

Fusing Environmental Sensing and Computer Vision: A Machine Learning Pipeline for Soil Nutrient and Moisture Mapping From Hyperspectral Imagery

V.Pushpalatha¹, Dr.T. Kamaleshwar²

¹Research scholar, Department of computer science, Vel Tech Rangarajan Dr.Sagunthala R&D Institute of Science and Technology, Chennai, Tamilnadu, India. espushpa7.6.93@gmail.com

²Associate Professor, Department of Computer Science & Engineering, Vel Tech Rangarajan Dr.Sagunthala R&D Institute of Science and Technology, Chennai, Tamil Nadu, India. kamalesh4u2@gmail.com

ABSTRACT:

Traditional soil nutrient and moisture analysis is labor-intensive, costly, and lacks spatial granularity, hindering precision agriculture. This paper proposes an integrated machine learning pipeline that fuses hyperspectral imagery (HSI) a powerful environmental sensing technology with advanced computer vision techniques to address this gap. We hypothesize that deep learning models can decode the complex, non-linear spectral signatures in HSI data to predict key soil properties accurately. Our methodology encompasses HSI data preprocessing, feature extraction using a Convolutional Neural Network (CNN), and regression modeling. Using a public dataset, we demonstrate that our proposed CNN-based model outperforms traditional spectral indices and machine learning models like Support Vector Regression (SVR) in predicting soil organic carbon (SOC) and moisture content. The results indicate the high potential of this pipeline for generating high-resolution, actionable soil maps to optimize resource use in agriculture.

Keywords: Precision Agriculture, Hyperspectral Imaging, Deep Learning, Convolutional Neural Networks, Soil Mapping, Nutrient Management.

1. INTRODUCTION

Soil health, characterized by its nutrient content (e.g., Nitrogen-N, Phosphorus-P, Potassium-K, Organic Carbon-SOC) and moisture, is a fundamental determinant of agricultural productivity and environmental sustainability [1]. Conventional soil testing methods are destructive, time-consuming, and provide only point-based data, making it impossible to visualize the significant spatial variability present within a single field [2]. This limitation is a major barrier to the adoption of precision agriculture, which aims to apply inputs (water, fertilizer) at variable rates based on precise spatial data.

The advent of remote sensing, particularly hyperspectral imaging (HSI), has opened new avenues for non-destructive, large-scale soil monitoring. HSI sensors capture the reflectance of materials across hundreds of narrow, contiguous spectral bands, creating a detailed spectral fingerprint [3]. This fingerprint is influenced by molecular absorption features, which are directly related to soil constituents like organic matter, water, and certain minerals [4].

While the connection between spectra and soil properties is established, the relationship is highly complex and non-linear. Traditional methods like spectral indices (e.g., Normalized Difference Water Index - NDWI) or linear models often fail to capture this complexity. This paper proposes a novel machine learning pipeline that leverages computer vision, specifically deep learning, to automatically learn these intricate patterns from HSI data for accurate prediction and mapping of soil nutrients and moisture.

2. LITERATURE REVIEW

Previous research has established the viability of HSI for soil characterization. For instance, [5] used partial least squares regression (PLSR) to predict clay and SOC content from airborne HSI with moderate success. Similarly, [6] demonstrated that specific wavelength regions in the visible-near-infrared (VNIR) and short-wave infrared (SWIR) are sensitive to soil moisture changes.

Machine learning models like Support Vector Machines (SVM) and Random Forests (RF) have shown improvement over linear models. [7] used an SVM to classify soil types using hyperspectral data, while [8] applied RF to estimate nitrogen content.

However, these models often rely on manual feature selection or pre-defined indices. Deep learning models, particularly CNNs, excel at automatically learning hierarchical features directly from raw or minimally processed data. Their application in agriculture is growing [9], but their use for direct soil

property prediction from hyperspectral data is still an emerging field. [10] used a 1D-CNN on spectral signatures to estimate SOC, showing superior performance over PLSR. Our work aims to extend this by designing an end-to-end pipeline that processes georeferenced HSI cubes, emphasizing spatial-spectral feature learning for precise geospatial mapping.

3. PROPOSED METHODOLOGY

The proposed end-to-end machine learning pipeline is designed to transform raw hyperspectral imagery into high-resolution predictive maps for soil nutrients and moisture. The architecture, depicted in Figure 1, consists of four integrated stages: (1) Data Acquisition & Preprocessing, (2) Ground Truth Integration, (3) Core CNN Model Architecture, and (4) Geospatial Prediction & Mapping.

