

Enhancing Inflation Forecasting in Indonesia Using N-BEATSx with Exogenous Factors

Talitha Adilla Fujisai Panglima Putri ^{1*}, Mohammad Idhom ^{2*}, Muhammad Nasrudin ^{3*}

* Data Science, Universitas Pembangunan Nasional "Veteran" Jawa Timur

22083010079@student.upnjatim.ac.id ¹, idhom@upnjatim.ac.id ², nasrudin.fasilkom@upnjatim.ac.id ³

Article Info

Article history:

Received 2026-04-14

Revised 2026-05-09

Accepted 2026-05-25

Keyword:

*Deep Learning,
Exogenous,
Forecasting,
Inflation,
N-BEATSx.*

ABSTRACT

Accurate inflation forecasting is crucial for economic stability and effective policymaking, particularly in emerging economies such as Indonesia, where monetary policy, global commodity markets, exchange rate fluctuations, and recurring religious seasonal events simultaneously influence price dynamics. This study proposes an inflation forecasting framework using the N-BEATSx (Neural Basis Expansion Analysis for Time Series with Exogenous Variables) deep learning model, incorporating macroeconomic variables, global oil prices, BI Rate, and the USD/IDR exchange rate, alongside Ramadan and Eid al-Fitr calendar dummy variables as exogenous inputs. The dataset comprises 153 monthly observations spanning January 2013 to September 2025, split into training, validation, and test sets, with a forecasting horizon of six months. The N-BEATSx model is benchmarked against SARIMAX, LSTM, and Prophet. Results on the test set show that N-BEATSx achieves competitive performance (RMSE 0.0067, MAE 0.0058, SMAPE 51.77%) outperforming SARIMAX (RMSE 0.0297, MAE 0.0266) and LSTM (RMSE 0.0098, MAE 0.0084). Although Prophet yields marginally lower absolute errors, the MAE gap is minimal (0.0002), while N-BEATSx offers superior interpretability through an explicit decomposition of forecasts into trend, seasonality, and exogenous components. Component decomposition analysis reveals that macroeconomic exogenous variables dominate the forecast output, confirming their theoretical relevance as inflation drivers. Six-month-ahead forecasts project inflation in the range of 2.47% - 3.54% for October 2025 to March 2026, approaching Bank Indonesia's upper target corridor, suggesting the need for preemptive monetary policy measures.



This is an open access article under the [CC-BY-SA](https://creativecommons.org/licenses/by-sa/4.0/) license.

I. INTRODUCTION

Inflation plays a crucial role in determining economic stability, particularly in emerging economies where price fluctuations directly affect purchasing power and economic growth [1]. In Indonesia, inflation is recognized as one of the most important macroeconomic indicators influencing economic stability [2]. Accurate inflation forecasting, therefore, becomes essential for policymakers in designing effective monetary and fiscal policies. However, predicting inflation remains a significant challenge because price dynamics are influenced not only by macroeconomic factors but also by recurring seasonal events.

Many countries around the world celebrate major holidays according to the lunar calendar. As the impact of a particular holiday on the time series is often similar year after year, it is important to incorporate it into the forecast [3]. In Indonesia, inflation fluctuations are often associated not only with macroeconomic variables but also with seasonal patterns tied to religious holidays such as Ramadan and Eid al-Fitr, which significantly increase inflation rates [4]. Empirical evidence indicates that inflation tends to rise during the Ramadan and Eid al-Fitr periods, driven by surges in household consumption, food demand, and transportation costs that characteristically accompany these religious observances. Data from Badan Pusat Statistik (BPS) indicate that monthly inflation often rises during these months compared to regular

periods. For instance, inflation recorded a deflation of -0.09% in February 2025, but increased to 1.03% in March 2025 during Ramadan and further rose to 1.95% in April 2025 during the Eid al-Fitr period [5]. This pattern suggested that seasonal consumption associated with religious holidays may contribute to short-term inflationary pressures.

In addition to seasonal effects, this study selects macroeconomic variables based on their strong theoretical and empirical relationship with inflation dynamics. The first variable is the BI Rate, which serves as a monetary policy instrument of Bank Indonesia. According to the monetary transmission mechanism framework, changes in the policy interest rate affect credit costs, aggregate demand, and ultimately the general price level [6]. When the central bank raises the BI Rate, borrowing becomes more expensive, reducing consumption and investment, thereby dampening inflationary pressures. Conversely, a rate cut stimulates demand and can accelerate inflation. This theoretical linkage is supported by demonstrating the significance of policy interest rates in inflation modeling [6]. The second variable is the world oil price, which is theoretically connected to inflation through both cost-push and demand-pull channels. From a cost-push perspective, rising oil prices increase production and transportation costs across sectors, which are subsequently transmitted to consumer prices. In the Indonesian context, where fuel subsidies and energy prices are sensitive to global oil market movements, this transmission is particularly relevant [7]. The third variable is the USD/IDR exchange rate, grounded in the theory of exchange rate pass-through (ERPT) to domestic prices. A depreciation of the Rupiah against the US Dollar raises the cost of imported goods, raw materials, and energy, thereby creating upward pressure on domestic inflation [8]. These three macroeconomic variables are therefore incorporated as exogenous inputs to capture both domestic monetary factors and external economic influences.

These fluctuations indicate the importance of developing reliable forecasting models to anticipate future inflation trends. Traditional time series forecasting methods, such as econometric and statistical models including ARIMA (AutoRegressive Integrated Moving Average) and SARIMA (Seasonal AutoRegressive Integrated Moving Average), have long been used to predict inflation [9] [10]. However, these approaches often struggle to capture complex nonlinear relationships [11]. As artificial intelligence has rapidly developed, deep learning models have emerged as promising alternatives for time series forecasting due to their ability to learn complex patterns from historical data. Ananda and Sutopo compared the performance of several deep learning models, such as LSTM (Long Short Term Memory), GRU (Gated Recurrent Unit), and BiGRU (Bidirectional Gated Recurrent Unit), in predicting stock prices for mining companies in the metal subsector (BRMS, MDKA, ARCI, and AMMN) using historical data from 2015 to 2024. Evaluation results show that the GRU model delivers the best overall performance, demonstrating better training stability

and achieving the highest accuracy rates for most stocks, particularly BRMS (RMSE 11.87) and AMNN (RMSE 252.13). Meanwhile, the LSTM model performed better with data having long historical records, such as MDKA (RMSE 85.56). The BiGRU model was effective at capturing specific trend patterns in data with certain characteristics, such as ARCI (RMSE 7.66). Overall, the study highlights that model performance is characterized by data characteristics and volume. Nevertheless, this study focuses on stock price prediction and does not consider macroeconomic variables or calendar-related effects that are essential in modeling inflation dynamics [2]. Then, a study by Hadi et al. uses an LSTM model to predict monthly inflation in Indonesia. It uses inflation data from Bank Indonesia covering January 2003 to May 2025, comprising 269 observations. The model was built using the Cross-Industry Standard Process for Data Mining (CRISP-DM) approach. The evaluation results demonstrate the LSTM model's ability to deliver excellent predictive performance, with an MSE of 0.001573, an MAE of 0.029463, and a coefficient of determination (R^2) of 0.96075. These results suggest that deep learning methods effectively capture nonlinear patterns and temporal dependencies of inflation data. However, the study relies solely on historical inflation data without incorporating macroeconomic indicators or calendar-related effects that may influence inflation dynamics in Indonesia [12]. Another study by Riswanda et al. applied the Neural Basis Expansion Analysis for Time Series (N-BEATS) with Optuna optimization. The results show that the optimized N-BEATS model significantly improved prediction accuracy compared to conventional models and the unoptimized N-BEATS model. It successfully reduced the Mean Absolute Percentage Error (MAPE) from 11.05% to 8.97%, outperforming several other models, such as ARIMA (16.95%) and LSTM (12.23%). Despite its strong performance, the study adopts a univariate approach that does not account for external variables or calendar effects that may contribute to inflation fluctuations [13].

