Applied Science Innovation (ASI)
Published by College of Physical and Applied Sciences, MOUAU

- *Received:* 03 May 2026
- *Revised*: 14 June 2026
- *Accepted:* 25 Jue 2026
- *Published:* 4 July 2026

FORECASTING NIGERIA'S INFLATION RATE: A COMPARATIVE ANALYSIS OF 1D-CNN AND LSTM DEEP LEARNING ALGORITHMS

Uchechukwu George[1*], Kingsley Uchendu[2], Theophilus O. Nelson[3], Kenneth K. Ibekwe[4], and Henry C. Igwegbe[5]

[12345]**Department of Statistics, Michael Okpara University of Agriculture, Umudike Abia State, Nigeria.**

***Corresponding author:** george.uchechukwu@mouau.edu.ng

ABSTRACT

This study compares the forecasting performance of a one-dimensional convolutional neural network (1D-CNN) and a long short-term memory (LSTM) network for predicting Nigeria's monthly headline inflation rate using data from January 2004 to December 2024 (252 observations). Models were trained on the 2004–2021 sub-sample and evaluated on the 2022–2024 hold-out period using mean squared error (MSE) and root mean squared error (RMSE). The LSTM achieved superior out-of-sample accuracy (testing RMSE = 10.21 vs. 18.52 for the 1D-CNN), attributable to its gating mechanisms for capturing long-term temporal dependencies. Both models effectively track the general inflation trajectory but struggle with the extreme volatility recorded in 2023–2024. Twelve-month-ahead forecasts for 2025 project a sharply decelerating inflationary path, with average annual rates of 13.40% (LSTM) and 5.89% (1D-CNN), implying continued but moderating monetary policy challenges. The principal limitation of the study is its univariate scope; exogenous drivers such as exchange rates, oil prices, and money supply are not incorporated, and the 252-month sample is modest relative to the scale typically favored for deep-learning training. Future work should extend the framework to multivariate LSTM variants and transformer-based architectures to improve generalizability. For monetary authorities, the LSTM's forecasts imply sustained but moderating price pressures throughout 2025, supporting a cautious approach to policy easing in the second half of the year. The findings advocate integrating LSTM models into central bank forecasting systems while retaining complementary traditional approaches for interpretability and robustness.

Keywords: inflation forecasting; Nigeria; 1D-CNN; LSTM; deep learning; time series; monetary policy; machine learning.

1. Introduction

Inflation is among the most closely watched macroeconomic indicators, with sustained deviations from target rates threatening purchasing power, investment stability, and aggregate welfare. Nigeria, Africa's largest economy by GDP, has experienced exceptionally turbulent price dynamics over the past two decades. From single-digit rates in the mid-2010s, headline inflation climbed to a historical high exceeding

60% by late 2024, driven by the removal of fuel subsidies, exchange-rate depreciation, and global commodity shocks (Central Bank of Nigeria [CBN], 2024).

The Central Bank of Nigeria (CBN) relies primarily on conventional econometric models, including autoregressive integrated moving average (ARIMA), vector autoregression (VAR), and error correction models (ECM), for inflation monitoring and projection. While these approaches are well-understood and interpretable, they rest on assumptions of linearity and stationarity that ill-suit the non-linear, structurally-breaking nature of Nigeria's inflation series. The repeated failure of traditional forecasts to anticipate the 2023–2024 inflationary surge underscores the need for methodological alternatives.

Machine learning (ML), particularly deep learning, offers a data-driven complement to conventional approaches. Convolutional neural networks (CNNs) and long short-term memory (LSTM) networks have demonstrated strong performance in financial and economic time-series tasks, including stock-price prediction (Majid et al., 2024; Fitrianingsih & Kharisudin, 2025), energy demand forecasting (Kharisudin et al., 2025; Sukatmo et al., 2025), and flood prediction (Atashi et al., 2023). Despite these extensive applications, comparative empirical evaluations of CNN and LSTM architectures specifically for inflation forecasting in sub-Saharan African emerging economies are conspicuously absent from the literature. This gap motivates the present study.

This study compares two deep learning architectures, 1D-CNNs and LSTMs, for forecasting Nigeria's inflation rate, asking which model better handles the temporal dependencies and volatility characteristics of the series, which hyperparameter settings work best in this emerging-market context, and what the resulting forecasts imply for monetary policy. To answer these questions, the study collects and preprocesses monthly inflation data from January 2004 to December 2024, and then builds, tunes, and evaluates both architectures using MSE and RMSE on a hold-out test set. The models are used to generate 12-month-ahead forecasts for 2025, with the ultimate goal of translating these forecasting results into concrete policy recommendations for Nigeria's monetary management.

The study is confined to Nigeria's official headline inflation series published by the CBN and the National Bureau of Statistics (NBS), covering January 2004 to December 2024. The analysis is univariate; exogenous drivers such as money supply, exchange rates, and oil prices are not incorporated. The 252-month sample, while adequate for training, is modest for deep learning at scale. Deep learning models are also not readily interpretable, limiting the economic narrative extractable from their internal representations. Future extensions should explore multivariate LSTM variants, Bayesian neural networks, and transformer-based architectures.

The remainder of the paper is organized as follows: Section 2 reviews related literature; Section 3 describes data and methodology; Section 4 presents results and discussion; Section 5 concludes with policy recommendations and directions for future research.

2. Literature Review

The application of deep learning to economic and financial time-series forecasting has grown rapidly. Collectively, the empirical record reveals two recurring findings: LSTM networks tend to excel at capturing long-range temporal dependencies and adapting to structural breaks, while 1D-CNNs offer competitive local pattern extraction at substantially lower computational cost. What remains absent is a head-to-head comparison focused on macroeconomic inflation forecasting in sub-Saharan African economies, the gap this study addresses.

2.1 LSTM-Based Forecasting

Atashi et al. (2023) compared LSTM and 1D-CNN for real-time flood prediction along the Red River of the North using streamflow records from 1994 to 2021. LSTM achieved an RMSE of 76.23 cfs against the 1D-CNN's 193.09 cfs, demonstrating the recurrent architecture's superior capacity to model complex hydrological dynamics. Though the domain differs, this performance gap corroborates LSTM's advantage in highly non-stationary series, a characteristic shared by Nigeria's inflation data.

