MIKAILALSYS
Journal of Multidisciplinary Sciences

e-ISSN : 2987-2286
p-ISSN : 2987-3924
Index: Harvard, Boston, Sydney, Copernicus, Sinta 4, Scilit, Semantic, Google, etc

https://doi.org/10.58578/mikailalsys.v4i3.12032

Application of Temporal Fusion Transformer for Multivariate Forecasting of Indonesian Composite Stock Price Index

Dzikra Rahmi Deni & Nonong Amalita
Universitas Negeri Padang, Indonesia
rahmidenidzikra@gmail.com; nongmat@fmipa.unp.ac.id

Submitted:	Revised:	Accepted:	Published:
Jul 27, 2026	Aug 24, 2026	Sep 5, 2026	Sep 10, 2026

Abstract

Accurate stock market index forecasting remains challenging because of complex temporal dependencies, high volatility, and uncertainty arising from dynamic financial conditions. This study investigates the application of the Temporal Fusion Transformer (TFT) to multivariate, multi-horizon forecasting of the Indonesian Composite Stock Price Index (IHSG). The proposed framework integrates historical IHSG movements, global financial indicators, technical indicators, and calendar features to generate point and probabilistic forecasts. Daily financial data from January 2000 to December 2025 were analyzed using a 252-trading-day historical input window to forecast the subsequent 30 trading days. Point-forecasting performance was evaluated using Mean Absolute Percentage Error (MAPE), whereas probabilistic performance was assessed using Quantile Risk (q-Risk). The TFT produced a MAPE of 13.723% over the 30-trading-day forecasting horizon. Its probabilistic forecasts maintained consistent quantile ordering, although further calibration is required to improve the reliability of uncertainty estimates. Variable selection analysis identified historical IHSG movements as the most influential predictive information, followed by the Relative Strength Index (RSI), lagged crude oil prices, volatility, and global market indicators, including movements in the S&P 500. The study contributes to financial forecasting research by demonstrating an interpretable deep learning framework that integrates point prediction, uncertainty representation, and variable importance analysis within a unified forecasting approach. The findings indicate that TFT can provide not only multi-horizon index forecasts but also transparent insights into the information structure underlying IHSG movements, thereby supporting more interpretable financial forecasting applications.

Keywords: IHSG; Model Interpretability; Multivariate Forecasting; Probabilistic Forecasting; Temporal Fusion Transformer

INTRODUCTION

The capital market plays a vital role in the financial system because it serves as a source of funding for companies and the government and as an investment vehicle for the public. Otoritas Jasa Keuangan (2019) explains that the capital market serves to bring together parties with excess funds and those in need of financing. In the Indonesian capital market, the Indonesia Stock Exchange provides various stock indices to reflect market conditions, one of which is the Composite Stock Price Index (IHSG), which represents the movement of all stocks listed on the Indonesia Stock Exchange.

The IHSG is a key indicator of the overall condition of the Indonesian stock market. Marselinus et al. (2025) explain that the IHSG reflects the performance of all stocks listed on the Indonesia Stock Exchange and can therefore be used to identify trends of market strength or weakness. Fitria and Komara (2026) explain that stock market performance can also reflect investment activity, financial stability, and market participants' confidence in macroeconomic conditions. However, the JCI's movements are not always stable because they are influenced by various economic and political factors, government policies, and global market developments.

Fluctuations in the IHSG indicate that forecasting index values is a complex issue in financial data analysis. The 2008 global financial crisis put pressure on the IHSG, with Fatmawati et al. (2024) demonstrating a significant negative impact in the short term. Tjen (2025) also explains that the COVID-19 pandemic caused changes in the volatility characteristics of the IHSG in the periods before, during, and after the pandemic. These conditions indicate that changes in the stock market can occur rapidly when faced with economic pressures or global uncertainty.

Forecasting the IHSG requires a broader range of information than just historical index data. Galena et al. (2025) explain that the IHSG's movement patterns cannot be fully explained based solely on historical information because global economic conditions and market sentiment also influence forecast outcomes. Hidayat and Sudjono (2022) demonstrate that variables such as global gold prices, global oil prices, and the USD/IDR exchange rate are correlated with the IHSG. Najibullah (2023) also found that the S&P 500 is correlated with IHSG movements, while Andriyani and Pimada (2025) showed that the VIX is related to changes in the Jakarta Composite Index as an indicator of global market uncertainty.

Various methods have been used to forecast the IHSG, including both statistical and machine learning approaches. Fauziyah et al. (2021) applied the ARIMAX-TARCH model, taking into account exogenous variables and changes in volatility, while Patriya et al. (2023) used Long Short-Term Memory (LSTM) to forecast the IHSG closing value. Tjen (2025) compared ARIMA-GARCH and LSTM across several market conditions. The results of this study indicate that different methods have varying capabilities in capturing complex patterns in IHSG data.

The Temporal Fusion Transformer (TFT) is one of the deep learning approaches suitable for multivariate time series forecasting problems. Lim et al. (2021) introduced the TFT as an attention-based architecture capable of learning both short-term and long-term temporal relationships and generating multi-horizon forecasts. The TFT combines a recurrent network, an attention mechanism, a Variable Selection Network, and a gating mechanism, enabling it to select relevant information during the forecasting process. In addition to generating forecasts, the TFT also provides information regarding the importance of variables and the temporal patterns used by the model during the learning process (Lim et al., 2021).

Although the application of TFT to stock market data is beginning to gain traction, its use specifically for forecasting the closing value of the IHSG remains limited. Hartanto et al. (2024) applied TFT to forecast stock prices in several industrial sectors, while Nurhasan and Ghufron (2026) compared TFT with LSTM on several individual stocks in Indonesia. Previous research on the IHSG has predominantly used ARIMA, ARIMAX, GARCH, and LSTM models and has focused more on prediction accuracy, while the use of models capable of generating multi-horizon forecasts while also providing interpretations of variables and temporal patterns remains limited.

Based on this gap, this study applies the Temporal Fusion Transformer to forecast the IHSG closing value using historical IHSG data, external variables such as the S&P 500, global gold prices, global oil prices, the USD/IDR exchange rate, and the VIX, as well as technical indicators derived from IHSG movements. The TFT was selected because it is capable of processing multivariate time series, generating quantile-based multi-horizon forecasts, and providing interpretable results through variable selection mechanisms. This study aims to determine the optimal TFT configuration, evaluate point forecasting performance and probabilistic forecasting quality over a 30-trading-day horizon, and analyze the contribution of input variables used during the forecasting process.

METHODS

This study employs a quantitative approach using a predictive research design based on secondary time-series data. The daily financial market data covers the period from January 2000 to December 2025. The IHSG closing value is used as the target, while predictor information is derived from global market indicators and technical indicators derived from the IHSG's historical data. The model is designed to generate multi-horizon probabilistic forecasts over 30 trading days and is evaluated out-of-sample using chronological data splitting.

The model's input structure consists of one target variable, five external variables, four technical indicators, and calendar features. The IHSG is the primary target. The external variables consist of the S&P 500, Gold Futures, WTI Crude Oil Futures, USD/IDR, and the CBOE Volatility Index (VIX). All five external variables are used in their lag-1 form. The technical indicators derived from the IHSG include the 20-day moving average (MA20), the 50-day moving average (MA50), the 14-day Relative Strength Index (RSI), and 20-day historical volatility. The calendar features include the time index, day, month, and quarter.

