

A LIGHTWEIGHT TRANSFER LEARNING FRAMEWORK FOR REAL-TIME IMAGE CLASSIFICATION IN RESOURCE-CONSTRAINED SYSTEMS

Alketa Hyso, Dezdemonia Gjylapi*

Department of Computer Sciences, University “Ismail Qemali”, Vlorë
Albania

* Corresponding Author, e-mail: alketa.hyso@univlora.edu.al

Abstract: This study investigates an EfficientNetB0-based transfer-learning approach for low-latency image classification on computing platforms with limited resources. The approach is evaluated on a four-class maize leaf disease dataset containing Blight, Common Rust, Gray Leaf Spot, and Healthy samples. The experimental protocol includes stratified dataset partitioning, training-time image augmentation, class-weighted optimization, and a two-stage training strategy in which the pretrained backbone is first frozen and then partially fine-tuned. Model performance is assessed using accuracy, macro-F1, balanced accuracy, ablation experiments, robustness tests, and deployment-related indicators. On the held-out test set, the model achieved 96.34% accuracy and a macro-F1 score of 0.95. Robustness evaluation under brightness changes, contrast shifts, Gaussian noise, and blur showed only limited performance degradation. The final model required approximately 28 MB of storage and achieved an inference time of about 3 ms per image. These results suggest that a compact transfer-learning model can support accurate and computationally efficient image classification in environments where memory, processing capacity, and response time are constrained.

Key words: real-time image classification, transfer learning, EfficientNetB0, resource-constrained systems, maize leaf disease detection.

1. INTRODUCTION

Recent advances in artificial intelligence have enabled intelligent systems to perform complex visual recognition tasks in real time, including image classification, object recognition, and automated detection across different application domains [1–3]. However, strong predictive accuracy alone is not sufficient for practical use. In mobile, embedded, and edge-based settings, an image-classification model must operate within limited memory and processing budgets, remain stable under input variability, and support fast inference on constrained hardware [4–6].

Deep learning has improved visual recognition accuracy, but large network designs can be difficult to use in real-time or mobile settings because they require considerable

computation. For this reason, a lightweight transfer-learning setup can reuse visual features learned from large image datasets and adapt them to the target task with lower training cost [7–9].

For applied visual recognition systems, evaluation should also consider system-level factors such as robustness to input variability, class imbalance, model size, and inference latency. Agricultural image classification offers a useful validation setting because images may vary due to illumination, background complexity, sensor differences, blur, noise, and heterogeneous visual patterns [10–12].

This study develops an EfficientNetB0-based transfer-learning pipeline intended for fast image recognition on devices with limited computational resources. The experimental validation focuses on maize-leaf imagery, where the target labels correspond to three disease conditions, Common Rust, Gray Leaf Spot, and Blight, together with a Healthy class [13]. The workflow brings together stratified dataset partitioning, training-time augmentation, class-weighted learning, two-phase optimization, ablation experiments, robustness evaluation, and deployment-focused measurements.

This study makes three main contributions. First, it designs a compact transfer-learning workflow for low-latency image classification. Second, it uses ablation experiments to measure the influence of augmentation, class weighting, and fine-tuning. Third, it examines practical deployment through robustness tests, network size, and per-image processing time.

The remainder of the paper reviews related work, describes the experimental methodology, presents the classification, robustness, and deployment results, and discusses limitations and future research directions.

2. RELATED WORK

Deep neural models are widely used in visual recognition because they can extract low-level patterns and progressively combine them into higher-level image features. Convolutional architectures and their variants have been successfully applied across different intelligent recognition domains, including food image classification, multi-label visual recommendation, detection-oriented information systems, and large-scale image recognition [1–3, 14, 15]. These advances established deep neural networks as a core technology for modern image classification systems.

Recent work in visual recognition has increasingly focused on model designs that can operate with limited computational resources. In edge and mobile vision settings, the model must fit within restricted memory and compute budgets while still producing predictions quickly [4]. Architectures designed for efficient vision, including MobileNet, MobileNetV2, and EfficientNet, reduce computational demand while retaining competitive recognition accuracy, which makes them suitable candidates for mobile and real-time visual applications [5, 6, 7]. In parallel, efficient DNN processing, model compression, and quantization have been widely studied as strategies for reducing model size and accelerating inference on constrained hardware [16–18].

