IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 6, October 2026
Impact Factor: 8.2

TerraSense AI: Soil Quality Classification Using Image Processing, Deep Learning and an AI-Based Soil Advisory System

Savita Mogare, Sushma Zade, Swastika Raul, Pranita Deshpande, Siddesh Nikam
Department of Information Technology
Sandip Institute of Technology and Research Centre, Nashik, India
(An Autonomous Institute Affiliated to Savitribai Phule Pune University, Pune)
mogares1311@gmail.com, sushdg@gmail.com, raulswastika20@gmail.com,
pranitadeshpande2005@gmail.com, siddeshnikam06@gmail.com

Abstract: *Knowing the soil type of a field is one of the first steps in agricultural planning, yet judging it by eye is subjective and laboratory analysis is slow. This paper presents TerraSense AI, a web-based prototype that sorts soil photographs into seven classes (Alluvial, Arid, Black, Laterite, Mountain, Red and Yellow) using a convolutional neural network (CNN) built with TensorFlow and Keras and served through Flask. The CyAUG-Dataset held 5,097 images before we audited it for exact duplicates, which turned up 54 duplicate groups and 55 extra files. We split the data only after handling the duplicates and resized every image to 224 × 224 RGB. On the held-out test split the model reached 93.42% accuracy, 92.89% weighted F1 and 83.07% macro F1, with Arid and Mountain soils noticeably weaker than the other classes. Around the classifier, the prototype adds a provisional, dataset-relative NPK rating drawn from a 257-row measurement table, a SQLite prediction history with analytics, and an AI Soil Advisor that is still under development. The CNN predicts soil type from an image; it does not measure fertility, nutrient concentration, crop suitability or overall soil health, and the NPK ratings are not validated agronomic thresholds. All results come from one dataset, so performance on other data remains untested. The contribution of this work is an integrated, transparent prototype together with an honest account of what its evaluation does and does not show*

Keywords: Soil Classification, Convolutional Neural Network, Image Processing, Deep Learning, Smart Agriculture, Flask, Soil Nutrients, AI-Based Advisory System

I. INTRODUCTION

Soil underpins crop production, and recognising the soil type of a field is an early input to decisions about how it should be managed [3]. In practice this is done either by looking at the soil or by sending it to a laboratory. Looking depends on the observer's experience and on the light and moisture at the time, while laboratory methods are more reliable but cost time and effort [3].

Deep learning is now widely used in agriculture, and recent surveys report that deep learning methods have become increasingly effective for soil image classification because they learn texture, colour and spatial patterns from images [2]. For soil in particular, reviews describe pipelines that run from image acquisition through segmentation and feature extraction to classification [5], [6], and deep learning techniques for soil image classification have been compared with one another [8].

It helps to keep two ideas apart. Soil *type* is a category such as alluvial or laterite. Soil *quality* or *health* depends on measured chemical, physical and biological properties such as pH, organic carbon and nutrient content, none of which a

Copyright to IJARSCT
www.ijarsct.co.in
DOI: 10.48175/IJARSCT-40365

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 6, October 2026
Impact Factor: 8.2

photograph classifier estimates. The system described here classifies the first and only hints at the second through a separate, provisional module.

A. Problem Statement

Farmers and students lack a quick, low-cost way to get a first-level soil-type label and see it next to basic nutrient context, and existing prototypes rarely say clearly what their outputs do and do not mean.

B. Objectives

The objectives of this work are: (1) to build a CNN that classifies soil images into seven classes; (2) to audit the dataset for exact duplicates before splitting; (3) to deploy the model through a Flask web dashboard with stored prediction history;
(4) to provide a clearly labelled, provisional NPK rating; and (5) to define safety boundaries for an AI Soil Advisor.

C. Contributions

TerraSense AI makes four contributions:
Working prototype. An image-classification system with a documented duplicate audit.
Clear separation. Image-based soil-type classification is kept apart from the provisional, dataset-relative nutrient rating.
History layer. A prediction-history and analytics layer built on SQLite.
Safe advisor design. An AI Soil Advisor with stated safety boundaries, labelled as work in progress.

II. LITERATURE SURVEY

Work relevant to this paper falls into three groups: soil image classification, soil analysis from chemical composition or sensor data, and systems that link soil classification to crop recommendation. Table I summarises seven sources taken from the project synopsis [TBV: details to be checked against the papers].

