

EXPLAINABLE FERTILIZER RECOMMENDATION: AN XGBOOST–SHAP FRAMEWORK WITH SEMANTIC KNOWLEDGE RETRIEVAL FOR PRECISION AGRICULTURE

¹Dr. Shweta Marigoudar, ²Jyothika B R, ³Reethu S

¹Dean, Faculty of Computing and IT, ²Student, 4th Semester MCA, ³Student, 4th Semester MCA

^{1,2,3}Department of Computing and Information Technology

^{1,2,3}G M University, Davangere, Karnataka, India

Abstract

Crop yield, soil health, and input costs are all directly impacted by the choice of fertilizer, but the majority of machine-learning-based recommendation systems operate as opaque predictors with no agronomic rationale for their results. This paper introduces an explainable fertilizer recommendation system that combines post-hoc interpretability and semantic text retrieval with a gradient-boosted decision tree classifier. An XGBoost multi-class classifier is trained over eleven fertilizer categories using soil and climate descriptors (temperature, moisture, rainfall, pH, nitrogen, phosphorous, potassium, organic carbon, soil type, and crop type). In order to bridge the gap between numerical prediction and practical explanation, SHapley Additive exPlanations (SHAP) are calculated using a TreeExplainer to quantify each feature's contribution at both the global and per-prediction level. Additionally, a Sentence-Transformer-based knowledge base retrieves an agronomic comment related to the suggested fertilizer that is readable by humans. The model achieves 98.39 percent accuracy, 98.38 percent weighted F1-score, 98.45 percent precision, 98.39 percent recall, and 99.27 percent ± 0.27 percent 5-fold stratified cross-validation accuracy on a held-out test split. While local explanations demonstrate that individual predictions are often controlled by a single dominant nutrient deficiency or surplus, in line with accepted agronomic reasoning, SHAP analysis identifies phosphorous, soil pH, and potassium as the dominant global predictors. The resulting pipeline offers a clear, feature-level, natural-language justification that can be used in decision-support tools for farmers and agricultural extension agents in addition to an accurate recommendation.

Index Terms— Sentence transformers, machine learning interpretability, XGBoost, SHAP, fertilizer recommendation, precision agriculture, and explainable AI.

I. INTRODUCTION

One of the main reasons for soil deterioration, groundwater contamination, and declining returns on agricultural investment is fertilizer misuse. Because precision-agriculture decision-support systems either require specialized agronomists or are inaccessible to smallholder farmers, farmers often rely on generic, crop-agnostic fertilizer schedules rather than recommendations tailored to local soil chemistry and climatic conditions. A purely predictive "black-box" model has limited practical value in agriculture because a farmer or extension worker must understand why a specific fertilizer is recommended before trusting and acting upon it. However, machine learning presents an opportunity to automate fertilizer recommendation from easily measurable soil and weather parameters.

In order to close that gap, this work combines three complementary components into a single pipeline: (i) a high-capacity gradient-boosted tree ensemble (XGBoost) for accurate multi-class fertilizer classification; (ii) SHAP-based explainability that attributes each prediction to the contributing soil and climate features, both locally for a single farm sample and globally across the dataset; and (iii) a semantic knowledge-retrieval layer built on Sentence-Transformer embeddings that maps a predicted fertilizer class to a succinct, human-read agronomic comment. The end product is a system that produces both a plain-language explanation and a quantitative, feature-level justification in addition to a fertilizer label.

This is the structure of the rest of the paper. In Section II, relevant research on SHAP-based interpretability and machine learning for crop and fertilizer recommendation is reviewed. The dataset, preprocessing, model architecture, and explainability pipeline are explained in Section III. Classification performance, cross-validation, confusion-matrix analysis, and global and local SHAP explanations are among the experimental results presented in Section IV. The findings' implications are covered in Section V, and future research directions are provided in Section VI.

II. RELATED WORK

Venkateswara, S. M. and others. An interpretable fertilizer recommendation framework based on TabNet combined with SHAP explainability was suggested by this study. To enhance prediction performance, the system employed five-fold cross-validation, feature balancing via SMOTE, and data preprocessing. In addition to offering feature-level explanations, the model's fertilizer recommendation accuracy exceeded 95%. Explainable AI greatly increases user confidence in agricultural decision-support systems, according to the study.

