**NUANSA INFORMATIKA** ***p-ISSN :1858-3911*** **,** ***e-ISSN : 2614-5405***
Volume 20 Number 2 July 2026 https://journal.fkom.uniku.ac.id/ilkom

# Evaluating a Geopolitical-Regime-Aware Hybrid Ensemble for IHSG Forecasting

**Imam Prasodjo**
Universitas Al-Khairiyah, Cilegon, Banten, Indonesia
E-mail: iprasodjo@gmail.com

***Abstract***

*This study evaluates a geopolitical-regime-aware hybrid ensemble (GRAHE) for one-day-ahead forecasting of the Indonesia Composite Index (IHSG) using observations from 2015 to 2025. The framework combines long short-term memory, extreme gradient boosting, and random forest forecasts through a softmax gate conditioned on one-month-lagged geopolitical risk and rolling market volatility. Hyperparameters were selected with Bayesian optimization on the 2022 validation period, and the locked test used monthly expanding-window retraining from 2 January 2023 to 5 December 2025 (697 forecasts). GRAHE obtained a return RMSE of 0.009334, directional accuracy of 50.50%, and Theil's U of 1.001. It did not outperform the strongest benchmark: Ridge regression achieved an RMSE of 0.009316, while XGBoost produced the highest directional accuracy of 54.23%. The RMSE difference between Ridge and GRAHE was not statistically significant according to the Diebold–Mariano test ($p = 0.744$). In high-GPR observations, GRAHE recorded an RMSE of 0.009626, 0.30% above the best baseline. Removing GPR features increased ensemble RMSE slightly to 0.009350, but market-only XGBoost remained competitive at 0.009326. These findings do not support a strong claim that geopolitical conditioning improves daily IHSG point forecasts, although the framework provides a transparent basis for regime-dependent risk analysis and further robustness testing..*

***Keywords -*** *Geopolitical risk, hybrid ensemble, IHSG, LSTM, machine learning, stock-price forecasting, XGBoost.*

*Submission: July 30, 2026* *Accepted: July 31, 2026* *Published: July 31, 2026*
*DOI Article: https://doi.org/.....*

## 1. INTRODUCTION

Forecasting equity-market movements remains difficult because prices reflect nonlinear interactions among market microstructure, investor expectations, macroeconomic conditions, and external shocks. The problem is especially demanding in emerging markets, where sensitivity to global risk, exchange-rate movements, capital flows, and commodity-price shocks can produce abrupt regime changes. The IHSG provides an important case because it aggregates the performance of equities listed on the Indonesia Stock Exchange and is commonly used as a broad indicator of Indonesian market conditions.

Geopolitical uncertainty is a plausible source of such regime changes. Caldara and Iacoviello developed a news-based Geopolitical Risk (GPR) index and showed that elevated geopolitical risk is associated with weaker investment and employment and greater downside risk [1]. Subsequent research finds that geopolitical risk helps explain stock-market volatility in emerging economies and that the effects may be stronger in emerging than in developed markets [2]–[4]. These findings motivate the inclusion of geopolitical information in a forecasting model rather than treating market history as the only predictive source.

Existing machine-learning studies demonstrate that LSTM networks can learn long temporal dependencies, while XGBoost and random forests effectively model nonlinear interactions and heterogeneous predictors [5]–[8]. Hybrid designs often outperform individual learners because their inductive biases are complementary [9]-[12]. However, three gaps remain. First, many studies forecast price levels after random train–test splitting, creating a risk of look-ahead leakage and overstated accuracy. Second, geopolitical uncertainty is often examined econometrically but rarely integrated directly into an adaptive hybrid forecasting architecture. Third, static ensemble weights do not recognize that the relative performance of a sequential model and a tree-based model may change during geopolitical stress.

To address these gaps, this study proposes the Geopolitical-Regime-Aware Hybrid Ensemble (GRAHE). Its main contributions are:

- A hybrid architecture combining LSTM, XGBoost, and random forest predictions through out-of-fold constrained stacking.
- A regime-aware softmax gate that changes ensemble weights as a function of lagged GPR and realized market volatility.

