A Comparative Analysis of Deep Learning Architectures for the Recognition of Endangered Flower Species in Kazakhstan

Zhanar Oralbekova
Astana IT University, Astana, Kazakhstan
zh.oralbekova@astanait.edu.kz

Tomiris Tapen
Astana IT University, Astana, Kazakhstan
230068@astanait.edu.kz

Ainur Shekerbek
L.N. Gumilyov Eurasian National University, Astana, Kazakhstan
shekerbek_aa@enu.kz

Makpal Zhartybayeva
Astana IT University, Astana, Kazakhstan
m.zhartybayeva@astanait.edu.kz (corresponding author)

Alibek Orynbek
Astana IT University, Astana, Kazakhstan
alibek.orynbek@astanait.edu.kz

Received: 1 March 2026 | Revised: 16 April 2026 and 23 April 2026 | Accepted: 30 April 2026

 DOI: https://doi.org/10.48084/etasr.18465

ABSTRACT

Deep learning was used in this study to recognize endangered flower species in Kazakhstan. The results can help monitor biodiversity and protect those species. Convolutional Neural Network (CNN) models were used to classify different flower species. In this project, a publicly available Kaggle dataset was used, that included five flower species, namely sunflower, tulip, dandelion, rose, and daisy and two endangered species from the Red Book of Kazakhstan, Paeonia anomala and Crocus korolkowii. Employing transfer learning with pre-trained ImageNet weights, eight deep learning architectures were trained and evaluated: AlexNet, VGG16, GoogLeNet, ResNet50, DenseNet121, EfficientNet-B0, MobileNetV2, and YOLOv8-cls. To enhance generalization and address class imbalance, data augmentation and class equalization methods were applied. The evaluation criteria included accuracy, precision, recall, F1-score, and mean Average Precision (mAP). When comparing the models, YOLOv8-cls yielded the highest accuracy, whereas MobileNetV2 provided a noteworthy balance between speed and precision. This research reveals that image recognition systems rooted in deep learning can benefit ecological studies conducted in Kazakhstan by automatically recognizing common and endangered flowers, therefore playing a role in cataloging biodiversity and environmental monitoring.

Keywords-convolutional neural networks; transfer learning; image classification; biodiversity monitoring; endangered flora

I. INTRODUCTION

Kazakhstan is home to a diverse range of flora. However, some species are threatened because of challenges like climate change, overgrazing, and industrial land use. Local plants such as Paeonia anomala and Crocus korolkowii are listed in the Red Book of Kazakhstan as endangered species [1-2]. Accurately documenting and classifying such species is vital for environmental planning and nature conservation.

Artificial intelligence and computer vision are becoming more important in environmental research. Deep learning, especially convolutional neural networks (CNN), has made it feasible to automatically identify plants and animals from digital images, reducing the need for extensive field surveys done manually by botanists [3-4].

This paper focuses on developing a reliable flower classification system based on modern CNN architectures. The dataset was sourced from an open Kaggle database. It initially included five common flower species, then two rare local species were added, resulting in a total of seven flower classes. The study used eight deep learning models: AlexNet, VGG16, GoogLeNet, ResNet50, DenseNet121, EfficientNet-B0, MobileNetV2, and YOLOv8. All models were further trained on this dataset using pre-trained ImageNet weights. To improve results and reduce the impact of different numbers of images per class, data augmentation and weighted sampling were used. Model performance was assessed using accuracy, precision, recall, F1-score, and mAP. The comparison revealed which models are better at recognizing flower species.

Contemporary research applying deep learning methodologies for flora classification and flower pattern identification has demonstrated significant progress in automated biodiversity monitoring [3-4]. Models based on CNNs have demonstrated their performance in accurately identifying different plant species under real-world lighting conditions and heterogeneous environments. Preliminary studies have mainly focused on simplified architectures or the use of pre-trained models on limited datasets.

