

# Crop Yield Prediction using Computationally Efficient Machine Learning Models

**Sanjay SINGLA**

Post Doc Research Scholar, Multimedia University, Cyberjaya, Malaysia. Department of Computer Science & Engineering, University Institute of Engineering, Chandigarh University, Gharuan, Punjab, India  
[dr.ssinglacs@gmail.com](mailto:dr.ssinglacs@gmail.com)

**Chinnaiyan SENTHILPARI**

Centre for Advanced Devices and Systems, Faculty of Artificial Intelligence and Engineering (FAIE), Multimedia University, Cyberjaya, Malaysia  
[c.senthilpari@mmu.edu.my](mailto:c.senthilpari@mmu.edu.my) (corresponding author)

**Chu-Liang LEE**

Centre for Advanced Devices and Systems, Faculty of Artificial Intelligence and Engineering (FAIE), Multimedia University, Cyberjaya, Malaysia  
[cllee@mmu.edu.my](mailto:cllee@mmu.edu.my)

**Md Zaka Al Shahariar TANVIR**

Department of Computer Science & Engineering, University Institute of Engineering, Chandigarh University, Gharuan, Punjab, India  
[mzastanvir.2001@gmail.com](mailto:mzastanvir.2001@gmail.com)

**Abstract**—Within the domain of agricultural precision farming, dependable crop yield forecasting constitutes a fundamental requirement for augmenting productivity, optimizing resource allocation, and facilitating evidence-based decision-making processes. However, the development of a comprehensive representation that adequately captures the multifaceted interactions among climatic, environmental, and agronomic variables remains a considerable obstacle. This investigation presents a crop yield prediction framework utilizing a TensorFlow-implemented Deep Neural Network (DNN) applied to the Large Scale Agriculture Dataset (LSAD), encompassing 19,689 observations across 30 Indian states, 55 crop varieties, and multiple cultivation cycles. To ensure rigorous and impartial model assessment, standardized data preprocessing, hyperparameter tuning, and validation protocols were implemented for comparative analysis with four established machine learning techniques, namely CatBoost, LightGBM, XGBoost, and Random Forest. The empirical findings demonstrate that the proposed DNN delivers superior predictive capacity, attaining the minimum RMSE of 103.89 and the maximum  $R^2$  value of 0.9873, thereby demonstrating robust generalization performance across test datasets. Notwithstanding the elevated computational and inference expenses associated with CatBoost, the proposed DNN consistently produces more precise yield estimations. Moreover, the most influential feature ensemble identified through SHAP-based interpretability analysis comprises cultivated area, precipitation volume, temperature, fertilizer application rate, and pesticide application rate, with stability validation confirming the consistency of this feature configuration across repeated analyses. These findings substantiate that the proposed model constitutes a dependable and precise instrument for automated crop yield forecasting and holds considerable applicability in precision agriculture and agricultural policy formulation contexts.

**Index Terms**—Crop yield prediction, Deep Neural Network, TensorFlow, Machine learning, Precision agriculture, Explainable artificial intelligence (XAI), SHAP.

## I. INTRODUCTION

Predicting crop yield accurately is a key aspect in sustainable agriculture that can be useful for food security, efficient resource use, and farm management decision-making. Accurate predictions help farmers, agricultural policymakers, and other agricultural groups plan irrigation water use, fertilizer applications, pest control, and harvest schedules, as well as implement cost-effective, environmentally friendly crop production methods. Crop yield, however, is influenced by a complicated interaction between climatic conditions, soil characteristics, crop management, and geographical factors, which result in prediction being a challenging task.

In recent years, machine learning (ML) techniques have made substantial progress in modelling the complex, non-linear relationship in agricultural data, enabling crop yield prediction. Random Forest, XGBoost, LightGBM and CatBoost are among ensemble learning algorithms that have shown good predictive accuracy in various agricultural use cases. However, the majority of the current research mainly aims at achieving high prediction accuracy and very little is devoted to computational efficiency, inference latency, model size, and deployment feasibility. Moreover, different preprocessing pipelines and experimental setups are used in comparison, and the evaluation is not always fair.

The advantage of deep learning is that it can be used to automatically learn a hierarchy of feature representations from high-dimensional data, which is a potential solution. While deep neural networks have emerged as promising approaches for precision agriculture, extensive research that equitably benchmarks deep learning against the best machine learning models to assess the predictive capabilities,

efficiency, explainability, and robustness of each remains scarce. These assessments are necessary to formulate viable decision support systems that can be applicable to real-life agriculture.

