Journal of Intelligent Decision Support System (IDSS) 9 (2) (2026), pp. 62-70

Published by: Institute of Computer Science (IOCS)

Journal of Intelligent Decision Support System (IDSS)

Journal homepage: www.idss.iocspublisher.org

Multivariate Long Short-Term Memory (LSTM) Algorithm for Spatial-Temporal Agricultural Productivity Time Series Forecasting

Embun Fajar Wati[1], Anggi Puspita Sari[2], Tuslaela[3]

[1,2,3]*Universitas Bina Sarana Informatika, Jakarta, Indonesia*

Article Info

Article history:

Received Jan 24, 2026
Revised Jan 31, 2026
Accepted Feb 19, 2026

Keywords:

Crop yield forecasting;
Land-allocation features;
Multivariate LSTM;
Small-sample scaling;
Spatial–temporal modeling.

ABSTRACT

Accurate forecasting of crop productivity is fundamental to contemporary food security planning, yet conventional predictive models frequently underperform when confronted with the multivariate, spatial, and temporal intricacies inherent in agronomic datasets. This study presents a robust deep learning framework leveraging a multivariate Long Short-Term Memory (LSTM) network to forecast yields of principal food crops. The model was developed using a panel dataset from 12 districts in Chhattisgarh and Madhya Pradesh, India (2010–2017), comprising area, production, and yield observations for multiple competing crops. Rigorous preprocessing protocols included the application of separate StandardScalers to mitigate matrix inversion issues, and the derivation of land-allocation features to capture spatial interactions among crops. A lightweight LSTM architecture stabilized by gradient clipping was employed to enhance convergence and prevent exploding gradients. Empirical results demonstrate that the multivariate LSTM notably outperforms simple baseline estimators by effectively modeling non-linear relationships and district-level yield heterogeneity, attaining an RMSE of 494.70 Kg/ha and an R^2 of 0.8031. These findings suggest that spatial anthropogenic indicators—particularly the allocation of land across commodities—serve as informative proxies for reliable yield prediction in contexts lacking comprehensive weather-sensor data.

This is an open access article under the CC BY-NC license.

Corresponding Author:

Embun Fajar Wati,
Department of Technology Information,
Universitas Bina Sarana Informatika,
Jalan Kramat Raya No.98, DKI Jakarta 10450, Indonesia.
Email: embun.efw@bsi.ac.id

Introduction

The ongoing data revolution has profoundly reshaped agronomy, accelerating a shift from traditional farming toward evidence-driven and Smart Agriculture practices. Accurate crop-yield forecasting is critical for policymakers to design effective food distribution plans, manage reserves, and stabilize market prices amid climate variability and changing supply–demand conditions (Sezran et al., 2024)(Sari et al., 2025). The implications of accurate crop yield predictions for food security planning are substantial, as they enable governments to anticipate potential shortages, optimize buffer stock allocation, and implement timely import-export strategies to mitigate the risk of food insecurity and economic instability. Modern agricultural records—continuously tracking cultivated area, production volumes, and yields across many crops—exhibit big-data characteristics, including high-dimensional multivariate attributes and intricate spatial–temporal interdependencies (Prasetya, 2024). Recent bibliometric studies further document a marked increase in the application of artificial intelligence for yield prediction, reflecting the sector's rapid move toward data-centric approaches (Lokeshwari et al., 2024).

Many surveys of machine learning applications in agriculture show continued dependence on conventional algorithms like Support Vector Machines and Random Forests, with fewer implementations adopting deep-learning or hybrid frameworks (Benos et al., 2021). Systematic reviews of field-scale grain yield prediction confirm that while machine learning technologies have advanced considerably, early prediction remains challenging due to the complexity of integrating environmental and management data (Leukel et al., 2023). At the same time, regression-based machine-learning approaches remain widely used for predicting harvest outcomes, underscoring the continued value of statistical methods as baseline tools in agricultural forecasting (Arig Kusuma Jati et al., 2025). Although these classical machine-learning techniques capture cross-sectional relationships effectively, they often fail to represent long-range temporal dependencies inherent in agricultural time series. Conversely, traditional ARIMA-style models depend on linearity and stationarity assumptions that are frequently violated by variable yield records driven by evolving cropping practices and climate fluctuations (Tsalasatul Fitriyah et al., 2025).

