models for predicting XAU/USD prices using long-term, multi-

timeframe financial data. By systematically examining multiple learning architectures, the study aims to identify modeling strategies capable of capturing the complex temporal dependencies and nonlinear behaviors inherent in gold price dynamics.

A central focus of the study lies in the integration of metaheuristic optimization algorithms to enhance model performance through automated feature selection and hyperparameter optimization. Feature selection is employed to reduce dimensionality and eliminate redundant information, while hyperparameter optimization seeks to identify configurations that balance model flexibility and generalization. The combined application of these strategies is intended to improve learning efficiency, stability, and scalability within a unified modeling framework.

Through comprehensive model development and optimization, the study seeks to advance methodological practices in financial time-series forecasting and to provide a robust analytical foundation suitable for practical financial decision-making contexts.

This study contributes to the existing literature on financial forecasting by proposing a comprehensive machine learning framework tailored to the prediction of gold prices across multiple temporal resolutions. The framework emphasizes the systematic integration of metaheuristic optimization techniques with advanced learning models to address key challenges related to feature redundancy, hyperparameter sensitivity, and model generalization.

A distinguishing contribution of the study lies in the joint treatment of feature selection and hyperparameter optimization as complementary components of model enhancement rather than as isolated preprocessing steps. By embedding optimization mechanisms directly within the modeling pipeline, the proposed approach aims to improve both predictive stability and computational efficiency. Furthermore, the use of long-term, multi-timeframe financial data enables a thorough examination of model behavior across diverse economic environments, thereby strengthening the empirical relevance and applicability of the framework.

The remainder of this paper is organized to provide a coherent and systematic exposition of the proposed methodology and analytical framework. Following this introduction, the next section describes the dataset, data preprocessing procedures, machine learning models, and optimization techniques employed in the study. Subsequent sections present the experimental design and analytical evaluation, followed by a discussion that situates the findings within the broader context of financial forecasting research. The paper concludes with a summary of key insights and directions for future investigation.

2. LITERATURE REVIEW

Gold price forecasting has been extensively investigated in the literature due to gold's strategic importance as a safe-haven asset, an inflation hedge, and a core component of global financial markets. The XAU/USD exchange rate, in particular, acts as a high-frequency barometer of macro-financial stress because it embeds information about the U.S. dollar's strength, real interest rate expectations, liquidity conditions,

and shifts in global risk sentiment. These characteristics make gold price dynamics intrinsically nonlinear, strongly regime-dependent, and sensitive to abrupt shocks, thereby challenging standard assumptions of stationarity and linearity that underlie many classical time-series models. Consequently, the recent forecasting literature increasingly prioritizes robust modeling pipelines that (i) capture long- and short-memory temporal dependencies, (ii) exploit multivariate information from macroeconomic and market indicators, and (iii) incorporate mechanisms for adaptivity and uncertainty quantification to support real-world decision-making under risk.

A major strand of the literature is the systematic advancement of deep learning architectures for commodity and precious-metal price prediction. A broad comparative study evaluates multiple deep learning systems for forecasting crude oil, gold, and silver prices under different temporal window configurations [16]. Its findings show that forecasting performance depends jointly on the selected architecture, the target commodity, and the length of the historical input sequence. This evidence is particularly relevant to XAU/USD because it cautions against assuming that one network configuration will generalize uniformly across markets or horizons. Instead, model selection and temporal context should be evaluated together, especially when price series exhibit nonlinear dynamics, volatility clustering, and rapidly changing signal-to-noise ratios.

A complementary literature directly compares gated recurrent architectures for gold price prediction. A 2023 study evaluates Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) models on gold price data, demonstrating the practical suitability of both architectures for learning nonlinear temporal dependencies [17]. The comparison is methodologically useful because LSTM provides a more elaborate memory structure, whereas GRU offers a more compact parameterization that can reduce training complexity. The results reinforce the importance of selecting recurrent architectures empirically rather than solely on theoretical complexity, with evaluation based on consistent data partitions and error measures.

Beyond architecture selection, an important theme concerns feature enrichment and the integration of macro-financial drivers. A study comparing LSTM, Bi-LSTM, and MLP incorporates variables such as economic indicators, interest rates, currency exchange rates, and news emotion, thereby recognizing that gold prices respond to a complex information set rather than to their own lag structure alone. The reported results favor LSTM, indicating that gated recurrent memory mechanisms provide a practical advantage when fusing heterogeneous predictors into a unified temporal representation [18]. Importantly, this work also signals that, although bidirectionality can be beneficial in offline sequence labeling tasks, it may be less advantageous in forecasting settings where causal ordering and forward-only information constraints dominate deployment realities.

Heterogeneous ensemble learning provides another route to robust gold price prediction. A 2024 study combines structurally different learning algorithms within a unified ensemble so that the final forecast benefits from complementary representations of the same market series [19]. The reported results support the premise that aggregating diverse learners can re-

duce dependence on the weaknesses of any single model and improve predictive stability. For gold markets, this is particularly valuable because trend-dominated, range-bound, and high-volatility periods may favor different inductive biases and therefore different component learners.

A further maturation in the literature is the explicit emphasis on interpretability and model transparency, especially given the high-stakes nature of financial forecasting. The Temporal Fusion Transformer (TFT) study applied to the Indian gold market adopts an end-to-end preprocessing and tuning workflow, benchmarks against classical and neural baselines, and then leverages attention mechanisms and variable importance analysis to explain model behavior. The results indicate that TFT not only improves predictive performance relative to ARIMA, linear regression, and LSTM, but also provides an interpretable account of how macroeconomic indicators such as CPI and exchange rates contribute to predictions over time [20]. From an applied standpoint, this interpretability is valuable because it supports scenario analysis, policy assessment, and model governance, enabling stakeholders to scrutinize whether forecasts align with plausible economic narratives rather than merely minimizing error *ex post*.

Ensemble learning constitutes another major research direction, motivated by the observation that individual models often specialize in different regimes. A deep learning ensemble developed for financial time-series forecasting integrates complementary linear and nonlinear modeling components and reports improved accuracy and robustness relative to individual benchmarks [21]. The broader implication for gold forecasting is that ensemble strategies may reduce model risk by diversifying across inductive biases: one component can represent persistent linear structure while another captures nonlinear temporal behavior. Such combinations are especially attractive when the data-generating process changes across market regimes.

Hybrid deep learning architectures are also widely explored in gold-specific settings, where combining convolutional feature extraction with recurrent temporal modeling is often advantageous. In the Iranian market, a comparison between GRU and CNN-LSTM for 18-carat gold price forecasting finds that CNN-LSTM yields lower prediction errors, supporting the hypothesis that convolutional layers can efficiently extract robust short-horizon representations (e.g., local trends and micro-patterns) before recurrent layers model longer dependencies [22]. Relatedly, a CNN-BiGRU hybrid framework is proposed for gold forecasting and reported to outperform competing hybrid and standalone models across a broad error metric set, indicating that bidirectional gated recurrence may further enhance representation learning in contexts where training can exploit richer temporal context, even if deployment remains forward-looking [23]. Collectively, these hybrid deep networks reflect a design pattern in the literature: use convolution for denoising and local pattern extraction, and use gated recurrence to integrate information across time.

Evidence from other highly volatile financial assets also informs model design for gold. A comparative study of cryptocurrency price forecasting evaluates ensemble learning alongside deep learning methods and shows that predictive performance varies materially across algorithms and market conditions [24]. Although cryptocurrencies and gold dif-

fer economically, both can display pronounced nonlinearity, abrupt regime changes, and sensitivity to investor sentiment. The study therefore supports broad model comparison and out-of-sample validation rather than reliance on a single fashionable architecture.

A related decision-support direction combines numerical market histories with textual information. A stacking ensemble for stock-market prediction integrates time-series data and news headlines, allowing different base learners to capture distinct signals before a higher-level model combines their outputs [25]. This contribution is relevant to gold because XAU/USD reacts to both quantitative market conditions and rapidly arriving macroeconomic or geopolitical news. The results motivate multimodal ensembles that incorporate information sources with different frequencies and noise characteristics.

Another direction expands gold forecasting beyond purely price-based inputs by incorporating multimodal information such as news content and sentiment. A multimodal regression framework uses concept-based news representations, mood signals from specialized groups, and technical indicators, and employs recurrent convolutional and recurrent layers to model temporal dynamics across modalities. By explicitly addressing nonstationarity—via learning how predicted values drift around local averages—the model achieves substantial error reduction relative to baselines [26]. The methodological implication is that gold price movements may be driven not only by lagged price dynamics and macroeconomic variables, but also by information diffusion processes that unfold at different speeds across news topics, thereby requiring architectures designed for heterogeneous temporal response patterns.