Considering these limitations, this paper proposes a deep neural network (DNN) model using the TensorFlow (TF) framework for predicting crop yields from enormous agricultural data on various crops, seasons and geographical areas. To ensure equitable assessment, the proposed model undergoes systematic evaluation against Random Forest, XGBoost, LightGBM, and CatBoost models, which are subjected to identical preprocessing procedures and trained utilizing consistent hyperparameter configurations prior to performance evaluation. In addition to predictive precision, the evaluation framework incorporates supplementary performance indicators including computational efficiency, inference latency, model footprint, interpretability through SHAP analysis, and statistical stability to enable a holistic assessment of model efficacy suitable for practical deployment in operational environments.

The principal conclusions derived from this investigation are presented as follows:

- A deep neural network architecture implemented via TensorFlow is proposed as a viable approach for forecasting crop yield through the integration of diverse agricultural datasets.
- Comparative analysis of all models, encompassing Random Forest, XGBoost, LightGBM, and CatBoost, is conducted under uniform experimental conditions to facilitate fair evaluation.
- To measure the practical deployment performance, the predictive accuracy, computational efficiency, inference time, and model size are evaluated in combination.
- Explainability and feature importance analysis using SHAP enhances models' interpretability.

## II. LITERATURE REVIEW

Machine learning (ML) has revolutionized crop yield prediction as it has performed well in modelling complex agricultural systems, consistently outscoring traditional statistical and process-based crop models. Initial research confirmed ensemble learning techniques, such as Random Forest, as effective predictors in different crops and locations [1]. Later systematic reviews found that the methods most commonly adopted were Random Forest, Neural Networks, and Linear Regression, with climatic variables like temperature and rainfall as well as soil properties being the most important predictors [2]. The more recent reviews indicate a shift towards using more complex ensemble and hybrid learning structures for these tasks, such as gradient boosting, stacking, and deep learning-based techniques like CNN and LSTM, which more fully exploit the spatial and temporal nature of agricultural data [3]–[5]. Notwithstanding, previous reviews invariably focus on prediction accuracy but offer scant consideration of computational efficiency, memory requirements, usage of energy and inference latency, whose deployment aspects remain largely overlooked.

The most recent ML algorithms are classed into gradient-boosting algorithms, which are used for crop yield-based prediction due to effective prediction accuracy, proper scalability and computational efficiency. LightGBM, XGBoost, and CatBoost have proven to be effective models with tree-based learning, feature optimisation, and the ability

to deal with heterogeneous data on a range of agricultural datasets, with their remarkable performance in numerous fields showcasing the remarkable ability to make accurate predictions [6]–[11]. Recent works also apply explainable artificial intelligence methods like SHAP and LIME to make the models more readable, and stack and hybrid ensemble approaches to boost prediction robustness and generalisation ability. The developments are further evidence of the ongoing high competitiveness of boosting-based models for agricultural forecasting. But the majority of model evaluation in the model space is still on accuracy, while other aspects of the model, like computational cost, memory requirement and real-time inference performance, have not been considered so much in practical applications. With AI's environmental and computational impacts, green AI and energy-based efficient machine learning models have become more top of mind. As a result, the field of Green AI and energy-efficient machine learning has become more popular due to the rising awareness of the environmental and computational footprint of AI.

A recent survey differentiates between leveraging AI in a way that supports sustainability and developing "efficient AI" systems; model pruning, quantisation, knowledge distillation and low-precision computation are the main approaches cited to mitigate computational costs while maintaining good predictive performance in the context of both approaches [12]–[16]. Empirical results show that these optimisation methods can significantly decrease model size, energy consumption, and inference time without compromising accuracy, and across various learning architectures. The results offer useful insights in the design of efficient models but do not yet fully describe their use within crop yield prediction, nor are full evaluations based on both predictive performance and efficiency very common.