Based on the illustrated pipeline, this figure 1 outlines an end-to-end machine learning workflow for predicting soil properties from hyperspectral imagery. The process begins with raw data acquisition and preprocessing, including essential corrections such as atmospheric adjustment (e.g., FLAASH [12]) and vegetation masking (e.g., NDVI). Lab-analyzed soil samples provide ground truth labels, which are used to train a 1D convolutional neural network (1D-CNN) architecture. This model extracts features from spectral signatures and performs regression to estimate target variables like soil organic carbon (SOC). Finally, predictions are mapped geospatially to produce high-resolution soil property maps, supporting precision agriculture applications [1, 10].




Figure 1: Proposed End-to-End Machine Learning Pipeline for Soil Property Mapping

3.1. Data Acquisition and Preprocessing

The initial stage involves preparing the raw hyperspectral data for analysis.

3.1.1. Data Acquisition: Hyperspectral imagery is captured using airborne (e.g., UAV-mounted) or spaceborne sensors (e.g., PRISMA, EnMAP). The data is collected as a three-dimensional (3D) "data cube," with two spatial dimensions (x, y) and one spectral dimension (λ), comprising hundreds of narrow, contiguous bands (e.g., from 400 nm to 2500 nm).

3.1.2. Preprocessing Chain: The raw data cube undergoes a series of corrections to convert sensor readings into accurate surface reflectance values, which are crucial for reproducible model training.

Radiometric Correction: Converts digital numbers (DNs) to at-sensor radiance values using calibration coefficients specific to the sensor.

Atmospheric Correction: Compensates for atmospheric absorption and scattering effects (e.g., water vapor, aerosols). We will employ a physics-based model like the **Fast Line-of-sight Atmospheric Analysis of Spectral Hypercubes (FLAASH)** algorithm to derive surface reflectance [12].

Geometric Correction & Registration: Corrects for spatial distortions due to sensor motion and terrain relief. The imagery is co-registered with high-precision GPS coordinates of ground truth sampling points.

Masking: Non-soil pixels (e.g., vegetation, water, man-made structures) are identified and masked out using spectral indices like the **Normalized Difference Vegetation Index (NDVI)** to ensure the model trains exclusively on pure soil spectra.

3.2. Ground Truth Data Integration

Geo-located soil samples are collected concurrently with the HSI flight campaign. These samples are analyzed in a laboratory using standard procedures (e.g., Walkley-Black method for SOC, Gravimetric method for moisture) to obtain precise reference measurements. These values form the ground truth labels (Y) for the supervised learning task. Each labeled sample is paired with its corresponding pixel spectrum from the preprocessed HSI data, creating the dataset (X_i, Y_i) , where X_i is a vector of reflectance values across all bands.

3.3. Core 1D-CNN Model Architecture

We propose a 1D-Convolutional Neural Network (1D-CNN) architecture to model the continuous spectral signature of each pixel. This design is chosen for its superior ability to automatically extract hierarchical spectral features without manual intervention.

The proposed 1D-CNN architecture begins by accepting a preprocessed spectral vector, typically comprising approximately 200 reflectance values after the removal of noisy bands. Input spectra are first normalized using techniques such as Standard Normal Variate (SNV) to mitigate scaling variations and improve convergence during training. The core feature extraction blocks consist of multiple sequential 1D convolutional layers equipped with small kernel sizes (e.g., 3–5) to detect local spectral absorption features and their complex interactions. Each convolutional layer is followed by a Rectified Linear Unit (ReLU) activation function to introduce non-linearity, and batch normalization is applied to enhance training stability. Subsequently, 1D max-pooling layers reduce spectral dimensionality and increase feature invariance. The flattened features are then passed through fully connected dense layers for high-level regression reasoning. The final output layer uses a linear activation function to produce a continuous estimate of the target soil property, such as SOC percentage. The model is trained end-to-end by minimizing the mean squared error (MSE) between predictions and laboratory-measured values using the Adam optimization algorithm.




Figure 2 : Hyperspectral Data Processing Workflow with a 1D-CNN Architecture

This figure 2 illustrates a comprehensive, end-to-end machine learning pipeline for predicting soil properties from hyperspectral imagery. The workflow begins with Data Acquisition & Preprocessing, where raw hyperspectral data undergoes critical corrections, such as atmospheric adjustment and vegetation masking. Following this, Ground Truth Integration is performed, using lab-analyzed soil

samples to provide the labels needed for model training. The core of the pipeline is the Core CNN Model Architecture, a specialized 1D-CNN that processes the spectral signatures to extract features and perform regression to estimate target variables like soil organic carbon (SOC). The final stage, Geospatial Prediction & Mapping, utilizes the trained model to generate high-resolution predictive maps of soil properties, which are essential for precision agriculture and sustainable land management.

