JNEST EISSN: 2961-8738 ● PISSN: 2961-8916 ● DOI: 10.56741/jnest.v5i03.2511
Article Received 15 July 2026, Revised 25 August 2026

LASSO-PCA-LSTM Framework for Predicting Indonesian Government Bond Yields with Bayesian-Optimized LSTM

Amanda Kayla Putri Wibowo*, Deni Saepudin

***Abstract*—Accurate government bond yield forecasting is important for investment decision-making and financial risk management, yet it remains challenging because of nonlinear yield behavior and complex macroeconomic interactions. This study proposes a hybrid forecasting framework combining Least Absolute Shrinkage and Selection Operator (LASSO), Principal Component Analysis (PCA), Bayesian Optimization, and Long Short-Term Memory (LSTM) to predict Indonesian fixed-rate government bond yields (FR0074, FR0075, and FR0076). Historical bond prices were converted into yield-to-maturity values and integrated with nine domestic and global macroeconomic indicators. Four forecasting scenarios—LSTM, LASSO-LSTM, PCA-LSTM, and LASSO-PCA-LSTM—were evaluated across three prediction horizons (t+1, t+5, and t+30). LASSO identified four influential predictors: the USD/IDR exchange rate, BI Rate, foreign exchange reserves, and the Fed Funds Rate. The LASSO-PCA-LSTM framework achieved the strongest predictive performance at the t+1 and t+5 horizons, with average R^2 values of 0.8685 and 0.1760, respectively, whereas PCA-LSTM performed relatively better at t+30. These findings indicate that the effectiveness of feature selection and dimensionality reduction varies across forecasting horizons. By combining predictor selection, dimensionality reduction, and Bayesian hyperparameter optimization, the proposed framework preserves relevant macroeconomic information while reducing redundancy. The results demonstrate its potential for government bond yield forecasting and for supporting investment analysis, portfolio management, and monetary policy assessment in emerging financial markets.**

***Index Terms*—Bayesian Optimization, Bond Yield Forecasting, Feature Selection, Government Securities, LASSO, Long Short-Term Memory (LSTM), Principal Component Analysis (PCA), Time Series Forecasting.**

I. Introduction

Government bonds are among the main financial instruments used by sovereign governments to finance fiscal deficits, infrastructure projects, and public expenditure, while also supporting financial market stability [1]–[2]. In Indonesia, government securities (Surat Utang Negara or SUN) are issued in several forms, including fixed-rate (FR), variable-rate (VR), and sharia-compliant bonds (sukuk) [3]. Fixed-rate (FR) government bonds are particularly attractive to investors because they provide relatively predictable returns, involve comparatively low investment risk, and can be purchased with accessible minimum investment amounts by retail investors [4]–[5]. Government bonds are also generally considered relatively safe assets during periods of economic uncertainty and therefore play an important role in portfolio diversification and fixed-income investment strategies [2], [6].

Amanda Kayla Putri Wibowo and **Deni Saepudin** are with the Faculty of Informatics, Telkom University, Bandung, Indonesia.
Corresponding author: *amandakayla696@gmail.com

Yield to Maturity (YTM) is an important measure of bond performance as it represents the return an investor is expected to receive when a bond is held until maturity [7]. Bond yields are not only relevant to investment decisions but also provide information for portfolio allocation, risk management, and monetary policy assessment [7]–[8]. At the same time, government bond yields are affected by a range of domestic and international macroeconomic conditions, including exchange rates, inflation, monetary policy, global interest rates, and broader financial market conditions [9]–[10]. Changes in interest rates, in particular, tend to be accompanied by changes in bond yields, with yields generally increasing when interest rates rise and declining when monetary policy becomes more accommodative [11]. The interaction between these factors creates dynamic and nonlinear patterns in bond yield movements, which makes accurate forecasting a challenging task [12].

Traditional forecasting methods, including the Autoregressive Integrated Moving Average (ARIMA) model, have been widely used to predict financial time-series data [13]. However, these approaches are generally more suited to relatively simple temporal relationships and may have difficulty representing the nonlinear dependencies and longer-term patterns observed in bond yield movements [13]. This limitation has encouraged the use of deep learning models that can learn complex patterns directly from sequential data. One of the most widely used approaches is Long Short-Term Memory (LSTM), a specialized Recurrent Neural Network (RNN) architecture. Through its memory-cell and gating mechanisms, LSTM can retain relevant information over longer sequences and is less susceptible to the vanishing-gradient problem associated with conventional RNNs [14]–[15]. These characteristics make LSTM particularly suitable for financial time-series forecasting, where past observations may influence future movements over different time periods.

A number of recent studies have examined the application of LSTM-based models to bond yield forecasting in different financial markets. Nunes et al. [16], for example, developed an LSTM-LagLasso framework to forecast Euro government bond yields and reported better predictive performance than multilayer perceptron and linear regression models. Shu and Chou [12] applied LSTM to U.S. Treasury yields and found that historical yield information could be used to capture complex movements in the yield series. In the Chinese bond market, Ping et al. [17] combined LASSO, PCA, and Bayesian

© 2026 IISTR

Optimization with LSTM and obtained higher prediction accuracy than several other deep learning models, including CNN, RNN, and Transformer. Li et al. [18] also found that incorporating macroeconomic indicators into an LSTM model improved forecasting performance compared with conventional recurrent neural networks. Taken together, these studies suggest that combining appropriate feature processing with deep learning can improve the modeling of complex financial time-series data.

However, several aspects of bond yield forecasting remain insufficiently addressed. First, most existing LSTM-based studies have examined developed financial markets, particularly the United States and Europe, or larger emerging markets such as China, while research specifically focused on Indonesian government bond yields is still relatively limited [12], [16]–[18]. Second, existing studies tend to emphasize overall forecasting accuracy, with less attention given to the role of feature selection and dimensionality reduction in forecasting performance across different horizons. This issue is particularly relevant when macroeconomic variables contain overlapping or redundant information. Third, research on Indonesian government bond yields has largely relied on conventional econometric approaches, which may not fully account for the nonlinear and multivariate characteristics of financial time-series data [20]. Therefore, the application of an integrated deep learning framework that combines feature selection, dimensional.

To address these limitations, this study proposes an integrated forecasting framework that combines Least Absolute Shrinkage and Selection Operator (LASSO) for feature selection, Principal Component Analysis (PCA) for dimensionality reduction, Bayesian Optimization for hyperparameter tuning, and Long Short-Term Memory (LSTM) for bond yield prediction. The proposed framework is evaluated using three actively traded Indonesian government bond series (FR0074, FR0075, and FR0076) under three prediction horizons (daily, weekly and monthly). Four modeling scenarios which are LSTM, LASSO-LSTM, PCA-LSTM, and LASSO-PCA-LSTM are systematically compared to evaluate the individual and combined effects of feature engineering on forecasting performance. Furthermore, this study identifies the macroeconomic variables that most strongly influence Indonesian government bond yields and analyzes how different feature engineering strategies affect prediction performance across daily, weekly and monthly m forecasting horizons.

The main contributions of this study are threefold. First, it develops an integrated LSTM-based forecasting framework for Indonesian government bond yields by incorporating LASSO feature selection, PCA dimensionality reduction, and Bayesian Optimization into a unified modeling pipeline. Second, it provides a comprehensive comparative evaluation of four forecasting frameworks across multiple government bond series and forecasting horizons. Third, it offers empirical evidence regarding the effectiveness of feature engineering techniques in improving bond yield prediction while identifying the macroeconomic variables that play dominant roles in the Indonesian government bond market.

II. Related Work

Government bond yields are affected by various domestic and global macroeconomic conditions, such as monetary policy, exchange rates, inflation, and financial market developments [21]. Wulandari et al. [9] found that both domestic and external macroeconomic variables have a significant influence on Indonesian government bond yields. Their results suggest that forecasting models may benefit from considering a broader set of economic indicators rather than relying on a single group of variables. Akram and Uddin [10] also found that movements in long-term interest rates in emerging markets are closely related to monetary policy, particularly changes in policy interest rates. Taken together, these studies suggest that bond yields are shaped by several interacting economic factors and are therefore more appropriately examined using a multivariate time-series approach.

The use of deep learning has provided another way to model financial time-series data, particularly when the relationships between variables are nonlinear and extend over different time periods. LSTM has received considerable attention in this area because of its ability to retain information from previous observations. Shu and Chou [12] applied LSTM to U.S. Treasury yield forecasting and showed that the model was able to capture the temporal movements of the yield series with high predictive accuracy. Nunes et al. [16] developed an LSTM-LagLasso model for Euro government bond yields using financial market and macroeconomic variables. Their model produced better forecasts than multilayer perceptron and conventional linear regression models. Li et al. [18] incorporated macroeconomic indicators into an LSTM model for U.S. bond yield forecasting and reported better performance than conventional recurrent neural networks. These studies provide evidence that LSTM can be useful for forecasting bond yields when the data contain complex temporal relationships.

The selection and transformation of input variables are also important when developing forecasting models. Financial datasets often contain predictors with overlapping information, which can increase model complexity without necessarily improving prediction accuracy. Ping et al. [17] addressed this issue by combining LASSO, PCA, and Bayesian Optimization with LSTM to forecast Chinese government bond yields. The resulting framework performed better than BP, CNN, RNN, and Transformer models. Yang et al. [19] similarly applied LASSO before LSTM modeling for stock price prediction and found that removing less relevant variables improved prediction accuracy. Kostyra [19] examined a PCA-LSTM approach for daily yield curve forecasting and reported better performance than Random Walk and autoregressive models, suggesting that dimensionality reduction can be useful while retaining the information most relevant to forecasting.