The TinyML applications and edge computing of AI are also becoming more popular in the field of smart agriculture, thanks to their energy efficiency. Lightweight machine learning models deployed on embedded devices have shown the potential for doing agricultural decision-making in the field without requiring a large amount of power to make the predictions, still achieving an acceptable amount of prediction performance [17, 18]. Other technologies that are on the rise in the field of agricultural sensing, including federated learning, edge intelligence, and Internet of Things (IoT), are also seen as promising solutions to these challenges of distributed and agile farming systems and cloud dependency [19]–[20]. Most of the existing applications, however, apply to classification tasks (crop recommendations) or other irrigation management tasks and do not work for precision yield forecasting.

In general, the current literature either focuses on the increase of prediction accuracy by developing sophisticated ensemble and deep learning systems or computational efficiency by optimizing AI systems with Green AI and model compression. There is limited literature that considers both accuracy, energy consumption, memory requirements, and inference latency in one crop yield prediction model. This study will fill this gap by comparing LightGBM, XGBoost, and a compressed model that can be deployed using TensorFlow Lite and a dashboard using Streamlit.

## III. METHODOLOGY

### A. Overview

This study introduces a computational efficient framework for the crop yield prediction with four stages

namely data preprocessing, model development, model performance evaluation and model explainability as shown in Fig. 1. Under a common experimental protocol, four state of the art machine learning models are trained with the agricultural dataset and a feature representation (structured representation) is used before the deep neural network (DNN) is trained. We use the typical regression measures to evaluate model performance, and model size and inference time to assess computational efficiency. Model interpretation is done using Shapley Additive exPlanations (SHAP). The proposed framework integrates aspects of prediction performance, computational efficiency, and model explainability in a common experimental framework, as opposed to traditional crop yield prediction studies which focus mostly on predictive accuracy.




Fig. 1. Overall framework of the proposed crop yield prediction methodology

### B. Data Preprocessing

Agricultural data was reviewed to identify missing, inconsistent, and extraneous data. The Production feature was dropped before training to avoid leakage to the targets. Continuous variables underwent standardization via StandardScaler, whilst categorical variables including Crop, Season, and State were processed through one-hot encoding. The resulting feature representation comprised 99 dimensions, with this consistent feature set utilized across all

analytical models. Subsequently, the preprocessed dataset was partitioned into distinct training, validation, and test subsets to facilitate an impartial assessment of model performance and generalization capacity. Detailed partitioning is given in the Experimental Setup.

### C. Baseline Machine Learning Models

A selection of four prominent regression algorithms served as comparative baseline models: RF, XGB, LGBM, and CatB. RF constitutes an ensemble methodology grounded in bagging principles, whereas XGB and LGBM represent gradient boosting frameworks that prioritize both predictive accuracy and computational efficiency optimization. CatB functions as a specialized machine learning algorithm designed for effective categorical variable processing, demonstrating substantial predictive performance. The combination of these models serves as an overall reference for evaluating the proposed DNN fairly. These models provide an overall reference for the proposed DNN to be fairly evaluated. The data sets, data preprocessing and evaluation criteria were all the same for all models so as to make a fair comparison.

### D. Proposed Deep Neural Network

The model proposed is a lightweight fully connected DNN for the prediction and computational efficiency trade-off. The neural network architecture comprises an input layer that processes 99 encoded features, succeeded by a sequence of hidden layers containing 512, 256, 128, and 64 neurons respectively. Batch Normalization and Dropout regularization techniques were applied following the initial pair of hidden layers to enhance training stability and mitigate overfitting. The network employs ReLU activation functions throughout the hidden layers. Crop yield estimation is formulated as a continuous regression task, implemented through a single output neuron with linear activation.

$$f(x) = \max(0, x) \quad (1)$$

This hierarchical reduction of hidden layer dimensionality allows to learn hierarchical feature abstraction without increasing the model complexity. Batch normalization helps to stabilize optimization, while dropout helps to enhance generalization. The lightweight architecture also makes it easier to implement in agricultural decision support systems with limited resources.

### E. Model Training

The proposed DNN underwent optimization through the Adam optimizer, with MSE serving as the loss function to be minimized, utilizing TensorFlow for implementation.

$$MSE = \frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2 \quad (2)$$

where  $y_i$  and  $\hat{y}_i$  represent the observed and predicted crop yields. Early stopping with the validation loss parameter was used to improve model generalization and mitigate overfitting tendencies. In the Experimental Setup section, hyperparameter values, software versions, and hardware specifications are provided. To ensure the reproducibility and a fair comparison of models, all experiments were run under the same preprocessing and evaluation protocol.