TABLE I: LITERATURE SURVEY OF DIFFERENT PAPERS

Research Paper	Methodology Used	Advantages	Research Gap / Our Contribution
Soil Quality Classification from Chemical Composition Using Machine Learning Methods with SHAP-Based Explanation (Humeniuk *et al.*, 2026) [1]	Machine learning using chemical soil properties with SHAP explanation.	High prediction accuracy; explainable results; useful for soil analysis.	Requires laboratory chemical testing. Our project uses a CNN on soil images for quick classification without lab tests.
Soil Image Classification and Segmentation: A Survey from Deep Learning, Multimodal Data and Hybrid Models (Boulos *et al.*, 2026) [2]	Survey of CNN, hybrid models and multimodal techniques for soil image classification.	Covers latest deep learning methods; provides comparison of techniques.	It is only a survey paper. Our project develops and implements a CNN-based soil classification system.
An Evaluation of Machine Learning for Soil Analysis in IoT-Enabled Smart Farming (Chaudhary *et al.*, 2026) [3]	Machine learning with IoT sensors for soil analysis.	Supports smart farming; improves agricultural decisions.	Depends on IoT sensors. Our project performs image-based classification using a CNN without additional hardware.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 6, October 2026
Impact Factor: 8.2

Hybrid CNN–ML Framework for Soil Classification and Crop Recommendation (Agarwal *et al.*, 2026) [4]	CNN for feature extraction combined with ML classifiers.	High classification accuracy; automatic feature extraction.	The hybrid model is more complex. Our project uses a simple CNN model with fast prediction and future fertilizer recommendation.
Machine Learning Approaches for Soil Image Classification: A Systematic Review (Chate & Bhamare, 2025) [6]	Reviews different machine learning and deep learning approaches for soil classification.	Compares various algorithms; identifies research trends.	No practical implementation. Our project implements CNN-based automated soil classification.
Classification of Soil Images Using Convolutional Neural Network (Chate & Bhamare, 2025) [7]	CNN-based classification of soil images.	High accuracy; automatic feature learning.	Focuses only on classification. Our project also provides fast prediction and future fertilizer recommendation.
Machine Learning and Computer Vision Technology to Analyze and Discriminate Soil Samples (Kaplan *et al.*, 2024) [9]	Uses computer vision and machine learning to analyze soil samples.	Accurate soil analysis; efficient image processing.	Mainly focuses on soil analysis. Our project offers an easy-to-use CNN-based classification system with automated feature extraction.

Three observations follow. First, public soil image datasets are limited, and lighting, camera quality, soil moisture and geography make classification harder [2], [3]. Second, some studies analyse soil from chemical composition [1] or apply computer vision to soil samples [9], which is a different task from soil-type classification. Third, systems that link soil classification to crop recommendation or yield prediction [4], [10] should be checked for whether their image and tabular data are paired at sample level [TBV]. A survey also lists limited data among the open difficulties of deep learning for soil images [2], and explainable AI has been applied to soil fertility classification [11]. For advisory text, hallucination in agricultural recommendations is a recognised risk [citation TBV].

A. Research Gaps

(i) Few prototypes state plainly that a soil-type label is not a fertility or quality measurement. (ii) Dataset audits for duplicates and leakage are rarely reported. (iii) Language-model advice in agriculture needs explicit safety boundaries. TerraSense AI addresses a practical part of these gaps: it audits duplicates before splitting, keeps the image classifier and the provisional nutrient module separate, and defines what the advisor may and may not say. It does not close the scientific gap between an image class and soil quality.

B. Additional Related Work

Beyond Table I, further studies informed this work. Dwivedi *et al.* review soil identification and classification using machine learning [5]. Chate and Bhamare review machine-learning approaches to soil image classification [6], classify soil images with a CNN [7] and compare deep learning techniques [8]. Sedimo *et al.* describe the SoilNet framework, which fuses deep learning models to classify soil types and predict crop yield [10], and Roufaida *et al.* study soil fertility classification with explainable AI [11]. [Authors: TBV. The descriptions of [5], [8], [10] and [11] are based on titles only and should be checked against the papers.]