For example, authors in [5] implemented transfer learning using CNNs for flower identification, showing that pre-trained models greatly outperform simpler model architectures. Authors in [3] applied deep neural networks for independent flower species identification, showing that convolutional layers can more effectively differentiate overlapping flower patterns [3]. Authors in [6] proposed a method for pineapple image classification using pre-trained CNNS. Authors in [7] used a hybrid method based on deep learning for classifying rose leaf diseases. Subsequent studies provided more sophisticated frameworks that enhanced precision and processing speed. Authors in [8] improved MobileNetV2 for efficient image categorization [8] acquiring performance levels similar to larger CNNs like VGG16 and ResNet50. Authors in [9] showcased immediate flower identification through CNNs designed for mobile technology [9].

Combining knowledge transfer with data augmentation enhances resilience in environmental image analysis [3-4]. Studies focusing on ecological factors have emphasized the necessity for AI-driven systems in preserving regional biodiversity. Authors in [2] studied natural populations of Paeonia anomala in Eastern Kazakhstan, whereas Mahmudov [10] examined the bio-ecological characteristics of Crocus korolkowii. These results highlight the significance of automated monitoring systems for at-risk species.

Advancements in methodology, particularly in validating model precision, have highlighted the significance of k-fold cross-validation for obtaining reliable evaluations of CNN performance with restricted data quantities [11-12].

The ImageNet dataset has played a foundational role in developing transferable visual representations for such models, providing large-scale pretraining that significantly enhances downstream accuracy on limited datasets [13].

Contemporary convolutional architectures have progressed toward enhanced computational effectiveness and feature reuse. MobileNetV2 presented inverted residuals and linear bottlenecks, allowing for lightweight but high-performance models that are ideal for mobile or embedded applications [14]. DenseNet enhanced gradient flow and feature transmission via dense connections among layers, attaining impressive results with a reduced number of parameters [15]. EfficientNet introduced a composite scaling method that equally adjusts network depth, width, and resolution to enhance accuracy while reducing computational overhead [16].

Despite the rapid development of deep CNNs for plant and flower classification, several important gaps remain in the existing literature.

First, most prior studies focus on globally available benchmark datasets containing common flower species, while region-specific ecological datasets–particularly those including endangered or Red Book [1] species–remain underexplored. There is a lack of research addressing biodiversity monitoring in geographically localized ecosystems, such as Central Asia and Kazakhstan.

Second, previous works typically evaluate either classification-based CNN architectures (e.g., VGG, ResNet, DenseNet) or detection-oriented frameworks (e.g., YOLO) separately. The systematic comparative analysis between modern detection-based and classification-based deep learning models within a unified experimental framework remains limited. In particular, the adaptability of detection-based architectures (such as YOLOv8-cls) for pure classification tasks in ecological contexts has not been sufficiently investigated.

Third, many studies report high classification accuracy but do not explicitly connect their findings to real-world ecological deployment scenarios. The feasibility of integrating deep learning models into biodiversity monitoring systems – especially for field identification of endangered plant species – is rarely discussed from a practical implementation perspective.

The literature demonstrates that CNNs and transfer learning are cutting-edge methods for detailed flower classification. This study addresses the above-mentioned gaps by evaluating eight CNN architectures on a region-specific flower dataset, including endangered Red Book species from Kazakhstan [1].

II. METHODOLOGY

A. Dataset Description

The foundation for the experiments was established using the Flowers Dataset [17]. This resource is provided under the MIT License that permits unlimited use for academic and educational purposes. The compiled dataset includes five different varieties of flowers, namely daisy, dandelion, rose, sunflower, and tulip. Two additional flower species (anomalous peony and crocus korolkowii) were independently added. The pictures were retrieved from publicly available databases to expand the range represented. The original dataset comprised 4,317 images across five flower categories. After adding two classes to the dataset, the completed dataset comprises 4,713 images and seven classes. The dataset was split into 70% training, 20% validation, and 10% testing, and employed a pre-defined random seed. This was done to ensure consistency and strategic category distribution. An early review of the dataset was primarily intended to confirm that each subset preserved accurate label allocation and consistent normalization metrics. Nevertheless, the dataset showed an uneven distribution of categories, where certain types (like dandelion and tulip) featured a greater number of instances than others (anomalous peony and crocus korolkowii).