***p-ISSN :1858-3911*** **,** ***e-ISSN : 2614-5405***
https://journal.fkom.uniku.ac.id/ilkom

- A leakage-resistant empirical protocol based on return forecasting, nested time-series validation, expanding-window testing, and publication-lag controls for GPR data.
- A stress-period evaluation that measures whether accuracy gains persist when geopolitical risk is in its upper historical quantiles.

The study addresses three research questions: RQ1: Does lagged geopolitical-risk information improve out-of-sample IHSG forecasting? RQ2: Does the proposed hybrid ensemble outperform its individual components and conventional benchmarks? RQ3: Are any gains concentrated in, or robust to, high-GPR market regimes?

## 2. RESEARCH METHOD

### *2.1.1. Sample and Data Sources*

The empirical implementation uses daily IHSG open, high, low, close, and volume observations from a public daily-IHSG dataset derived from Yahoo Finance, combined with the official monthly GPR dataset of Caldara and Iacoviello [1], [23], [26]. The downloaded IHSG series covers 2 January 2015–5 December 2025. After constructing 60-day rolling features and one-day-ahead targets, the usable sample begins on 30 March 2015 and contains 2,356 observations. Because the public market dataset ends on 5 December 2025, the present results are preliminary and require replication with complete official IDX data before submission.

To prevent publication-time leakage, the baseline model uses the previous month's GPR value for every trading day in the current month. A robustness specification may use a daily GPR series only when the release timestamp and data vintage are documented. Missing market observations are not interpolated across non-trading days. Monthly GPR observations are carried forward only after the lag rule has been applied.

*Table I. Candidate feature set. Every feature at date t must be computable using information available no later than t.*

| Feature group | Variables | Frequency | Status |
|---|---|---|---|
| Market prices | Open, High, Low, Close, Adjusted Close | Daily | Yahoo Finance-derived daily-IHSG dataset |
| Liquidity | Volume; log volume change; 5-, 20-day volume ratios | Daily | Yahoo Finance-derived and constructed |
| Returns | 1-, 2-, 5-, 10-, and 20-day log returns | Daily | Derived |
| Trend | SMA/EMA 5, 10, 20, 50; MACD; price-to-average ratios | Daily | Derived |
| Momentum | RSI-14; stochastic oscillator; rate of change | Daily | Derived |
| Risk | ATR-14; Parkinson range; rolling volatility 5, 10, 20, 60 | Daily | Derived |
| Geopolitical | Lagged GPR level; change; z-score; high-GPR dummy | Monthly→daily | Official Caldara–Iacoviello data |
| Interactions | GPR × volatility; GPR × USD/IDR change; GPR × oil return | Mixed | Optional robustness |
| Target | Next-day IHSG log return; reconstructed next-day close | Daily | Derived |

### *2.1.2. Target Definition and Transformation*

The primary target is the one-day-ahead log return:

$$r(t+1) = \ln[P(t+1) / P(t)] \quad (1)$$

Forecasting returns rather than the nonstationary price level reduces the risk that a model appears accurate simply because adjacent price levels are close. The predicted closing price is reconstructed as

P-hat(t+1) = P(t) exp[r-hat(t+1)]. Price-level metrics are reported for interpretability, while model selection is based primarily on return loss and directional performance.

### *2.1.3. Data Partitioning*

The chronological partition used in the empirical run is training through 31 December 2021 (1,433 usable observations), validation during 2022 (226 observations), and locked out-of-sample testing from 2 January 2023 through 5 December 2025 (697 observations). The training sample determines the GPR thresholds: Q75 = 105.201 and Q90 = 118.211. Model hyperparameters and ensemble coefficients were fixed using the validation period before the test forecasts were generated.

***p-ISSN :1858-3911* , *e-ISSN : 2614-5405***
https://journal.fkom.uniku.ac.id/ilkom

*2.2.1. PROPOSED GEOPOLITICAL-REGIME-AWARE HYBRID ENSEMBLE*

Figure 1 summarizes the proposed architecture. The model has three base learners and a regime-aware meta-layer. Base learners are trained on identical information sets where possible, but the LSTM receives ordered sequences while tree models receive the final feature vector and selected lag summaries.

