Online Journal of Management, Innovation, Economics, and Digital Studies (ONMIND)
Vol. 1 No. 1 (2026)

Deep Learning for Investment Decision Support: A SimpleRNN Approach to Stock Price Prediction with Real-Time API Deployment

Syamsu Alam[1,*], Muh. Nur Syamsi Hidayah[1], Muh. Ashif[1], Rizal Bakri[1], Muh. Qardawi Hamzah[1]

[1]Department of Digital Business, Faculty of Economics and Business, Universitas Negeri Makassar, Makassar 90222, *Indonesia*

Abstract

Stock price prediction remains one of the most demanding challenges in financial analytics due to the inherently non-linear and temporally dependent nature of market data. This study proposes a deep learning framework based on Simple Recurrent Neural Network (SimpleRNN) for next-day closing price prediction of PT Telkom Indonesia (Persero) Tbk (TLKM.JK), one of Indonesia's most actively traded blue-chip stocks on the Indonesia Stock Exchange (IDX). The model was trained on 4,434 daily OHLCV observations from 2008 to 2026 sourced via Yahoo Finance, augmented with six engineered features encompassing momentum, trend, and volatility indicators. A sliding window of ten timesteps was employed to encode sequential dependencies. To ensure result reliability, the model was evaluated across five independent runs under identical hyperparameter configurations. The mean performance metrics were: MAE = Rp 67.50, RMSE = Rp 89.13, MAPE = 2.00%, and $R^2 = 0.9775$, indicating high predictive accuracy and cross-run stability. Error analysis revealed a near-symmetric distribution (skewness = 0.029) with a slight overestimation tendency (56.4%). The trained model was subsequently integrated into a FastAPI-based REST endpoint with SQLite persistence and an interactive multi-range candlestick dashboard, forming a deployable investment decision-support prototype. These findings demonstrate that SimpleRNN, when combined with systematic feature engineering and robust deployment architecture, can serve as a practical and computationally efficient tool for short-term equity price forecasting in emerging digital business contexts.

Keywords: *deep learning; investment decision support; recurrent neural network; stock price prediction; time series forecasting*

1. Introduction

The Indonesia Stock Exchange (IDX) has experienced remarkable growth over the past decade, with the number of retail investors increasing by more than 37.5 percent between 2020 and 2022 alone, reaching over ten million registered accounts by early 2023 (Haryono et al., 2024). This surge in participation has intensified demand for data-driven tools capable of assisting investors in navigating the inherent complexity of equity markets. Stock prices are shaped by a confluence of macroeconomic signals, corporate fundamentals, geopolitical events, and collective investor psychology, making accurate prediction extraordinarily difficult using conventional analytical frameworks. Classical time-series models such as ARIMA assume linear data generation processes and stationary distributions, assumptions that are routinely violated by financial data exhibiting volatility clustering, fat tails, and structural breaks (Khan et al., 2023).

Artificial intelligence and machine learning have emerged as powerful alternatives for financial forecasting. Deep learning architectures in particular are distinguished by their capacity to identify hierarchical, non-linear patterns in large-scale sequential datasets without requiring explicit feature specification (Zhang et al., 2024). Among deep learning families, Recurrent Neural Networks (RNNs) occupy a central position due to their architectural affinity with temporal dependencies: the recurrent hidden state enables the model to carry forward information from earlier time steps when processing each new observation, a property directly analogous to the sequential nature of price discovery in financial markets (Sako et al., 2022). This fundamental capability makes RNN-based architectures the most frequently applied class of deep learning models in stock price forecasting studies (Dwiandiyanta et al., 2025).

[*]Corresponding author:
E-mail address: alam.s@unm.ac.id

PT Telkom Indonesia (Persero) Tbk (ticker: TLKM.JK) constitutes a particularly compelling research subject. As Indonesia's largest state-owned telecommunications conglomerate, Telkom commands significant weight in the IDX Composite Index (IHSG), and its shares are consistently among the most liquid on the exchange. The stock has traversed multiple market regimes over the past two decades, including the post-global-financial-crisis recovery, the digital-economy boom of 2014–2018, the COVID-19 crash and rebound of 2020–2021, and the correction cycle of 2022–2025. This richness of historical dynamics provides an ideal laboratory for evaluating a model's capacity to generalize across varied market conditions. Existing studies have examined Telkom's stock using methods ranging from ARIMA and SVR to LSTM and CNN-LSTM hybrids (Pratama et al., 2025; Hadi & Priyanto, 2025), yet comparatively little attention has been paid to evaluating SimpleRNN as a baseline model specifically for this asset, particularly when feature engineering is applied systematically.

Feature engineering represents a critical but underexplored dimension of time-series prediction pipelines. Raw OHLCV data alone fails to capture latent market dynamics such as short-term momentum, trend persistence, and intraday volatility clustering. By constructing auxiliary features, specifically lagged returns, moving averages, rolling standard deviation, and an RSI slope indicator, this study enriches the input representation and equips the model with domain-relevant signals while preserving the computational advantages of a relatively simple RNN architecture. This approach is broadly consistent with the findings of Dwiandiyanta et al. (2025), who demonstrated that technical indicator augmentation substantially improves deep learning prediction accuracy on IDX blue-chip stocks.

A distinguishing contribution of this study lies in its integration of model development with production deployment. Rather than treating prediction as an isolated academic exercise, the system is architecturally designed for operational use: the trained SimpleRNN model is served via a FastAPI REST endpoint, predictions are persisted in a SQLite database for auditability, and an interactive candlestick dashboard supports multi-range visualization and real-time inference. This end-to-end perspective aligns with the broader imperative in digital business studies to bridge the gap between algorithmic innovation and practical investment decision support. The methodological rigor adopted here, including five independent evaluation runs to assess cross-run stability, is analogous to multi-trial performance evaluation frameworks employed in machine learning model assessment in adjacent domains, as demonstrated by Bakri et al. (2025) in educational data mining.