The economic role of gold also depends on whether it behaves as a hedge or a safe haven under changing equity-market conditions. Evidence based on conditional comoments shows that this relationship is state-dependent rather than constant across markets and periods [27]. This finding matters for forecasting because the usefulness of equity indicators as predictors of gold may strengthen, weaken, or reverse during stress episodes. It therefore supports regime-aware and multivariate specifications that allow cross-market dependence to evolve over time.

Cross-asset evidence further demonstrates that gold's defensive properties are conditional on both the comparison asset and the crisis period. A study of oil, gold, and cryptocurrencies for the G7 equity markets compares their hedging, diversification, and safe-haven roles before and after the COVID-19 episode [28]. The changing results across assets and subperiods reinforce the need to model structural breaks explicitly. They also suggest that a gold forecasting framework should incorporate broader risk indicators rather than treating the gold series as an isolated process.

Chaos theory has also been integrated with deep learning to improve general time-series forecasting. A 2024 study develops a framework in which chaotic structure informs the representation and prediction of complex temporal processes, reporting gains across nonlinear forecasting tasks [29]. The methodological relevance to gold lies in recognizing that apparently irregular fluctuations may still contain structured dynamics that conventional linear models fail to capture. This connects naturally with broader discussions of sensitivity to

initial conditions, regime switching, and nonlinear dependence in financial markets.

Finally, methodological transfer from agricultural commodity markets underscores the value of systematic optimization. A comparative study applies several optimized machine learning techniques to commodity-price prediction and demonstrates that tuning strategy can materially affect performance across model families [30]. The broader lesson is that no single forecasting architecture universally dominates across commodities; results depend on the target series, feature set, optimization procedure, and validation design. This evidence supports the present study's emphasis on metaheuristic hyperparameter search rather than manual configuration.

Overall, the reviewed literature—spanning 15 studies—demonstrates a clear shift from isolated statistical forecasting toward deep, hybrid, ensemble, multimodal, and regime-aware frameworks. Across diverse markets and datasets, the evidence suggests that gold price prediction benefits from (i) architectures capable of learning nonlinear temporal dependencies, such as LSTM and GRU; (ii) ensemble designs that combine complementary model structures; (iii) multimodal and macroeconomic feature integration, including news, sentiment, CPI, and exchange rates; (iv) interpretable mechanisms such as attention and variable importance; and (v) validation procedures that account for changing market regimes. Taken together, these directions form a coherent foundation for advanced XAU/USD forecasting and for decision-support tools used in trading, hedging, and macro-financial analysis.

3. MATERIALS AND METHODS

3.1 Dataset Description

The empirical analysis conducted in this study is based on an extensive historical dataset of XAU/USD prices, which represents the value of gold quoted in United States dollars within international financial markets. The dataset spans from June 2004 to February 2025, providing more than twenty years of continuous market observations. This prolonged temporal coverage is particularly valuable for financial modeling, as it encompasses multiple economic cycles, shifts in monetary policy regimes, and episodes of pronounced market turbulence. Over this period, gold markets have been influenced by a wide range of macroeconomic forces, including inflationary and deflationary environments, financial crises, geopolitical tensions, and changes in global liquidity conditions. As such, the dataset offers a rich empirical foundation for developing forecasting models that are resilient to structural changes and evolving market dynamics.

A key feature of the dataset is its multi-timeframe structure, which captures gold price movements across both high-frequency and low-frequency temporal horizons. Intraday market behavior is represented through observations recorded at 5-minute, 15-minute, 30-minute, 1-hour, and 4-hour intervals, allowing the analysis to incorporate short-term price fluctuations and microstructural effects. In parallel, longer-term dynamics are captured through daily, weekly, and monthly data, which reflect broader macroeconomic trends and strategic investment behavior. This multi-resolution design acknowledges the fact that gold price formation mechanisms differ across time horizons, with short-term prices often driven

by liquidity and order flow, while long-term prices are more closely linked to macroeconomic fundamentals and investor expectations.

Across all timeframes, the dataset maintains a consistent set of financial attributes derived from standard market price information. Each observation includes a timestamp and corresponding calendar date, ensuring precise temporal ordering and facilitating alignment across different frequencies. The core price variables consist of the Open, High, Low, and Close prices, expressed in United States dollars, which collectively describe the intraperiod price range and final settlement price. These variables form the foundation of most financial time-series analyses and are widely used in both academic research and professional market practice. In addition to price information, tick volume is included as a proxy measure of trading activity and market participation. Although tick volume does not represent actual traded volume in decentralized markets, it has been shown to capture meaningful information regarding market intensity and liquidity conditions.

The target variable in this study is defined with reference to the gold closing price, which serves as a benchmark indicator in financial valuation, portfolio assessment, and forecasting applications. The closing price is particularly relevant because it reflects the consensus market valuation at the end of each trading interval and is commonly used in the computation of returns, risk measures, and performance metrics. Depending on the modeling configuration, the forecasting task may focus on predicting the future closing price directly or on modeling the directional movement of prices over a specified horizon. Both formulations are economically meaningful, as they support decision-making in contexts such as trading, hedging, and strategic asset allocation.

Importantly, the dataset spans a wide range of economic and financial regimes, including periods of relative stability, systemic crises, and transitional phases associated with changes in monetary policy. This diversity enhances the suitability of the data for training machine learning models that aim to generalize beyond specific historical episodes. By exposing models to heterogeneous market conditions, the study seeks to mitigate the risk of regime-dependent overfitting and to improve robustness across different economic environments.

For model development and evaluation, the dataset is partitioned into distinct training, validation, and testing subsets. The training subset is used to estimate model parameters and learn underlying data patterns, while the validation subset supports model selection and optimization processes. The testing subset is reserved exclusively for out-of-sample assessment, ensuring that model evaluation reflects genuine predictive capability rather than in-sample fitting. This partitioning strategy adheres to established best practices in financial machine learning and is designed to preserve temporal integrity while preventing information leakage across datasets.

Understanding the intraday dynamics of market activity is essential for analyzing liquidity patterns, volatility transmission, and trader behavior across different trading sessions. Market participation is not uniformly distributed throughout the day; instead, it exhibits pronounced temporal heterogeneity driven by overlapping trading hours, information releases, and institutional trading strategies. To provide an empirical illustration of these intraday variations, Figure 1 presents

the distribution of trading volume across the 24-hour trading cycle. The visualization highlights substantial differences in both central tendency and dispersion of volume at different hours, offering valuable insights into periods of heightened market engagement and elevated liquidity risk.

A comprehensive understanding of asset return behavior requires examining both their distributional characteristics and their temporal variation across trading days. Financial return series are well known for exhibiting non-normal features such as leptokurtosis, asymmetry, and clustering, which have important implications for risk assessment and predictive modeling. To capture these stylized facts from multiple statistical perspectives, Figure 2 provides an integrated visualization framework that combines kernel density estimation, swarm plots, violin plots, and box plots. Together, these graphical representations offer complementary insights into the overall density of returns, cross-day dispersion, central tendency, and the presence of extreme observations, thereby enabling a nuanced assessment of return dynamics and day-of-week effects.

Examining the relationship between opening and closing prices is fundamental for understanding intraday price formation, market efficiency, and short-term price dynamics in financial markets. The degree of alignment between these two price points reflects the extent to which new information is incorporated during the trading session and provides insights into intraday volatility and directional persistence. As illustrated in Figure 3, the scatter distribution of opening and closing prices reveals a strong linear association, suggesting high price continuity and limited intraday price dislocation over the observed period. Such behavior is indicative of relatively stable trading conditions and supports the use of opening prices as informative predictors in intraday and end-of-day forecasting frameworks.

3.2 Data Preprocessing

Prior to model training, a comprehensive data preprocessing pipeline is implemented to ensure data quality, numerical stability, and compatibility with deep learning architectures. Given the scale and complexity of the multi-timeframe financial dataset, preprocessing plays a critical role in facilitating effective learning and reliable forecasting performance.

The first stage of preprocessing addresses the presence of missing values and irregularities in the time series. Long-horizon financial datasets frequently contain gaps arising from market closures, holidays, or data acquisition inconsistencies. Such irregularities, if left untreated, can distort temporal dependencies and degrade model performance. Appropriate handling strategies are therefore employed to maintain continuity of the time series while preserving its underlying statistical characteristics.