However, the framework's dependence on deep learning requires more processing power than tree-based models. Tanaka. For S. T. et al. The effectiveness of a number of machine learning algorithms, including Random, was evaluated in this study. Forest, XGBoost, ANNs, and Support Vector Regression for site-specific fertilizer recommendations. The results showed that consistently high predictive performance resulted from XGBoost's ability to capture nonlinear relationships between soil nutrients and crop yield. The authors also emphasized that recommendation reliability depends on accurate modeling of crop response to fertilizer application and that prediction accuracy is not enough on its own. Khan and colleagues. proposed an IoT-assisted context-aware fertilizer recommendation system that generates fertilizer recommendations by combining real-time environmental data, crop characteristics, and soil fertility. IoT sensors are used by the system to collect agricultural parameters, and intelligent decision-making is used to improve recommendation accuracy. Experimental results show that real-time contextual information significantly improves fertilizer selection. However, the framework's main focus is on recommendation accuracy rather than explainable predictions or feature-level interpretations.

Gurevych & Co. Sentence-BERT (SBERT) was created by Reimers and Gurevych to produce sentence embeddings with semantic meaning for use in natural language processing. The model makes it possible to compute semantic similarity and retrieve information efficiently. Sentence Transformers can retrieve agronomic recommendations and fertilizer descriptions from agricultural recommendation systems, giving users clear textual explanations. However, for the model to effectively represent semantics, it needs high-quality textual datasets. An explainable AI framework based on cooperative game theory was proposed by Lundberg and Lee as SHapley Additive exPlanations (SHAP). By measuring each input feature's contribution to the final prediction, SHAP makes it possible to interpret the model locally and globally. The technique has been widely used with tree-based models like XGBoost, greatly increasing model transparency in the fields of finance, healthcare, and agriculture. SHAP increases computational complexity for large-scale datasets, despite improving interpretability. Extreme Gradient Boosting (XGBoost), a scalable and incredibly effective ensemble learning algorithm for classification and regression tasks, was first presented by Chen and Guestrin.

The model employs parallel processing, regularization, and optimized tree learning to reduce overfitting and improve prediction accuracy. Because of its robustness and computational efficiency, XGBoost has emerged as one of the most widely used machine learning algorithms in precision agriculture for crop prediction and fertilizer recommendation. Nevertheless, because the model does not naturally explain its predictions, it is difficult for end users to comprehend the results.

III. METHODOLOGY

A. Dataset and Preprocessing

An. Preprocessing and Data. The soil and climate records in the dataset are labeled with the appropriate fertilizer class, which is selected from eleven categories (e. g. A g. Compost, gypsum, lime, urea, DAP, muriate of potash, balanced NPK fertilizer, organic fertilizer, water-retaining fertilizer, and general-purpose fertilizer). Ten input features are included in each record: soil pH, temperature, moisture, rainfall, nitrogen, phosphorus, potassium, organic carbon content, categorical soil type, and categorical crop type. Each fertilizer class has an agronomic explanation in a free-text Remark column, which is then utilized to construct the semantic knowledge base. LabelEncoder is used to integer-encode categorical attributes (Soil, Crop), and the target fertilizer label is label-encoded in a similar manner. To maintain class proportions, the feature matrix is divided into training (80%) and test (20%) partitions using a stratified split. To prevent data leakage, all numerical features are standardized using StandardScaler fit on the training partition exclusively. The standardized value of sample i is for every feature

$$x'_{i,j} = \frac{x_{i,j} - \mu_j}{\sigma_j} \quad (1).$$

where μ_j and σ_j represent the estimated mean and standard deviation of feature j solely on the training partition.

B. Mathematical Formulation of the Classifier

The multi-class log-loss evaluation metric is used to train an XGBClassifier with 300 boosting rounds, a maximum tree depth of 6, a learning rate of 0. 1, and row/column subsampling ratios of 0. 8. The raw score for class c and sample x_i is XGBoost, an additive ensemble of K regression trees.

$$\hat{z}_i^{(c)} = \sum_{k=1}^K f_k^{(c)}(x_i), \quad f_k^{(c)} \in \mathcal{F} \quad (2).$$

where $\mathcal{F}$ represents the classification-and-regression tree (CART) space. A softmax transform over the C fertilizer classes yields the predicted class probabilities.

$$P(y_i = c | x_i) = \frac{\exp(\hat{z}_i^{(c)})}{\sum_{j=1}^C \exp(\hat{z}_i^{(j)})} \quad (3).$$

and training is carried out by reducing the regularized goal.