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 6, October 2026
Impact Factor: 8.2

III. SYSTEM ARCHITECTURE AND METHODOLOGY

A. Overview and Dataset

TerraSense AI has two data paths. On the image path, an uploaded photograph is resized and passed to the CNN, which returns one of seven soil classes with its class probabilities. On the nutrient path, the user enters nitrogen (N), phosphorus (P) and potassium (K) values and gets ratings relative to the distribution in a local measurement table. The two paths are not trained jointly, and the nutrient ratings are not derived from the image.

The dataset is the CyAUG-Dataset. The literature describes CyAUG as a CycleGAN-based augmentation variant of an original photographic soil dataset with the same seven classes [citation TBV against the dataset page and licence].

B. System Architecture

Each component below carries a status: **[Implemented]** works in the prototype, **[Partial]** works but is incomplete, and **[Proposed]** is future work. The overall architecture is shown in Fig. 1.

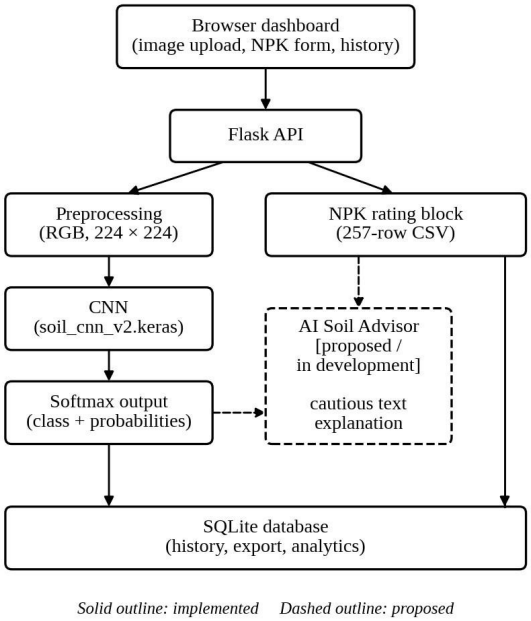


Fig. 1. TerraSense AI system architecture. Solid outline: implemented; dashed outline: proposed.

C. Data Collection and Dataset Audit [Implemented]

The CyAUG-Dataset provided 5,097 images in seven classes. Before splitting, we audited the files for exact duplicates and found 54 duplicate groups containing 55 extra files. The source, licence and capture conditions of the underlying photographs are [TBV].

D. Duplicate Handling and Data Splitting [Implemented]

Duplicates were handled first and the dataset was split afterwards (Fig. 2). The split ratios, whether the split was stratified, the random seed, and whether synthetic images are grouped with their source images are [TBV].

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 6, October 2026
Impact Factor: 8.2

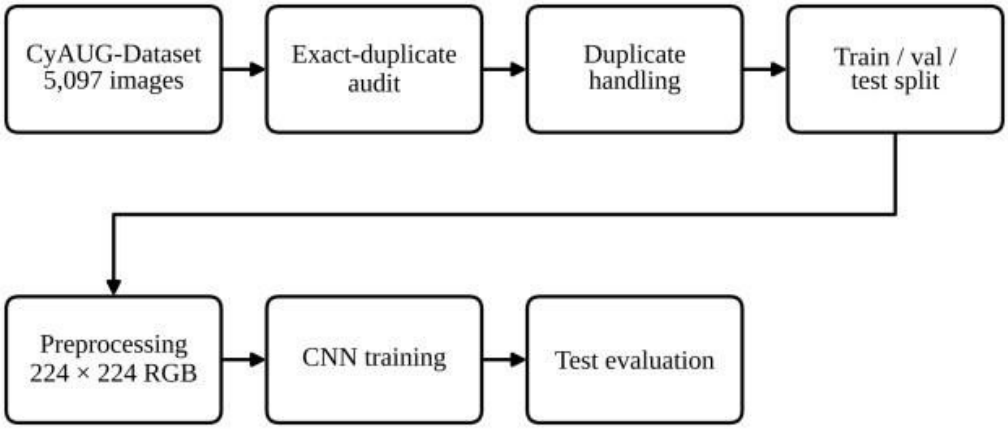


Fig. 2. Dataset preparation and CNN training workflow.

E. Image Preprocessing [Implemented]

Images are converted to RGB and resized to 224 × 224 pixels. Whether pixel values are rescaled (for example to 0– 1) inside or outside the model, and which augmentations were used in training, are [TBV].

