# N-BEATS for Multi-Asset Forecasting: Superior Performance Without Recurrence or Convolution

Nguyen Hoang Ha[1], Thanh Manh Le[1], Cuong H. Nguyen Dinh[2], Khanh Hoang Bao[1], Hoa Tran Thai[1,3,✉]

[1]University of Sciences, Hue University, Viet Nam.
[2]Phu Xuan University, Viet Nam.
[3]University of Economics, Hue University, Viet Nam.
nguyenhoangha@hueuni.edu.vn, lmthanh@hueuni.edu.vn, cuong.nguyen@pxu.edu.vn, hbkhanh0706@gmail.com, tranthaihoa@hueuni.edu.vn

**Abstract.** Forecasting financial time series is challenging due to nonlinear dynamics and market volatility. This study evaluates the Neural Basis Expansion Analysis for Time Series (N-BEATS) model for predicting the prices of Bitcoin (BTC), Gold, Crude Oil (WTI), and Apple (AAPL). Its performance is compared with 1D-CNN, LSTM, GRU, Bi-LSTM, and hybrid CNN–LSTM model. Daily historical data from 2014 to 2024 were collected from Yahoo Finance through the Python library yfinance. Each model uses the previous 30 days of data to predict the next day's price. Experimental results show that N-BEATS achieves the highest predictive accuracy across all datasets, producing lower MSE, RMSE, and MAE, and higher $R^2$ values than competing models. The model captures both short-term variations and long-term trends in multiple financial markets. Visualization confirms that N-BEATS closely follows actual price movements, even during high-volatility periods. These results demonstrate that N-BEATS is an effective framework for financial forecasting without relying on recurrent or convolutional mechanisms. Future work will extend the approach to multi-step prediction and integrate macroeconomic indicators to enhance interpretability and generalization.
**Keywords:** N-BEATS, Time-series forecasting, Bitcoin, Gold, Crude Oil, Apple stock.

## 1. Introduction

Financial time-series forecasting is a critical topic in computational finance and machine learning with significant implications for investment strategies, risk management, and portfolio optimization. The nonlinear dynamics, structural breaks, and stochastic nature of financial data make accurate prediction a challenging task [1]. Traditional statistical models such as ARIMA and GARCH rely on linear assumptions and are often unable to capture the complex dependencies and volatility clustering in modern markets[2]. These classical econometric approaches, while foundational in financial analytics, have increasingly shown limitations when confronted with the inherent non-stationarity and regime shifts characteristic of contemporary financial markets [3].

Deep learning models, including recurrent and convolutional architectures, have recently shown superior ability to learn temporal dependencies and nonlinear mappings directly from raw data [4]. LSTM networks and their variants (GRU, Bi-LSTM) have become dominant paradigms for sequential data analysis in finance, demonstrating empirical advantages over traditional methods [5], [6]. However, most of these models, such as LSTM or GRU, depend on recurrent structures that may suffer from gradient instability, vanishing gradient problems, and limited interpretability [4], [3], which hamper their application in high-frequency trading and real-time forecasting scenarios requiring rapid inference times [7].

The N-BEATS model introduces a different approach based on backward and forward residual blocks. It decomposes the time series into trend and seasonality components without relying on recurrence or convolution [8], [9]. N-BEATS has achieved strong results in generic forecasting benchmarks and has demonstrated state-of-the-art performance in the M4 forecasting competition [10], but its performance in the financial domain, where price movements are highly irregular and volatile, remains underexplored. Recent research has shown the promise of Transformer-based architectures such as Informer and Autoformer for financial forecasting, which employ attention

✉ Coresponding author: tranthaihoa@hueuni.edu.vn.

mechanisms to capture period-based dependencies; however, these models often require substantial computational resources and large-scale datasets for effective training [11], [12].

This study investigates the effectiveness of N-BEATS in forecasting four representative financial assets: Bitcoin, Gold, Crude Oil, and Apple stock. The performance of N-BEATS is compared with five established deep-learning models: 1D-CNN, LSTM, GRU, Bi-LSTM, and a hybrid CNN–LSTM. The empirical results show that N-BEATS achieves the lowest forecasting errors and the highest coefficient of determination across four datasets in financial markets. The study also provides visualization analyses that confirm the model's stability and responsiveness to market fluctuations. Collectively, these findings establish N-BEATS as an accurate and architecture-agnostic framework for financial forecasting.

The remainder of this paper is structured as follows: Section 2 provides a comprehensive review of related work in deep learning for financial time-series forecasting. Section 3 details the methodology, including data description, preprocessing techniques, and the N-BEATS architecture. Section 4 presents the experimental setup, quantitative results, and comparative performance analysis across diverse asset classes. Finally, Section 5 concludes the study by summarizing key findings and suggesting directions for future research.

## 2. Related work

Financial time series forecasting has long remained a fundamental challenge in computational finance and econometrics due to the inherent nonlinearity, volatility, and chaotic nature of market data. While traditional statistical methods such as ARIMA and GARCH provided the initial foundation for market analysis [13], the paradigm has shifted significantly toward Deep Learning (DL). This transition is driven by the superior ability of DL models to automatically extract features and model complex nonlinear dependencies without the rigid a priori assumptions required by classical linear models [14].

Among deep learning architectures, Recurrent Neural Networks (RNNs) and their gated variants, particularly LSTM networks, have become the dominant standard for sequential data analysis . Numerous recent studies have empirically demonstrated the superiority of LSTM over traditional econometric benchmarks. For instance, Gülmez [5] established that optimized Deep LSTM networks significantly outperform classical regression methods in stock price prediction. Similarly, Ateeq et al. [6] developed an enhanced LSTM framework for Bitcoin forecasting, highlighting its robustness in handling the extreme volatility characteristic of cryptocurrency markets. Cho et al. [15] introduced the GRU as a more computationally efficient alternative to LSTM that preserves the ability to model temporal dependencies, while Bi-LSTM improves contextual representation by processing sequences in both forward and backward directions, albeit at the cost of increased computational complexity [16]. However, despite their popularity, pure recurrent architectures often suffer from computational inefficiencies due to sequential processing and are prone to the vanishing gradient problem when modeling long-term dependencies, necessitating the development of hybrid and convolutional alternatives.