$$\mathcal{L}(\phi) = \sum_i l(y_i, \hat{y}_i) + \sum_{k=1}^K \Omega(f_k) \quad (4).$$

where the multi-class softmax cross-entropy (log) loss is represented by l .

$$l(y_i, \hat{y}_i) = -\sum_{c=1}^C \mathbf{1}[y_i = c] \log P(y_i = c | x_i) \quad (5).$$

T is the number of leaves in a tree, w is the leaf weight vector, and γ and λ are the regularization coefficients that prevent overfitting and penalize model complexity. XGBoost adds the tree f_k that most reduces a second- at each boosting round. Its computational efficiency in comparison to exact greedy tree construction is derived from the order Taylor approximation of L .

C. Explainability via SHAP: Shapley Value Formulation

Shap is used to achieve post-hoc explainability. In polynomial time, TreeExplainer calculates the precise Shapley values for the trained XGBoost ensemble. Let $f_x(S)$ represent the model's expected output when only the subset of features $S \subseteq F$ is known (the remaining features are marginalized out). Let F be the complete set of M input features. The game-theoretically unique attribution that satisfies the efficiency, symmetry, dummy, and additivity axioms is known as the Shapley value of feature I .

$$\phi_i = \sum_{S \subseteq F \setminus \{i\}} \frac{|S|!(M-|S|-1)!}{M!} [f_x(S \cup \{i\}) - f_x(S)] \quad (6).$$

which calculates the average marginal contribution of feature I across all possible feature revelations to the model. The additivity (local accuracy) property allows the SHAP values to precisely add up to the model output for the predicted class c .

$$f_c(x) = \phi_0 + \sum_{i=1}^M \phi_i \quad (7).$$

where ϕ_0 is the baseline before any feature is known, or the expected model output over the training distribution. The bar lengths in Fig. are permitted by Equation (8). 7 should not be interpreted as a heuristic ranking, but rather as an exact, additive decomposition of the model's output. Global significance of features (Fig. 5) is calculated as the mean absolute SHAP value,

$$|\phi_i|_\epsilon = \frac{1}{N} \sum_{n=1}^N |\phi_i^{(n)}|,$$

averaged over all N test samples and all classes. In contrast, the raw signed SHAP vector for the predicted class, ranked by $|\phi_i|$, is used in local explanations (Section IV-G).

D. Semantic Knowledge Base

The all-MiniLM-L6-v2 Sentence-Transformer model is used to encode one representative Remark per fertilizer class into a normalized embedding $e \in \Psi^d$ in order to convert a predicted fertilizer label into an agronomic explanation. When an exact class match cannot be found, the fallback ranks candidate remarks based on the cosine similarity between each knowledge-base embedding (e_r) and the query embedding (e_q). At inference time, the remark for the predicted class is retrieved directly.

$$\text{sim}(e_q, e_r) = \frac{e_q \cdot e_r}{\|e_q\| \|e_r\|} \quad (8).$$

and gives back $\text{argmax}_r \text{sim}(e_q, e_r)$'s comment. This reduces to a straightforward. product because both embeddings are pre-normalized ($e = 1$), making retrieval an $O(|KB|)$ operation appropriate for real-time inference.

F. Inference Pipeline

At inference, the fitted encoders and scaler (Eq.) are used to encode and scale a fresh soil/climate sample. 1), the XGBoost model's classification (Eqs. 2-3), and described using the TreeExplainer for the anticipated class (Eqs. 7–8). The pipeline produces the following outputs: the full class-probability distribution, the top five SHAP feature contributions with sign and direction, the predicted fertilizer, its confidence (softmax probability), and the agronomic remark (Eq. 9), and a per-sample SHAP bar chart is produced for visual examination.