Majid et al. (2024) benchmarked LSTM and Bidirectional LSTM (BiLSTM) for forecasting the Indonesia Stock Exchange Composite Index. The BiLSTM (6-9-1 configuration) achieved a MAPE of 0.572%, outperforming the best LSTM model (0.743%), underscoring the advantage of bidirectional context for capturing dependencies in both temporal directions.

Kharisudin et al. (2025) evaluated Vanilla LSTM, Stacked LSTM, BiLSTM, and CNN-LSTM for short-term electricity load forecasting using five years of daily data from the Central Java–DIY system. The CNN-LSTM hybrid delivered the best accuracy across MSE, MAE, RMSE, MAPE, and R^2, while offering a favourable trade-off between training speed and predictive precision.

2.2 1D-CNN-Based Forecasting

Kalange (2025) applied a 1D-CNN to forecast next-day closing prices for five major stocks on the Bombay Stock Exchange (2015–2023). With Min-Max normalisation and a 60-day sliding window, the model achieved MAPE values of 1.5–4.2%, outperforming ARIMA and providing competitive accuracy relative to LSTM at lower computational cost.

Namdari et al. (2025) tested 1D-CNN against DFNN, Simple RNN, LSTM, and GRU for short-term urban water demand forecasting. The 1D-CNN attained the best MAPE (3.52%) while requiring substantially less training time than recurrent models, highlighting its suitability for resource-constrained operational settings.

Listyanto and Yustanti (2025) compared CNN, MLP, and LSTM for daily iron-sales forecasting across an 80/20 train–test split. CNN outperformed both alternatives, achieving MAE = 0.293, RMSE = 0.357, MAPE = 0.081, and $R^2 = 0.9989$, confirming strong pattern extraction from short, univariate retail series.

2.3 Hybrid CNN–LSTM Models

Sayed et al. (2024) proposed a hybrid CNN–LSTM autoencoder for short-term photovoltaic power forecasting. The combined architecture improved RMSE and MAE by 5–25% relative to standalone CNN, GRU, and LSTM baselines. Thai et al. (2025) developed a multi-window 1D-CNN–LSTM framework for cryptocurrency price forecasting that consistently outperformed standalone architectures across medium- and long-term horizons. Fitrianingsih and Kharisudin (2025) employed a CNN-BiLSTM with an attention mechanism, delivering superior accuracy on MAE, MAPE, RMSE, and R^2. Sukatmo et al. (2025) showed that while LSTM achieved marginally lower MSE (3.35×10^{-4} vs. 4.68×10^{-4}), the 1D-CNN trained approximately 84% faster, a practically relevant trade-off for real-time applications.

Collectively, the literature indicates that LSTM excels at capturing long-range temporal dependencies, while 1D-CNN offers computational efficiency. Head-to-head comparisons in macroeconomic inflation forecasting for sub-Saharan African economies are absent from this body of work, motivating the present study.

3. Materials and Methods

3.1 Data

The dataset comprises official monthly headline inflation rates for Nigeria from January 2004 to December 2024 (252 observations), sourced from the Central Bank of Nigeria (CBN) statistical bulletin and the National Bureau of Statistics (NBS). Headline inflation measures the year-on-year percentage change in the Consumer Price Index (CPI) for all items. No missing values were identified in the series.

3.2 Analytical Framework

The analysis proceeds in three stages (Figure 1): (i) data acquisition and preprocessing; (ii) model construction and training; and (iii) evaluation, forecasting, and comparison. Two deep learning models, a 1D-CNN and an LSTM are compared using MSE and RMSE on a held-out test sample. The model with the lower RMSE on the test set is deemed the superior forecasting tool.

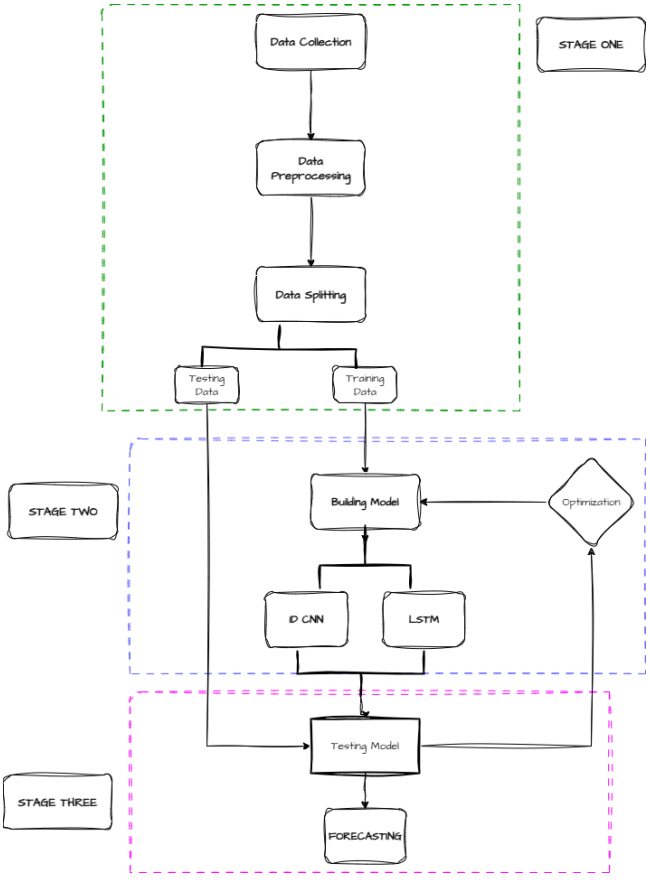


Figure 1. *Three-stage methodology framework for model development and evaluation.*

3.3 Data Splitting

3.3.1 Chronological Train–Test Split

To preserve temporal ordering and avoid look-ahead bias, data are partitioned chronologically at an 80:20 ratio:

- Training set: January 2004–December 2021 (216 observations, 80%)
- Test set: January 2022–December 2024 (36 observations, 20%)

3.3.2 Sliding-Window Feature Construction

Supervised learning inputs are created via a sliding-window procedure. Given a window size ω, each sample consists of ω consecutive observations as features (X) and the immediately following value as the target (Y). A window size of ω = 12 was selected to embed annual seasonal structure into the model inputs, consistent with practice for monthly macroeconomic series (Namdari et al., 2025; Thai et al., 2025). Table 1 illustrates the procedure for an illustrative series of length 21.