3.4. Geospatial Prediction and Mapping

Once trained and validated, the model is deployed on the entire preprocessed HSI scene. The spectrum of every single soil pixel is fed into the model, generating a predicted value for the target property. These predictions are then compiled into a new 2D geospatial raster layer, where the value of each pixel represents the estimated soil property concentration. This final output is a high-resolution map that can be imported into a Geographic Information System (GIS) for v.

4. DATASET AND DATA DESCRIPTION

To validate the proposed methodology, this study utilizes the LUCAS 2018 Topsoil Hyperspectral Dataset, a large-scale, publicly available, and well-curated dataset provided by the European Soil Data Centre (ESDAC) [11, 13]. The Land Use/Cover Area frame Statistical survey (LUCAS) is a pan-European program designed to monitor changes in land cover and soil properties. The 2018 topsoil survey is particularly valuable as it includes both comprehensive laboratory analyses and corresponding hyperspectral data, making it an ideal benchmark for this research.

4.1. Data Acquisition and Sources

The LUCAS 2018 topsoil dataset is the result of a standardized campaign across the European Union.

Soil Sampling: Approximately 20,000 topsoil samples (0-20 cm depth) were collected from representative locations using a standardized methodology to ensure consistency [13]. Each sample was precisely geolocated using GPS.

Laboratory Analysis: The collected soil samples were analyzed in a single dedicated laboratory for a wide range of physico-chemical properties. For this study, the following key analytes are used as ground truth labels:

Soil Organic Carbon (SOC): Measured using the dry combustion method (elemental analysis), expressed in g/kg or %.

Soil Moisture Content: Calculated gravimetrically by weighing before and after drying at 105°C, expressed as a percentage of the soil's dry weight.

Additional Nutrients: Data for Nitrogen (N) content and pH are also available for multi-task learning extensions.

Hyperspectral Data Acquisition: The hyperspectral data for the LUCAS 2018 topsoil samples were measured in a laboratory setting under controlled conditions using a FOSS XDS Near Infrared Spectrophotometer. This ensures a high signal-to-noise ratio and avoids atmospheric interference present in airborne imagery. Each sample was scanned to obtain its reflectance spectrum across the 400–2500 nm wavelength range at a fine spectral resolution, resulting in a vector of 4200 reflectance values per sample [11].

4.2. Data Characteristics and Preprocessing

The raw spectral data from the LUCAS dataset required specific preprocessing to match the input requirements of the proposed 1D-CNN model.

Spectral Range: The full spectrum (400-2500 nm) was utilized. Noisy bands at the extremes of the detector ranges (particularly around 400-450 nm and 2350-2500 nm) and regions with strong water vapor absorption (e.g., around 1400 nm and 1900 nm) were removed, resulting in a refined set of approximately 2000 relevant bands for model input.

Data Splitting: The dataset was randomly partitioned into three subsets to ensure robust evaluation:

Training Set (70%): Used to train the 1D-CNN model parameters.

Validation Set (15%): Used for hyperparameter tuning and to monitor for overfitting during training.

Test Set (15%): A held-out set used only for the final evaluation to report unbiased performance metrics (RMSE, R^2).

Spectral Preprocessing: Each spectral curve underwent Standard Normal Variate (SNV) normalization. This technique scales each individual spectrum to have a mean of zero and a standard deviation of one, effectively removing multiplicative effects of scatter and particle size, which is a common preprocessing step in spectroscopy [14].

Property	Number of Samples	Key Variable Range	Measurement Method
Spectral Data	~20,000	400 - 2500 nm	Lab-based NIR Spectrometry
Soil Organic Carbon (SOC)	~20,000	0.01 - 40.2 %	Dry Combustion
Moisture Content	~20,000	1.5 - 65.8 %	Gravimetric
pH	~20,000	3.2 - 9.5	CaCl ₂ extraction
Total Nitrogen (N)	~20,000	0.01 - 5.6 g/kg	Dry Combustion

Table 1: Summary of the LUCAS 2018 Topsoil Dataset Used in This Study

The Table 1 analysis in this study is based on the extensively recognized LUCAS 2018 Topsoil Dataset, which comprises approximately 20,000 soil samples collected from across the European Union [13]. Each sample includes a high-resolution hyperspectral reflectance spectrum measured under laboratory conditions across the 400–2500 nm range using NIR spectrometry, ensuring high signal quality and minimal environmental noise. Corresponding ground truth measurements for key soil properties were obtained through standardized laboratory methods: Soil Organic Carbon (SOC) and Total Nitrogen were analyzed via dry combustion, Moisture Content was determined gravimetrically, and pH was measured using CaCl₂ extraction. The wide and ecologically relevant ranges of these variables (e.g., SOC: 0.01–40.2%; Moisture: 1.5–65.8%) make this dataset highly suitable for training and evaluating robust machine learning models aimed at predicting soil attributes from spectral data [11]. This large sample size and methodological consistency support the generalizability of the proposed 1D-CNN model.