F. CNN Architecture and Training [Implemented; details TBV]

The classifier is a CNN built with TensorFlow and Keras and saved as models/soil_cnn_v2.keras, with class names in models/soil_class_names.txt. Layer types and counts, filter sizes, parameter count, regularisation, loss function, optimiser, learning rate, batch size, epochs and callbacks will be copied from the training notebook (model summary and compile call) rather than inferred. The meaning of the v2 suffix should also be explained.

G. Model Evaluation [Implemented; partial reporting]

Accuracy, weighted F1 and macro F1 are reported in Table III. Per-class metrics, the confusion matrix, calibration and external validation are not yet reported and are listed as future work.

H. Flask Backend and Prediction API [Implemented]

A Flask application loads the saved model once and exposes a prediction endpoint that accepts an uploaded image and returns the predicted class and probabilities. Route names, file-size and file-type limits, and error handling are [TBV]. Latency was not measured.

I. Web Dashboard [Implemented]

An HTML, CSS and JavaScript dashboard supports image upload and shows the prediction (Fig. 3).

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 6, October 2026
Impact Factor: 8.2

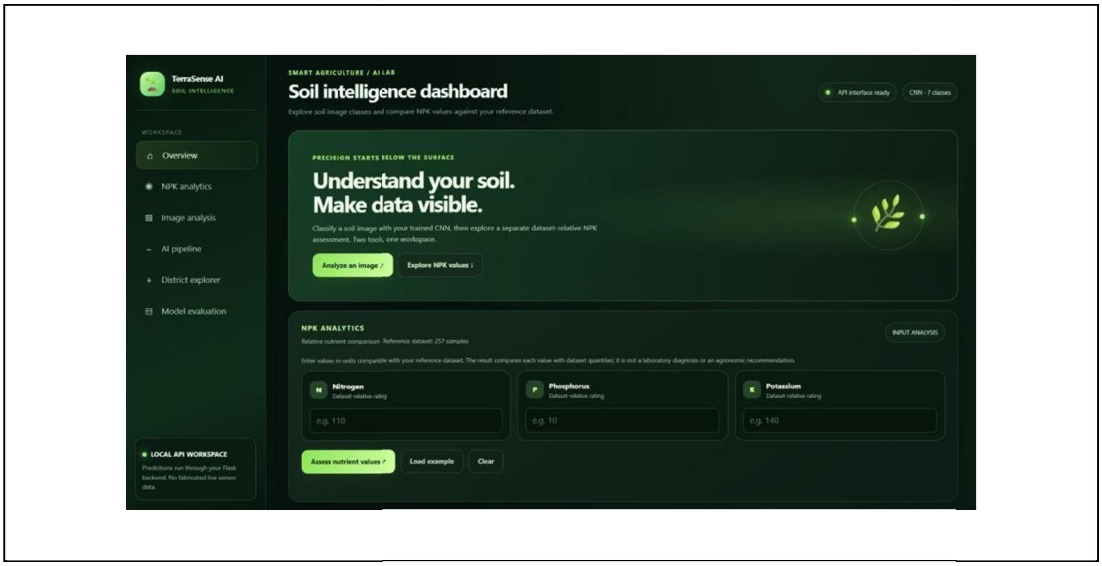


Fig. 3. Soil image prediction dashboard.

J. Provisional NPK Analysis [Partial]

An endpoint accepts N, P and K values and returns ratings relative to the distribution of the 257-row table soil_measurements_clean.csv (columns Sample, pH, EC, OC, N, P, K, S, Zn, B, Fe, Mn, Cu). A rating only says where an entered value sits within that table; it is not an agronomic sufficiency threshold. The units, sampling method, laboratory method and provenance of the table are [TBV], and so is the exact rating rule, which is why Section IV gives no equation for it.

K. SQLite Prediction History and Analytics [Implemented]

Predictions are stored in SQLite, with endpoints for history, export and analytics. The stored fields and the retention of uploaded images are [TBV], and any privacy handling of uploads should be stated.

L. AI Soil Advisor and Safety Boundaries [Proposed / under development]

The advisor is meant to combine the model's actual outputs and user-provided soil information with a generated explanation (Fig. 4). Because language models can give confident but unsupported agricultural advice [citation TBV], we propose these boundaries: the advisor must quote the model's class and confidence rather than invent them; it must say that the class is not a fertility measurement; it must describe NPK ratings as dataset-relative; it must not give fertiliser doses, pesticide advice or laboratory-equivalent recommendations; and it must point users to laboratory testing and local agricultural extension advice for decisions. The language model used, the prompts and any retrieval source are [TBV].