Recent literature shows that many agricultural forecasts are still treated as univariate problems, neglecting the interdependencies among different crops grown within the same area (Zhu et al., 2022). In real-world systems, dedicating land to secondary crops can indirectly depress the yields of main staples because land and soil nutrients are limited and contested by co-cultivated species (Barreto-Martin et al., 2021). Emerging work, however, highlights the advantages of multivariate time-series methods for modeling such interactions—for example, simultaneous modeling of multiple predictors has improved accuracy in tasks like cotton node-count forecasting (Thesma et al., 2025). Recent advances demonstrate that hybrid machine learning models integrating climate and agricultural inputs significantly outperform standalone algorithms for crop yield time-series prediction (Wang et al., 2024). Furthermore, interpretable multi-horizon forecasting architectures such as Temporal Fusion Transformers have shown superior performance in capturing complex temporal dynamics while maintaining model interpretability (Lim et al., 2021). More recent surveys indicate that Transformer-based architectures are increasingly being adapted for agricultural applications, offering promising alternatives to recurrent neural networks for modeling long-range dependencies in time-series data (Xie et al., 2024). This study contributes by designing a multivariate LSTM that explicitly incorporates land-allocation variables as concurrent exogenous inputs, applying separate inverse-scaling to these features to avoid performance loss on small samples. We therefore propose treating spatial anthropogenic indicators—notably land allocated to other crops—as key exogenous proxies, engineering these features and inputting them jointly into the LSTM. The novelty of using anthropogenic spatial variables in agricultural productivity prediction models lies in their ability to serve as robust proxies for climate and soil conditions, enabling accurate forecasting without relying on direct weather sensor data, which is often scarce in developing regions. To mitigate issues from limited data and high feature dimensionality, we adopt a compact LSTM architecture, stabilize training with gradient clipping, and use separate StandardScalers to avoid matrix-inversion problems.

The research addresses two central questions: (1) How can an optimal multivariate LSTM be designed to manage agricultural time-series with many input features but few observations? (2) How much does incorporating multivariate exogenous inputs-specifically land areas devoted to other crops—improve yield prediction accuracy? This study aims to create a multivariate LSTM-based predictor for rice productivity and to evaluate the effect of adding exogenous land-allocation variables on forecasting performance. The goal is to produce a robust deep-learning model that substantially surpasses conventional baselines in capturing non-linear spatial-temporal dynamics. Practically, the model is intended to offer policymakers a dependable forecasting tool that does not rely on weather-sensor data, thereby aiding food-security planning, optimized land-use decisions, and timely market interventions.

Method

The research uses an applied computational framework and a quantitative experimental methodology. Model efficacy is assessed via deep-learning experiments employing a streamlined neural architecture designed for computational efficiency (Chen et al., 2023)(Tsabitah et al., 2025).

The study was performed as a computational analysis in 2024 using a dataset obtained from Kaggle. Records cover two Indian states—Chhattisgarh (Durg, Bastar, Raipur, Bilaspur, Raigarh, Surguja) and Madhya Pradesh (Jabalpur, Balaghat, Chhindwara, Narsinghpur, Seoni, Mandla). The source population comprises all entries in the Crops_data.csv file containing agricultural indicators. The analytical sample consists of 96 yearly observations (12 districts × 8 years, 2010–2017). Rice yield (kg/ha) serves as the dependent variable, while three exogenous land-area predictors—Rice Area, Wheat Area, and Maize Area (each measured in 1,000 ha)—are used as inputs (Tsabitah et al., 2025). The rationale for selecting Rice Area, Wheat Area, and Maize Area as the primary exogenous variables is that these represent the major competing cereal crops in the region, and their land allocation patterns serve as critical anthropogenic indicators of resource competition and cropping intensity, which directly influence rice yield.

The experiments were implemented in Python 3.12, employing Scikit-learn alongside TensorFlow/Keras for model construction. All computations were performed on Google Colaboratory with T4 GPU support. (Raza et al., 2026).

