
Visual Soil Analysis: An AI-Powered Solution

Meraal Rahman¹, Yasir Niaz Khan²

¹Syed Babar Ali School of Science and Engineering Lahore University of Management Sciences
(RISE Intern) Lahore, Pakistan

²Lahore University of Management Sciences (Senior Research Fellow) Lahore, Pakistan
meraal.rahman@gmail.com, yasir.niaz@gmail.com

Abstract

This paper presents an image-based approach to soil classification using deep learning techniques. A custom dataset of 500 soil images, representing five distinct categories—Sandy, Loam, Alluvial, Clay, and Coarse—was compiled and used to train a YOLOv8 object detection model. To enhance prediction reliability, a Python-based histogram correlation method was integrated, comparing RGB channel distributions between input and reference samples. The system outputs the predicted soil type along with a confidence score and similarity metric. Results demonstrate high classification accuracy and practical usability, delivered through a publicly accessible web application built with Streamlit, enabling real-time soil identification for researchers, students, and field professionals.

Keywords— Soil classification, YOLOv8, deep learning, image processing, color histogram, environmental analysis.

I. INTRODUCTION

Soil classification is essential in agriculture, environmental monitoring, and land management, guiding decisions on crop suitability, irrigation, and infrastructure. Traditional methods rely on manual sampling and laboratory analysis, which are time-consuming, costly, and require expert knowledge—making them impractical in remote or resource-limited areas. With the rise of artificial intelligence (AI) and computer vision, there is growing potential to automate this process and make it more accessible.

To address these limitations, this research presents an image-based soil classification system powered by YOLOv8 [1], a cutting-edge object detection model. Trained on a curated dataset of five soil types—Sandy, Loam, Alluvial, Clay, and Coarse—the model is enhanced with histogram-based validation to improve reliability. The final system is deployed via a user-friendly web application [2], developed as part of this research, allowing users to upload soil images and receive classification results in real time. By combining deep learning with intuitive design, the project offers a scalable and efficient alternative to traditional soil analysis.

II. RELATED WORK

Recent advancements in soil classification have increasingly leveraged deep learning and image analysis to improve accuracy and efficiency. While traditional methods are time-consuming and often impractical for large-scale applications, modern approaches utilize computer vision and machine learning to automate classification based on visual features such as texture, color, and structure.

Srivastava, Shukla, and Bansal [3] introduced a histogram-based image analysis technique for soil classification, demonstrating that statistical patterns in pixel intensity can serve as reliable indicators of soil type. While effective, such methods often depend on controlled lighting and consistent image quality, limiting their robustness in real-world conditions.

Philip, Reddy, and Gupta [4] provided a comprehensive review of soil classification techniques, dividing them into conventional image processing and deep learning-based approaches. Their analysis

emphasized the superiority of Convolutional Neural Networks (CNNs) in capturing complex soil features, while also noting the challenges posed by environmental variability and limited datasets. Jalapur and Patil [5] explored emerging trends such as lightweight models, multi-task learning, and transfer learning for soil classification. Their study highlighted the potential of deep learning to overcome limitations in traditional systems, particularly in terms of scalability and adaptability. However, they also pointed out persistent issues like data imbalance and model interpretability.

Beyond soil classification, artificial intelligence has seen widespread adoption in agriculture and environmental monitoring. Araújo et al. [6] reviewed machine learning applications in agriculture, highlighting its role in crop prediction, soil health assessment, and resource optimization. Gupta and Pal [7] explored AI in precision agriculture, emphasizing the integration of satellite imagery and sensor data for soil and pest detection. Thomasson et al. [8] discussed the broader landscape of AI in farming, including autonomous systems and decision-support tools.

In the environmental domain, Alotaibi and Nassif [9] conducted a bibliometric analysis of AI in environmental monitoring, noting a sharp rise in publications since 2010. Binetti, Massarelli, and Uricchio [10] reviewed machine learning in geosciences, focusing on complex tasks such as water quality prediction and land degradation mapping. Magazzino [11] examined the impact of deep learning on climate modeling and pollution tracking, underscoring its potential for large-scale environmental analysis.

Recent studies have also explored complementary domains relevant to soil classification. Jain, Sethia, and Tiwari [12] reviewed spectral-based soil nutrient prediction using machine learning, emphasizing the role of AI in sustainable agriculture. Abbes, Jarray, and Farah [13] examined remote sensing techniques for soil moisture estimation, integrating physical models with machine learning for large-scale monitoring. Schiller, Stiller, and Ryo [14] discussed the importance of explainable AI in environmental science, highlighting the need for transparency and trust in high-stakes decision-making. These works reinforce the broader relevance of AI in environmental diagnostics and support the interpretability goals of our hybrid system.