Table 1. Sliding-window feature construction (illustrative example, ω = 3)

S/N	X-Features (inputs)	Y-Label (target)
1	0, 1, 2	3
2	1, 2, 3	4
3	2, 3, 4	5
…	…	…
18	17, 18, 19	20

Note. For a series of length T and window size ω, the resulting design matrix has shape (T − ω) × ω.

3.3.3 Normalisation

Min-Max scaling transforms all observations to the interval [0, 1] prior to training, preventing gradient instability: $x' = (x - x_{min})\frac{1}{(x_{max} - x_{min})}$

where x is the raw observation, $x_{\min}$ and $x_{\max}$ are the series minimum and maximum, and ω' ∈ [0, 1] is the scaled value. Inverse scaling is applied to model outputs to recover percentage-point units. No differencing was applied; the models operate on scaled inflation levels to preserve level information needed for multi-step forecasting.

3.4 Model Architectures

3.4.1 One-Dimensional Convolutional Neural Network (1D-CNN)

A 1D-CNN applies learned convolutional filters along the time axis of sequential data to extract local temporal patterns. The convolution operation at position n is:

$$Y_n = \sum_{k=0}^{K-1} x(n-k)h_k \quad (1)$$

where x is the input signal, h_k is the kth filter weight, and K is the kernel length.

The ReLU activation function (Nair & Hinton, 2010) introduces non-linearity after each convolutional layer:

$$ReLU(z) = max(0, z) \quad (2)$$

Because the forecasting task is regression, the output layer uses a linear activation and produces a single scalar. The 1D-CNN architectural layout is shown in Figure 2.

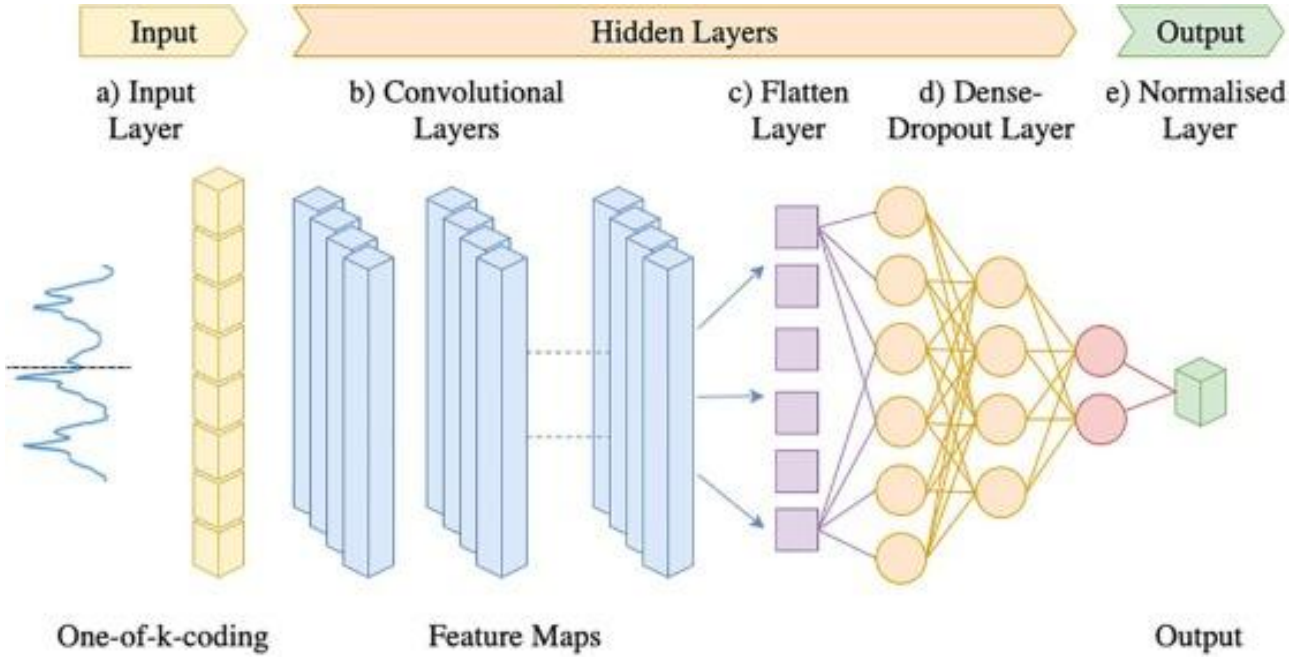


Figure 2. *1D-CNN architectural layout (adapted from Sánchez-Reolid et al., 2022).*

3.4.2 Long Short-Term Memory (LSTM)

The LSTM network (Hochreiter & Schmidhuber, 1997) is a gated recurrent architecture designed to learn long-range dependencies while mitigating the vanishing-gradient problem. It maintains a cell state C_t carrying information across time steps, modulated by three learnable gates:

Forget gate: determines what information from the previous cell state C_{t-1} is discarded:

$$f_t = \sigma\left(W_f \cdot [h_{t-1}, x_t] + b_f\right) \quad (3)$$

Input gate: controls what new information enters the cell state:

$$i_t = \sigma(W_i \cdot [h_{t-1}, x_t] + b_i) \quad (4)$$

$$C_t = tanh(W_C \cdot [h_{t-1}, x_t] + b_C) \quad (5)$$

Cell state update: integrates the forget and input gate signals:

$$C_t = f_t \odot C_{t-1} + i_t \odot C_t \quad (6)$$

Output gate: determines the hidden state passed to the next time step:

$$o_t = \sigma(W_o \cdot [h_{t-1}, x_t] + b_o) \quad (7)$$

$$h_t = o_t \odot tanh(C_t) \quad (8)$$

The LSTM architectural layout is shown in Figure 3.