A crucial preprocessing step involves the alignment of timestamps across multiple timeframes. Since observations are recorded at different temporal resolutions, careful synchronization is required to ensure that features derived from various timeframes correspond to the correct chronological context. This alignment process is essential when integrating information from intraday and longer-term data sources, as misalignment could introduce spurious relationships or lead to information leakage. By enforcing strict temporal consis-

tency, the preprocessing pipeline ensures that all model inputs reflect information available up to the forecasting point.

Feature scaling and normalization are applied to address disparities in the numerical ranges of price and volume variables. Financial time-series data often exhibit large variations in magnitude, particularly when long historical periods are considered. Such variations can hinder gradient-based optimization and slow model convergence. Through appropriate scaling techniques, input features are transformed to comparable ranges, thereby improving numerical stability and enhancing the efficiency of the learning process.

To capture temporal dependencies and dynamic market behavior, lagged features and rolling window representations are constructed from the original price and volume series. These transformations enable the models to incorporate historical context and to learn patterns related to momentum, volatility persistence, and trend evolution. Rolling window features, in particular, allow the extraction of localized statistical information that reflects recent market conditions, which is especially relevant in nonstationary financial environments.

Given the extensive use of overlapping price indicators and multiple timeframes, correlation analysis is conducted to identify redundant and highly correlated features. Excessive multicollinearity can reduce learning efficiency, increase computational burden, and complicate model interpretation. By systematically analyzing feature correlations, the preprocessing stage supports dimensionality reduction and prepares the data for subsequent feature selection procedures.

Finally, the processed data are organized into structured input formats suitable for deep learning models. This includes reshaping time-series data into sequential representations and ensuring consistency with the architectural requirements of the learning algorithms employed in the study. Through this carefully designed preprocessing pipeline, the raw financial data are transformed into a form that supports robust, scalable, and economically meaningful machine learning-based forecasting of XAU/USD prices.

3.3 Deep Learning Models

The selection of deep learning models for financial time-series forecasting requires careful consideration of the unique statistical and structural properties of market data. Financial time series, particularly those associated with commodity markets such as gold, exhibit pronounced nonlinearity, temporal dependence, volatility clustering, and regime shifts driven by macroeconomic and geopolitical factors. These characteristics necessitate modeling approaches that are capable of learning complex sequential patterns while remaining robust to noise and nonstationarity. In the context of XAU/USD price forecasting, the chosen models must effectively process multi-timeframe inputs, accommodate high-dimensional feature spaces, and generalize across heterogeneous market conditions.

Model selection in this study is guided by both theoretical suitability and empirical relevance within the financial forecasting literature. Priority is given to architectures that have demonstrated strong capabilities in sequential learning, representation extraction, and nonlinear function approximation. Rather than relying on a single modeling paradigm, a diverse set of deep learning models is employed to reflect different




Figure 1. Market Activity by Hour (Volume Distribution).




Figure 2. Statistical and Distributional Analysis of Asset Returns.

inductive biases and learning mechanisms. This diversity enables a comprehensive examination of how alternative architectures capture gold price dynamics and provides a robust baseline for subsequent optimization procedures.

The Deep Pyramid Recurrent Neural Network (DPRNN) is included due to its hierarchical recurrent structure, which is designed to capture temporal dependencies at multiple levels of abstraction. By organizing recurrent layers in a pyra-




Figure 3. Relationship Between Opening and Closing Prices.

midial fashion, the DPRNN progressively reduces temporal resolution while retaining essential sequential information. This hierarchical compression allows the model to integrate short-term fluctuations with longer-term trends, making it particularly suitable for multi-timeframe financial data. In the context of gold price forecasting, such an architecture facilitates the simultaneous modeling of intraday volatility and broader market movements. However, the depth and structural complexity of the DPRNN can increase sensitivity to architectural design choices and training dynamics, highlighting the importance of systematic configuration and optimization.

The Variational Autoencoder (VAE) is employed as a probabilistic learning framework capable of extracting latent representations from financial time series. By modeling the underlying data distribution through a latent variable structure, the VAE enables dimensionality reduction and noise filtering in a principled manner. In financial applications, these latent representations can capture hidden market states and structural patterns that are not immediately observable in raw price data. This characteristic is particularly valuable in gold markets, where prices often reflect unobserved macroeconomic expectations and investor sentiment. Nonetheless, because the primary objective of the VAE is reconstruction rather than direct prediction, its application to forecasting tasks involves a trade-off between representational richness and deterministic forecasting accuracy.

The Encoder–Attention LSTM (EA-LSTM) extends conventional recurrent architectures by integrating an attention mechanism that dynamically weights different segments of the input sequence. This mechanism allows the model to focus selectively on temporally relevant information, thereby enhancing its ability to capture salient market events and structural changes. In the context of XAU/USD forecasting, attention mechanisms are especially useful for identifying periods of heightened informational content, such as rapid price movements or shifts in volatility. Additionally, attention-enhanced models offer improved interpretability by providing insight

into which temporal features contribute most significantly to the learning process. However, the increased architectural complexity of EA-LSTM models can lead to higher computational demands and greater sensitivity to hyperparameter settings.

Graph Neural Networks (GNNs) are incorporated to explore the modeling of structured dependencies that extend beyond linear temporal sequences. Although financial time series are inherently ordered in time, relationships among features, timeframes, or derived indicators can exhibit complex interdependencies that are naturally represented in graph form. GNNs enable the propagation and aggregation of information across interconnected nodes, allowing the model to learn relational patterns that may not be captured by purely sequential architectures. In gold price forecasting, such relational modeling can provide a flexible framework for integrating heterogeneous sources of information. However, the effectiveness of GNNs is closely tied to the quality of graph construction, and inappropriate structural assumptions may limit their practical utility.

Convolutional Neural Networks (CNNs) are included due to their demonstrated effectiveness in extracting local patterns and hierarchical representations from structured data. When applied to time-series inputs, convolutional filters can identify recurring short-term motifs, such as local trends, reversals, or volatility bursts. These localized patterns often carry important information for financial forecasting, particularly in high-frequency market environments. CNNs also offer advantages in terms of parameter efficiency and parallel computation, which can be beneficial when processing large-scale financial datasets. Nevertheless, the inherently local receptive fields of CNNs may constrain their ability to capture long-range temporal dependencies unless combined with additional architectural components.

Bidirectional Recurrent Neural Networks (BiRNNs) are employed to leverage contextual information from both past and future directions within a sequence. By processing time-series data in forward and backward order, BiRNNs generate richer sequence representations that incorporate surrounding temporal context. This bidirectional perspective can enhance pattern recognition and improve the modeling of symmetric temporal structures. In financial applications, such representations can be useful for retrospective analysis and sequence modeling. However, from an operational forecasting standpoint, the reliance on future information imposes practical limitations, as future observations are not available at prediction time. Consequently, the applicability of BiRNNs in real-time financial decision-making requires careful consideration.

Taken together, the deep learning models selected in this study represent a broad spectrum of architectural approaches to financial time-series forecasting. Each model exhibits distinct strengths in terms of temporal modeling, feature representation, and computational characteristics, while also presenting specific limitations when applied to the complex dynamics of gold markets. By establishing this diverse set of baseline models, the study lays a rigorous foundation for subsequent enhancements through advanced optimization and feature selection strategies, thereby supporting a systematic investigation of improved forecasting methodologies for XAU/USD

prices.

3.4 Metaheuristic Algorithms

The rapid growth in the complexity of deep learning architectures has significantly increased the difficulty of configuring models for financial time-series forecasting tasks. In markets such as gold, where price dynamics are driven by nonlinear interactions among macroeconomic conditions, investor behavior, and liquidity fluctuations, the performance of forecasting models is highly sensitive to hyperparameter settings. These hyperparameters govern essential aspects of the learning process, including convergence speed, stability, and generalization capability. As a result, conventional trial-and-error or grid-based tuning approaches are no longer sufficient for identifying optimal configurations, particularly when large-scale, multi-timeframe datasets are involved. Within this context, metaheuristic optimization algorithms provide a flexible and efficient framework for automating the hyperparameter tuning process in deep learning-based financial forecasting systems.

Metaheuristic algorithms are designed to explore complex and high-dimensional search spaces using stochastic and population-based strategies. Unlike gradient-based optimization methods, metaheuristics do not rely on differentiability assumptions and are therefore well suited for optimizing objective functions that are noisy, nonconvex, or computationally expensive to evaluate. These characteristics are common in financial forecasting problems, where model evaluation often depends on sequential learning processes and where small parameter changes can lead to substantial differences in predictive behavior. By iteratively refining candidate solutions through adaptive search mechanisms, metaheuristics enable the systematic discovery of high-quality hyperparameter configurations that enhance model robustness and learning efficiency.