The configuration of the LSTM model itself can also affect its predictive performance. Bayesian Optimization provides a systematic way to search for suitable

hyperparameter combinations without having to evaluate every possible configuration. Idham Maulana et al. [20] applied Bayesian Optimization to LSTM and obtained higher forecasting accuracy than several machine learning methods, including RNN, Bi-LSTM, K-Nearest Neighbor, Decision Tree, and Random Forest. Overall, the existing literature points to three relevant aspects of financial forecasting: the importance of incorporating macroeconomic information, the ability of LSTM to model temporal dependencies, and the potential benefits of combining feature selection, dimensionality reduction, and hyperparameter optimization. These aspects provide the basis for developing an integrated forecasting framework for Indonesian government bond yields.

Existing research has established the usefulness of LSTM and various feature engineering and optimization techniques in financial forecasting. However, their combined application to government bond yield forecasting, particularly in the Indonesian market, remains relatively limited. Several related research gaps therefore warrant further investigation. Most existing studies focus on bond markets in the United States, Europe, and China [12], [16]–[18], while empirical evidence for Indonesian government bond yields remains limited. Furthermore, previous studies generally evaluate a single modeling framework or emphasize overall prediction accuracy without systematically investigating how feature selection and dimensionality reduction interact across different forecasting horizons [17], [19]–[20]. Motivated by these limitations, this study develops an integrated forecasting framework that combines LASSO, PCA, Bayesian Optimization, and LSTM for Indonesian government bond yield prediction and systematically compares four modeling scenarios across daily, weekly and monthly term forecasting horizons.

III. Research Methods

A. Overall Framework

This study proposes an integrated forecasting framework for predicting Indonesian government bond yields using Long Short-Term Memory (LSTM), supported by LASSO feature selection, Principal Component Analysis (PCA), and Bayesian Optimization. The framework is designed to systematically evaluate the individual and combined effects of feature engineering techniques on forecasting performance across different prediction horizons. The study is conducted through five main stages: data collection, data preprocessing, feature engineering, model development, and performance evaluation. The workflow of the proposed framework, including the sequence of these stages, is presented in Fig. 1.

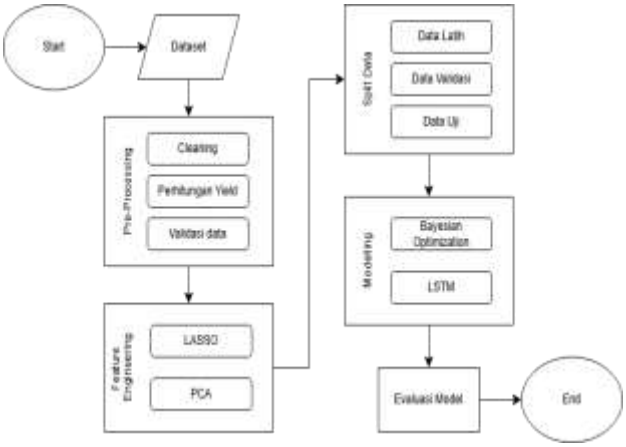


Fig. 1. System workflow for Indonesian government bond yield prediction using LSTM with LASSO feature selection, PCA dimensionality reduction, and Bayesian hyperparameter optimization

As shown in Fig. 1, the study starts with the collection of historical government bond prices and macroeconomic variables from publicly available sources. The data are prepared for modeling by calculating bond yields, interpolating missing values, and applying normalization to maintain consistency across variables. The next stage involves feature engineering, where LASSO is used to select relevant predictors and PCA is applied to reduce the dimensionality of the input data. Based on these steps, four forecasting scenarios are developed: LSTM, LASSO-LSTM, PCA-LSTM, and LASSO-PCA-LSTM. All four models are optimized using Bayesian Optimization and tested under the same experimental conditions to provide a consistent basis for performance comparison.

To provide a clearer representation of the proposed prediction framework, the detailed modeling pipeline is presented in Fig. 2.

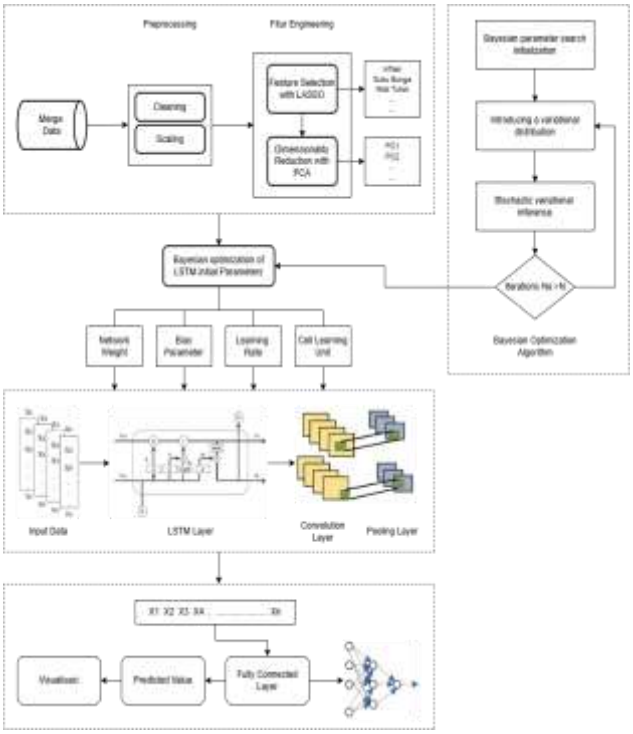


Fig. 2. Detailed modeling pipeline integrating data preprocessing, LASSO, PCA, Bayesian Optimization, LSTM training, and evaluation

Fig. 2 illustrates the complete implementation process of the proposed forecasting system. Historical bond prices are

first transformed into Yield to Maturity (YTM) values before being integrated with macroeconomic variables. Depending on the experimental scenario, the input features are either directly used, filtered through LASSO, transformed using PCA, or processed using both techniques simultaneously. The processed features are then supplied to the LSTM model whose hyperparameters are optimized using Bayesian Optimization. Finally, model performance is assessed using multiple evaluation metrics over three forecasting horizons (daily, weekly and monthly).

B. Dataset

The proposed forecasting framework utilizes two categories of data, namely Indonesian government bond data and macroeconomic variables. Government bond data consist of three actively traded fixed-rate government securities, FR0074, FR0075, and FR0076, which were selected because they represent different maturity characteristics while maintaining sufficient historical observations for time-series forecasting. Historical bond prices were converted into Yield to Maturity (YTM), which was used as the target variable in the forecasting models. The YTM observations are available on a daily basis and cover the period from September 2017 to April 2025.

Nine macroeconomic variables were included as explanatory variables to account for the potential effects of domestic and global economic conditions on bond yield movements. The selection of these variables was guided by previous research on the determinants of government bond yields and broader financial market dynamics [9]–[10]. A detailed description of the variables included in the analysis is provided in Table I.

TABLE I. MACROECONOMIC VARIABLES USED IN THIS STUDY

Variable	Symbol
USD/IDR Exchange Rate	KURS
Inflation	INF
BI Rate	BIR
Composite Stock Price Index	IHSG
Gross Domestic Product	PDB
Foreign Exchange Reserves	CDV
Fed Funds Rate	FFR
World Oil Price (WTI)	OIL
U.S. Treasury Yield	USY

As presented in Table I, the explanatory variables include both domestic and international macroeconomic indicators. The domestic variables describe Indonesia's economic and monetary conditions, while the international variables represent external financial conditions that may affect capital flows and government bond markets. Including both groups of variables allows the forecasting models to account for the different economic factors associated with movements in Indonesian government bond yields.

C. Data Preprocessing

Data preprocessing was carried out to improve data quality and maintain consistency before the forecasting models were developed. The process included YTM computation, missing-value treatment, feature scaling, and chronological data partitioning. The same preprocessing procedures were applied to all experimental scenarios so that differences in forecasting results could be attributed to the modeling frameworks rather than differences in data preparation.

The initial dataset contained historical prices of Indonesian government bonds together with several macroeconomic variables obtained from publicly available sources. Because bond prices were used as the primary financial measure, they were first converted into Yield to Maturity (YTM), which represents the return expected from holding a bond until its maturity date. The resulting YTM values were then compared with published market yields to assess the accuracy of the calculation.

The calculated yields showed close agreement with the published market yields. This indicates that the YTM calculation was sufficiently accurate for use in the subsequent forecasting experiments. Among the three-bond series, FR0076 had the lowest average absolute error, followed by FR0075 and FR0074.

The macroeconomic variables were obtained from different institutions and therefore did not always follow the same publication schedule. As a result, some observations contained missing values. These missing observations were treated using interpolation to maintain the continuity of the time series and reduce the loss of information.

After the missing values were addressed, the variables were scaled before being used for model training. The explanatory variables were standardized using StandardScaler, resulting in a mean of zero and a standard deviation of one. The target yield values, on the other hand, were transformed to a range between 0 and 1 using Min-Max Scaling. Standardization helps place the explanatory variables on a comparable numerical scale, while Min-Max normalization is useful for the gradient-based optimization used during LSTM training.

The processed data were divided chronologically into three subsets: 60% for training, 20% for validation, and 20% for testing. Random shuffling was not applied because doing so could disrupt the temporal order of the observations. The training data were used to fit the forecasting models, while the validation data were used for model development and hyperparameter optimization. The final 20% of observations were kept as an unseen test set and were used only for the final evaluation. In this arrangement, the validation period follows the training period, and the test period follows the validation period, preventing future observations from being introduced during model development. The same 60:20:20 chronological split was maintained across all forecasting scenarios and horizons to ensure a consistent basis for comparison.

D. Feature Engineering

Feature engineering was used to examine whether removing less relevant information and reducing redundancy among predictors could improve bond yield forecasting. Two techniques were applied for this purpose: Least Absolute Shrinkage and Selection Operator (LASSO) for feature selection and Principal Component Analysis (PCA) for dimensionality reduction.

LASSO regression uses L1 regularization to shrink the regression coefficients of less informative variables toward zero, allowing the more relevant predictors to be retained. The regularization parameter was selected through cross-validation based on the value that produced the lowest prediction error. The resulting cross-validation performance is shown in Fig. 3.