A critical gap emerges from reviewing these prior studies, while deep learning applications for inflation forecasting have grown substantially, no existing study has simultaneously integrated the N-BEATSx decomposition-based deep learning architecture, theoretically grounded macroeconomic transmission variables specific to the Indonesian economy, and explicit lunar calendar effects of Ramadan and Eid al-Fitr within a single unified forecasting framework. Studies that incorporate macroeconomic variables, such as exchange rates, interest rates, and commodity prices, typically employ recurrent architectures like LSTM or GRU that lack interpretable decomposition outputs. Conversely, decomposition-based models such as N-BEATS have been applied in a univariate setting, forgoing the rich exogenous information available. This gap is particularly consequential for Indonesia, where inflation dynamics are simultaneously shaped by monetary policy transmission, global commodity prices, pass-through exchange rate depreciation effects, and recurring demand surges tied to the Islamic lunar calendar, a

combination of drivers that no single prior study has addressed comprehensively.

Therefore, this study proposes using the N-BEATSx (Neural Basis Expansion Analysis for Time Series with Exogenous Variables) deep learning model, incorporating macroeconomic variables and Ramadan-Eid al-Fitr calendar effects, to address this identified gap. The N-BEATSx framework is uniquely suited for this purpose because it combines the nonlinear learning capacity of deep learning with an interpretable basis expansion mechanism that explicitly decomposes forecasts into trend, seasonality, and exogenous components, enabling not only accurate prediction but also transparent attribution of inflation drivers. This interpretability advantage is particularly valuable for monetary policymakers who require not just point forecasts but also an understanding of the structural factors underlying inflation movements.

In the context of Indonesia, where inflation is simultaneously driven by monetary policy, global commodity markets, exchange rate dynamics, and religious seasonal patterns, an univariate with exogenous approach is both theoretically justified and empirically necessary. However, many existing studies focus primarily on macroeconomic variables and pay limited attention to calendar-related effects, potentially reducing forecasting accuracy. The N-BEATSx model offers a flexible architecture that enables the integration of both historical time series data and multiple exogenous variables within a single deep learning framework. By incorporating macroeconomic indicators alongside seasonal calendar effects, this model is expected to better capture both structural and seasonal patterns in inflation data. Therefore, the novelty of this study lies in the theoretically grounded integration of macroeconomic variables and religious calendar effects into the N-BEATSx deep learning framework for inflation forecasting in Indonesia. This research is expected to contribute to the existing literature by providing a more comprehensive modeling approach and improving the accuracy of inflation predictions, particularly in capturing seasonal economic dynamics.

The novelty of this study is therefore threefold. First, it applies N-BEATSx specifically to Indonesian inflation forecasting, extending the univariate N-BEATS framework of Riswanda et al. [13] by incorporating multiple exogenous variables. Second, it explicitly models the inflationary impact of Ramadan and Eid al-Fitr through binary calendar encoding grounded in the Indonesian Islamic calendar, addressing the calendar-effect gap identified across all reviewed deep learning studies. Third, it provides a rigorous multi-model comparison, including SARIMAX, LSTM, and Prophet, that benchmarks N-BEATSx against both statistical and deep learning baselines under identical data conditions, contributing a reproducible evaluation framework to the existing literature on Indonesian inflation forecasting.

II. METHODS

This study implements the N-BEATSx model as a deep learning approach for inflation forecasting. The training and evaluation processes are conducted simultaneously through hyperparameter tuning using Optuna integrated with time series cross-validation to ensure optimal model generalization. The steps are detailed in the following subsections in Figure 1.

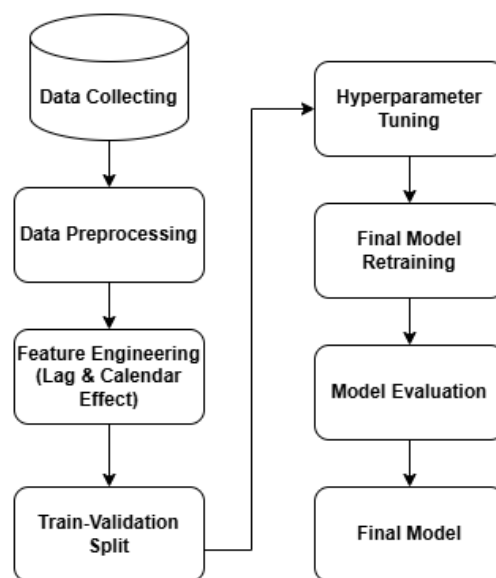


Figure 1. Research Flow

A. Data Collecting

This study uses inflation data from the BPS website, as well as macroeconomic variables such as global oil prices, the BI rate, and the Rupiah exchange rate against the US Dollar, obtained from the Bank Indonesia website and [investing.com](https://www.investing.com) [5] [14] [15]. A total of 165 data from January 2012 to September 2025 were obtained and stored in Excel format, ready for subsequent data processing and analysis.

B. Data Preprocessing

Data preprocessing ensures the quality and consistency of the dataset prior to model development. The process begins with loading the dataset, followed by an initial inspection using descriptive statistics to understand its structure and variable types. Duplicate data and missing values are then identified and addressed to prevent potential bias in the analysis. Outlier detection is also performed to identify extreme values that could affect the model's performance. Next, the date column is converted into a standard datetime format to ensure proper temporal representation. The dataset is then indexed by the date variable and sorted chronologically to preserve the time-series structure required for modelling. These preprocessing steps are essential to ensure that the data is clean, consistent, and suitable for subsequent feature engineering and model training [16].

C. Feature Engineering

Feature engineering enhances the model's ability to capture temporal dependencies and seasonal patterns in the data. In this study, lag-based features are constructed from the target variable, including lag1, lag3, lag6, and lag12, which represent inflation values from the previous 1, 3, 6, and 12 periods, respectively. The estimation of parameters depends on the lag length used [17]. These lag features capture short and long-term historical effects that influence current inflation levels. Additionally, calendar-based dummy variables indicate the occurrence of Ramadan and Eid al-Fitr. These are used to capture the recurring seasonal patterns of increased consumption and economic activity that occur during these periods, and which are known to contribute to inflationary fluctuations in Indonesia. By incorporating both lag features and calendar effect variables, the model is expected to learn the underlying temporal structure and improve forecasting accuracy.

D. Splitting Data

The dataset is divided into training, validation, and test sets using a time-based split. The last 36 observations are separated from the dataset, with 24 allocated for validation and the remaining 12 used to learn the underlying patterns in the data. The validation set is utilized during the model development phase, particularly for hyperparameter tuning, to evaluate model performance on unseen data and prevent overfitting. Meanwhile, the test set is reserved for the final evaluation to assess the model's predictive performance on completely unseen data [18]. Additionally, the training and validation sets are combined during hyperparameter tuning to maximize the use of available data. The splitting strategy preserves the temporal order of the data, ensuring that the model is trained on past observations and evaluated on future periods, thereby avoiding data leakage and providing a reliable estimate of model performance [19].