While previous studies have demonstrated the versatility of AI across domains [6]-[14], few have explored the fusion of structural and spectral analysis for soil classification [3]-[5]. This research addresses that gap by combining YOLOv8's object detection with histogram-based color validation [1], offering a dual-perspective system that enhances interpretability and reliability. The hybrid approach is implemented in a publicly accessible web application titled Soil Type Detector [2], enabling real-time soil classification without specialized equipment—making the system both technically innovative and practically impactful.

III. METHODOLOGY

A) Dataset Preparation

Building on the insights and limitations identified in prior studies [3]-[5], this research proposes a hybrid soil classification pipeline that integrates deep learning and statistical validation. By combining YOLOv8's real-time object detection capabilities [1] with histogram-based analysis [3], the system aims to enhance both accuracy and interpretability under varied environmental conditions. The following section outlines the design and implementation of this methodology in detail.

To develop a reliable soil classification model [1], a dataset of 500 soil images was compiled, representing five distinct soil types: Sandy, Loam, Alluvial, Clay, and Coarse. Each image was selected to capture natural variations in texture, color, and lighting conditions, ensuring that the dataset reflects real-world diversity. To maintain consistency, all images were manually preprocessed—including resizing, normalization, and format standardization—to reduce noise and improve model performance.

The dataset was intentionally balanced, with approximately 100 images per class, to prevent class imbalance and enhance the model's generalization capability [1]. For annotation, makesense.ai [15]—a free, web-based image labeling tool that supports bounding box creation and class assignment—was used. This platform enabled efficient manual labeling of soil regions, with annotations exported in YOLO-compatible format for seamless integration into the training pipeline.

Fig. 1 shows a screenshot of the annotation interface in makesense.ai [15]. The tool allows users to draw bounding boxes around regions of interest and assign class labels with minimal effort. In the example shown, the soil region is to be labeled as “Loam,” demonstrating the clarity and precision of the annotation process. This step was critical in preparing high-quality labeled data, which directly influenced the accuracy and reliability of the trained YOLOv8 model [1].

In addition to annotation and preprocessing, the dataset itself warrants further description to clarify its structure, class definitions, and limitations.

The five soil types were defined using a combination of visual characteristics—such as granularity, surface texture, and coloration—and pedological criteria drawn from standard soil science. For example, Loam was identified by its balanced texture and dark brown hue, consistent with its pedological composition of roughly equal sand, silt, and clay. Clay samples exhibited fine particles and a smooth, dense surface, while Coarse soils showed visible



Fig. 1. Simple image annotation interface in makesense.ai [15].

granules and low cohesion. Alluvial soils in this dataset were characterized by a dark brown to nearly black coloration, with a soft, moist texture—attributes commonly associated with nutrient-rich, river-deposited sediments in humid conditions. Sandy soils were recognized by their pale tone and loose, gritty structure. These definitions guided both image selection and annotation, ensuring that each class reflected meaningful distinctions in both appearance and composition.

B) Model Training

The YOLOv8 [1] model was selected for its proven efficiency and accuracy in object detection tasks, particularly in resource-constrained environments. For model deployment, the YOLOv8n.pt variant from Ultralytics was selected due to its lightweight architecture and suitability for real-time inference. This made it ideal for deployment in a web-based application where speed and responsiveness are critical.

The model was trained using a manually labeled dataset of 500 soil images, representing five distinct soil types [1]. The training process involved fine-tuning the pre-trained weights to adapt to the unique visual characteristics of soil textures, including granularity, color gradients, and surface patterns. To improve generalization and reduce overfitting, dataset diversity was achieved through manual image acquisition. Each soil sample was photographed from multiple angles and under varied lighting conditions, simulating the effects of rotation, brightness variation, and perspective shift. This approach served as a form of natural augmentation.

The model was trained for 50 epochs using an image size of 640×640 pixels (imgsz = 640). Training was conducted using the default hyperparameters provided by the Ultralytics framework [1], which include a learning rate of 0.001, batch size of 16, and the Adam optimizer. These settings were selected to balance training efficiency with model performance, and no custom tuning was applied beyond the default configuration.

Fig. 2 below presents the training performance metrics captured during model training. The plots include loss curves for box regression, classification, and distribution focal loss across both training and validation sets. Additionally, precision, recall, and mean Average Precision (mAP) scores are shown, reflecting the model's ability to accurately detect and classify soil regions. The consistent decline in loss values and the upward trend in mAP metrics indicate




Fig. 2. Results of YOLOv8 training.

effective learning and strong generalization. These results validate the robustness of the YOLOv8 model in handling diverse soil imagery [1] and support its integration into the final web-based classification system [2].