To mitigate the limitations of recurrent networks, a substantial body of literature has focused on hybridizing CNNs with LSTMs. In this architecture, CNN layers are typically employed to extract local invariant features, while LSTM layers capture temporal dependencies. Recent work by Singh et al. [17] and Alam et al. [7] reported that such CNN-LSTM hybrid models yield statistically significant improvements in accuracy for portfolio management and stock market closing price prediction compared to single-architecture baselines. In the commodities sector, Mohsin et al. [11] successfully applied deep CNN architectures to crude oil forecasting, confirming the efficacy of convolutions in filtering noise and identifying short-term price patterns. Another prominent direction involves integrating signal decomposition techniques, such as CEEMDAN or Variational Mode Decomposition (VMD), with deep learning. Li et al. [12] proposed a CEEMDAN-Informer-LSTM hybrid, demonstrating that isolating high- and low-frequency components enhances forecasting precision for stock indices compared to standalone Transformer or LSTM models.

While hybrid models and Transformer-based architectures have gained traction due to their attention mechanisms, they often incur high computational costs and require massive datasets for training, which may not be optimal for all financial applications [18]. In this context, the N-BEATS model has emerged as a disruptive architecture. Unlike its predecessors, N-BEATS eliminates recurrence and convolution entirely, relying instead on a pure deep neural architecture based on backward and forward residual links. Seminal work by Oreshkin et al. [10] and subsequent validation by Kasprzyk et al. [8] in the energy sector demonstrated that N-BEATS outperforms both statistical and complex hybrid models in electricity load forecasting, primarily due to its interpretable decomposition of trend and seasonality components.

Despite its theoretical advantages, the application of N-BEATS in financial forecasting remains relatively underexplored and fragmented compared to the saturation of LSTM-based studies. Existing research attempting to utilize N-BEATS for finance often modifies the core architecture with complex auxiliary mechanisms rather than evaluating the robustness of the baseline model. For example, recent studies have combined N-BEATS with Generative Adversarial Networks (GANs) for risk awareness or integrated it into attention-based unified networks. Furthermore, the majority of current literature restricts its scope to single asset classes, focusing independently on cryptocurrencies, commodities, or specific equity markets [19], [20], [21]. There is a notable scarcity of comprehensive, cross-domain evaluations that test model generalizability across heterogeneous financial markets. Given the strong connectedness between assets like Bitcoin, Gold, Oil, and the USD identified by Köse et al. [22], a robust forecasting framework must demonstrate stability across these diverse volatility regimes.

This study addresses these gaps by conducting a systematic empirical evaluation of the pure N-BEATS architecture across four distinct asset classes: Cryptocurrency (Bitcoin), Precious Metals (Gold), Energy (Crude Oil), and Equities (Apple). Unlike prior works that prioritize architectural complexity through hybridization, this research investigates the hypothesis that the generic, interpretable backcast-forecast mechanism of N-BEATS is sufficient to achieve state-of-the-art performance. By benchmarking N-BEATS directly against established deep learning baselines (LSTM, GRU, CNN, and CNN-LSTM) under identical experimental conditions, this work provides evidence of its superior cross-domain generalization and filling the critical gap in the comparative evaluation of architecture-agnostic models in finance.

## 3. Methodology

### 3.1 Data Description

The study uses daily historical prices of four representative financial assets: Bitcoin (BTC), Gold, Crude Oil (WTI), and Apple Inc. (AAPL). Data were collected from Yahoo Finance via the Python library yfinance, covering the period from September 2014 to October 2024. Each dataset includes the variables Open, High, Low, Close, and Volume. Only the closing price was used for forecasting, as it reflects the final consensus of market value each day. All time series were normalized to the range [0, 1] using Min–Max scaling to ensure numerical stability and convergence during training, which is defined as Equation (1)

$$x_t^{norm} = \frac{x_t - x_{min}}{x_{max} - x_{min}} \tag{1}$$

where $x_t$ denotes the original value at time step $t$, and $x_{min}$ and $x_{max}$ represent the minimum and maximum values of the time series, respectively.

The data were divided into 70% for training, 15% for validation, and 15% for testing, maintaining chronological order to preserve temporal dependencies. A sliding window of 30 days was used as input to predict the next day's closing price. This configuration allows the model to learn short-term temporal patterns while remaining computationally efficient.

### 3.2 N-BEATS Architecture

The N-BEATS is a fully connected neural architecture composed of stacked residual blocks that perform backcast-forecast decomposition, enabling the model to learn trend and seasonality components without relying on temporal recurrence or convolution [23], as illustrated in **Fig. *1***.

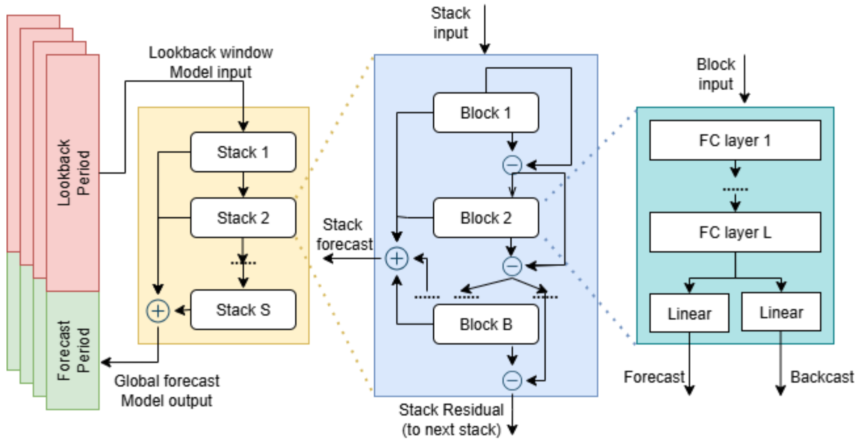


**Fig. 1.** Architecture of N-BEATS

Let $x \in R^T$ denote the input time series of length $T$ (lookback window), and let $\hat{y} \in R^H$ be the forecast vector with horizon $H$ (with $H = 1$ for one-step-ahead prediction). The N-BEATS architecture consists of $S$ stacked residual blocks, where each block refines the prediction by operating on the residuals of the previous block.