E. Model Training

The N-BEATSx architecture extends the original N-BEATS framework by explicitly incorporating exogenous variables into the forecasting process. The model receives as input a combined vector consisting of historical target values and exogenous variables over a lookback window of length L , defined as:

$$z_t = [y_{t-L+1}, \dots, y_t, x_{t-L+1}, \dots, x_t] \quad (1)$$

where y_t represent the target variable (inflation) and x_t denotes the exogenous variables, which in this study include world oil prices (*harga minyak dunia*), the BI Rate, the USD/IDR exchange rate (*Kurs USD/IDR*), and Ramadan-Eid al-Fitr calendar dummy variables. The exogenous variables are concatenated directly with the lagged target values and passed as a single input vector to the network, enabling each block to jointly learn from both endogenous dynamics and external influences.

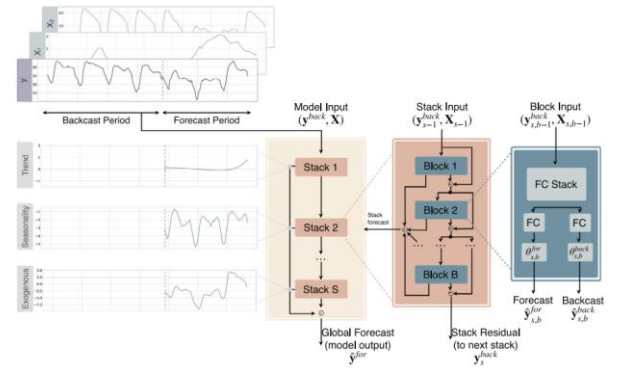


Figure 2. N-BEATSx Architecture [20]

The overall architecture is organized into a hierarchy of stacks and blocks. Each block receives the residual backcast from the previous blocks as its input and produces a backcast-forecast pair:

$$(y^{back}, y^{for}) = f_{\theta}(z_t) \quad (2)$$

The backcast y^{back} captures the portion of the input signal explained by that block and is subtracted from the current input before being passed to the next block, acting as a residual connection that progressively removes explainable patterns. The forecast y^{for} contributes a partial prediction of the future horizon. The global forecast is obtained by aggregating contributions across all "S" stacks and "B" blocks:

$$\hat{y}_{t+1:t+H} = \sum_{s=1}^S \sum_{b=1}^B y_{s,b}^{for} \quad (3)$$

To achieve interpretability, each block decomposes its forecast output through basis expansion into three interpretable components:

$$y_{s,b}^{for} = T\theta_{s,b}^{(T)} + S\theta_{s,b}^{(S)} + X\theta_{s,b}^{(X)} \quad (4)$$

where T , S , and X are the basis matrices associated with the trend, seasonality, and exogenous components, respectively. The coefficient θ denotes the corresponding learned coefficient vectors [21]

The integration of exogenous variables operates specifically through the exogenous basis expansion term $X\theta_{s,b}^{(X)}$. Within each block, the concatenated input vector z_t is first passed through a stack of fully connected (FC) layers with ReLU activations, producing a shared hidden representation. This representation is then projected through two separate FC outputs. The exogenous coefficients vector $\theta_{s,b}^{for}$ for the forecast output and one $\theta_{s,b}^{back}$ for the backcast output. The exogenous coefficient vector $\theta_{s,b}^{(X)}$ is extracted from $\theta_{s,b}^{for}$ and multiplied by the exogenous basis matrix X , which encodes the projected influence of the external

variables over the forecast horizon. This design allows the model to learn how each exogenous variable, such as oil price shocks or seasonal effects from Ramadan, individually contributes to and modifies the inflation forecast, while remaining separable from the trend and seasonality components. The coefficients θ are optimized end-to-end during training by minimizing the prediction error between actual and forecasted values [21]. To enhance the model's ability to capture complex temporal dynamics, lagged features of the target variable are included alongside external variables as input features too [20]

F. Hyperparameter Tuning

Hyperparameter tuning is performed to obtain the optimal configuration of the N-BEATSx model, improve forecasting performance, and avoid the problem of overfitting or underfitting [22]. In this study, the tuning process is conducted using the Optuna optimization framework, which applies a Bayesian optimization approach to efficiently explore the hyperparameter search space [13]. Hyperparameter tuning is focused on key parameters that significantly influence model performance. Several key hyperparameters are optimized, including the input window size, learning rate, number of training steps, and hidden layer size. A study has shown that not all hyperparameters significantly impact model performance, and some can be kept at their default values without substantial performance degradation [23]. The search space is defined using a combination of categorical and continuous distributions, allowing the model to evaluate different configurations across multiple trials.

After completing the hyperparameter tuning process, the prediction results from the cross-validation procedure are extracted for further analysis. The output generated by the model is structured into a new dataset consisting of the timestamp (ds), actual values (y), and predicted values (NBEATSx). For clarity and ease of interpretation, the columns are renamed to represent the actual observations and model predictions. This structured result is then used to evaluate the model's performance by comparing predicted values against the corresponding actual values. The comparison enables the calculation of evaluation metrics and provides insight into how well the model captures the underlying patterns in the data during the validation phase [24].

G. Final Model Retraining

After obtaining the optimal hyperparameter configuration from the tuning process, the final model is retrained using the combined training and validation dataset to maximize the use of available data. This retraining step ensures that the model learns from a larger dataset, thereby improving its ability to capture underlying patterns and enhancing generalization performance [25]. The selected hyperparameters, including input size, learning rate, number of training steps, and hidden layer size, are fixed based on the best results obtained from

the optimization process. The retrained model is then used for final evaluation on the test dataset and for generating future forecasts. This approach ensures that the final model benefits from both optimal parameter selection and maximum data utilization, leading to more reliable and robust prediction results [25]. It is common practice to retrain the model using the entire dataset to maximize the information available for learning [26]. Therefore, the retrained model serves as the final model in this study and is expected to provide accurate and consistent predictions for real forecasting applications.

H. Model Evaluation

Model performance is assessed using evaluation metrics, namely Mean Absolute Error (MAE), Root Mean Squared Error (RMSE) [21]. In addition to MAE and RMSE, this study also uses Symmetric Mean Absolute Percentage Error (SMAPE) to evaluate model performance in the form of a symmetric percentage. SMAPE is more stable than MAPE when actual values are close to zero, making it a suitable metric for analyzing data such as inflation [27]. To provide a comprehensive evaluation of forecasting performance, it is recommended that multiple evaluation metrics be used, as different metrics capture different aspects of model accuracy [28].

$$MAE = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i| \quad (5)$$

$$RMSE = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2} \quad (6)$$

$$MAPE = \frac{100\%}{n} \sum \frac{|y_i - \hat{y}_i|}{y_i} \quad (7)$$

$$SMAPE = \frac{1}{n} \sum_{i=1}^n \frac{|y_i - \hat{y}_i|}{\frac{|y_i| + |\hat{y}_i|}{2}} \times 100\% \quad (8)$$

III. RESULT AND DISCUSSION

A. Data Collecting

Inflation data were obtained from the Badan Pusat Statistik (BPS) website, ensuring the use of official and credible sources. To enrich the analysis, additional macroeconomic variables were collected from investing.com, as well as the BI rate obtained from Bank Indonesia's website. The dataset consists of 165 monthly observations spanning from January 2012 to September 2025, providing a sufficiently long time horizon to capture both short-term fluctuations and long-term economic patterns. The variables collected are summarized in Table I below, covering five dimensions: general inflation (*Inflasi Umum*), world oil prices (*Harga Minyak Dunia*), the BI Rate, and the USD/IDR exchange rate (*Nilai Tukar Rupiah USD*).