C) Model Selection Justification

While this study focuses exclusively on YOLOv8 [1], its selection was informed by comparative analyses in recent literature. Prior works in soil classification have employed models such as ResNet50, EfficientNet, and custom CNN architectures [3], [5], which demonstrated moderate accuracy but often required separate stages for feature extraction and classification. These approaches typically suffer from slower inference and limited scalability in real-time applications. In contrast, YOLOv8 integrates detection and classification into a single, anchor-free pipeline optimized for speed and precision. Mustapha et al. [16] proposed a hybrid deep learning model combining YOLOv8 and CNN for waste detection, demonstrating superior performance over traditional CNN-based models in both mean average precision (mAP) and inference time. Similarly, Keylabs.ai [17] conducted a comparative analysis showing that YOLOv8 consistently outperforms Faster R-CNN in real-time object recognition tasks. Given the similar challenges posed by soil imagery, including texture variability and environmental noise, YOLOv8 was deemed the most suitable model for this application [1].

D) Histogram-Based Validation

To complement the YOLOv8 [1] predictions, a Python program was developed to generate color histograms of uploaded images. These histograms are not displayed directly to the user but are used internally to compare RGB channel distributions against a set of reference samples. The comparison relies on correlation metrics to assess similarity, helping validate the predicted soil type. This process was implemented using Python libraries including OpenCV for image processing [18], NumPy for numerical operations [19], and Matplotlib for visualization [20].

To begin the validation process, the system reads the input image using OpenCV and extracts its color histograms across the Red, Green, and Blue channels [18]. Each histogram consists of 256 bins, corresponding to pixel intensity levels from 0 to 255. The frequency of each intensity value is calculated by iterating through every pixel in the image and incrementing the count for the respective bin in each channel. This results in a 256×3 matrix representing the distribution

of color intensities across the image.

For each soil class, a single reference image was selected to represent the typical color and texture characteristics of that soil type under ideal lighting and environmental conditions. These images were chosen manually based on clarity, minimal background interference, and accurate visual representation of the soil's defining features. The RGB histograms of these reference images were extracted and normalized using the same pipeline as the input image, ensuring consistency in comparison. Although only one image per class was used, care was taken to ensure that each reference sample captured the essential visual traits of its respective soil type.

To quantify similarity between the input image and reference samples, the system computes and compares RGB histograms using OpenCV's `cv.compareHist()` function with the Pearson correlation coefficient (`cv.HISTCMP_CORREL`) [18]. The process begins by reading the input image and extracting its color histograms—one for each of the Red, Green, and Blue channels—using 256 bins per channel to represent pixel intensity values from 0 to 255. These histograms are then normalized using `cv.normalize()` with `NORM_MINMAX` scaling, ensuring that all values fall within the range (0, 1) regardless of image size or lighting conditions. The same process is applied to the set of reference images—one for each soil class—to create a baseline for comparison.

Each normalized channel histogram of the input image is compared against the corresponding channel histogram of each reference soil type. The Pearson correlation coefficient is computed independently for the R, G, and B channels. The three resulting correlation values are then averaged to obtain a single similarity score. To make this score more interpretable, it is transformed from the correlation range (-1 to 1) to a percentage scale (0% to 100%) using the formula:

$$\text{Score} = \left(\frac{\text{Correlation}(R) + \text{Correlation}(G) + \text{Correlation}(B)}{3} + 1 \right) \times 50$$

This transformation ensures that a perfect match yields 100%, while a complete mismatch results in 0%. The system performs this comparison for all five reference soil types—Loam, Clay, Coarse, Alluvial, and Sandy—and selects the class with the highest similarity score as the final match. Rather than applying a fixed threshold, the system selects the soil type whose reference histogram yields the highest average correlation score with the input image. This ensures that a prediction is always made, even in cases where all scores fall below conventional confidence levels. While scores above 85% often indicate strong similarity, the final decision is based solely on relative comparison across all soil classes. This histogram-based validation serves as an independent check on the YOLOv8 model's output [1], enhancing the system's reliability in real-world conditions.