5. IMPLEMENTATION AND RESULTS

5.1. Experimental Setup and Implementation

The proposed 1D-CNN model was implemented using the TensorFlow and Keras deep learning frameworks. All experiments were conducted on a high-performance computing node equipped with an NVIDIA Tesla V100 GPU with 32GB memory. The model was trained for 200 epochs with a batch size of 64. The Adam optimizer was employed with an initial learning rate of 0.001, which was reduced by a factor of 0.5 if the validation loss plateaued for 10 consecutive epochs. Early stopping was implemented with a patience of 15 epochs to prevent overfitting.

The architecture consisted of three 1D convolutional layers with 64, 128, and 256 filters respectively, each with a kernel size of 5 and followed by ReLU activation and batch normalization. A max-pooling layer with pool size 2 was added after each convolutional block. The flattened features were passed through two dense layers (128 and 64 units) before the final output layer.

Parameter	Value/Description
Framework	TensorFlow 2.8, Keras
Input Dimensions	2000 (after band selection)
Convolutional Layers	3 (64, 128, 256 filters)
Kernel Size	5
Pooling	MaxPooling1D (size 2)
Dense Layers	128, 64 units
Optimizer	Adam (lr=0.001)
Loss Function	Mean Squared Error (MSE)

Table 2: Model Configuration and Hyperparameter

5.2. RESULTS AND PERFORMANCE ANALYSIS

The proposed 1D-CNN model was evaluated against three benchmark methods: Partial Least Squares Regression (PLSR), Support Vector Regression (SVR) with radial basis function kernel, and Random Forest (RF) regression. Performance was assessed using Root Mean Square Error (RMSE) and Coefficient of Determination (R²).

Model	RMSE	R ²	Training Time (min)
PLSR	2.89	0.72	0.5
SVR	2.45	0.78	12.3
RF	2.12	0.82	8.7

1D-CNN (Ours)	1.58	0.89	23.5
---------------	------	------	------

Table 3: Performance Comparison for SOC Prediction

The performance comparison for Soil Organic Carbon (SOC) prediction, as detailed in Table 3, demonstrates the clear superiority of the proposed 1D-CNN model over established benchmark methods. While Partial Least Squares Regression (PLSR) achieved a baseline performance (RMSE = 2.89, R^2 = 0.72), more advanced machine learning models like Support Vector Regression (SVR) and Random Forest (RF) showed improved accuracy, with RF attaining an R^2 of 0.82. However, the 1D-CNN architecture significantly outperformed all others, achieving a markedly lower RMSE of 1.58 and the highest explained variance with an R^2 of 0.89. This substantial improvement in predictive accuracy, though accompanied by a longer training time (23.5 minutes), underscores the model's enhanced capability to capture the complex, non-linear relationships within hyperspectral data, making it a highly effective tool for precise SOC quantification.




Figure 3: Comparative Performance Evaluation of Machine Learning Models for Soil Property Prediction
Figure 3 visually summarizes the comparative performance of machine learning models in predicting soil organic carbon (SOC) from hyperspectral data, illustrating the critical trade-off between computational efficiency and predictive accuracy. The proposed 1D-CNN model (ours) demonstrates a substantial reduction in RMSE and the highest R^2 value among all methods, confirming its superior capability in decoding complex spectral-soil relationships [10]. While traditional methods like PLSR offer minimal training time, their limited accuracy (higher RMSE, lower R^2) restricts practical utility for precision agriculture. The enhanced performance of the 1D-CNN aligns with recent advances in deep learning for hyperspectral analysis [9], validating its use cases where prediction fidelity such as generating high-resolution nutrient maps for variable-rate applications outweighs computational overhead [1, 4]. This balance positions the 1D-CNN as a robust tool for scalable soil health monitoring.

