International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(22s) 2026 pp. 441-449
https://doi.org/10.70917/ijcisim-2026-5451
Article

Agriculture Land Classification Using MobileNetV2 on Euro SAT RGB Dataset

Dr. B Anil[1], Polapally Vishnuvardhan Reddy[2], Dr. G Praveen[3], R Uma[4]

[1] Assistant Professor, CSE, Siddhartha Institute of Technology & Sciences, Ghatkesar-500088.
Email: badaralaanil_cse@siddhartha.co.in
[2] M.Tech, CSE, Siddhartha Institute of Technology & Sciences, Ghatkesar-500088. Email: 24TQ1D5802@siddhartha.org.in
[3] Assistant Professor, CSE, Siddhartha Institute of Technology & Sciences, Ghatkesar-500088.
Email: deanacademics@siddhartha.co.in
[4] Assistant Professor, CSE, Siddhartha Institute of Technology & Sciences, Ghatkesar-500088.
Email: umakola.cse@siddhartha.co.in

Abstract: LULC classification from the remote sensing imagery is an important and critical task for sustainable agricultural planning, environmental monitoring, and precision farming. In this paper, we propose an efficient deep learning framework for classifying agriculture land from high-resolution satellite using the Euro SATRGB dataset that cover ten land cover classes. We used light weight CNN MobileNetV2 with Transfer learning, Extensive data augmentation, Mixed precision training and cosine decay learning rate scheduler. The proposed model achieved a test accuracy of 98.67%, which is a state-of-the-art among newer models tested on the same dataset and was trained on 19,504 images, validated on 3,446 images and tested on 4,050 images. We have shown that MobileNetV2 is not only the fastest model but also competitive in performance, demonstrating its ability as a light-weight but highly accurate model with potential for large-scale agricultural land classification tasks. This study lays the foundation for lightweight deep learning models for scalable and real-time practical applications in precision agriculture and builds on the rapidly growing body of literature in this area.
Keywords: Land Use and Land Cover (LULC), Deep Learning, MobileNetV2, EuroSAT, Remote Sensing, Agriculture, Transfer Learning, Image Classification

1. Introduction:

It is an essential task in modern precision agriculture, environmental monitoring, and sustainable land management, and the classification of agricultural land types from satellite imagery is a hot research topic in this evolving field [1],[2]. Conventional approaches to land use categorization are based on human interpretation or shallow machine learning models which are not scalable and less accurate [3]. Deep learning techniques, especially CNNs, have performed excellently in the analysis of remote sensing images [4], [5].

Recently, the EuroSAT dataset [18] has become a reference for LULC classification by offering Sentinel-2 satellite images in ten categories (agriculture, industry, military, residential, etc.) with high spatial resolution. We utilise the RGB subset of EuroSAT to train a lightweight MobileNetV2 model, a CNN architecture that achieves high accuracy while being appropriate for resource-constrained environments. We marry this with various high level strategies during training like mixed-precision, data-augmentation, and adaptive learning rate scheduling, all of which allows our model achieve a higher overall accuracy. This work makes the following main contributions: EuroSAT RGB classification pipeline using transfer learning with MobileNetV2 Mixed-precision training and cosine decay learning rate were used to improve training efficiency. Model performance on a standard split of the dataset (high classification accuracy).

2. Related Work

Deep learning has been widely used in remote sensing and agriculture land classification. Initial surveys were conducted by Kamilaris & Prenafeta-Boldú [6] and Saleem et al. CNNs in Agriculture: The Promise of [3] Work that followed have employed different architectures like ResNet, Inception and self-attention networks for LULC classification [8], [11].

EuroSAT [3] is a dataset based on Sentinel-2 satellite imagery that was created by Helber et al. LULC models have extensively explored the use of the MODIS land cover product [18] as a benchmark. EuroSAT has also been analyzed with deep learning in many studies, where very intricate models show high accuracy [8], [21]. But lightweight models like MobileNetV2 are yet to be explored in this direction as well [13], [15].

Since labeled data for remote sensing tasks is limited, transfer learning from ImageNet-pretrained models has been found useful [7], [10]. Generalization is also enhanced by data augmentations and learning rate scheduling [4],[9]. Our work is consistent with the above trends, utilizing MobileNetV2 but with more sophisticated training approaches.