### F. Performance Evaluation

The efficacy of the model was subsequently assessed through three distinct metrics: RMSE, MAE, and  $R^2$ .

$$RMSE = \sqrt{\frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2} \quad (3)$$

$$MAE = \frac{1}{N} \sum_{i=1}^N |y_i - \hat{y}_i| \quad (4)$$

$$R^2 = 1 - \frac{\sum_{i=1}^N (y_i - \hat{y}_i)^2}{\sum_{i=1}^N (y_i - \bar{y})^2} \quad (5)$$

Beyond predictive performance, model size and model inference time were also measured to assess the computational efficiency and deployment suitability of the models.

### G. Model Explainability

The importance of each feature in the decisions made by the model was measured using SHAP. SHAP summary plots and feature importance rankings were used to explain the behaviour of the model at the global level and explanations at the local level were used to understand the prediction of the model for an individual observation. The significance of the features created from the Random Forest, XGBoost, Light GBM and CatBoost models was also compared and contrasted to identify common features across the learning models. The explainability analysis confirms the model behaviour, also allowing to point out the parameters involved in the agricultural production and to give an insight of the model results in each specific case. All these methodological components constitute an easily reproducible, transparent and computation-efficient method for an application-oriented crop yield prediction.

## IV. EXPERIMENTAL SETUP

The proposed framework was tested using a cleaned dataset of crop yield, which has 55 crop types, 6 agricultural seasons, and 30 states in India with 19,689 samples. The data was preprocessed first with standard procedures such as feature scaling and one-hot encoding, and then randomly split into training (80%), Validation (10%), and testing (10%). This proposed DNN was then evaluated with the same experimental setup as Random Forest, XGBoost, and LightGBM and CatBoost. All experiments were run on Kaggle on an NVIDIA Tesla T4 GPU and model performance was measured in terms of RMSE, MAE,  $R^2$ , inference time, model size, and interpretability using SHAP. The whole experimental setup is summarized in Table I.

TABLE I  
EXPERIMENTAL CONFIGURATION

| Category        | Configuration              |
|-----------------|----------------------------|
| Dataset         | Cleaned Crop Yield Dataset |
| Total Samples   | 19,689                     |
| Time Period     | 1997–2020                  |
| Crop Categories | 55                         |

|                               |                                                                       |
|-------------------------------|-----------------------------------------------------------------------|
| States                        | 30                                                                    |
| Agricultural Seasons          | 6                                                                     |
| Target Variable               | Yield                                                                 |
| Removed Feature               | Production (to prevent data leakage)                                  |
| Numerical Preprocessing       | Median Imputation + StandardScaler                                    |
| Categorical Preprocessing     | Most Frequent Imputation + One-Hot Encoding                           |
| Encoded Feature Dimension     | 99                                                                    |
| Data Partition                | Training: 80% (15,751), Validation: 10% (1,969), Testing: 10% (1,969) |
| Validation Strategy           | Separate Validation Set                                               |
| Baseline Models               | Random Forest, XGBoost, LightGBM, CatBoost                            |
| Proposed Model                | TensorFlow-based Deep Neural Network                                  |
| Network Architecture          | 512–256–128–64 Fully Connected                                        |
| Activation Function of Layers | ReLU                                                                  |
| Output-based Layer            | Linear (Regression)                                                   |
| Optimizer                     | Adam-Optimizer(Learning Rate = 0.001)                                 |
| Loss-Function                 | Mean Squared Error (MSE)                                              |
| Batch and it's Size           | 256                                                                   |
| Maximum attempted Epochs      | 200                                                                   |
| Callbacks                     | EarlyStopping, ReduceLROnPlateau, ModelCheckpoint                     |
| Random Seed                   | 42                                                                    |
| Programming Language          | Python 3.12                                                           |
| Development Libraries         | TensorFlow, Scikit-learn, XGBoost, LightGBM, CatBoost, SHAP, Optuna   |
| Hyperparameter Optimization   | Optuna (20 Trials per Tree-Based Model)                               |
| Computing Platform            | Kaggle                                                                |
| Hardware                      | NVIDIA Tesla T4 GPU (16 GB), 4 CPU Cores, 31.35 GB RAM                |
| Evaluation Metrics            | RMSE, MAE, $R^2$ , Inference Time, Model Size                         |
| Explainability Method         | SHAP (Summary Plot and Feature Importance)                            |

## V. RESULTS AND DISCUSSION

A comparative analysis was conducted utilizing four widely-adopted machine learning algorithms like Random Forest, XGBoost, LightGBM, and CatBoost & alongside the proposed Deep Neural Network (DNN) model. To ensure methodological rigor and fair evaluation, all models were subjected to identical experimental parameters and conditions during the training and testing phases. The performance assessment encompassed multiple dimensions, including prediction accuracy, computational efficiency, interpretability of results, and statistical stability as key evaluation criteria.