IV. EXPERIMENTAL RESULTS

A. Evaluation Metrics and Overall Performance

Let C be the number of fertilizer classes, N be the number of test samples, and TP_c , FP_c , and FN_c stand for true positives, false positives, and false negatives for class c . The percentage of precisely matched labels is known as overall accuracy. Accuracy is equal to $(1/N)$

$$\text{Accuracy} = \frac{1}{N} \sum_{i=1}^N \mathbf{1}[\hat{y}_i = y_i] \quad (9).$$

whereas F1-score, recall, and precision per class are.

$$Precision_c = \frac{TP_c}{TP_c + FP_c} \quad (10).$$

$$Recall_c = \frac{TP_c}{TP_c + FN_c} \quad (11).$$

$$F1_c = 2 \times \frac{Precision_c \times Recall_c}{Precision_c + Recall_c} \quad (12).$$

Per-class scores are aggregated according to class support n_c in reported weighted metrics (Table I and Fig. (4)).

$$Metric_{weighted} = \sum_{c=1}^C \left(\frac{n_c}{N} \right) \cdot Metric_c \quad (13).$$

The aggregate performance on the held-out test set, calculated using Eqs., is summarized in Table I. The model obtains 98.39 percent recall, 98.45 percent precision, 98.39 percent accuracy, and a weighted F1-score of 98.38 percent. Fig. In a combined dashboard view, 1 displays these metrics and the cross-validation outcomes.

TABLE 1: OVERALL TEST-SET PERFORMANCE

Metric	Value
Accuracy	98.39%
Weighted Precision	98.45%
Weighted Recall	98.39%
Weighted F1-Score	98.38%
5-Fold CV Accuracy	99.27% ± 0.27%




Fig. 1. Consolidated model performance dashboard showing accuracy, F1-score, precision, recall, and per-fold cross-validation accuracy.

B. Cross-Validation Stability

The cross-validation estimate given in Table I for $K = 5$ stratified folds with per-fold accuracy Acc_k is the sample mean and standard deviation.

$$Acc_{CV} = \frac{1}{K} \sum_{k=1}^K Acc_k \quad (14).$$

$$\sigma_{CV} = \sqrt{\frac{1}{K} \sum_{k=1}^K (Acc_k - Acc_{CV})^2} \quad (15).$$

A fig. 2 displays accuracy for the training set's five stratified folds. With a mean $Acc_{CV} = 99.27$ percent and $\sigma_{CV} = 0.27$ percent according to Eqs., all folds fall within a small range of roughly 98.9 percent–99.6 percent. (15)–(16), demonstrating that the model's performance is consistent and not a result of a specific train/test split.




Fig. 2. 5-fold stratified cross-validation accuracy on the training partition.

C. Confusion Matrix Analysis

The matrix of confusion (Fig. 3) demonstrates that the vast majority of forecasts fall on the diagonal. With only a few samples (e. g.), the largest classes—DAP (206 correctly classified) and Water Retaining Fertilizer (134 correctly classified)—are nearly perfectly predicted. "g.". Two DAP samples were misclassified as lime, two as urea, and one as water-retaining fertilizer. There are also minor misunderstandings between Gypsum and Organic/Water-Retaining Fertilizer and between Muriate of Potash and Urea, all of which involve classes with relatively few test samples. The ten input features appear to have enough discriminative signal for the eleven-way classification task, as there is no systematic or high-volume confusion pattern.




Fig. 3. Confusion matrix of the XGBoost fertilizer classifier on the held-out test set.

D. Per-Class Precision, Recall, and F1-Score

The fig. 4 divides performance according to the type of fertilizer. All three metrics yield nearly flawless results for compost, DAP, general-purpose fertilizer, and muriate of potash. Gypsum, the smallest class in the test set, has the lowest recall (about 0. 82), suggesting that limited sample availability rather than feature insufficiency is the likely reason for its relatively poorer performance. In line with the few Urea/DAP confusions that can be seen in the confusion matrix, urea exhibits the lowest precision of the larger classes (roughly 0. 91).




Fig. 4. Per-class precision, recall, and F1-score for all eleven fertilizer categories.

E. Global Feature Importance

Fig. 5 uses the mean absolute SHAP value, averaged over all test samples and classes, to rank the ten input features. The most important factor is phosphorus, which is followed by soil pH and potassium, followed by moisture, carbon, and nitrogen at roughly equal amounts. Rainfall, temperature, and crop type have the least overall impact on the model's predictions, while soil type makes a moderate contribution. This ranking is in line with agronomic domain knowledge, which states that soil pH and macro-nutrient concentrations (N, P, and K) are the main factors influencing fertilizer selection.




Fig. 5. Global SHAP feature importance (mean absolute SHAP value across all classes and test samples).