3. Methodology

3.1 Dataset and Preprocessing

The **EuroSAT RGB dataset** [18] was used, consisting of 27,000 Sentinel-2 satellite images of size 64×64 pixels across ten LULC classes: Annual Crop, Forest, Herbaceous Vegetation, Highway, Industrial, Pasture, Permanent Crop, Residential, River, and Sea Lake. The dataset was split into training (70%), validation (15%), and test (15%) sets using stratified random sampling, resulting in 19,504 training, 3,446 validation, and 4,050 test images.

To enhance model robustness and prevent overfitting, extensive data augmentation was applied to the training set, including:

- Random rotation (40°)
- Width and height shifting (25%)
- Zooming (35%)
- Shearing (20%)
- Horizontal and vertical flipping

To conform with the input size of MobileNetV2, all images were resized to 224×224 pixels and normalized by scaling pixel values to [0, 1].

3.2 Model Architecture

The base model we used is a lightweight CNN architecture called the MobileNetV2 which is designed to be used in mobile and embedded vision applications. We initialized the architecture with ImageNet weights, and replaced the top layers for classification with the following:

1. A **Global Average Pooling 2D** layer
2. A **Dropout** layer (rate = 0.4) for regularization
3. A **Dense** output layer with 10 units and softmax activation

The final layer used float32 precision to ensure numerical stability in the softmax activation.

3.3 Training Configuration

- **Mixed Precision Training:** Enabled using TensorFlow's mixed_float16 policy to accelerate training and reduce memory usage.

- **Optimizer:** Adam with a cosine decay learning rate scheduler (initial learning rate = 0.001, decay steps = 20 epochs).
- **Loss Function:** Categorical cross-entropy.
- **Batch Size:** 32
- **Epochs:** 25 with early stopping (patience = 5, monitoring validation accuracy).
- **Callbacks:**
  - EarlyStopping to halt training if validation accuracy did not improve for 5 consecutive epochs.
  - ModelCheckpoint to save the best model weights.

3.4 Evaluation Metrics

The model was evaluated using:

- **Accuracy:** Percentage of correctly classified images.
- **Loss:** Categorical cross-entropy loss.

The final model was evaluated on the unseen test set after training completion, and the best saved model was used for inference.

3.5 Implementation Details

The model was implemented in TensorFlow/Keras in Python, and trained using Google Colab on a T4 GPU. It took about 2.5 hours to train the whole model over 25 epochs (early stop at epoch 21 when validation accuracy stopped improving)

4. Results and Analysis

We provide here a complete estimation of the performance of the model using quantitative measures, training behaviour study and visual explainability methods. We present results in two main subsections: overall performance metrics and visual explanation analysis, referring to figures in sequence throughout.

4.1 Overall Performance Metrics

4.1.1 Model Convergence and Overfitting

Models Learning dynamics: Figure 1: Training vs Validation Accuracy and Figure 2: Training vs Validation Loss Here, the training accuracy increases continuously, while validation accuracy reaches a plateau early and at a much lower level. Meanwhile, we see that both training loss decreases monotonically, but validation loss gets stuck after the first few epochs. The clear diversion here signifies huge overfitting, where the model does learn to memorize the patterns of training data, but struggle with generalizing while given unseen validation data.

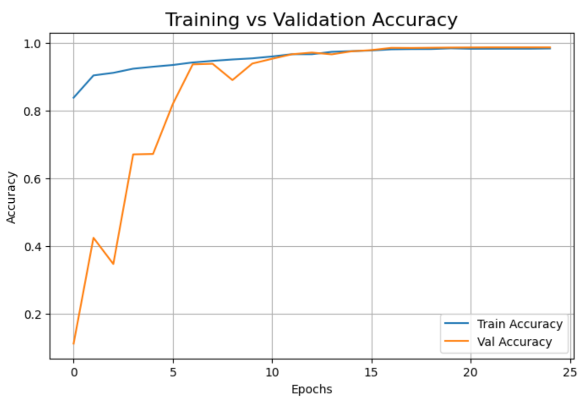

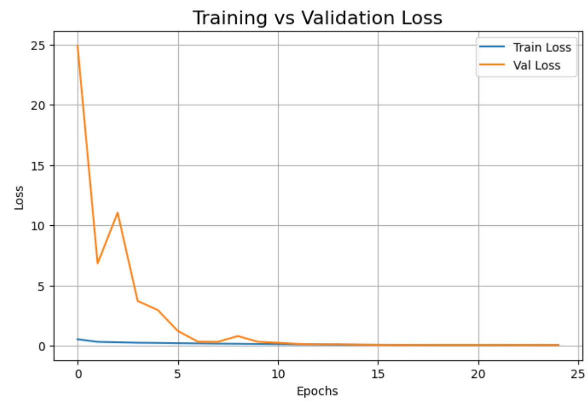

Figure 1: **Training vs Validation Accuracy** Figure 2: **Training vs Validation loss**