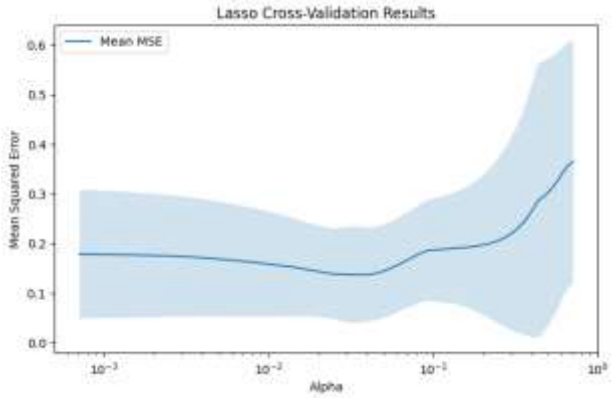


Fig. 3. LASSO cross-validation curve for selecting the optimal regularization parameter

Fig. 3 presents the cross-validation results used to determine the optimal regularization parameter. At the selected parameter value, LASSO produces a smaller set of macroeconomic variables by shrinking the coefficients of less relevant predictors toward zero. This reduces the number of inputs used in the subsequent forecasting models while retaining the predictors selected by the LASSO procedure. The variables retained after LASSO selection are listed in Table II.

TABLE II. MACROECONOMIC VARIABLES SELECTED BY LASSO

No.	Variable
1.	USD/IDR Exchange Rate
2.	BI Rate
3.	Foreign Exchange Reserves
4.	Fed Funds Rate

Table II shows the four macroeconomic variables retained by LASSO and subsequently used in the forecasting analysis. The selected variables include indicators related to both domestic monetary conditions and external financial factors. This combination allows the models to consider information from the Indonesian economy as well as conditions in the international financial market when forecasting government bond yields.

PCA was applied after the LASSO selection step to examine whether the input data could be represented using fewer dimensions without losing much of the information contained in the original variables. In this process, the original variables are transformed into a smaller set of principal components that are not correlated with one another. This transformation can help simplify the model inputs and reduce problems caused by highly correlated predictors. For the different experimental scenarios, PCA was applied either to all available explanatory variables or only to the variables retained by LASSO.

E. LSTM Architecture

LSTM was selected as the main forecasting model because it can learn temporal patterns and nonlinear relationships from sequential financial data [14]–[15]. This is particularly relevant to government bond yields, where current movements may be related to information contained in previous observations as well as changes in macroeconomic conditions. Compared with a conventional Recurrent Neural Network (RNN), LSTM uses memory cells together with input, forget, and output gates to control which information is retained or discarded as the sequence progresses. This mechanism allows the model to preserve useful information over longer sequences and reduces the vanishing-gradient issue that can limit standard RNNs.

The architecture used in the proposed forecasting framework is presented in Fig. 4.

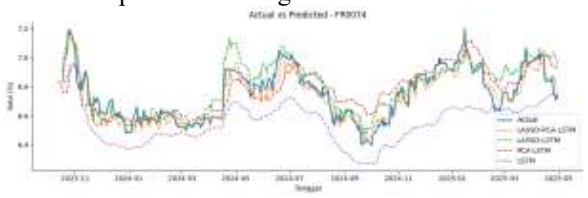


Fig. 4. Proposed LSTM architecture for Indonesian government bond yield forecasting

Fig. 4 shows how the input layer receives sequential observations of historical bond yields and macroeconomic variables. The variables provided to the LSTM differ across the four experimental scenarios. The model either uses the original set of variables, the variables retained by LASSO, the principal components obtained from PCA, or the reduced set produced by applying LASSO and PCA together. These sequential inputs are processed through one or more LSTM layers to extract temporal dependencies before being forwarded to a fully connected output layer that generates the predicted bond yield. The same network architecture was adopted for all experimental scenarios to ensure that performance differences originated from feature engineering rather than architectural modifications.

The prediction framework was evaluated under three forecasting horizons, namely daily ahead, weekly ahead, and monthly days ahead. This experimental design enables the proposed framework to assess the robustness of each modeling strategy across daily, weekly and monthly term forecasting tasks.

F. Experimental Setup

To objectively evaluate the proposed forecasting framework, four experimental scenarios were constructed incrementally by combining feature engineering techniques with the LSTM prediction model. The baseline scenario employed a conventional LSTM model using all available input variables. The remaining scenarios progressively incorporated LASSO feature selection, PCA dimensionality reduction, and the combination of both techniques. This experimental design enables the individual and combined contributions of feature engineering to be systematically evaluated under identical forecasting conditions. The four forecasting scenarios evaluated in this study are summarized in Table III.

As presented in Table III, the experimental scenarios were designed incrementally to isolate the contribution of each feature engineering technique. The baseline model serves as the reference, whereas the remaining models evaluate whether removing irrelevant variables, reducing

feature dimensionality, or combining both strategies improves forecasting performance.

TABLE III. EXPERIMENTAL SCENARIOS

Scenario	Description
LSTM	Baseline LSTM using all input variables
LASSO-LSTM	LSTM with LASSO feature selection
PCA-LSTM	LSTM with PCA dimensionality reduction
LASSO-PCA-LSTM	LSTM with LASSO feature selection followed by PCA

To ensure a fair comparison, all experiments were performed using identical training and testing procedures. Forecasting performance was evaluated for three prediction horizons, namely daily, weekly and monthly. All models were implemented using the same preprocessing pipeline and evaluation metrics.

G. Hyperparameter Optimization

The predictive performance of LSTM is highly dependent on the selection of hyperparameters. Therefore, Bayesian Optimization was employed to automatically determine the optimal hyperparameter configuration instead of relying on manual trial-and-error procedures. Bayesian Optimization iteratively explores the search space by balancing exploration and exploitation, thereby reducing computational cost while efficiently identifying high-performing parameter combinations.

The optimization process considered several hyperparameters associated with the LSTM network, including the number of hidden units, lookback window size, dropout rate, batch size, learning rate, and training epochs. For each experimental scenario, Bayesian Optimization searched for the hyperparameter combination that minimized the forecasting error on the validation set. The validation set consisted of the chronologically subsequent 20% of observations following the training set and was used exclusively for hyperparameter selection. The test set was not involved in the optimization process and was reserved for the final evaluation of the selected model.

The optimal hyperparameter configuration obtained for each forecasting scenario is presented in the corresponding result tables in Section IV. By employing Bayesian Optimization, all forecasting models were trained under optimized conditions, allowing a fair comparison among the four proposed modeling scenarios while minimizing the influence of manually selected parameter values.

H. Performance Evaluation

The forecasting performance of each model was evaluated using four complementary evaluation metrics, namely Mean Squared Error (MSE), Mean Absolute Error (MAE), Coefficient of Determination (R^2), and Explained Variance Score (EVS). Multiple evaluation metrics were employed because each metric captures different aspects of prediction performance.

Mean Squared Error (MSE) measures the average squared difference between predicted and actual values, making it sensitive to large prediction errors. Mean Absolute Error (MAE) quantifies the average magnitude of prediction errors regardless of direction, providing a more interpretable measure of forecasting accuracy. The Coefficient of Determination (R^2) evaluates the proportion of variance explained by the forecasting model, whereas Explained Variance Score (EVS) measures the ability of the model to preserve variations within the observed data.

The mathematical formulations of these evaluation metrics are presented in Equations (1)–(4).

$$R^2 = 1 - (\Sigma(y_i - \hat{y}_i)^2) / (\Sigma(y_i - \bar{y})^2) \quad (1)$$

where y_i is the actual value, $\hat{y}_i$ is the predicted value, $\bar{y}$ is the mean of actual values, and n is the number of observations.

Mean Squared Error (MSE) quantifies the average squared prediction error:

$$MSE = (1/n)\ \Sigma(y_i - \hat{y}_i)^2 \quad (2)$$

where y_i is the actual value, $\hat{y}_i$ is the predicted value, and n is the number of observations.

Mean Absolute Error (MAE) provides an interpretable measure of average absolute deviation in yield units:

$$MAE = (1/n)\ \Sigma|y_i - \hat{y}_i| \quad (3)$$

where y_i is the actual value, $\hat{y}_i$ is the predicted value, and n is the number of observations.

Explained Variance Score (EVS) evaluates the proportion of target variance captured by predictions:

$$EVS = 1 - Var(y - \hat{y}) / Var(y) \quad (4)$$

where $Var(\cdot)$ denotes the variance of the actual or residual values.

To comprehensively evaluate the proposed framework, all four forecasting scenarios were assessed using identical evaluation metrics under three prediction horizons (daily, weekly and monthly). This evaluation strategy enables an objective comparison of forecasting accuracy while providing insight into the effectiveness of feature selection and dimensionality reduction under different forecasting periods.

IV. RESULT AND DISCUSSION

A. Yield-to-Maturity (YTM) Calculation

Since historical bond prices were used as the primary financial data, they were first converted into Yield-to-Maturity (YTM), which served as the prediction target throughout this study. To verify the reliability of the implemented calculation, the computed YTM values were compared with the published market yields for each government bond series. The average absolute error between the calculated and published yields is presented in Table IV.

TABLE IV. AVERAGE ABSOLUTE ERROR OF YTM CALCULATION

Bond Series	Average Absolute Error
FR0074	0.0105
FR0075	0.0044
FR0076	0.0018

As presented in Table IV, the calculated YTM values closely match the published market yields for all three government bond series, indicating that the implemented yield calculation provides satisfactory accuracy. Among the evaluated bonds, FR0076 produced the smallest average absolute error (0.0018), followed by FR0075 (0.0044) and FR0074 (0.0105). These results confirm that the calculated YTM values are sufficiently accurate to be used as target variables in the subsequent forecasting experiments.