Table 1. Input Variables and Feature Characteristics for IHSG Forecasting

Variable	Instrument Name	Data Type	Unit	Role
IHSG	IDX Composite (JKSE)	Daily closing value	Index points	Target
S&P 500	S&P 500 (SPX)	Daily closing value	Index points	External
Gold Price	Gold Futures	Daily closing value	USD per troy ounce	External
Oil Price	Crude Oil WTI Futures	Daily closing value	USD per barel	External
USD/IDR Exchange Rate	USD/IDR	Daily closing value	Rupiah per USD	External
VIX	CBOE Volatility Index (VIX)	Daily closing value	Volatility index (%)	External

Variable	Instrument Name	Data Type	Unit	Role
MA20	20-Day Moving Averag	Technical indicator	Index points	Historical
MA50	50-Day Moving Average	Technical indicator	Index points	Historical
RSI	Relative Strength Index (14-day)	Technical indicator	Index (0–100)	Historical
Volatility	20-Day Historical Volatility	Standard deviation of daily returns	Ratio	Historical

Research data was collected through documentation and the extraction of historical financial market data. The processing stages include data quality checks, synchronization based on the IHSG trading calendar, creation of lag-1 external variables, creation of technical indicators and calendar features, chronological data partitioning, standardization based on training data parameters, and the creation of a sliding window. All stages are designed to ensure that future information does not enter the training or model selection processes.

1. Data Preprocessing

Data preprocessing is performed to ensure that the time-series data used in TFT modeling has the structure, chronological order, and data quality required for the analysis. This stage includes data verification and sorting, data quality checks, synchronization based on the IHSG trading calendar, and the preparation of data ready for use in modeling. The stages of time-series data preprocessing in this study are described as follows.

a. Data Verification and Sorting

Data checks are performed before the data is prepared as model input. The items reviewed include date formats, data types, units, duplicate observations, missing values, and period completeness. Extremely high or low values also need to be examined in their chronological order to ensure that genuine market changes are not mistaken for recording errors (Tsai & Hsiao, 2010).

After verification, the data is sorted from the earliest to the latest date. This order must be maintained because relationships in a time series depend on the position of each observation. Sorting errors can create a misleading historical window and allow future information to be included in earlier periods (Tsai & Hsiao, 2010).

b. Synchronizing the Trading Calendar

The observation dates for the IHSG, S&P 500, gold, oil, USD/IDR, and VIX are not always the same. Differences in holidays, time zones, and trading calendars make it impossible to merge the data based solely on row order. Alignment must be performed based on dates so that each value represents a corresponding point in time (Tsai & Hsiao, 2010).

The JCI trading calendar is used as a reference because the JCI is the target variable. After the lag is applied, external data is matched to the JCI trading dates. Observations that do not yet have complete pairs are checked before window formation, so that the model does not accept empty values or data from mismatched dates (Tsai & Hsiao, 2010).

c. Application of External Variable Lag

A lag shifts the value of a variable to the next time period, so that past information is used to forecast the current value. All external variables are given a lag of one trading day so that the input at time t is derived from previously available information. For the jth external variable, the one-period lag is formulated as follows:

$$x_{j,t}^{\text{lag}} = X_{j,t-1}, \tag{1}$$

The external variables are denoted as $X_{i,t}^{lag}$, representing the value of the i-th external variable used as an input feature at time t. To capture temporal dependencies and the delayed effects of external market conditions, the lagged values of external variables from the previous trading day are represented as $X_{j,t-1}$. Here, $t$indicates the observation period or trading day, while $j$represents the index of external variables included in the forecasting framework. The external variables considered in this study include global market indicators and macro-financial factors, namely the S&P 500 index, global gold prices, crude oil prices, USD/IDR exchange rates, and the CBOE Volatility Index (VIX).

Equation (1) shows that the values of the external variables used at time t are derived from the values available on the previous trading day. The use of lags is intended to prevent the use of information that was not yet available when the IHSG closing value at time t was forecasted.

d. Formation of Technical Features, Temporal Features, and TFT Input Clustering

After lag-1 is established, four technical indicators are calculated from the IHSG history: the 20-day moving average (MA20) to represent the short-term trend, the 50-day moving average (MA50) for the medium-term trend, the 14-day RSI for momentum, and 20-day volatility calculated from the standard deviation of daily returns. The establishment of these indicators results in a warm-up period, during which observations that do not yet have complete indicator values are excluded. Next, `time_idx` is created as a sequential index, and the calendar features used include day, month, and quarter.

TFT inputs are grouped according to the model structure. IHSG, SP500_lag1, Gold_lag1, Oil_lag1, USDIDR_lag1, VIX_lag1, MA20, MA50, RSI, and Volatility are treated as observed historical inputs. time_idx is used as a known future real input, while day, month, and quarter are used as known future categorical inputs. series_id is used as the IHSG series identifier. This grouping ensures that the decoder only uses information that is known or required as the temporal structure for the future horizon.

e. Chronological Organization of Data

The data was divided chronologically into a 70% training set, a 15% validation set, and a 15% testing set. The 70% of data from the initial period was used to train the model and determine the TFT parameters or weights. Next, 15% of the data is used as the validation set to select the best hyperparameter configuration and monitor the early stopping process. The remaining 15% of the data from the final period is used as the testing set to evaluate the model's out-of-sample performance on data not used in the training process or in model configuration selection (Tashman, 2000).

Data partitioning is performed without shuffling to preserve the temporal order. Shuffling time-series data can cause observations from future periods to be included in the training data, resulting in data leakage. Therefore, the temporal boundaries between the training, validation, and test data are maintained throughout all modeling stages, including when calculating standardization parameters and forming a sliding window. With this procedure, the training, configuration selection, and model evaluation processes continue to use information in the correct chronological order (Tashman, 2000).

f. Data Standardization

Numeric variables have different scales and are therefore standardized using z-scores before being used in the model. Standardization is applied to the IHSG, five external lag-1 variables, MA20, MA50, RSI, and Volatility. The mean and standard deviation parameters are calculated only from the training data; these same parameters are then used to transform the validation and test data. This approach ensures that statistical information from the validation and test periods does not influence the training process. The quantile forecast results are converted back to the original IHSG scale before calculating MAE, RMSE, MAPE, q-Risk, coverage intervals, and presenting the results.

$$z_{t,j} = \frac{x_{t,j} - \mu_j^{\text{train}}}{\sigma_j^{\text{train}}} \quad (2)$$

In the standardization process, $z_{j,t}$represents the standardized value of the j-th variable at time t, while $x_{j,t}$denotes the original value of the j-th variable at time t. The term μ_j^{train}represents the mean value of the j-th variable calculated from the training data, whereas σ_j^{train}represents the standard deviation of the j-th variable estimated from the training data. The mean and standard deviation obtained from the training dataset are used as reference parameters for the standardization process to prevent information leakage from the test dataset.

The mean and standard deviation are calculated using only the training data. The same parameters are then used to standardize the validation and test data to prevent information leakage from future periods. Forecast results on the standardized scale are converted back to the original scale before being presented and evaluated. The inverse transformation is defined as follows:

$$\hat{y}_t = \hat{z}_t{}^{(q)} \sigma_y{}^{train} + \mu_y{}^{train} \quad (3)$$

In the forecasting framework, $\hat{y}_t^{(q)}$represents the IHSG forecast value at quantile $q$in the original scale, while $\hat{z}_t^{(q)}$denotes the IHSG forecast value at quantile $q$in the standardized scale. The term μ_y^{train}represents the mean value of IHSG calculated from the training dataset, whereas σ_y^{train}represents the standard deviation of IHSG estimated from the training dataset. These parameters are used to transform the standardized forecast outputs back into their original scale. The quantile levels considered in this study are defined as $q \in \{0.10, 0.50, 0.90\}$, representing the lower, median, and upper predictive intervals generated by the model. Equation (3) is used to convert the forecast results from the standardized scale back to the original IHSG scale. The denormalization process is performed before the forecast results are presented and evaluated using MAPE and q-Risk.

g. Sliding Window Method

A sliding window transforms a single long time series into a set of historical input-target pairs. Each window contains k observations for the encoder and the next h observations for the target. The window is then shifted by one step to form a new sample without altering the temporal order (Lim et al., 2021).