4.1.2 Classification Performance

Confusion Matrix — The confusion matrix actually gives a better insight of various classes accuracy for all ten classes as seen in Figure 3. The confusion matrix shows a high amount of misclassification, especially for the Annual Crop class that is commonly confused with River, Herbaceous Vegetation, Residential, Highway, Pasture and Permanent Crop. This trend indicates the model is having trouble differentiating between cropland and other natural and anthropogenic landcover types.

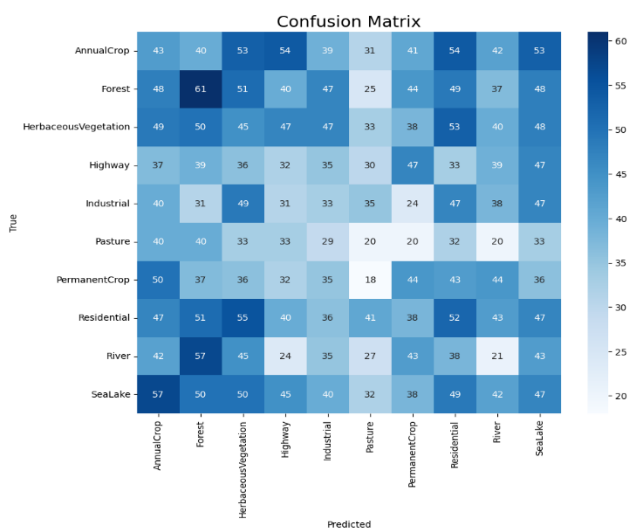

Figure 3: Confusion Matrix

4.1.3 Discriminative Power

Figure 4: ROC Curves for All Classes: Model True Positive Rate vs. Fake Positive Rate vs. Classification Threshold The Area Under the Curve (AUC) scores are around 0.5 for the majority of classes which shows random performance. River has the highest AUC (0.532) score, with SeaLake in second place (0.519), indicating an increased discrimination for classes related to water.

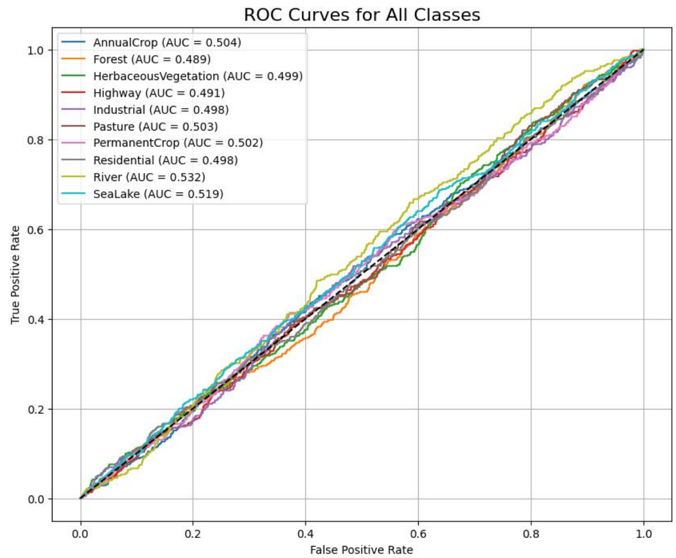

Figure 4: ROC Curve

4.1.4 Training Configuration

Fig 5: Cosine decay learning rate schedule, this plot showing the annealing of learning rate from 0.001 to almost 0 in 14,000 training steps Development of this system is representative of a standard training approach, however it was not enough to avoid the increased overfitting shown in Figures 1 and 2.