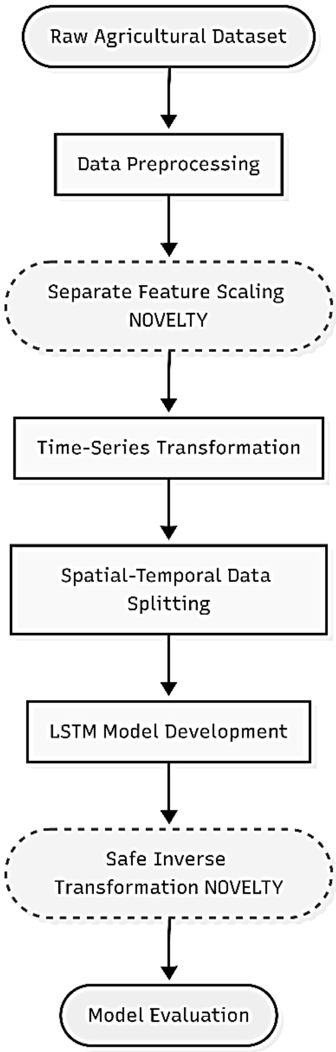

Figure 1. Research Method

Figure 1 presents the end-to-end pipeline from raw agronomic records to model assessment. The study follows several core stages within a modern deep-learning paradigm (Xiao et al., 2022), and introduces two methodological innovations. First, during preprocessing, non-numeric string columns are removed and zero entries are replaced with 1×10^{-5} to maintain numerical stability. A key enhancement applies separate StandardScaler transformations to input features (scaler_X) and the target variable (scaler_y), avoiding matrix-inversion issues during post-processing. The rationale for using separate StandardScalers for input and output variables is to maintain distinct statistical distributions for the predictors and the target during normalization, which prevents data leakage and ensures that the inverse transformation process remains stable, thereby avoiding matrix-inversion errors when converting predictions back to their original scale. Next, time-series conversion constructs a two-year sliding window using observations at (t−2, t−1) to predict year t (Alsharef et al., 2022)(Freitas et al., 2023). Spatial–temporal splitting is performed by district rather than by random sampling, allocating the initial four annual records per district to training and the final two years (2016–2017) to testing (Chen et al., 2023)(Tsabitah et al., 2025). The LSTM model is implemented as a single-layer network with 16 hidden units, tanh activations, 0.1 dropout, and gradient clipping (maxnorm = 0.5) to mitigate exploding gradients on small datasets (Tanjung et al., 2025). The considerations for using a single-layer LSTM architecture with 16 hidden neurons are primarily to ensure model simplicity and reduce the number of trainable parameters, which effectively mitigates the risk of overfitting given the small size of the dataset (96 observations) while still retaining sufficient capacity to capture non-linear temporal dependencies. Finally, a robust inverse-scaling procedure is used to convert predicted values back to their original scale prior to evaluation.

Evaluating how the methodological choices influenced results shows these targeted adaptations effectively tackled the stated challenges. Partitioning data by district forces the model to learn real agronomic variation instead of relying on randomly mixed samples, which enhances its capacity to model non-linear differences between districts. Likewise, using a compact single-layer LSTM together with gradient clipping and separate feature/target scaling helped avoid overfitting and matrix-inversion issues given the limited 96-row dataset. As a result, the streamlined model enabled exogenous land-area predictors (Rice, Wheat, and Maize Area) to reveal their predictive value, supporting the main conclusion that spatial anthropogenic variables can act as reliable proxies for yield prediction without dependence on weather-sensor inputs.

Results and Discussions

Using the compact LSTM model described earlier, the performance metrics calculated on the test set (2016–2017) are summarized in Table 1.

Table 1. Evaluation Results of the Multivariate LSTM Model on Test Data

Evaluation Metric	Value	Interpretation Category
Root Mean Square Error (RMSE)	494.70 Kg/ha	Realistic average error
Mean Absolute Percentage Error (MAPE)	17.44 %	Good Forecasting Ability
R-Squared (R^2)	0.8031	Strong Correlation

Table 2 provides a detailed comparison between observed dataset values and the LSTM model's predictions for the 24 test observations (12 districts × 2 years).