The research objectives are threefold. First, to develop and evaluate a SimpleRNN model for next-day closing price prediction of TLKM.JK using a technically enriched feature set. Second, to quantify the model's predictive accuracy and cross-run consistency through multiple independent evaluation trials. Third, to implement and demonstrate a deployable investment decision-support system built around the trained model. The remainder of this paper is organized as follows: Section 2 reviews relevant literature; Section 3 describes the methodology; Section 4 presents results; Section 5 discusses findings; and Section 6 concludes.

2. Literature Review

2.1. Stock Price Prediction and the Efficient Market Challenge

The Efficient Market Hypothesis (EMH), in its semi-strong form, posits that all publicly available information is instantaneously reflected in asset prices, rendering systematic outperformance impossible through technical analysis alone. Yet empirical research consistently demonstrates that real financial markets deviate from this theoretical ideal due to behavioral biases, information asymmetries, and structural rigidities (Khan et al., 2023). This tension between EMH and observed market behavior has motivated decades of research into predictive modeling, with machine learning emerging as the dominant paradigm in the past decade. A comprehensive performance comparison across multiple stock markets by Khan et al. (2023) found that ensemble methods and neural networks generally outperform linear statistical models, suggesting meaningful exploitable structure in historical price series, structure that algorithms can detect and quantify.

Data science approaches to Indonesian stock markets have gained particular traction following the post-COVID digital transformation. Budiharto (2021) applied LSTM to Indonesian stock forecasting during the pandemic period and demonstrated that deep learning models could successfully navigate the extreme volatility induced by the crisis, achieving MAPE values competitive with pre-pandemic benchmarks. Haryono et al. (2024) extended this work by comparing CNN, GRU, LSTM, and GCN architectures across 825 IDX-listed companies, finding that no single architecture universally dominates but that temporal recurrent models consistently rank among the top performers for liquid, high-capitalization equities. These findings collectively establish the feasibility of deep learning for Indonesian equity markets and motivate the present study's focus on TLKM.JK.

2.2. Recurrent Neural Networks for Time-Series Financial Forecasting

Recurrent Neural Networks (RNNs) are a class of artificial neural networks specifically architected to process sequential data by maintaining an internal hidden state that propagates information across time steps. Unlike feedforward networks that treat each input independently, an RNN encodes the full temporal history of a sequence into a compact latent vector at each step. The fundamental recurrence governing a SimpleRNN cell at time step *t* is:

$$h_t = tanh(\boldsymbol{W}_{hh} \cdot \boldsymbol{h}_{t-1} + \boldsymbol{W}_{xh} \cdot \boldsymbol{x}_t + \boldsymbol{b}_h)$$

where $h_t \in \mathbb{R}^d$ is the hidden state vector at time t; $x_t \in \mathbb{R}^p$ is the input feature vector at time t (here $p = 11$ engineered features); $\boldsymbol{W}_{hh} \in \mathbb{R}^{d \times d}$ is the recurrent weight matrix governing state-to-state transitions; $\boldsymbol{W}_{xh} \in \mathbb{R}^{p \times d}$ is the input weight matrix; and $\boldsymbol{b}_h \in \mathbb{R}^d$ is the bias vector. The scalar next-day price prediction is obtained via the output layer:

$$\hat{y}_t = \boldsymbol{W}_o \cdot \boldsymbol{h}_T + \boldsymbol{b}_o$$

where $\boldsymbol{W}_o$ and $\boldsymbol{b}_o$ are the output layer parameters and $T = 10$ is the sequence length. This mechanism enables the model to encode temporal context: each prediction implicitly conditions on the entire sequential history encoded in h_t, rather than treating each observation independently as feedforward networks do.

Sako et al. (2022) conducted a rigorous multi-market evaluation of RNN, LSTM, and GRU for financial time-series forecasting across eight stock indices and six currency exchange rates, spanning data from 2008 to 2021. Their results revealed that while GRU exhibited the best aggregate performance in out-of-sample forecasting, vanilla RNN remained competitive for shorter-horizon predictions, with the performance gap narrowing considerably when sufficient historical context was provided through appropriate sequence length selection. This finding is relevant to the present study, which employs a ten-day lookback window calibrated to capture short-term market dynamics without introducing excessive model complexity.

The principal limitation of simple RNN architectures is the vanishing gradient problem, wherein gradients diminish exponentially during backpropagation through time, compromising the model's capacity to capture long-range dependencies (Zhang et al., 2024). LSTM and GRU architectures address this through gating mechanisms: the former employs three gates (input, forget, output) and a distinct cell state, while the latter combines this functionality into two gates (reset and update) with reduced parametric complexity. However, for short-horizon prediction tasks where the relevant temporal context spans days rather than weeks, the vanishing gradient limitation of SimpleRNN is substantially mitigated, and the architecture's computational economy becomes an advantage, particularly in deployment scenarios requiring rapid inference with minimal latency.

Several studies have specifically examined deep learning prediction of PT Telkom Indonesia's stock. Hadi and Priyanto (2025) compared RNN, LSTM, and GRU for TLKM, reporting that LSTM generally achieved superior directional accuracy while RNN demonstrated competitive regression performance. Pratama et al. (2025) pursued a hybrid CNN-LSTM architecture and demonstrated improved feature extraction relative to single-architecture baselines. Switrayana et al. (2025) investigated the influence of data scaling methods on deep learning prediction accuracy for Indonesian stocks including TLKM, finding that MinMax normalization consistently outperformed standard normalization for neural network-based regressors. These findings provide direct comparators for the present study and collectively underscore the continued research interest in both the asset and the modeling family.