3.4.1 Role of Metaheuristics in Hyperparameter Optimization

In deep learning-based financial forecasting, hyperparameter optimization plays a critical role in shaping the bias–variance characteristics of predictive models. Hyperparameters such as learning rates, network depth, and regularization controls directly influence the trade-off between model flexibility and generalization. An overly flexible model may closely fit historical gold price movements but fail to generalize to new market conditions, while an overly constrained model may underrepresent the complex dynamics inherent in financial data. Metaheuristic algorithms address this challenge by enabling adaptive exploration of hyperparameter spaces, allowing the learning process to balance bias and variance in a data-driven manner.

Automated hyperparameter search is particularly important in financial contexts where datasets span multiple economic regimes and structural changes. Manual tuning strategies often rely on subjective judgment and limited experimentation, which may lead to suboptimal configurations that perform well only under specific conditions. Metaheuristic optimization mitigates this risk by maintaining a population of candidate solutions and continuously updating them based on performance feedback. Through mechanisms such as selection, mutation, and information sharing, these algorithms can

escape local optima and adapt to the underlying structure of the optimization landscape.

Computational efficiency is another crucial consideration in large-scale financial datasets. Deep learning models trained on multi-timeframe XAU/USD data require substantial computational resources, as each training cycle involves processing long sequences and high-dimensional inputs. Exhaustive search methods quickly become infeasible as the number of hyperparameters increases. Metaheuristic algorithms offer a computationally efficient alternative by focusing search efforts on promising regions of the parameter space and reducing redundant evaluations. This efficiency enables the practical deployment of automated optimization strategies even when computational budgets are constrained, thereby supporting scalable and reproducible financial forecasting workflows.

3.5 Greylag Goose Optimization (GGO) Algorithm

The Greylag Goose Optimization (GGO) algorithm is a population-based metaheuristic inspired by the collective foraging, migration, and vigilance behavior of greylag geese. The algorithm models the adaptive balance between exploration and exploitation observed in goose flocks, where individuals dynamically switch between wide-area searching and focused local refinement depending on environmental conditions and collective feedback. This balance is particularly well suited for hyperparameter optimization in deep learning-based financial forecasting, where the optimization landscape is nonlinear, multimodal, and sensitive to small parameter perturbations.

In GGO, each candidate solution represents a search agent, denoted as X_i , within the population of size n . The agents collectively explore the search space over a predefined number of iterations $t_{\max}$ with the objective of minimizing or maximizing a fitness function F_n , depending on the optimization goal. At each iteration, the population is divided into two adaptive groups: an exploration group (n_1) responsible for global search and an exploitation group (n_2) dedicated to local refinement around promising regions. The best-performing agent at iteration t is denoted as $X^*(t)$ and serves as a reference point for guided movements.

Exploration Phase. The exploration mechanism in GGO is designed to encourage diversity and global search during the optimization process. Exploration updates are applied primarily to agents in the exploration group when the iteration index satisfies the condition $t \bmod 2 = 0$. In this phase, agent movements are governed by adaptive coefficients A and C , which control step size and directional influence.

When the absolute value of A is less than one, agents are guided toward the current best solution according to

$$X(t+1) = X^*(t) - A \cdot |C \cdot X^*(t) - X(t)|, \quad (1)$$

which promotes directed exploration around promising regions while maintaining stochastic variability.

When $|A| \geq 1$, exploration is intensified through interactions among randomly selected agents, denoted as X_{Paddle1} , X_{Paddle2} , and X_{Paddle3} . A time-dependent control parameter z is com-

puted as

$$z = 1 - \left(\frac{t}{t_{\max}} \right)^2, \quad (2)$$

which gradually shifts the algorithm from exploration toward exploitation as iterations progress. The agent position is then updated using a weighted combination of neighboring solutions:

$$X(t+1) = w_1 X_{\text{Paddle1}} + z w_2 (X_{\text{Paddle2}} - X_{\text{Paddle3}}) + (1-z) w_3 (X - X_{\text{Paddle1}}). \quad (3)$$

This formulation enhances population diversity and enables long-range movements across the search space, which is essential for avoiding premature convergence in complex optimization landscapes.

Additionally, when stochastic control parameters satisfy alternative conditions, exploration is performed using a spiral-shaped movement inspired by goose flock dynamics:

$$X(t+1) = w_4 |X^*(t) - X(t)| e^{bl} \cos(2\pi l) + 2w_1 (r_4 + r_5) X^*(t), \quad (4)$$

where b and l regulate the spiral amplitude and frequency. This mechanism allows agents to explore the vicinity of elite solutions while preserving stochastic motion.

Exploitation Phase. The exploitation phase focuses on intensifying the search around high-quality solutions to refine convergence. Agents assigned to the exploitation group update their positions using multiple reference points, denoted as sentry agents X_{Sentry1} , X_{Sentry2} , and X_{Sentry3} . Candidate positions are computed as

$$X_k = X_{\text{Sentry}k} - A_k |C_k \cdot X_{\text{Sentry}k} - X|, \quad k \in \{1, 2, 3\}, \quad (5)$$

and the final position update is obtained by averaging these candidates:

$$X(t+1) = \bar{X}_i|_0^3. \quad (6)$$

This strategy enables localized refinement around multiple elite solutions, enhancing stability and convergence accuracy.

In cases where alternative conditions are met, exploitation agents update their positions according to

$$X(t+1) = X(t) + D(1+z)w(X - X_{\text{Flock1}}), \quad (7)$$

which models coordinated flock movement and reinforces convergence toward optimal regions.

A distinctive feature of GGO is its adaptive adjustment of exploration and exploitation group sizes. When the best fitness value remains unchanged for consecutive iterations, the algorithm increases the number of exploration agents while reducing the exploitation group. This adaptive mechanism enhances global search capability in stagnation scenarios and improves robustness when optimizing deep learning hyperparameters for nonstationary financial datasets.

Through the integration of these exploration and exploitation strategies, GGO provides a flexible and adaptive optimization framework capable of efficiently navigating complex, high-dimensional search spaces. This makes it particularly suitable for tuning deep learning models applied to financial time-series forecasting tasks such as XAU/USD price prediction.

3.5.1 State-of-the-Art Metaheuristic Optimization Algorithms

To ensure a comprehensive evaluation of optimization strategies, this study employs a diverse set of state-of-the-art metaheuristic algorithms, each characterized by distinct search behaviors and adaptive mechanisms. The use of multiple optimizers allows for the assessment of how different exploration and exploitation strategies influence hyperparameter tuning outcomes in financial forecasting applications.

The Genetic Algorithm (GA) is grounded in evolutionary theory and simulates the process of natural selection through iterative cycles of reproduction, crossover, and mutation. Candidate solutions are evaluated based on a fitness function, and higher-performing configurations are more likely to propagate to subsequent generations. In the context of deep learning hyperparameter optimization, GA supports global exploration of the search space while preserving solution diversity. This evolutionary process is particularly effective in avoiding premature convergence and in handling the complex, multimodal optimization landscapes commonly encountered in financial modeling tasks.

Greytag Goose Optimization (GGO) is inspired by the collective migratory behavior and social interactions of greytag geese. The algorithm emphasizes cooperation and information exchange among agents, enabling the population to converge toward high-quality solutions through shared exploration. GGO incorporates adaptive movement strategies that balance exploration of new regions with exploitation of promising areas. In financial forecasting applications, this cooperative behavior supports stable convergence and enhances robustness when optimizing deep learning hyperparameters under noisy and nonstationary conditions.

The Bat Algorithm (BA) draws inspiration from the echolocation capabilities of bats, modeling frequency modulation, loudness adjustment, and pulse emission rates to guide the search process. BA dynamically adapts its exploration and exploitation behavior based on feedback from candidate solutions, allowing it to efficiently navigate complex search spaces. This adaptability makes BA well suited for optimizing deep learning models whose performance landscapes may vary across different regions of the hyperparameter space.

Differential Evolution (DE) is a population-based optimization algorithm that generates new candidate solutions by combining the differences between randomly selected individuals. This vector-based mutation strategy enables efficient exploration of continuous search spaces and has been shown to perform well in high-dimensional optimization problems. In the context of financial forecasting, DE offers a simple yet powerful mechanism for tuning hyperparameters, providing a balance between exploration efficiency and algorithmic simplicity.

The Whale Optimization Algorithm (WOA) is motivated by the bubble-net hunting strategy of humpback whales. It employs adaptive mechanisms that alternate between encircling prey and spiral-shaped movements to explore and exploit the search space. This dual-mode search behavior allows WOA to dynamically adjust its focus based on the optimization stage, making it effective in navigating the nonlinear and multimodal fitness landscapes associated with deep learning hyperparameter tuning.