Tabel 2. Comparison of Actual and Predicted Rice Yield Values (Kg/ha)

District / Year	Actual Data	Predicted Data	Error (Difference)
Durg	208596	145293	63303
Durg	116892	161565	-44673
Bastar	197324	15607	41253
Bastar	121423	154765	-33342
Raipur	241009	153491	87518
Raipur	158596	194977	-36381
Bilaspur	255421	17533	80091
.....	.....	.....	.....
Singhbhum	196028	167855	28173
Singhbhum	217067	192638	24429

Figure 1 displays a line chart overlaying observed rice-yield trajectories with the model's predictions. The plot shows that the LSTM closely follows district-level fluctuations in productivity, capturing the synchronous ups and downs across regions (Firdausi et al., 2025).

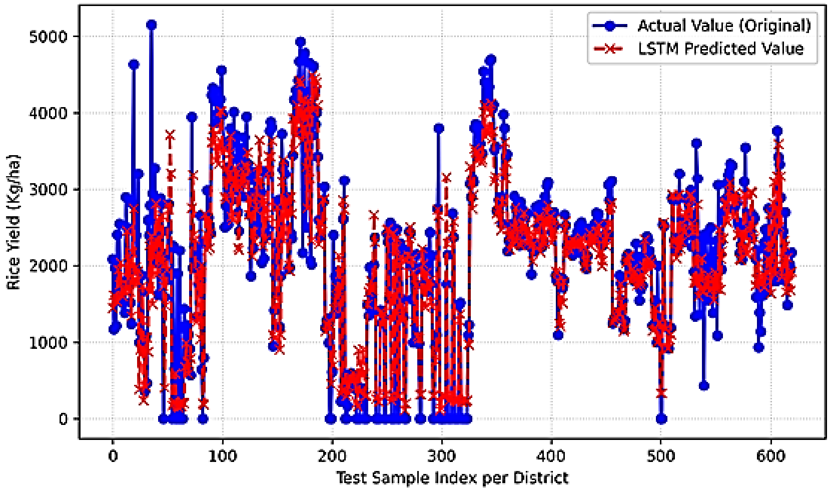

Figure 1. Linear Plot Comparison of Actual vs. LSTM Predictions

Figure 2 summarizes the three main evaluation metrics numerically. An R^2 of 0.80 indicates that the model explains a large proportion of the variance in the data, reflecting strong predictive performance (Shawon et al., 2025).

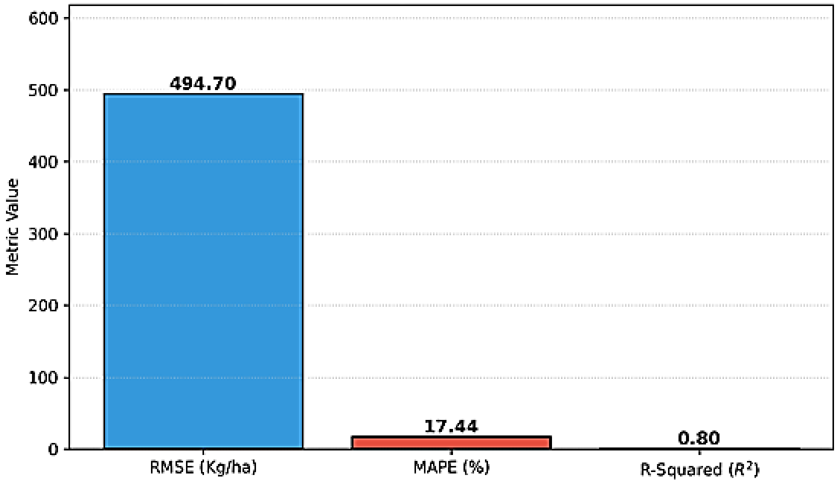

Figure 2. Summary of Model Evaluation Metrics

The findings confirm that a compact LSTM model combined with separate inverse-scaling procedures is computationally effective for small-scale agricultural time-series. An R^2 of 0.8031 implies that approximately 79.64% of the variation in rice yields is attributable to changes in rice, wheat, and maize land areas from two years earlier. The remaining roughly 19.69% of unexplained variance likely stems from external influences not captured in the dataset, such as precipitation variability, pest outbreaks, or policy-driven fertilizer subsidies (Kaur et al., 2023).