TABLE I
SAMPLE DATA

Date	Inflasi Umum	Harga Minyak Dunia (USD/barrel)	BI Rate	Nilai Tukar Rupiah USD (USD/IDR)
01-01-2012	0.0365	98.46	0.060	8980
01-02-2012	0.0356	107.08	0.057	9010
...
2025-08-01	0.0231	64.21	0.050	16665
2025-09-01	0.0265	63.17	0.047	16490

B. Data Preprocessing

Prior to analysis, the dataset underwent a systematic preprocessing pipeline to ensure data quality and structural integrity. The dataset was first examined for missing values and duplicate data. No significant missing values or duplicates were identified across all 165 observations and five variables, confirming that the raw dataset was complete and non-redundant. Outlier detection was subsequently conducted, revealing nine outlier observations. These observations were retained rather than removed, as they represent valid fluctuations in inflation data, deviating only slightly (by approximately 0.04) from typical values. In the context of economic time series, such variations are considered natural and may provide valuable insights into underlying economic dynamics, removing them could result in the loss of meaningful patterns. The datetime column was then verified, standardised into a consistent monthly format (YYYY-MM-DD), and set as the chronological index of the database. This temporal is vital for accurate time series forecasting, as it ensures the sequential structure of observations is preserved throughout all subsequent processing stages. The result of this preprocessing stage is a clean, well-structured dataset ready for feature engineering.

C. Feature Engineering

Several transformations were applied in the feature engineering stage to enhance the model's ability to capture

temporal patterns and external influences on inflation. Calendar effects were encoded using binary dummy variables to represent the occurrence of Ramadan and Eid al-Fitr. Two separate binary columns were constructed, the Ramadan variables take a value of 1 for months in which Ramadan falls, and 0 otherwise. Similarly, the Eid al-Fitr variables take a value of 1 for the month in which Eid al-Fitr is observed, and 0 otherwise. The determination of which months correspond to Ramadan and Eid al-Fitr was based on the official Indonesian Islamic calendar for each respective year over the study period (2012 - 2025). Since the Islamic calendar is lunar-based, the Gregorian months associated with these events shift by approximately 10-11 days each year, making explicit binary dummy encoding necessary rather than relying on fixed seasonal indices. These two variables are treated as independent binary exogenous inputs to the N-BEATSx model, allowing it to separately learn the inflationary contribution of each religious event.

Lag features were also generated for the target variable, specifically lag-1, lag-3, lag-6, and lag-12, to capture short-term and seasonal dependencies in the time series data. These intervals were selected based on established time-series analysis practices. Lag-1 captures the immediate temporal dependence, lag-3 reflects quarterly patterns, lag-6 captures semi-annual trends, and lag-12 represents yearly seasonality, particularly relevant for Indonesia inflation data.

The dataset was then reformatted to comply with the input structure required by the NeuralForecast framework, ensuring consistent column names (unique_id, ds, y). As a consequence of generating lag features, particularly lag-12, the first twelve observations contained missing values. These rows were therefore removed to maintain data integrity, resulting in a final dataset of 153 observations spanning from January 2013 to September 2025. These feature engineering steps provide the model with richer information, enabling it to learn temporal relationships more effectively and improve its forecasting performance. The final dataset can be seen in Table II below.

TABLE II
FINAL DATASET

unique_id	ds	y	lag1	lag3	lag6	lag12	Harga Minyak Dunia	BI Rate	Nilai Tukar Rupiah USD	Ramadan	Eid al-Fitr
inflasi	01-01-2013	0.0457	0.0430	0.0461	0.0456	0.0365	97.65	0.0575	9690	0	0

<i>inflasi</i>	01-02-2013	0.0531	0.0457	0.0432	0.0458	0.0356	92.03	0.0575	9660	0	0
...
<i>inflasi</i>	01-08-2025	0.0231	0.0237	0.0160	-0.0009	0.0212	64.21	0.0500	16665	0	0
<i>inflasi</i>	01-09-2025	0.0265	0.231	0.0187	0.0103	0.0184	63.17	0.0475	16490	0	0

D. Splitting Data

After feature engineering, the final dataset was divided into training, validation, and test sets using a time-based splitting approach, which strictly preserves the chronological order of the data to prevent data leakage and reflect real-world forecasting conditions.

TABLE III
SPLITTING DATA

Subset	Data Range	Proportions
Training Set	Jan 2013 - Sept 2022	117 Observations (76,5%)
Validation Set	Oct 2022 - Sept 2024	24 Observations (15,7%)
Test Set	Oct 2024 - Sept 2025	12 Observations (7,8%)
Total	153 Observations	100%

The training set provides a substantial historical period for the model to learn underlying macroeconomic and seasonal patterns. Validation set used for hyperparameter tuning and model optimization over two years of unseen data. Test set reserved exclusively for final evaluation to assess the model's generalization capability on the most recent period. This time-based splitting strategy ensures that the model is always trained on past data and evaluated on future periods, thereby preventing any form of look-ahead bias. The training and validation sets were also combined at certain stages, specifically during cross-validation and final retraining, to maximize data utilization while ensuring the test set remained entirely unseen for unbiased evaluation. The forecasting horizon was set to six months ($h=6$), meaning the model generates predictions for up to six periods ahead in a single pass. This design reflects a realistic short to medium-term forecasting scenario relevant for monetary policy planning.

E. Model Training

The model training process was conducted using the N-BEATSx architecture within the NeuralForecast framework. Hyperparameter optimisation was performed using the Optuna library with a Bayesian optimisation approach, specifically employing the Tree-structured Parzen Estimator (TPE) as the default sampler. Unlike grid search or random search, TPE builds a probabilistic model of the objective function and uses it to select the most promising hyperparameter candidates in each subsequent trial, thereby searching the hyperparameter space more efficiently. The optimisation was conducted over 20 trials, with the objective of minimising the RMSE of the cross-validation forecasts.

The search space was defined across five hyperparameters. The input size was sampled categorically from $\{12, 18\}$, representing different lengths of historical observations used as model input, allowing the model to capture varying temporal dependencies. The learning rate was sampled within the range of 8×10^{-4} to 1.2×10^{-3} on a logarithmic scale, enabling efficient exploration of both smaller and larger gradient update steps during optimisation. The maximum number of training steps was sampled as an integer between 800 and 1100 in increments of 100, allowing the model to balance underfitting and overfitting. The hidden layer size was chosen categorically from $\{64, 128\}$, representing different levels of model capacity for capturing nonlinear patterns. Finally, the number of blocks per stack (n_blocks) was sampled categorically from $\{[1,1], [2,2]\}$, where each element corresponds to the number of blocks in the trends and seasonality stacks, respectively, controlling the depth of the decomposition hierarchy.

To ensure robust model evaluation consistent with time series forecasting principles, a rolling forecast origin strategy was used with three windows ($n_windows=3$) and a step size of six periods ($step_size=6$), matching the forecasting horizon of $h=6$. This approach evaluates the model across multiple forecast origins rather than a single holdout split, thereby providing a more reliable estimate of

forecasting performance. The model incorporates both historical and future exogenous variables. Historical exogenous variables include lag features and macroeconomic indicators, while future exogenous variables consist of the Ramadan and Eid al-Fitr binary dummy variables, which are known in advance for future periods. The performance of each hyperparameter configuration was evaluated using RMSE as the optimisation objective. Upon completion of the 20-trial Bayesian search, the best hyperparameter configuration was identified as presented in Table IV below.