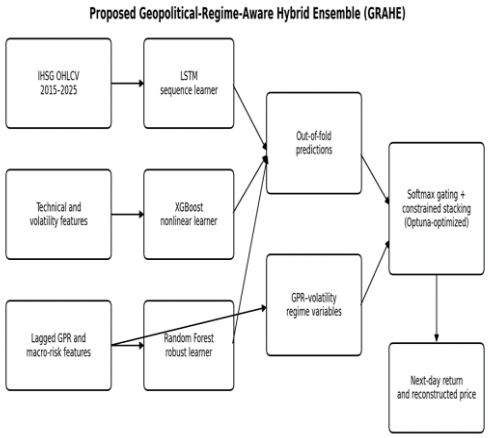



Fig. 1. Proposed GRAHE architecture.

*2.2.2. LSTM Base Learner*

The LSTM receives rolling sequences of standardized market and GPR features. Candidate sequence lengths are 10, 20, 30, and 60 trading days. The architecture contains one or two LSTM layers, dropout, and a linear output neuron for next-day return. Early stopping is based on validation loss. All scalers are fitted separately inside each training window.

*2.2.3. XGBoost and Random-Forest Base Learners*

XGBoost models nonlinear threshold effects and interactions through regularized boosted trees [5]. Random forest supplies a lower-variance bagging component [7]. Candidate hyperparameters include tree depth, number of estimators, learning rate, subsampling, column sampling, minimum child weight, regularization terms, and maximum feature fraction. Feature screening is performed only inside the training window.

*2.2.4. Regime-Aware Constrained Stacking*

Out-of-fold predictions from the three base learners form the meta-level inputs. Let m index the base model and z(t) contain the lagged GPR z-score and rolling volatility. Dynamic weights are defined by:

$$w(m,t) = \exp[a(m) + b(m)\cdot GPRz(t) + c(m)\cdot Vol(t)] / \Sigma j \exp[a(j) + b(j)\cdot GPRz(t) + c(j)\cdot Vol(t)] \quad (2)$$

The final return forecast is:

$$\text{r-hat}(t+1) = \Sigma m\ w(m,t)\ \text{r-hat}(m,t+1) \quad (3)$$

The softmax transformation guarantees nonnegative weights that sum to one. A static constrained stack is included as an ablation benchmark by setting all regime coefficients to zero. Coefficients and base-model hyperparameters are optimized only with training and validation data.

*2.2.5. Bayesian Hyperparameter Optimization*

Bayesian hyperparameter optimization was executed with Optuna's tree-structured Parzen estimator and random seed 20260730. The validation objective was one-day-ahead return RMSE on 2022. The studies used 10 Ridge trials, 12 SVR trials, 12 random-forest trials, 15 XGBoost trials, and 8 LSTM trials. The regime gate was subsequently estimated on the same locked validation predictions with an L2-regularized squared-error objective.

$$J(\Theta) = \text{RMSEvalidation}(\Theta) + \lambda||\Theta\text{gate}||^2 \quad (4)$$

The selected settings were: Ridge $\alpha$ = 230.635; random forest with 300 trees, depth 3, minimum leaf size 4, and maximum-feature fraction 0.675; XGBoost with 350 trees, depth 4, learning rate 0.0388, subsample 0.811, column sample 0.751, minimum child weight 7, L1 regularization 0.837, and L2 regularization 0.603; and a two-layer LSTM with sequence length 30, 12 hidden units, dropout 0.071, learning rate 0.00426, batch size 32, and early stopping. These values were not altered after observing the test results.