Multiverse Optimization (MVO) is inspired by cosmological theories and conceptualizes candidate solutions as universes connected through probabilistic exchange mechanisms. White holes, black holes, and wormholes facilitate information transfer among universes, enabling both exploration and exploitation of the search space. In financial forecasting applications, MVO provides a flexible and adaptive framework capable of handling complex objective functions and interacting hyperparameter effects.

Particle Swarm Optimization (PSO) models the collective movement of social organisms, such as bird flocks, through iterative updates of particle velocities and positions. Each particle adjusts its trajectory based on personal experience and shared knowledge from the swarm. PSO is widely recognized for its simplicity and rapid convergence, making it an attractive option for optimizing deep learning models in large-scale financial datasets where computational efficiency is paramount.

Stochastic Fractal Search (SFS) is inspired by fractal growth phenomena observed in natural systems. The algorithm combines diffusion processes with adaptive updating mechanisms to explore the search space in a stochastic yet structured manner. By maintaining diversity and encouraging exploration across multiple scales, SFS is effective in avoiding local optima and sustaining robust search behavior. This property is particularly valuable in financial forecasting, where noisy data and evolving market dynamics can complicate optimization efforts.

Together, these metaheuristic algorithms form a comprehensive optimization framework for enhancing deep learning-based gold price forecasting models. Their complementary search strategies and adaptive behaviors enable systematic and efficient exploration of hyperparameter spaces, supporting the development of robust and scalable financial forecasting systems suitable for complex and dynamic market environments.

4. EXPERIMENTAL RESULTS

4.1 Baseline Model Performance (Before Optimization)

This subsection evaluates the predictive performance of the baseline deep learning models prior to the application of any feature selection or metaheuristic-based hyperparameter optimization. The purpose of this analysis is to establish a quantitative benchmark that reflects the intrinsic forecasting capability of each model architecture when applied to XAU/USD price data under identical experimental conditions. Such a baseline assessment is essential for isolating the true impact of optimization strategies introduced in subsequent stages of the study.

The comparative results, reported in Table 1, reveal noticeable variation in model behavior across all evaluation metrics, underscoring the heterogeneous suitability of different deep learning architectures for gold price forecasting. Error-based metrics, correlation measures, and agreement indices are jointly considered to ensure a comprehensive assessment of predictive accuracy, bias, and robustness.

Among the evaluated baseline models, DPRNN exhibits the strongest overall performance. Specifically, DPRNN achieves an MSE of 0.0589 and an RMSE of 0.2426, indicating

relatively low dispersion between predicted and observed gold prices. The MAE value of 0.0802 further confirms the model's ability to limit absolute forecasting deviations, while the MBE of 0.0715 suggests a moderate but controlled systematic bias. From a statistical association perspective, DPRNN records a correlation coefficient of 0.778 and an R^2 value of 0.79, demonstrating a strong linear relationship and high explanatory power. Additionally, the model achieves an NSE of 0.804 and a Willmott Index (WI) of 0.774, reflecting a high level of agreement between predictions and observed values.

The VAE-based model demonstrates competitive but slightly weaker baseline performance relative to DPRNN. While the RMSE of 0.2419 is marginally lower than that of DPRNN, the MSE increases to 0.07195 and the MAE rises to 0.0935, indicating greater sensitivity to absolute prediction errors. The MBE value of 0.0806 suggests a higher degree of systematic deviation. In terms of correlation, the VAE model attains an r value of 0.754 and an R^2 of 0.767, reflecting a moderate reduction in explanatory capability. The NSE and WI values of 0.78 and 0.776, respectively, indicate acceptable agreement, though with reduced stability compared to DPRNN.

The EALSTM model exhibits moderate baseline performance, with an MSE of 0.08074 and an RMSE of 0.2681. The MAE of 0.1067 and MBE of 0.1352 reveal increased absolute and directional errors, suggesting that while the attention mechanism enhances temporal focus, it does not fully mitigate error accumulation over the forecasting horizon. The correlation coefficient of 0.753 and R^2 value of 0.765 indicate a reasonable level of association, though inferior to both DPRNN and VAE. Agreement-based metrics further decline, with an NSE of 0.769 and WI of 0.723, highlighting limitations in predictive consistency.


The GNN-based model shows a further deterioration in baseline forecasting performance. The model records an MSE of 0.16032 and an RMSE of 0.2841, accompanied by an MAE of 0.1458 and an MBE of 0.1521. These values indicate elevated prediction errors and systematic bias. Although the correlation coefficient remains relatively moderate at 0.738 and the R^2 value reaches 0.759, agreement metrics such as NSE (0.754) and WI (0.693) reveal reduced alignment with observed gold price movements. This outcome suggests that relational modeling alone may be insufficient for capturing dominant temporal dependencies in financial price series without extensive structural refinement.

CNN and BiRNN models demonstrate the weakest baseline performance among the evaluated architectures. The CNN model reports an MSE of 0.3489 and an RMSE of 0.5907, with MAE and MBE values increasing to 0.2156 and 0.2913, respectively. These results reflect substantial forecasting deviations and limited robustness. Correspondingly, the correlation coefficient declines to 0.724 and R^2 to 0.733, while NSE and WI fall to 0.733 and 0.659. The BiRNN model performs even less favorably, recording the highest error values with an MSE of 0.5714 and an RMSE of 0.7559. The MAE of 0.2504 and MBE of 0.3841 indicate pronounced absolute and directional errors. Despite achieving an r value of 0.704 and an R^2 of 0.714, the agreement metrics further deteriorate, with NSE and WI reaching only 0.711 and 0.657, respectively.

Overall, the baseline results clearly demonstrate that while

Table 1. Baseline performance comparison of deep learning models for XAU/USD price forecasting before optimization.


Model	MSE	RMSE	MAE	MBE	r	R^2	RRMSE	NSE	WI
DPRNN	0.0589	0.2426	0.0802	0.0715	0.778	0.790	6.08	0.804	0.774
VAE	0.07195	0.2419	0.0935	0.0806	0.754	0.767	6.29	0.780	0.776
EALSTM	0.08074	0.2681	0.1067	0.1352	0.753	0.765	6.43	0.769	0.723
GNN	0.16032	0.2841	0.1458	0.1521	0.738	0.759	6.55	0.754	0.693
CNN	0.3489	0.5907	0.2156	0.2913	0.724	0.733	6.86	0.733	0.659
BiRNN	0.5714	0.7559	0.2504	0.3841	0.704	0.714	7.18	0.711	0.657

**Figure 4.** Distribution of Model Performance Metrics Using Violin Plots with Jittered Observations.

certain architectures, particularly DPRNN and VAE, exhibit relatively strong inherent forecasting capability, none of the evaluated models achieve consistently low error and high agreement across all metrics. The observed performance dispersion highlights the sensitivity of deep learning models to architectural design and parameter configuration in financial time-series forecasting. These findings provide a strong empirical justification for the subsequent application of feature selection and metaheuristic-based hyperparameter optimization to enhance predictive accuracy, stability, and generalization in XAU/USD price forecasting.

A rigorous evaluation of predictive model performance requires assessing multiple error- and efficiency-based metrics simultaneously in order to capture both accuracy and robustness. Single-point summary statistics may obscure variability across evaluation runs or validation folds, thereby limiting the interpretability of comparative model performance. To address this, Figure 4 presents the distributional characteristics of key performance metrics using violin plots augmented with jittered data points. This visualization highlights the central tendency, dispersion, and density structure of error measures (such as mean squared error, root mean squared error, mean absolute error, and mean bias error) alongside goodness-of-fit and efficiency indicators, enabling a comprehensive and transparent assessment of model reliability and stability.

Comparative analysis of predictive models benefits from understanding not only their individual performance levels but also the structural similarities in their evaluation outcomes. Hierarchical clustering provides a systematic framework for identifying groups of models that exhibit comparable be-

**Figure 5.** Dendrogram for Hierarchical Clustering of Predictive Models.

havior across multiple performance metrics, thereby revealing latent relationships among modeling approaches. In this context, Figure 5 illustrates a dendrogram resulting from hierarchical clustering applied to the evaluated models. The branching structure and linkage distances reflect the degree of similarity between models, enabling a clear interpretation of clusters that share analogous predictive characteristics and offering insights into model complementarity and redundancy.