The input residual to block $(i-1)$ is defined in Equation (2).

$$r_{i-1} = x - \sum_{j=1}^{i-1} \hat{y}_j^b \tag{2}$$

with $r_0 = x$ for the first block.

Each block contains a stack of fully connected layers with ReLU activations The hidden representation produced by this stack is given by Equation (3).

$$h_i = FC_{stack}(r_{i-1}) \tag{3}$$

where $d$ is the hidden dimension.

The hidden vector $h_i$ is then projected into two sets of coefficients corresponding to the backcast and forecast components, as defined in Equations (4) and (5).

$$\theta_i^b = Dense\,(h_i) \in R^{P_b} \tag{4}$$

$$\theta_i^f = Dense\,(h_i) \in R^{P_f} \tag{5}$$

where $P_b$ and $P_f$ denote the numbers of basis coefficients for the backcast and forecast, respectively.

These coefficients are mapped back to the time domain through basis functions. For the generic N-BEATS variant used in this study, linear basis functions are employed, yielding the backcast and forecast signals in Equations (6) and (7).

$$\widehat{y_i^b} = G^b \theta_i^b \tag{6}$$

$$\widehat{y_i^f} = G^f \theta_i^f \tag{7}$$

where $G^b \in R^{T \times P_b}$ and $G^f \in R^{H \times P_f}$ are the basis matrices for backcast and forecast, respectively.

The residual is then updated by removing the estimated backcast, as defined in Equation (8).

$$r_i = r_{i-1} - \widehat{y_i^b} \tag{8}$$

The final prediction is obtained by aggregating the forecast outputs from all $S$ blocks, as shown in Equation (9).

$$\hat{y} = \sum_{i=1}^{S} \widehat{y_i^f} \tag{9}$$

For the interpretable N-BEATS variant, the basis functions can be constrained to represent specific components, such as polynomial bases for trends and trigonometric bases for seasonality. However, in the generic variant adopted in this work, all basis functions and coefficients are learned end-to-end without structural constraints. The model is trained by minimizing the Mean Absolute Error (MAE) loss, defined in Equation (10).

$$\mathcal{L} = \frac{1}{N}\sum_{n=1}^{N}\left|\hat{y}^{(n)} - y^{(n)}\right| \tag{10}$$

where $N$ is the number of training samples, $\hat{y}^{(n)}$ is the predicted value, and $y^{(n)}$ is the corresponding ground-truth target.

The N-BEATS model was implemented in PyTorch and trained using the Adam optimizer with a learning rate of $\alpha = 0.001$. An early stopping strategy based on validation loss with a patience of 15 epochs was applied to prevent overfitting.

### 3.3 Evaluation Metrics

Model performance was assessed using four standard metrics, which are defined as follows in Table 1.

**Table 1.** Definitions and interpretations of evaluation metrics used to assess model performance.

<table>
<tr><th>Metric</th><th>Formula</th><th></th></tr>
<tr><td>Mean Squared Error: (MSE)</td><td>$MSE = \frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2$</td><td rowspan="4">where<br>- $N$ is the total samples;<br>- $y$ and $\hat{y}$ are true and predicted values, respectively;<br>- $\underline{y} = \frac{1}{N}\sum_{i=0}^{N-1} y_i$.</td></tr>
<tr><td>Root Mean Squared Error: (RMSE)</td><td>$RMSE = \sqrt{MSE}$</td></tr>
<tr><td>Mean Absolute Error: (MAE)</td><td>$MAE = \frac{1}{n}\sum_{i=1}^{n}\lvert y_i - \hat{y}_i\rvert$</td></tr>
<tr><td>Coefficient of Determination: ($R^2$)</td><td>$R^2 = 1 - \frac{\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{n}(y_i - \bar{y}_i)^2}$</td></tr>
</table>

Lower MSE, RMSE, and MAE indicate better predictive accuracy, while $R^2$ closer to 1 indicates higher explanatory power.

## 4. Experiment

This study evaluates the effectiveness of the N-BEATS architecture in forecasting financial time series by benchmarking it against established deep learning models—specifically LSTM, GRU, Bi-LSTM, 1D-CNN, and a Hybrid CNN-LSTM. All experiments were conducted in a Google Colab environment using Python 3.10 and PyTorch 2.0, ensuring a standardized computational framework. To guarantee a rigorous comparison, each model was trained under identical conditions, utilizing a 30-day lookback window to predict the subsequent day's closing price for four distinct asset classes: Bitcoin (BTC), Gold, Crude Oil (WTI), and Apple Inc. (AAPL).

The datasets were partitioned chronologically into training (70%), validation (15%), and testing (15%) sets, with an early stopping mechanism based on validation loss implemented to mitigate overfitting.The quantitative results, summarized in Table 2, demonstrate that the N-BEATS architecture consistently outperforms all baseline models across standard metrics, including MSE, RMSE, MAE, and $R^2$.

In the high-volatility domain of cryptocurrency, N-BEATS achieved the lowest MSE of $3.30 \times 10^6$ and an $R^2$ of 0.9944. Visual analysis confirms that the model closely tracks rapid price reversals in Bitcoin, whereas recurrent architectures (LSTM, GRU) exhibit a noticeable lag, and convolutional models (1D-CNN, Hybrid) tend to overshoot during momentum shifts.

This superior performance extends to the commodities market; for Gold, which is characterized by more stable long-term trends, N-BEATS achieved an almost perfect fit ($R^2 = 0.9967$), significantly outperforming the 1D-CNN model, which yielded an MSE approximately 6.5 times higher.