TABLE IV
BEST HYPERPARAMETER

Rank	1	2	3
RMSE	0.3150	0.3207	0.3249
Input Size	12	12	12
Learning Rate	0.00096	0.00080	0.00176
Max Steps	900	900	1000
Hidden Size	64	64	64
<i>n_blocks</i>	[1,1]	[1,1]	[1,1]

As shown in Table IV, the optimal configuration yielded a cross-validation RMSE of 0.3150 with an input size of 12 months, a learning rate of approximately 0.00096, a maximum of 900 training steps, a hidden size of 64 units, and 1 block per stack ([1,1]). Notably, the top three configurations converged on the same structural settings (input size = 12, hidden size = 64, *n_blocks* = [1,1]), differing only marginally in learning rate and maximum steps, suggesting that these structural hyperparameters are the most influential in determining model performance for this dataset. The best hyperparameter (rank 1) was subsequently used to retrain

the final N-BEATSx model on the combined training and validation dataset, prior to evaluation on the held-out test set.

F. Model Evaluation

The trained N-BEATSx model was first evaluated on the validation set spanning October 2022 to September 2024, comprising 24 monthly observations. As shown in Figure 3 below, the model demonstrates a reasonably good ability to follow the general downward trend of inflation during this period, achieving an MAE of 0.00539, an RMSE of 0.00799, and a SMAPE of 18.22%. The model successfully captured the broad directional movement of inflation, particularly the gradual decline from late 2022 through mid-2023. The prediction curve closely follows the actual values during relatively stable periods, such as November 2023 through February 2024, where both lines converge with minimal deviation. However, the model tends to overestimate inflation during periods of sharp downward movement, most notably around September 2023, where actual inflation dropped to approximately 2.3% while the model predicted approximately 4.1%. This overestimation likely reflects the model’s learned tendency to anchor predictions to recent higher inflation values, as the sharp decline in September 2023 was an abrupt structural shift that was difficult to anticipate from historical patterns alone. A similar overestimation pattern is observed towards the end of the validation period (July-September 2024), where actual inflation continued to decline while predictions remained relatively elevated.

The SMAPE of 18.22%, while moderate, should be interpreted with caution. In periods when actual inflation is low, particularly in the 2%-3% range observed in 2024, even small absolute prediction errors yield relatively large percentage deviations because of the small denominator in the SMAPE formula. This is a well-documented limitation of SMAPE in near-zero or low-magnitude time series contexts. The MAE and RMSE, being scale-dependent metrics, provide a more stable and representative picture of the model’s absolute forecasting accuracy, both of which remain at acceptably low levels relative to the scale of the target variable

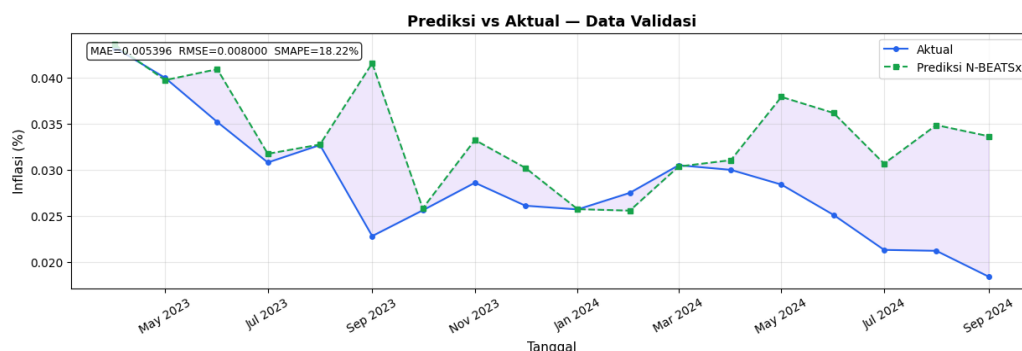


Figure 3. Prediction vs Actual - Validation Set

On the test set covering October 2024 to September 2025, the model yielded an MAE of 0.00577, RMSE of 0.00671, and a notably elevated SMAPE of 51.77%, as visualised in Figure 4 (*Prediksi vs Aktual - Data Test*). Inspection of the prediction curve reveals that the model correctly identified the general low-inflation regime during the October-November 2024 period, where predicted and actual values converge closely at approximately 1.6% and 1.5%, respectively. However, the model failed to capture the anomalous deflation event of approximately -0.09% observed in February 2025, instead predicting a value near 0.003. This extreme near-zero and negative actual value is the primary driver of the elevated SMAPE, as the metric is mathematically unbounded when actual values approach zero, a well-documented limitation of SMAPE in near-zero time series contexts. Excluding this single observation, the

SMAPE for the remaining eleven months would be substantially lower, more accurately reflecting the model's true forecasting capability.

Furthermore, from March 2025 onward, the model consistently underestimated the upward recovery in inflation, predicting values in the range of 0.007-0.015, while actual inflation rose progressively to 0.027 by September 2025. This persistent underestimation suggests that the model, which was predominantly trained on the 2013-2022 inflation regime, may have difficulty generalising to the structural upward shift in inflation observed throughout 2025. This structural break, where inflation recovered from near-deflation in February 2025 to 2.65% by September 2025, represents a rapid regime change that is inherently difficult for any time series model to anticipate without incorporating forward-looking macroeconomic signals.




Figure 4. Prediction vs Actual - Test Set

These observations raise a broader question regarding the model's generalizability to periods of economic shock and structural disruption. The N-BEATSx model, as with most data-driven deep learning approaches, is fundamentally trained to recognize and extrapolate patterns observed in historical data. Its forecasting reliability is therefore contingent on the assumption that future inflation dynamics remain broadly consistent with the distributional characteristics of the training period. This assumption may not hold during episodes of severe economic disruption. The February 2025 deflation episode serves as a direct empirical demonstration of this limitation, a demand-driven shock of sufficient magnitude can push inflation into a distributional region that falls outside the model's learned feature space, resulting in substantial prediction error regardless of model architecture or hyperparameter configuration.

More broadly, while the training dataset encompasses the COVID-19 pandemic suppression period (2020-2021) and the post-pandemic commodity price surge (2022), providing the model with some exposure to abnormal inflation regimes, each economic crisis exhibits distinct characteristics. A future shock transmitted through the oil price channel, one of the exogenous variables explicitly incorporated in this model, would partially activate the model's learned response to that variable. However, shocks

driven by mechanisms outside the model's set of exogenous variables, such as sudden capital-flow reversals, sovereign credit events, or geopolitical supply disruptions, would likely fall outside the model's generalization capacity. In such scenarios, the model's point forecasts should be treated as conditional predictions contingent on the absence of major unforeseen structural breaks, rather than unconditional projections. Future research may address this limitation by incorporating regime-switching mechanisms, augmenting the exogenous variable set with financial stress indicators such as capital flow volatility indices, or combining N-BEATSx with macroeconomic scenario analysis to better account for tail-risk events in inflation forecasting

G. Model Interpretability Analysis

One of the key advantages of the N-BEATSx architecture over conventional deep learning models, such as LSTM, is its inherent interpretability through explicit component decomposition. Unlike black-box models that produce a single undifferentiated output, N-BEATSx decomposes each forecast into three interpretable components, trend, seasonality, and exogenous contributions, enabling transparent attribution of inflation drivers. This decomposition is particularly valuable in the context of monetary policy, where understanding the

source of inflationary pressure is as important as the forecast value itself, and where policymakers require not merely point predictions but also structural insights into which transmission channels are actively shaping price dynamics

As visualised in Figure 5, the decomposition results reveal several important insights into the structure of Indonesia’s inflation dynamics during October 2024 to September 2025. The trend component exhibits a distinct two-phase pattern over the test horizon. During the first phase (October 2024 - February 2025), the trend declines monotonically from approximately 0.0154 to near 0.0001, consistent with the broader post-pandemic disinflation process in which inflationary pressures gradually eased as supply chains normalized and Bank Indonesia’s monetary tightening cycle took effect. The second phase (April 2025 - September 2025) shows a stabilization of the trend component, hovering in the range of 0.0104 - 0.0131, suggesting that the model learned a new, more stable inflation regime following the sharp deflation episode of February 2025. The inflection point at approximately April 2025 (trend rebounding to 0.0131) reflects the model’s learned response to the structural recovery in inflation observed from March 2025 onward.