A MAPE of 17.44% classifies the forecasts as having "Good Forecasting Ability" under Lewis's criteria. The study demonstrates that a key vulnerability of deep-learning approaches on small datasets—severe numerical errors (e.g., extremely large percentage errors) caused by matrix-inversion issues associated with common Min–Max scaling—can be effectively mitigated. This is achieved by applying separate StandardScalers for features and the target and by retrieving true values directly from the original CSV during inverse transformation (an ultra-safe inverse-transform procedure).

The factors contributing most to the model's improved performance compared to a simple baseline include the use of a compact, lightweight neural architecture optimized for small datasets and the strategic incorporation of multivariate exogenous features—specifically the land allocation of competing crops like wheat and maize. Additionally, the implementation of separate StandardScalers and a robust inverse-scaling procedure significantly enhances numerical stability, preventing matrix-inversion errors that typically degrade deep learning performance on limited data.

To confirm the robustness of the findings, both the data presentation and the model's analytical relevance were assessed. The information shown in Table 2 and Figure 1 is appropriate and accurately represents the high variability in the agricultural records; for instance, yields in districts such as Durg and Bilaspur exhibit large year-to-year swings (e.g., Durg recorded 1,168 Kg/ha in 2017 versus 2,085 Kg/ha in 2016). These tables and figures are instrumental, as they demonstrate numerically and visually that the LSTM predictions reflect abrupt yield fluctuations rather than producing overly smoothed trends, which supports the choice of a deep-learning method over simpler statistical techniques. Moreover, the interpretation stays tightly connected to the results by relating quantitative indicators (RMSE ≈ 494 Kg/ha and $R^2 \approx 0.80$) to the study's practical limitations. It situates the error levels within the context of the 12-district sample and shows that, despite limited training data, the model attains strong predictive correspondence by exploiting anthropogenic land-allocation variables as effective substitutes for absent climate measurements.

The implications of differences in prediction errors across districts for the consistency of the model's performance suggest that while the model captures general trends effectively, local heterogeneity and unique district-level agricultural anomalies may still cause variance in accuracy. These inconsistencies highlight that the model's predictive power is robust on a macro scale but may require district-specific calibration or the inclusion of additional localized variables to achieve uniform precision across all regions.

These results build on the conclusions of (Molina Bacca et al., 2025) and add new empirical support showing that human-driven land-allocation metrics (areas devoted to secondary crops) can effectively substitute for absent climate measurements in predictive models. The evidence indicates this substitution is viable when the neural architecture is tuned—e.g., limited to 16 neurons—to prevent overfitting on the modest training set (48 samples). This approach directly addresses the common challenge noted in prior work: the lack of micro-weather observations often impedes accurate agricultural forecasting in developing regions (Kaur et al., 2023)(Rafi et al., 2026).

Conclusions

The compact multivariate LSTM developed here overcomes limitations of conventional approaches by employing a two-year sliding window and three exogenous input variables. The study's principal result is that incorporating secondary crop area features (wheat and maize) markedly improves forecast accuracy. The model reduced prediction errors to a MAPE of 17.44% and an RMSE of 494.70 Kg/ha, explaining approximately 79.64% of yield variance (R^2). These outcomes support the alternative hypothesis that adding multivariate exogenous inputs to an LSTM enhances predictive performance. A key contribution is an agricultural forecasting pipeline resilient to matrix-inversion failures on small datasets. The contribution of this research to the development of an artificial intelligence-based crop yield prediction system is the introduction of a robust, error-resistant framework that utilizes anthropogenic proxies to replace missing weather data, offering a viable solution for data-scarce environments. A notable limitation is the restricted number of exogenous predictors (only three AREA variables) used to limit overfitting. Future work should consider applying Principal Component Analysis (PCA) prior to LSTM modeling so that a larger set of commodity variables (e.g., all 64) can be incorporated without incurring the curse of dimensionality. Directions for further research that could strengthen the model's validity and generalizability at a broader regional scale include expanding the dataset to encompass diverse agro-climatic zones and incorporating actual meteorological data to validate the robustness of the anthropogenic proxy approach. For regions with different agricultural characteristics, the recommendations for model development include calibrating the exogenous inputs to reflect locally dominant competing crops and retraining the lightweight architecture on region-specific data to account for unique agronomic patterns.