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 6, October 2026
Impact Factor: 8.2

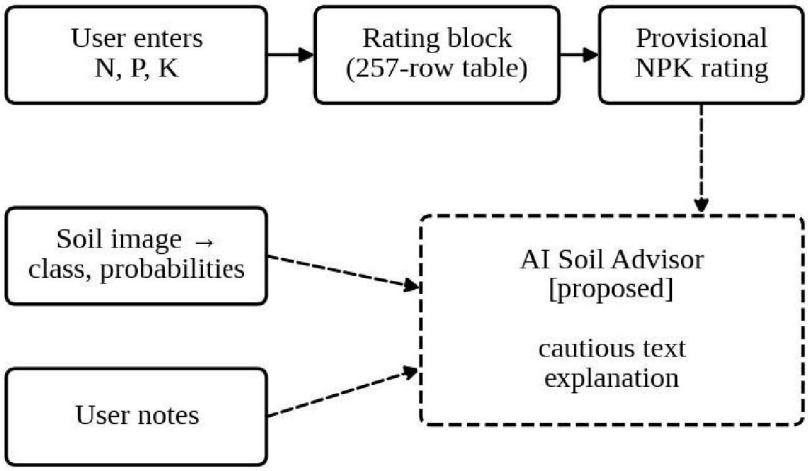


Fig. 4. NPK analysis and AI Soil Advisor workflow.

IV. MATHEMATICAL FORMULATION

The image task is modelled as seven-class classification. Equations are given only where they describe the implemented pipeline or a standard evaluation measure; items still awaiting confirmation are marked.

Image resizing. An input image x is converted to RGB and resized to 224 × 224 pixels, where R is the resize operator (interpolation method [TBV]).

$$x' = R(x), x' \in \mathbb{R}^{224\times224\times3} \quad (1)$$

Normalisation [to be kept only if confirmed]. If pixel values are rescaled by a constant, x'' is the scaled image.

$$x'' = x'/255 \quad (2)$$

Softmax classification. The network maps x'' to $K = 7$ logits $z_1, \ldots, z_K$, and the probability of class k is

ezk

Categorical cross-entropy [to be kept only if confirmed as the training loss]. For a one-hot label y and predicted probabilities p, over a batch of N samples:

$$\mathcal{L} = -\frac{1}{}\sum_{i=1}^{N}\sum_{N^K k=1} y_{i,k} log\ p_{i,k} \quad (4)$$

Evaluation metrics. For class k, TP_k, FP_k and FN_k are the true positives, false positives and false negatives, and N is the number of test images.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 6, October 2026
Impact Factor: 8.2

$$Accuracy = \frac{1}{N_k}\sum TP_k$$

$$Precision_k = \frac{TP_k}{TP_k + FP_k}$$

$$Recall_k = TP_k + FN_k \quad (7)$$

$$F1_k = Precision \quad 2\ Precision_k + k\ Recall\ Recall_{kk} \quad (8)$$

(9)

$$Macro - F1 = \frac{1}{K}\sum_{k=1}^{K} F1_{Kk} \quad (10)$$

Here n_k is the number of test images of class k (its support). Macro-F1 weights every class equally, while WeightedF1 weights each class by its share of the test set, which explains the gap between the two reported values.

Nutrient rating. The dataset-relative rule used by the prototype has not been confirmed, so no equation is claimed. Once confirmed, it should be written out together with the percentile or threshold values and the units.

V. RESULTS AND DISCUSSION

A. Dataset Class Distribution

Table II is to be filled from the dataset audit log. Its class counts decide how far accuracy can be trusted.