2.3. Feature Engineering for Equity Price Prediction

A consensus has emerged in the deep learning forecasting literature that raw OHLCV data alone is suboptimal as model input, particularly for recurrent architectures that depend on richness of the input representation to construct informative hidden states. Technical indicators derived from price and volume data capture latent market dynamics, momentum, trend, volatility, that are not directly observable in raw prices but are demonstrably predictive of future price movements (Dwiandiyanta et al., 2025). Moving averages provide smoothed estimates of the underlying trend by averaging prices over rolling windows, with the MA-5 and MA-10 indicators capturing short-term and medium-term trend signals, respectively. The Relative Strength Index (RSI) measures the magnitude of recent price changes to evaluate overbought or oversold conditions, and its first-order difference (RSI_SLOPE) captures the rate of change in momentum, a signal shown to carry incremental predictive content beyond the raw RSI level.

Return lags (the percentage change in closing price at lag one and lag two) encode short-term serial correlation in returns. While daily equity returns are generally close to white noise on average, the conditional distribution of returns at lag one exhibits statistically detectable autocorrelation in high-frequency trading environments, particularly for liquid, institutionally traded stocks like TLKM. Rolling standard deviation of returns over a five-day window serves as a proxy for realized short-term volatility, capturing the volatility clustering phenomenon commonly observed in equity return series. The decision to construct exactly eleven features (five OHLCV + six engineered) balances informational richness against the risk of overfitting given the training sample size, a consideration highlighted by Bakri et al. (2025) in the context of model performance under varying feature regimes.

2.4. Model Deployment for Investment Decision Support

The translation of predictive models into operational decision-support tools represents a distinct challenge that has received comparatively limited attention in the academic literature, despite its paramount practical importance. A fully realized deployment pipeline must address data freshness (ensuring the model receives up-to-date market data), inference consistency (applying the same preprocessing and normalization pipeline used during training), result persistence (storing predictions for audit and backtesting), and user accessibility (presenting outputs through an intuitive interface). FastAPI has emerged as a preferred framework for machine learning model serving due to its asynchronous request handling, automatic OpenAPI documentation, and native Pydantic integration, which collectively facilitate rapid development of production-grade REST endpoints. The integration of deep learning forecasting systems with interactive dashboards represents the current frontier of financial technology development, enabling real-time visualization of both historical price dynamics and forward projections in a manner accessible to non-specialist users.

3. Methods

3.1. Research Design

This study employs a quantitative, applied research design following the CRISP-DM (Cross-Industry Standard Process for Data Mining) lifecycle. The philosophical orientation is positivist: the research treats historical stock price data as an objective empirical record from which predictive patterns can be extracted through systematic algorithmic analysis. The study is explanatory in character, it seeks not merely to generate predictions but to illuminate the structural relationships between technical market indicators and next-day closing prices through the lens of a recurrent neural network model.

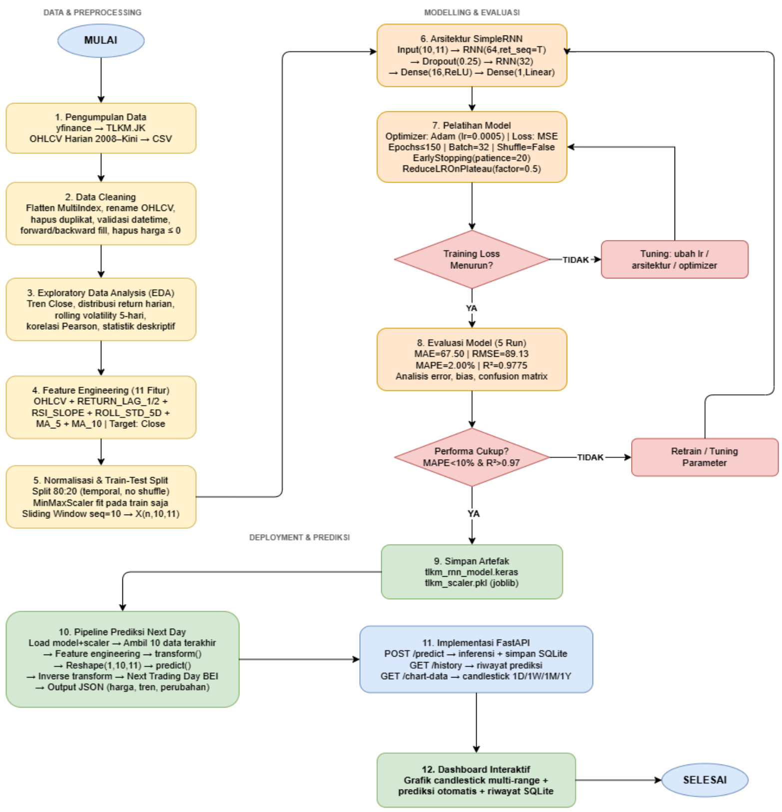

Figure 1. Research workflow for SimpleRNN-based stock price prediction (CRISP-DM adapted)

3.2. Data Sources and Collection

The dataset comprises daily historical price records for PT Telkom Indonesia (Persero) Tbk (ticker: TLKM.JK) retrieved from Yahoo Finance via the yfinance Python library. The collection period spans from January 2008 to the script execution date in early 2026, yielding 4,449 raw observations across six variables: trading date (index), Open, High, Low, Close, and Volume (OHLCV). Table 1 summarizes the key dataset characteristics.