The seasonality component, which captures both recurring periodic patterns and the calendar effects of Ramadan and Eid al-Fitr, is encoded through binary dummy variables and demonstrates meaningful variation throughout the test period. Seasonality contributions decline from 0.0167 in October 2024 to a trough of approximately 0.0046 in January 2025, before recovering to around 0.0099 in March 2025. This recovery in February-March 2025 is economically interpretable as the model’s learned anticipation of Ramadan 2025, which fell in March of that year, confirming that the explicit calendar encoding successfully informs the seasonality stack. Notably, the seasonality component turns negative in September 2025 (-0.0046), suggesting a seasonal dampening effect in that month that partially offsets the upward pressure from other components, a pattern consistent with the historically lower inflation typically observed in the post-Eid adjustment period. This bidirectional behavior of the seasonality component demonstrates that the model has learned nuanced calendar dynamics beyond a simple additive positive effect.




Figure 5. N-BEATSx Component Decomposition - Test Set

The decomposition analysis consistently shows that the exogenous component dominates throughout the 12-month test period. Exogenous contributions range from a minimum of 0.0643 in October 2024 to a maximum of 0.0883 in September 2025, remaining substantially and

persistently larger in magnitude than both the trend and seasonality components across all twelve forecast periods. This finding carries two important implications. First, it empirically validates the core research hypothesis that macroeconomic transmission variables, specifically the BI Rate, world oil prices, and the USD/IDR exchange rate,

alongside lag features, collectively exert the strongest quantitative influence on Indonesia’s near-term inflation, justifying their theoretical inclusion as outlined in the introduction. Second, the gradual increase in exogenous contributions from 0.0643 in October 2024 to 0.0883 in September 2025 suggests that the influence of macroeconomic external factors on inflation intensified over the test period, potentially reflecting the growing impact of global commodity price movements and exchange rate depreciation pressures on domestic prices during this period. From a policy standpoint, this decomposition provides Bank Indonesia with empirical evidence that external macroeconomic channels, rather than purely domestic seasonal or trend dynamics, were the primary drivers of inflation during 2024 - 2025, suggesting that monetary policy responses during this period should prioritize exchange rate stabilization and commodity price pass-through management over-demand side interventions alone.

H. Future Forecasting

Following evaluation, the N-BEATSx model was retrained on the full available dataset and used to generate a six-month-ahead forecast for the period October 2025 to March 2026. As presented in Figure 6 and Table V, the model projects an upward inflation trajectory, rising from approximately 2.47% in October 2025 to a peak of

approximately 3.54% in January 2026, before moderating slightly to around 2.97% by March 2026. The projected upward trend from October 2025 through January 2026 is economically consistent with several seasonal factors. The model’s learned patterns suggest that the year-end period, encompassing increased government spending, holiday, consumption, and the approaching Ramadan season in early 2026, tends to exert upward pressure on prices. The peak projected in January 2026 may reflect the model’s anticipation of pre-ramadan demand effects, given that Ramadan 2026 is expected to fall in February-March 2026. The subsequent moderation from February to March 2026 aligns with the typical post-peak adjustment after Eid al-Fitr. All six projected values fall within or near Bank Indonesia’s inflation target corridor $2.5\% \pm 1\%$, with the January 2026 peak at 3.54%, marginally approaching the upper bound of this corridor. This suggests that while inflationary pressure is expected to increase in the near term, it remains broadly manageable within the central bank’s policy framework. The forecast is accompanied by a $\pm 15\%$ uncertainty interval to acknowledge the inherent uncertainty in medium-term economic forecasting, particularly given the structural volatility observed in the 2024-2025 test period. As the interval widens progressively from October 2025 toward March 2026, this appropriately reflects the compounding uncertainty of longer forecast horizons.

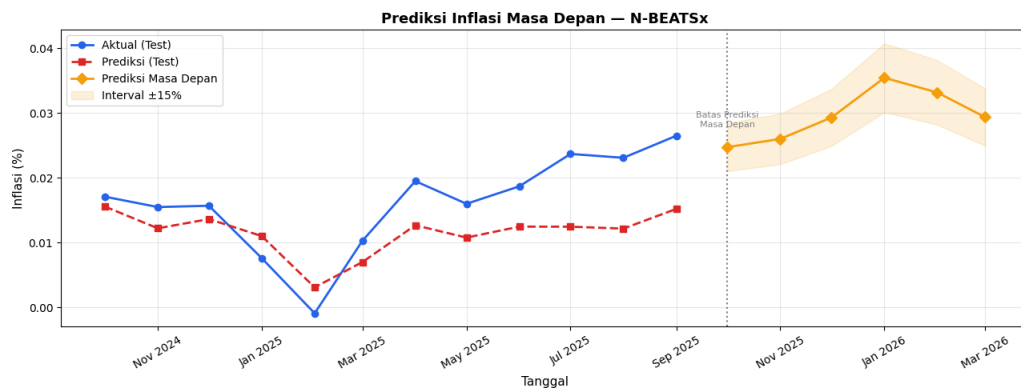


Figure 6. Future Forecasting

TABLE V
FUTURE FORECASTING

Date	Forecast
October 2025	0.0247
November 2025	0.0260
December 2025	0.0293
January 2026	0.0354

February 2026	0.0332
March 2026	0.0294

I. Model Comparison and Analysis

To further evaluate the effectiveness of the proposed approach, a comparative analysis was conducted by benchmarking the N-BEATSx model against three baseline models: LSTM, SARIMAX, and Prophet. The comparison results are presented in Table VI below

TABLE VI
MODEL COMPARISON RESULTS

Model	RMSE	MAE	SMAPE (%)
N-BEATSx	0.0057	0.0067	51.76
LSTM	0.0098	0.0084	96.48
SARIMAX	0.0297	0.0266	73.95
Prophet	0.0048	0.0056	44.61

In terms of absolute error metrics, N-BEATSx outperforms both SARIMAX and LSTM across all evaluation metrics. SARIMAX records the highest error values (RMSE 0.0297, MAE 0.0266), reflecting the inherent limitation of linear statistical models in capturing the nonlinear and regime-shifting dynamics of Indonesian inflation. LSTM, while more flexible as a deep learning architecture, yields a higher MAE (0.0084) and an extremely elevated SMAPE (96.48%), suggesting that the model produces unstable predictions, particularly in low inflation periods, a behaviour likely attributable to its lack of an explicit decomposition mechanism and its sensitivity to the vanishing gradient problem over longer sequences.