*2.3. EXPERIMENTAL DESIGN*

*2.3.1. Benchmark Models*

*Table II. Forecasting benchmarks.*

***p-ISSN :1858-3911*** **,** ***e-ISSN : 2614-5405***
https://journal.fkom.uniku.ac.id/ilkom

| Code | | Model | Purpose |
|---|---|---|---|
| B0 | | Random walk / zero-return | P-hat(t+1)=P(t) |
| B1 | | ARIMA | Linear time-series benchmark |
| B2 | | Ridge regression | Regularized linear feature benchmark |
| B3 | | Support vector regression | Nonlinear kernel benchmark |
| B4 | | Random forest | Standalone bagging learner |
| B5 | | XGBoost | Standalone boosting learner |
| B6 | | LSTM | Standalone sequential learner |
| B7 | | Static stack | Fixed nonnegative ensemble weights |
| P | | GRAHE | Dynamic GPR–volatility conditioned weights |

*2.3.2. Walk-Forward Evaluation*

An expanding-window evaluation simulated deployment. Ridge, SVR, random forest, XGBoost, and LSTM models were retrained at the start of each test month, producing 36 monthly forecast blocks. The training window expanded to include all observations available before each block. ARIMA(1,0,1) forecasts were updated sequentially with realized returns and refitted approximately quarterly. All feature scalers were fitted within the relevant training window, and no test outcome was used for hyperparameter selection or ensemble fitting.

*2.3.3. Evaluation Metrics and Statistical Tests*

Price-level accuracy is reported with RMSE, MAE, MAPE, and symmetric MAPE. Return forecasts are evaluated with RMSE, MAE, coefficient of determination, and directional accuracy. Theil's U compares the proposed model with the random-walk benchmark. Pairwise Diebold–Mariano tests assess whether loss differentials are statistically significant [14]. Directional predictability can additionally be evaluated with the Pesaran–Timmermann test [15].

Economic evaluation is secondary to statistical accuracy and should use a transparent long/flat or long/short rule based on the sign of predicted return. Transaction-cost sensitivity must be reported over a prespecified range, with turnover, annualized return, volatility, Sharpe ratio, and maximum drawdown. These results must not be described as investable performance without accounting for execution constraints and data availability.

*2.3.4. High-Geopolitical-Risk Regimes*

A high-GPR regime is defined before model evaluation, preferably as months where lagged GPR exceeds the 75th percentile estimated from the training sample. The threshold must not be recomputed using final-test data. Performance is reported separately for calm and high-GPR observations. Alternative thresholds at the 80th and 90th percentiles form robustness checks.

*2.3.5. Explainability and Diagnostics*

SHAP values were calculated for an XGBoost reference model trained through 2022 and evaluated on test-period features. The interpretation is model-specific and does not establish causality. The reported importance ranking is therefore used as a diagnostic complement to the forecast comparison rather than as evidence of a structural geopolitical transmission mechanism.

## 3. RESEARCH RESULT

The locked test contains 697 one-day-ahead forecasts from 2 January 2023 to 5 December 2025. Table III reports every benchmark without selective omission. Return RMSE is the primary ranking criterion; price-level metrics are included for interpretation. The zero-return benchmark has undefined directional accuracy and is therefore shown with a dash.

Table III. Locked out-of-sample forecasting performance, 2 January 2023–5 December 2025.

***p-ISSN :1858-3911*** **,** ***e-ISSN : 2614-5405***
https://journal.fkom.uniku.ac.id/ilkom

| Model | Price RMSE | Price MAE | Return RMSE | DA (%) | Theil U | DM p-value vs GRAHE |
|---|---|---|---|---|---|---|
| Random walk | 64.998 | 47.166 | 0.009328 | — | 1.000 | 0.574 |
| ARIMA | 64.925 | 47.058 | 0.009322 | 52.65 | 0.999 | 0.364 |
| Ridge | 64.957 | 47.279 | 0.009316 | 52.94 | 0.999 | 0.744 |
| SVR | 72.264 | 53.461 | 0.010285 | 50.07 | 1.103 | <0.001 |
| Random forest | 65.250 | 47.321 | 0.009367 | 53.08 | 1.004 | 0.111 |
| XGBoost | 64.976 | 47.091 | 0.009325 | 54.23 | 1.000 | 0.354 |
| LSTM | 64.999 | 47.145 | 0.009331 | 49.78 | 1.000 | 0.880 |
| Static stack | 65.250 | 47.321 | 0.009367 | 53.08 | 1.004 | 0.111 |
| GRAHE | 65.023 | 47.154 | 0.009334 | 50.50 | 1.001 | — |