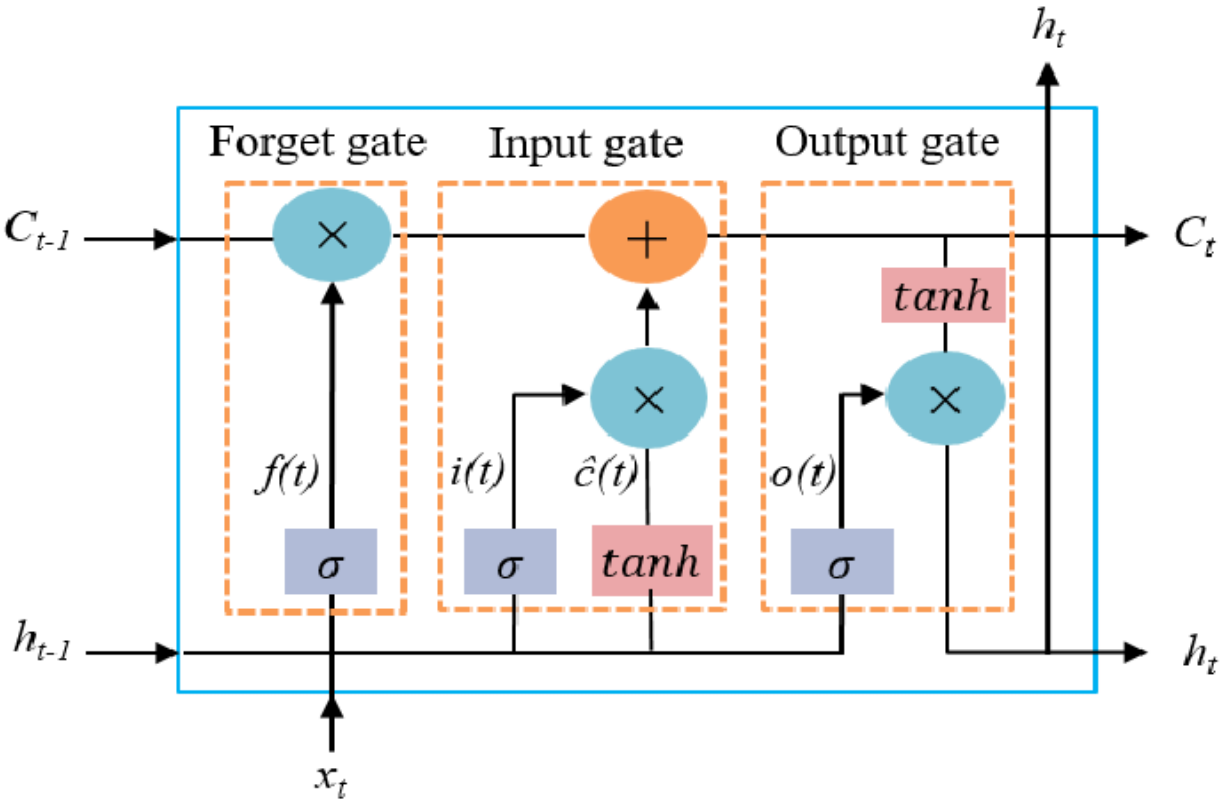


Figure 3. LSTM architectural layout (adapted from Mohsen et al., 2021).

3.5 Loss Function and Optimisation

Both models are trained to minimise the mean squared error (MSE) loss:

$$L(\theta) = \frac{1}{2N}\sum_{i=1}^{N}(\hat{y}_i - y_i)^2 \quad (9)$$

where N is the number of observations, y_i is the observed inflation value, and $\hat{y}_i$ is the model prediction. Parameter updates are performed with the Adam optimiser (Kingma & Ba, 2015):

$$\gamma_i \leftarrow \gamma_i - \alpha\frac{\hat{m_i}}{\sqrt{\hat{v_i}}+\varepsilon} \quad (10)$$

where γ_i are model parameters, $\varphi = 0.001$ is the learning rate, $\hat{m}_i$ and $\hat{v}_i$ are bias-corrected first and second moment estimates of the gradient, and ε is a small constant preventing division by zero.

3.6 Hyperparameter Selection and Validation

A window size of 12 months was selected to embed the annual seasonal cycle in input features, a standard choice for monthly macroeconomic series (Namdari et al., 2025). Training epochs were set to 100 following preliminary experiments confirming convergence within this budget. Dropout rates of 0.20 were applied after each convolutional and recurrent block as regularization, consistent with common practice for small-sample deep learning. The Adam learning rate of 0.001 is the widely used default (Kingma & Ba, 2015). A 20% within-training validation split was used for loss monitoring. A chronological hold-out design was preferred over full cross-validation to preserve temporal structure and avoid look-ahead bias.

3.7 Evaluation Metrics

Model performance is assessed on the held-out test set using MSE and RMSE. RMSE is expressed in the original units (percentage points) and serves as the primary ranking criterion; lower values indicate higher out-of-sample accuracy:

$$MSE = \frac{1}{N}\sum_{i=1}^{N}(y_i - \hat{y}_i)^2 \quad (11)$$

$$RMSE = \sqrt{MSE} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(y_i - \hat{y}_i)^2} \quad (12)$$

4. Results and Discussion

4.1 Exploratory Data Analysis

The dataset spans 252 monthly inflation observations (January 2004–December 2024). Table 2 presents summary statistics for the full series.