B. Data Pre-Processing and Augmentation

Before model training, images had to be prepared. For this purpose, preprocessing and augmentation methods were applied to all data, ensuring consistency, improving model effectiveness, and reducing the chances of overfitting. Due to the original dataset containing images of varying sizes and qualities, all images were adjusted to a standardized size of 224×224 pixels, consistent with the expected input dimensions needed by various pre-trained CNNs. Each image was converted from its RGB format into a tensor format and normalized according to the statistical properties of the ImageNet dataset, using mean values of [0.485, 0.456, 0.406] and standard deviation values of [0.229, 0.224, 0.225]. This phase ensured that the input data size matched the weights of the pre-trained model, leading to faster and more reliable convergence during optimization.

To improve the model's capacity for broad application and replicate typical changes in illumination and viewpoint, several methods of data enhancement were used only on the training portion of the dataset. These methods involved randomly mirroring images horizontally, rotation by a maximum of ±10 degrees, and subtly changing intensity, contrast, and saturation. The models were made more resistant to variations in image alignment, color arrangement, and lighting situations because of these adjustments.

The validation and testing subsets, on the other hand, were not subjected to random augmentations. They were only resized, converted to tensors, and normalized to maintain consistency for evaluation and to provide an unbiased assessment of model accuracy.

To maintain consistent results across all experimental runs, a single, unchanging random seed was employed during the stages of both data loading and data partitioning. The data were handled in smaller groups, each containing 32 images, and the training set was shuffled.

C. Training and Validation of CNN Architectures

The main goal of this phase was to assess and contrast various contemporary CNN architectures through transfer learning to identify the most efficient model for flower species recognition. Eight neural network models underwent training and evaluation: AlexNet, VGG16, GoogLeNet, ResNet50, DenseNet121, EfficientNet-B0, MobileNetV2, and YOLOv8-cls. To speed up the learning process and minimize overfitting on the comparatively limited dataset, every CNN architecture began with pre-trained weights from ImageNet. The sole changes implemented involved substituting the existing final layers with newly incorporated fully connected layers that feature seven output neurons, each representing one of the flower categories outlined in the baseline dataset. Every network functioned as an independent experiment, but they all followed the same training procedures to guarantee a fair comparison. The feature extraction modules (convolutional backbones) remained unchanged during the initial phase of training, focusing updates exclusively on the newly integrated classification heads.

Adam optimizer was used to train the models. Since the number of images per class varied, a weighted cross-entropy loss function was used. This helped the model better account for all classes. We also used the ReduceLROnPlateau scheduler, which automatically reduced the learning rate if the error on the test data stopped decreasing. Each model was trained for 10 epochs with a batch size of 32. After training, the model version that demonstrated the best accuracy on the validation data was retained.

A uniform evaluation procedure was used for each experiment: following the training phase, the model's performance was assessed on the held-out test set. This assessment produced measures for accuracy, precision, recall, and F1-score (weighted).

In the third phase of the investigation, the main goal was to develop and train a simple CNN as a foundation for categorizing flower images. The execution was performed utilizing the PyTorch platform within the Python environment. The purpose was to obtain initial results for further research with more complex models. A simple CNN model was created with five convolutional layers. Each layer used Batch normalization, ReLU activation function, and Max pooling. The features were then passed to two fully connected layers. Dropout ($p = 0.4$) was also used to reduce overfitting. The final layer had seven output neurons. Each output corresponded to one flower species: daisy, dandelion, rose, sunflower, tulip, anomalous peony, and crocus korolkowii.

The simple five-layer CNN demonstrated 67.9% accuracy on the validation data. This result was used as a baseline for comparison with the other models. However, since the accuracy was below 70%, the limitations of the baseline models justified the use of advanced, pre-trained CNN architectures that would provide reliable classification.