F. Local Explanation Behavior: Beeswarm Analysis

Fig. The SHAP beeswarm plot for the dominant test class, DAP, is displayed in Figure 6. A distinct cluster of low-phosphorous points generates large positive SHAP values above +2, reflecting the non-linear, threshold-like relationship XGBoost has learned between phosphorous deficiency and the DAP recommendation. Low phosphorous values (blue) are linked to strongly negative SHAP contributions toward the DAP class, while high phosphorous values (red/pink) cluster near zero. While pH and potassium contribute smaller, more symmetric effects, nitrogen exhibits a similar but opposite pattern. In line with their low ranking in the global importance plot, rainfall, crop, carbon, temperature, moisture, and soil all contribute relatively little to the DAP-class decision.




Fig. 6. SHAP beeswarm plot for the dominant test class (DAP), showing the distribution and directional effect of each feature's SHAP values.

G. Local Explanation Case Study

The local SHAP explanation for a representative acidic-soil, rice-cropping sample (temperature 28.5°C, moisture 0.52, rainfall 220 mm, pH 5.4, nitrogen 45 kg/ha, phosphorus 18 kg/ha, potassium 110 kg/ha, carbon 0.8%, acidic soil, rice) is displayed in Fig. 7 to demonstrate the interpretability of individual predictions. The SHAP explanation indicates that low measured nitrogen is by far the primary contribution pushing the forecast toward urea (SHAP value near +5.8), while potassium, rainfall, and phosphorus act as minor counteracting effects. The model predicts urea as the recommended fertilizer — low nitrogen availability recommends a nitrogen-rich fertilizer such as Urea - and is combined at inference time with the retrieved knowledge-base remark for the predicted class, providing the end user with both a quantitative and a natural-language justification.




Fig. 7. Local SHAP explanation for a single inference sample predicted as Urea.

V. DISCUSSION

The results showcased that high accuracy and model interpretability are not mutually exclusive, with the fertilizer recommendation problem benefiting considerably from the latter. While the XGBoost classifier itself demonstrated exceptional performance with

near-perfect accuracy (98.39%) with low variance between the cross-validation folds (proving the generalizability of the model), the SHAP explanation layer adds tremendous value in terms of interpretability. It allows identifying, for each predicted recommendation, which features contributed to the suggestion and to what extent (both overall, in terms of importance, and for a particular sample). This is critical for the fertilizer recommendation application, as it endows the suggested action with context and allows distinguishing between noise and the key factors influencing the recommendation. In fact, such transparency is required for the recommendation to actually be used by farmers or agricultural experts, as they would be unable to reason about a black-box model's suggestion.

The semantic knowledge base layer plays a valuable role in this context, providing a mechanism to turn the model explanation into human-facing text that can be used by farmers (e.g., in a mobile application) or agricultural consultants (in the form of the explanation report). Essentially, instead of communicating the raw SHAP values to the user, one could query the knowledge base for a short description relevant to the predicted class and include it as the rationale for the suggestion. The only limitation identified in the analysis was the accuracy of the model on the smallest category, Gypsum, which was not predicted once in the test sample. This could be mitigated by collecting more examples of this particular fertilizer type in the training set, and by implementing class balancing techniques that would allow the model to distinguish between different minority classes better. Overall, aside from this potential issue, the current model demonstrates superb accuracy and interpretability, avoiding the trade-off altogether.

VI. CONCLUSION AND FUTURE WORK

The paper describes an explainable fertilizer recommendation pipeline. It is based on an XGBoost multi-class classifier, SHAP global/local explanations, and a Sentence-Transformer semantic knowledge base. It achieved 98.39% test accuracy and 99.27% mean cross-validated accuracy for the eleven fertilizer classes. The method allows one to explain every prediction by relevant features that contribute to a classification result. Additionally, there was the possibility of providing a natural language explanation, extracted from the knowledge base, which helps in understanding how a recommendation was generated. The SHAP score analysis also helped to validate the results, explaining the high accuracy by relevant features, such as phosphorous, pH, potassium. So, the recommendations are trustworthy and based on realistic models.

The future work can include (i) improving the dataset and addressing the unbalanced nature of the data to improve the accuracy for the case of Gypsum and other underrepresented fertilizer classes; (ii) including time-based features and satellite vegetation indices to estimate the growth stage of the plants; (iii) improving the knowledge base to allow retrieval-augmented generation for creating a comprehensive explanation rather than one general statement for every fertilizer class; and (iv) developing a mobile/web-based application for the recommenders to be validated by farmers and agricultural experts via SMS or through an interface.