### *3.1. Overall Forecasting Accuracy*

GRAHE achieved a return RMSE of 0.009334 and directional accuracy of 50.50%. The lowest RMSE was obtained by Ridge regression (0.009316), making GRAHE 0.19% worse on this criterion. Among the nonlinear standalone learners, XGBoost produced an RMSE of 0.009325 and the highest directional accuracy, 54.23%. GRAHE was 0.09% worse than XGBoost. The Diebold–Mariano comparisons did not identify a significant difference between GRAHE and Ridge ($p = 0.744$), XGBoost ($p = 0.354$), random forest ($p = 0.111$), LSTM ($p = 0.880$), ARIMA ($p = 0.364$), or the random-walk benchmark ($p = 0.574$). Only SVR was significantly inferior to GRAHE ($p < 0.001$).

All competitive models produced Theil's U values close to 1.000, indicating that improvements over the zero-return benchmark were economically small. Ridge achieved U = 0.999, XGBoost U = 1.000, and GRAHE U = 1.001. Consequently, the evidence does not justify a claim of clear predictive superiority. The pattern is consistent with the low signal-to-noise ratio of one-day equity-index returns and supports emphasizing statistical uncertainty rather than price-level visual fit.

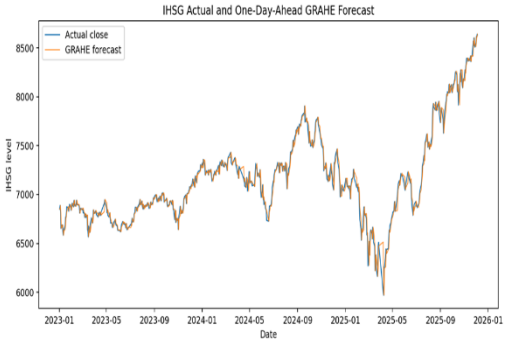



*Fig. 2. Actual IHSG closing level and reconstructed one-day-ahead GRAHE forecast, 2023–5 December 2025.*

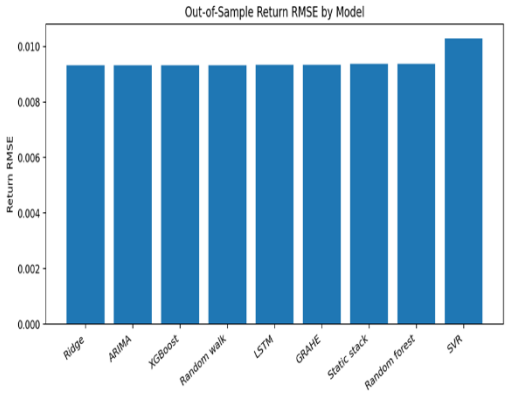



*Fig. 3. Locked-test return RMSE across benchmark models.*

*Table IV. Regime-specific performance.*

| Regime | Obs. | GRAHE RMSE | Best baseline RMSE | Improvement (%) | GRAHE DA (%) |
|---|---|---|---|---|---|
| Calm GPR | 88 | 0.006987 | 0.006951 (LSTM) | -0.53 | 54.55 |
| High GPR (≥ training Q75) | 609 | 0.009626 | 0.009597 (Ridge) | -0.30 | 49.92 |
| Extreme GPR (≥ training Q90) | 450 | 0.010106 | 0.010076 (Ridge) | -0.30 | 49.78 |

### *3.2. Performance under Geopolitical Stress*

GRAHE did not obtain an error advantage under geopolitical stress. In the 609 high-GPR observations, its RMSE was 0.009626

***p-ISSN :1858-3911*** **,** ***e-ISSN : 2614-5405***
https://journal.fkom.uniku.ac.id/ilkom

versus 0.009597 for Ridge, corresponding to a −0.30% improvement. In the 450 extreme-GPR observations, GRAHE recorded 0.010106 compared with 0.010076 for Ridge, again approximately 0.30% worse. In the smaller calm subsample (88 observations), LSTM was the best baseline and GRAHE was 0.53% worse. The test sample is heavily concentrated above the training Q75 threshold, so the calm/high comparison is unbalanced and should be reassessed with alternative rolling or annual thresholds.