III. RESULTS

AlexNet provided a reliable deep-learning baseline, showing stable convergence and balanced learning without overfitting (Figure 1).

VGG16 demonstrated good classification capability. Performance metrics were balanced across all classes of flowers in the dataset. The model worked well on both training and validation data, and its hierarchical feature extraction design confirmed its reliability and robustness.

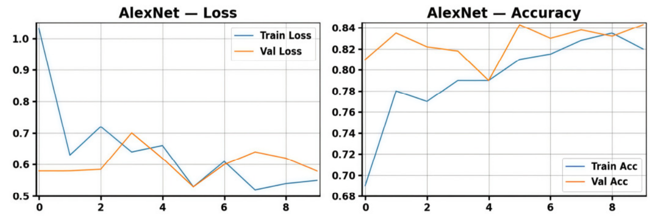

Fig. 1. Training and Validation Curves for the AlexNet model.

By visually inspecting random predictions, findings showed that the model correctly recognized even difficult samples, with varying lighting conditions and background. VGG16 showed stable accuracy for all classes and a high F1-score, which indicates that the model was able to learn detailed flower features. However, this model is more complex and requires more computation time than lightweight networks. (Figure 2).

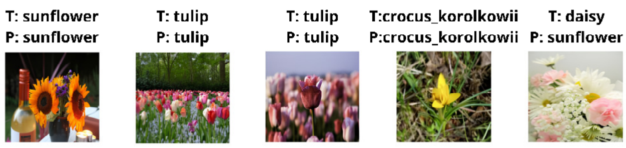

Fig. 2. VGG16 model predictions.

ResNet50 utilizes residual connections that can help the model to learn faster and also avoid the problem of vanishing gradients. The model was trained for 10 epochs, with a learning rate set at 8×10^{-4} and a batch size of 32. To reduce the impact of uneven classes, a weighted cross-entropy loss was used.

Figure 3 shows how the loss decreased, and accuracy and validation F1 per epoch increased. This indicates that the model generalized without overfitting (Figure 3). ResNet50 accurately classified the majority of daisy, dandelion, and rose samples, while most errors occurred between tulip and sunflower, classes that share similar color palettes and textures.

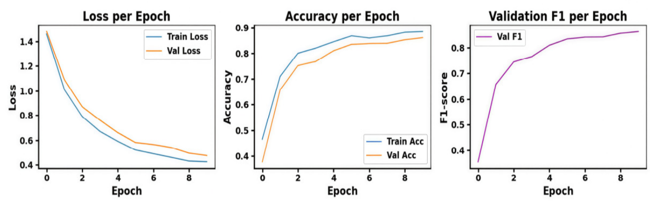

Fig. 3. Training and validation of ResNet50.

Visual inspection of random test predictions confirmed that the network was capable of distinguishing between morphologically distinct flowers, such as crocus korolkowii and daisy, with high confidence. But there were partial misclassifications in categories with overlapping color distributions (rose, tulip, sunflower, and dandelion), demonstrating the natural similarity of flowers, rather than the model's instability.

Advanced Feature Extraction and Scaled Networks Models incorporating complex block designs showed improved efficiency. GoogLeNet achieved balanced performance across classes, displaying a progressive increase in accuracy during training (Figure 4). DenseNet121 demonstrated enhanced feature propagation through dense connectivity, maintaining a consistent upward trend in validation accuracy (Figure 5).

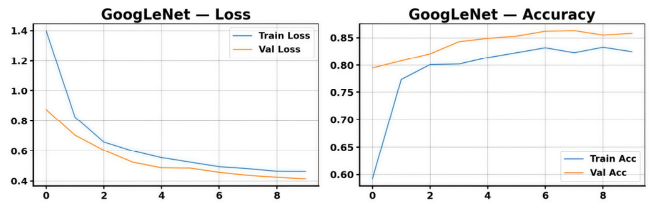

Fig. 4. Training and validation curves for the GoogLeNet model.