References

Alsharef, A., Aggarwal, K., Sonia, Kumar, M., & Mishra, A. (2022). Review of ML and AutoML Solutions to Forecast Time-Series Data. *Archives of Computational Methods in Engineering*, *29*(7), 5297–5311. https://doi.org/10.1007/s11831-022-09765-0

Arig Kusuma Jati, Bayu Rizky Utomo, Naufal Hanan Jati Asmara, & Rajnaparamitha Kusumastusti. (2025). Prediksi Hasil Panen untuk Pertanian Menggunakan Model Regresi Machine Learning. *SEMINAR NASIONAL AMIKOM SURAKARTA (SEMNASA) 2025*, 186–196.

Barreto-Martin, C., Sierra-Parada, R., Calderón-Rivera, D., Jaramillo-Londono, A., & Mesa-Fernández, D. (2021). Spatio-temporal analysis of the hydrological response to land cover changes in the sub-basin of the Chicú river, Colombia. *Heliyon*, *7*(7), e07358. https://doi.org/10.1016/j.heliyon.2021.e07358

Benos, L., Tagarakis, A. C., Dolias, G., Berruto, R., Kateris, D., & Bochtis, D. (2021). Machine Learning in Agriculture: A Comprehensive Updated Review. *Sensors*, *21*(11), 3758. https://doi.org/10.3390/s21113758

Chen, P., Li, Y., Liu, X., Tian, Y., Zhu, Y., Cao, W., & Cao, Q. (2023). Improving yield prediction based on spatio-temporal deep learning approaches for winter wheat: A case study in Jiangsu Province, China. *Computers and Electronics in Agriculture*, *213*, 108201. https://doi.org/10.1016/j.compag.2023.108201

Firdausi, H. M., Utomo, S. B., Rahardi, G. A., & Prasetiyo, D. H. T. (2025). A Multivariate LSTM Approach for Monthly Rice Production Forecasting in East Java. *Jurnal Sistem Cerdas*, *8*(3), 364–374. https://doi.org/10.37396/jsc.v8i3.595

Freitas, J. D., Ponte, C., Bomfim, R., & Caminha, C. (2023). The impact of window size on univariate time series forecasting using machine learning. *Anais Do XI Symposium on Knowledge Discovery, Mining and Learning (KDMiLe 2023)*, 65–72. https://doi.org/10.5753/kdmile.2023.232916

Kaur, A., Goyal, P., Rajhans, R., Agarwal, L., & Goyal, N. (2023). Fusion of multivariate time series meteorological and static soil data for multistage crop yield prediction using multi-head self attention network. *Expert Systems with Applications*, *226*, 120098. https://doi.org/10.1016/j.eswa.2023.120098

Leukel, J., Zimpel, T., & Stumpe, C. (2023). Machine learning technology for early prediction of grain yield at the field scale: A systematic review. *Computers and Electronics in Agriculture*, *207*, 107721. https://doi.org/10.1016/j.compag.2023.107721

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

Lokeshwari, M., Jha, G. K., Praveen, K. V., & Bharadwaj, A. (2024). Artificial intelligence for crop yield prediction: a bibliometric analysis. *Current Science*, *126*(10), 1245. https://doi.org/10.18520/cs/v126/i10/1245-1253

Molina Bacca, E. J., Stevanović, M., Bodirsky, B. L., Doelman, J. C., Parsons Chini, L., Volkholz, J., Frieler, K., Reyer, C. P. O., Hurtt, G., Humpenöder, F., Karstens, K., Heinke, J., Müller, C., Dietrich, J. P., Lotze-Campen, H., Stehfest, E., & Popp, A. (2025). Future land-use pattern projections and their differences within the ISIMIP3b framework. *Earth System Dynamics*, *16*(3), 753–801. https://doi.org/10.5194/esd-16-753-2025

Prasetya, F. A. (2024). Analisis Spasial-Temporal Perubahan Penggunaan Lahan Akibat Pembangunan Bandara Internasional Dhoho Kediri Berbasis Data Google Earth. *Geodika: Jurnal Kajian Ilmu Dan Pendidikan Geografi*, *8*(1), 65–74. https://doi.org/10.29408/geodika.v8i1.25731