The most critical performance disparity was observed in the Crude Oil dataset, a time series fraught with sharp structural breaks and irregular oscillations. Here, traditional recurrent networks failed to capture the underlying variance effectively, with LSTM and Bi-LSTM models dropping to $R^2$ values of 0.4814 and 0.4581, respectively. In contrast, N-BEATS maintained robust stability with an $R^2$ of 0.9761, avoiding the divergence often seen in RNNs when processing non-stationary sequences.

Finally, in the equity market (Apple Inc.), N-BEATS demonstrated high precision with the lowest residual error (MSE = 11.22) and an $R^2$ of 0.9975, indicating a tighter fit and better generalization compared to the slightly underfitting tendencies of the CNN and LSTM baselines.

**Table 2.** Performance comparison of forecasting models on BTC, Gold, Oil, and AAPL datasets.

| Type | Model | MSE | MAE | RMSE | $R^2$ |
|---|---|---|---|---|---|
| **BTC** | N-BEATS | 3.3 x $10^6$ | 12.21 x $10^2$ | 18.16 x $10^2$ | 0.9944 |
| | Bi-LSTM | 5.6 x $10^6$ | 17.18 x $10^2$ | 23.65 x $10^2$ | 0.9905 |
| | GRU | 7.19 x $10^6$ | 18.92 x $10^2$ | 26.81 x $10^2$ | 0.9877 |
| | LSTM | 1.31 x $10^6$ | 25.06 x $10^2$ | 36.18 x $10^2$ | 0.9777 |
| | 1D-CNN | 1.31 x $10^6$ | 26.42 x $10^2$ | 36.25 x $10^2$ | 0.9776 |
| | Hybrid | 1.95 x $10^6$ | 31.94 x $10^2$ | 44.13 x $10^2$ | 0.9668 |
| **AAPL** | N-BEATS | 11.22 | 2.19 | 3.35 | 0.9975 |
| | GRU | 13.43 | 2.38 | 3.67 | 0.9970 |
| | Bi-LSTM | 13.84 | 2.55 | 3.72 | 0.9970 |
| | LSTM | 37.45 | 4.17 | 6.12 | 0.9918 |
| | 1D-CNN | 50.72 | 5.18 | 7.12 | 0.9889 |
| | Hybrid | 53.15 | 5.14 | 7.29 | 0.9883 |
| **Gold** | N-BEATS | 5.81 x $10^2$ | 16.94 | 24.10 | 0.9967 |
| | LSTM | 7.18 x $10^2$ | 19.41 | 26.80 | 0.9959 |
| | GRU | 9.57 x $10^2$ | 21.82 | 30.94 | 0.9945 |
| | Bi-LSTM | 19.43 x $10^2$ | 34.00 | 44.08 | 0.9888 |
| | Hybrid | 29.11 x $10^2$ | 41.33 | 53.96 | 0.9832 |
| | 1D-CNN | 38.24 x $10^2$ | 46.76 | 61.84 | 0.9780 |
| **Oil** | N-BEATS | 4.45 | 1.53 | 2.11 | 0.9761 |
| | 1D-CNN | 43.38 | 5.52 | 6.59 | 0.7671 |
| | GRU | 72.32 | 7.90 | 8.50 | 0.6117 |
| | Hybrid | 87.56 | 7.54 | 9.36 | 0.5299 |
| | LSTM | 96.60 | 9.20 | 9.83 | 0.4814 |
| | Bi-LSTM | 100.93 | 9.36 | 10.05 | 0.4581 |

For all datasets, N-BEATS attains $R2$>0.97, showing that it captures most of the variance in observed prices. The low and steady residuals confirm the model's stable learning behavior and robust generalization across different financial environments.

Across all datasets, N-BEATS consistently achieved the best results, with $R2$ values exceeding 0.97 in every case. This confirms its strong explanatory capability and stable generalization under different market dynamics.

**Bitcoin (BTC)**

The cryptocurrency market is known for high volatility and non-stationary trends. As shown in Fig. 3, N-BEATS achieved the highest accuracy with an MSE of $3.30 \times 10^6$ and an $R^2$ of 0.99444.

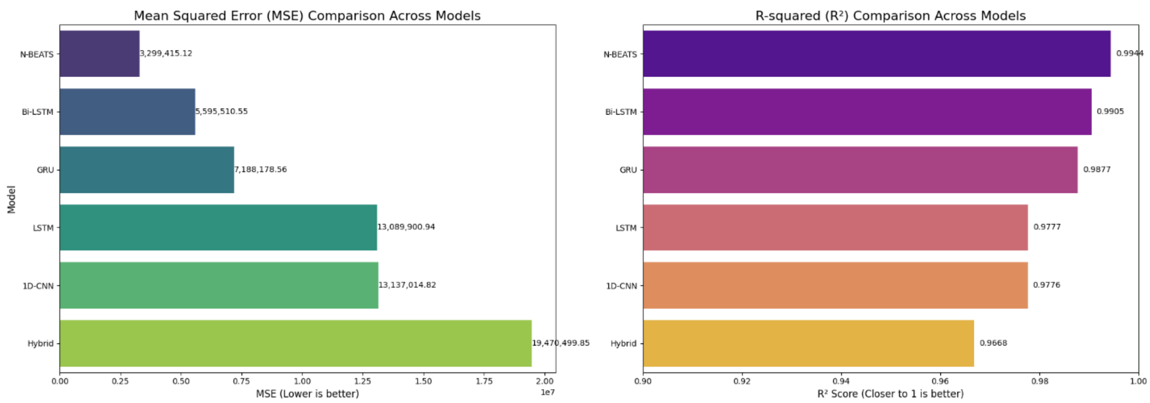


**Fig. 2**.Quantitative comparison of MSE and $R^2$ scores for Bitcoin forecasting. N-BEATS significantly outperforms the Hybrid and CNN baselines.