The calculated YTM values for each Indonesian government bond series are further illustrated in Fig. 5.



Fig. 5. Comparison of Yield-to-Maturity values for FR0074, FR0075, and FR0076

Fig. 5 illustrates the temporal movement of the calculated YTM values for the three government bond series during the observation period. Although the three bond series tend to move in the same general direction, the size and volatility of their yield movements are not identical. Differences in maturity, coupon characteristics, and investor demand may contribute to these variations. For this reason, using several government bond series provides a broader basis for evaluating the forecasting framework than focusing on a single bond.

The fluctuations observed in the yield series also show that bond yields do not follow a simple and stable pattern over time. Their movements can change in response to evolving market and economic conditions, resulting in dynamic relationships within the time series. This provides a basis for using LSTM models, which are designed to learn temporal patterns from sequential observations.

B. LSTM Model

This model served as the reference for assessing whether the additional feature engineering techniques could improve forecasting performance. Before training, Bayesian Optimization was used to search for a suitable combination of LSTM hyperparameters. The resulting configuration is reported in Table V.

TABLE V. OPTIMAL HYPERPARAMETERS OF THE LSTM MODEL

Hyperparameter	Daily	Monthly	Weekly
Lookback	10	30	12
LSTM Units	50	47	34
Dropout	0.2099	0.3465	0.3330
Batch Size	37	45	26
Learning Rate	0.000982	0.000268	0.002913

Table V presents the hyperparameter configuration selected through Bayesian Optimization for the baseline LSTM model. These parameters were used during model training across the three forecasting horizons. Using an automated optimization procedure provides a consistent basis for selecting the model configuration and avoids relying on manually chosen parameter values.

The baseline LSTM forecasting results for the t+1 horizon are presented in Table VI.

TABLE VI. PERFORMANCE OF THE LSTM MODEL FOR THE DAILY PREDICTION HORIZON

Bond Series	MSE	MAE	R^2	EVS
FR0074	0.051483	0.205282	-0.6915	0.6687
FR0075	0.019283	0.119156	0.2207	0.7341
FR0076	0.017732	0.116592	-0.2805	0.6189

Table VII shows that the baseline LSTM produced different results across the three-bond series at the t+1 horizon. FR0075 recorded a positive R^2 of 0.2207, while FR0074 and FR0076 produced negative R^2 values of −0.6915 and −0.2805, respectively. The MSE and MAE values were relatively moderate, but the negative R^2 values for FR0074 and FR0076 indicate that the model was not able to explain the observed yield variation consistently across all three-bond series. The baseline LSTM can therefore serve as a reference model, although its predictive performance differs depending on the bond series.

The results also suggest that the model was able to capture part of the short-term movement in bond yields, particularly for FR0075. However, the negative R^2 values obtained for the other two series show that this performance was not consistent across the dataset. Therefore, the baseline results provide a useful point of comparison for assessing whether the addition of LASSO, PCA, and their combination can improve the forecasting results.

The forecasting performance at the t+5 horizon is summarized in Table VII.

TABLE VII. PERFORMANCE OF THE LSTM MODEL FOR THE WEEKLY PREDICTION HORIZON

Bond Series	MSE	MAE	R^2	EVS
FR0074	0.021982	0.114364	0.2392	0.2895
FR0075	0.021962	0.117283	0.0617	0.0932
FR0076	0.014357	0.097435	-0.0167	0.2312

Table VII shows a decline in forecasting performance as the horizon moves from t+1 to t+5. Compared with the one-day horizon, the model records lower R^2 values and higher prediction errors at t+5, indicating that forecasting becomes more difficult as the prediction period gets longer. One possible explanation is that bond yields over a longer period are affected by changes in macroeconomic conditions and financial market movements that cannot be fully captured by recent historical observations alone.

The forecasting performance for the monthly prediction monthly horizon is presented in Table VIII.

TABLE VIII. PERFORMANCE OF THE LSTM MODEL FOR THE MONTHLY PREDICTION HORIZON

Bond Series	MSE	MAE	R^2	EVS
FR0074	0.040079	0.164401	-0.3672	0.0509
FR0075	0.038521	0.151972	-0.6682	-0.0645
FR0076	0.014783	0.101663	-0.0342	0.0749

According to Table VIII, the baseline LSTM model achieves the lowest forecasting performance for the longest prediction horizon. The deterioration in R^2 and the corresponding increase in prediction errors indicate that long-term bond yield forecasting remains a challenging task because of the highly nonlinear and dynamic interactions among macroeconomic variables. For this reason, feature selection and dimensionality reduction were considered as additional steps that could help the forecasting models handle redundant information and focus on a smaller set of useful predictors.

The baseline LSTM serves as a reference point for assessing the effect of these feature engineering techniques. Its performance was relatively satisfactory at the shorter forecasting horizons but became weaker as the prediction period increased. This pattern provides a basis for examining whether the addition of feature selection and dimensionality reduction can improve forecasting performance, particularly at the t+5 and t+30 horizons.

C. LASSO–LSTM Model

To examine the role of feature selection in bond yield forecasting, LASSO was applied before the LSTM models were trained. Through L1 regularization, LASSO shrinks the coefficients of less informative variables toward zero, leaving a smaller set of predictors for the forecasting model. The variables retained by LASSO were then used as the LSTM inputs. Bayesian Optimization was applied separately to identify a suitable hyperparameter configuration for each forecasting horizon. The selected macroeconomic variables are presented in Table IX.

TABLE IX. MACROECONOMIC VARIABLES SELECTED BY LASSO

No.	Variable
1.	USD/IDR Exchange Rate
2.	BI Rate
3.	Foreign Exchange Reserves
4.	Fed Funds Rate

Table IX presents the four macroeconomic variables retained by LASSO for the forecasting models. The selected variables cover both domestic and external economic conditions. The USD/IDR exchange rate and the BI Rate provide information on domestic monetary and economic conditions, while foreign exchange reserves and the Fed Funds Rate reflect developments in the external financial environment. Reducing the number of input variables is expected to improve learning efficiency while minimizing the impact of irrelevant information on the forecasting model.

The optimal hyperparameter configuration obtained after applying LASSO feature selection is summarized in Table X.

TABLE X. OPTIMAL HYPERPARAMETERS OF THE LASSO-LSTM MODEL

Hyperparameter	Daily	Monthly	Weekly
Lookback	5	30	8
LSTM Units	77	77	39
Dropout	0.2475	0.2475	0.2895
Batch Size	20	20	22
Learning Rate	0.001281	0.001281	0.000987

Table X shows that Bayesian Optimization selected different hyperparameter configurations for each forecasting horizon after feature selection was performed. The short-term model employed a shorter lookback window with a larger number of LSTM units to capture recent temporal dependencies, whereas the medium- and long-term models adopted different parameter settings to accommodate the increasing complexity of bond yield movements. These optimized hyperparameters were subsequently used in all LASSO-LSTM experiments.

The forecasting performance of the LASSO-LSTM model for the daily prediction horizon (t+1) is presented in Table XI.

TABLE XI. PERFORMANCE OF THE LASSO-LSTM MODEL FOR THE DAILY PREDICTION HORIZON

Bond Series	MSE	MAE	R^2	EVS
FR0074	0.008704	0.073575	0.7133	0.7899
FR0075	0.016087	0.111183	0.3510	0.8236
FR0076	0.016593	0.111125	-0.2004	0.5927

As shown in Table XI, incorporating LASSO feature selection substantially improved short-term forecasting performance compared with the baseline LSTM model. FR0074 achieved the highest predictive capability with an R^2 value of 0.7133, while FR0075 also maintained a positive coefficient of determination ($R^2 = 0.3510$) and the highest EVS value (0.8236). In contrast, FR0076 still produced a negative R^2 value, indicating that although feature selection reduced prediction errors, it was insufficient to fully capture the underlying yield dynamics of this bond series.

The forecasting performance for the weekly prediction weekly horizon is summarized in Table XII.

TABLE XII. PERFORMANCE OF THE LASSO-LSTM MODEL FOR THE WEEKLY PREDICTION HORIZON

Bond Series	MSE	MAE	R^2	EVS
FR0074	0.031958	0.143888	-0.0573	0.1474
FR0075	0.021979	0.123494	0.1045	0.1697
FR0076	0.024265	0.132263	-0.7514	0.1720

According to Table XII, forecasting performance declined as the prediction horizon increased to weekly. Only FR0075 maintained a positive R^2 value, whereas FR0074 and FR0076 produced negative coefficients of determination. These results suggest that although LASSO successfully removed irrelevant features, feature selection alone was insufficient to preserve predictive performance as the forecasting horizon became longer.

The forecasting performance for the monthly prediction monthly horizon is presented in Table XIII.

TABLE XIII. PERFORMANCE OF THE LASSO-LSTM MODEL FOR THE MONTHLY PREDICTION HORIZON

Bond Series	MSE	MAE	R^2	EVS
FR0074	0.144803	0.343700	-4.0116	-0.8821
FR0075	0.037020	0.155403	-0.5816	-0.4978
FR0076	0.020180	0.109311	-0.4291	0.0290

As presented in Table XIII, the LASSO-LSTM model experienced a substantial decline in performance for the longest forecasting horizon. All bond series produced negative R^2 values, with FR0074 exhibiting the greatest deterioration. Although FR0076 achieved the lowest prediction errors among the three series, its coefficient of determination remained negative, indicating limited explanatory capability. These findings suggest that aggressive feature selection may discard macroeconomic information that becomes increasingly important for long-term forecasting. Therefore, the next experiment investigates whether dimensionality reduction using Principal Component Analysis (PCA) can preserve more informative features while improving long-horizon forecasting performance.