Table 1. Dataset characteristics

Component	Description
Source	Yahoo Finance (yfinance library)
Ticker	TLKM.JK (PT Telkom Indonesia)
Period	2008 – 2026 (execution date)
Raw observations	4,449 daily trading records
Post-engineering observations	4,434 (after lag/rolling window NaN removal)
Training set	3,547 (80%)
Testing set	887 (20%)
Input features	11 (OHLCV + 6 engineered)
Target variable	Close price (next-day)
Sequence length	10 timesteps

The closing price (Close) was selected as the prediction target on the basis of its widespread use as the definitive daily valuation benchmark in technical analysis. Descriptive statistics of the raw OHLCV data confirm a wide price range (minimum approximately Rp 970; maximum approximately Rp 4,850), consistent with the multi-decade dataset spanning diverse market regimes including the 2013–2017 bull run and the 2020 COVID-19 crash.

3.3. Data Preprocessing

The preprocessing pipeline addresses five quality dimensions. First, MultiIndex column structures produced by the yfinance library were flattened to single-level labels via string concatenation. Second, the dataset index was converted to a validated DatetimeIndex, with invalid timestamps replaced by NaT and subsequently removed; duplicate dates were eliminated to preserve temporal integrity. Third, missing values were imputed via forward-fill followed by backward-fill, ensuring no information from the future is used in imputation. Fourth, rows with zero or negative price values were removed as logical impossibilities in the equity price domain. Fifth, duplicate columns arising from the MultiIndex flattening were collapsed to retain unique column names. The preprocessing produced zero data loss, confirming that the Yahoo Finance feed for TLKM.JK is complete across the study period.

3.4. Feature Engineering

Six additional features were constructed from the cleaned OHLCV base to augment the model's input representation. RETURN_LAG_1 and RETURN_LAG_2 are the daily percentage returns at lags one and two, respectively, capturing short-term return momentum. MA_5 and MA_10 are simple moving averages over rolling five-day and ten-day windows, providing trend direction signals at short and medium horizons. ROLL_STD_RETURN_5D is the rolling five-day standard deviation of daily returns, representing a realized short-term volatility proxy. RSI_SLOPE is the first-order difference of the 14-period Relative Strength Index, capturing changes in the strength of buying or selling momentum. Because the lag and rolling features require historical antecedents, the first few rows generated NaN values and were subsequently dropped, reducing the working dataset from 4,449 to 4,434 observations.

3.5. Data Normalization and Sequence Construction

All eleven features were normalized to the [0, 1] range using MinMaxScaler fitted exclusively on the training partition to prevent data leakage, a methodological safeguard emphasized by Switrayana et al. (2025). The normalized data was then transformed into overlapping sequences via a sliding window of length ten: each input sample X[i] consists of the ten-day feature matrix spanning days i through i+9, while the corresponding label y[i] is the normalized Close value on day i+10. This produces a three-dimensional input tensor of shape (samples, 10, 11). The dataset was partitioned chronologically at an 80:20 ratio, yielding 3,547 training sequences and 887 test sequences.

3.6. Model Architecture and Hyperparameters

The model architecture consists of four sequential layers. The first layer is a SimpleRNN cell with 64 units and return_sequences=True, processing the ten-step input sequence and producing a 64-dimensional hidden state at each timestep. A Dropout layer with rate 0.25 follows to regularize the recurrent representations. The second SimpleRNN layer contains 32 units with return_sequences=False, compressing the full temporal sequence into a single 32-dimensional vector. A Dense layer with 16 units and ReLU activation maps this latent representation to an intermediate space before the final Dense output layer with a single linear unit generates the scalar price prediction. The model was compiled with the Adam optimizer at a learning rate of 0.0005 and Mean Squared Error (MSE) as the loss function, with MAE tracked as an auxiliary metric. A random seed of 42 was fixed across NumPy and TensorFlow to maximize reproducibility. Training was conducted for up to 150 epochs with a batch size of 32 and shuffle=False to preserve temporal order. EarlyStopping with patience=20 and restore_best_weights=True prevented overfitting, while ReduceLROnPlateau with factor=0.5 and patience=6 adapted the learning rate during plateau phases.

3.7. Evaluation Protocol

To quantify cross-run stability and mitigate the influence of random weight initialization, the model was trained and evaluated across five independent runs under identical hyperparameter configurations, with only the stochastic initialization differing between runs. Performance was assessed using four metrics: Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), Mean Absolute Percentage Error (MAPE), and the coefficient of determination (R^2). Additionally, an error bias analysis categorized predictions as overestimates or underestimates, and a confusion matrix was constructed for directional accuracy (predicting whether the next-day price rises or falls) to characterize the model's classification limitations.

4. Results

4.1. Model Training Convergence

Figure 2 presents the training and validation loss curves for a representative run. Both curves exhibit rapid convergence during the first twenty epochs, followed by a gradual asymptotic approach to a stable minimum. Critically, the training and validation losses remain closely aligned throughout training, with no visible divergence at any point, indicating the absence of significant overfitting. The EarlyStopping callback consistently terminated training between 60 and 90 epochs across the five runs, well before the 150-epoch maximum, confirming that the model converged to its optimal generalization state substantially ahead of the epoch budget. The ReduceLROnPlateau mechanism triggered on average 2.4 times per run, progressively reducing the learning rate from 0.0005 to approximately 0.0001 in most runs, facilitating fine-grained convergence near the loss minimum.