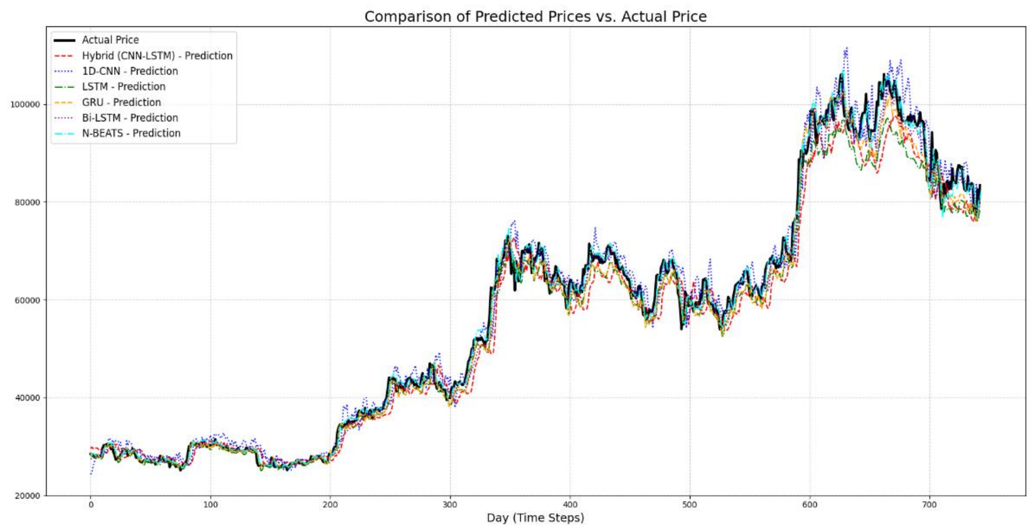


**Fig. 3.** Comparison of Predicted Prices vs. Actual Price for Bitcoin. The N-BEATS model (Cyan) closely follows the actual price (Black) during sharp peaks, whereas RNN-based models exhibit lag.

The Fig. 3 reveals that RNN-based models (LSTM, GRU) exhibit a "lag" phenomenon, reacting slowly to sudden price spikes. Conversely, convolutional models (1D-CNN) and the Hybrid CNN-LSTM tend to "overshoot," amplifying noise during high-volatility periods. N-BEATS maintains a balance, showing smoother convergence and rapid adaptation to momentum shifts due to its residual stacking architecture

**Gold**

For the Gold dataset, which exhibits more stable long-term trends compared to crypto, N-BEATS yielded the lowest MSE (581.02) and highest $R^2$ (0.9967), as shown in Fig. 5.

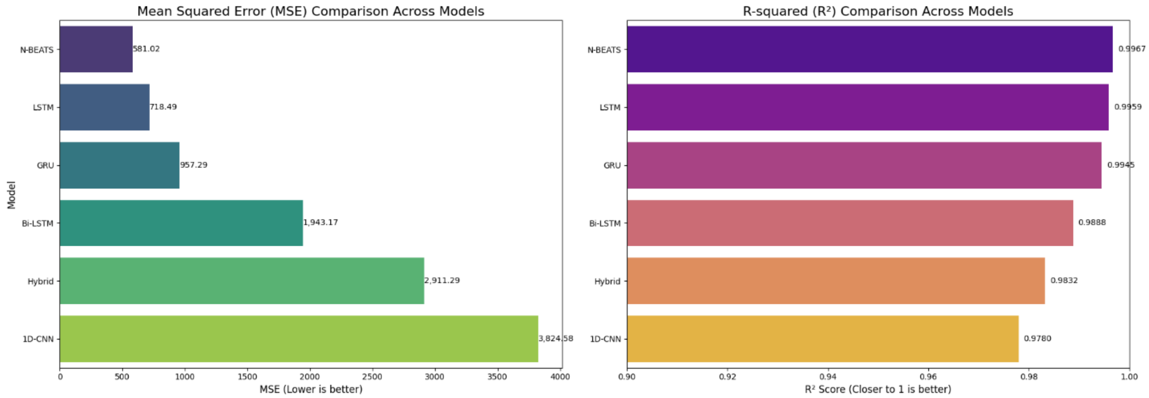


**Fig. 4.** Metric comparison for Gold. The error rates (MSE) for 1D-CNN and Hybrid models are nearly 5-6 times higher than that of N-BEATS.

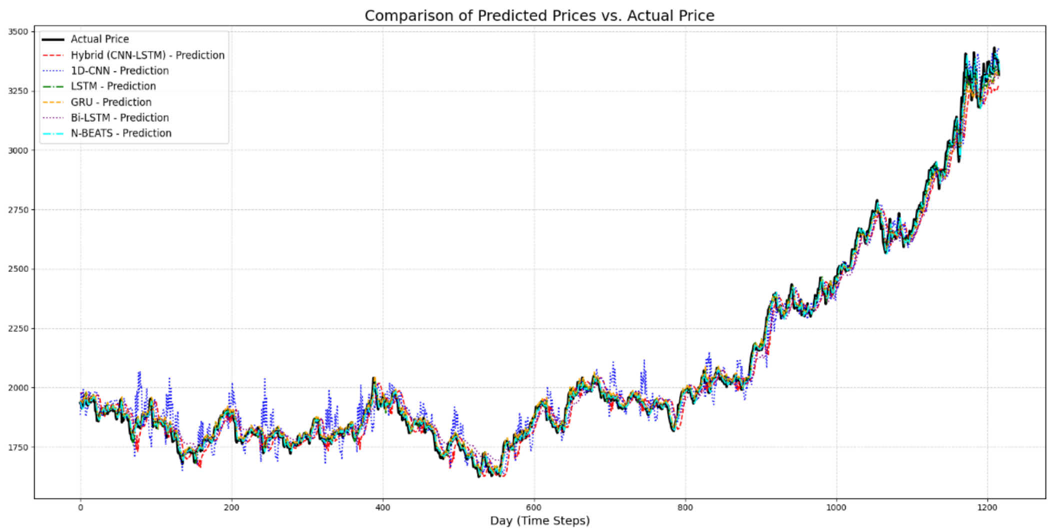


**Fig. 5.** Comparison of Predicted Prices vs. Actual Price for Gold. N-BEATS demonstrates superior tracking of the trend compared to the noisy predictions of the Hybrid model.