D. PCA–LSTM Model

To evaluate the contribution of dimensionality reduction without feature selection, Principal Component Analysis (PCA) was applied before training the LSTM model. Unlike LASSO, which selects a smaller set of predictors by shrinking less relevant coefficients toward zero, PCA retains the information from the original macroeconomic variables by transforming them into a smaller number of principal components. The resulting components are orthogonal to one another and together retain most of the variation in the original data. This transformation reduces the correlation among the input variables and produces a more compact set of features for the LSTM model. The cumulative explained variance from the PCA procedure is presented in Fig. 6.

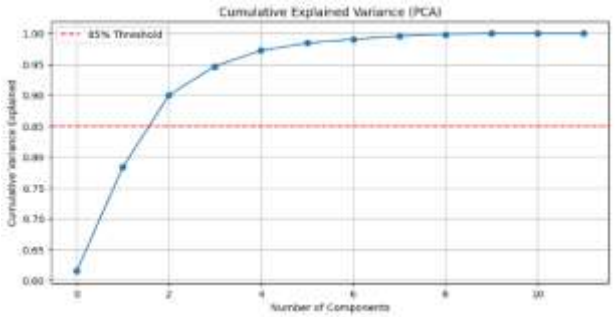


Fig. 6. Cumulative explained variance of PCA for the PCA–LSTM model

Fig. 6 shows that the cumulative explained variance reached the predefined threshold with the third principal component. The first three components accounted for approximately 90% of the total variance in the original dataset and were therefore retained for the subsequent modeling stage. This reduction means that much of the variation contained in the nine macroeconomic variables could be represented using only three principal components. The contribution of each principal component to the total variance is illustrated in Fig. 7.

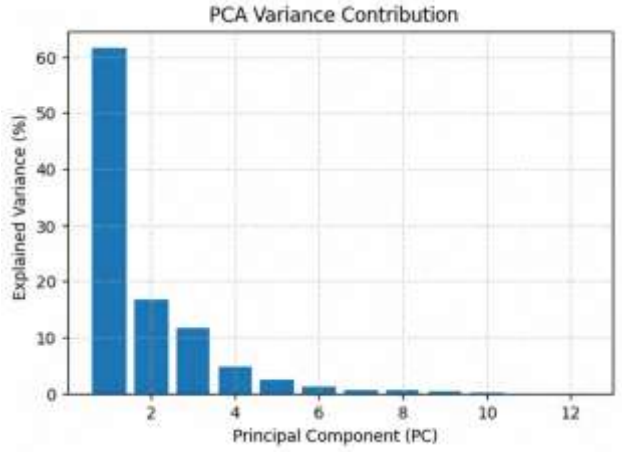


Fig. 7. Explained variance ratio of each principal component for the PCA–LSTM model

Fig. 7 shows that the first principal component accounted for the largest share of the explained variance. The second and third components added further information, and together the three components represented most of the variation in the original dataset. These components were therefore retained as the input features for the LSTM model.

To validate the number of retained principal components, a parallel analysis was conducted. The corresponding scree plot is presented in Fig. 8.

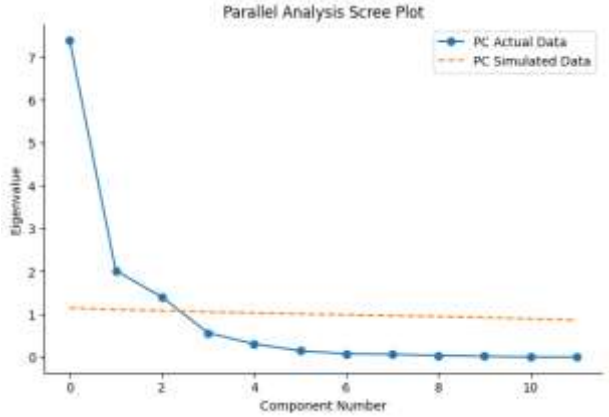


Fig. 8. Parallel analysis scree plot for the PCA–LSTM model

Fig. 8 confirms that the first three principal components possess eigenvalues greater than those generated through random simulation, indicating that only these components contain meaningful information. Consequently, three principal components were retained for the subsequent forecasting experiments.

The optimal hyperparameter configuration obtained for the PCA–LSTM model is summarized in Table XIV.

TABLE XIV. OPTIMAL HYPERPARAMETERS OF THE PCA-LSTM MODEL

Hyperparameter	Daily	Monthly	Weekly
Lookback	9	11	5
LSTM Units	34	32	68
Dropout	0.3168	0.2804	0.2013
Batch Size	19	23	40
Learning Rate	0.002241	0.001436	0.000572

As presented in Table XV, Bayesian Optimization produced different hyperparameter configurations for each prediction horizon, reflecting the varying temporal characteristics of daily, weekly and monthly term forecasting. These optimized parameters were subsequently adopted throughout the PCA–LSTM experiments to ensure optimal model performance.

The forecasting performance of the PCA–LSTM model for the daily prediction horizon is presented in Table XV.

TABLE XV. PERFORMANCE OF THE PCA-LSTM MODEL FOR THE DAILY PREDICTION HORIZON

Bond Series	MSE	MAE	R^2	EVS
FR0074	0.012770	0.091527	0.5770	0.5953
FR0075	0.008952	0.078873	0.6371	0.7768
FR0076	0.019501	0.116492	-0.4140	0.5155

Table XV shows that PCA–LSTM performed relatively well for FR0074 and FR0075 at the t+1 horizon, with R^2 values of 0.5770 and 0.6371, respectively. In contrast, the model recorded a negative R^2 of −0.4140 for FR0076, indicating that its predictions for this bond series were less effective in explaining the observed variation in yield. Thus, the PCA transformation improved predictive performance for some bond series but did not provide consistent explanatory capability across all three series.

TABLE XVI. PERFORMANCE OF THE PCA-LSTM MODEL FOR THE WEEKLY PREDICTION HORIZON

Bond Series	MSE	MAE	R^2	EVS
FR0074	0.037909	0.168322	-0.2542	0.3713
FR0075	0.070159	0.239853	-1.8585	0.4380
FR0076	0.018318	0.108413	-0.3221	0.0185

According to Table XVI, forecasting accuracy decreased considerably for the weekly prediction horizon. All three-bond series produced negative R^2 values, indicating that dimensionality reduction alone was insufficient to preserve predictive capability as the forecasting horizon increased. Although PCA successfully reduced feature redundancy, the transformed feature space may have omitted information required to capture medium-term bond yield fluctuations.

The forecasting performance of the PCA–LSTM model for the monthly prediction horizon is presented in Table XVII.

As presented in Table XVII, the PCA–LSTM model exhibited relatively better long-term forecasting performance than the previous LASSO-LSTM scenario.

TABLE XVII. PERFORMANCE OF THE PCA-LSTM MODEL FOR THE t+30 PREDICTION HORIZON

Bond Series	MSE	MAE	R^2	EVS
FR0074	0.048444	0.169911	-0.6525	-0.2879
FR0075	0.023848	0.120258	-0.0328	0.0892
FR0076	0.013378	0.095835	0.0641	0.0877

FR0074 and FR0075 still recorded negative R^2 values, whereas FR0076 achieved a positive R^2 of 0.0641. This made FR0076 the only bond series in the evaluated scenarios to obtain a positive coefficient of determination at the t+30 horizon. The result suggests that retaining information from the broader set of macroeconomic variables through PCA may be more suitable for this longer forecasting horizon than relying solely on the variables selected by LASSO. However, the relatively low R^2 values also show that the model still had limited explanatory power at t+30. Based on this result, the next experiment combines LASSO and PCA to examine whether using feature selection and dimensionality reduction together can provide more consistent forecasting performance across the different horizons.

E. LASSO–PCA–LSTM Model

The final experimental scenario combines LASSO feature selection and Principal Component Analysis (PCA) within a single forecasting framework before LSTM model training. In this approach, LASSO first eliminates less relevant macroeconomic variables, after which PCA transforms the retained variables into a lower-dimensional feature space. The objective of this integrated framework is to simultaneously reduce irrelevant information and multicollinearity while preserving the most informative characteristics of the input data. The selected variables remained identical to those identified in the previous LASSO experiment; therefore, the feature selection results refer to Table III, while the corresponding LASSO cross-validation curve and coefficient path are presented in Fig. 3 and Fig. 4, respectively.

Following feature selection, PCA was applied to the retained variables. The cumulative explained variance is illustrated in Fig. 9.

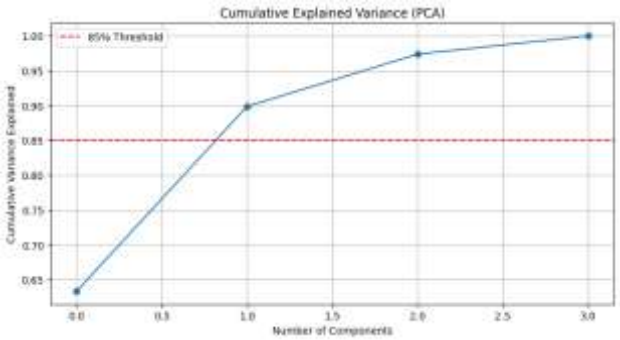


Fig. 9. Cumulative explained variance of PCA for the LASSO–PCA–LSTM model

Figure 9 shows that the cumulative explained variance passed the predefined threshold with the inclusion of the second principal component. The first two components together accounted for approximately 89.86% of the total variance in the selected variables, so two components were retained for the subsequent analysis. Compared with the PCA-LSTM scenario, fewer components were required because redundant variables had already been removed during the LASSO feature selection stage.

The explained variance contributed by each principal component is presented in Fig. 10.