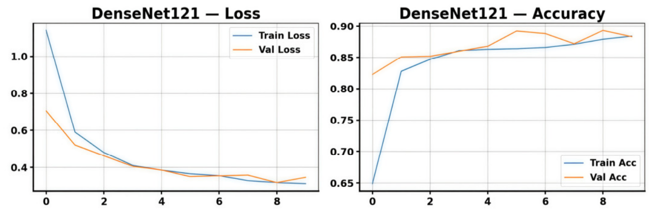

Fig. 5. DenseNet121 loss and accuracy curves.

Notably, DenseNet121 exhibited perfect recall (1.000) for Paeonia anomala and Crocus korolkowii. EfficientNet-B0 applied a compound scaling method, reliably decreasing loss across epochs (Figure 6), making it highly effective for the medium-sized dataset.

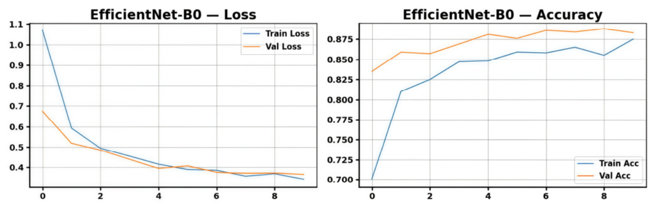

Fig. 6. Training and validation curves for the EfficientNet-B0 model.

The lightweight framework MobileNetV2 achieved an accuracy of 0.871, representing the optimal trade-off between computational efficiency and high recognition performance. Sample predictions generated by MobileNetV2 confirmed its high precision in identifying both common and endagered flower types (Figure 7).

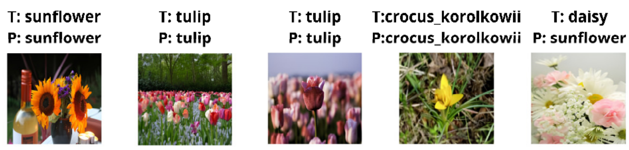

Fig. 7. Sample predictions generated by the MobileNetV2 model.

However, YOLOv8-cls showed the highest performance, achieving the highest accuracy, precision, and recall of 0.919.

The confusion matrix (Figure 8) confirms that YOLOv8 maintained nearly perfect precision across all classes, with highly discriminative feature representation.

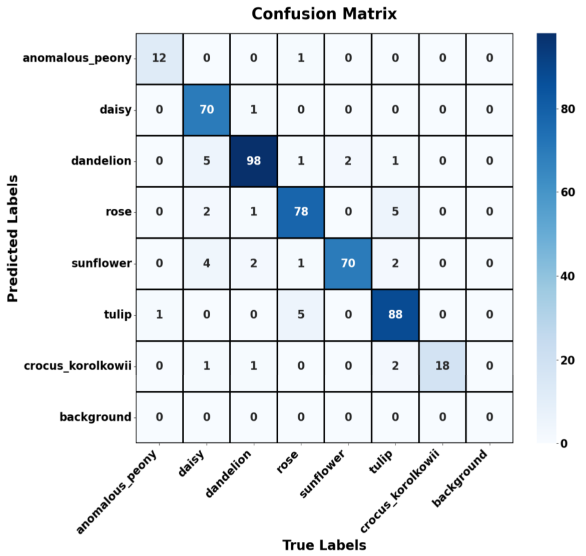

Fig. 8. Confusion matrix for YOLOv8 model on the test dataset.

Observation of test predictions in Figure 9 revealed that the model accurately predicts difficult cases, for example, with bad lighting or different perspectives. Crocus korolkowii, tulip, and daisy were all correctly classified with confidence scores close to 1.0. This shows that YOLOv8-cls is able to distinguish spatial and color cues inherent to each floral species. The inference time of this model was low. This confirms the model's suitability for real-time use.