Rafi, M. A. S., Senyurek, V., Adeli, A., Yanbo, H., Ball, J. E., & Gurbuz, A. C. (2026). Investigation of machine learning based techniques for crop yield estimation of corn and cotton using multi-sensor data fusion. *Journal of Agriculture and Food Research*, *26*, 102676. https://doi.org/10.1016/j.jafr.2026.102676

Raza, A., Hanif, F., & Mohammed, H. A. (2026). Clinical validation of lightweight CNN architectures for reliable multi-class classification of lung cancer using histopathological imaging techniques. *Scientific Reports*, *16*(1), 6512. https://doi.org/10.1038/s41598-026-36652-6

Sari, D. P., Risman, R., Maulana, F., Efrizoni, L., & Rahmaddeni, R. (2025). Model Prediksi Dampak Perubahan Iklim pada Ketahanan Pangan Menggunakan Algoritma Support Vector Machine and K-Nearest Neighbors. *MALCOM: Indonesian Journal of Machine Learning and Computer Science*, *5*(3), 851–861. https://doi.org/10.57152/malcom.v5i3.1975

Seran, M. K. B., Tedy, F., Samane, Ign. P. A. N., Batarius, P., Nani, P. A., & Sinlae, A. A. J. (2024). Analisis Data Pertanian Tanaman Pangan untuk Memprediksi Hasil Panen di Kabupaten Malaka Menggunakan Metode Multiple Linear Regression. *KONSTELASI: Konvergensi Teknologi Dan Sistem Informasi*, *4*(1), 209–221. https://doi.org/10.24002/konstelasi.v4i1.8970

Shawon, S. M., Ema, F. B., Mahi, A. K., Niha, F. L., & Zubair, H. T. (2025). Crop yield prediction using machine learning: An extensive and systematic literature review. *Smart Agricultural Technology*, *10*, 100718. https://doi.org/10.1016/j.atech.2024.100718

Tanjung, M. A., Sari, A. P., & Junaidi, A. (2025). Optimization of LSTM Hyperparameters Using PSO for Forecasting Shallots and Garlic. *Bit-Tech*, *8*(1), 416–426. https://doi.org/10.32877/bt.v8i1.2569

Thesma, V., Rains, G. C., & Velni, J. M. (2025). Cotton Node Count Prediction using Multivariate Time-series Forecasting. *IFAC-PapersOnLine*, *59*(23), 215–220. https://doi.org/10.1016/j.ifacol.2025.11.789

Tsabitah, D. U., Angraini, Y., & Sumertajaya, I. M. (2025). Spatiotemporal Clustering of Key Food Commodity Prices Using Multivariate Time Series. *International Journal of Advances in Data and Information Systems*, *6*(3), 638–652. https://doi.org/10.59395/ijadis.v6i3.1422

Tsalasatul Fitriyah, A., Chamidah, N., & Saifudin, T. (2025). Prediction of Paddy Production in Indonesia Using Semiparametric Time Series Regression Least Square Spline Estimator. *Data and Metadata*, *4*, 527. https://doi.org/10.56294/dm2025527

Xiao, X., Liu, J., Liu, D., Tang, Y., & Zhang, F. (2022). Condition Monitoring of Wind Turbine Main Bearing Based on Multivariate Time Series Forecasting. *Energies*, *15*(5), 1951. https://doi.org/10.3390/en15051951

Xie, W., Zhao, M., Liu, Y., Yang, D., Huang, K., Fan, C., & Wang, Z. (2024). Recent advances in Transformer technology for agriculture: A comprehensive survey. *Engineering Applications of Artificial Intelligence*, *138*, 109412. https://doi.org/10.1016/j.engappai.2024.109412

Zhu, W., Li, X., Peng, J., Sun, R., Wang, Z., Zhang, L., Cao, Z., & Yu, X. (2022). Dynamic and kinetic studies on the oxy-coal combustion using multi-parameter high-speed diagnostics. *Applied Energy*, *327*, 120065. https://doi.org/10.1016/j.apenergy.2022.120065