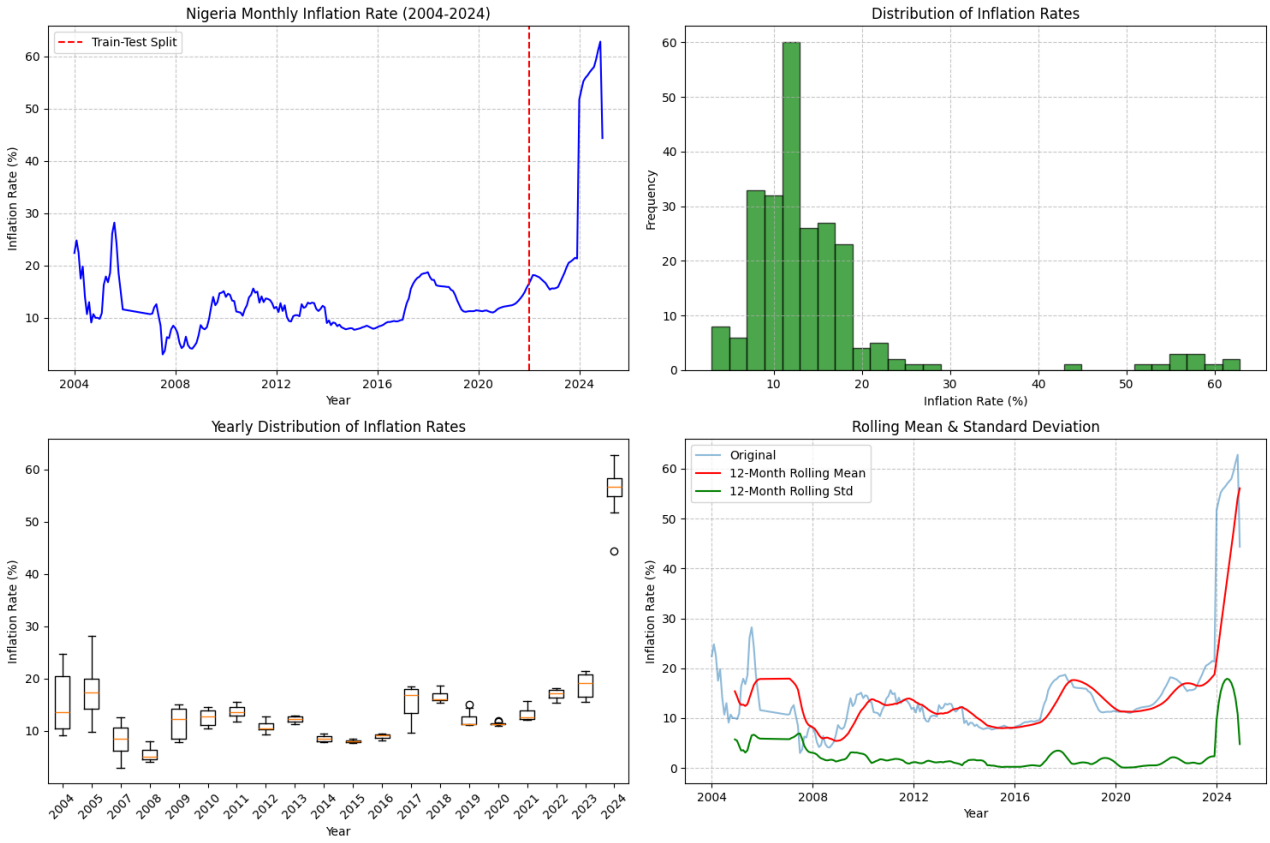


Figure 4. *Exploratory data analysis of Nigeria's monthly inflation rate (2004–2024): time-series plot (top-left), histogram (top-right), annual box plots (bottom-left), and 12-month rolling mean and standard deviation (bottom-right).*

Table 2. Descriptive statistics of Nigeria's monthly headline inflation rate, 2004–2024

Statistic	Value
Number of Observations	252
Mean (%)	15.42
Standard Deviation (pp)	11.32
Minimum (%)	3.00
Maximum (%)	62.80
25th Percentile (%)	9.30
Median / 50th Percentile (%)	12.60
75th Percentile (%)	16.48

Note. pp = percentage points. Values are monthly headline CPI inflation rates (year-on-year, %). Source: Central Bank of Nigeria / National Bureau of Statistics.

The mean of 15.42% with a standard deviation of 11.32 pp reflects the series' high dispersion. The histogram is markedly right-skewed, with most observations clustering between 5% and 20% but a substantial tail reflecting the post-2022 acceleration. The full EDA is presented in Figure 4.

The time-series plot (Figure 4, top-left) reveals four distinct regimes: (i) high volatility with peaks above 20% from 2004 to 2007; (ii) relative stability below 10% from 2008 to 2015; (iii) a 2016–2017 spike to approximately 18.5% followed by gradual moderation; and (iv) an unprecedented exponential surge from 2022 through late 2024, surpassing 60%. The rolling standard deviation (bottom-right) confirms volatility clustering and a structural break from 2023 onwards, posing the principal modelling challenge.

4.2 Model Architectures and Hyperparameters

Table 3 summarises the final architectures and shared training configuration. Both networks use the same loss function, optimiser, batch size, epoch budget, and validation fraction, ensuring a fair comparison.

Table 3. Model architectures and hyperparameters

Category	Layer / Parameter	Configuration
1D-CNN	Input Layer	(batch_size, 1, 12)
	Conv1D – L1	64 filters, kernel = 3, ReLU
	MaxPooling1D	pool_size = 2
	Dropout	rate = 0.20
	Conv1D – L2	32 filters, kernel = 3, ReLU
	MaxPooling1D	pool_size = 2
	Dropout	rate = 0.20
	Flatten	–
	Dense	50 units, ReLU
	Dropout	rate = 0.20

Category	Layer / Parameter	Configuration
	Output	1 unit, Linear
LSTM	Input Layer	(batch_size, 12, 1)
	LSTM Layer 1	50 units, return_sequences = True
	Dropout	rate = 0.20
	LSTM Layer 2	50 units
	Dropout	rate = 0.20
	Dense / Output	1 unit, Linear
Shared	Loss	MSE
	Optimiser	Adam (φ = 0.001)
	Batch size	32
	Epochs	100
	Validation split	20% of training set

Note. Conv1D = one-dimensional convolutional layer; ReLU = rectified linear unit; φ = learning rate.

4.3 Training Convergence

Figure 5 shows training loss curves over 100 epochs for both models. Both architectures converge smoothly without divergence between training and validation loss, indicating the absence of substantial overfitting a favourable outcome given the limited sample size.