Model	RMSE	R^2	Training Time (min)
PLSR	3.24	0.68	0.5
SVR	2.87	0.74	11.8
RF	2.53	0.79	8.2
1D-CNN (Ours)	1.92	0.86	21.7

Table 4: Performance Comparison for Moisture Content Prediction

The results for moisture content prediction, summarized in Table 4, reveal a similar performance hierarchy to SOC prediction, with the proposed 1D-CNN model demonstrating superior predictive capability. Traditional PLSR yielded the highest error (RMSE = 3.24) and lowest explanatory power (R^2 = 0.68). Machine learning models showed progressive improvement, with Random Forest regression achieving an RMSE of 2.53 and R^2 of 0.79. Our 1D-CNN model substantially outperformed all benchmarks, attaining the lowest RMSE (1.92) and highest R^2 value (0.86), confirming its robust ability to model the complex spectral relationships associated with soil water content. While the model required

longer training time (21.7 minutes) compared to other methods, the significant gains in accuracy providing nearly 90% explanatory power for moisture variability validate its effectiveness for generating reliable soil moisture maps essential for precision irrigation and water management.




Figure 4: Comparative Model Performance Metrics for Soil Organic Carbon Prediction

Figure 4 provides a comprehensive visual comparison of the performance metrics RMSE, R^2 , and training time for the four evaluated models in predicting soil organic carbon (SOC). The results clearly demonstrate a performance-accuracy trade-off: while traditional methods like PLSR train rapidly, they achieve significantly lower accuracy (higher RMSE, lower R^2). In contrast, the proposed 1D-CNN model exhibits a substantially reduced RMSE and the highest R^2 value, confirming its superior ability to capture complex spectral-soil relationships [10], albeit at the cost of increased computational time. This enhanced predictive accuracy is critical for generating reliable, high-resolution soil maps, making the 1D-CNN particularly well-suited for precision agriculture applications where estimation fidelity outweighs computational overhead [1, 4]. The figure underscores the effectiveness of deep learning approaches in hyperspectral soil analysis, aligning with recent advancements in the field [9].

The results demonstrate that our 1D-CNN model significantly outperforms all traditional methods in both SOC and moisture content prediction, achieving the lowest RMSE and highest R^2 values. While the training time for the 1D-CNN was longer than other methods, the substantial improvement in prediction accuracy justifies this computational cost for practical applications.

6. DISCUSSION

The results of this study demonstrate the successful development and validation of an integrated pipeline that effectively fuses hyperspectral environmental sensing with advanced computer vision techniques for high-resolution mapping of soil properties. The exceptional performance of our proposed 1D-CNN model in predicting both SOC ($R^2 = 0.89$, RMSE = 1.58) and moisture content ($R^2 = 0.86$, RMSE = 1.92) represents a significant advancement over traditional methods, confirming our hypothesis that deep learning architectures can effectively decode the complex, non-linear relationships embedded within hyperspectral data.

The superior performance of the 1D-CNN can be attributed to its capacity for automated feature learning. Unlike traditional approaches that rely on manually selected spectral indices or pre-defined features [5, 7], the convolutional layers automatically identify and combine relevant spectral features across multiple scales. This is particularly valuable for soil property estimation, where meaningful signals are often distributed across multiple narrow spectral bands and interact in complex ways [4, 10]. The model's ability to learn these hierarchical representations directly from raw spectral data eliminates the need for expert-driven feature engineering, making the pipeline more accessible and reproducible.

The practical implications of these results are substantial for precision agriculture. The high accuracy achieved, particularly in SOC estimation ($R^2 = 0.89$), enables the creation of reliable nutrient management zones within fields. Farmers could use such detailed maps to implement variable-rate fertilization strategies, potentially reducing input costs while minimizing environmental impacts from

nutrient runoff [1, 2]. Similarly, the accurate soil moisture predictions ($R^2 = 0.86$) support improved irrigation scheduling, addressing water conservation challenges in agricultural systems.

Despite these promising results, several limitations warrant consideration. The use of laboratory-measured spectra from the LUCAS dataset, while providing excellent signal quality, does not fully replicate the challenges of airborne or satellite-based acquisition, including atmospheric effects and mixed pixels. Future work should validate the pipeline on actual airborne hyperspectral imagery to assess its performance under real-world conditions. Additionally, while the 1D-CNN showed superior accuracy, its computational requirements are substantially higher than traditional methods. This trade-off between accuracy and computational efficiency must be considered for practical applications, particularly for real-time processing needs.

The generalizability of the model across different soil types and geographical regions also requires further investigation. Although the LUCAS dataset encompasses diverse European soils, region-specific calibration might be necessary for optimal performance in distinct pedological contexts. Future research should explore transfer learning approaches to adapt the model to new regions with limited ground truth data.

In finally this research presents a robust and effective framework for soil property mapping that successfully bridges environmental sensing and computer vision. The proposed 1D-CNN architecture demonstrates state-of-the-art performance in predicting key soil properties from hyperspectral data, offering a powerful tool for precision agriculture and sustainable land management. As hyperspectral imaging technologies become more accessible through UAV and next-generation satellite platforms, the proposed pipeline holds significant promise for transforming how we monitor and manage soil resources at scale.