Transfer learning is another important strategy for efficient model development, particularly when domain-specific datasets are limited or moderately imbalanced. In this

setting, a model can start from previously learned image features instead of learning all representations from the target dataset alone [8, 9]. In addition, data augmentation and class-imbalance handling are commonly used to improve generalization and balanced learning in image classification tasks [19–22]. These methods are particularly relevant for deployment-oriented pipelines, where performance stability must be maintained across different input conditions.

Practical deployment also requires evaluation beyond standard accuracy metrics. For deployment-oriented image classification, the evaluation should also examine how predictions change when images are affected by blur, noise, contrast shifts, or lighting variation [23]. Agricultural image classification provides a practical validation domain because real-world images often contain variability in lighting, background, sensor quality, disease appearance, and acquisition conditions [10–12, 24–26]. This makes it suitable for evaluating whether lightweight transfer learning pipelines remain accurate, robust, and efficient under conditions relevant to resource-aware intelligent systems.

Compared with studies that primarily emphasize classification accuracy, the present work combines ablation analysis, robustness testing, and deployment-oriented metrics within a compact transfer learning framework. In this way, the agricultural validation task is used not only to measure classification accuracy, but also to examine whether the proposed pipeline remains practical under deployment-related constraints.

3. EXPERIMENTAL METHODOLOGY

3.1. Dataset, Preprocessing and Experimental Setup

The experiments use the Kaggle Corn or Maize Leaf Disease Dataset [13], which contains 4,188 RGB leaf images distributed across four labels: Blight (1,146 images), Common Rust (1,306 images), Gray Leaf Spot (574 images), and Healthy (1,162 images). The dataset is moderately imbalanced, with Gray Leaf Spot as the minority class.

All images are provided in JPEG format and organized in class-wise directories. During data loading, image paths are automatically indexed and mapped to numeric labels based on alphabetical class ordering. The resulting mapping is stored to ensure consistent label interpretation across training and inference.

Images were divided into training, validation, and test subsets with a 70:15:15 ratio, and the split was stratified so that each subset retained the original class proportions. The implementation was carried out in TensorFlow/Keras, and the random seed was fixed at 42 so that the dataset partitions and training runs could be reproduced consistently.

Before being passed to the network, each image is scaled to 224×224 pixels, the spatial resolution expected by EfficientNetB0, and then transformed with the ImageNet-specific preprocessing routine used by the backbone.

During data loading, augmentation is restricted to the training subset and includes horizontal flips, rotations of up to $\pm 8^\circ$, zoom changes of up to $\pm 10\%$, and contrast adjustments of up to $\pm 10\%$. This increases visual diversity during model fitting, while the validation and test subsets are kept unchanged for independent evaluation. Examples of the augmentation process are presented in Figure 1.



Figure 1. Training-time augmentation example for a maize leaf image: original input, transformed image after rotation, zoom, and contrast changes, and the corresponding pixel-wise difference map.

3.2. Class Imbalance Handling

To address class imbalance, class weights are computed using an inverse-frequency scheme:

$$w_i = \frac{N}{K \cdot n_i} \quad (1)$$

In this expression, N is the size of the training subset, K is the total number of target categories, and n_i is the number of training images assigned to category i .

These weights are incorporated into the categorical cross-entropy loss to improve balanced learning across classes.

3.3. Model Architecture and Training Strategy

EfficientNetB0 initialized with ImageNet weights is used as the base network [7, 14], and its original top classification layers are excluded. The remaining backbone serves as the feature-extraction component, allowing the model to reuse pretrained image representations within a transfer-learning setup [8, 9].

The custom output head contains a global-average-pooling layer, a dropout layer with rate 0.25, and a four-unit softmax layer for maize-leaf category prediction. This design keeps the model compact while allowing adaptation to the target classification task.

The optimization procedure was organized as a two-phase transfer-learning process. During the initial phase, the EfficientNetB0 feature extractor was kept non-trainable, while the newly added classification layers were optimized with Adam at a learning rate of 1×10^{-3} for 12 epochs. After this adaptation step, training continued by enabling updates in the final 30 layers of the backbone and lowering the learning rate to 1×10^{-4} for a further 10 epochs.