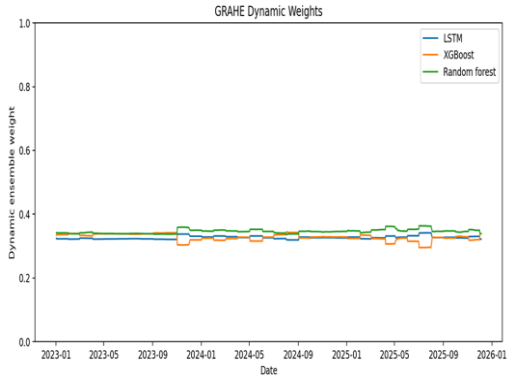



*Fig. 4. GRAHE softmax weights over the locked test period. Mean weights were 32.67% for LSTM, 32.75% for XGBoost, and 34.58% for random forest.*

### *3.3. Ablation Study*

The ablation evidence is mixed. Removing all GPR features and using a volatility-only gate increased RMSE from 0.009334 to 0.009350, a small 0.17% deterioration. Dynamic weighting improved on the static stack by 0.36%, but the static solution had collapsed to a 100% random-forest weight on the 2022 validation sample. More importantly, removing random forest improved RMSE to 0.009323, while market-only XGBoost achieved 0.009326 and 54.23% directional accuracy. Thus, neither the three-learner composition nor the geopolitical gate is consistently supported by the present evidence.

*Table V. Ablation analysis*

| Specification | GPR | Weights | Learners | Return RMSE | DA (%) |
|---|---|---|---|---|---|
| Full GRAHE | Yes | Dynamic GPR + volatility | LSTM+XGB+RF | 0.009334 | 50.50 |
| No GPR features | No | Dynamic volatility only | LSTM+XGB+RF | 0.009350 | 51.22 |
| Static stack | Yes | Static | LSTM+XGB+RF | 0.009367 | 53.08 |
| No random forest | Yes | Dynamic | LSTM+XGB | 0.009323 | 49.35 |
| No LSTM | Yes | Dynamic | XGB+RF | 0.009345 | 53.37 |
| Market-only XGBoost | No | — | XGB | 0.009326 | 54.23 |

### *3.4. Feature Importance and Model Behavior*

The XGBoost SHAP diagnostic ranked intraday range relative to close (range_rel), the 14-day stochastic oscillator (stoch14), the fourth lag of daily return (ret_lag4), and 10-day volatility (vol10) as the leading predictors. The strong regularization selected during validation produced near-zero SHAP contributions for many other variables, including most GPR features. This pattern indicates that short-horizon market-state variables dominated the fitted XGBoost model and that geopolitical variables did not emerge as major daily point-forecast drivers.

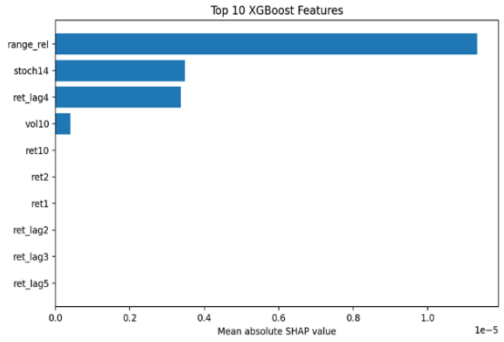



*Fig. 5. Mean absolute SHAP importance for the XGBoost reference model.*

## 4. DISCUSSION

Bagian Hypothesis 1 receives, at most, weak support: the full GPR specification marginally outperformed the no-GPR ensemble, but market-only XGBoost remained slightly better than GRAHE and the GPR features were not prominent in the SHAP ranking. Hypothesis 2 is not supported because GRAHE did not outperform Ridge or XGBoost on return RMSE. Hypothesis 3 is also not supported in its strong form: dynamic gating improved on the static stack, but not on the best simpler model, and its relative performance did not improve during high-GPR periods. The appropriate interpretation is therefore a transparent negative or mixed result rather than confirmation of all proposed hypotheses.