The Prophet, however, presents a more nuanced comparison. Prophet achieves lower RMSE (0.00483) and MAE (0.0056) than N-BEATSx, as well as a lower SMAPE (44.61%), making it the best-performing model by raw metrics on this test set. This result is not entirely surprising, as Prophet is specifically engineered for decomposition-based forecasting with native support for custom holiday effects, precisely the calendar dynamics (Ramadan and Eid al-Fitr) that this study targets. The relatively strong performance of Prophet further validates the core hypothesis of this study, that explicitly modeling calendar effects is crucial for accurate inflation forecasting in Indonesia. Nevertheless, the marginal metric difference between Prophet and N-BEATSx must be interpreted critically in the context of the February 2025 deflation anomaly. As previously discussed, the SMAPE values of both models are disproportionately inflated by this single near-zero observation. When considering MAE, which is less sensitive to extreme values, the gap between Prophet (0.0056) and N-BEATSx (0.0058) is minimal (0.0002), suggesting that both models achieve comparable absolute forecasting accuracy.

Furthermore, N-BEATSx offers several architectural advantages that raw metrics alone do not capture. First, N-BEATSx provides an interpretable decomposition of forecast outputs into explicit trend, seasonality, and exogenous components, enabling policymakers to understand which factors drive inflation movements, a feature Prophet does not natively provide at the same granularity for multiple simultaneous exogenous regressors. Second, N-BEATSx integrates macroeconomic

variables (BI Rate, oil prices, exchange rate) and lag features as historical exogenous inputs within a deep learning framework, allowing it to model complex nonlinear interactions between these variables and inflation, interactions that Prophet models only through additive linear regressors. Third, the N-BEATSx architecture is more scalable and extensible for future work involving additional exogenous variables or longer forecast horizons. Overall, while Prophet demonstrates a slight quantitative edge on this particular test period, N-BEATSx remains the more comprehensive and theoretically grounded model for inflation forecasting, particularly given its capacity for nonlinear modeling of multiple exogenous variables and interpretable component decomposition. The results collectively confirm that incorporating calendar effects and exogenous variables within the N-BEATSx architecture substantially improves forecasting accuracy over traditional approaches such as SARIMAX and LSTM.

IV. CONCLUSION

The objective of this study is to develop an accurate inflation forecasting model using deep learning techniques, specifically the N-BEATSx model, and to compare its performance with that achieved using traditional and alternative approaches, such as SARIMAX, LSTM, and Prophet. The results demonstrate that the N-BEATSx achieves competitive forecasting accuracy with an RMSE of 0.0067 and MAE of 0.0057 on the test set, outperforming both SARIMAX (RMSE 0.0297, MAE 0.0266) and LSTM (RMSE 0.0098, MAE 0.0084) across all evaluation metrics. While Prophet achieves marginally lower absolute errors (RMSE 0.00483, MAE 0.0056), the difference in MAE between the two models is minimal (0.0002), and N-BEATSx offers superior architectural advantages in terms of nonlinear univariate modeling and interpretable component decomposition that raw metrics alone do not capture. The elevated SMAPE values observed across all models are primarily attributable to the anomalous February event of February 2025, where near-zero actual values mathematically inflate percentage-based error metrics rather than reflecting genuine model failure.

These findings confirm that N-BEATSx is more effective at capturing both the temporal dynamics and external influences of inflation data. The decomposition analysis further reveals that the exogenous component, comprising the BI Rate, world oil prices, USD/IDR exchange rate, and lag features, consistently dominates the forecast output throughout the test period, with contributions ranging from 0.0643 to 0.0883, substantially exceeding those of the trend and seasonality components. This empirically validates the theoretical justification for incorporating macroeconomic transmission variables as exogenous inputs. The incorporation of Ramadan and Eid al-Fitr calendar effects through binary dummy encoding allows the model to capture seasonal demand shocks

characteristic of Indonesia's inflation behaviour, as evidenced by the seasonality component's recovery in February - March 2025 in anticipation of Ramadan, a pattern absent in models that do not explicitly account for Islamic calendar effects.

The practical implications of this study's findings are directly relevant to monetary and fiscal policymakers in Indonesia. The six-month forecast projects an upward inflation trajectory from 2.47% in October 2025 to a peak of 3.54% in January 2026, before moderating to approximately 2.97% by March 2026. These projections approach the upper bound of Bank Indonesia's inflation target corridor $2,5\% \pm 1\%$, suggesting that preemptive monetary policy measures, such as careful recalibration of the BI Rate, may be warranted in the fourth quarter of 2025 to prevent inflation from breaching the target ceiling during the pre-Ramadan demand surge of early 2026. Next, the dominant exogenous component observed throughout 2024-2025 indicates that Indonesia's inflation during this period was more strongly driven by external macroeconomic channels than by domestic seasonal or trend dynamics. This finding suggests that Bank Indonesia should prioritize stabilization rate stabilization and commodity price pass-through management as primary policy instruments, as these are the channels exerting the greatest quantifiable influence on near-term inflation outcomes. Lastly, the significant inflationary impact of Ramadan and Eid al-Fitr captured by the model underscores the importance of anticipatory fiscal measures, such as market operations, price stabilization programs, and supply-side interventions, being implemented at least one to two months prior to the onset of Ramadan, rather than reactively during the holiday period itself.

This study demonstrates that N-BEATSx, when integrated with theoretically grounded exogenous variables and religious calendar effects, is a robust, accurate, and interpretable approach to inflation forecasting in Indonesia. It highlights the critical importance of combining deep learning architectures with domain-specific economic knowledge to improve both predictive accuracy and policy relevance. Nevertheless, this study acknowledges several limitations. The model's generalizability to periods of severe economic shock or structural break remains constrained, as demonstrated by the February 2025 deflation. Future research may address this by incorporating regime-switching mechanisms, augmenting the exogenous variable set with financial stress indicators, extending the forecasting horizon beyond six months, or exploring hybrid approaches that combine N-BEATSx with macroeconomic scenario analysis to better account for tail-risk events in Indonesia's inflation dynamics.