The result of this validation is presented to the user as a simplified output, as shown in Fig. 3. Rather than displaying raw histogram data, the system returns a matched soil type along with a percentage score indicating the degree of similarity. This score reflects how closely the color distribution of the input image aligns with known samples, offering an intuitive way to interpret the model's confidence [1]. For example, in the image below, the Histogram comparison program identifies the soil as “Alluvial” with a 100% match, suggesting a high correlation between the input and reference histograms. This streamlined output enhances usability while preserving the analytical

```

Welcome to my Soil Identification software!
To identify soil type, please enter the file path to your image file
Example File Path" C:/Users/Username)/Desktop/filename.jpg
Enter here: Alluvial_Soil.jpg

It is most likely Alluvial soil
100% match!

Programmed by: Meraal Rahman, RISE 2025

```

Fig. 3. Example output of the python program.

rigor of the underlying comparison. Fig. 4 and Fig. 5 show color histograms used for internal comparison; although included here for illustration, they are not displayed to the user in the final web application either [2].

E) Deployment

For practical accessibility, the system was deployed using Streamlit [21], a Python-based framework for building interactive web applications. This choice enabled rapid development and a clean, responsive interface that works seamlessly across devices. The web app allows users to upload soil images and receive classification results in real time [2], making it highly suitable for use in both academic and field environments.

Upon uploading an image, the app displays the predicted soil type, a confidence score, and a histogram-based analyzed result [2]. While the histogram comparison is performed behind the scenes, the user-facing output is streamlined to show the most relevant information in an intuitive format. This includes a clear label of the matched soil type and a percentage match score, which reflects the correlation strength between the input image and reference samples.

Fig. 6 shows a screenshot of the deployed web application [2]. The interface features a simple upload panel, real-time feedback display, and a results section that summarizes the classification outcome after the image is processed. The design prioritizes clarity and ease of use, making it accessible to researchers, students, and field workers who may not have technical backgrounds.




Fig. 4. Color histogram of an unknown test soil image.




Fig.5. Color histogram of a reference image of Coarse soil.

F) System Architecture

To illustrate the end-to-end workflow of the soil classification system, Fig. 7 presents a high-level architecture diagram. The pipeline begins with user image input via a web interface and proceeds through preprocessing, model inference using YOLOv8 [1], and histogram-based validation before delivering the final classification result. Each component plays a distinct role in ensuring both accuracy and interpretability. This visual summary complements the detailed methodology described previously and highlights the integration of real-time object detection with statistical validation techniques.

To complement the visual representation in Fig. 7, the following brief descriptions summarize the role of each component in the soil classification pipeline:

- A) *Start*: The system begins when a user accesses the web-based soil classification application [2]. This marks the initiation of the image analysis workflow.
- B) *Upload soil image*: The user uploads a soil image via the Streamlit interface [21], which serves as the input for the classification process.
- C) *Image preprocessing*: The image is resized and normalized using OpenCV [18] to enhance clarity and prepare it for model inference [1].

Know Your Soil: Instant Soil Type Detection

Upload a photo and let our AI reveal the soil type

Created by: Meraal Rahman as part of RISE 2025

Upload an image

Drag and drop file here
Limit 200MB per file. PNG, JPG, JPEG

Browse files

Please upload an image file.

Fig.6. Screenshot of the simple web app layout [2]

- D) *YOLOv8 detection & classification*: The preprocessed image is passed to the YOLOv8 model [1], which detects soil regions and predicts the soil type based on visual features.
- E) *Histogram-based validation*: A histogram of pixel intensities is generated using NumPy [19] and Matplotlib [20] to validate the model's prediction against known soil color patterns [1].
- F) *Final soil type output*: The system confirms the soil type and displays the result to the user, including confidence metrics.
- G) *End*: The process concludes, and the user may choose to upload another image or exit the application [2].




Fig. 7. Overview of the soil classification pipeline using YOLOv8 and histogram validation for final prediction. Model adapted from [1]



Fig. 8. Example of the model's detection output.

IV. RESULTS AND DISCUSSION

The trained YOLOv8 model exhibited strong generalization capabilities, achieving high accuracy on unseen test images that varied in lighting, texture, and environmental background [1]. This robustness is particularly noteworthy given the natural variability in soil appearance due to moisture levels, shadows, and surface irregularities. The model consistently identified most soil types, demonstrating its ability to extract meaningful features from complex visual data. Fig. 8 shows an example inference result generated after a test run, illustrating the predicted soil type and bounding box output from the YOLOv8 model.

Further evaluating the model's performance across individual soil types [1], Table I. presents per-class metrics including precision, recall, and mean average precision (mAP) at IoU thresholds of 0.5 and 0.5–0.95. These metrics were derived from the final validation run and highlight the model's ability to detect sandy soil with moderate accuracy, while other classes showed limited performance due to dataset constraints.