***p-ISSN :1858-3911* , *e-ISSN : 2614-5405***
https://journal.fkom.uniku.ac.id/ilkom

The framework may still be useful for risk monitoring because it exposes how model weights and errors change across uncertainty regimes. However, its value in the current run is diagnostic rather than predictive. The results suggest that future work should test lower-frequency targets, volatility forecasting, richer macro-financial variables, alternative publication-lag assumptions, and gate regularization selected with multiple validation regimes. No result in this study should be interpreted as an investment recommendation.

## 5. CONCLUSION

This study implemented GRAHE as a dynamic ensemble of LSTM, XGBoost, and random forest forecasts conditioned on lagged geopolitical risk and rolling volatility. In 697 locked forecasts from 2 January 2023 to 5 December 2025, GRAHE achieved return RMSE of 0.009334, directional accuracy of 50.50%, and Theil's U of 1.001. It did not outperform Ridge (RMSE 0.009316) or XGBoost (RMSE 0.009325; directional accuracy 54.23%), and the Diebold–Mariano differences against these models were not significant. GRAHE also remained approximately 0.30% worse than Ridge in high and extreme GPR regimes. Although removing GPR features slightly worsened the ensemble, a simpler two-model or market-only specification performed at least as well. The central conclusion is therefore that geopolitical-regime-aware weighting did not provide robust daily IHSG point-forecast gains in this implementation. The negative result is informative, but full official market data, alternative targets, repeated seeds, and additional robustness tests are required before journal submission.

## 6. SUGGESTION

Future research should evaluate GRAHE using complete and officially verified IHSG data over a longer out-of-sample period. The stability of the findings should be examined using repeated neural-network seeds, block-bootstrap confidence intervals, alternative walk-forward windows, and dynamically updated geopolitical-risk thresholds. Further studies should also investigate whether geopolitical information is more useful for forecasting weekly returns, cumulative returns, volatility, downside risk, or market direction than for predicting one-day-ahead returns. Country-specific GPR indicators, geopolitical threat and act subindices, exchange rates, commodity prices, foreign portfolio flows, and global financial-stress variables may be incorporated with appropriate publication-lag controls. Alternative ensemble mechanisms, including online weighting, regime-switching models, and regularized mixture-of-experts architectures, should be compared with simpler benchmarks. Finally, the economic relevance of the forecasts should be evaluated using realistic transaction costs, turnover, Sharpe ratios, and maximum drawdown without interpreting the results as investment recommendations.

## REFERENCES

[1] D. Caldara and M. Iacoviello, "Measuring geopolitical risk," American Economic Review, vol. 112, no. 4, pp. 1194–1225, 2022, doi: 10.1257/aer.20191823.

[2] A. A. Salisu, A. Ogbonna, L. Lasisi, and A. Olaniran, "Geopolitical risk and stock market volatility in emerging markets: A GARCH–MIDAS approach," The North American Journal of Economics and Finance, vol. 62, art. 101755, 2022, doi: 10.1016/j.najef.2022.101755.

[3] D. Das, M. Kannadhasan, and M. Bhattacharyya, "Do the emerging stock markets react to international economic policy uncertainty, geopolitical risk and financial stress alike?" The North American Journal of Economics and Finance, vol. 48, pp. 1–19, 2019, doi: 10.1016/j.najef.2019.01.008.

[4] M. Yang, J. Zhang, and H. Li, "Geopolitical risk and stock market volatility in emerging economies: Evidence from GARCH-MIDAS model," Discrete Dynamics in Nature and Society, art. 1159358, 2021, doi: 10.1155/2021/1159358.

[5] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining, 2016, pp. 785–794, doi: 10.1145/2939672.2939785.

[6] S. Hochreiter and J. Schmidhuber, "Long short-term memory," Neural Computation, vol. 9, no. 8, pp. 1735–1780, 1997, doi: 10.1162/neco.1997.9.8.1735.

[7] L. Breiman, "Random forests," Machine Learning, vol. 45, pp. 5–32, 2001, doi: 10.1023/A:1010933404324.

[8] M. Nabipour, P. Nayyeri, H. Jabani, S. S., and A. Mosavi, "Deep learning for stock market prediction," Entropy, vol. 22, no. 8, art. 840, 2020, doi: 10.3390/e22080840.