Model optimization relies on categorical cross-entropy as the objective function. Training stability was supported through early stopping, adaptive reduction of the learning rate after validation plateaus, and checkpoint saving of the best-performing

model. A batch size of 32 was used in every training configuration. The overall architecture and training strategy are illustrated in Figure 2.




Figure 2. Model architecture overview

3.4. Evaluation Metrics and Performance Assessment

The held-out test subset was assessed with several metrics rather than accuracy alone. Because the class distribution is uneven, the evaluation also reports macro-F1 and balanced accuracy to reflect performance across individual maize-leaf categories.

A four-class confusion matrix was used to inspect the distribution of correct and incorrect predictions. Rows correspond to the reference maize-leaf labels, columns correspond to model outputs, and off-diagonal entries reveal which disease categories are mistaken for one another. The same counts support the calculation of precision, recall, and F1-score for each class.

Accuracy measures the overall proportion of correctly classified samples:

$$Accuracy = \frac{TP+TN}{TP+TN+FP+FN} \quad (2)$$

Precision, recall, and F1-score are computed as:

$$Precision = \frac{TP}{TP+FP} \quad (3)$$

$$Recall = \frac{TP}{TP+FN} \quad (4)$$

$$F1 = 2 \cdot \frac{Precision \cdot Recall}{Precision + Recall} \quad (5)$$

Macro-F1 is reported to give equal weight to each maize-leaf category:

$$Macro - F1 = \frac{1}{K} \sum_{i=1}^K F1_i \quad (6)$$

where K is the number of classes.

Balanced accuracy is calculated by averaging the recall values obtained for the individual classes:

$$Balanced Accuracy = \frac{1}{K} \sum_{i=1}^K Recall_i \quad (7)$$

Test-set results are reported for the checkpoint that reached the best validation accuracy during training.

3.5. Component Contribution Analysis

An ablation study is conducted to evaluate the contribution of key components. Four configurations are tested under identical settings:

Baseline: augmentation + class weighting + fine-tuning

NoAug: without augmentation

NoCW: without class weighting

NoFT: without fine-tuning

All ablation variants are compared on the held-out test partition, with accuracy, macro-F1, and balanced accuracy reported to capture both overall and class-balanced behavior.

3.6. Robustness Analysis

The robustness experiment uses the trained model without additional fitting and applies predefined image corruptions to the test subset. The tested conditions include brightness changes, contrast modification, Gaussian noise, and blur.

For every corrupted version of the test subset, accuracy, macro-F1, and balanced accuracy are recalculated to measure the effect of each input change. The perturbation settings are kept fixed so that the conditions can be compared consistently.

3.7. System-Level Design and Deployment Considerations

The proposed framework is designed as a deployment-oriented image-classification pipeline rather than only as a predictive model. The workflow includes image input, preprocessing, feature extraction through a pretrained EfficientNetB0 backbone, lightweight classification, and output visualization through a prototype interface.

From a system-level perspective, the framework considers computational efficiency, latency, scalability, and robustness. The EfficientNetB0 backbone keeps the network footprint modest, and latency was measured on a per-image basis to evaluate whether the pipeline can respond within practical time constraints. The modular pipeline can be adapted to other image-classification tasks by modifying the dataset and output layer. Robustness is evaluated through controlled perturbations such as noise, blur, and illumination variation. A Gradio prototype is used to test predictions interactively and to illustrate how the pipeline could be evaluated before use on constrained devices.

4. EXPERIMENTAL RESULTS AND ANALYSIS

4.1. Overall Classification Performance

On the held-out test partition, the EfficientNetB0-based classifier obtained 96.34% accuracy with a loss of 0.114. This outcome shows that the trained model transfers well to images excluded from model fitting, even though the dataset contains moderate class imbalance and visually similar disease symptoms. A macro-F1 value of 0.95 indicates that the classifier maintains strong performance at the class level, including for Gray Leaf Spot, which has fewer samples than the other categories. Similarly, the balanced accuracy of 0.95 demonstrates consistent recall across all four disease categories. Overall, the agreement between macro-F1 and balanced accuracy suggests that the classifier preserves reliable recognition across all maize-leaf categories, including the smaller Gray Leaf Spot class.

4.2. Per-Class Performance Analysis

For each maize-leaf category, the predicted class labels were compared with the corresponding test annotations, and precision, recall, and F1-score were then calculated. The class-level results are reported in Table 1.