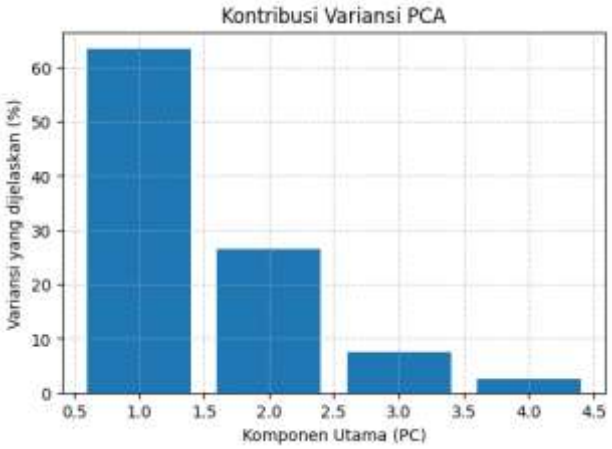


Fig. 10. Explained variance ratio of each principal component for the LASSO–PCA–LSTM model

Figure 10 shows that the first principal component contributed the largest proportion of the total variance, whereas the second principal component provided complementary information required to exceed the cumulative variance threshold. Consequently, these two principal components were selected as inputs to the LSTM model.

To validate the retained components, a parallel analysis was conducted. The corresponding scree plot is shown in Fig. 11.

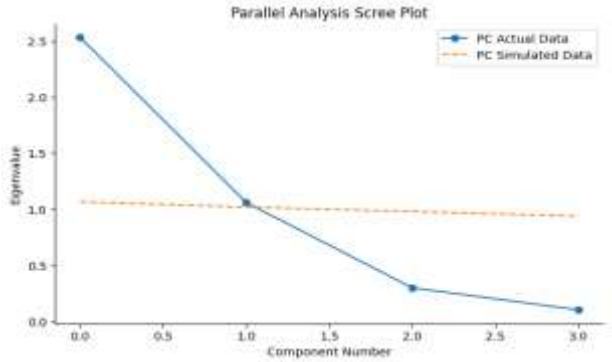


Fig. 11. Parallel analysis scree plot for the LASSO–PCA–LSTM model

As illustrated in Fig. 11, only the first two principal components exhibited eigenvalues greater than those generated from random simulation. This result confirms that retaining two principal components provides an appropriate representation of the selected macroeconomic variables while minimizing redundant information.

The optimal hyperparameter configuration obtained through Bayesian Optimization is summarized in Table XVIII.

TABLE XVIII. OPTIMAL HYPERPARAMETERS OF THE LASSO-PCA-LSTM MODEL

Hyperparameter	Daily	Monthly	Weekly
Lookback	8	8	12
LSTM Units	76	76	34
Dropout	0.2076	0.2076	0.3330
Batch Size	22	22	26
Learning Rate	0.002324	0.002324	0.002913

As presented in Table XVIII, Bayesian Optimization selected similar hyperparameter configurations for the

daily and weekly term forecasting horizons, whereas a different configuration was obtained for the monthly prediction horizon. These optimized parameters were subsequently used throughout the LASSO–PCA–LSTM experiments.

The forecasting performance of the proposed LASSO–PCA–LSTM model for the daily prediction horizon (t+1) is presented in Table XIX.

TABLE XIX. PERFORMANCE OF THE LASSO-PCA-LSTM MODEL FOR THE t+1 PREDICTION HORIZON

Bond Series	MSE	MAE	R^2	EVS
FR0074	0.004508	0.051493	0.8441	0.8867
FR0075	0.004309	0.052893	0.8159	0.8165
FR0076	0.001621	0.031586	0.8820	0.8829

As shown in Table XIX, the LASSO–PCA–LSTM model achieved the strongest performance among the evaluated scenarios at the t+1 horizon. All three-bond series produced positive R^2 values ranging from 0.8159 to 0.8820, together with relatively low MSE and MAE values. These results indicate that the combined feature-engineering approach provided strong predictive performance for the one-day-ahead horizon across all three- bond series.

The forecasting performance for the weekly prediction horizon is summarized in Table XX.

TABLE XX. PERFORMANCE OF THE LASSO-PCA-LSTM MODEL FOR THE t+5 PREDICTION HORIZON

Horizon	Bond Series	MSE	MAE	R^2	EVS
t+5	FR0074	0.030210	0.142865	0.0005	0.6181
	FR0075	0.016813	0.105585	0.3150	0.3701
	FR0076	0.010911	0.086309	0.2124	0.4828

According to Table XX, the LASSO–PCA–LSTM model achieved the highest average R^2 (0.1760) and EVS (0.4903) among the four evaluated models at the t+5 horizon, while also producing the lowest average MSE. The improvement at t+5 was noticeably smaller than that observed at t+1. The average R^2 was still relatively low, while both LASSO–LSTM and PCA–LSTM recorded negative average R^2 values. Although the combined feature-engineering approach performed better overall at t+5, the results show that the models had considerably less explanatory power at this horizon than at t+1.

The forecasting performance for the monthly prediction horizon (t+30) is presented in Table XXI.

TABLE XXI. PERFORMANCE OF THE LASSO-PCA-LSTM MODEL FOR THE MONTHLY PREDICTION HORIZON

Bond Series	MSE	MAE	R^2	EVS
FR0074	0.074990	0.217466	-1.5739	-0.9053
FR0075	0.050789	0.180017	-1.1685	-1.1611
FR0076	0.030129	0.135862	-1.1179	-0.9108

As presented in Table XXI, forecasting performance declined substantially for the monthly prediction horizon. All bond series produced negative R^2 values, indicating that the combined feature engineering strategy was less effective for long-term prediction than for daily and weekly term forecasting. This result may indicate that LASSO is useful for reducing less relevant variables at shorter forecasting horizons, but removing too many predictors may also reduce the amount of macroeconomic information available to the model when forecasting over longer periods. Consequently, the comparative analysis presented in the following subsection evaluates the overall performance of all four forecasting scenarios to determine the most suitable modeling strategy for each prediction horizon.

F. Comparative Performance Analysis

To comprehensively evaluate the proposed forecasting framework, the average performance of the four forecasting scenarios was compared across the three prediction horizons using Mean Squared Error (MSE), Mean Absolute Error (MAE), Coefficient of Determination (R^2), and Explained Variance Score (EVS). The comparison aims to identify the most suitable feature engineering strategy for daily, weekly and monthly term Indonesian government bond yield forecasting.

The average evaluation results for the daily prediction horizon are presented in Table XXII.

TABLE XXII. AVERAGE PERFORMANCE COMPARISON OF ALL MODELS FOR THE t+1 PREDICTION HORIZON

Model	MSE	MAE	R^2	EVS
LASSO–PCA–LSTM	0.002862	0.041217	0.8685	0.8926
LASSO–LSTM	0.013795	0.098628	0.2880	0.7354
PCA–LSTM	0.013741	0.095631	0.2667	0.6292
LSTM	0.029499	0.147010	-0.2504	0.6739

As presented in Table XXII, the proposed LASSO–PCA–LSTM model achieved the best overall performance for the one-day forecasting horizon. Compared with the baseline LSTM model, the proposed framework reduced the average MSE by more than 90% while increasing the average coefficient of determination from −0.2504 to 0.8685. These results demonstrate that integrating feature selection and dimensionality reduction substantially improves short-term bond yield forecasting by eliminating redundant variables and reducing multicollinearity before LSTM training.

In contrast, LASSO–LSTM and PCA–LSTM individually improved prediction accuracy but were unable to match the performance achieved by their combined implementation. The comparison of actual and predicted bond yields for all forecasting models at the daily prediction horizon is illustrated in Fig. 12.

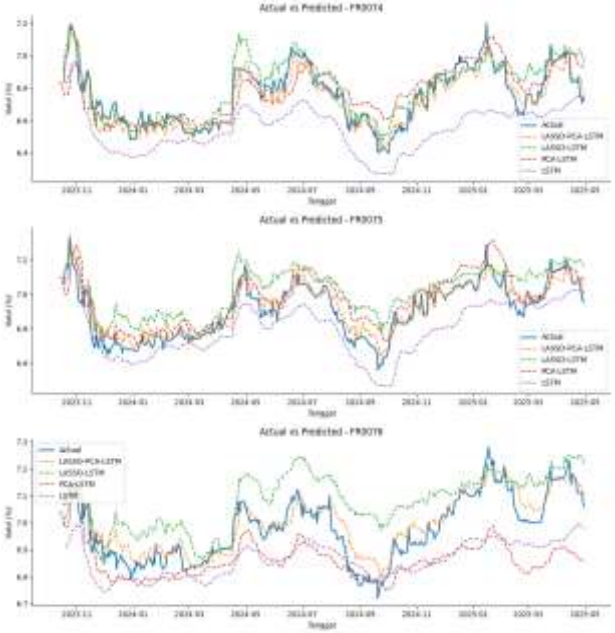


Fig. 12. Comparison of actual and predicted bond yields for all forecasting models at the daily prediction horizon

Fig. 12 provides a visual comparison of the predictions reported in Table XXII. Across the three government bond series, the LASSO–PCA–LSTM predictions generally follow the movement of the actual yields more closely. The baseline LSTM and PCA–LSTM models show larger differences from the observed values, particularly when the yields change more sharply, while LASSO–LSTM falls between the other models in terms of visual agreement. The pattern in the figure is consistent with the quantitative results and suggests that combining feature selection with dimensionality reduction can be useful for capturing short-term movements in bond yields.

The average forecasting performance at the t+5 horizon is reported in Table XXIII.

TABLE XXIII. AVERAGE PERFORMANCE COMPARISON OF ALL MODELS FOR THE t+5 PREDICTION HORIZON

Model	MSE	MAE	R^2	EVS
LASSO–PCA–LSTM	0.019311	0.111586	0.1760	0.4903
LSTM	0.019434	0.109694	0.0947	0.2046
LASSO–LSTM	0.026067	0.133215	-0.2347	0.1630
PCA–LSTM	0.042128	0.172196	-0.8116	0.2759

Table XXIII shows that LASSO–PCA–LSTM achieved the strongest overall results at the t+5 horizon, with the highest average R^2 of 0.1760 and EVS of 0.4903. The baseline LSTM model had a slightly lower average MAE, but its average R^2 was lower than that of LASSO–PCA–LSTM. Thus, although the difference in prediction error was relatively small, the integrated model was better at explaining the variation in bond yields at this forecasting horizon.