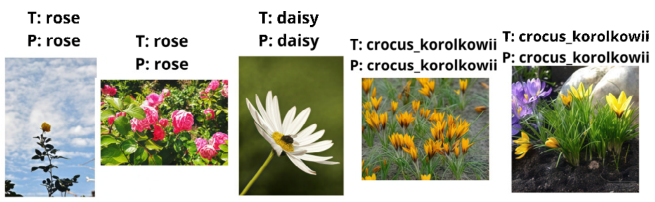

Fig. 9. Sample predictions generated by the YOLOv8 model.

Overall, YOLOv8-cls demonstrated high performance when compared with classical convolutional architectures, achieving the highest accuracy, precision, and recall among all models evaluated in this study. Its combination of speed, robustness, and generalization makes it the most efficient solution for automated flower classification within the proposed experimental framework.

IV. MODEL SELECTION AND COMPARISON RATIONALE

In this study, a total of eight CNN architectures were analyzed for the task of flower image classification. The selection was based on their diversity in structure, computational complexity, and representativeness of different generations of CNNs, from classical deep CNNs to more recent efficient architectures. Thus, the final set of eight models reflects a broad spectrum of CNN evolution, allowing for a comprehensive comparison between traditional deep learning approaches and modern architectures focused on efficiency, scalability, and improved feature extraction.

V. COMPARATIVE PERFORMANCE ANALYSIS

The performance metrics (recall, accuracy, precision, and F1-score) of all considered models can be seen in Table I. These metrics indicate how performance improved with different models.

TABLE I. MODEL COMPARISON BASED ON TEST METRICS

Model	Accuracy	Precision	Recall	F1-score
YOLOv8-cls	0.919	0.919	0.919	0.919
MobileNetV2	0.871	0.875	0.871	0.871
VGG16	0.856	0.858	0.856	0.856
DenseNet121	0.845	0.853	0.845	0.845
EfficientNet-B0	0.837	0.845	0.837	0.838
GoogLeNet	0.837	0.843	0.837	0.837
ResNet50	0.807	0.833	0.807	0.806
AlexNet	0.805	0.818	0.805	0.802

The experimental results of Table I reveal several critical insights into the behavior of different CNN architectures when applied to regional biodiversity data.

Although YOLOv8 is primarily recognized as an object detection framework, its classification variant (YOLOv8-cls) achieved the highest accuracy of 0.919. This performance stems from its advanced internal attention mechanisms and optimized convolutional blocks, which allow the model to effectively isolate floral features from complex and "noisy" natural backgrounds typical of Kaggle datasets. In ecological field photos, where background clutter (grass, soil, or other vegetation) is prevalent, YOLOv8-cls demonstrates superior discriminative power.

A vital finding for conservation efforts is the performance of DenseNet121 regarding endangered flora. While its overall accuracy (0.845) was not the highest, it achieved a perfect recall score (1.000) in Red Book species classification. In the context of environmental monitoring, high recall is more critical than overall precision, as it ensures that no endangered specimens are "missed" by the automated system, thereby providing more reliable data for biodiversity protection. MobileNetV2 as the "Golden Mean": From a practical implementation perspective, MobileNetV2 represents the optimal balance between performance and resource consumption. Despite having a lower accuracy (0.871) compared to YOLOv8, its lightweight architecture makes it the most suitable candidate for integration into mobile applications used by biologists in the field. The low computational overhead and reduced CPU requirements of MobileNetV2 allow for real-time identification in remote areas where high-performance hardware is unavailable.

The progression from AlexNet (0.805) to YOLOv8-cls (0.919) illustrates the evolution of deep learning toward better feature reuse and scalability. While traditional models like VGG16 (0.856) remain robust, modern architectures optimized for specific constraints such as recall or efficiency for mobile deployment offer more targeted solutions for regional ecological research.