TABLE II: DATASET CLASS DISTRIBUTION (TO BE FILLED FROM THE DATASET AUDIT LOG)

Class	Before	After	Train	Val	Test
Alluvial Soil	[TBV]	[TBV]	[TBV]	[TBV]	[TBV]
Arid Soil	[TBV]	[TBV]	[TBV]	[TBV]	[TBV]
Black Soil	[TBV]	[TBV]	[TBV]	[TBV]	[TBV]
Laterite Soil	[TBV]	[TBV]	[TBV]	[TBV]	[TBV]
Mountain Soil	[TBV]	[TBV]	[TBV]	[TBV]	[TBV]
Red Soil	[TBV]	[TBV]	[TBV]	[TBV]	[TBV]
Yellow Soil	[TBV]	[TBV]	[TBV]	[TBV]	[TBV]
Total	5,097	[TBV]	[TBV]	[TBV]	[TBV]

Before / After = images before and after duplicate handling. If all 55 extra files were removed the total would be 5,042 [TBV]. *B. Classification Performance*

TABLE III: CLASSIFICATION PERFORMANCE OF THE IMAGE CLASSIFIER (SOIL_CNN_V2)

Model	Evaluation Set	Acc.	Wtd F1	Macro F1
CNN (soil_cnn_v2.keras)	Held-out CyAUG test split after duplicate handling; 7 classes; size [TBV]	93.42%	92.89%	83.07%

Values are as recorded by the authors and are to be confirmed against the experiment log [TBV]. Per-class precision and recall, the confusion matrix, inference latency and cross-validation results are not available and are not reported.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 6, October 2026
Impact Factor: 8.2

C. Interpretation of Results

Accuracy is the share of test images classified correctly. Weighted F1 averages the per-class F1 scores in proportion to class size, so it follows the large classes. Macro F1 gives every class the same weight regardless of size. The macro F1 of 83.07% sits about 9.8 points below the weighted F1 of 92.89%, which means at least one class scores clearly lower than the average, and that class is probably smaller or harder. This fits the weakness we saw for Arid and Mountain soils, but the exact classes, their F1 values and their sizes cannot be pinned down without the per-class report [TBV].

Macro F1 exposes these weaker classes where accuracy hides them. Mountain soil is not defined by one colour or texture, and photographs of arid ground can look like other light-coloured soils. These are candidate explanations to test, not established causes; the confusion matrix, the class counts and a review of misclassified images are the obvious next checks.

D. Class Imbalance and Duplicates

The class counts in Table II decide how far accuracy can be trusted. Exact duplicates were found and handled before splitting, so identical files cannot sit in both training and test data. The audit only catches byte-identical or hashidentical copies, however, so near-duplicates are not covered.

E. Generalization

If the CyAUG variant contains synthetic images generated from original photographs, images derived from the same source photograph could land on both sides of a random split, and exact-duplicate removal would not detect that. Until the split is shown to be source-grouped, or the model is tested on independent real photographs, 93.42% should be read as in-dataset performance and not as expected field accuracy [TBV].

F. Accuracy versus Usefulness

A correct soil-type label is useful only if the user needs that label. It does not tell a farmer whether the soil is fertile, which nutrients are lacking, or which crop to grow.

G. Dataset Bias

If images within a class share backgrounds, cameras, lighting or generation artefacts, the model may be learning those cues rather than soil properties. We tested on one dataset only.

H. Existing Features and Limitations

Table IV summarises the status and current limitation of each component.

TABLE IV: EXISTING FEATURES AND LIMITATIONS

Component	Status	Current Limitation
Soil image classifier (7 classes)	Implemented	In-dataset evaluation only; Arid and Mountain weaker; no per-class report
Duplicate audit	Implemented	Exact duplicates only; near-duplicates and synthetic siblings unchecked
Flask API and dashboard	Implemented	Latency, load and security not evaluated
NPK rating	Partial	Dataset-relative; units and provenance unconfirmed; no agronomic thresholds
SQLite history, export, analytics	Implemented	Schema and privacy handling to be documented

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 6, October 2026
Impact Factor: 8.2

AI Soil Advisor	Proposed	Not validated; must not issue fertiliser or crop recommendations

VI. CONCLUSION

This paper described TerraSense AI, a Flask-based prototype that classifies soil photographs into seven classes with a CNN, stores prediction history in SQLite, offers a provisional dataset-relative NPK rating, and includes an AI Soil Advisor that is still under development. After exact-duplicate handling, the classifier reached 93.42% accuracy, 92.89% weighted F1 and 83.07% macro F1 on the held-out test split, with Arid and Mountain soils weaker.

A. Limitations

The CNN predicts soil type from an image, not fertility, nutrient content, crop suitability or soil health. The NPK ratings are not validated thresholds. Evaluation covers one dataset whose augmented images may not be independent of their sources, so performance on real field photographs is unknown. Per-class metrics, a confusion matrix and latency have not been reported, and the advisor is unvalidated.