PCA–LSTM and LASSO–LSTM, in contrast, showed a decline in performance at t+5. This result suggests that using PCA or LASSO alone did not provide the same benefit as combining the two techniques for the five-day forecasting task. The prediction results for the four models at the t+5 horizon are presented in Fig. 13.



Fig. 13. Comparison of actual and predicted bond yields for all forecasting models at the weekly prediction horizon

Fig. 13 demonstrates that prediction accuracy generally decreased as the forecasting horizon increased. Nevertheless, the LASSO–PCA–LSTM model continued to follow the actual yield movements more closely than the remaining forecasting models. Larger deviations are observed for the PCA–LSTM model, particularly during periods of substantial yield fluctuations, confirming that dimensionality reduction alone is insufficient to maintain forecasting accuracy over medium-term prediction horizons.

The average forecasting performance for the monthly prediction horizon is presented in Table XXIV.

TABLE XXIV. AVERAGE PERFORMANCE COMPARISON OF ALL MODELS FOR THE t+30 PREDICTION HORIZON

Model	MSE	MAE	R^2	EVS
PCA–LSTM	0.028557	0.128668	-0.2071	-0.0370
LSTM	0.031127	0.139345	-0.3565	0.0204
LASSO–PCA–LSTM	0.051969	0.177781	-1.2868	-0.9924
LASSO–LSTM	0.067334	0.202805	-1.6741	-0.4503

As presented in Table XXIV, PCA–LSTM achieved the lowest average MSE and MAE and the highest average R^2 among the four evaluated models at the t+30 horizon. However, its average R^2 remained negative (−0.2071), indicating that the model did not outperform a mean-based benchmark in terms of explained variance on the test data. The other three models also produced negative average R^2 values, with LSTM, LASSO–PCA–LSTM, and LASSO–LSTM obtaining −0.3565, −1.2868, and −1.6741, respectively. Therefore, PCA–LSTM should be interpreted as the relatively best-performing model among the evaluated alternatives for the t+30 horizon, rather than as a model with strong absolute predictive capability. The results indicate that forecasting performance deteriorates substantially as the prediction horizon increases.

The comparison of actual and predicted bond yields at the monthly prediction horizon is illustrated in Fig. 14.

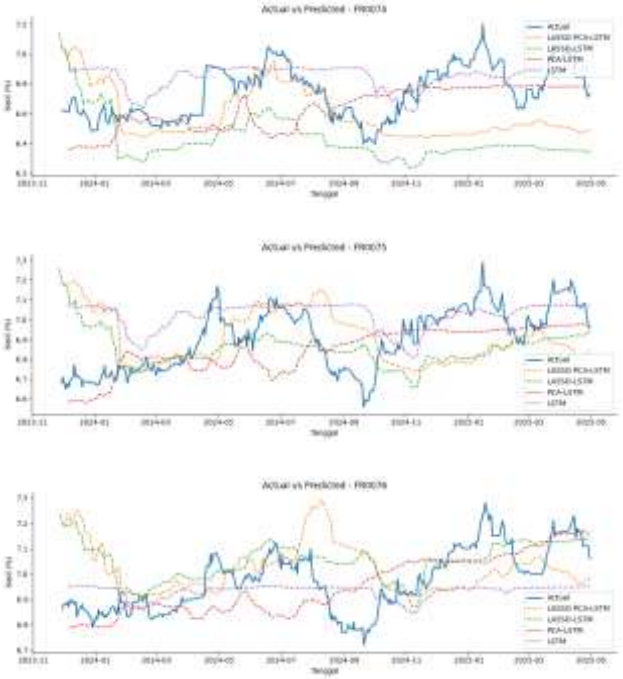


Fig. 14. Comparison of actual and predicted bond yields for all forecasting models at the monthly prediction horizon

Fig. 14 shows that all forecasting models experienced greater deviations from the actual bond yields as the prediction horizon increased. Nevertheless, the PCA–LSTM model generated prediction curves that more closely followed the observed yield movements than the remaining forecasting scenarios. In contrast, both LASSO–LSTM and LASSO–PCA–LSTM exhibited

substantially larger deviations, indicating that the elimination of input variables through LASSO reduced the availability of long-term macroeconomic information required by the forecasting model.

The comparison across the four forecasting frameworks shows that the contribution of feature engineering varies with the forecasting horizon. LASSO–PCA–LSTM performed best at the t+1 and t+5 horizons, while PCA–LSTM gave the strongest relative results at t+30. This difference suggests that there is no single feature engineering approach that is equally suitable for all forecasting horizons. The way input variables are selected and transformed may need to be adjusted according to the length of the prediction period. For government bond yield forecasting, the results therefore emphasize the need to consider both the structure of the input data and the temporal characteristics of the forecasting task when developing LSTM-based models.

G. Discussion

The experimental results demonstrate that the effectiveness of feature engineering depends on the forecasting horizon. The proposed LASSO–PCA–LSTM framework consistently achieved the best performance for the daily and weekly prediction horizons, whereas PCA–LSTM provided relatively better results for the monthly horizon. These findings indicate that combining feature selection and dimensionality reduction enables the LSTM model to learn more representative daily and weekly-term temporal patterns while reducing redundant information contained in macroeconomic variables.

The superior performance of the proposed LASSO–PCA–LSTM model for short-term forecasting is consistent with previous studies reporting that feature selection and dimensionality reduction improve deep learning performance by eliminating irrelevant variables and mitigating multicollinearity before model training [17],[20],[21]. Similarly, Ping et al. [17] demonstrated that integrating LASSO, PCA, and Bayesian Optimization within an LSTM framework produced higher forecasting accuracy than several conventional deep learning models. These findings support the results obtained in this study, where the integrated framework consistently outperformed the baseline LSTM model for the t+1 and t+5 prediction horizons.

For the t+30 horizon, PCA–LSTM produced the best relative performance among the four evaluated models based on the lowest average MSE and MAE and the highest average R^2. Nevertheless, its negative average R^2 indicates that the model still had limited absolute predictive capability. This finding suggests that preserving information from the complete macroeconomic feature set through PCA may be more suitable for longer-horizon forecasting than aggressive feature selection, although neither approach achieved strong explanatory performance at the t+30 horizon. This observation is consistent with the findings of Kostyra [20] who reported that PCA effectively preserves essential information for yield curve forecasting while reducing feature dimensionality.

The experimental findings indicate that LSTM-based models can provide meaningful predictive performance for shorter forecasting horizons, particularly when supported by appropriate feature engineering [16]–[17], [21]. However, the results also show that forecasting capability deteriorates as the horizon increases. At the t+30 horizon, all four models produced negative average R^2 values, indicating limited explanatory capability relative to a mean-based benchmark. Therefore, the effectiveness of LSTM-based forecasting in this study is horizon-dependent and should not be generalized across all forecasting periods [22].

Overall, this study demonstrates that selecting an appropriate feature engineering strategy is as important as choosing the forecasting algorithm itself. Rather than applying a single preprocessing technique across all forecasting tasks, the experimental results suggest that the choice of feature engineering should be adjusted according to the intended prediction horizon. This finding provides practical guidance for future research on Indonesian government bond yield forecasting and contributes empirical evidence regarding the application of integrated feature engineering techniques in emerging financial markets. The stronger performance of LASSO–PCA–LSTM at the shorter forecasting horizons is in line with earlier studies showing that feature selection and dimensionality reduction can help deep learning models by removing less informative variables and reducing multicollinearity before model training [17], [19], [23]. Ping et al. [17] for example, combined LASSO, PCA, and Bayesian Optimization with LSTM and reported better forecasting accuracy than several other deep learning models. A similar pattern can be observed in the present study, where LASSO–PCA–LSTM achieved better results than the baseline LSTM at both the t+1 and t+5 horizons. This suggests that reducing the input space can be particularly useful when the model is forecasting relatively close to the current observation.

A different pattern was observed at the t+30 horizon. PCA–LSTM gave the best relative results among the four models, based on its lower average MSE and MAE and higher average R^2. However, the average R^2 remained negative, meaning that the model did not provide strong explanatory performance compared with a mean-based benchmark. One possible explanation is that PCA allowed information from the broader set of macroeconomic variables to remain in the model, whereas LASSO removed several predictors during the feature selection stage. For a longer forecasting horizon, retaining a wider range of information may be more useful because the factors affecting bond yields can change over time. This result is broadly consistent with Kostyra [20], who found that PCA could reduce the dimensionality of yield curve data while retaining information relevant to forecasting. Nevertheless, this interpretation should be treated cautiously because comparative forecasting performance can be sensitive to the forecast horizon, evaluation design, and benchmark used [24].

The results also show a clear decline in predictive performance as the forecasting horizon becomes longer. The LSTM-based models performed more effectively at t+1 and t+5, particularly when feature engineering was incorporated into the modeling process. At t+30, however, all four models produced negative average R^2 values, indicating that their predictions were less effective than a mean-based benchmark in explaining the observed variation. This pattern suggests that the relationship between the available macroeconomic predictors and government bond yields becomes more difficult to capture

over a longer forecasting period. Similar considerations have been reported in previous research on financial time-series forecasting, where predictive performance tends to depend on the length of the forecasting horizon [16], [17], [22]. Hewamalage et al. [24] further emphasized that forecasting performance should be evaluated carefully across different forecast steps and against appropriate benchmarks because model performance may vary substantially with the evaluation horizon and design.