The superior performance of N-BEATS in this domain highlights its ability to model global temporal dependencies. While LSTM and GRU performed reasonably well ($R^2 > 0.99$), the Convolutional and Hybrid models struggled to recover subtle cyclical patterns, resulting in significantly higher RMSE values.

**Crude Oil (WTI)**

The Crude Oil dataset presented the most significant challenge due to sharp reversals and structural breaks. Here, the performance gap between N-BEATS and the baselines was most pronounced. N-BEATS achieved an $R^2$ of 0.9761, whereas the second-best model (1D-CNN) dropped to 0.7671, and recurrent models failed to capture the variance ($R^2 < 0.62$).

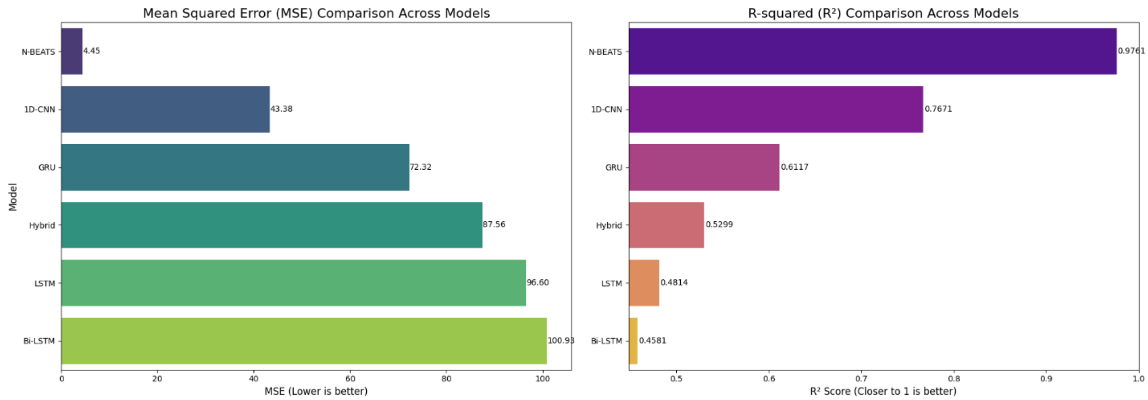


**Fig. 6.** Comparison of Predicted Prices vs. Actual Price for Crude Oil. Note the significant divergence of LSTM and Bi-LSTM (Green/Purple) from the actual price during the downward trend, while N-BEATS (Cyan) remains aligned.

As visualized in Fig. 7, N-BEATS adapts to irregular oscillations and avoids overfitting, whereas recurrent models diverge near sudden price spikes.

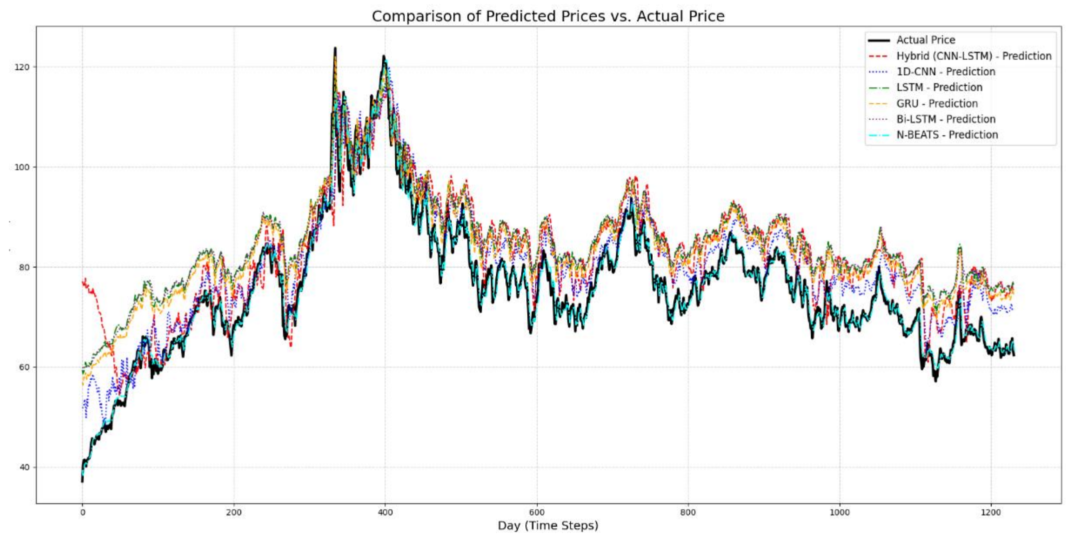


**Fig. 7.** Comparison of Predicted Prices vs. Actual Price for Crude Oil. Note the significant divergence of LSTM and Bi-LSTM (Green/Purple) from the actual price during the downward trend, while N-BEATS (Cyan) remains aligned.

The residual analysis further confirms stable learning without excessive variance, a critical factor for forecasting energy markets.

**Apple Stock (AAPL)**

The results for AAPL (Fig. 8) highlight the superior precision of N-BEATS in equity price forecasting. The model reaches an MSE of 11.22 and $R2$ = 0.9975, with small amplitude deviations between predicted and observed values. Other architectures show gradual error increases, particularly CNN and LSTM, which underfit slower-moving trends.

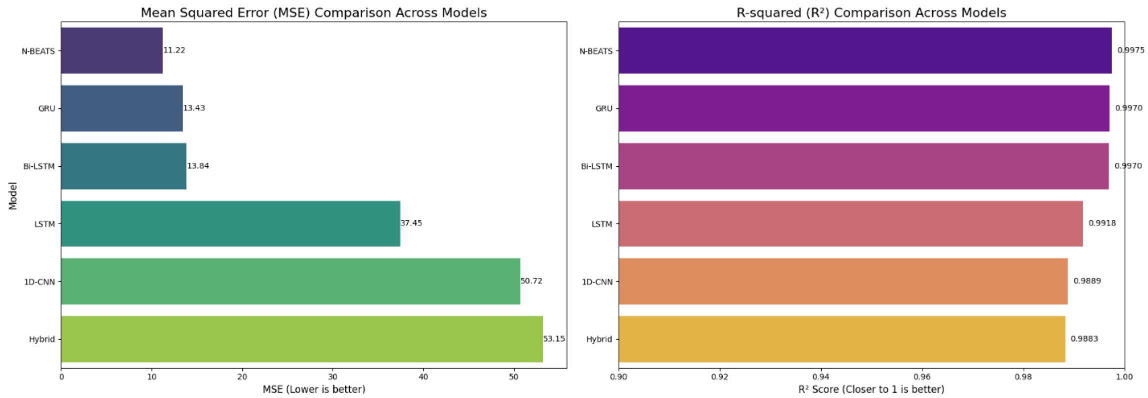


**Fig. 8.** Metric comparison for AAPL. All models performed relatively well, but N-BEATS consistently minimized the error margin.