B. Future Work (Not Yet Implemented)

(1) Report per-class precision, recall and a confusion matrix, and evaluate with source-grouped splits. (2) Test on an independent external set of real field photographs. (3) Add explainability such as Grad-CAM to show which image regions drive a prediction. (4) Integrate soil sensors and laboratory-validated nutrient data with documented units. (5) Develop soil-quality models trained on measured properties. (6) Ground the advisor in verified extension sources and evaluate its factual accuracy. (7) Add crop-specific recommendations only after expert and field validation.

ACKNOWLEDGMENT

The authors thank their project guide and the faculty of the Department of Information Technology, Sandip Institute of Technology and Research Centre, Nashik, for their guidance, and the creators of the datasets and open-source tools used in this work. [Wording to be confirmed with the project guide.]

REFERENCES

[1]. H. Humeniuk, D. Tymoshchuk, and A. Sverstiuk, "Soil quality classification from chemical composition using machine learning methods with SHAP-based explanation," *Environmental Challenges*, vol. 22, Art. no. 101404, 2026, doi: 10.1016/j.envc.2025.101404.

[2]. J. Boulos, V. Eglin, B. Kerautret, E. Larue, and J.-M. Côme, "Soil image classification and segmentation: A survey from deep learning, multimodal data and hybrid models," *Geodata and AI*, Art. no. 100053, 2026, doi: 10.1016/j.geoai.2026.100053.

[3]. P. Chaudhary, P. Gulia, N. S. Gill, N. Alduaiji, P. K. Shukla, and S. Vishnoi, "An evaluation of machine learning for soil analysis in internet of things-enabled smart farming," *Scientific Reports*, vol. 16, Art. no. 10318, 2026. [Online]. Available: https://www.nature.com/articles/s41598-026-36017-z

[4]. S. Agarwal, H. Bansal, and R. Singh, "Hybrid CNN–ML framework for soil classification and crop recommendation," *Int. J. Electrical and Electronics Research (IJEER)*, vol. 14, no. 2, pp. 389–402, Jun. 2026, doi: 10.37391/IJEER.140216.

[5]. R. K. Dwivedi, N. Kumari, A. Bishnoi, and R. P. Pandey, "Soil identification and classification using machine learning: A review," in *Proc.*
2022 11th Int. Conf. System Modeling & Advancement in Research Trends (SMART), Moradabad, India, Dec. 2022, pp. 958–963. [Online]. Available: https://ieeexplore.ieee.org/document/10047446

[6]. G. D. Chate and S. S. Bhamare, "Machine learning approaches for soil image classification: A systematic review," *Int. J. Innovative Science and Research Technology*, vol. 10, no. 4, pp. 825–829, Apr. 2025, doi: 10.38124/ijisrt/25apr045.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 6, October 2026
Impact Factor: 8.2

[7]. G. D. Chate and S. S. Bhamare, "Classification of soil images using convolutional neural network," *Int. J. Image, Graphics and Signal Processing (IJIGSP)*, 2025, doi: 10.5815/IJIGSP.2025.05.03. [Volume and pages TBV.]

[8]. G. D. Chate and S. S. Bhamare, "Comparative performance analysis of deep learning techniques for soil image classification," 2025, doi: 10.14445/22315381/IJETT-V73I4P101. [Venue TBV: the synopsis lists IJCA, but the DOI contains IJETT.]

[9]. S. Kaplan, E. Ropelewska, S. Günaydın, K. Sabancı, and N. Çetin, "Machine learning and computer vision technology to analyze and discriminate soil samples," *Scientific Reports*, vol. 14, Art. no. 19945, 2024. [Online]. Available: https://www.nature.com/articles/s41598024-69464-7

[10]. K. Sedimo *et al.*, "Design of SoilNet framework to classify soil types and predict crop yield using fusion deep learning models," 2025, doi: 10.51983/IJISS-2025.IJISS.15.3.36. [Journal name TBV; the DOI suggests IJISS, vol. 15, no. 3.]

[11]. K. Roufaida, M. Abdelhak, and R. Khaled, "Unveiling soil fertility classification with explainable AI," in *Proc. IEEE 3ICT*, 2024. [Online]. Available: https://ieeexplore.ieee.org/document/10824403 [Full proceedings title TBV.]