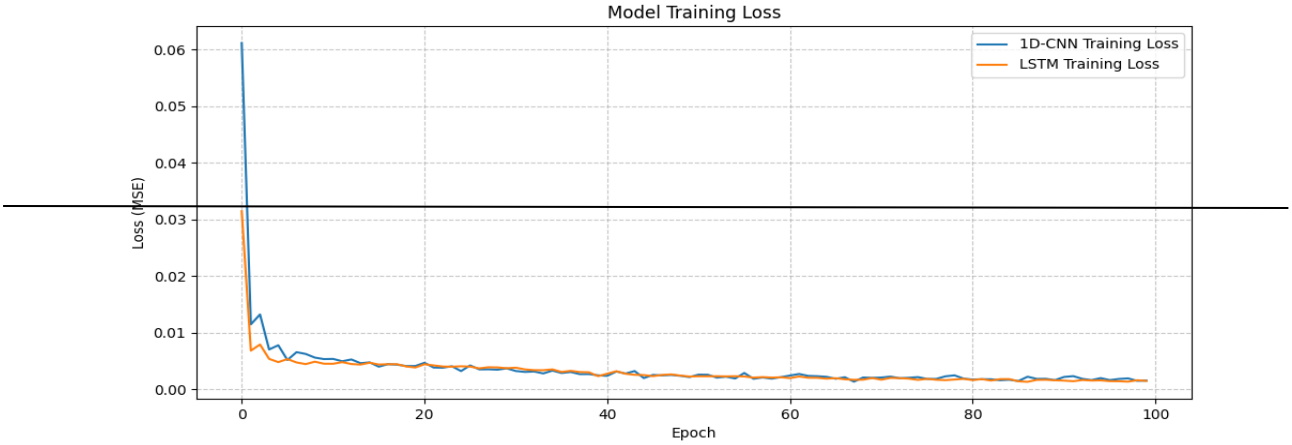


Figure 5. *Training loss histories for the 1D-CNN and LSTM over 100 epochs.*

4.4 Out-of-Sample Performance

Table 4 reports MSE and RMSE on training and test sets. On the training set the two models are closely matched (RMSE: 2.01 vs. 2.12 pp), confirming comparable in-sample fit. On the test set, the LSTM demonstrates markedly superior generalisation: its RMSE of 10.21 pp is nearly half that of the 1D-CNN (18.52 pp), and its MSE (104.18) is less than one-third of the 1D-CNN's (343.07). The large gap between training and testing RMSE for the 1D-CNN (2.12 vs. 18.52) suggests that fixed-window convolutional

filters optimised on the relatively stable 2004–2021 sub-sample struggle to generalise to the unprecedented 2022–2024 volatility regime. These results are consistent with findings by Atashi et al. (2023) and Sukatmo et al. (2025) in analogous non-stationary settings.

Table 4 ***Model performance metrics (training and testing)***

Model	Training MSE (pp²)	Training RMSE (pp)	Testing MSE (pp²)	Testing RMSE (pp)
1D-CNN	4.4933	2.1197	343.071	18.522
LSTM	4.0530	2.0132	104.180	10.207*

Note. MSE = mean squared error; RMSE = root mean squared error; pp = percentage points. Asterisk (*) denotes the best-performing model on testing RMSE.

4.5 Visual Comparison of Forecasts

Figure 6 presents the full-sample and test-period forecast visualisation.

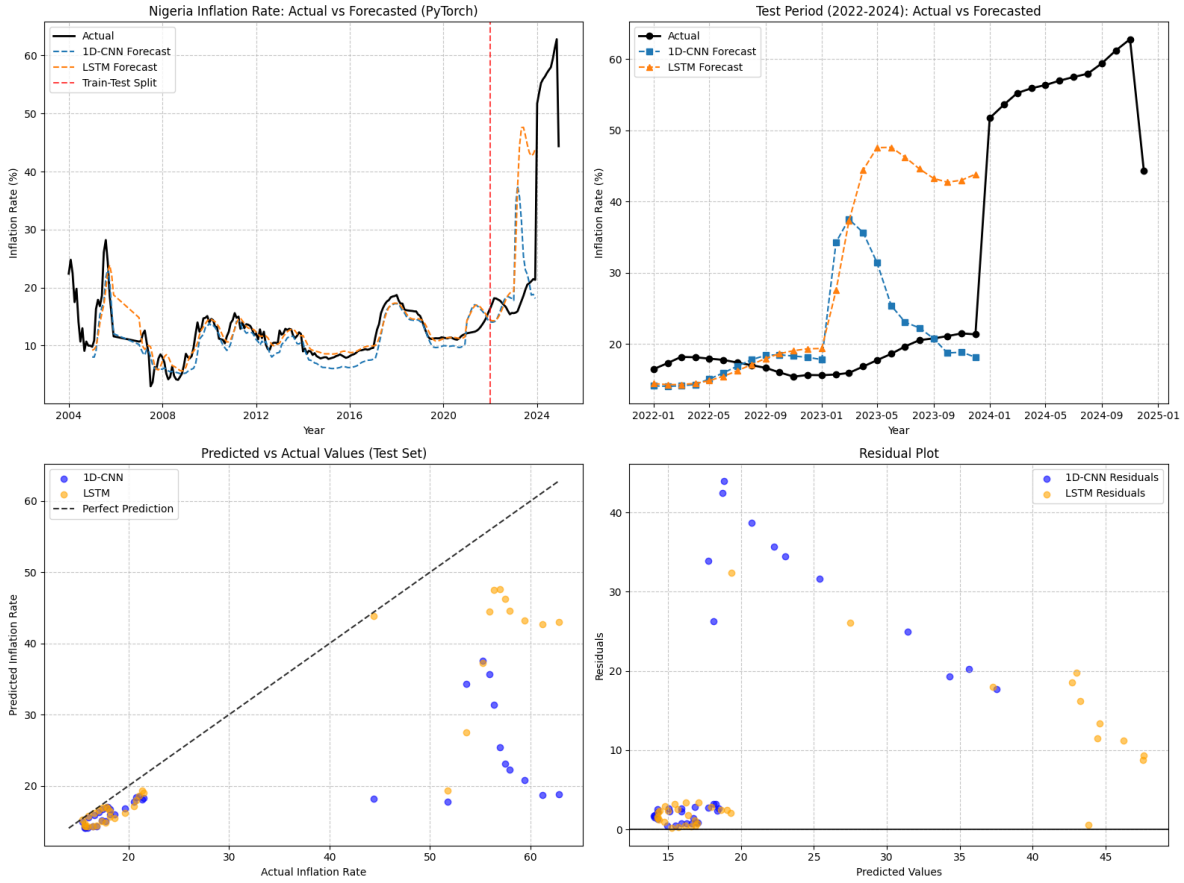


Figure 6. *Actual versus predicted inflation rates: full sample 2004–2024 (top-left), test period 2022–2024 (top-right), prediction error scatter (bottom-left), and residual plots (bottom-right). Black line = actual; blue = 1D-CNN; orange = LSTM.*

Both models capture the broad inflationary trajectory, but LSTM forecasts align more closely with actual values, particularly during the 2023–2024 acceleration. The 1D-CNN exhibits a systematic lag in responding to rapid directional changes. Residual plots confirm that LSTM errors are more randomly

distributed around zero, while 1D-CNN residuals display systematic patterns suggesting unmodelled temporal autocorrelation.

4.6 Twelve-Month-Ahead Forecasts for 2025

4.6.1 Monthly Projections

Using a recursive multi-step prediction strategy, where each predicted value feeds back as input for the next step, both models generate 12-month-ahead projections for January–December 2025 (Table 5).