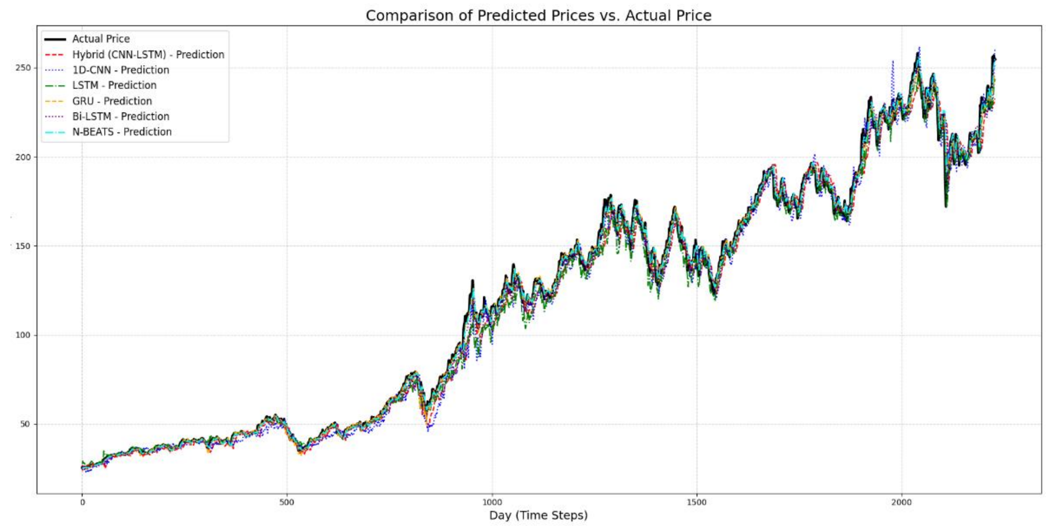


Fig. 9. Comparison of Predicted Prices vs. Actual Price for Apple Stock. N-BEATS exhibits minimal amplitude deviation from the ground truth.

The model's performance stability across asset classes indicates strong cross-domain generalization, enabling consistent accuracy for both volatile and smooth price series. The superior performance of N-BEATS is primarily due to its backcast–forecast residual design, which decomposes the input sequence into trend and seasonal components. This mechanism enables the model to generalize effectively across assets with distinct volatility regimes.

Unlike RNN-based models, N-BEATS does not rely on memory states, thus avoiding vanishing gradients and instability in long-sequence training. Compared with CNNs, it models global temporal dependencies rather than localized patterns, improving both short-term precision and long-horizon coherence.

The small and stable residuals observed across all assets confirm that N-BEATS learns underlying dynamics efficiently and converges without oscillation. Its consistent $R2 > 0.97$ indicates that it captures the majority of variance in the actual prices, achieving reliable forecasting across heterogeneous financial domains. This combination of accuracy, stability, and simplicity makes N-BEATS a strong candidate for integration into financial analytics pipelines and decision-support systems.

## 5 Conclusion and Future Work

This study assessed the effectiveness of the N-BEATS model in forecasting financial time series using daily data of Bitcoin, Gold, Crude Oil, and Apple stock collected from Yahoo Finance via yfinance. Each model used 30 days of historical data to predict the next day's price. Experimental results showed that N-BEATS consistently achieved the lowest MSE, RMSE, and MAE and the highest $R2$ (above 0.97) across all assets, outperforming 1D-CNN, LSTM, GRU, Bi-LSTM, and CNN–LSTM. The results confirm that N-BEATS effectively captures both short-term variations and long-term trends without relying on recurrent or convolutional structures, providing stable and accurate forecasts across diverse markets. Future work will extend this approach to multi-horizon prediction, integrate macroeconomic and sentiment variables, and apply explainable AI techniques to enhance model interpretability and practical use in financial analytics.

## Declarations

## Funding

This research was funded by University of Sciences, Hue University, Vietnam under the research project code ĐHKH2026D-002.

## Data availability

The datasets used in this study were collected from Yahoo Finance and are publicly available at: https://finance.yahoo.com/

## Code Availability

The code used to train and evaluate the models in this study is available upon request from the corresponding author.

## Conflict of interest

The authors declare that they have no competing interests relevant to the content of this article.

## References

[1] S. Giantsidi and C. Tarantola, "Deep learning for financial forecasting: A review of recent trends," *Int. Rev. Econ. Finance*, vol. 104, p. 104719, Dec. 2025, doi: 10.1016/j.iref.2025.104719.

[2] G. Sun and S. Deng, "Financial Time Series Forecasting: A Comparison Between Traditional Methods and AI-Driven Techniques," *J. Comput. Signal Syst. Res.*, vol. 2, no. 2, pp. 86–93, Mar. 2025, doi: 10.71222/339b9812.

[3] F. Spinnato, A. Ceni, A. Cossu, R. Guidotti, C. Gallicchio, and D. Bacciu, "Enhancing randomized recurrent neural networks with explainable attribution methods," *Neurocomputing*, vol. 666, p. 132318, Feb. 2026, doi: 10.1016/j.neucom.2025.132318.

[4] M. D. Nguyen, "Advanced investing with deep learning for risk-aligned portfolio optimization", doi: 10.1371/journal.pone.0330547.