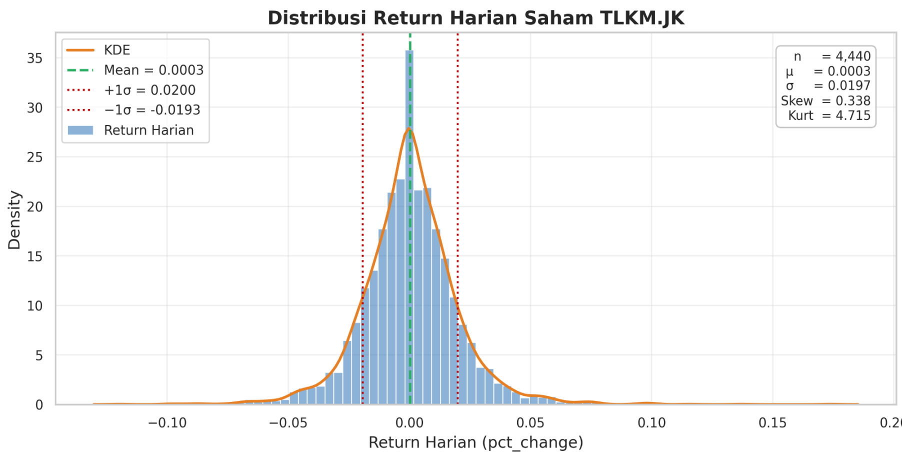

Figure 2. Training and validation loss curves (representative run, MSE scale). Convergence without divergence confirms absence of overfitting.

4.2. Prediction Accuracy Across Five Runs

Table 2 presents the four evaluation metrics across the five independent runs and their arithmetic mean. The results demonstrate high predictive accuracy and, crucially, strong cross-run consistency. The mean MAE of Rp 67.50 indicates that the model's absolute prediction error is, on average, less than Rp 70 on a price series that routinely trades between Rp 2,500 and Rp 4,500, representing a relative error of approximately 2.0%. The RMSE of Rp 89.13 slightly exceeds the

MAE, consistent with a distribution of errors that is generally well-behaved with a modest number of larger deviations. The mean MAPE of 2.00% falls well below the commonly cited 10% threshold for acceptable forecasting accuracy, placing the model in the 'highly accurate' range by standard financial forecasting benchmarks. The mean R^2 of 0.9775 implies that the model explains 97.75% of the variance in the test set's closing prices, a figure that falls within the narrow band of 0.9765 to 0.9784 across all five runs (a span of only 0.0019), confirming the stability of the learned representations. Figure 3 illustrates the model's predictions plotted against actual closing prices for the full test period, visually confirming the high correspondence between predicted and realized values.

Table 2. Model performance across five independent evaluation runs

Run	MAE (Rp)	RMSE (Rp)	MAPE (%)	R^2
1	65.58	87.51	1.96	0.9784
2	68.65	90.43	2.05	0.9769
3	69.93	91.11	2.07	0.9765
4	67.41	88.84	1.98	0.9777
5	65.95	87.77	1.96	0.9782
Mean	**67.50**	**89.13**	**2.00**	**0.9775**

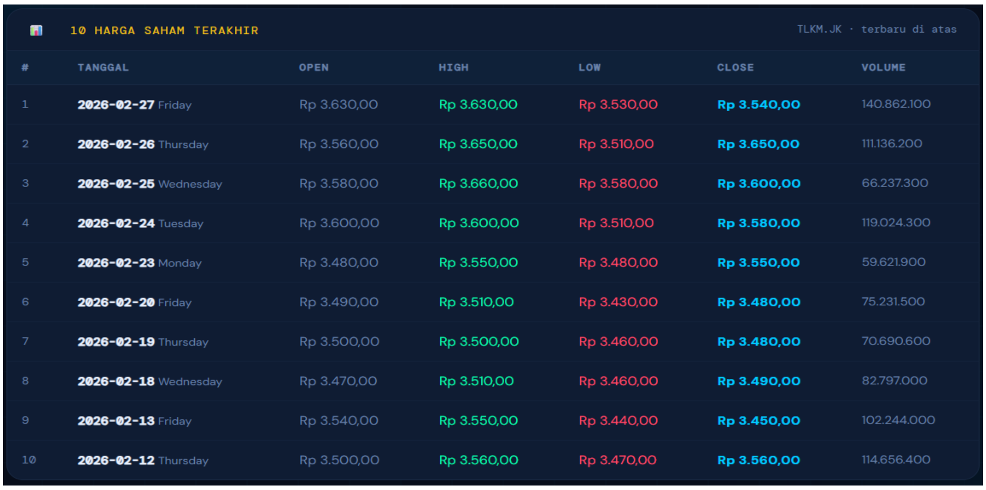

10 HARGA SAHAM TERAKHIR — TLKM.JK · terbaru di atas

#	TANGGAL	OPEN	HIGH	LOW	CLOSE	VOLUME
1	2026-02-27 Friday	Rp 3.630,00	Rp 3.630,00	Rp 3.530,00	Rp 3.540,00	140.862.100
2	2026-02-26 Thursday	Rp 3.560,00	Rp 3.650,00	Rp 3.510,00	Rp 3.650,00	111.136.200
3	2026-02-25 Wednesday	Rp 3.580,00	Rp 3.660,00	Rp 3.580,00	Rp 3.600,00	66.237.300
4	2026-02-24 Tuesday	Rp 3.600,00	Rp 3.600,00	Rp 3.510,00	Rp 3.580,00	119.024.300
5	2026-02-23 Monday	Rp 3.480,00	Rp 3.550,00	Rp 3.480,00	Rp 3.550,00	59.621.900
6	2026-02-20 Friday	Rp 3.490,00	Rp 3.510,00	Rp 3.430,00	Rp 3.480,00	75.231.500
7	2026-02-19 Thursday	Rp 3.500,00	Rp 3.500,00	Rp 3.460,00	Rp 3.480,00	70.690.600
8	2026-02-18 Wednesday	Rp 3.470,00	Rp 3.510,00	Rp 3.460,00	Rp 3.490,00	82.797.000
9	2026-02-13 Friday	Rp 3.540,00	Rp 3.550,00	Rp 3.440,00	Rp 3.450,00	102.244.000
10	2026-02-12 Thursday	Rp 3.500,00	Rp 3.560,00	Rp 3.470,00	Rp 3.560,00	114.656.400

Figure 3. *Actual versus predicted TLKM.JK closing prices on the test set (887 observations). The prediction series closely tracks the actual trajectory across all market phases.*

4.3. Error Distribution and Bias Analysis

Table 3 summarizes the statistical properties of the prediction error distribution. Of the 887 test sequences, 878 observations were included in the bias analysis after excluding the nine edge-case observations where the reference closing price was unavailable for percentage calculation.