[9] P. H. Vuong, T. T. Dat, T. K. Mai, and P. H. Uyen, "Stock-price forecasting based on XGBoost and LSTM," Computer Systems Science and Engineering, vol. 40, no. 1, pp. 237–246, 2022, doi: 10.32604/CSSE.2022.017685.

[10] Z. Shi, Y. Hu, G. Mo, and J. Wu, "Attention-based CNN-LSTM and XGBoost hybrid model for stock prediction," arXiv:2204.02623, 2022, doi: 10.48550/arXiv.2204.02623.

[11] K. K. Yun, S. W. Yoon, and D. Won, "Prediction of stock price direction using a hybrid GA-XGBoost algorithm with a three-stage feature engineering process," Expert Systems with Applications, vol. 186, art. 115716, 2021, doi: 10.1016/j.eswa.2021.115716.

***p-ISSN :1858-3911* , *e-ISSN : 2614-5405***
https://journal.fkom.uniku.ac.id/ilkom

[12] H. Oukhouya, H. Kadiri, K. El Himdi, and R. Guerbaz, "Forecasting international stock market trends: XGBoost, LSTM, LSTM-XGBoost, and backtesting XGBoost models," Statistics, Optimization & Information Computing, vol. 12, no. 1, pp. 200–209, 2024, doi: 10.19139/soic-2310-5070-1822.

[13] T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama, "Optuna: A next-generation hyperparameter optimization framework," in Proc. 25th ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining, 2019, pp. 2623–2631, doi: 10.1145/3292500.3330701.

[14] F. X. Diebold and R. S. Mariano, "Comparing predictive accuracy," Journal of Business & Economic Statistics, vol. 13, no. 3, pp. 253–263, 1995, doi: 10.1080/07350015.1995.10524599.

[15] M. H. Pesaran and A. Timmermann, "A simple nonparametric test of predictive performance," Journal of Business & Economic Statistics, vol. 10, no. 4, pp. 461–465, 1992, doi: 10.1080/07350015.1992.10509922.

[16] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in Advances in Neural Information Processing Systems 30, 2017, pp. 4765–4774.

[17] R. J. Hyndman and A. B. Koehler, "Another look at measures of forecast accuracy," International Journal of Forecasting, vol. 22, no. 4, pp. 679–688, 2006, doi: 10.1016/j.ijforecast.2006.03.001.

[18] R. F. Engle, "Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation," Econometrica, vol. 50, no. 4, pp. 987–1007, 1982, doi: 10.2307/1912773.

[19] T. Bollerslev, "Generalized autoregressive conditional heteroskedasticity," Journal of Econometrics, vol. 31, no. 3, pp. 307–327, 1986, doi: 10.1016/0304-4076(86)90063-1.

[20] S. Deng, X. Huang, Y. Zhu, Z. Su, Z. Fu, and T. Shimada, "Stock index direction forecasting using an explainable eXtreme Gradient Boosting and investor sentiments," The North American Journal of Economics and Finance, vol. 64, art. 101848, 2023, doi: 10.1016/j.najef.2022.101848.

[21] S. Mehtab, J. Sen, and A. Dutta, "Stock price prediction using machine learning and LSTM-based deep learning models," in Proc. Int. Conf. Decision Aid Sciences and Application, 2020, pp. 447–453.

[22] Z. Shi, Y. Hu, G. Mo, and J. Wu, "Source code for attention-based CNN-LSTM and XGBoost stock prediction," associated repository, accessed for methodological comparison.

[23] M. Iacoviello, "Geopolitical Risk (GPR) Index: Data and replication materials," official project data page, accessed July 2026.

[24] Indonesia Stock Exchange, "Composite Stock Price Index and market statistics," official market-data portal, accessed July 2026.

[25] A. A. Salisu et al., "Geopolitical risk and stock market volatility: A global perspective," Finance Research Letters, vol. 53, art. 103620, 2023, doi: 10.1016/j.frl.2022.103620.

[26] G. Harrison, "Daily IHSG: Daily Jakarta Composite Index historical data derived from Yahoo Finance," Hugging Face Datasets, file ihsg_daily.csv, accessed July 30, 2026.