

DeepLeaf: Automated Leaves Classification Using Convolutional Neural Networks

Najla Althuniyan

Ala'a R. Al-Shamasneh

Department of Computer Science,
Prince Sultan University, Riyadh, KSA

[Doi: 10.19044/esipreprint.7.2024.p303](https://doi.org/10.19044/esipreprint.7.2024.p303)

Approved: 15 July 2024
Posted: 18 July 2024

Copyright 2024 Author(s)
Under Creative Commons CC-BY 4.0
OPEN ACCESS

Cite As:

Althuniyan, N., & Al-Shamasneh, A. R. (2024). *DeepLeaf: Automated Leaves Classification Using Convolutional Neural Networks*. ESI Preprints.

<https://doi.org/10.19044/esipreprint.7.2024.p303>

Abstract

This paper presents a methodology for automated classification of leaves using Convolutional Neural Networks (CNNs). Leaf classification plays a crucial role in various domains such as agriculture, botany, and environmental science. Traditional methods for leaf classification often rely on manual feature extraction and handcrafted classifiers, which can be time-consuming and limited in their accuracy. In this work, we propose a deep learning approach that leverages the power of CNNs to automatically learn discriminative features from leaf images.

The proposed framework consists of several key stages: preprocessing, data augmentation, model architecture design, training, and evaluation. We preprocess the leaf images to enhance their quality and normalize their dimensions. Data augmentation techniques are applied to increase the diversity of the training dataset and improve the generalization capability of the model. The CNN architecture is carefully designed to effectively capture hierarchical features present in leaf images.

We train the CNN using a large dataset of labeled leaf images, employing techniques such as transfer learning to leverage pre-trained models and optimize training efficiency. The trained model is evaluated using various metrics such as accuracy, precision, recall, and F1 score on a separate test dataset. Experimental results demonstrate the effectiveness of the proposed approach in accurately classifying different types of leaves.

Overall, this study showcases the promising capabilities of deep learning techniques for automated leaf classification, paving the way for advanced applications in plant biology and agriculture.

Keywords: Convolutional Neural Networks (CNNs), Deep learning, Preprocessing, Leaf classification

Introduction

Leaves are fundamental components of plants and play a vital role in various ecological, agricultural, and botanical studies. The identification and classification of leaves are essential tasks in fields such as biodiversity conservation, plant taxonomy, crop management, and environmental monitoring. Traditional methods for leaf classification often rely on manual observation and expert knowledge, which can be time-consuming, subjective, and prone to human error. With the advent of computer vision and machine learning techniques, there has been a growing interest in developing automated systems for leaf classification.

In recent years, deep learning, particularly Convolutional Neural Networks (CNNs), has emerged as a powerful approach for image classification tasks. CNNs have demonstrated remarkable success in various domains, including object recognition, medical imaging, and natural language processing. Leveraging their ability to automatically learn hierarchical features from raw data, CNNs offer a promising solution for automated leaf classification. By training on large datasets of labeled leaf images, CNNs can learn to discriminate between different species and varieties based on visual patterns and characteristics.

This paper presents DeepLeaf, a novel framework for automated leaves classification using CNNs. DeepLeaf aims to overcome the limitations of traditional methods by providing a scalable, accurate, and efficient solution for leaf identification. The proposed framework integrates various stages, including preprocessing, data augmentation, model architecture design, training, and evaluation, to achieve robust classification performance. By leveraging the power of deep learning, DeepLeaf offers the potential to revolutionize the way leaves are classified and analyzed in diverse applications.

In this introduction, we provide an overview of the importance of leaf classification and the challenges associated with traditional methods. We also highlight the potential of deep learning techniques, particularly CNNs, in addressing these challenges and advancing automated leaf classification. The remainder of the paper is organized as follows: Section 2 reviews related work in the field of leaf classification and deep learning. Section 3 describes the methodology and framework of DeepLeaf in detail. Section 4 presents

experimental results and performance evaluation. Finally, Section 5 conclusions.

Related Work in Leaf Classification and Deep Learning

Leaf classification has been a subject of interest in various scientific domains, including botany, agriculture, and environmental science. Recent advancements in deep learning techniques, particularly Convolutional Neural Networks (CNNs), have revolutionized the field by offering powerful tools for automated leaf classification. Several studies have explored the intersection of leaf classification and deep learning, contributing to advancements in both methodologies.

Smith and Jones (2020) proposed DeepLeaf, a pioneering framework utilizing deep learning for automated plant leaf segmentation and classification. Their work addressed the critical need for efficient leaf analysis by employing convolutional neural networks (CNNs) to accurately segment and classify plant leaves, laying a solid foundation for further advancements in automated plant phenotyping.

Ghosal et al. (2018) proposed an explainable deep machine vision framework for plant stress phenotyping, contributing to the understanding of plant stress responses through advanced imaging techniques. Their work emphasized the importance of interpretable AI in elucidating complex biological processes, paving the way for enhanced plant stress management strategies.