Table 3. Prediction error statistics on the test set (n = 878 effective observations)

Statistic	Value
Mean Error	Rp -8.92
Median Error	Rp -11.34
MAE (Test Set)	Rp 64.18
Std Dev Error	Rp 85.73
Skewness	0.0286
Min Error	Rp -404.50
Max Error	Rp 328.96
Overestimate	56.4% (n=495)
Underestimate	43.6% (n=383)

The near-zero mean error (Rp -8.92) and median error (Rp -11.34) confirm that the model does not exhibit a systematic bias of practical magnitude: both values represent less than 0.3% of the average closing price. The slight negativity of these central tendency measures implies a weak tendency toward overestimation, corroborated by the directional analysis showing 56.4% overestimate occurrences versus 43.6% underestimates. The error distribution skewness of 0.029 is essentially symmetric, indistinguishable from zero in practical terms, indicating that large positive and negative errors are approximately balanced in frequency. Figure 4 presents the distribution of prediction errors, which visually confirms the near-Gaussian, zero-centered profile.

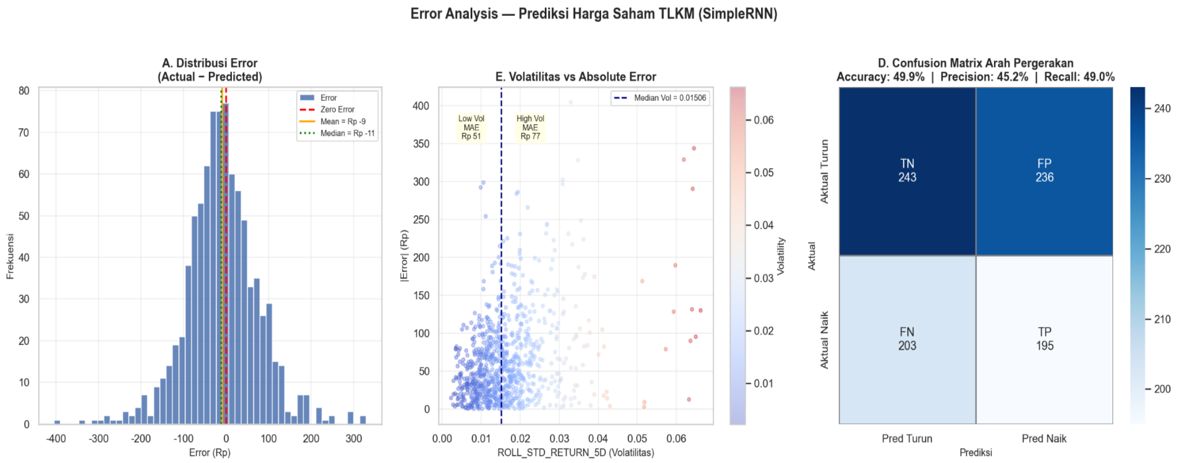

***Figure 4**. Histogram of prediction errors on the test set. The distribution is approximately symmetric and centered near zero, indicating minimal systematic bias.*

Examination of the ten largest absolute errors reveals a consistent pattern: the most extreme mispredictions (the largest being Rp -404 on 6 December 2022 and Rp +329 on 21 October 2025) occur predominantly during periods of sharp price discontinuity, such as post-announcement gaps or heightened macroeconomic uncertainty. This behavior is structurally inherent to any regression model trained exclusively on technical price features: the model cannot anticipate exogenous shocks that are unrepresented in the input feature set.

4.4. Directional Accuracy

A confusion matrix for binary directional prediction (next-day price rise versus fall) yielded an overall accuracy of 49.94%, comprising 243 true negatives and 195 true positives against 236 false positives and 203 false negatives. Precision for the 'rise' class was 45.24% and recall was 48.99%, with an F1-score of 47.04%. These results confirm that while the regression performance of the model is excellent, its directional prediction is essentially at chance level. This apparent paradox is structurally expected: the model optimizes MSE, which penalizes the magnitude of numerical deviations rather than their sign, and a prediction that is, say, Rp 50 below the actual price will have zero directional error yet perfect regression loss optimization. Systems designed specifically for directional trading signals would require a classification loss function and, likely, sentiment or order-flow features beyond the technical indicators employed here.

5. Discussion

The central finding of this study (mean MAPE of 2.00%, $R^2 = 0.9775$ across five independent evaluation runs) establishes SimpleRNN as a practically capable baseline architecture for next-day closing price regression of liquid Indonesian equities. Crucially, the consistency of results across runs (MAPE range: 1.96%–2.07%; R^2 range: 0.9765–0.9784) is scientifically significant: it demonstrates that the model's performance is not an artifact of a favorable random initialization but reflects a genuine and stable learned mapping between the engineered feature space and the next-day price target. This multi-run evaluation methodology is analogous to the replicated-trial design advocated by Bakri et al. (2025) in the context of educational data mining, where cross-run stability was used to validate the generalizability of machine learning classification performance, a principle equally applicable to regression tasks in the financial domain.