$\mathbf{X}_t = \{\mathbf{x}_{t-k+1}, \mathbf{x}_{t-k+2}, \dots, \mathbf{x}_t\}$	(4)

Here, $X_t$represents the historical input window at the initial forecasting point t, $x_t$denotes the input vector at time t, and $k$represents the encoder length. The input vector $x_t$may consist of multiple features, expressed as:

$$\mathbf{x}_t = [y_t, , s_t, , e_t, , m_t \, u_t \, d_t \, b_t]$$

where $y_t$represents the historical IHSG values, while the remaining components represent external variables and time-related features included in the forecasting framework. The target values for future periods are defined as:

$\mathbf{Y}_t = \{y_{t+1}, y_{t+2}, \dots, y_{t+h}\}$	(5)

In this formulation, $Y_t$denotes the forecast target sequence after the initial time point t, $y_{t+\tau}$represents the actual IHSG value at forecasting horizon τ, $h$indicates the forecasting horizon length, and $\tau = 1,2, \dots, h$represents each future prediction step.

Given $n$effective observations, the number of input–target window pairs that can be generated is determined based on the available historical sequence length, encoder length, and forecasting horizon.

$N_{\text{window}} = n - k - h + 1$	(6)

In Equations (4)–(6), k is the encoder length, h is the horizon length, and n is the number of effective observations. This study sets k = 252 and h = 30. After constructing the lags and technical indicators, the effective data begins on March 20, 2000. The validation and testing windows are defined so that the target is entirely within each data segment, while the encoder context may use previous historical observations without extending beyond the target's starting point.

2. Temporal Fusion Transformer (TFT)

The Temporal Fusion Transformer (TFT) is an attention-based deep learning architecture designed for multivariate and multi-horizon time-series forecasting(Lim et al., 2021). TFT is designed to process three main types of inputs: static covariates, time-varying observed variables, and time-varying known future variables. Observed variables are only available during the historical period, whereas known variables are available throughout the forecasting horizon.

TFT integrates several key components, including Gated Residual Networks (GRN), Variable Selection Networks (VSN), static covariate encoders, LSTM-based temporal processing, interpretable multi-head attention, and quantile output layers. The combination of these components enables TFT to identify relevant input variables, capture both short-term and long-term temporal dependencies, and generate probabilistic multi-horizon forecasts.

The general formulation of TFT quantile forecasting is expressed as:

$$\hat{y}_{i,t+\tau}^{(q)} = f_q(\tau, y_{i,t-k:t}, \boldsymbol{z}_{i,t-k:t}, \boldsymbol{x}_{i,t-k:t+\tau}, \boldsymbol{s}_i \quad (7)$$

where $\hat{y}_{i,t+\tau}^{(q)}$ represents the predicted value at quantile level $q$and forecasting horizon τ. The term $y_{i,t-k:t}$represents the historical target values, $z_{i,t-k:t}$ represents time-varying observed variables available only during the historical period, $x_{i,t-k:t+\tau}$ represents time-varying known future variables, and $s_i$represents static covariates.

TFT directly generates forecasts for multiple future horizons, formulated as:

$$\tau \in \{1, 2, ., .. \tau_{max}\} \quad (8)$$

Furthermore, TFT applies quantile regression to produce probabilistic prediction intervals using different quantile levels:

$$q \in \{0.10, 0.50, 0.90\} \quad (9)$$

This approach enables TFT to simultaneously forecast multiple future horizons while providing prediction intervals that represent forecasting uncertainty. TFT also provides interpretability through variable selection weights and attention weights. Variable selection weights indicate the relative contribution of each input variable to the forecasting process, whereas attention weights highlight historical time steps that receive greater importance from the model. However, these weights represent predictive associations and should not be interpreted as causal relationships

The overall TFT architecture used in this study is presented in Figure 1. Temporal Fusion Transformer Architecture.

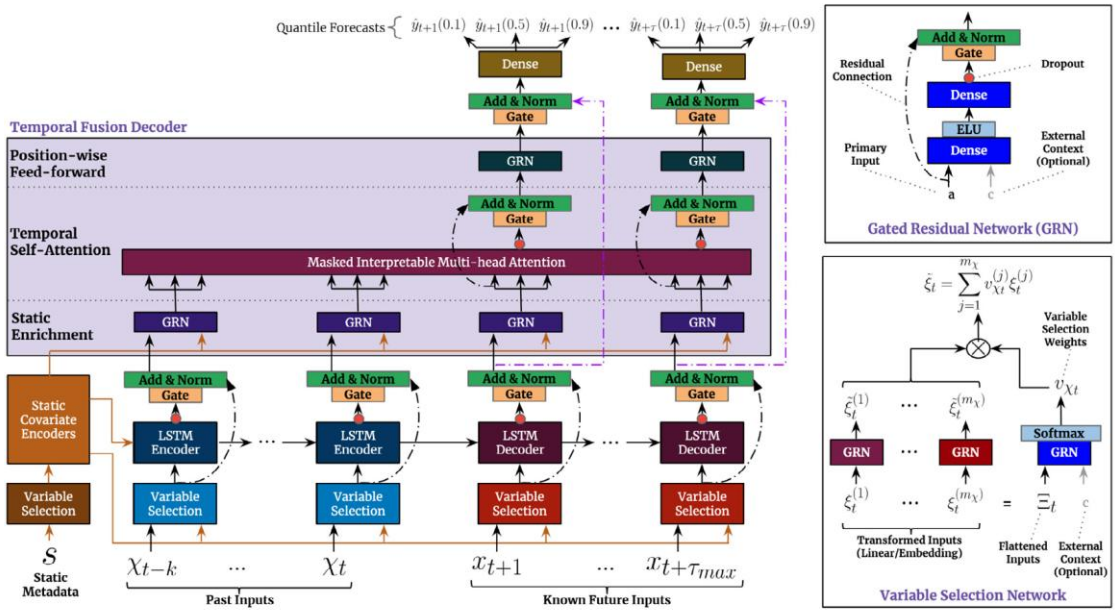

Figure 1. Temporal Fusion Transformer Architecture

Figure 1. Temporal Fusion Transformer Architecture illustrates the complete workflow of TFT, including variable selection, static covariate encoding, LSTM-based temporal processing, interpretable multi-head attention, and quantile forecasting output. The architecture integrates different types of input information, including static covariates, time-varying observed variables, and known future variables to generate multi-horizon forecasts. Through the combination of gating mechanisms, attention layers, and variable selection networks, TFT is able to capture short-term and long-term temporal dependencies while providing interpretable forecasting results.

a. Gated Residual Network (GRN)

The Gated Residual Network (GRN) is a fundamental component of TFT that controls information flow and enables adaptive model complexity. GRN selectively passes useful information while reducing unnecessary transformations during the learning process.