The class-level scores show reliable recognition across the four categories. In the evaluated test split, the Healthy category is classified without errors, which indicates that non-symptomatic leaves are clearly separated from diseased samples. This behavior is important for screening use cases, where false warnings should be minimized.

Common Rust is recognized with very few errors, and Blight also retains high precision and recall. Gray Leaf Spot is the most difficult category, with recall decreasing to 0.85, mainly because some samples resemble Blight lesions in the confusion matrix. Even so, its F1-score of 0.88 shows that the minority class remains detectably represented despite the imbalance in the dataset.

Table 1. Class-wise precision, recall, and F1-score for the held-out test subset

Class	Precision	Recall	F1-score
Blight	0.93	0.95	0.94
Common Rust	0.98	0.99	0.98
Gray Leaf Spot	0.92	0.85	0.88
Healthy	1	1	1

4.3. Class-Level Error Analysis

Figure 4 shows the class-level prediction pattern. Most samples fall on the diagonal of the matrix, while the few errors are concentrated between Blight and Gray Leaf Spot, whose lesions can look similar.

In the evaluated test split, no confusion occurred between Healthy and diseased leaves, supporting the classifier’s reliability for screening purposes. The remaining errors were concentrated between visually similar disease classes, particularly Blight and Gray Leaf Spot, indicating localized rather than random uncertainty.




Figure 4. Confusion matrix for the four-class maize leaf classification task

4.4. Ablation Study Results

Table 2 reports on the effect of removing individual elements from the training pipeline. The largest performance loss is observed for the NoFT configuration, indicating that updating the upper EfficientNetB0 layers is the key factor behind the domain adaptation achieved in this experiment. In contrast, excluding augmentation or class weighting changes the results only slightly, suggesting that these components mainly provide secondary stabilization for the present dataset.

Table 2. Ablation study results across pipeline configurations

Variant	Accuracy	Macro-F1	Balanced Accuracy
NoCW	0.9634	0.9528	0.9468
NoAug	0.9618	0.9522	0.9486
Baseline	0.9603	0.9489	0.9453
NoFT	0.9269	0.9004	0.8893

From a practical standpoint, these findings indicate that partial backbone fine-tuning is the dominant factor enabling domain adaptation, whereas augmentation and class weighting act as secondary stabilizing components under the present dataset characteristics.

4.5. Robustness Analysis

Field-like image variation was simulated by introducing predefined corruptions into the test images, and the resulting performance changes are reported in Table 3.

Table 3. Performance under synthetic image corruptions

Condition	Accuracy	Macro-F1	Balanced Accuracy
Clean	0.955	0.943	0.938
Brightness	0.955	0.943	0.938
Contrast	0.952	0.939	0.935
Gaussian Noise	0.955	0.943	0.938
Blur	0.944	0.927	0.919

The model remains highly stable under brightness variation and Gaussian noise, with negligible performance degradation. Contrast adjustments produce only minor impact. Blur causes the largest decline, which is consistent with the removal of fine lesion details that help separate visually similar disease classes. Even under the corrupted inputs, macro-F1 stays higher than 0.92, suggesting that the classifier relies on stable disease-related patterns rather than isolated pixel-level cues.

4.6. Deployment and Qualitative Evaluation

Beyond quantitative metrics, qualitative evaluation was performed using the Gradio-based deployment prototype. Figure 5 presents an example prediction.




Figure 5. Example prediction from the Gradio-based deployment interface showing the predicted class and corresponding softmax probabilities for a maize leaf image.

For the illustrated maize leaf image, the model predicts Blight with a probability of 0.9686, while alternative classes receive substantially lower probabilities. The separation between the predicted class probability and the remaining class probabilities indicates confident model output in the prototype setting.

Deployment metrics further indicate practical feasibility: Model size: 27.75 MB; Latency: 3.02 ms per image; Throughput: 331 images per second. These characteristics support real-time inference capability, making the system suitable for integration into mobile-assisted or web-based agricultural advisory tools.

5. DISCUSSION

The proposed EfficientNetB0-based transfer learning framework achieved strong and balanced performance, indicating effective multi-class image recognition under moderate class imbalance. The ablation results show that partial backbone fine-tuning is the main contributor to performance gains, while augmentation and class weighting provide additional stabilization.