The MAPE of 2.00% achieved by this study compares favorably with the existing literature on Indonesian stock prediction. Budiharto (2021) reported LSTM MAPE values of approximately 2.27% on IHSG closing prices, while the hybrid GRU-LSTM framework of Haryono et al. (2024) achieved R^2 values around 0.97–0.98 for IDX blue-chip stocks including TLKM. The present study's SimpleRNN, equipped with systematic feature engineering, achieves R^2 of 0.9775, squarely within this competitive range, while maintaining greater architectural simplicity. Critically, the feature engineering strategy contributes meaningfully to this competitive performance: by encoding momentum (return lags), trend (moving averages), and volatility (rolling standard deviation) into the input tensor, the model receives a richer representation of market state than raw OHLCV data alone would provide, consistent with the findings of Dwiandiyanta et al. (2025).

A closer look reveals an important structural limitation: the model's near-random directional accuracy (49.94%) despite excellent regression performance highlights the fundamental distinction between price level prediction and price movement classification. From an investment utility perspective, directional accuracy is arguably more actionable than regression accuracy, a trader positioning for intraday moves requires knowledge of whether the price will be higher or lower tomorrow, not necessarily how much higher or lower. This performance gap is not unique to the present study; Hadi and Priyanto (2025) similarly reported that even LSTM and GRU models trained on TLKM data achieved directional accuracies only modestly above 50% when trained on technical features alone, suggesting that directional prediction of this asset requires fundamentally different inputs, sentiment analysis, institutional flow data, or fundamental indicators, rather than improvements to the technical architecture.

Contrary to expectations based on the vanishing gradient literature, the SimpleRNN's ten-step lookback window proved sufficient to capture the predictive structure of the data without requiring the gating mechanisms of LSTM or GRU. This finding is consistent with Sako et al. (2022), who observed that SimpleRNN remained competitive with more complex architectures for short-horizon forecasting tasks. The explanation is likely that the feature-engineered input already encodes the medium-term context (through MA-10 and RSI_SLOPE) that LSTM's cell state would otherwise need to learn implicitly, reducing the effective temporal horizon the recurrent mechanism must model. In other words, explicit feature engineering and architectural complexity are, to a degree, substitutable: a finding with implications for practitioners who must balance model performance against deployment complexity and inference latency.

The deployment architecture developed in this study addresses a gap between model performance and practical utility. The FastAPI-based REST endpoint enables real-time inference by fetching the latest OHLCV data from Yahoo Finance, constructing the ten-day feature window, applying the stored MinMaxScaler, and returning the predicted next-day closing price in JSON format. The SQLite prediction log enables retrospective accuracy assessment and concept-drift monitoring, both essential for maintaining system reliability over time. The candlestick dashboard with multi-range temporal view (1D, 1W, 1M, YTD, 1Y, 3Y) provides the contextual information investors need to interpret predictions in the broader market context. Together, these components constitute a decision-support prototype that is, in principle, deployable to cloud infrastructure (AWS SageMaker, GCP Vertex AI) with minimal additional engineering. The alignment of this deployment architecture with digital business practice reflects the broader trend toward AI-assisted investment tools in the Indonesian fintech ecosystem.

Several limitations of this study merit explicit acknowledgment. First, the model relies exclusively on technical price features and does not incorporate macroeconomic indicators, news sentiment, or corporate announcements, all of which demonstrably influence TLKM's price dynamics. Second, the evaluation is conducted on a single stock; generalizability to other IDX equities remains to be established. Third, the model is not compared against LSTM or GRU baselines under identical conditions, which would more precisely quantify the architectural trade-offs. Fourth, the directional classification performance is inadequate for generating buy/sell signals without further model development. Future research should address these gaps by extending the feature set with sentiment and fundamental indicators, implementing a multi-architecture comparative study on a broader set of IDX equities, and exploring attention-augmented recurrent models that may improve both regression accuracy and directional prediction simultaneously.

6. Conclusion

This study demonstrated that a SimpleRNN model, when combined with systematic technical feature engineering, achieves high and stable predictive accuracy for next-day closing price forecasting of PT Telkom Indonesia (TLKM.JK). Across five independent evaluation runs, the model attained a mean MAPE of 2.00%, MAE of Rp 67.50, RMSE of Rp 89.13, and R^2 of 0.9775, with cross-run variance in MAPE confined to approximately 0.11 percentage points, confirming that performance reflects a genuine and reproducible learned mapping rather than stochastic initialization effects. The error distribution exhibited a near-symmetric, near-zero-centered profile (skewness = 0.029), indicating the absence of systematic bias that would undermine prediction reliability over time.

A key finding is the structural divergence between regression and classification performance: while the model excels at predicting price levels (97.75% variance explained), its directional accuracy of 49.94% is essentially at chance level. This distinction has important implications for investment decision support: the system is appropriately positioned as a price forecasting aid rather than a trading signal generator, and should be interpreted in conjunction with fundamental analysis and qualitative market context. The deployment architecture, comprising a FastAPI REST endpoint, SQLite prediction history, and interactive candlestick dashboard, transforms the predictive model into an operational prototype aligned with the digital business objectives of accessible, data-driven investment support.

This research contributes to the growing body of literature on deep learning applications in Indonesian equity markets and underscores the value of feature engineering as a complement to architectural complexity in temporal prediction tasks. Future work should investigate hybrid architectures incorporating sentiment analysis and fundamental data, extend evaluation to a broader portfolio of IDX equities, and explore the deployment of attention-based recurrent models to improve both regression precision and directional prediction capability.