GRN receives a primary input $a$and an optional context vector c. The GRN formulation is defined as:

$$GRN_{\omega}(a,c) = LayerNorm\left(a + GLU_{\omega}(\eta_1)\right) \quad (10)$$

where:

$$\eta_1 = W_{1,\omega}\eta_2 + b_{1,\omega} \quad (11)$$

and:

$$\eta_2 = ELU\left(W_{2,\omega}a + W_{3,\omega}c + b_{2,\omega}\right) \quad (12)$$

where a represents the primary input, $c$represents the context vector, Wrepresents learnable weight matrices, and $b$represents bias parameters.

The Exponential Linear Unit (ELU) activation function is used to introduce nonlinear transformations into the network. In addition, TFT employs a Gated Linear Unit (GLU) mechanism to regulate information flow through the network.

The GLU formulation is expressed as:

$$GLU_\omega(\gamma) = \sigma\left(W_{4,\omega}\gamma + b_{4,\omega}\right) \odot \left(W_{5,\omega}\gamma + b_{5,\omega}\right) \quad (13)$$

where σdenotes the sigmoid activation function and $\odot$ represents element-wise multiplication.

The gating mechanism enables TFT to suppress irrelevant components while preserving useful information through residual connections. When a specific component does not contribute significantly to the forecasting process, the gate output can approach zero. Conversely, important information can be maintained and propagated through residual pathways.

b. Variable Selection Network (VSN)

The Variable Selection Network (VSN) is an essential component of TFT that identifies the most relevant input variables at each time step. This mechanism is important because the predictive contribution of variables may vary over time, particularly in complex financial time-series forecasting. TFT applies separate VSN modules for three categories of inputs:

1. Static covariates;
2. Time-varying observed variables;
3. Time-varying known future variables.

The variable selection weights are generated using GRN followed by softmax normalization:

$$v_{\chi_t} = Softmax\left(GRN_{v_\chi}(\Xi_t, c_s)\right) \quad (14)$$

where v_{χ_t} represents the variable selection weights, Ξ_t represents the combined feature representations at time t, and c_s represents the static context vector.

Each input variable is subsequently processed through an individual GRN:

$$\tilde{\xi}_t^{(j)} = GRN_{\xi^{(j)}}\left(\xi_t^{(j)}\right) \quad (15)$$

The processed variables are then combined based on their selection weights:

$$\tilde{\xi}_t = \sum_{j=1}^{m_\chi} v_{\chi_t}^{(j)}\, \tilde{\xi}_t^{(j)} \quad (16)$$

The value of $v_{\chi_t}^{(j)}$ represents the relative predictive contribution of the j-th variable at time t. Since the weights are generated through softmax normalization, they represent the relative importance of variables within the forecasting process. However, these weights indicate predictive relevance rather than causal influence.

c. Static Covariate Encoders

Static Covariate Encoders integrate time-invariant information into the forecasting process. Static variables are transformed into context vectors that influence variable selection, LSTM initialization, and static enrichment. TFT uses static context vectors to provide additional information throughout temporal processing, allowing static characteristics to affect the learned representations. The Static Covariate Encoders integrate time-invariant information into the forecasting process. Unlike time-varying variables, static covariates remain unchanged throughout the forecasting period. In TFT, static variables are transformed into context vectors that influence variable selection, temporal processing, and static enrichment stages. The static covariate encoder generates four context vectors that are used in different components of the TFT architecture:

$$(c_s, c_e, c_c, c_h) = GRN_s(s_i) \quad (17)$$

where s_i represents the static covariate vector, while c_s, c_e, c_c, and c_h denote different static context vectors.

The generated context vectors have different roles within the TFT architecture. The vector c_s is used as a context input for the temporal variable selection network. The vectors c_c and c_h are used to initialize the LSTM cell state and hidden state, respectively. Meanwhile, c_e is used during the static enrichment stage to incorporate static information into temporal representations.

For the IHSG forecasting task, the static covariate component has limited variation because the model focuses on a single time-series entity. Therefore, the main contribution of this component is providing structural information for the forecasting architecture rather than distinguishing between multiple entities.

d. Temporal Processing

The Temporal Processing component is responsible for learning temporal dependencies from historical observations and known future information. TFT combines a sequence-to-sequence architecture based on LSTM for local temporal pattern extraction with an interpretable multi-head attention mechanism for capturing long-range temporal dependencies (Lim et al., 2021).

1) LSTM Encoder–Decoder

The LSTM Encoder–Decoder component is used to capture local temporal patterns from sequential data. The encoder processes historical target values and observed variables within the input window, while the decoder processes known future variables throughout the forecasting horizon.

The historical input sequence received by the encoder is expressed as:

$$\{y_{t-k:t}, z_{t-k:t}\} \quad (18)$$

where $k$represents the length of the historical look-back window, $y_{t-k:t}$represents historical target values, and $z_{t-k:t}$represents time-varying observed variables. The known future inputs received by the decoder are formulated as:

$$\{x_{t+1:t+\tau}\} \quad (19)$$

where $x_{t+1:t+\tau}$represents time-varying known variables available during the forecasting horizon.

The LSTM encoder generates hidden representations by learning historical temporal dependencies:

$$h_t^{enc} = LSTM_{enc}(y_{t-k:t}, z_{t-k:t}) \quad (20)$$

Meanwhile, the decoder generates future temporal representations based on known future information:

$$h_t^{dec} = LSTM_{dec}(h_t^{enc}, x_{t+1:t+\tau}) \quad (21)$$

The encoder captures historical patterns, whereas the decoder produces temporal representations for future forecasting horizons. However, the LSTM output is not directly used as the final prediction. The representation is further processed through static enrichment, attention mechanisms, gating layers, and quantile output layers.

The LSTM representation is combined with the output from the variable selection network through gated residual connections:

$$\tilde{x}_t = LayerNorm\left(\xi_t + GLU(\xi_t)\right) \quad (22)$$

This mechanism enables TFT to preserve the original feature representation while incorporating learned temporal patterns.

2) Static Enrichment

After temporal processing using LSTM, the temporal representations are enriched using static context information. This process allows static characteristics to influence the learned temporal representation before applying the attention mechanism.

The static enrichment process is formulated as:

$$\tilde{h}_t = GRN_e(h_t, c_e) \quad (23)$$

where $h_t$represents the temporal representation generated by the LSTM layer and $c_e$represents the static enrichment context vector.

Static enrichment allows TFT to incorporate time-invariant information into dynamic temporal features. In a single IHSG forecasting series, static enrichment has limited variability because the series identity remains unchanged; however, it maintains the general TFT architecture and enables extension to multi-series forecasting problems.

3) Interpretable Multi-Head Attention

The Interpretable Multi-Head Attention mechanism enables TFT to capture long-term temporal dependencies by identifying important historical periods that influence future predictions. Unlike conventional attention mechanisms, TFT applies an interpretable attention design by sharing value projections across attention heads.

The scaled dot-product attention mechanism is defined as:

$$Attention(Q, K, V) = A(Q, K)V \quad (24)$$

where Q, K, and Vrepresent query, key, and value matrices, respectively.

The attention weight matrix is calculated as:

$$A(Q, K) = Softmax\left(\frac{QK^T}{\sqrt{d_{attn}}}\right) \quad (25)$$

where d_{attn} represents the attention dimension.

For multi-head attention, TFT applies multiple attention heads with different query and key projections while maintaining a shared value representation:

$$H_m = Attention\left(Q, W_Q^{(m)}, KW_K^{(m)} VW_V\right) \quad (26)$$

The final attention output is obtained by combining all attention heads:

$$\widetilde{H} = [H_1, H_2, ., ..H_m]W_H \quad (27)$$

The use of shared value projections enables attention weights from different heads to be aggregated and interpreted more easily. These attention weights indicate which historical time steps receive greater importance during forecasting.