### A. Overall Performance Comparison




Fig. 2. Performance comparison.

A brief summary of the overall prediction performance of the models evaluated is provided in Table II and Fig. 2. The proposed TensorFlow DNN was found to have the lowest RMSE (103.89) and the highest  $R^2$  (0.9873), which confirmed that the proposed DNN model has not only the best predictive performance but also a good generalization ability. This excellent performance is due to the fact that it is able to learn nonlinear relationships between climatic, temporal, and agricultural variables in a hierarchical representation of features. CatBoost showed the best local predictive accuracy with the lowest MAE (10.18). XGBoost and LightGBM were also competitive. In contrast, Random Forest showed significantly higher prediction errors, indicating that it did not effectively capture the intricate patterns within the data set. Overall, the results show the effectiveness of deep feature learning in predicting the crop yield accurately and reliably.

TABLE II  
PREDICTIVE PERFORMANCE COMPARISON

| Model                   | RMSE ↓ | MAE ↓ | R² ↑   |
|-------------------------|--------|-------|--------|
| Proposed TensorFlow DNN | 103.89 | 17.19 | 0.9873 |
| CatBoost                | 121.91 | 10.18 | 0.9825 |
| XGBoost                 | 150.53 | 12.39 | 0.9733 |
| LightGBM                | 151.86 | 13.53 | 0.9728 |
| Random Forest           | 283.72 | 30.55 | 0.9051 |

### B. Computational Efficiency

Results of the computational performance of the models evaluated are shown in Table III and in Fig. 3. In terms of inference time, CatBoost outperformed with 0.014 s, and model size, CatBoost also showed the smallest size at 3.57 MB, so it is very feasible for real-time deployment. The proposed TensorFlow DNN had a multi-layer structure, which needed a relatively higher amount of computation resources, but it was still efficient in terms of computation and achieved the highest prediction accuracy. Random Forest, on the other hand, had the most complicated model and was not easily implementable on resource-limited systems. In short, these results show that there is a trade-off between computational efficiency and prediction accuracy.




Fig. 3. Accuracy vs. speed trade-off.

TABLE III  
COMPUTATIONAL PERFORMANCE COMPARISON

| Model                   | Inference Time (s) | Model Size (MB) |
|-------------------------|--------------------|-----------------|
| CatBoost                | 0.014              | 3.57            |
| XGBoost                 | 0.023              | 4.38            |
| LightGBM                | 0.077              | 5.31            |
| Proposed TensorFlow DNN | 0.197              | 4.38            |
| Random Forest           | 0.230              | 127.27          |

### C. Prediction Performance Analysis

The performance for prediction of the proposed model is shown in Fig. 4 and Fig. 5. The results of the actual and predicted yields for the evaluation set are in good agreement, which shows the high accuracy of the proposed DNN model. Furthermore, there is no apparent systematic pattern with respect to prediction errors around the zero value on the residual plot, indicating that there is not a large prediction bias. The result of this study also shows that the presented framework is reliable, stable, accurate and robust in predicting the crop yield for different agricultural conditions.

### D. Explainability Analysis

The results of the explainability analysis (see Fig. 6 and Fig. 7) are consistent and indicate that the most important variables are cultivated area, crop category, fertilizer use, pesticide use, annual rain and temperature-related variables. In the SHAP summary, it can be seen that the model prediction is influenced by the features. In the feature importance analysis, it can be seen that the feature ranks are consistent across different machine learning models. This high agreement level indicates that the proposed framework is not tailored to the datasets but rather to meaningful features in the field of agriculture, which makes the proposed model more transparent, reliable and applicable to the prediction of crop yields and agricultural decision support.




Fig. 4. Actual vs. predicted crop yield




Fig. 5. Residual Plot




Fig. 6. Feature Importance.