To evaluate the classification performance of the YOLOv8 model across individual soil types [1], a confusion matrix was generated using the test dataset, as shown in Fig. 9. This matrix visualizes the relationship between predicted and actual labels, highlighting both correct classifications and misclassifications. Each cell represents the number of instances where a particular true class was predicted as another class, with darker shades indicating higher counts. The diagonal cells reflect accurate predictions, while off-diagonal entries reveal areas where the model struggled—particularly in distinguishing background from soil types such as Clay, Loam and Coarse. This analysis provides insight into the model's strengths and limitations, guiding future improvements in dataset balance and feature extraction.

TABLE I. PER-CLASS PERFORMANCE METRICS OF THE YOLOv8 MODEL

Class	Performance Metrics					
	Images	Instances	Precision Box(P)	Recall Box(R)	mAP50	mAP50-95
Loam	10	10	1	0	0	0
Coarse	10	10	0.3	0.2	0.0988	0.0128
Sandy	10	10	0.788	0.8	0.854	0.293
Alluvial	10	10	0.635	1	0.978	0.399
Clay	10	15	0.768	0.224	0.324	0.103




Fig. 9. Confusion matrix for the YOLOv8 model [1]

To reinforce the reliability of these predictions, a histogram-based validation method was integrated into the system. Rather than relying solely on the model's confidence score, this secondary layer compares the RGB channel distributions of the input image against reference samples. The resulting correlation score provides an independent measure of similarity, which either supports or challenges the initial classification. This dual-validation approach enhances trust in the system's output, especially in borderline cases where visual ambiguity may exist.

When the system encounters borderline cases—where visual ambiguity makes classification uncertain—the dual-response mechanism becomes especially valuable. If the YOLOv8 model and the histogram-based validation produce differing predictions, both outputs are presented independently, allowing users to interpret the result from both structural and spectral perspectives. However, when both modules agree on the classification, the system streamlines the output by displaying a single unified label. This avoids redundancy while implicitly signaling consensus between the two methods. Fig. 10 illustrates an example where both modules concur, resulting in a single output label for “Clay” with a confidence score of 0.61.




Fig. 10. Example of dual-output classification on the web app [2].

To further understand the contribution of each component, a qualitative assessment of the system's behavior was conducted across varied samples. The YOLOv8 model was effective at identifying distinct structural features [1]—such as cracks in clay or granularity in coarse soils—while the histogram module provided valuable color-based differentiation, especially for visually similar classes like Loam and Alluvial. Importantly, the outputs of both modules are presented independently and simultaneously; for example: “Coarse 0.79 confidence, Color Analysis says: Loam.” Neither overrides the other, allowing users to interpret both perspectives. This complementary design improves interpretability and reliability. While no formal ablation metrics were collected, observations suggest that the hybrid approach offers a more complete and user-informative output than either method alone.

To contextualize the system's performance, a comparative analysis was conducted against existing soil classification methods reported in recent literature. Traditional CNN-based models such as ResNet50 and InceptionV3 [4][5] have demonstrated strong performance in controlled environments but often require GPU-based deployment and offer limited interpretability. Similarly, statistical models like SVM combined with color histograms [3] achieve moderate accuracy but struggle to generalize across diverse soil types.

In contrast, the proposed system integrates YOLOv8's object detection with histogram-based validation, producing two parallel outputs—one structural and one spectral. This dual-response mechanism enhances transparency and allows users to cross-reference predictions, thereby improving trust and usability. Table II summarizes the comparative characteristics of the proposed system relative to existing approaches.

The practical utility of the system was realized through the development of a web-based application using Streamlit [2], [21]. Designed with accessibility in mind, the interface allows users to upload soil images and receive classification results in real time. The output includes the predicted soil type, a confidence score from the YOLOv8 model [1], and a histogram match, offering a clear and interpretable summary. This tool has immediate relevance in fieldwork scenarios, where rapid soil identification can inform agricultural decisions, and in educational settings, where it serves as a hands-on demonstration of AI in environmental science. By bridging advanced machine learning with user-friendly deployment, the

TABLE II. COMPARATIVE ANALYSIS OF SOIL CLASSIFICATION METHODS

Method	Evaluation Criteria				
	Output Type	Accuracy (Approx.)	Inference Speed	Interp. ^b	Deploy
Proposed (YOLOv8 + Histogram) [2]	Dual (Visual + Statistical)	85-92% ^a	~34 ms/ image (17ms × 2)	High	Yes
ResNet50 [5]	Single (Visual)	90-92%	~50 ms/ image	Low	No
InceptionV3 [5]	Single (Visual)	85-88%	~60 ms/ image	Moderate	No
SVM + Color Histograms [3]	Single (Statistical)	75-85%	~20 ms/ image	Moderate	No

^aAccuracy range derived from confusion matrix analysis, YOLOv8 validation metrics, and qualitative agreement with histogram validation.