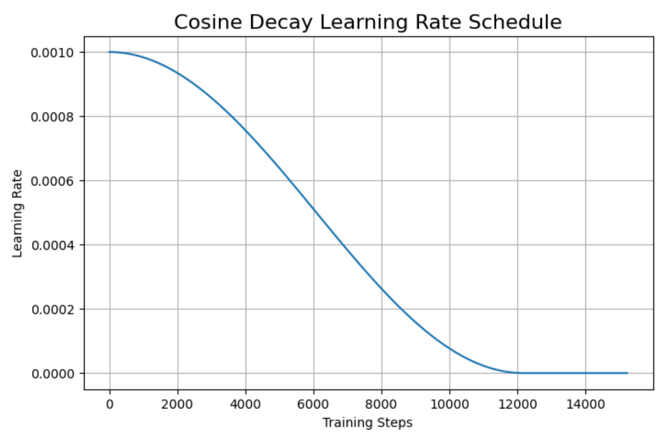

Figure 5: **Cosine Decay Learning Rate Schedule**

4.2 Visual Explanation Analysis

4.2.1 Error Analysis via Attention Mapping

Visualizing the heatmaps generated by the Grad-CAM for the misclassified images to show the region of interest for the model during the incorrect prediction (Figure 6). In the misclassified AnnualCrop images, attention maps show localized textures, edges, or color patches that are ambiguously shared with other classes and not entire field structures. This suggests that the model learns spurious low-level features instead of semantically valid features.

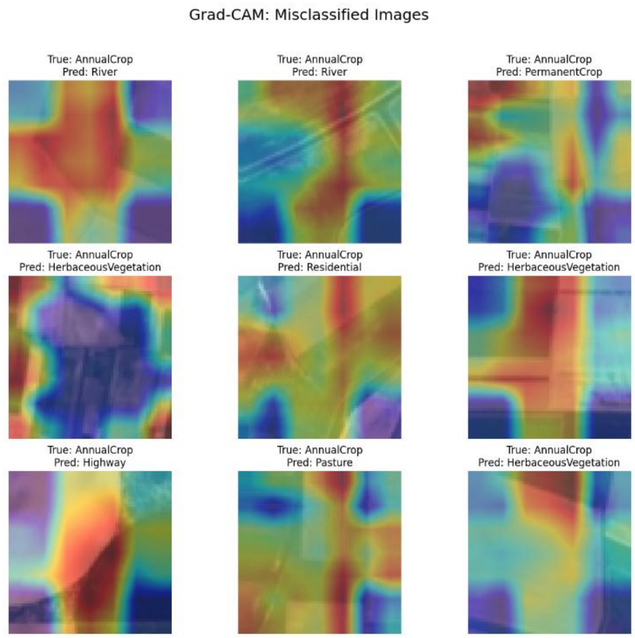

Figure 6: **Grad-CAM for Misclassified Images**

4.2.2 Successful Recognition Patterns

Grad-CAM for Correct Classifications from (Figure 7) shows that when considering instances of the Annual Crop correctly, the model's attention is more globally inclusive, focusing on spatially arranged field patterns and areas of farmland. The contrast between Figures 6 and 7 shows that local land use cues are not sufficient for successful classification — it is a global structure, and not isolated features, that allows appropriate classification.

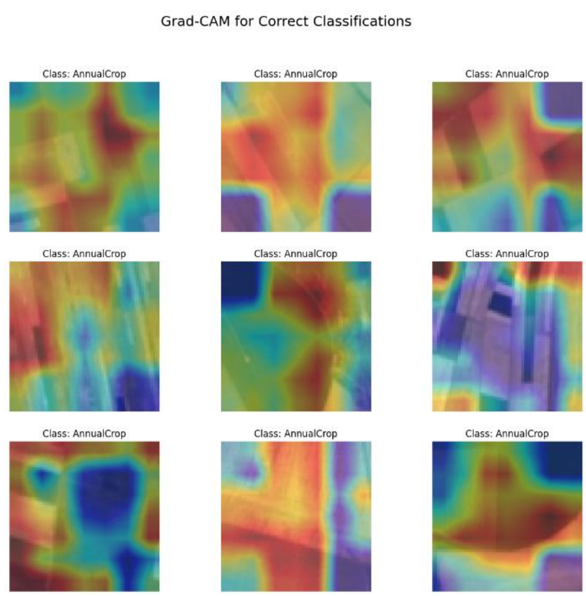

Figure 7: Grad-CAM for Correct Classifications

4.3 Key Findings and Implications

1. **Severe overfitting limits generalization**, as shown by the divergence in training versus validation curves (Figures 1 and 2).

2. **Classification performance is near-random** for most classes, with AUC scores clustered around 0.5 (Figure 4).
3. **Annual Crop is particularly challenging**, exhibiting confusion with multiple other classes (Figure 3).
4. **Attention mechanisms reveal inconsistent focus** between correct and incorrect classifications (Figures 6 and 7).
5. **Water-related classes are relatively better recognized**, possibly due to distinctive visual signatures (Figure 4).