Fig. 7. SHAP summary

### E. Statistical Robustness

The results of statistical robustness, presented in Fig. 8 and Fig. 9, prove the reliability and stability of the proposed framework. The estimated prediction variability from repeated resampling experiments is relatively low, as evidenced by the narrow width of the corresponding bootstrap confidence interval, and confirms the strength and consistent generalization ability of the proposed model.




Fig. 8. Bootstrap Confidence Interval.

The correlation heatmap also offers insights into the interrelationships between the input variables, identifying meaningful interactions and suggesting that there is no serious multicollinearity that would negatively impact the prediction outputs. The results confirm the stability of the benchmark dataset and its suitability in representing various crops, environments, and agricultural activities. The general results of the statistical analyses show the reliability, robustness, and repeatability of the proposed crop yield prediction framework.




Fig. 9. Correlation Heatmap.

### F. Discussion

The empirical findings demonstrate a fundamental tension between predictive precision and computational resource utilization. TensorFlow DNN exhibited superior capability in discerning intricate nonlinear patterns, thereby achieving the most robust predictive outcomes; conversely, CatBoost demonstrated the most expeditious inference performance among all examined models and generated the most compact model architecture, rendering it particularly suitable for deployment scenarios with constrained computational resources. SHAP-based interpretation analysis revealed that the proposed framework incorporates agronomically pertinent features, thereby enhancing the conceptual transparency of the methodology. The investigation concludes that the recommended DNN architecture represents an optimal equilibrium among predictive accuracy, computational efficiency, and practical applicability for the purpose of crop yield forecasting within the scope of this research

### VI. CONCLUSION

In this study, a large-scale agricultural dataset was used, with multiple crops, seasons, and geographical regions, to predict crop yield using a TensorFlow-based Deep Neural Network (DNN). The suggested model was thoroughly tested and compared against the Random Forest, XGBoost, Light-GBM, and CatBoost models with the same experimental setup. Our proposed DNN model was found to be the best overall predictive model as it yielded the lowest RMSE and the highest R2 when compared with all the models considered. The proposed DNN gave the most accurate and reliable yield predictions, while CatBoost was faster to infer and the cost of computation was lower. Additionally, explainability via SHAP showed that the model used meaningful variables for agronomists, and robustness analysis showed that the model was robust across multiple repetitions. Future work aims at integrating multi-modal data sources, such as satellite imagery, soil information, and real-time weather observations, and transformer-based architectures and uncertainty-aware learning to further enhance the prediction accuracy and deployment in intelligent precision agriculture systems.

### REFERENCES

[1] J. H. Jeong, J. P. Resop, N. D. Mueller, D. H. Fleisher, K. Yun, E. E. Butler, *et al.*, "Random forests for global and regional crop yield predictions," *PLoS ONE*, vol. 11, no. 6, p. e0156571, Jun. 2016, doi: 10.1371/journal.pone.0156571.

[2] T. van Klompenburg, A. Kassahun, and C. Catal, "Crop yield prediction using machine learning: A systematic literature review," *Comput. Electron. Agric.*, vol. 177, p. 105709, Oct. 2020, doi: 10.1016/j.compag.2020.105709.

[3] M. A. Javed and M. A. Azmi Murad, "Crop yield prediction in agriculture: A comprehensive review of machine learning and deep learning approaches, with insights for future research and sustainability," *Heliyon*, vol. 10, no. 24, p. e40836, Dec. 2024, doi: 10.1016/j.heliyon.2024.e40836.

[4] S. M. Shawon, F. B. Ema, A. K. Mahi, F. L. Niha, and H. T. Zubair, "Crop yield prediction using machine learning: An extensive and systematic literature review," *Smart Agric. Technol.*, vol. 10, p. 100718, 2025, doi: 10.1016/j.atech.2025.100718.

[5] A. Oikonomidis, C. Catal, and A. Kassahun, "Deep learning for crop yield prediction: A systematic literature review," *N. Z. J. Crop Hortic. Sci.*, vol. 51, no. 1, pp. 1–26, 2023.

[6] V. Parganiha and M. Verma, "An effective soil analysis and crop yield prediction based on optimised Light GBM in smart agriculture," *J. Agron. Crop Sci.*, vol. 210, no. 4, p. e12726, 2024, doi: 10.1111/jac.12726.