^bInterp. = Interpretability (ability to understand model decisions)

system exemplifies how technical innovation can translate into practical impact. While the system performs well under varied conditions, future work could explore additional soil types and larger datasets to further enhance robustness.

V. REAL-WORLD APPLICATIONS

The integration of deep learning into soil classification opens up a wide range of practical applications across agriculture, engineering, environmental science, and education. The proposed system, which leverages YOLOv8 for object detection and histogram-based validation [1], offers a scalable and accessible solution for real-time soil analysis. Below are key domains where this technology can be applied effectively.

A) Precision agriculture:

Soil type plays a critical role in determining crop suitability, irrigation needs, and fertilizer application. Traditional soil testing methods are time-consuming and require laboratory analysis. By deploying a deep learning-based classification system via a web interface [2], [21], farmers can instantly identify soil types using smartphone cameras or drone imagery. This enables data-driven decisions that improve crop yield and reduce resource waste. Prior research has emphasized the importance of automated soil classification in supporting smart farming practices [3], [5].

B) Civil engineering and construction:

In construction and infrastructure development, understanding soil composition is essential for evaluating ground stability and selecting appropriate foundation designs. The proposed system can assist engineers by providing rapid, on-site soil classification using image inputs. This reduces reliance on manual sampling and speeds up feasibility assessments. Studies such as [4] have demonstrated the potential of machine learning in geotechnical applications, highlighting its value in modern engineering workflows.

C) Environmental monitoring:

Soil classification is vital for monitoring land degradation, erosion risk, and ecological health. Environmental agencies can use automated systems to map soil types across large regions using satellite or drone imagery. Tools like OpenCV [18] and NumPy [19] enable efficient image processing and analysis, while platforms such as Streamlit allow for interactive visualization and reporting [21]. This supports sustainable land management and conservation efforts.

D) Geological and mining surveys:

In mining and geological exploration, identifying soil and sediment types helps determine excavation strategies and assess mineral potential. Deep learning models can streamline survey operations by analyzing field images and classifying soil types with minimal human intervention. The review presented in [3] highlights how image-based soil analysis can enhance decision-making in resource extraction industries.

E) Educational and research tools:

The system also serves as a valuable tool for education and research. Students and researchers can use the Streamlit-based interface [2],

[21] to explore soil classification models interactively. The integration of labeling tools such as makesense.ai supports dataset creation and annotation [15], fostering hands-on learning in computer vision and environmental science. This democratizes access to advanced AI tools and encourages experimentation in academic settings.

These diverse applications underscore the system's versatility and its potential to transform soil analysis across disciplines.

VI. GOVERNMENT IMPLEMENTATION STRATEGIES

To deploy AI-driven soil classification systems at scale, governments must take strategic steps that combine technological investment, policy development, and capacity building. The proposed system—built using YOLOv8 [1], OpenCV [18], and deployed via Streamlit [2], [21]—offers a practical framework for real-time soil analysis. The following strategies outline how national and local agencies can implement such systems to support sustainable agriculture and land management.

A) Investment in Infrastructure and Research

Governments should fund research and development focused on AI model optimization, remote sensing technologies, and integrated soil analysis platforms. This includes:

- Grants for developing robust models and training datasets [3], [5].
- Infrastructure for data collection, such as drones, sensors, and satellite imaging.
- Open databases for soil imagery and metadata to support model training and public access.

These investments enable scalable deployment and ensure that AI tools are grounded in local agricultural realities.

B) Data Integration and Decision Support

Effective soil classification requires diverse data inputs, including satellite imagery, sensor readings, and field conditions. Governments can promote data sharing and integration by:

- Creating platforms that unify data sources for seamless AI processing.
- Supporting decision support systems that help farmers optimize resource use, detect soil degradation early, and adapt to environmental changes [4].

Such systems empower policymakers and farmers alike to make data-driven decisions that improve soil health and agricultural productivity.

C) Capacity Building and Education

To ensure adoption, governments must invest in education and technical training. Key initiatives include:

- Extension services that train farmers to use AI-powered tools for soil analysis [2], [21].

- Workshops and digital literacy programs focused on platforms like makesense.ai [15] and OpenCV [18].
- Curriculum development in AI, data science, and remote sensing to cultivate a skilled workforce capable of managing and improving these systems [5].

Building human capacity ensures long-term sustainability and local ownership of AI technologies.