A detailed comparison of predictive models requires not only evaluating their average performance but also examining the variability and stability of their evaluation metrics across different configurations or validation instances. Error-based measures and goodness-of-fit indicators may exhibit heterogeneous dispersion patterns that are not fully captured by point estimates alone. To address this issue, Figure 6 presents a combined visualization framework in which box plots are overlaid with swarm plots for individual performance metrics, complemented by explicit annotations of the mean and standard deviation. This integrated representation facilitates a nuanced assessment of central tendency, dispersion, and outlier behavior across models, thereby supporting a more transparent interpretation of model accuracy, robustness, and comparative reliability.

Analyzing the joint distributional behavior of multiple performance metrics provides deeper insight into model accuracy, efficiency, and stability than isolated metric evaluations. Kernel density estimation offers a flexible, non-parametric approach for capturing the underlying probability structure of each metric, while stacked visualizations facilitate simultane-




Figure 6. Distribution of Individual Performance Metrics Across Models Using Box and Swarm Plots with Mean and Standard Deviation.




Figure 7. Stacked Kernel Density Estimation Streamgraph of Model Performance Metrics.

ous comparison across heterogeneous scales. In this regard, Figure 7 presents a stacked kernel density streamgraph of key error-based and goodness-of-fit metrics, enabling a holistic assessment of their relative concentration, overlap, and dispersion. This representation highlights dominant density regions and reveals potential trade-offs among metrics, thereby supporting a more integrated interpretation of predictive model performance.

4.2 Optimized Model Performance

This subsection evaluates the predictive performance of the optimized forecasting models obtained by integrating metaheuristic algorithms with the DPRNN architecture. The objective of this analysis is to quantify the extent to which metaheuristic-driven optimization enhances forecasting accuracy, stability, and agreement relative to the baseline models. By employing multiple optimization strategies under identical experimental conditions, this section provides a comprehensive assessment of how different metaheuristics influence the performance of deep learning models in XAU/USD price forecasting.

The optimized performance results are summarized in Table 2, which reports error-based metrics, correlation measures, and agreement indices for each metaheuristic–DPRNN combination. Across all evaluated configurations, substantial reductions in prediction error and systematic bias are observed, accompanied by marked improvements in correlation strength and overall model agreement.

Among the optimized models, the GGO+DPRNN configuration achieves the strongest overall performance. This model records an exceptionally low MSE of 2.05×10^{-5} and an RMSE of 4.52×10^{-3} , indicating a dramatic reduction in forecasting error magnitude. The MAE value of 2.05×10^{-4} further confirms the model's ability to maintain minimal ab-

solute deviations, while the MBE of 5.40×10^{-5} reflects negligible systematic bias. From a statistical association perspective, the GGO-optimized model attains a correlation coefficient of 0.987 and an R^2 value of 0.983, demonstrating an exceptionally strong alignment between predicted and observed gold prices. Agreement-based metrics further reinforce this outcome, with an NSE of 0.986 and a Willmott Index of 0.988, indicating near-perfect predictive consistency.

The GA+DPRNN model also demonstrates strong optimization benefits, achieving an MSE of 3.12×10^{-5} and an RMSE of 5.58×10^{-3} . The MAE and MBE values of 2.60×10^{-4} and 6.80×10^{-5} , respectively, indicate low absolute and directional errors. The correlation coefficient reaches 0.979, with an R^2 of 0.975, while the NSE and WI values of 0.978 and 0.981 confirm a high degree of agreement. Although slightly inferior to the GGO-based configuration, the GA-optimized model exhibits substantial improvements over all baseline results.

The BA+DPRNN configuration yields further competitive performance, with an MSE of 4.18×10^{-5} and an RMSE of 6.50×10^{-3} . The MAE of 3.00×10^{-4} and MBE of 8.20×10^{-5} reflect controlled error behavior, while the correlation coefficient of 0.973 and R^2 value of 0.97 indicate strong explanatory power. Agreement metrics remain high, with an NSE of 0.972 and WI of 0.977, confirming the effectiveness of bat-inspired optimization in refining deep learning hyperparameters for financial forecasting.

The DE+DPRNN model also exhibits notable performance gains, recording an MSE of 5.42×10^{-5} and an RMSE of 7.36×10^{-3} . The MAE and MBE values of 3.36×10^{-4} and 9.60×10^{-5} , respectively, indicate slightly increased error levels relative to GA- and BA-based configurations. Nonetheless, the model maintains a strong correlation coefficient of 0.966 and an R^2 of 0.963, supported by an NSE of 0.965 and WI of 0.973, reflecting robust predictive agreement.

The WOA+DPRNN configuration records an MSE of 6.98×10^{-5} and an RMSE of 8.38×10^{-3} , with corresponding MAE and MBE values of 3.76×10^{-4} and 1.13×10^{-4} . Although error levels increase slightly, the model retains a strong correlation coefficient of 0.959 and an R^2 of 0.957. The NSE and WI values of 0.958 and 0.968, respectively, indicate consistent predictive behavior and effective convergence.

The MVO+DPRNN model achieves an MSE of 8.88×10^{-5} and an RMSE of 9.44×10^{-3} . The MAE and MBE increase to 4.08×10^{-4} and 1.28×10^{-4} , while the correlation coefficient and R^2 values reach 0.952 and 0.95, respectively. Agreement metrics remain high, with an NSE of 0.95 and WI of 0.962, confirming the capability of multiverse-based optimization to enhance model stability and forecasting accuracy.

The PSO+DPRNN configuration yields an MSE of 1.08×10^{-4} and an RMSE of 1.04×10^{-2} . The MAE and MBE values of 4.36×10^{-4} and 1.43×10^{-4} reflect moderate error increases relative to other optimized models. Nonetheless, the correlation coefficient of 0.946 and R^2 of 0.944, together with an NSE of 0.944 and WI of 0.957, indicate substantial improvements over baseline performance.

Finally, the SFS+DPRNN model records an MSE of 1.33×10^{-4} and an RMSE of 1.15×10^{-2} , with MAE and MBE values of 4.63×10^{-4} and 1.58×10^{-4} . Although this con-




Figure 8. Parallel Coordinates Plot of Normalized Performance Metrics for Hybrid Predictive Models.

figuration exhibits the highest error among the optimized models, it still achieves a correlation coefficient of 0.939 and an R^2 of 0.937. The NSE and WI values of 0.937 and 0.952, respectively, confirm meaningful performance gains relative to all baseline architectures.

Overall, the optimized results presented in Table 2 demonstrate that integrating metaheuristic algorithms with DPRNN leads to substantial improvements in forecasting accuracy, bias reduction, and predictive agreement. The consistent enhancement across all optimization strategies highlights the effectiveness of metaheuristic-driven hyperparameter tuning in addressing the complexity and nonlinearity of gold price dynamics.


Evaluating predictive models across multiple performance criteria requires a visualization framework capable of simultaneously capturing trade-offs among error-based measures and goodness-of-fit indicators. Parallel coordinates plots are particularly well suited for this purpose, as they enable the comparison of high-dimensional performance profiles within a unified visual space. In this context, Figure 8 presents a parallel coordinates representation of normalized model performance metrics for different hybrid modeling configurations. Each polyline corresponds to a distinct model, allowing for a direct visual assessment of relative strengths, weaknesses, and overall balance across accuracy, bias, correlation, and efficiency metrics.

A comprehensive evaluation of predictive model performance necessitates examining both the central tendency and dispersion of multiple accuracy and efficiency metrics. While summary statistics provide a concise overview, distribution-based visualizations offer deeper insight into variability, robustness, and the presence of extreme values across competing models. In this respect, Figure 9 presents a box plot representation of key performance metrics overlaid with swarm plots of individual observations. This combined visualization enables a clear assessment of median performance, interquartile ranges, and point-level dispersion, thereby facilitating a nuanced comparison of error magnitudes, correlation strength, and overall model reliability.

Evaluating predictive models across a diverse set of performance indicators requires a metric-wise examination of both central tendency and dispersion in order to ensure robust and reliable conclusions. Different error, bias, and efficiency metrics may exhibit distinct distributional characteristics that reflect trade-offs between accuracy and stability. In this context, Figure 10 presents a series of horizontal box plots over-

Table 2. Performance comparison of metaheuristic-optimized DPRNN models for XAU/USD price forecasting.