REFERENCES

- [1] A. N. Paradhita, "Prediksi Inflasi di Indonesia Menggunakan Algoritma Fuzzy dengan Bahasa Pemrograman Python," *Jurnal Penelitian Inovatif*, vol. 4, no. 2, pp. 457–464, Apr. 2024, doi: 10.54082/jupin.339.
- [2] N. Hadi, S. Alam, and I. Kurniawan, "Penerapan Model Lstm (Long Short-Term Memory) Dalam Peramalan Inflasi Month On Month Di Indonesia," *JATI (Jurnal Mahasiswa Teknik Informatika)*, vol. 9, no. 5, pp. 9009–9015, Jul. 2025, doi: 10.36040/jati.v9i5.15119.
- [3] S. J. Taylor and B. Letham, "Forecasting at scale," Sep. 27, 2017. doi: 10.7287/peerj.preprints.3190v2.
- [4] S. Saleh, D. T. K. Wardani, and M. A. Ivantri, "Ramadan, Eid Ul Fitr, And Inflation: Lesson From Indonesian Subnational Data," *Jurnal Ekonomi & Studi Pembangunan*, vol. 20, no. 2, 2019, doi: 10.18196/jesp.20.2.5020.
- [5] Badan Pusat Statistik, "Inflasi Year-On-Year (y-o-y)," 2025. Accessed: Sep. 20, 2025. [Online]. Available: <https://www.bps.go.id/id/pressrelease?page=1&keyword=inflasi&year=2025>
- [6] A. Nurfarhana and T. Anita, "Pengaruh BI Rate dan Nilai Tukar Mengambang oleh Bank Indonesia terhadap Tingkat Inflasi di Indonesia Tahun 2008–2015," *Sosio e-kons*, vol. 9, no. 3, pp. 194–199, Mar. 2018, doi: 10.30998/sosioekons.v9i3.2248.
- [7] R. Wulan, N. Nurpadilah, and R. Pebrian, "Pengaruh Inflasi, Harga Minyak Dunia, Dan Suku Bunga (Bi Rate) Terhadap Indeks Harga Saham Gabungan (Ihsg) (Data Per Bulan Periode 2011-2020)," *PMB*, vol. 1, no. 2, pp. 130–143, Feb. 2023.
- [8] S. P. Sari and S. Nurjannah, "Analisis Pengaruh Nilai Tukar, Jumlah Uang Beredar dan BI Rate Terhadap Inflasi di Indonesia dan Dampaknya Terhadap Daya Beli Masyarakat," *AKTIVA: Journal of Accountancy and Management*, vol. 1, no. 1, pp. 21–29, Jan. 2023, doi: 10.24260/aktiva.v1i1.1015.
- [9] A. Yusnitasari, "Forecasting Indonesian Inflation Using The Box-Jenkins Arima Method," *Jurnal Paradigma Multidisipliner (JPM)*, vol. 1, no. 2, pp. 152–159, 2020.
- [10] M. I. Rizki and T. A. Taqiyuddin, "Penerapan Model SARIMA untuk Memprediksi Tingkat Inflasi di Indonesia," *Jurnal Sains Matematika dan Statistika*, vol. 7, no. 2, Aug. 2021, doi: 10.24014/jsms.v7i2.13168.
- [11] A. Mutiara, N. Fitriyati, and Mahmudi, "Analisis Laju Prediksi Inflasi Di Indonesia: Perbandingan Model Garch/Arch Dengan Long Short Term Memory," *Jurnal Lebesgue : Jurnal Ilmiah Pendidikan Matematika, Matematika dan Statistika*, vol. 4, no. 1, pp. 94–112, 2023.
- [12] M. Kristanaya, A. T. Darmaliana, and S. S. M. Wara, "Implementation of the Stacked LSTM Model in Predicting Inflation in Indonesia Based on the USD/IDR Exchange Rate," *Discover - Data Science, Informatics, Statistics, & Electrical Research Conference (2025)*, vol. 1, no. 1, pp. 1–8, 2025.
- [13] M. N. Riswanda, T. Trimono, and W. S. J. Saputra, "Prediksi Laju Inflasi di Jawa Timur Menggunakan Model N-BEATS dan Optimasi Optuna," *MALCOM: Indonesian Journal of Machine Learning and Computer Science*, vol. 5, no. 3, Jul. 2025, doi: 10.57152/malcom.v5i3.2141.
- [14] B. Indonesia, "BI Rate." Accessed: Sep. 19, 2025. [Online]. Available: <https://www.bi.go.id/en/statistik/indikator/BI-Rate.aspx>
- [15] "Data Historis Minyak Mentah WTI Berjangka," Investing.com. Accessed: Sep. 20, 2025. [Online]. Available: <https://id.investing.com/commodities/crude-oil-historical-data>
- [16] M. Kristanaya, A. T. Darmaliana, and S. S. M. Wara, "Implementation of the Stacked LSTM Model in Predicting Inflation in Indonesia Based on the USD/IDR Exchange Rate," *Discover - Data Science, Informatics, Statistics, & Electrical Research Conference (2025)*, vol. 1, no. 1, pp. 1–8, 2025.
- [17] M. Heru Widiyanto, R. Mayasari, and G. Garno, "Implementasi Time Series Pada Data Penjualan Di Gaikindo Menggunakan Algoritma Seasonal Arima," *JATI (Jurnal Mahasiswa Teknik*

- Informatika*), vol. 7, no. 3, pp. 1501–1506, Oct. 2023, doi: 10.36040/jati.v7i3.6879.
- [18] Kaka Kalam Djati Permana, Esa Firmansyah, and Beben Sutara, “Analisis Komparatif Linear Regression dan Support Vector Regression pada Prediksi Penduduk Kabupaten Sumedang,” *Data Sciences Indonesia (DSI)*, vol. 5, no. 2, pp. 84–92, Dec. 2025, doi: 10.47709/dsi.v5i2.7423.
- [19] R. H. Nindy, H. Iftiah, A. Waladi, and Y. Perdana, “Implementation of Machine Learning for Stock Price Prediction Using the LSTM Algorithm,” *MISTIC: Media Applied Science and Engineering*, vol. 1, no. 2, pp. 31–37, 2025.
- [20] K. G. Olivares, C. Challu, G. Marcjasz, R. Weron, and A. Dubrawski, “Neural basis expansion analysis with exogenous variables: Forecasting electricity prices with NBEATSx,” *Int J Forecast*, vol. 39, no. 2, pp. 884–900, Apr. 2023, doi: 10.1016/j.ijforecast.2022.03.001.
- [21] Y. Shangguan *et al.*, “Transfer Learning-Enhanced N-BEATSx for Multivariate Forecasting of Tight Gas Well Production,” *Electronics (Basel)*, vol. 14, no. 19, p. 3875, Sep. 2025, doi: 10.3390/electronics14193875.
- [22] A. Handijono and Z. Suhatman, “Penerapan Machine Learning untuk Mempertahankan Loyalitas Pelanggan menggunakan Churn Prediction,” *Spectrum*, vol. 2, no. 3, pp. 259–266, Jul. 2025.
- [23] H. J. Weerts, A. C. Mueller, and J. Vanschoren, “Importance of tuning hyperparameters of machine learning algorithms,” *arXiv preprint arXiv:2007.07588*, 2020. doi: 10.48550/arXiv.2007.07588.
- [24] M. S. Alfa, A. Ibrahim, and M. I. Umar, “The Comparison of Various Exponential Smoothing Models and Their Significance in Modern Time Series Forecasting,” *Afropolitan Journals*, vol. 8, no. 1, pp. 8–17, 2022.
- [25] M. A. Fawaid, R. A. Pramunendar, N. A. Sri Winarsih, M. S. Rohman, and D. O. Ratmana, “Optimasi Model U-Net Backbone Resnet50 Pada Segmentasi Citra Banjir Menggunakan Sequential Hyperparameter Tuning,” *rabit*, vol. 11, no. 1, pp. 1478–1488, Jan. 2026.
- [26] Prapas, I., Derakhshan, B., Mahdiraji, A.R. et al. Continuous Training and Deployment of Deep Learning Models. *Datenbank Spektrum* 21, 203–212 (2021). <https://doi.org/10.1007/s13222-021-00386-8>
- [27] M. A. Akmal, Trimono, and A. R. Pratama, “Penerapan LSTM untuk Prediksi Jangka Pendek Penjualan Paket Data (Studi Kasus: Kota Malang),” *Discover - Data Science, Informatics, Statistics, & Electrical Research Conference (2025)*, vol. 1, no. 1, pp. 19–27, 2025.
- [28] N. Mehdiyev, D. Enke, P. Fettke, and P. Loos, “Evaluating Forecasting Methods by Considering Different Accuracy Measures,” *Procedia Comput Sci*, vol. 95, pp. 264–271, 2016, doi: 10.1016/j.procs.2016.09.332