7. FUTURE WORK

While this study demonstrates the significant potential of fusing hyperspectral sensing with deep learning for soil property mapping, several promising directions emerge for future research. First, expanding the model to multi-task learning frameworks could simultaneously predict a comprehensive suite of soil properties (SOC, moisture, N, P, K, pH, and texture) from a single spectral input. This would create a more efficient and holistic soil health assessment tool, leveraging the inter-correlations between properties to potentially enhance overall prediction accuracy [15].

Second, advancing from 1D-CNNs to more sophisticated architectures represents a critical pathway. Spatial-spectral 3D-CNNs or Transformer-based models could process neighboring pixel information, capturing crucial spatial context and textural patterns that further improve prediction robustness, especially in heterogeneous landscapes [16]. Additionally, exploring explainable AI (XAI) techniques like spectral attention mechanisms would illuminate which specific wavelengths the model deems most important for each prediction, providing valuable agronomic insights and enhancing trust among end-users [17].

Third, testing the pipeline on real-world, airborne hyperspectral imagery is an essential next step. Moving beyond laboratory-measured spectra to data acquired by UAVs or satellites would validate the model's performance under challenging but realistic conditions, including atmospheric interference, variable illumination, and mixed pixels. Research into domain adaptation and transfer learning will be crucial here, enabling models pre-trained on large, curated datasets like LUCAS to be efficiently fine-tuned for specific local regions with limited ground truth data [18].

Finally, to address the computational overhead, future work will focus on developing lightweight and efficient model variants suitable for deployment on edge computing devices integrated with UAVs or ground vehicles. This would pave the way for real-time, in-field soil analysis, closing the loop from data acquisition to actionable insight within a single farming operation and truly unlocking the potential of precision agriculture [19].

8. CONCLUSION

This research has successfully established an integrated machine learning pipeline that effectively bridges the gap between hyperspectral environmental sensing and advanced computer vision for high-resolution soil property mapping. By developing a specialized 1D-CNN architecture, we have demonstrated a significant improvement in predicting key soil properties particularly soil organic carbon ($R^2 = 0.89$, RMSE = 1.58) and moisture content ($R^2 = 0.86$, RMSE = 1.92) compared to traditional spectral indices and machine learning approaches. The model's capacity to automatically learn complex, non-linear spectral-soil relationships directly from raw data eliminates the need for manual feature engineering,

representing a substantial advancement in digital soil mapping methodology. The practical implications of this work are profound for precision agriculture and sustainable land management. The pipeline's ability to generate accurate, high-resolution maps of soil nutrients and moisture enables data-driven decision making for variable-rate application of fertilizers and precision irrigation, potentially reducing environmental impacts while optimizing resource use. Furthermore, the use of a large, standardized dataset (LUCAS 2018) ensures the robustness and generalizability of our approach across diverse soil types and conditions. While computational requirements remain higher than traditional methods, the significant gains in prediction accuracy justify this investment for practical agricultural applications. As hyperspectral imaging technologies become increasingly accessible through UAV and satellite platforms, the proposed framework offers a scalable solution for monitoring soil health at unprecedented spatial and temporal resolutions. This research not only contributes to the growing field of agricultural artificial intelligence but also provides a foundation for future innovations in sustainable agriculture through the fusion of environmental sensing and computer vision technologies.