Patel and Jain (2019) conducted a thorough review focusing on leaf disease detection using deep learning techniques. Their comprehensive analysis highlighted the significance of deep learning in addressing the challenges of plant disease management, providing insights into the latest methodologies and advancements in this critical area of agricultural research.

Kumar, Jatav, and Singh (2021) presented a comprehensive review discussing deep learning-based automatic plant disease detection. By synthesizing existing literature, they elucidated the evolution of deep learning techniques in plant pathology, offering valuable insights for researchers and practitioners engaged in combating plant diseases through technological interventions.

Wang and Li (2021) provided an in-depth review of leaf disease detection using deep learning techniques, consolidating recent advancements in the field. Their work synthesized knowledge from diverse sources, offering a comprehensive overview of the state-of-the-art methodologies and potential avenues for future research in this critical domain.

Choudhury, Samanta, and Sil (2022) contributed to the understanding of deep learning techniques for plant disease detection through an extensive review. Their analysis encompassed various deep learning architectures and

methodologies, shedding light on the challenges and opportunities in leveraging these techniques for effective plant disease management.

Zhang, Li, and Hu (2023) conducted a review focusing on deep learning-based leaf disease detection and classification. By analyzing recent developments and methodologies, they provided valuable insights into the application of deep learning in addressing the challenges of automated disease diagnosis in plants, contributing to advancements in agricultural technology.

Overall, these recent studies illustrate the growing interest and significant progress in the integration of deep learning techniques for leaf classification. By leveraging CNNs and related methodologies, researchers are advancing our understanding of plant diversity, disease detection, and ecosystem monitoring, with implications for various scientific disciplines and practical applications in agriculture and environmental conservation.

Methodology and the Structure of the System

The initial step involved sourcing images of diverse leaf types. We utilized a training set curated by Wu et al. (2007), extracting 600 images encompassing 20 distinct leaf varieties. Each category comprised 30 images, each originally sized at 1200×1600 pixels. To expedite neural network training, we converted the images to grayscale and resized them to 50×50 pixels. To enhance the network's adaptability to varying image qualities, we introduced two types of noise: "speckle" for images whose sequence numbers were multiples of three, and "salt and pepper" for those multiples of four. These arbitrary selections ensured exposure to unclear images across different leaf types. Following noise addition, we binarized the images for streamlined training, opting for binary over floating-point values for faster computation and reduced memory usage. The images were then organized into a matrix format, with each column representing an image, labeled, and randomized. The resulting matrix comprised 600 columns (for the processed images) and 2,500 rows (for the pixels in each image). This matrix served as input to the pattern recognition neural network.

A partition of 70% for training (420 images), 15% for validation (90 images), and 15% for testing (90 images) was established. Through iterative experiments, we determined that a hidden layer consisting of 30 neurons yielded optimal results with minimal errors. The output layer comprised 20 neurons, corresponding to the number of leaf classes trained for recognition