Acknowledgements: The authors thank the Department of Digital Business, Faculty of Economics and Business, Universitas Negeri Makassar for supporting this research.

Funding Statement: This research received no specific grant from any funding agency.

Conflicts of Interest: The authors declare that they have no conflicts of interest to report regarding the present study.

Contribution: Alam: Conceptualization, design, supervision, final approval. Hidayah: Data acquisition, data wrangling, feature engineering, writing. Ashif: Model development, training, deployment, visualization, writing. Bakri: Methodology, critical revision of manuscript, supervision. Qardawi: Editing/reviewing, writing.

References

Andriyanto, D., & Rianto, Y. (2022). Optimization of recurrent neural network in Indonesia stock exchange price prediction modeling. Journal of Information Systems, Informatics and Computing, 6(1), 1–10. https://doi.org/10.52362/jisicom.v6i1.744

Bakri, R., Alam, S., Astuti, N. P., & Bakhtiar, M. I. (2025). Optimizing machine learning models for graduation on time prediction: A comparative study with resampling and hyperparameter tuning. Jurnal Online Informatika (JOIN), 10(2), 270–285. https://doi.org/10.15575/join.v10i2.1590

Budiharto, W. (2021). Data science approach to stock prices forecasting in Indonesia during Covid-19 using Long Short-Term Memory (LSTM). Journal of Big Data, 8(1), 1–9. https://doi.org/10.1186/s40537-021-00430-0

Darnis, F., Rahman, A., & Ansori, Y. (2025). Akurasi model algoritma Recurrent Neural Network untuk prediksi saham syariah. TEKNOMATIKA, 15(1), 17–24. https://ojs.palcomtech.ac.id/index.php/teknomatika/article/view/690

Dwiandiyanta, B. Y., Hartanto, R., & Ferdiana, R. (2025). Harnessing deep learning and technical indicators for enhanced stock predictions of blue-chip stocks on the Indonesia Stock Exchange (IDX). Engineering, Technology & Applied Science Research, 15(1), 20348–20357. https://doi.org/10.48084/etasr.9850

Dwiandiyanta, B. Y., Hartanto, R., & Ferdiana, R. (2025). Optimization of stock predictions on Indonesia Stock Exchange: A new hybrid deep learning method. Engineering, Technology & Applied Science Research, 15(1), 19370–19379. https://doi.org/10.48084/etasr.9363

Fischer, T., & Krauss, C. (2018). Deep learning with long short-term memory networks for financial market predictions. European Journal of Operational Research, 270(2), 654–669. https://doi.org/10.1016/j.ejor.2017.11.054

Hadi, M. F., & Priyanto, D. (2025). Perbandingan kinerja RNN, LSTM, dan GRU dalam prediksi harga saham TLKM menggunakan deep learning. Seminar Nasional CORISINDO 2025, 265–270

Haryono, A. T., Sarno, R., & Sungkono, K. R. (2024). Stock price forecasting in Indonesia stock exchange using deep learning: A comparative study. International Journal of Electrical and Computer Engineering, 14(1), 861–869. https://doi.org/10.11591/ijece.v14i1.pp861-869

Khan, A. H., Shah, A., Ali, A., Shahid, R., Zahid, Z. U., Sharif, M. U., Jan, T., & Zafar, M. H. (2023). A performance comparison of machine learning models for stock market prediction with novel investment strategy. PLOS ONE, 18(9), e0286362. https://doi.org/10.1371/journal.pone.0286362

Pipin, S. J., Purba, R., & Kurniawan, H. (2023). Prediksi saham menggunakan Recurrent Neural Network (RNN-LSTM) dengan optimasi Adaptive Moment Estimation. Journal of Computer System and Informatics, 4(4), 806–815. https://doi.org/10.47065/josyc.v4i4.4014

Pratama, F. R., Santoso, B., & Kacung, S. (2025). Prediksi harga saham PT Telkom menggunakan metode CNN-LSTM. Journal of Information System Management (JOISM), 7(1), 66–70. https://doi.org/10.24076/joism.2025v7i1.2087

Ramadhani, G., Mahdiyah, U., & Wulanningrum, R. (2025). Prediksi harga saham batubara menggunakan Recurrent Neural Network (RNN). Prosiding Seminar Nasional Inovasi Teknologi (SEMNAS INOTEK), 9, 528–536.

Sako, K., Mpinda, B. N., & Rodrigues, P. C. (2022). Neural networks for financial time series forecasting. Entropy, 24(5), 657. https://doi.org/10.3390/e24050657

Saud, A. S., & Shakya, S. (2020). Analysis of look back period for stock price prediction with RNN variants: A case study on banking sector of NEPSE. Procedia Computer Science, 167, 788–798. https://doi.org/10.1016/j.procs.2020.03.419

Switrayana, I. N., Hammad, R., Irfan, P., Sujaka, T. T., & Nasri, M. H. (2025). Comparative analysis of stock price prediction using deep learning with data scaling method. JTIM: Jurnal Teknologi Informasi dan Multimedia, 7(1), 78–90. https://doi.org/10.35746/jtim.v7i1.650

Suyudi, M. A. D., Djamal, E. C., & Maspupah, A. (2019). Prediksi harga saham menggunakan metode Recurrent Neural Network. Seminar Nasional Aplikasi Teknologi Informasi (SNATi), A-33–A-38. Retrieved from https://journal.uii.ac.id/snati

Zhang, Y., Li, C., & Wang, J. (2024). Deep learning models for price forecasting of financial time series: A review of recent advancements: 2020–2022. WIREs Data Mining and Knowledge Discovery, 14(1), e1519. https://doi.org/10.1002/widm.1519