Table 5. *2025 monthly headline inflation rate forecasts (%)*

Month	1D-CNN Forecast (%)	LSTM Forecast (%)
January 2025	11.97	38.89
February 2025	7.75	32.10
March 2025	6.68	25.17
April 2025	6.25	18.90
May 2025	5.46	13.67
June 2025	5.17	9.65
July 2025	5.38	6.74
August 2025	5.75	4.76
September 2025	5.14	3.47
October 2025	4.02	2.72
November 2025	3.06	2.36
December 2025	4.05	2.32

Note. Forecasts generated using a recursive prediction strategy.

4.6.2 Forecast Summary Statistics

Table 6 *2025 forecast summary statistics*

Metric	1D-CNN	LSTM
Annual Average (%)	5.89	13.40
Q1 Average (%)	8.80	32.05
Q2 Average (%)	5.63	14.07
Q3 Average (%)	5.42	4.99
Q4 Average (%)	3.71	2.47
Minimum Monthly Forecast (%)	3.06	2.32
Maximum Monthly Forecast (%)	11.97	38.89
Trend Jan–Dec (pp change)	−66.15	−94.02

Note. pp = percentage points. Trend = December 2025 forecast minus January 2025 forecast.

Both models project steep deceleration in 2025, with the steepest decline in Q1, consistent with base effects from the 2024 surge. The LSTM's substantially higher Q1 average (32.05% vs. 8.80%) reflects its retention of the elevated 2024 price level. By year-end both models converge (LSTM: 2.32%; 1D-CNN: 4.05%), though divergence is widest in H1.

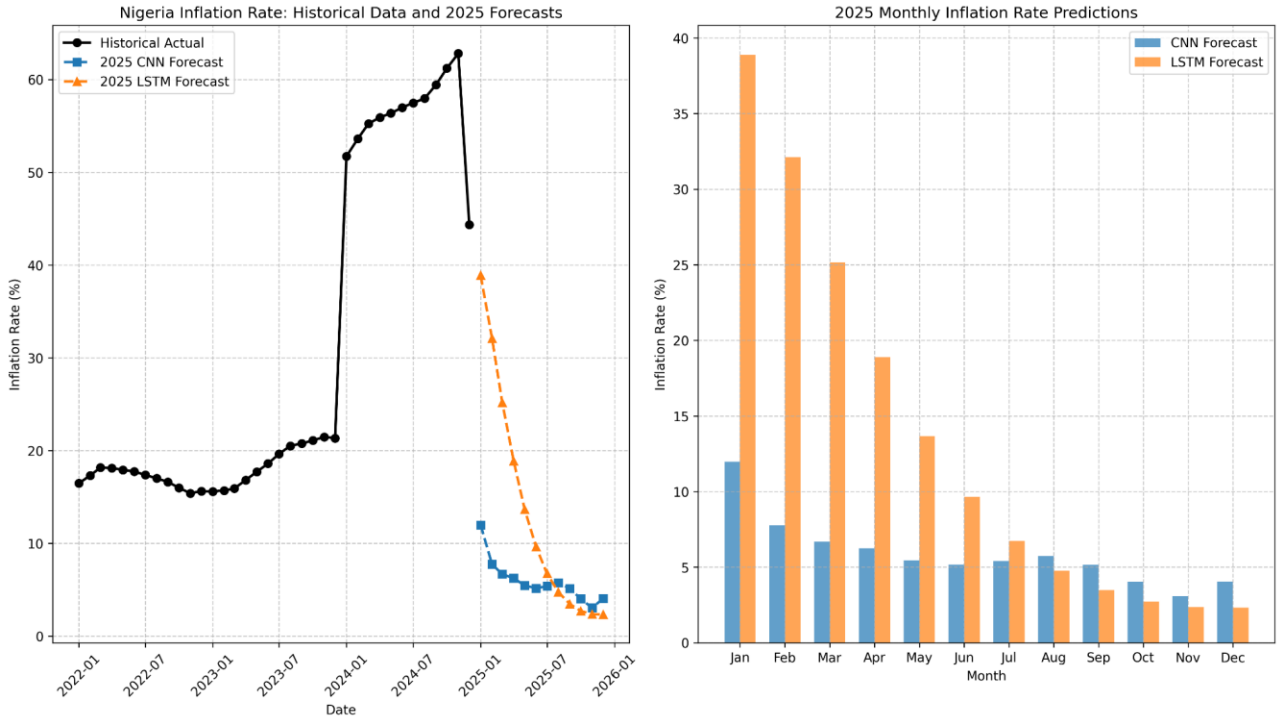


Figure 7. 2025 inflation forecasts: historical context and model projections (left panel) and monthly comparison bar chart (right panel). Black line = historical actual; blue = 1D-CNN; orange = LSTM.

4.7 Policy Implications

The H1 2025 forecast divergence is practically significant for monetary policy. The LSTM's higher near-term projection implies inflation remains well above the CBN's single-digit target for most of the year, arguing for sustained tightening and cautious easing only in H2. Premature rate cuts could risk re-anchoring inflation expectations upward. The 1D-CNN's more aggressive disinflation path would imply earlier scope for rate reductions, but its weaker accuracy warrants caution.

Both models agree that cost-push drivers, supply-chain constraints, exchange-rate pass-through, and energy price effects will not have fully resolved by year-end 2025. Complementary fiscal measures targeting agricultural productivity, fuel distribution, and foreign-exchange liquidity are therefore likely necessary alongside monetary tightening. The LSTM should be deployed as a complementary nowcasting tool alongside, not in replacement of, interpretable VAR and ECM frameworks within the CBN's forecasting suite.

4.8 Discussion

LSTM's superior performance stems from four complementary mechanisms rooted in its architecture. Its forget, input, and output gates allow selective retention of relevant historical information, which is essential for capturing the persistent regimes and structural breaks that characterize Nigeria's inflation series. This gating structure also enables the cell state to adapt more quickly to shifts in volatility, such as the high-inflation 2023–2024 period, compared to the CNN's fixed receptive field, which constrains how much

context it can incorporate at once. Additionally, the cell-state design lets LSTM retain signals from distant time steps without the decay problem that limits plain RNNs, preserving long-range dependencies important for a series with slow-moving macroeconomic dynamics. Finally, these properties translate into more plausible forecasts: LSTM's lower 2025 annual inflation projection (13.40%) implies a more gradual disinflation path, which is arguably more consistent with expected base-effect dynamics than a sharper decline would be.