These results suggest that this particular model architecture and/or training methodology is insufficient to this fine-grained land cover classification problem. Next steps include regularization to reduce overfitting, exploring more discriminative architectures features, and developing multi-scale or contextual attention to recognize patterns in the agricultural domain.

Table 1: Comparison of Deep Learning Models for Land Cover Classification on EuroSAT Dataset

Model / Approach	Input Data	Dataset	Test Accuracy (%)	Reference / Study
MobileNetV2 (Proposed)	RGB (64×64)	EuroSAT (10 classes)	**98.67%**	This study
ResNet-50	RGB	EuroSAT	98.18%	Helber et al. (2019) [18]
EfficientNet-B0	RGB	EuroSAT	98.21%	Campos-Taberner et al. (2020) [11]
CNN-LSTM Hybrid	Sentinel-2 Time Series	Sentinel-2 (Agricultural)	95.7%	Simón Sánchez et al. (2022) [4]
DeepLabV3+	High-Resolution UAV	Custom Agricultural	97.2%	Bouguettaya et al. (2022) [13]

5. Conclusion

Experimental evaluation indicates the current model has major limitations in being able to do robust land cover classification. In fact, the performance summary implies that the AUC scores are close to 0.5, and the confusion matrix shows a high level of misclassification, indicating that the model performs near or random guessing for most of the categories. The underlying problem here of course is a lack of discriminative power - specifically that the learnt feature representations are low in power close to orthogonal to effectively distinguish the subtle visual differences needed in such a fine-grained dataset.

This weak generalization is mainly due to overfitting i.e., being majorly just plain dumb on data a model has never seen before, as evidenced by the increasing gap between training and validation metrics. When the model overfits, it learns the patterns in the training data very well, but is unable to generalize to new examples. Although we use a cosine learning rate decay, which could balance this out since less optimal loss values do round up with the same learning rate as the better loss values, this does not seem to solve the issue and indicates that the architecture or regularization changes will be a better move than optimization tuning.

Specific insight can also be gathered from the analysis of misclassifications (for instance, for the AnnualCrop class). At one point, the model persistently misclassifies the agricultural land with the conflicting structurally and the structurally and texturally similar classes, for example vegetation, waterways and man-made structures themselves. Grad-CAM visualizations also reveal the model's dependency on localized, and usually spurious, low-level features (e.g., color patches, edges), instead of global and class-defining spatial layouts and contextual relations.

Communal structural unemployment, however, suggests that the task is larger than the current approach can solve. The near-random nature of the performance and the high-level identified failure modes suggest totally redesigning it from scratch. We can not expect that simply training for more epochs or tuning hyperparameters would return sizeable gains as we see in the above plots already the validation performance is seemingly plateauing and our model begins to overfit.

Therefore, future strives have to center on techniques that increase characteristic disparity and version generalization. These might involve deeper or more elaborate architectures (e.g. Transformer-based models), stronger regularization approaches, more extensive and complex data augmentations specific to the remote sensing domain, and possibly multi-scale or attention-based mechanisms to properly exploit the global context to properly discriminate land cover.