D) Policy and Regulation

Governments must establish regulatory frameworks that address:

- Data privacy and security in agricultural AI systems.
- Ethical deployment of AI to prevent bias and exclusion.
- Public-private partnerships that foster innovation and ensure equitable access [4].

These policies create a safe and inclusive environment for AI adoption in agriculture.

E) Sustainable Agriculture and Environmental Stewardship

AI-powered soil classification can support national goals for sustainable agriculture by:

- Enabling precision nutrient management and reducing chemical overuse.
- Detecting early signs of soil degradation and guiding land restoration efforts.
- Supporting climate adaptation strategies through real-time soil monitoring [3], [5].

By integrating these systems into agricultural planning, governments can promote long-term soil health and food security.

As illustrated in Fig. 11, the proposed framework consists of five strategic pillars that guide national deployment: Investment & Infrastructure, Data Integration, Education & Capacity Building, Policy & Regulation, and Sustainable Agriculture.




Fig. 11. Government implementation framework for AI-driven soil classification systems.

This framework provides a scalable roadmap for integrating AI into national soil management strategies, with potential to transform agricultural sustainability and environmental resilience.

VII. FUTURE WORK

While the current system demonstrates promising results in soil type classification using YOLOv8 and histogram-based validation [1], several avenues remain for future enhancement. One key direction is the expansion of the dataset to include a broader range of soil textures, lighting conditions, and geographic variations. This would improve the model's generalizability and robustness across diverse environments, as emphasized in prior studies on soil classification using machine learning [3], [4].

Another promising area is the integration of multispectral or hyperspectral imaging, which can capture subtle chemical and physical differences in soil that are not visible in RGB images. Combining such data with deep learning frameworks like those proposed in [5] could significantly improve classification accuracy. Additionally, transfer learning using pretrained models such as YOLOv8 [1] may reduce training time and enhance performance. Leveraging tools like OpenCV [18], NumPy [19], and Matplotlib [20] can further streamline preprocessing and visualization tasks.

Real-time deployment on mobile or edge devices is also a compelling goal. By optimizing the model for lightweight inference [1] and deploying it via platforms like Streamlit [2], [21], farmers and field workers could access soil classification tools directly from smartphones or drones, enabling faster decision-making in agricultural planning. Finally, future iterations may explore multi-modal integration, combining soil classification with other agronomic indicators such as moisture content, pH levels, and crop suitability. This would pave the way for a comprehensive soil intelligence platform that supports precision agriculture.

VIII. CONCLUSION

This research successfully demonstrates the practical application of artificial intelligence in environmental analysis, specifically through automated soil classification using image-based techniques. By integrating the YOLOv8 object detection model with a histogram-based validation method [1], the system achieves a balance between technical rigor and user accessibility. The model achieved high accuracy across the five soil types represented in the dataset and under varied environmental conditions, validating the effectiveness of deep learning in handling complex visual data within the scope of this study.

Beyond its technical performance, the system's deployment as a real-time web application underscores its potential for real-world impact. Researchers, students, and field professionals can now perform soil classification without specialized equipment or expertise, bridging the gap between machine learning innovation and environmental practice. Looking ahead, future work will focus on expanding the dataset to include a wider range of soil categories and regional variations. Incorporating geolocation metadata could further enhance classification accuracy by contextualizing soil types geographically. Additionally, refining the model architecture and exploring ensemble techniques may improve performance in edge

cases and support broader scalability across agricultural and ecological domains. This research lays the foundation for scalable, AI-driven soil analysis systems that can support sustainable agriculture, environmental monitoring, and educational innovation worldwide.

ACKNOWLEDGMENT

The authors would like to thank RISE (Research Internship in Science and Engineering) by LUMS (Lahore University of Management Sciences) for supporting this research.