Taken together, the results suggest that feature engineering should not necessarily be treated as a fixed preprocessing strategy for every forecasting task. The relative performance of LASSO–PCA–LSTM at t+1 and t+5 and PCA–LSTM at t+30 indicates that the usefulness of a particular feature-processing strategy can vary with the forecasting horizon. This interpretation is also supported by evidence from financial forecasting showing that LASSO-based variable selection can improve LSTM performance by retaining informative predictors while reducing the influence of redundant variables [25]. At the same time, preprocessing procedures such as feature selection and dimensionality reduction need to be incorporated appropriately within the model-validation pipeline, because data-dependent preprocessing performed before validation can bias estimates of predictive performance [26]. In the context of Indonesian government bond yields, these findings suggest that model development should consider not only the forecasting algorithm but also how input variables are selected, transformed, and validated for the intended prediction horizon. The results therefore provide empirical evidence that can inform further research on feature engineering and deep learning for bond yield forecasting in emerging financial markets.

V. Conclusion

This study developed an integrated forecasting framework that combines LSTM, LASSO, PCA, and Bayesian Optimization to predict Indonesian government bond yields. Four forecasting scenarios were examined, namely LSTM, LASSO-LSTM, PCA-LSTM, and LASSO-PCA-LSTM, using three forecasting horizons: t+1, t+5, and t+30.

The results show that the benefit of feature engineering varies across forecasting horizons. LASSO-PCA-LSTM produced the highest overall performance at the t+1 and t+5 horizons, with average R^2 values of 0.8685 and 0.1760, respectively. At t+30, PCA-LSTM performed relatively better than the other frameworks. This result suggests that combining feature selection and dimensionality reduction can be useful for short- and medium-horizon forecasting, although the contribution of each technique is not necessarily consistent across longer horizons.

The lower predictive performance at t+30 also suggests that longer-term government bond yield movements are more difficult to capture using the variables and modeling approaches considered in this study. Changes in macroeconomic conditions and financial markets over longer periods may introduce additional variation that is not fully reflected in the available predictors. Future research could therefore consider a broader range of macroeconomic and market-based variables, including market sentiment indicators, as well as alternative deep learning architectures such as Transformer-based models to examine whether they can provide better performance at longer forecasting horizons.

References and Footnotes

[1] D. Schoenmaker and W. Schramade, *Corporate Finance for Long-Term Value*. Berlin: Springer, 2023.

[2] R. Wulandari, U. Farrohah, and S. Sumriyah, "Penerbitan obligasi negara sebagai alternatif pendanaan defisit APBN Indonesia," *J. Ris. Rumpun Ilm. Sos. Polit. Hum.*, vol. 2, no. 2, pp. 18–29, 2023. [Online]. Available: https://prin.or.id/index.php/JURRISH/article/view/1249

[3] S. Sutarmin, C. F. Ambari, F. Anisah, and M. Ma'rifah, "Penilaian obligasi pemerintah untuk tujuan investasi," *Indones Rich J.*, vol. 3, no. 1, pp. 54–60, 2022. [Online]. Available: https://irich.pknstan.ac.id/irj/article/view/48

[4] V. P. Nur and I. M. S. Sayekti, "Investor ritel mulai borong SUN FR, ini keunggulannya dibanding ORI," Kontan.co.id, 2026. [Online]. Available: https://investasi.kontan.co.id/news/investor-ritel-mulai-borong-sun-fr-ini-keunggulannya-dibanding-ori

[5] Bibit, "Apa saja keuntungan obligasi FR? Kenali manfaat investasi yang menguntungkan," 2024. [Online]. Available: https://artikel.bibit.id/investasi1/apa-saja-keuntungan-obligasi-fr-kenali-manfaat-investasi-yang-menguntungkan

[6] Y. D. Hapsari and A. Hamid, "Safe haven assets in emerging market (evidence from gold, government bonds, and crude oil in Indonesia 2019–2024)," *Int. J. Appl. Fin. Bus. Stud.,* vol. 13 no. 4, 668–677, 2026. [Online]. Available: https://www.ijafibs.pelnus.ac.id/index.php/ijafibs/article/view/464

[7] N. V. Popova, "On the yield to maturity of a coupon bond," *Fin. Theory Pract.,* vol. 27, no. 3, pp. 126–138, 2023, doi: 10.26794/2587-5671-2023-27-3-126-138.

[8] D. Xu, Y. Hu, S. Corbet, and C. Lang, "Return connectedness of green bonds and financial investment channels in China: Implications for hedging and regulation," *Res. Int. Bus. Fin.,* vol. 70, e102329, 2024, doi: 10.1016/j.ribaf.2024.102329.

[9] H. Wulandari, N. A. Achsani, and Moch. H. Santoso, "Determinants of Indonesian government bond yield," *J. Apl. Bis. Manaj.,* vol. 10, no. 3) e895, doi: 10.17358/jabm.10.3.895.

[10] T. Akram and S. A.-H. Uddin, "An empirical analysis of long-term Brazilian interest rates," *PLoS One,* vol. 16, no. 9, e0257313, 2021, doi: 10.1371/journal.pone.0257313.

[11] G. Amisano and O. Tristani, "Monetary policy and long-term interest rates," *Quant. Econ.*, vol. 14, no. 2, pp. 689–716, 2023, doi: 10.3982/QE1287.

[12] L. Shu and J.-K. Chou, "Using deep learning techniques to predict 10-year us treasury yield," in *2021 11th International Conference on Information Science and Technology (ICIST)*, pp. 545–552, 2021, doi: 10.1109/ICIST52614.2021.9440560.

[13] V. I. Kontopoulou, A. D. Panagopoulos, I. Kakkos, and G. K. Matsopoulos, "A review of ARIMA vs. machine learning approaches for time series forecasting in data driven networks," *Future Internet*, vol. 15, no. 8, e255, 2023, doi: 10.3390/fi15080255.

[14] S. Saleti, L. Y. Panchumarthi, Y. R. Kallam, L. Parchuri, and S. Jitte, "Enhancing forecasting accuracy with a moving average-integrated hybrid ARIMA-LSTM model," *SN Comput Sci.,* vol. 5, no. 6, e704, 2024, doi: 10.1007/s42979-024-03060-4.

[15] A. Yadav, C. K. Jha, and A. Sharan, "Optimizing LSTM for time series prediction in Indian stock market," *Procedia Comput Sci.*, vol. 167, pp. 2091–2100, 2020, doi: 10.1016/j.procs.2020.03.257.

[16] M. Waqas and U. W. Humphries, "A critical review of RNN and LSTM variants in hydrological time series predictions," *MethodsX*, vol. 13, e102946, 2024, doi: 10.1016/j.mex.2024.102946.

[17] M. Nunes, E. Gerding, F. McGroarty, M. Niranjan, and G. Sermpinis, "Deep learning for bond yield forecasting: The LSTM-LagLasso," *Int. J. Fin. Econ.,* vol. 31, no. 1, pp. 1269–1283, 2025, doi: 10.1002/ijfe.3116.

[18] W. Ping, Y. Hu, and L. Luo, "Price forecast of treasury bond market yield: Optimize method based on deep learning model," *IEEE Access*, vol. 12, pp. 194521–194539, 2024, doi: 10.1109/ACCESS.2024.3519438.

[19] Z. Li, B. Wang, and Y. Chen, "Incorporating economic indicators and market sentiment effect into US Treasury bond yield prediction with machine learning," *J. Infrastructure Policy Dev*., vol. 8, no. 9, e7671, 2024, doi: 10.24294/jipd.v8i9.7671.

[20] P. W. Santosa, "Macroeconomic indicators and yield curve of Indonesian government bond," *Bus. Manag. Econ. Eng.,* vol. 19, no. 1, pp. 34–48, 2021, doi: 10.3846/bmee.2021.13167.

[21] E. Koroleva and M. Kopeykin, "Understanding of macro factors

that affect yield of government bonds," *Risks,* vol. 10, no. 8, e166, 2022, doi: 10.3390/risks10080166.
[22] T. P. Kostyra, "Forecasting the yield curve for Poland with the PCA and machine learning," *Bank i Kredyt.*, vol. 55, no. 4, pp. 459–478, 2023, doi: 10.5604/01.3001.0054.8580.
[23] M. M. Idam, S. Anam, and H. A. Bukhori, "Bayesian-Optimized LSTM framework for accurate stock price prediction," *J. Appl. Inf. Comput.*, vol. 10, no. 2, pp. 1282–1290, 2026, doi: 10.30871/jaic.v10i2.11756.
[24] H. Hewamalage, K. Ackermann, and C. Bergmeir, "Forecast evaluation for data scientists: Common pitfalls and best practices," *Data Min. Knowl. Discov.*, vol. 37, no. 2, pp. 788–832, 2023, doi: 10.1007/s10618-022-00894-5.
[25] A. Moscovich and S. Rosset, "On the cross-validation bias due to unsupervised preprocessing," *J. R. Stat. Soc. Ser. B Stat. Methodol.*, vol. 84, no. 4, pp. 1474–502, 2022, doi: 10.1111/rssb.12537.
[26] J. Yang, Y. Wang, and X. Li, "Prediction of stock price direction using the LASSO-LSTM model combines technical indicators and financial sentiment analysis," *PeerJ Comput. Sci.*, vol. 8, e1148, 2022, doi: 10.7717/peerj-cs.1148.



Amanda Kayla Putri Wibowo received her bachelor's degree in Data Science from Telkom University, Bandung, Indonesia. Her research interests include machine learning, deep learning, financial time-series forecasting, data analytics, and artificial intelligence. She is particularly interested in applying data-driven computational methods to analyze complex datasets and develop predictive models for financial and other real-world applications within interdisciplinary research and innovation. (email: amandakayla696@gmail.com)



Deni Saepudin is a lecturer in the Faculty of Informatics at Telkom University, Bandung, Indonesia. His research interests include artificial intelligence, machine learning, data mining, financial technology, and intelligent information systems. His academic work focuses on developing and applying computational methods to support data-driven decision-making, predictive modeling, intelligent systems, and technology-based solutions across various application domains in contemporary digital environments. (email: denisaepudin@telkomuniversity.ac.id)