Furthermore, TFT applies a causal mask within the attention mechanism to prevent future information from influencing earlier predictions. This process is essential to avoid information leakage in time-series forecasting.

4) Position-Wise Feed-Forward and Residual Connections

The After the attention layer, TFT applies a position-wise feed-forward network based on GRN to introduce additional nonlinear transformations at each time step. This layer allows the model to refine temporal representations before generating the final forecast output.

The feed-forward transformation is formulated as:

$$\tilde{y}_t = GRN_y(\widetilde{H}_t) \quad (28)$$

where $\widetilde{H}_t$ represents the output from the interpretable multi-head attention layer.

TFT also applies residual connections around the feed-forward layer:

$$\theta_t = LayerNorm\left(\widetilde{H}_t + GLU(\tilde{y}_t)\right) \quad (29)$$

The residual pathway allows the model to retain useful temporal representations when additional nonlinear transformations do not provide significant improvement. This mechanism improves training stability and enables adaptive information flow throughout the TFT architecture.

3. Model Evaluation

The model was evaluated using the hold-out test dataset based on point forecasting accuracy and probabilistic forecasting performance. The median quantile forecast ($Q_{0.50}$) was evaluated using Mean Absolute Percentage Error (MAPE), while the quality of probabilistic forecasts was assessed using Quantile Risk (q-Risk) at quantile levels $Q_{0.10}$, $Q_{0.50}$, and $Q_{0.90}$.

a. Mean Absolute Percentage Error (MAPE)

Mean Absolute Percentage Error (MAPE) measures forecasting error based on the absolute difference between predicted and actual values relative to the actual observations. The errors are averaged across all forecasting periods and expressed as percentages (Prayudani et al., 2019). MAPE provides a scale-independent measure, allowing comparison between different forecasting methods. A lower MAPE value indicates smaller relative forecasting errors and better model accuracy.

MAPE is formulated as:

$$MAPE = \frac{100}{n}\sum_{t=1}^{n}\left|\frac{y_t - \hat{y}_t}{y_t}\right| \quad (30)$$

where y_t represents the actual IHSG value at time t, $\hat{y}_t$ denotes the predicted IHSG value, and n represents the number of observations in the evaluation dataset. The interpretation of MAPE values follows the classification presented by Lubis et al. (2023).

Table 2. MAPE Interpretation Criteria

MAPE Value	Forecasting Performance Interpretation
< 10%	Very Good
10–20%	Good
20–50%	Acceptable
> 50%	Poor

b. Quantile Risk (q-Risk)

Quantile Risk (q-Risk) is a normalized form of quantile loss used to evaluate probabilistic forecasts. Unlike training loss, which is used to optimize model parameters, q-Risk is calculated on the out-of-sample test data to assess the quality of generated quantile predictions (Lim et al., 2021).

The q-Risk metric is defined as:

$$qRisk(q) = \frac{2\sum_{i=1}^{N} QL_q\left(y_i, \hat{y}_i^{(q)}\right)}{\sum_{i=1}^{N} |y_i|} \quad (31)$$

where N represents the total number of evaluated observations, QL_q denotes the quantile loss function at quantile q, y_i represents the actual value, and $\hat{y}_i^{(q)}$ denotes the predicted value at quantile q.

In this study, q-Risk is calculated for $q = \{0.10, 0.50, 0.90\}$. Lower q-Risk values indicate that the predicted quantiles are closer to the observed values after considering the target scale, reflecting better probabilistic forecasting performance.

RESULTS

1. Dataset Characteristics and Forecasting Framework

The forecasting framework was developed using the Indonesian Composite Stock Price Index (IHSG) closing value as the prediction target by integrating global financial indicators, technical indicators, and temporal features. The input variables were designed to represent both internal market dynamics and external financial conditions influencing IHSG movements. The input features consisted of historical IHSG information, technical indicators including moving averages, Relative Strength Index (RSI), and volatility, as well as global market indicators including the S&P 500 index, gold prices, crude oil prices, USD/IDR exchange rates, and the VIX index. The Temporal Fusion Transformer (TFT) model was configured using 252 historical trading days as encoder input and generated probabilistic forecasts for the subsequent 30 trading days.

Table 3. Summary of Input Variables Used in IHSG Forecasting

Variable Group	Variables	Role
Target variable	IHSG closing value	Forecast target
Historical variables	IHSG lag values, MA20, MA50, RSI, volatility	Observed historical inputs
External variables	S&P 500, Gold, Oil, USD/IDR, VIX	External predictors
Calendar features	Trading day information	Known future inputs

2. TFT Hyperparameter Optimization

The TFT model was optimized using random search to identify the optimal model configuration for multivariate and multi-horizon IHSG forecasting. The optimization process evaluated multiple combinations of hyperparameters based on validation quantile loss, which measures the quality of probabilistic forecast estimation. The configuration with the lowest validation quantile loss was selected as the final model configuration and subsequently used for hold-out testing and forecasting evaluation.

Table 4. Optimal TFT Hyperparameter Configuration

Parameter	Value
Encoder length	252 trading days
Forecast horizon	30 trading days
Hidden size	240
LSTM layers	2
Dropout	0.20
Learning rate	0.01
Batch size	128
Best Trial	126
Validation quantile loss	0.081024

The optimal TFT configuration was obtained from trial 126, which achieved the lowest validation quantile loss of 0.081024 among the evaluated configurations. The selected configuration utilized a 252-trading-day encoder length, representing approximately one year of historical market information, allowing the model to capture both short-term fluctuations and longer-term temporal dependencies in IHSG movements. The optimized model was then applied to the hold-out test dataset for forecasting evaluation.

3. Hold-Out Forecasting Performance

The TFT model was evaluated using the hold-out test dataset with a 30-trading-day forecasting horizon. The median quantile output (Q0.50) was used as the point forecast and evaluated using Mean Absolute Percentage Error (MAPE). This evaluation was conducted to measure the point forecasting performance of TFT in predicting IHSG movements across the forecasting horizon.

Table 5. TFT Point Forecasting Performance Based on MAPE

Model	Forecast Horizon	MAPE (%)
TFT	30 trading days	13.723

The TFT model achieved a MAPE value of 13.723% for the 30-trading-day forecasting horizon. Based on the adopted evaluation criteria, this result indicates good forecasting performance, demonstrating that TFT is capable of capturing temporal patterns within IHSG movements.