[5] B. Gülmez, "Stock price prediction with optimized deep LSTM network with artificial rabbits optimization algorithm," *Expert Syst. Appl.*, vol. 227, p. 120346, Oct. 2023, doi: 10.1016/j.eswa.2023.120346.

[6] K. Ateeq, A. A. A. Zarooni, A. Rehman, and M. A. Khan, "A Mechanism for Bitcoin Price Forecasting using Deep Learning," *Int. J. Adv. Comput. Sci. Appl. IJACSA*, vol. 14, no. 8, Aug. 2023, doi: 10.14569/IJACSA.2023.0140849.

[7] K. Alam, M. H. Bhuiyan, I. U. Haque, M. F. Monir, and T. Ahmed, "Enhancing Stock Market Prediction: A Robust LSTM-DNN Model Analysis on 26 Real-Life Datasets," *IEEE Access*, vol. 12, pp. 122757–122768, 2024, doi: 10.1109/ACCESS.2024.3434524.

[8] M. Kasprzyk, P. Pełka, B. N. Oreshkin, and G. Dudek, "Enhanced N-BEATS for mid-term electricity demand forecasting," *Appl. Soft Comput.*, vol. 182, p. 113575, Oct. 2025, doi: 10.1016/j.asoc.2025.113575.

[9] B. N. Oreshkin, D. Carpov, N. Chapados, and Y. Bengio, "N-BEATS: Neural basis expansion analysis for interpretable time series forecasting," Feb. 20, 2020, *arXiv*: arXiv:1905.10437. doi: 10.48550/arXiv.1905.10437.

[10] B. N. Oreshkin, G. Dudek, P. Pełka, and E. Turkina, "N-BEATS neural network for mid-term electricity load forecasting," *Appl. Energy*, vol. 293, p. 116918, July 2021, doi: 10.1016/j.apenergy.2021.116918.

[11] M. Mohsin and F. Jamaani, "A novel deep-learning technique for forecasting oil price volatility using historical prices of five precious metals in context of green financing – A comparison of deep learning, machine learning, and statistical models," *Resour. Policy*, vol. 86, p. 104216, Oct. 2023, doi: 10.1016/j.resourpol.2023.104216.

[12] J.-C. Li, L.-P. Sun, X. Wu, and C. Tao, "Enhancing financial time series forecasting with hybrid Deep Learning: CEEMDAN-Informer-LSTM model," *Appl. Soft Comput.*, vol. 177, p. 113241, June 2025, doi: 10.1016/j.asoc.2025.113241.

[13] G. T. Wilson, "Time Series Analysis: Forecasting and Control, 5th Edition, by George E. P. Box, Gwilym M. Jenkins, Gregory C. Reinsel and Greta M. Ljung, 2015. Published by John Wiley and Sons Inc., Hoboken, New Jersey, pp. 712. ISBN: 978-1-118-67502-1," *J. Time Ser. Anal.*, vol. 37, no. 5, pp. 709–711, 2016, doi: 10.1111/jtsa.12194.

[14] N. Yemi Hussain, "Deep Learning Architectures Enabling Sophisticated Feature Extraction and Representation for Complex Data Analysis," *Int. J. Innov. Sci. Res. Technol. IJISRT*, pp. 2290–2300, Nov. 2024, doi: 10.38124/ijisrt/IJISRT24OCT1521.

[15] K. Cho *et al.*, "Learning Phrase Representations using RNN Encoder–Decoder for Statistical Machine Translation," in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, A. Moschitti, B. Pang, and W. Daelemans, Eds., Doha, Qatar: Association for Computational Linguistics, Oct. 2014, pp. 1724–1734. doi: 10.3115/v1/D14-1179.

[16] A. Aziz, Md. A. Hossain, A. N. Chy, Md. Z. Ullah, and M. Aono, "Leveraging contextual representations with BiLSTM-based regressor for lexical complexity prediction," *Nat. Lang. Process. J.*, vol. 5, p. 100039, Dec. 2023, doi: 10.1016/j.nlp.2023.100039.

[17] P. Singh, M. Jha, M. Sharaf, M. A. El-Meligy, and T. R. Gadekallu, "Harnessing a Hybrid CNN-LSTM Model for Portfolio Performance: A Case Study on Stock Selection and Optimization," *IEEE Access*, vol. 11, pp. 104000–104015, 2023, doi: 10.1109/ACCESS.2023.3317953.

[18] A. Hadizadeh, M. J. Tarokh, and M. M. Ghazani, "A novel transformer-based dual attention architecture for the prediction of financial time series," *J. King Saud Univ. Comput. Inf. Sci.*, vol. 37, no. 5, p. 72, June 2025, doi: 10.1007/s44443-025-00045-y.

[19] G. Asmat and K. M. Maiyama, "Bitcoin Price Prediction Using N-BEATs ML Technique," *EAI Endorsed Trans. Scalable Inf. Syst.*, vol. 12, no. 2, Apr. 2025, doi: 10.4108/eetsis.9006.

[20] X. Wang, C. Li, C. Yi, X. Xu, J. Wang, and Y. Zhang, "EcoForecast: An interpretable data-driven approach for short-term macroeconomic forecasting using N-BEATS neural network," *Eng. Appl. Artif. Intell.*, vol. 114, p. 105072, Sept. 2022, doi: 10.1016/j.engappai.2022.105072.

[21] B. S. Naik *et al.*, "Stock Price Forecasting using N-BEATS Deep Learning Architecture," *J. Sci. Res. Rep.*, vol. 30, no. 9, pp. 483–494, Sept. 2024, doi: 10.9734/jsrr/2024/v30i92373.

[22] N. Köse, Y. E. Gür, and E. Ünal, "Deep Learning and Machine Learning Insights Into the Global Economic Drivers of the Bitcoin Price," *J. Forecast.*, vol. 44, no. 5, pp. 1666–1698, 2025, doi: 10.1002/for.3258.

[23] B. Kaur *et al.*, "N-BEATS architecture for explainable forecasting of multi-dimensional poultry data," *PLOS ONE*, vol. 20, no. 4, p. e0320979, thg 4 2025, doi: 10.1371/journal.pone.0320979.