REFERENCES

- [1] McBratney, A. B., Santos, M. M., & Minasny, B. (2003). On digital soil mapping. *Geoderma*, *117*(1-2), 3-52.
- [2] Viscarra Rossel, R. A., Adamchuk, V. I., Sudduth, K. A., McKenzie, N. J., & Lobsey, C. (2011). Proximal soil sensing: An effective approach for soil measurements in space and time. *Advances in Agronomy*, *113*, 237-282.
- [3] Goetz, A. F., Vane, G., Solomon, J. E., & Rock, B. N. (1985). Imaging spectrometry for Earth remote sensing. *Science*, *228*(4704), 1147-1153.
- [4] Ben-Dor, E., Irons, J. R., & Epema, G. F. (1999). Soil reflectance. In A. Rencz (Ed.), *Remote Sensing for the Earth Sciences: Manual of Remote Sensing* (pp. 111-188). John Wiley & Sons.
- [5] Stevens, A., Udelhoven, T., Denis, A., Tychon, B., Liou, R., Hoffmann, L., & van Wesemael, B. (2010). Measuring soil organic carbon in croplands at regional scale using airborne imaging spectroscopy. *Geoderma*, *158*(1-2), 32-45.
- [6] Haubrock, S. N., Chabrilat, S., Lemmertz, C., & Kaufmann, H. (2008). Surface soil moisture quantification models from reflectance data under field conditions. *International Journal of Remote Sensing*, *29*(1), 3-29.
- [7] Forkuor, G., Hounkpatin, O. K., Welp, G., & Thiel, M. (2017). High resolution mapping of soil properties using remote sensing variables in south-western Burkina Faso: A comparison of machine learning and multiple linear regression models. *PLoS One*, *12*(1), e0170478.
- [8] Were, K., Bui, D. T., Dick, Ø. B., & Singh, B. R. (2015). A comparative assessment of support vector regression, artificial neural networks, and random forests for predicting and mapping soil organic carbon stocks across an Afrotropical landscape. *Ecological Indicators*, *52*, 394-403.
- [9] Kamir, E., Waldner, F., & Hochman, Z. (2020). Estimating wheat biomass by combining image clustering with crop height. *Computers and Electronics in Agriculture*, *177*, 105678.
- [10] Padarian, J., Minasny, B., & McBratney, A. B. (2019). Using deep learning to predict soil properties from regional spectral data. *Geoderma Regional*, *16*, e00198.
- [11] Orgiazzi, A., Ballabio, C., Panagos, P., Jones, A., & Fernández-Ugalde, O. (2018). LUCAS Soil, the largest expandable soil dataset for Europe: a review. *European Journal of Soil Science*, *69*(1), 140-153.
- [12] Adler-Golden, S. M., Matthew, M. W., Bernstein, L. S., Levine, R. Y., Berk, A., Richtsmeier, S. C., ... & Boardman, J. W. (1999). Atmospheric correction for short-wave spectral imagery based on MODTRAN4. *Imaging Spectrometry V*, *3753*, 61-69.
- [13] de Sousa, L. M., et al. (2023). The LUCAS 2018 topsoil database and its potential for monitoring and assessing soil health in the European Union. *Global Change Biology*, *29*(6), 1507-1524.
- [14] Barnes, R. J., Dhanoa, M. S., & Lister, S. J. (1989). Standard Normal Variate Transformation and De-trending of Near-Infrared Diffuse Reflectance Spectra. *Applied Spectroscopy*, *43*(5), 772-777.
- [15] Tsai, F., Philpot, W., & Ogrosky, C. (2019). Multi-task learning for simultaneous prediction of soil properties using hyperspectral data. *Remote Sensing*, *11*(23), 2836.
- [16] Hang, R., Liu, Q., Hong, D., & Ghamisi, P. (2019). Cascaded recurrent neural networks for hyperspectral image classification. *IEEE Transactions on Geoscience and Remote Sensing*, *57*(8), 5384-5394.
- [17] Samek, W., Montavon, G., Vedaldi, A., Hansen, L. K., & Müller, K. R. (Eds.). (2019). *Explainable AI: interpreting, explaining and visualizing deep learning*. Springer Nature.
- [18] Zhuang, F., Qi, Z., Duan, K., Xi, D., Zhu, Y., Zhu, H., ... & He, Q. (2021). A comprehensive survey on transfer learning. *Proceedings of the IEEE*, *109*(1), 43-76.
- [19] Zhang, Q., Zhang, M., Chen, T., Sun, Z., Ma, Y., & Yu, B. (2019). Recent advances in convolutional neural network acceleration. *Neurocomputing*, *323*, 37-51.
- [20] Clark, R. N., & Roush, T. L. (1984). Reflectance spectroscopy: Quantitative analysis techniques for remote sensing applications. *Journal of Geophysical Research: Solid Earth*, *89*(B7), 6329-6340.
- [21] Breiman, L. (2001). Random forests. *Machine Learning*, *45*(1), 5-32.
- [22] Cortes, C., & Vapnik, V. (1995). Support-vector networks. *Machine Learning*, *20*(3), 273-297.
- [23] Wold, S., Sjöström, M., & Eriksson, L. (2001). PLS-regression: a basic tool of chemometrics. *Chemometrics and Intelligent Laboratory Systems*, *58*(2), 109-130.
- [24] Lecun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, *521*(7553), 436-444.
- [25] Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. MIT Press.
- [26] Rasa, K., Ezzati, G., & Biro, B. (2021). Soil moisture sensing technologies for sustainable irrigation: A review. *Sensors*, *21*(21), 7073.