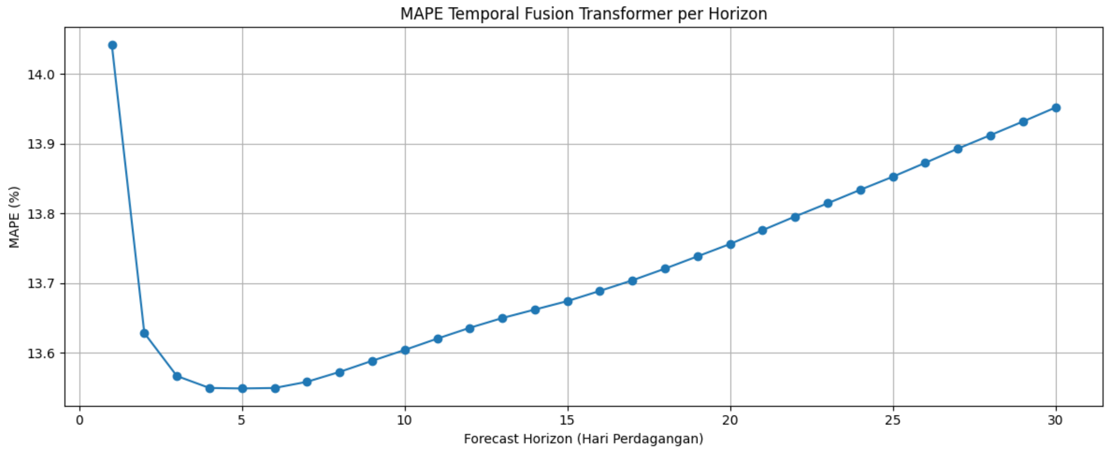

Figure 2. MAPE Performance Across 1–30 Trading-Day Forecast Horizons

Figure 2 illustrates the variation of MAPE values across forecasting horizons from 1 to 30 trading days. The forecasting error remained relatively stable throughout the prediction period, with MAPE values ranging approximately between 13.5% and 14.0%. This indicates that the model maintained consistent forecasting performance across different prediction horizons rather than experiencing substantial error degradation as the forecasting period increased.

4. Probabilistic Forecasting Performance

The TFT model generated probabilistic forecasts using three quantile outputs, namely Q0.10, Q0.50, and Q0.90, representing lower, median, and upper prediction estimates. The performance of probabilistic forecasts was evaluated using Quantile Risk (q-Risk), which measures the quality of quantile predictions by assessing the deviation between observed values and estimated uncertainty intervals.

Table 6. Quantile Risk Performance of TFT

Quantile	q-Risk
Q0.10	0.038587
Q0.50	0.140200
Q0.90	0.126794

The q-Risk evaluation demonstrates that TFT was able to generate probabilistic forecasts across different uncertainty levels. The Q0.10 quantile achieved the lowest q-Risk value (0.038587), indicating that the lower prediction boundary produced relatively smaller quantile errors. Meanwhile, the Q0.50 and Q0.90 quantiles represent the median prediction and upper uncertainty boundary, providing additional information regarding possible future IHSG variations.

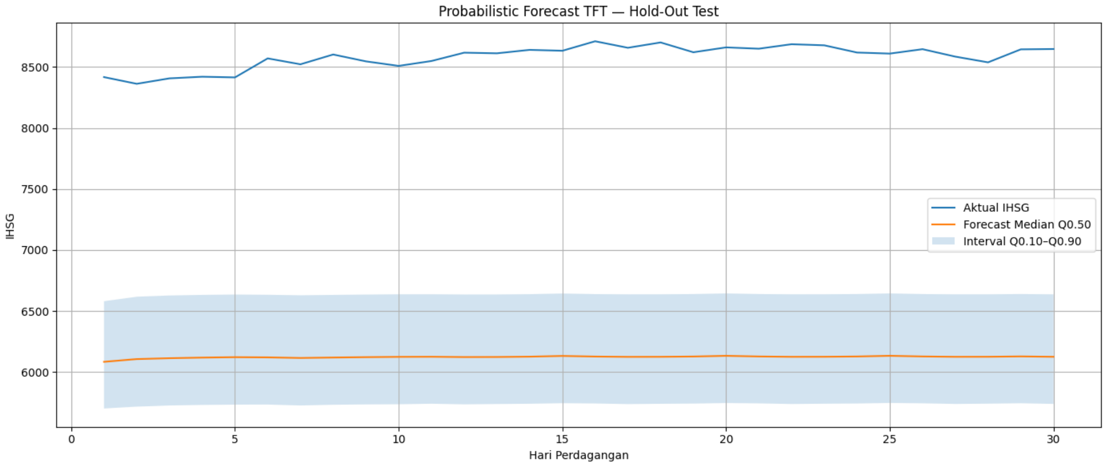

Figure 3. Probabilistic Forecasting Output with Quantile Prediction Intervals Generated by TFT

Figure 3 illustrates the probabilistic forecasting results generated by TFT on the hold-out test dataset. The model produced median forecasts (Q0.50) together with prediction intervals represented by Q0.10 and Q0.90 quantiles. This probabilistic representation allows the model to capture uncertainty associated with future IHSG movements rather than producing only a single deterministic prediction.

5. Model Interpretability Analysis Through Variable Importance

The variable selection network of TFT was used to identify the relative contribution of each input variable during the forecasting process. The results indicate that historical IHSG values received the highest importance weight, followed by technical and external market variables, including RSI, crude oil price lag, volatility, and S&P 500 lag. This finding demonstrates that the model primarily relies on historical market dynamics while incorporating additional information from technical indicators and global financial conditions to improve forecasting representation

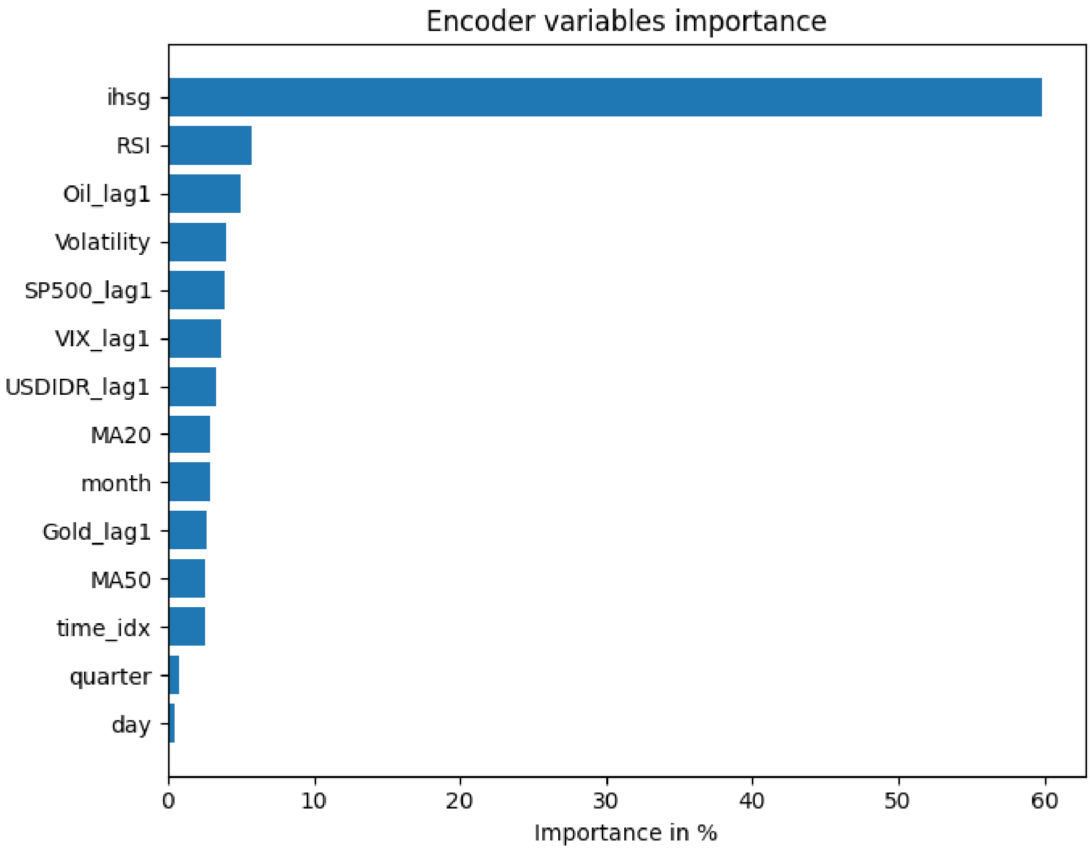

Figure 4. Variable Importance Scores Generated by TFT Variable Selection Network

Figure 4 presents the variable importance scores generated by the TFT variable selection network. Historical IHSG values contributed the largest predictive importance, indicating that previous index movements contain significant temporal information related to future IHSG behavior. Among the technical indicators, RSI showed relatively high importance, suggesting that momentum-related information contributes to forecasting market movements. Furthermore, the contribution of crude oil prices, volatility, and S&P 500 movements indicates that global financial conditions provide additional predictive information for representing IHSG dynamics.