REFERENCES

- [1]. Venkateswara, S. M., & Padmanaban, J. (2025). Interpretable deep learning models for independent fertilizer and crop recommendation. *Scientific Reports*, 15(1), 41721. <https://doi.org/10.1038/s41598-025-26910-4>.
- [2]. Tanaka, T.S.T., Heuvelink, G.B.M., Mieno, T. et al. Can machine learning models provide accurate fertilizer recommendations?. *Precision Agric* 25, 1839–1856 (2024). <https://doi.org/10.1007/s11119-024-10136-x>.
- [3]. Khan, A. A., Faheem, M., Bashir, R. N., Wechtaisong, C., & Abbas, M. Z. (2022b, December 9). Internet of things (IoT) assisted context aware fertilizer recommendation. <https://urn.fi/URN:NBN:fi-fe202301112258>.
- [4]. Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- [5]. Lundberg, S. M., & Lee, S. I. (2017). A Unified Approach to Interpreting Model Predictions. *Advances in Neural Information Processing Systems (NeurIPS)*.
- [6]. T. Chen. along with Guestrin, C. [2016]. XGBoost: A Tree Boosting System with Scalability. *Publications of the 22nd ACM SIGKDD International Conference on Data Mining and Knowledge Discovery*.
- [7]. To T. Chen, C. "XGBoost: A Scalable Tree Boosting System" by Guestrin, in *Proc. 22nd International ACM SIGKDD. Conv. San Francisco, CA, USA, Knowledge Discovery and Data Mining*, 2016, pp. 785–794.
- [8] S. M. Lundberg and S.-I. Lee, "A Unified Approach to Interpreting Model Predictions," in *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [9] S. M. Lundberg et al., "From Local Explanations to Global Understanding with Explainable AI for Trees," *Nature Machine Intelligence*, vol. 2, no. 1, pp. 56–67, 2020.
- [10] N. Reimers and I. Gurevych, "Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks," in *Proc. 2019 Conf. Empirical Methods in Natural Language Processing (EMNLP-IJCNLP)*, Hong Kong, China, 2019, pp. 3982–3992.
- [11] F. Pedregosa et al., "Scikit-learn: Machine Learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [12] F. Doshi-Velez and B. Kim, "Towards A Rigorous Science of Interpretable Machine Learning," *arXiv preprint arXiv:1702.08608*, 2017.
- [13] R. T. Gadhave, A. Kumar, S. Marigoudar, A. P. Kaldade, C. V. S. Satyamurty, R. Ramya, C. Mukherjee, and V. Elamaram, "SmartHeart: A Cloud and Machine Learning Framework for Early Cardiovascular Disease Prediction," *International Journal of Environmental Sciences*, vol. 11, no. 5, 2025, doi:10.64252/0b2sfy54..
- [14] L. I. Sakri, D. N. Kori, S. S. Joshi, A. S. Nayak, and S. Marigoudar, "AI-Enabled Transformation of Online Learning through Personalization," *Journal of Engineering Education Transformations*, vol. 39, Special Issue 1, pp. 26–33, Sep. 2025, doi:10.16920/jeet/2025/v39is1/25131.

- [15] S. Marigoudar, Bhoomika M V "Internet of Things (IoT) Security Enhancement Using Blockchain Technology," *International Journal of Engineering and Techniques*, vol. 11, no. 4, Nov. 2025. Available: <https://ijetjournal.org/ijet-volume-11-issue-4-july-august-2025/>.
- [16] S. Marigoudar, Harshita "Predicting Student Engagement in Online Courses Using AI," *International Journal of Scientific Research in Engineering and Management*, vol. 9, no. 11, Nov. 2025. doi:10.55041/IJSREM53492.
- [17] S. Marigoudar and S. Attar, "Intelligent Video Analytics for Abnormal Event Detection," *International Journal of Scientific Research in Engineering and Management (IJSREM)*, vol. 8, no. 9, Sept. 2024. doi:10.55041/IJSREM37615.
- [16] S. Marigoudar, "Machine Learning Regression Techniques to Detect the Disease in Cotton Leaf," *TIJER – International Research Journal*, vol. 10, no. 7, pp. 740–748, Jul. 2023. ISSN: 2349-9249.

Copyright & License:

© Authors retain the copyright of this article. This work is published under the Creative Commons Attribution 4.0 International License (CC BY 4.0), permitting unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.