Model	MSE	RMSE	MAE	MBE	r	R^2	RRMSE	NSE	WI
GGO + DPRNN	2.05×10^{-5}	4.52×10^{-3}	2.05×10^{-4}	5.40×10^{-5}	0.987	0.983	0.059	0.986	0.988
GA + DPRNN	3.12×10^{-5}	5.58×10^{-3}	2.60×10^{-4}	6.80×10^{-5}	0.979	0.975	0.081	0.978	0.981
BA + DPRNN	4.18×10^{-5}	6.50×10^{-3}	3.00×10^{-4}	8.20×10^{-5}	0.973	0.970	0.098	0.972	0.977
DE + DPRNN	5.42×10^{-5}	7.36×10^{-3}	3.36×10^{-4}	9.60×10^{-5}	0.966	0.963	0.128	0.965	0.973
WOA + DPRNN	6.98×10^{-5}	8.38×10^{-3}	3.76×10^{-4}	1.13×10^{-4}	0.959	0.957	0.178	0.958	0.968
MVO + DPRNN	8.88×10^{-5}	9.44×10^{-3}	4.08×10^{-4}	1.28×10^{-4}	0.952	0.950	0.228	0.950	0.962
PSO + DPRNN	1.08×10^{-4}	1.04×10^{-2}	4.36×10^{-4}	1.43×10^{-4}	0.946	0.944	0.278	0.944	0.957
SFS + DPRNN	1.33×10^{-4}	1.15×10^{-2}	4.63×10^{-4}	1.58×10^{-4}	0.939	0.937	0.408	0.937	0.952

**Figure 9.** Box Plot with Swarm Overlay of Model Performance Metrics.

laid with swarm plots for individual performance metrics. This visualization enables a detailed comparison of metric distributions across models, highlighting variability, median performance, and the presence of extreme values, thereby supporting a comprehensive assessment of model consistency and predictive reliability.

Assessing the relative effectiveness of competing predictive models requires a normalized comparison framework that highlights performance gains with respect to a common benchmark. Improvement ratios provide an intuitive and scale-independent measure of how closely each model approaches the best observed performance for a given metric. In this regard, Figure 11 presents an improvement ratio matrix in which each entry quantifies the relative performance of a model compared to the best-performing model for each evaluation criterion. This visualization enables a concise yet comprehensive comparison across error-based, correlation, and efficiency metrics, thereby facilitating the identification of dominant modeling strategies and systematic performance differentials.

5. DISCUSSION

The empirical findings of this study provide important insights into the behavior of gold prices in modern financial markets and the effectiveness of advanced forecasting frameworks when applied to long-horizon, multi-timeframe data. From a financial market perspective, the results highlight the intrinsic complexity of XAU/USD price dynamics, which are shaped by nonlinear interactions among macroeconomic conditions, investor expectations, and global risk sentiment.

The baseline analysis demonstrates that even sophisticated learning architectures exhibit substantial variation in predictive accuracy and stability when operating without systematic optimization, reflecting the challenges posed by nonstationarity, volatility clustering, and regime-dependent behavior in gold markets.

The pronounced improvement observed after optimization underscores the critical role of model configuration in financial forecasting applications. The optimized DPRNN-based models exhibit significantly reduced forecasting errors, enhanced correlation with observed prices, and stronger agreement metrics across all evaluated measures. These improvements suggest that optimization enables the model to more effectively capture both short-term market fluctuations and longer-term structural patterns embedded within the XAU/USD series. In economic terms, this enhanced predictive capability implies a more accurate representation of market information, which is essential for decision-making in trading, hedging, and portfolio allocation contexts. Improved forecast reliability can reduce estimation risk for investors and support more efficient pricing of gold as a financial asset.

The impact of optimization on forecasting accuracy is particularly relevant in the context of gold's role as a safe-haven asset. Gold prices often respond sharply to changes in macroeconomic uncertainty, monetary policy expectations, and financial stress. The ability of the optimized framework to substantially reduce bias and error indicates that the model is better equipped to adapt to such conditions and to reflect the underlying economic signals driving price movements. This adaptability is crucial for financial practitioners who rely on timely and accurate forecasts to manage exposure during periods of heightened volatility or systemic risk.

When compared with existing financial forecasting studies, the findings of this work are broadly consistent with the growing consensus that data-driven models outperform traditional econometric approaches in capturing nonlinear and regime-dependent market behavior. Prior research on commodity and foreign exchange forecasting has emphasized the limitations of linear models in environments characterized by frequent structural breaks and evolving market dynamics. The results presented here reinforce these observations by demonstrating that even among advanced learning architectures, performance gains are strongly contingent on appropriate optimization strategies. The substantial performance gap between baseline and optimized configurations suggests that methodological rigor in model tuning is as important as the




Figure 10. Horizontal Box Plots with Swarm Overlays for Individual Model Performance Metrics.




Figure 11. Improvement Ratio Matrix Relative to the Best Performance Achieved for Each Metric.

choice of forecasting architecture itself.

Despite these encouraging results, several limitations must be acknowledged. Financial markets are subject to sudden shocks arising from geopolitical events, policy announcements, and exogenous crises, which can induce abrupt structural changes in price behavior. While the long historical cov-

erage of the dataset allows the model to learn from multiple past regimes, no forecasting framework can fully anticipate unprecedented events or structural breaks that fundamentally alter market dynamics. Consequently, model performance may deteriorate during extreme events that fall outside the range of historical experience captured in the training data.

Another limitation relates to the assumption of temporal stability in the learned relationships. Although optimization improves generalization across observed regimes, financial markets remain adaptive systems in which behavioral patterns and institutional structures evolve over time. Changes in market microstructure, trading technology, or regulatory environments may affect the relevance of historical patterns for future forecasting. Addressing such challenges may require adaptive or online learning strategies that continuously update model parameters in response to new information.

Overall, the discussion highlights that while the proposed forecasting framework represents a significant advancement in modeling gold price dynamics, its effectiveness is inherently linked to the evolving nature of financial markets. Recognizing these limitations is essential for responsible application and interpretation of forecasting results in economic and financial decision-making contexts.

6. CONCLUSION AND FUTURE WORK

This study has presented a comprehensive investigation into the application of deep learning and metaheuristic optimization techniques for forecasting XAU/USD prices using long-term, multi-timeframe financial data. Gold occupies a central position in global financial markets as a safe-haven asset, an inflation hedge, and a portfolio diversification instrument. Accurate forecasting of gold prices is therefore of significant economic importance, with direct implications for investors, financial institutions, and policymakers operating in increasingly volatile and interconnected markets.

The empirical analysis demonstrates that baseline deep learning models, while capable of capturing certain nonlinear and temporal patterns in gold price dynamics, exhibit substantial variability in predictive accuracy, bias, and agreement. These limitations reflect the inherent complexity of financial time series, which are characterized by nonstationarity, regime shifts, and sensitivity to model configuration. The observed baseline performance underscores the challenges associated with applying deep learning architectures to financial forecasting tasks without systematic optimization.

By integrating metaheuristic optimization algorithms with deep learning models, the study establishes a robust framework for enhancing forecasting performance in the XAU/USD market. Metaheuristic-driven hyperparameter optimization enables automated exploration of complex configuration spaces, effectively addressing issues related to bias–variance trade-offs and convergence stability. The optimized models exhibit pronounced reductions in forecasting error, improved correlation with observed prices, and enhanced agreement metrics, indicating more reliable representation of gold price dynamics across diverse market conditions.

From an economic perspective, the findings highlight the value of advanced optimization techniques in improving the informational efficiency of forecasting models used in financial markets. More accurate and stable gold price predictions can support better-informed trading strategies, risk management practices, and portfolio allocation decisions. For institutional investors and asset managers, such improvements contribute to more effective hedging against inflation and systemic risk. At the macro-financial level, enhanced gold price forecasting may also provide policymakers with additional insights into market expectations, investor sentiment, and the credibility of monetary regimes.

Overall, this study contributes to the financial forecasting literature by demonstrating that the combination of deep learning and metaheuristic optimization constitutes a powerful methodological approach for modeling complex commodity price dynamics. The proposed framework is particularly well suited for long-horizon and multi-timeframe financial datasets, where traditional econometric methods often fall short in capturing nonlinear and regime-dependent behavior.

While the proposed framework delivers substantial methodological and empirical contributions, several avenues for future research remain open. One important direction involves the integration of additional macroeconomic and financial indicators, such as interest rates, inflation measures, exchange rate indices, and financial market volatility proxies. Incorporating such exogenous variables could further enhance the economic interpretability and predictive robustness of

gold price forecasting models, particularly during periods of macroeconomic stress.

Another promising avenue lies in extending the framework to real-time and high-frequency forecasting applications. As financial markets continue to evolve toward faster trading environments, the ability to deploy optimized deep learning models in near real-time settings becomes increasingly relevant. Future work may therefore focus on reducing computational overhead, improving model scalability, and assessing latency constraints associated with live market deployment.

The adaptability of the proposed optimization framework also enables its extension to other financial assets and markets. Applying similar methodologies to foreign exchange rates, equity indices, energy commodities, or fixed-income instruments would allow for a broader assessment of the generalizability of metaheuristic-optimized deep learning models across asset classes. Such extensions could contribute to the development of unified forecasting systems capable of supporting cross-asset portfolio management and systemic risk monitoring.