The variable importance results highlight the interpretability advantage of TFT, as the model provides information regarding the relative contribution of input variables used during forecasting. However, these importance scores represent predictive associations within the model and should not be interpreted as direct causal relationships.

6. Thirty-Day IHSG Forecast

The optimized TFT model was applied to generate a 30-trading-day ahead IHSG forecast using quantile-based prediction outputs. The forecast consists of the median prediction (Q0.50) and prediction intervals represented by the lower and upper quantiles (Q0.10 and Q0.90), providing an estimation of possible future IHSG movements under uncertainty conditions.

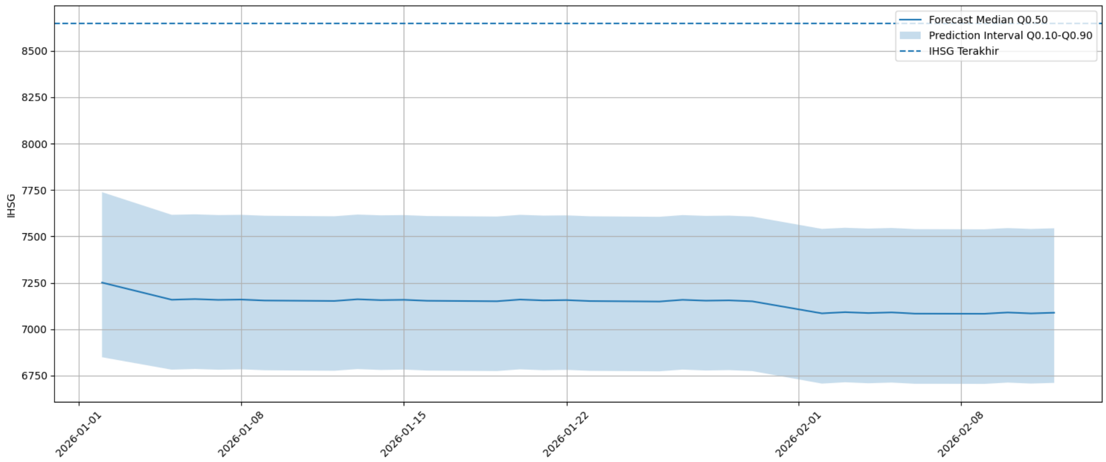

Figure 5. Thirty-Day Ahead IHSG Forecast with Quantile

Prediction Intervals Generated by TFT

Figure 5 illustrates the 30-trading-day ahead IHSG forecast generated by the optimized TFT model. The median forecast (Q0.50) represents the expected IHSG movement, while the prediction interval between Q0.10 and Q0.90 represents the uncertainty range associated with future market conditions. The relatively stable forecast pattern indicates that the model captures the underlying temporal structure of IHSG movements while maintaining uncertainty representation through quantile-based forecasting.

The prediction interval provides additional information regarding potential future variations in IHSG values, which is particularly important in financial markets characterized by volatility and unexpected external influences. Therefore, the TFT forecasting framework not only generates future index estimates but also provides information regarding the level of uncertainty associated with those predictions.

DISCUSSION

The findings demonstrate that the Temporal Fusion Transformer (TFT) provides a flexible framework for multivariate and multi-horizon forecasting of the Indonesian Composite Stock Price Index (IHSG). The model successfully integrates historical index movements, global financial indicators, technical indicators, and calendar features to generate quantile-based forecasts while simultaneously providing information regarding variable relevance and temporal dependencies. The variable selection mechanism indicates that historical IHSG values represent the most dominant predictive information, followed by technical and external market variables, including RSI, crude oil prices, volatility, and S&P 500 movements. These findings suggest that IHSG forecasting is strongly influenced by its own historical dynamics, while external financial conditions provide additional information for representing market behavior.

The dominant contribution of historical IHSG values is consistent with the characteristics of financial time-series data, where previous market movements contain important temporal patterns associated with future observations. However, the variable importance generated by TFT should not be interpreted as causal relationships, because the selection weights represent the predictive contribution of each variable within the model structure rather than direct economic influence. The contribution of RSI indicates the relevance of momentum-based technical information in forecasting index movements, while the importance of crude oil prices, volatility, and S&P 500 movements reflects the influence of global financial conditions on the Indonesian stock market. These results support previous findings that technical indicators and international market variables can improve the representation of complex financial market dynamics.

Although TFT provides advantages in interpretability and probabilistic forecasting, the results indicate that higher model complexity does not automatically guarantee improved forecasting performance. Financial time-series data such as IHSG are characterized by high volatility, structural changes, and uncertainty arising from various internal and external market factors. Consequently, forecasting performance depends not only on model architecture but also on market conditions, forecasting horizon, and the availability of relevant predictive information. Therefore, the contribution of TFT should not be evaluated solely based on numerical forecasting accuracy but also through its ability to provide additional insights regarding uncertainty and model decision-making processes.

The forecasting results demonstrate that TFT is capable of generating reliable point forecasts while simultaneously providing probabilistic forecasting outputs through quantile estimation. The integration of prediction intervals allows the model to represent uncertainty associated with future IHSG movements, which is particularly relevant in volatile financial environments. Furthermore, the variable selection network and attention mechanism enable the identification of important predictive variables and temporal patterns that influence the forecasting process. This capability provides an advantage over conventional black-box forecasting approaches because the model offers greater transparency regarding the information used to generate predictions.

The probabilistic forecasting evaluation provides additional insight into the capability and limitation of TFT. The absence of quantile crossing indicates that the model maintains consistent ordering among generated quantile forecasts, ensuring logical relationships between lower, median, and upper prediction intervals. However, the prediction interval coverage indicates that uncertainty estimation still requires further calibration. This finding highlights that probabilistic forecasting involves not only producing multiple forecast scenarios but also ensuring that the estimated uncertainty range accurately represents possible future variations. Therefore, future improvements should focus on enhancing uncertainty calibration while maintaining forecasting reliability.

From a methodological perspective, this study contributes to the application of interpretable deep learning models in Indonesian financial forecasting. Previous IHSG forecasting studies have primarily focused on improving point prediction performance using statistical models or conventional machine learning approaches. This study extends previous research by applying TFT as an interpretable forecasting framework that combines multi-horizon prediction, probabilistic forecasting, variable selection, and attention mechanisms. The findings demonstrate that understanding which variables contribute to model predictions is an important aspect of financial forecasting, particularly when dealing with complex and dynamic market environments.

The findings also have practical implications for financial analysis and market monitoring. The identification of historical IHSG movements, RSI, global commodity prices, volatility, and S&P 500 movements as important predictive inputs provides information regarding the market signals captured by the forecasting framework. However, TFT should not be interpreted as a standalone investment decision system because financial markets remain affected by unexpected events, policy changes, investor sentiment, and other external factors beyond the variables incorporated in this study. Instead, the model can serve as a decision-support framework by providing predictive information and insights into relevant market factors.