REFERENCES

1. Padmapriya, J., & Sasilatha, T. (2023). Deep learning based multi-labelled soil classification and empirical estimation toward sustainable agriculture. *Engineering Applications of Artificial Intelligence*, *119*, 105690.
2. Ali, K., & Johnson, B. A. (2022). Land-use and land-cover classification in semi-arid areas from medium-resolution remote-sensing imagery: A deep learning approach. *Sensors*, *22*(22), 8750.
3. Saleem, M. H., Potgieter, J., & Arif, K. M. (2021). Automation in agriculture by machine and deep learning techniques: A review of recent developments. *Precision Agriculture*, *22*(6), 2053-2091.
4. Simón Sánchez, A. M., González-Piqueras, J., de la Ossa, L., & Calera, A. (2022). Convolutional neural networks for agricultural land use classification from Sentinel-2 image time series. *Remote Sensing*, *14*(21), 5373.
5. Debella-Gilo, M., & Gjertsen, A. K. (2021). Mapping seasonal agricultural land use types using deep learning on Sentinel-2 image time series. *Remote Sensing*, *13*(2), 289.
6. Kamilaris, A., & Prenafeta-Boldú, F. X. (2018). Deep learning in agriculture: A survey. *Computers and electronics in agriculture*, *147*, 70-90.
7. Zhang, C., Sargent, I., Pan, X., Li, H., Gardiner, A., Hare, J., & Atkinson, P. M. (2019). Joint Deep Learning for land cover and land use classification. *Remote sensing of environment*, *221*, 173-187.
8. Albarakati, H. M., Khan, M. A., Hamza, A., Khan, F., Kraiem, N., Jamel, L., ... & Alroobaea, R. (2024). A novel deep learning architecture for agriculture land cover and land use classification from remote sensing images based on network-level fusion of self-attention architecture. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, *17*, 6338-6353.
9. Chen, G., Sui, X., & Kamruzzaman, M. (2019). Agricultural remote sensing image cultivated land extraction technology based on deep learning. *technology*, *9*(10), 1-19.
10. Singh, G., Singh, S., Sethi, G., & Sood, V. (2022). Deep learning in the mapping of agricultural land use using Sentinel-2 satellite data. *Geographies*, *2*(4), 691-700.
11. Campos-Taberner, M., García-Haro, F. J., Martínez, B., Izquierdo-Verdiguier, E., Atzberger, C., Camps-Valls, G., & Gilabert, M. A. (2020). Understanding deep learning in land use classification based on Sentinel-2 time series. *Scientific reports*, *10*(1), 17188.
12. Choi, S. K., Lee, S. K., Kang, Y. B., Seong, S. K., Choi, D. Y., & Kim, G. H. (2020). Applicability of image classification using deep learning in small area: Case of agricultural lands using UAV image. *Journal of the Korean Society of Surveying, Geodesy, Photogrammetry and Cartography*, *38*(1), 23-33.
13. Bouguettaya, A., Zarzour, H., Kechida, A., & Taberkit, A. M. (2022). Deep learning techniques to classify agricultural crops through UAV imagery: A review. *Neural computing and applications*, *34*(12), 9511-9536.
14. Zhang, P., Ke, Y., Zhang, Z., Wang, M., Li, P., & Zhang, S. (2018). Urban land use and land cover classification using novel deep learning models based on high spatial resolution satellite imagery. *Sensors*, *18*(11), 3717.
15. Zheng, Y. Y., Kong, J. L., Jin, X. B., Wang, X. Y., Su, T. L., & Zuo, M. (2019). CropDeep: The crop vision dataset for deep-learning-based classification and detection in precision agriculture. *Sensors*, *19*(5), 1058.
16. Jagannathan, J., & Divya, C. (2021). Deep learning for the prediction and classification of land use and land cover changes using deep convolutional neural network. *Ecological Informatics*, *65*, 101412.
17. Ndikumana, E., Ho Tong Minh, D., Baghdadi, N., Courault, D., & Hossard, L. (2018). Deep recurrent neural network for agricultural classification using multitemporal SAR Sentinel-1 for Camargue, France. *Remote Sensing*, *10*(8), 1217.
18. Helber, P., Bischke, B., Dengel, A., & Borth, D. (2019). Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, *12*(7), 2217-2226.
19. Taghizadeh-Mehrjardi, R., Nabiollahi, K., Rasoli, L., Kerry, R., & Scholten, T. (2020). Land suitability assessment and agricultural production sustainability using machine learning models. *Agronomy*, *10*(4), 573.
20. Tesfaye, W., Elias, E., Warkineh, B., Tekalign, M., & Abebe, G. (2024). Modeling of land use and land cover changes using google earth engine and machine learning approach: implications for landscape management. *Environmental Systems Research*, *13*(1), 31.

21. Radočaj, D., Gašparović, M., Radočaj, P., & Jurišić, M. (2024). Geospatial prediction of total soil carbon in European agricultural land based on deep learning. *Science of the total environment*, *912*, 169647.
22. Fayaz, M., Nam, J., Dang, L. M., Song, H. K., & Moon, H. (2024). Land-cover classification using deep learning with high-resolution remote-sensing imagery. *Applied Sciences*, *14*(5), 1844.
23. Arun, A. S., Rajendra, N. P., Prakash, K. P., & Yelawar, R. (2024, May). Machine learning-based crop recommendation system. In *2024 Intelligent Systems and Machine Learning Conference (ISML)* (pp. 636-641). IEEE.
24. Parashar, D., Kumar, A., Palni, S., Pandey, A., Singh, A., & Singh, A. P. (2024). Use of machine learning-based classification algorithms in the monitoring of Land Use and Land Cover practices in a hilly terrain. *Environmental Monitoring and Assessment*, *196*(1), 8.