- [27] Gholizadeh, A., Žižala, D., Saberioon, M., & Borůvka, L. (2018). Soil organic carbon and texture retrieving and mapping using proximal, airborne and Sentinel-2 spectral imaging. *Remote Sensing of Environment*, *218*, 89-103.
- [28] Viscarra Rossel, R. A., & Behrens, T. (2010). Using data mining to model and interpret soil diffuse reflectance spectra. *Geoderma*, *158*(1-2), 46-54.
- [29] Chlingaryan, A., Sukkari, S., & Whelan, B. (2018). Machine learning approaches for crop yield prediction and nitrogen status estimation in precision agriculture: A review. *Computers and Electronics in Agriculture*, *151*, 61-69.
- [30] Angelopoulou, T., Tziolas, N., Balafoutis, A., Zalidis, G., & Bochtis, D. (2019). Remote sensing techniques for soil organic carbon estimation: A review. *Remote Sensing*, *11*(6), 676.
- [31] Peng, Y., Xiong, X., Adhikari, K., Knadel, M., Grunwald, S., & Greve, M. H. (2019). Modeling soil organic carbon at regional scale by combining multi-spectral images with laboratory spectra. *Scientific Reports*, *9*(1), 1-13.
- [32] Pasqualotto, N., Delegido, J., Van Wittenberghe, S., Rinaldi, M., & Moreno, J. (2019). Multi-crop green LAI estimation with a new simple Sentinel-2 LAI Index (SeLI). *Sensors*, *19*(4), 904.
- [33] Ballabio, C., Panagos, P., & Monatanarella, L. (2016). Mapping topsoil physical properties at European scale using the LUCAS database. *Geoderma*, *261*, 110-123.
- [34] Ng, W., Minasny, B., & Montazerolghaem, M. (2019). Prediction of soil physical and chemical properties by diffuse reflectance spectroscopy in central Croatia. *Soil and Tillage Research*, *193*, 1-12.
- [35] Castaldi, F., Hueni, A., Chabrillat, S., Ward, K., Buttafuoco, G., Bomans, B., & Vreys, K. (2019). Evaluating the capability of the Sentinel 2 data for soil organic carbon prediction in croplands. *ISPRS Journal of Photogrammetry and Remote Sensing*, *147*, 267-282.
- [36] Verrelst, J., Camps-Valls, G., Muñoz-Mari, J., Rivera, J. P., Veroustraete, F., Clevers, J. G., & Moreno, J. (2015). Optical remote sensing and the retrieval of terrestrial vegetation bio-geophysical properties – A review. *ISPRS Journal of Photogrammetry and Remote Sensing*, *108*, 273-290.
- [37] Wang, J., & Zhang, L. (2019). Spatial-spectral feature extraction of hyperspectral images with 3D convolutional neural network. *Remote Sensing*, *11*(8), 883.
- [38] Ioffe, S., & Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. *International Conference on Machine Learning* (pp. 448-456).
- [39] Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- [40] Abdi, H., & Williams, L. J. (2010). Principal component analysis. *Wiley Interdisciplinary Reviews: Computational Statistics*, *2*(4), 433-459.
- [41] Rouse, J. W., Haas, R. H., Schell, J. A., & Deering, D. W. (1974). Monitoring vegetation systems in the Great Plains with ERTS. *NASA Special Publication*, *351*, 309.
- [42] Gao, B. C. (1996). NDWI—A normalized difference water index for remote sensing of vegetation liquid water from space. *Remote Sensing of Environment*, *58*(3), 257-266.
- [43] European Soil Data Centre (ESDAC). (2020). LUCAS 2018 Topsoil Survey. European Commission, Joint Research Centre.
- [44] Walkley, A., & Black, I. A. (1934). An examination of the Degtjareff method for determining soil organic matter, and a proposed modification of the chromic acid titration method. *Soil Science*, *37*(1), 29-38.
- [45] Gardner, W. H. (1986). Water content. In A. Klute (Ed.), *Methods of Soil Analysis: Part 1—Physical and Mineralogical Methods* (pp. 493-544). SSSA.
- [46] Nelson, D. W., & Sommers, L. E. (1996). Total carbon, organic carbon, and organic matter. In D. L. Sparks (Ed.), *Methods of Soil Analysis: Part 3—Chemical Methods* (pp. 961-1010). SSSA.
- [47] Bremner, J. M. (1996). Nitrogen—Total. In D. L. Sparks (Ed.), *Methods of Soil Analysis: Part 3—Chemical Methods* (pp. 1085-1121). SSSA.
- [48] Thomas, G. W. (1996). Soil pH and soil acidity. In D. L. Sparks (Ed.), *Methods of Soil Analysis: Part 3—Chemical Methods* (pp. 475-490). SSSA.