Several limitations should be considered when interpreting the findings. First, this study focuses on IHSG level forecasting using a single historical dataset covering the period from 2000 to 2025, which may limit the generalizability of the findings across different market regimes. Second, the forecasting framework relies on available external variables, while future market conditions may differ significantly due to unexpected economic and geopolitical events. Third, although TFT provides probabilistic forecasts, the uncertainty estimation still requires further improvement through calibration strategies. Future research may address these limitations by applying rolling-origin evaluation across different volatility regimes, incorporating additional macroeconomic and sentiment-related indicators, exploring alternative forecasting targets such as returns or volatility, and developing improved TFT-based calibration approaches.

CONCLUSION

This study investigated the application of the Temporal Fusion Transformer (TFT) for multivariate and multi-horizon forecasting of the Indonesian Composite Stock Price Index (IHSG) by integrating historical index movements, global financial indicators, technical indicators, and calendar features. The findings demonstrate that TFT is capable of generating interpretable quantile-based forecasts while identifying the relative contribution of different input variables throughout the forecasting process. Historical IHSG movements were identified as the most dominant predictive information, while RSI, crude oil prices, volatility, and S&P 500 movements provided additional information for representing complex market dynamics.

This study contributes to financial forecasting research by demonstrating the application of an interpretable deep learning framework that integrates multi-horizon forecasting, probabilistic prediction, and variable importance analysis within a single forecasting approach. Unlike conventional forecasting frameworks that primarily focus on numerical prediction accuracy, TFT provides additional insights into predictive variables and temporal patterns through its variable selection network and attention mechanism. The findings indicate that the value of TFT extends beyond point forecasting performance by providing a more transparent forecasting framework capable of representing uncertainty and explaining the information used during the prediction process.

The probabilistic forecasting results further demonstrate that TFT is able to generate quantile-based prediction intervals while maintaining consistent relationships among forecast quantiles. However, the uncertainty estimation produced by the model still requires further calibration to improve the reliability of prediction intervals under different market conditions. These findings highlight that financial forecasting models should not only focus on generating accurate predictions but also provide reliable uncertainty representation, particularly in highly dynamic environments such as stock markets.

Future research should focus on improving the robustness and reliability of TFT-based financial forecasting. Further studies may apply rolling-origin evaluation across different market regimes, incorporate additional macroeconomic and sentiment-related variables, explore alternative forecasting targets such as stock returns or volatility, and develop improved calibration strategies to enhance probabilistic forecast reliability.

REFERENCES

Andriyani, R. A., & Pimada, L. M. (2025). The influence of global risks on the Composite Index. *Journal of Development Economic and Social Studies, 4*(4), 1111–1124. https://doi.org/10.21776/jdess.2025.04.4.05

Fatmawati, A. K., Badriah, L. S., & Gunawan, D. S. (2024). Perkembangan Indeks Harga Saham Gabungan (IHSG) sebelum dan sesudah Krisis Ekonomi Global Tahun 2008. *Journal of Economic, Management, Accounting and Technology (JEMATech), 7*(1), 31–39. https://doi.org/10.32500/jematech.v7i1.5233

Fauziyah, E., Ispriyanti, D., & Tarno. (2021). Pemodelan dan Peramalan Indeks Harga Saham Gabungan (IHSG) menggunakan ARIMAX-TARCH. *Jurnal Gaussian, 10*(4), 595–604. https://doi.org/10.14710/j.gauss.10.4.595-604

Fitria, A. E., & Komara, E. F. (2026). Do macroeconomic factors matter for Indonesia's stock market? Evidence from the Jakarta Composite Index. *International Journal of Management, Entrepreneurship, Social Science and Humanities, 10*(1), 16–34. https://doi.org/10.31098/ijmesh.v10i1.4142

Galena, M. V., Sediono, Mardianto, M. F. F., & Pusporani, E. (2025). Prediction analysis of Jakarta Composite Index movement using support vector regression method. *G-Tech: Jurnal Teknologi Terapan, 9*(1), 68–77. https://doi.org/10.70609/gtech.v9i1.5879

Hartanto, S., & Gunawan, A. A. S. (2024). Temporal fusion transformers for enhanced multivariate time series forecasting of Indonesian stock prices. *International Journal of Advanced Computer Science and Applications, 15*(7), 140–148. https://doi.org/10.14569/IJACSA.2024.0150713

Hidayat, R., & Sudjono. (2022). The effect of world gold price, world oil price, USD/IDR exchange rate, and inflation on the Joint Stock Price Index (JCI) on the Indonesia Stock Exchange (IDX). *Scholars Bulletin, 8*(1), 33–41. https://doi.org/10.36348/sb.2022.v08i01.005

Lim, B., Arık, S. Ö., Loeff, N., & Pfister, T. (2021). Temporal Fusion Transformers for interpretable multi-horizon time series forecasting. *International Journal of Forecasting, 37*(4), 1748–1764. https://doi.org/10.1016/j.ijforecast.2021.03.012

Lubis, F. S., Hanifah, & Mayuta, F. P. (2023). Analisis Peramalan (*Forecasting*) Jumlah Pemakaian Bahan Bakar PLTD Bengkalis pada PT PLN UP3 Dumai. *Seminar Nasional Teknologi Informasi, Komunikasi dan Industri (SNTIKI 15)*, 254–264.

Lumbanbatu, M. T. G., Kartikasari, P., & Fakhriyana, D. (2025). Peramalan Indeks Harga Saham Gabungan (IHSG) dengan Metode *Fuzzy Time Series* Chen dan *Fuzzy Time Series* Cheng. *Jurnal Gaussian, 14*(1), 469–479. https://doi.org/10.14710/j.gauss.14.1.469-479

Najibullah. (2023). Pengaruh Indeks Global terhadap Indeks Harga Saham Gabungan (IHSG): Bukti Empiris menggunakan Data Harian. *Jurnal Mahasiswa Akuntansi Samudra (JMAS), 4*(4), 171–178.

Nurhasan, & Ghufron. (2026). Perbandingan Metode *Temporal Fusion Transformer* (TFT) dan *Long Short-Term Memory* (LSTM) dalam Prediksi Harga Saham Indonesia Berbasis Data Teknikal dan Fundamental. *Applied Information Technology and Computer Science, 5*(1), 118–125.

Otoritas Jasa Keuangan. (2019). *Buku 3: Pasar Modal (Seri Literasi Keuangan Perguruan Tinggi).*

Patriya, E., Latif, A., & Handayani. (2023). Peramalan Harga Saham Penutupan Indeks Harga Saham Gabungan (IHSG) menggunakan Algoritma *Long Short-Term Memory* (LSTM). *Jurnal Ilmiah Ekonomi Bisnis, 28*(2), 304–314. https://doi.org/10.35760/eb.2023.v28i2.7964

Tashman, L. J. (2000). Out-of-sample tests of forecasting accuracy: An analysis and review. *International Journal of Forecasting, 16*(4), 437–450. https://doi.org/10.1016/S0169-2070(00)00065-0

Tjen, Y. C. (2025). *Peramalan IHSG dengan Metode ARIMA-GARCH dan LSTM pada Periode Sebelum, Masa dan Sesudah COVID-19* [Thesis, Universitas Islam Indonesia].

Tsai, C.-F., & Hsiao, Y.-C. (2010). Combining multiple feature selection methods for stock prediction: Union, intersection, and multi-intersection approaches. *Decision Support Systems, 50*(1), 258–269. https://doi.org/10.1016/j.dss.2010.08.028