Several limitations merit acknowledgement. First, the univariate design excludes exogenous macroeconomic drivers, exchange rates, oil prices, money supply, and food-supply indices known determinants of Nigerian inflation. Second, both models underperform during extreme volatility periods (2023–2024), a fundamental constraint of univariate models trained on historical patterns. Third, the 252-observation sample is modest for deep learning. Fourth, both models lack direct interpretability, complicating communication to policymakers.

5. Conclusion

This study built and compared a 1D-CNN and an LSTM model to forecast Nigeria's monthly headline inflation rate using 252 monthly observations spanning January 2004 to December 2024, and found that the LSTM substantially outperformed the 1D-CNN on the test set (RMSE of 10.21 versus 18.52 percentage points), a gap attributed to the LSTM's gating mechanisms, which better capture long-term temporal dependencies in the series. Both models trained stably over 100 epochs without overfitting, indicating that dropout regularisation was effective. Looking ahead, the LSTM projects a decelerating but still elevated 2025 annual average inflation rate of 13.40%, suggesting that continued monetary policy vigilance remains warranted, whereas the 1D-CNN forecasts a much sharper disinflation path (5.89% annual average); given its weaker out-of-sample performance, however, the CNN's more optimistic forecast should be treated with caution as a basis for policy decisions. Based on these findings, the study recommends that the Central Bank of Nigeria integrate the LSTM model as a supplementary tool alongside its existing ARIMA and VAR forecasting frameworks, while pursuing several avenues to strengthen its predictive value: developing a multivariate LSTM that incorporates exchange rates, money supply (M2), oil prices, and food-supply indices; applying Bayesian hyperparameter optimization and data augmentation to improve model generalization; exploring hybrid CNN-LSTM and transformer-based architectures for further performance gains; and, rather than relying on single point forecasts, providing ensemble forecast ranges to communicate forecast uncertainty more transparently to policymakers.This study provides the first empirical comparison of 1D-CNN and LSTM architectures for inflation forecasting in Nigeria, establishing a benchmark for analogous deep-learning forecasting exercises across sub-Saharan African and emerging-market economies.

Acknowledgement

The authors thank the Central Bank of Nigeria and the National Bureau of Statistics for making the inflation data publicly available. No external funding was received for this study.

References

Atashi, V., Kardan, R., Gorji, H. T., & Lim, Y. H. (2023). Comparative study of deep learning LSTM and 1D-CNN models for real-time flood prediction in Red River of the North, USA. In *2023 IEEE International Conference on Electro Information Technology (eIT)* (pp. 22–28). IEEE. https://doi.org/10.1109/eIT57321.2023.10187358

Central Bank of Nigeria. (2024). *Annual report and statement of accounts.* CBN. https://www.cbn.gov.ng

Fitrianingsih, R. N., & Kharisudin, I. (2025). Time series modeling of stock price using CNN-BiLSTM with attention mechanism. *Unnes Journal of Mathematics, 13*(1), 11–20. https://doi.org/10.15294/ujm.v13i1.13451

Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation, 9*(8), 1735–1780. https://doi.org/10.1162/neco.1997.9.8.1735

Hoa Tran Thai, Le Thi My, & Nguyen-Dinh, C. H. (2025). Hybrid model of 1D-CNN and LSTM for forecasting Ethereum closing prices. *International Journal of Information Technology, 17*(7), 1–13. https://doi.org/10.1007/s41870-025-02472-6

Kalange, D. (2025). Time series forecasting in finance using convolutional neural network (CNN) model. *Journal of Mathematics and Applied Statistics, 3*(1). https://doi.org/10.35914/mathstat.v3i1.285

Kharisudin, I., Arissinta, I. O., Aulia, S. A., & Wijaya, G. K. (2025). Comparative study of LSTM-based models with hyperparameter optimization for short-term electricity load forecasting. *BAREKENG: Jurnal Ilmu Matematika dan Terapan, 20*(1), 105–122. https://doi.org/10.30598/barekengvol20iss1pp0105-0122

Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR).* https://arxiv.org/abs/1412.6980

Listyanto, N. P., & Yustanti, W. (2025). Comparative study of time series forecasting on iron sales using CNN, MLP, and LSTM. *Journal of Emerging Information System and Business Intelligence, 6*(3). https://doi.org/10.26740/jeisbi.v6i3.71361

Majid, M. A., Saputri, P. D., & Djoepri, S. (2024). Stock market index prediction using bi-directional long short-term memory. *Journal of Applied Informatics and Computing, 8*(1), 55–61. https://doi.org/10.30871/jaic.v8i1.7195

Mohsen, S., Elkaseer, A., & Scholz, S. G. (2021). Industry 4.0-oriented deep learning models for human activity recognition. *IEEE Access, 9*, 167128–167139. https://doi.org/10.1109/ACCESS.2021.3125733

Nair, V., & Hinton, G. E. (2010). Rectified linear units improve restricted Boltzmann machines. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10)* (pp. 807–814).

Namdari, H., Haghighi, A., & Ashrafi, S. M. (2025). Short-term urban water demand forecasting: Application of 1D-CNN in comparison with different deep learning schemes. *Stochastic Environmental Research and Risk Assessment, 39*, 4211–4226. https://doi.org/10.1007/s00477-023-02565-3

Sánchez-Reolid, R., López de la Rosa, F., López, M. T., & Fernández-Caballero, A. (2022). One-dimensional convolutional neural networks for low/high arousal classification from electrodermal activity. *Biomedical Signal Processing and Control, 71*(Part B), Article 103203. https://doi.org/10.1016/j.bspc.2021.103203

Sayed, M., Gharghory, S., & Kamal, H. A. (2024). A hybrid model of CNN and LSTM autoencoder-based short-term PV power generation forecasting. *Electrical Engineering, 106*(4), 1–17. https://doi.org/10.1007/s00202-023-02220-8

Sukatmo, S., Nugroho, H. A., Rusanto, B. H., & Soekirno, S. (2025). Performance comparison of 1D-CNN and LSTM deep learning models for time series-based electric power prediction. *ELKOMIKA: Jurnal Teknik Energi Elektrik, Teknik Telekomunikasi & Teknik Elektronika, 13*(1), 44. https://doi.org/10.26760/elkomika.v13i1.44

Thai, H. T., Le, T. M., & Nguyen-Dinh, C. H. (2025). A novel multi-window time-series forecasting of major cryptocurrencies using hybrid 1D-CNN–LSTM framework. https://doi.org/10.21203/rs.3.rs-8217257/v1