[7] P. Mahesh and R. Soundrapandian, "Yield prediction for crops by gradient-based algorithms," *PLOS ONE*, vol. 19, no. 8, p. e0291928, Aug. 2024, doi: 10.1371/journal.pone.0291928.

[8] M. Bouni *et al.*, "Integrated IoT approaches for crop recommendation and yield-prediction using machine-learning," *IoT*, vol. 5, no. 4, pp. 634–649, Sep. 2024, doi: 10.3390/iot5040028.

[9] R. N. V. J. Mohan, P. S. Rayanoothala, and R. P. Sree, "Next-gen agriculture: Integrating AI and XAI for precision crop yield predictions," *Front. Plant Sci.*, vol. 15, p. 1451607, Jan. 2025, doi: 10.3389/fpls.2024.1451607.

[10] U. Nikhil, A. Pandiyan, S. Raja, and Z. Stamenkovic, "Machine learning-based crop yield prediction in South India: Performance analysis of various models," *Computers*, vol. 13, no. 6, p. 137, 2024, doi: 10.3390/computers13060137.

[11] K. S. M. Anbananthan, S. Subbiah, D. Chelliah, P. Sivakumar, V. Somasundaram, K. H. Velshankar, M. K. A. Khan, P. M. D. R. Vincent, F. Alzami, and M. Pant, "An intelligent decision support system for crop yield prediction using hybrid machine learning algorithms," *F1000Research*, vol. 10, p. 1143, Nov. 2021, doi: 10.12688/f1000research.73009.1.

[12] V. Bolo'n-Canedo, L. Mora'n-Ferna'ndez, B. Cancela, and A. Alonso-Betanzos, "A review of green artificial intelligence: Towards a more sustainable future," *Neurocomputing*, vol. 599, p. 128096, Sep. 2024, doi: 10.1016/j.neucom.2024.128096.

[13] S. Santos, A. L. C. Ottoni, R. Borgo, D. D. Ferreira, *et al.*, "A systematic review of Green Machine Learning: Practices and challenges for sustainability," *Artif. Intell. Rev.*, vol. 59, no. 5, Feb. 2026, doi: 10.1007/s10462-026-11515-8.

[14] E. Paula, J. Soni, H. Upadhyay, and L. Lagos, "Comparative analysis of model compression techniques for achieving carbon efficient AI," *Sci. Rep.*, vol. 15, p. 23461, Jul. 2025, doi: 10.1038/s41598-025-07821-w.

[15] D. Liu, Y. Zhu, Z. Liu, Y. Liu, C. Han, J. Tian, R. Li, and W. Yi, "A survey of model compression techniques: Past, present, and future," *Front. Robot. AI*, vol. 12, 2025, doi: 10.3389/frobt.2025.1518965.

[16] S. Khan, N. S. Naz, T. Mazhar, M. U. Tariq, T. Shahzad, S. Guizani, and H. Hamam, "Green AI techniques for reducing energy consumption in AI systems," *Array*, vol. 29, p. 100652, Mar. 2026, doi: 10.1016/j.array.2025.100652.

[17] R. K. Dhanaraj, M. Maragatharajan, A. Sureshkumar, and S. P. Balakannan, "On-device AI for climate-resilient farming with intelligent crop yield prediction using lightweight models on smart agricultural devices," *Sci. Rep.*, vol. 15, p. 31195, Aug. 2025, doi: 10.1038/s41598-025-16014-4.

[18] M. Baishya and L. Dutta, "Tiny ML based crop recommendation system for precision agriculture 5.0," *Smart Agric. Technol.*, vol. 12, p. 101247, Dec. 2025, doi: 10.1016/j.atech.2025.101247.

[19] T. Miller, G. Mikiciuk, I. Durluk, M. Mikiciuk, A. Łobodzińska, and M. S. nieg, "The IoT and AI in agriculture: The time is now – A systematic review of smart sensing technologies," *Sensors*, vol. 25, no. 12, p. 3583, Jun. 2025, doi: 10.3390/s25123583.

[20] V. Hiremani, R. M. Devadas, Preethi, R. Sapna, T. Sowmya, P. Gujjar, N. S. Rani, and K. R. Bhavya, "Federated learning for crop yield prediction: A comprehensive review of techniques and applications," *MethodsX*, vol. 14, p. 103408, May 2025, doi: 10.1016/j.mex.2025.103408.