YOLOv8-cls showed the best performance, due to its modern design, efficient feature extraction, and deep convolutional blocks. Even though this model is initially known as an object detector, its classification handled the flower dataset very well. MobileNetV2 model achieved an accuracy of 0.871, which demonstrates its ability to maintain competitive accuracy while requiring significantly fewer computational resources. Therefore, this model can be used on mobile devices under real-time conditions. The VGG16 model achieved an accuracy of 0.856. DenseNet121 achieved a similar result of 0.845. It performs effectively because it transfers and utilizes learned features well. EfficientNet-B0 and GoogLeNet achieved an accuracy of 0.837. These models provide a good balance between accuracy and speed. ResNet50 achieved an accuracy of 0.807. Its structure helps make training more stable, but it doesn't always yield the highest accuracy, especially with a small dataset. AlexNet achieved an accuracy of 0.805, due to the fact that it is an older and simpler model.

From the comparison of different models, it is noticeable how, over time, performance improved gradually. A good balance is achieved with models like DenseNet121, EfficientNet-B0, and MobileNetV2, while YOLOv8-cls represents a new step in evolution. This highlights how CNN models become more optimized and scalable.

The outcomes of the experiment demonstrated that convolutional structures are capable of independently extracting significant attributes from unprocessed flower images without any form of manual preparation. The basic CNN's effectiveness was, however, constrained because of its somewhat basic arrangement, coupled with the unequal representation of flower types. The neural network occasionally gave too much importance to more common flower species like tulips and dandelions, but performed poorly on less common species like anomalous peonies.

In spite of these limitations, the experiment efficiently determined the initial benchmark accuracy rate (roughly 68%) and offered crucial observations that would be used to make subsequent improvements in this part, which included transfer learning, class equalization, and the assessment of several CNN structures that had previously been trained.

VI. CONCLUSION

This research demonstrates the transformative potential of deep convolutional neural networks (CNNs) in the automated classification of both common and endangered floral species. By evaluating eight distinct architectures, this study provides a comprehensive benchmark for biodiversity monitoring.

A. Contribution

The primary contribution of this work lies in the development of a region-specific dataset that integrates rare endemic flora from the Red Book of Kazakhstan [1], specifically Paeonia anomala and Crocus korolkowii. Unlike previous studies that focus on global benchmark datasets, this research addresses the critical gap in monitoring geographically localized ecosystems in Central Asia. Furthermore, the study validates the high efficacy of detection-origin architectures, such as YOLOv8-cls, for pure classification tasks in complex ecological contexts.

B. Key Findings and Comparative Analysis

- YOLOv8-cls emerged as the best-performing model, achieving a peak accuracy and precision of 0.919, which represents a significant evolution over traditional deep models.
- MobileNetV2 (0.871) and VGG16 (0.856) demonstrated strong performance, with MobileNetV2 offering the most efficient trade-off between computational speed and recognition precision for mobile deployment.
- Architectures such as DenseNet121 and EfficientNet-B0 proved the advantages of dense connectivity and composite scaling, particularly in maintaining high recall for rare species.

The results confirm that transfer learning and data augmentation are essential strategies for overcoming class imbalances when working with limited datasets of endangered species.

These findings underscore that modern CNNs are not merely theoretical tools but are ready for integration into portable ecological instruments and real-time biodiversity tracking systems. By automating the identification of Red Book species, this technology directly aids in the preservation of Kazakhstan's natural heritage and the creation of digital botanical collections

C. Future Work

Future work will focus on expanding the dataset to include more regional varieties and optimizing models for edge computing in field conditions.

DECLARATION OF COMPETING INTERESTS

The authors declare no conflicts of interest.

ACKNOWLEDGMENT

Not applicable in this study.

DATA AVAILABILITY

The data supporting the findings of this study are available in [17] and from the corresponding author upon reasonable request.

REFERENCES

[1] *The Red Data Book of Kazakhstan: Plants*. Institute of Botany and Phytointroduction, 2014.

[2] S. A. Kubentayev *et al.*, "Current State of Natural Populations of Paeonia anomala (Paeoniaceae) in East Kazakhstan," *Diversity*, vol. 15, no. 11, Oct. 2023, https://doi.org/10.3390/d15111127.