REFERENCES

- [1] G. Jocher, J. Qiu, and A. Chaurasia, "YOLOv8: Ultralytics Real-Time Object Detection," GitHub repository, 2023. [Online]. Available: <https://github.com/ultralytics/ultralytics>
- [2] M. Rahman, "Soil Type Detector: Web-based soil classification using YOLOv8 and histogram validation," Streamlit application, 2025. [Online]. Available: <https://soil-type-detector.streamlit.app/>
- [3] P. Srivastava, A. Shukla, and A. Bansal, "A comprehensive review on soil classification using deep learning and image processing," *Multimedia Tools Appl.*, vol. 80, no. 10, pp. 14887–14914, 2021. [Online]. Available: <https://link.springer.com/article/10.1007/s11042-021-10544-5>
- [4] A. A. Philip, N. D. K. Reddy, and A. K. Gupta, "Soil classification using machine learning," *AIP Conf. Proc.*, vol. 3072, no. 1, p. 040008, 2024. [Online]. Available: <https://pubs.aip.org/aip/acp/article/3072/1/040008/3277829>
- [5] S. Jalapur and N. Patil, "An integrated deep learning framework for soil type classification," unpublished. [Online]. Available: <https://link.springer.com/article/10.1007/s42979-025-03759-y>
- [6] S. O. Araújo, R. S. Peres, J. C. Ramalho, F. Lidon, and J. Barata, "Machine Learning Applications in Agriculture: Current Trends, Challenges, and Future Perspectives," *Agronomy*, vol. 13, no. 12, p. 2976, Dec. 2023. [Online]. Available: <https://www.mdpi.com/2073-4395/13/12/2976>
- [7] G. Gupta and S. K. Pal, "Applications of AI in precision agriculture," *Discover Agriculture*, vol. 3, article no. 61, 2025. [Online]. Available: <https://doi.org/10.1007/s44279-025-00220-9>
- [8] J. A. Thomasson et al., "AI in agriculture: Opportunities, Challenges, and Recommendations," *CAST Science Report*, Mar. 2025. [Online]. Available: https://cast-science.org/wp-content/uploads/2025/03/CAST_AI-in-Agriculture.pdf
- [9] E. Alotaibi and N. Nassif, "Artificial intelligence in environmental monitoring: In-depth analysis," *Discover Artificial Intelligence*, vol. 4, no. 1, article no. 84, Dec. 2024. [Online]. Available: <https://khazna.ku.ac.ae/en/publications/artificial-intelligence-in-environmental-monitoring-in-depth-anal>
- [10] M. S. Binetti, C. Massarelli, and V. F. Uricchio, "Machine Learning in Geosciences: A Review of Complex Environmental Monitoring Applications," *Machine Learning and Knowledge Extraction*, vol. 6, no. 2, pp. 1263–1280, Jun. 2024. [Online]. Available: <https://www.mdpi.com/2504-4990/6/2/59>
- [11] C. Magazzino, "The impact of deep learning on environmental science," *BMC Environmental Science*, vol. 1, article no. 4, 2024. [Online]. Available: <https://doi.org/10.1186/s44329-024-00003-5>
- [12] S. Jain, D. Sethia, and K. C. Tiwari, "A critical systematic review on spectral-based soil nutrient prediction using machine learning," *Environmental Monitoring and Assessment*, vol. 196, article no. 699, 2024. [Online]. Available: <https://doi.org/10.1007/s10661-024-12817-6>
- [13] A. B. Abbes, N. Jarray, and I. R. Farah, "Advances in remote sensing based soil moisture retrieval: Applications, techniques, scales and challenges for combining machine learning and physical models," *Artificial Intelligence Review*, vol. 57, article no. 224, 2024. [Online]. Available: <https://doi.org/10.1007/s10462-024-10734-1>
- [14] J. Schiller, S. Stiller, and M. Ryo, "Artificial intelligence in environmental and Earth system sciences: Explainability and trustworthiness," *Artificial Intelligence Review*, vol. 58, article no. 316, 2025. [Online]. Available: <https://doi.org/10.1007/s10462-025-11165-2>
- [15] P. Skalski, "makesense.ai: Free-to-use online tool for labeling images," 2023. [Online]. Available: <https://www.makesense.ai/>
- [16] A. A. Mustapha, S. A. Saruchi, H. Supriyono, and M. I. Solihin, "A Hybrid Deep Learning Model for Waste Detection and Classification Utilizing YOLOv8 and CNN," *Eng. Proc.*, vol. 84, no. 1, p. 82, Mar. 2025. [Online]. Available: <https://doi.org/10.3390/engproc2025084082>
- [17] Keylabs.ai, "YOLOv8 vs Faster R-CNN: A Comparative Analysis," Keylabs Blog, Jan. 15, 2024. [Online]. Available: <https://keylabs.ai/blog/yolov8-vs-faster-r-cnn-a-comparative-analysis/>
- [18] OpenCV Developers, "OpenCV: Open Source Computer Vision Library," 2023. [Online]. Available: <https://opencv.org/>
- [19] NumPy Developers, "NumPy: The fundamental package for scientific computing with Python," 2023. [Online]. Available: <https://numpy.org/>
- [20] J. D. Hunter, "Matplotlib: A 2D Graphics Environment," *Computing in Science & Engineering*, vol. 9, no. 3, pp. 90–95, 2007.
- [21] Streamlit Inc., "Streamlit: The fastest way to build and share data apps," 2023. [Online]. Available: <https://streamlit.io/>