Robustness testing under brightness, contrast, noise, and blur perturbations showed stable behavior, suggesting that the model captures meaningful visual features rather than relying only on fragile image artifacts. At the system level, the model footprint of

about 28 MB and the per-image inference time of approximately 3 ms indicate that the framework can meet the memory and latency requirements of edge and mobile computing platforms.

Despite these results, the study has limitations. The framework was validated on a single agricultural image dataset, and synthetic perturbations cannot fully represent real-world variability. Future work should evaluate the pipeline on additional datasets, real edge or mobile hardware, and more diverse image-classification tasks.

6. CONCLUSIONS AND FUTURE WORK

This study evaluated an EfficientNetB0-based pipeline for fast visual recognition on constrained computing platforms. In the four-class maize-leaf experiment, the model reached 96.34% test accuracy and a macro-F1 value of 0.95. The same trained model also retained consistent predictions across test images modified with noise, blur, contrast, and illumination changes.

The results indicate that compact transfer-learning pipelines can retain high classification accuracy while improving robustness and inference efficiency. The proposed framework provides a reproducible basis for intelligent, resource-aware visual recognition systems, with future work directed toward broader multi-domain and hardware-level validation.

REFERENCES

- [1] Patel, J., Talsania, H., Modi, K. ODNET: Optimized DenseNet for Indian food classification. *International Journal on Information Technologies and Security*, vol.15, no.4, 2023, pp. 27-36. DOI: 10.59035/FPBL3081.
- [2] Vlahova-Takova, M., Lazarova, M. CNN based multi-label image classification for presentation recommender system. *International Journal on Information Technologies and Security*, vol.16, no.4, 2024, pp. 73-84. DOI: 10.59035/PUYE7368.
- [3] Manev, K., Petrov, M., Aleksieva-Petrova, A. Enhancing phishing site detection and prevention through machine learning with image classification. *International Journal on Information Technologies and Security*, vol.17, no.4, 2025, pp. 87-98. DOI: 10.59035/SABE2922.
- [4] Zhou, Z., Chen, X., Li, E., Zeng, L., Luo, K., Zhang, J. Edge intelligence: Paving the last mile of artificial intelligence with edge computing. *Proceedings of the IEEE*, vol.107, no.8, 2019, pp. 1738-1762. DOI: 10.1109/JPROC.2019.2918951.
- [5] Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H. MobileNets: Efficient convolutional neural networks for mobile vision applications. *arXiv*, 2017. DOI: 10.48550/arXiv.1704.04861.
- [6] Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., Chen, L.C. MobileNetV2: Inverted residuals and linear bottlenecks. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Salt Lake City, USA, June 2018, pp. 4510-4520. DOI: 10.1109/CVPR.2018.00474.