[3] T. Abbas *et al.*, "Deep Neural Networks for Automatic Flower Species Localization and Recognition," *Computational Intelligence and Neuroscience*, vol. 2022, no. 1, 2022, Art. no. 9359353, https://doi.org/10.1155/2022/9359353.

[4] H. Hiary, H. Saadeh, M. Saadeh, and M. Yaqub, "Flower classification using deep convolutional neural networks," *IET Computer Vision*, vol. 12, no. 6, pp. 855–862, 2018, https://doi.org/10.1049/iet-cvi.2017.0155.

[5] C. Narvekar and M. Rao, "Flower classification using CNN and transfer learning in CNN- Agriculture Perspective," in *2020 3rd International Conference on Intelligent Sustainable Systems (ICISS)*, Sept. 2020, pp. 660–664, https://doi.org/10.1109/ICISS49785.2020.9316030.

[6] N. M. Rahim, M. A. Hairuddin, M. S. A. M. Ali, N. M. Tahir, A. A. Almisreb, and N. D. K. Ashar, "Pretrained Convolutional Neural Network for Fruit Classification Analysis of Pineapple Plantation Images," *Engineering, Technology & Applied Science Research*, vol. 15, no. 2, pp. 20819–20826, Apr. 2025, https://doi.org/10.48084/etasr.9249.

[7] S. Nuanmeesri, "A Hybrid Deep Learning and Optimized Machine Learning Approach for Rose Leaf Disease Classification," *Engineering, Technology & Applied Science Research*, vol. 11, no. 5, pp. 7678–7683, Oct. 2021, https://doi.org/10.48084/etasr.4455.

[8] A. I. R. H. Barat, W. S. Astuti, D. Hartama, A. P. Windarto, and A. Wanto, "Enhanced flower image classification using MobileNetV2 with optimized performance," *JITK (Jurnal Ilmu Pengetahuan dan Teknologi Komputer)*, vol. 11, no. 1, pp. 180–191, Aug. 2025, https://doi.org/10.33480/jitk.v11i1.6497.

[9] G. Rajkomar and S. Pudaruth, "A Mobile App for the Identification of Flowers Using Deep Learning," *International Journal of Advanced Computer Science and Applications*, vol. 14, no. 5, pp. 76–102, 2023, https://doi.org/10.14569/IJACSA.2023.0140508.

[10] M. A. Valijonovich, "Bioecological Features and Anatomical Structure of <i>Crocus korolkovii</i> Regel & Maw in the Introduction of Uzbekistan," *American Journal of Plant Sciences*, vol. 08, no. 03, pp. 471–481, 2017, https://doi.org/10.4236/ajps.2017.83032.

[11] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016, pp. 770–778, https://doi.org/10.1109/CVPR.2016.90.

[12] T.-M. Dutschmann, L. Kinzel, A. ter Laak, and K. Baumann, "Large-scale evaluation of k-fold cross-validation ensembles for uncertainty estimation," *Journal of Cheminformatics*, vol. 15, no. 1, Apr. 2023, Art. no. 49, https://doi.org/10.1186/s13321-023-00709-9.

[13] J. Deng, W. Dong, R. Socher, L.-J. Li, Kai Li, and Li Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *2009 IEEE Conference on Computer Vision and Pattern Recognition*, June 2009, pp. 248–255, https://doi.org/10.1109/CVPR.2009.5206848.

[14] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "MobileNetV2: Inverted Residuals and Linear Bottlenecks," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, June 2018, pp. 4510–4520, https://doi.org/10.1109/CVPR.2018.00474.

[15] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely Connected Convolutional Networks," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017, pp. 2261–2269, https://doi.org/10.1109/CVPR.2017.243.

[16] M. Tan and Q. Le, "EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks," in *Proceedings of the 36th International Conference on Machine Learning*, May 2019, pp. 6105–6114.

[17] tensorflow, "models/research/slim/datasets/flowers.py at master tensorflow/models," *GitHub*. https://github.com/tensorflow/models/blob/master/research/slim/datasets/flowers.py.