From a methodological standpoint, future research may explore hybrid optimization strategies that combine multiple metaheuristic algorithms or integrate reinforcement learning mechanisms to dynamically adjust optimization behavior in response to changing market conditions. These enhancements could further improve convergence efficiency and robustness in nonstationary financial environments.

Finally, future studies may investigate the integration of the proposed forecasting framework into intelligent financial decision-support systems. Embedding optimized predictive models within automated trading platforms, risk assessment tools, or policy analysis frameworks could facilitate more responsive and data-driven decision-making in complex financial markets. Such developments would strengthen the practical relevance of advanced forecasting methodologies and contribute to more resilient and efficient financial systems.

DATA AVAILABILITY

The dataset used in this study is publicly available on Kaggle at https://www.kaggle.com/datasets/novandraanugrah/xauusd-gold-price-historical-data-2004-2024?select=XAU_15m_data.csv.

REFERENCES

- [1] M. Ryan, S. Corbet, and L. Oxley, "Is gold always a safe haven?" *Finance Research Letters*, vol. 64, p. 105438, 2024.
- [2] I. Özkal, I. A. Özkan, and F. Başçiftçi, "Metaverse token price forecasting using artificial neural networks (ANNs) and adaptive neural fuzzy inference system (ANFIS)," *Neural Computing and Applications*, vol. 36, no. 7, pp. 3267–3290, 2024.
- [3] H. Kinatader, I. Gurrib, and T. Choudhury, "Navigating crises: Gold's role as a safe haven for U.S. sectors," *Finance Research Letters*, vol. 69, p. 106210, 2024.

- [4] W. Mensi, D. Maitra, R. Selmi, and X. V. Vo, "Extreme dependencies and spillovers between gold and stock markets: Evidence from MENA countries," *Financial Innovation*, vol. 9, p. 47, 2023.
- [5] H. Ben Ameur, F. Jamaani, and M. N. A. Abu Alfoul, "Examining the safe-haven and hedge capabilities of gold and cryptocurrencies: A GARCH and regression quantiles approach in geopolitical and market extremes," *Heliyon*, vol. 10, no. 22, p. e40400, 2024.
- [6] K. N. Konstantakis, P. Xidonas, P. G. Michaelides, and S. Goutte, "Gold and COVID-19: Uncovering the safe haven hypothesis with dynamic MSR modeling," *International Review of Financial Analysis*, vol. 89, p. 102858, 2023.
- [7] F. Yu, "Macroeconomic information, global economic policy uncertainty and gold futures return predictability," *Finance Research Letters*, vol. 55, p. 103789, 2023.
- [8] E. Feder-Sempach, P. Szczepocki, and J. Bogułębska, "Global uncertainty and potential shelters: Gold, bitcoin, and currencies as weak and strong safe havens for main world stock markets," *Financial Innovation*, vol. 10, p. 67, 2024.
- [9] U. A. Sheikh, M. Asad, Z. Ahmed, and U. Mukhtar, "Asymmetrical relationship between oil prices, gold prices, exchange rate, and stock prices during global financial crisis 2008: Evidence from pakistan," *Cogent Economics & Finance*, vol. 8, no. 1, p. 1757802, 2020.
- [10] O.-C. Chuang, R. Gupta, C. Pierdzioch, and B. Shu, "Financial uncertainty and gold market volatility: Evidence from a GARCH-MIDAS approach with variable selection," *Econometrics*, vol. 12, no. 4, p. 38, 2024.
- [11] A. A. Salisu, R. Gupta, S. Karmakar, and S. Das, "Forecasting output growth of advanced economies over eight centuries: The role of gold market volatility as a proxy of global uncertainty," *Resources Policy*, vol. 75, p. 102527, 2022.
- [12] A. Ampountolas, "Enhancing forecasting accuracy in commodity and financial markets: Insights from GARCH and SVR models," *International Journal of Financial Studies*, vol. 12, no. 3, p. 59, 2024.
- [13] F. Wang, J. Chang, W. Zuo, and W. Zhou, "Research on efficiency and multifractality of gold market under major events," *Fractal and Fractional*, vol. 8, no. 8, p. 488, 2024.
- [14] R. Gupta, S. Karmakar, and C. Pierdzioch, "Safe havens, machine learning, and the sources of geopolitical risk: A forecasting analysis using over a century of data," *Computational Economics*, vol. 64, no. 1, pp. 487–513, 2024.
- [15] A. A. Salisu, R. Gupta, S. Ntyikwe, and R. Demirel, "Gold and the global financial cycle," *Quantitative Finance and Economics*, vol. 7, no. 3, pp. 475–490, 2023.
- [16] P. Foroutan and S. Lahmiri, "Deep learning systems for forecasting the prices of crude oil and precious metals," *Financial Innovation*, vol. 10, no. 1, p. 111, 2024.
- [17] N. P. J. R. Dewi, N. L. W. S. R. Ginantra, I. W. A. S. Darma, and I. G. A. Indrawan, "Application of long short-term memory (LSTM) and gated recurrent unit (GRU) algorithm in gold price prediction," *Proceedings of the 2023 11th International Conference on Cyber and IT Service Management (CITSM)*, pp. 1–6, 2023.
- [18] R. Popli, I. Kansal, R. Kumar, V. Khullar, P. Sharma, and D. Parsad, "Gold price prediction in stock exchange using LSTM and Bi-LSTM multivariate deep learning approaches," *Computational Intelligence and Mathematical Applications*, pp. 1–5, 2024.
- [19] N. M. Trieu and N. T. Thinh, "A robust prediction for gold price using heterogeneous ensemble learning," *Proceedings of the 2024 13th International Conference on Control, Automation and Information Sciences (ICCAIS)*, pp. 1–5, 2024.
- [20] K. N. TP and S. Sultana, "Advanced techniques for daily gold price forecasting in india through statistical analysis and predictive modeling," *Proceedings of the 2024 8th International Conference on Computational System and Information Technology for Sustainable Solutions (CSITSS)*, pp. 1–6, 2024.
- [21] K. He, Q. Yang, L. Ji, J. Pan, and Y. Zou, "Financial time series forecasting with the deep learning ensemble model," *Mathematics*, vol. 11, no. 4, p. 1054, 2023.
- [22] A. H. Baradaran, M. Bohlouli, and M. R. J. Motlagh, "Decoding tomorrow's gold prices: A comparative study of GRU and CNN-LSTM in the iranian market," *Proceedings of the 2024 10th International Conference on Web Research (ICWR)*, pp. 329–334, 2024.
- [23] M. M. Billah and S. Das, "Analysis and prediction of gold price using CNN and Bi-GRU based neural network model," *Proceedings of the 2021 24th International Conference on Computer and Information Technology (ICCIT)*, pp. 1–6, 2021.
- [24] A. Bouteska, M. Z. Abedin, P. Hajek, and K. Yuan, "Cryptocurrency price forecasting—a comparative analysis of ensemble learning and deep learning methods," *International Review of Financial Analysis*, vol. 92, p. 103055, 2024.
- [25] R. Corizzo and J. Rosen, "Stock market prediction with time series data and news headlines: A stacking ensemble approach," *Journal of Intelligent Information Systems*, vol. 62, no. 1, pp. 27–56, 2024.
- [26] S. Anbaee Farimani, M. V. Jahan, and A. Milani Fard, "An adaptive multimodal learning model for financial market price prediction," *IEEE Access*, vol. 12, pp. 121 846–121 863, 2024.
- [27] L. Ming, P. Yang, and Q. Liu, "Is gold a hedge or a safe haven against stock markets? evidence from conditional comoments," *Journal of Empirical Finance*, vol. 74, p. 101439, 2023.

-
- [28] S. Tarchella, R. Khalfaoui, and S. Hammoudeh, “The safe haven, hedging, and diversification properties of oil, gold, and cryptocurrency for the G7 equity markets: Evidence from the pre- and post-COVID-19 periods,” *Research in International Business and Finance*, vol. 67, p. 102125, 2024.
- [29] B. Jia, H. Wu, and K. Guo, “Chaos theory meets deep learning: A new approach to time series forecasting,” *Expert Systems with Applications*, vol. 255, p. 124533, 2024.
- [30] M. Sari, S. Duran, H. Kutlu, B. Guloglu, and Z. Atik, “Various optimized machine learning techniques to predict agricultural commodity prices,” *Neural Computing and Applications*, vol. 36, no. 19, pp. 11 439–11 459, 2024.