- [7] Tan, M., Le, Q.V. EfficientNet: Rethinking model scaling for convolutional neural networks. *Proceedings of the 36th International Conference on Machine Learning (ICML)*, Long Beach, USA, June 2019, pp. 6105-6114. DOI: 10.48550/arXiv.1905.11946.
- [8] Zhuang, F., Qi, Z., Duan, K., Xi, D., Zhu, Y., Zhu, H., et al. A comprehensive survey on transfer learning. *Proceedings of the IEEE*, vol.109, 2020, pp. 43-76. DOI: 10.1109/JPROC.2020.3004555.
- [9] Hossen, M.I., Awrangjeb, M., Pan, S., et al. Transfer learning in agriculture: a review. *Artificial Intelligence Review*, vol.58, 2025, Article 97. DOI: 10.1007/s10462-024-11081-x.
- [10] Gui, P., Dang, W., Zhu, F., Zhao, Q. Towards automatic field plant disease recognition. *Computers and Electronics in Agriculture*, vol.191, 2021, Article 106523. DOI: 10.1016/j.compag.2021.106523.
- [11] Shafay, M., Hassan, T., Owais, M., Hussain, I., Khawaja, S.G., Seneviratne, L., Werghi, N. Recent advances in plant disease detection: challenges and opportunities. *Plant Methods*, vol.21, no.1, 2025, Article 140. DOI: 10.1186/s13007-025-01450-0.
- [12] Khan, A.T., Jensen, S.M., Khan, A.R., Li, S. Plant disease detection model for edge computing devices. *Frontiers in Plant Science*, vol.14, 2023, Article 1308528. DOI: 10.3389/fpls.2023.1308528.
- [13] Ghose, S. Corn or Maize Leaf Disease Dataset. *Kaggle*, 2021.
URL: <https://www.kaggle.com/datasets/smaranjitghose/corn-or-maize-leaf-disease-dataset>
(Visited on 14.04.2026).
- [14] Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L. ImageNet: A large-scale hierarchical image database. *2009 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Miami, USA, June 2009, pp. 248-255. DOI: 10.1109/CVPR.2009.5206848.
- [15] He, K., Zhang, X., Ren, S., Sun, J. Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Las Vegas, USA, June 2016, pp. 770-778. DOI: 10.1109/CVPR.2016.90.
- [16] Sze, V., Chen, Y.H., Yang, T.J., Emer, J.S. Efficient processing of deep neural networks: A tutorial and survey. *Proceedings of the IEEE*, vol.105, no.12, 2017, pp. 2295-2329. DOI: 10.1109/JPROC.2017.2761740.
- [17] Han, S., Mao, H., Dally, W.J. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. *International Conference on Learning Representations (ICLR)*, San Juan, Puerto Rico, May 2016. DOI: 10.48550/arXiv.1510.00149.
- [18] Gholami, A., Kim, S., Dong, Z., Yao, Z., Mahoney, M.W., Keutzer, K. A survey of quantization methods for efficient neural network inference. In book *Low-Power Computer Vision* (eds. Bhattacharyya, S.S., Deprettere, E.F., Leupers, R., Takala, J.), Chapman and Hall/CRC, USA, 2022, pp. 291-326. DOI: 10.1201/9781003162810-13.
- [19] Shorten, C., Khoshgoftaar, T.M. A survey on image data augmentation for deep learning. *Journal of Big Data*, vol.6, 2019, Article 60. DOI: 10.1186/s40537-019-0197-0.
- [20] Xu, M., Yoon, S., Fuentes, A., Park, D.S. A comprehensive survey of image augmentation techniques for deep learning. *Pattern Recognition*, vol.137, 2023, Article 109347. DOI: 10.1016/j.patcog.2023.109347.

- [21] Cui, Y., Jia, M., Lin, T.Y., Song, Y., Belongie, S. Class-balanced loss based on effective number of samples. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Long Beach, USA, June 2019, pp. 9260-9269. DOI: 10.1109/CVPR.2019.00949.
- [22] Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P. Focal loss for dense object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol.42, no.2, 2020, pp.318-327. DOI: 10.1109/TPAMI.2018.2858826.
- [23] Hendrycks, D., Dietterich, T. Benchmarking neural network robustness to common corruptions and perturbations. *International Conference on Learning Representations (ICLR)*, New Orleans, USA, May 2019. DOI: 10.48550/arXiv.1903.12261.
- [24] Liu, J., Wang, X. Plant diseases and pests detection based on deep learning: a review. *Plant Methods*, vol.17, 2021, Article 22. DOI: 10.1186/s13007-021-00722-9.
- [25] Xu, M., Park, J.E., Lee, J., Yang, J., Yoon, S. Plant disease recognition datasets in the age of deep learning: Challenges and opportunities. *Frontiers in Plant Science*, vol.15, 2024, Article 1452551. DOI: 10.3389/fpls.2024.1452551.
- [26] Wang, Y., Yang, Z., Kootstra, G., et al. The impact of variable illumination on vegetation indices and evaluation of illumination correction methods on chlorophyll content estimation using UAV imagery. *Plant Methods*, vol.19, 2023, Article 51. DOI: 10.1186/s13007-023-01028-8.

Information about the authors:

Alketa Hyso – PhD in Computer Engineering, Head of the Computer Science Department at the University “Ismail Qemali”, Vlora, Albania. Research interests: intelligent systems, machine learning applications, and applied computational modeling.

Dezdemonë Gjylapi – PhD in Computer Science, Full-time Lecturer in the Computer Science Department at the University “Ismail Qemali” of Vlora, Albania, and Head of the Teaching Group “Programming and Software Development”. Research interests: genetic algorithms, artificial neural networks, and their applications in various domains.

Manuscript received on 16 April 2026