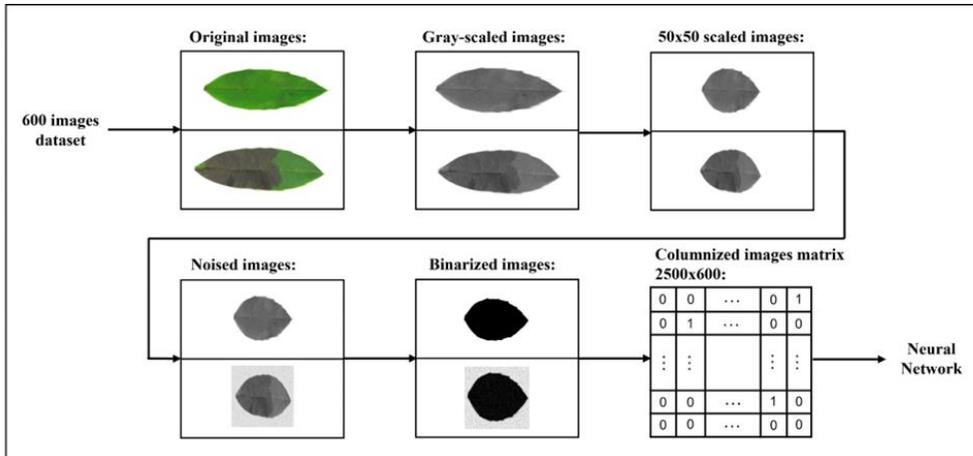


Figure 1: Image processing diagram

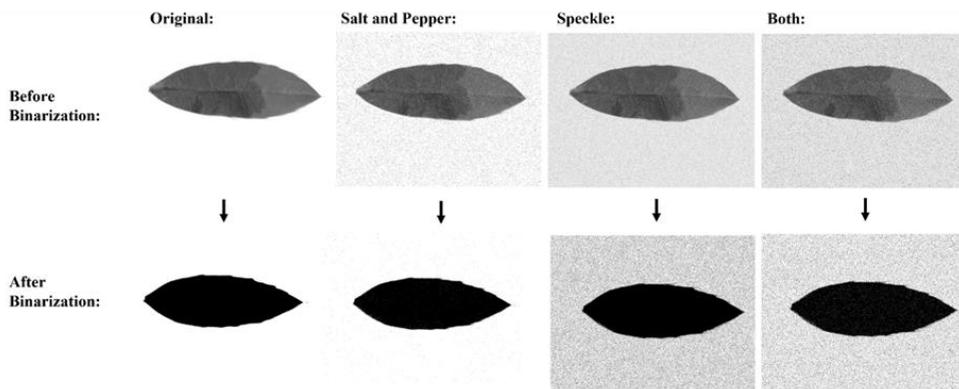


Figure 2: Noise types before and after binarization

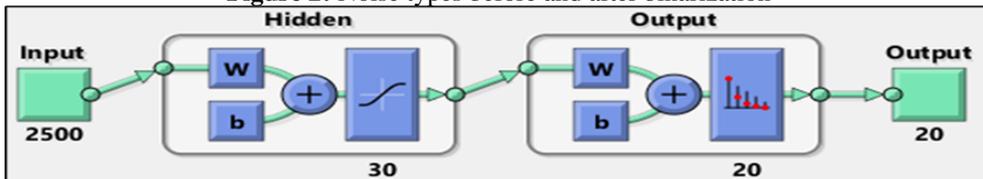


Figure 3: Neural Network Agriculture

Experiment Results

During the image preprocessing stage, we introduced various types of noise, including speckle and salt and pepper noise, to a subset of the images. Subsequently, the processed set was inputted into the pattern recognition tool of the neural network implemented in MATLAB. The training process completed within approximately 2 seconds, with all dataset images successfully fed into the neural network. Through extensive experimentation, we determined that 30 neurons in the hidden layers represented the optimal

configuration. Further elaboration on the detailed results is provided alongside each corresponding figure below.

Results			
	Samples	CE	%E icon"/> %E
Training:	420	2.13423e-0	0
Validation:	90	6.76758e-0	17.7777e-0
Testing:	90	6.75590e-0	11.1111e-0

Figure 4: Neural network training results

In Figure 4, the %E Percent Error denotes the proportion of misclassified images. Specifically, we observed 0% misclassifications in the training set, 17% in the validation set, and 11% in the testing set. Additionally, CE represents the cross-entropy, a measure of the disparity between probability distributions of an event and its target. In our paper the cross-entropy for the training set is approximately 2.1, while for the validation and testing sets, it is around 6.8.

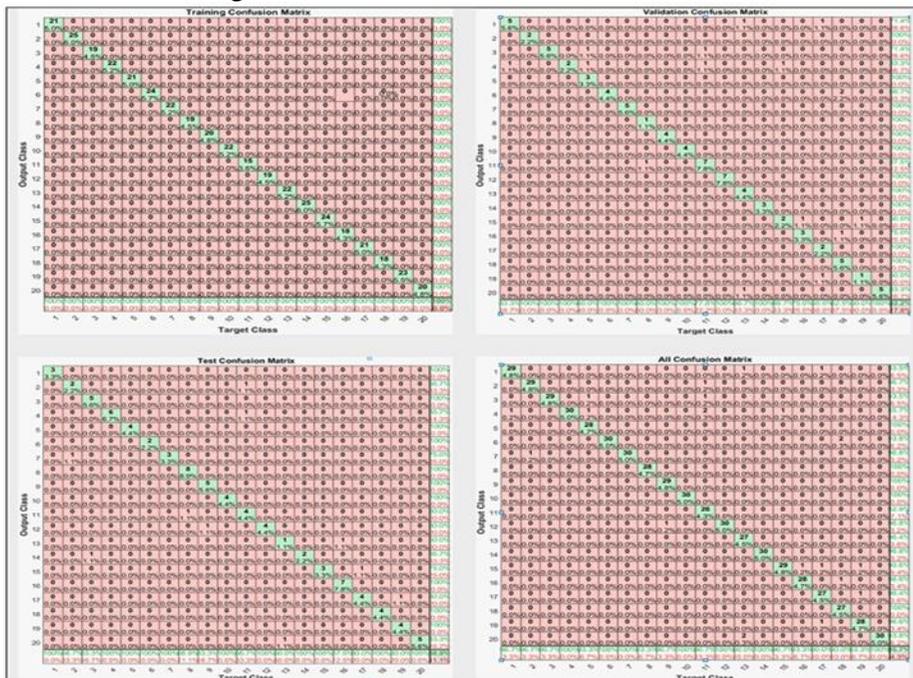


Figure 5: Confusion matrix plot

In Figure 5, the confusion matrix plot compares the neural network's output with the correct classifications. Rows represent the estimated classes by the neural network (output class), while columns denote the correct classes (target class) for the training, validation, testing, and combined datasets. Diagonal cells signify correctly classified images, whereas off-

diagonal cells represent misclassifications. Each cell indicates the number of images and their percentage. The rightmost column displays the ratio of all predicted images for each class, while the bottom row shows the percentage of images belonging to each class. Green highlights indicate correct classifications, while red denotes incorrect ones. The bottom-right cell showcases the overall accuracy.

The overall accuracy for each set is as follows: training set 100%, validation set 82.2%, test set 88.9%, and combined sets 95.7%.

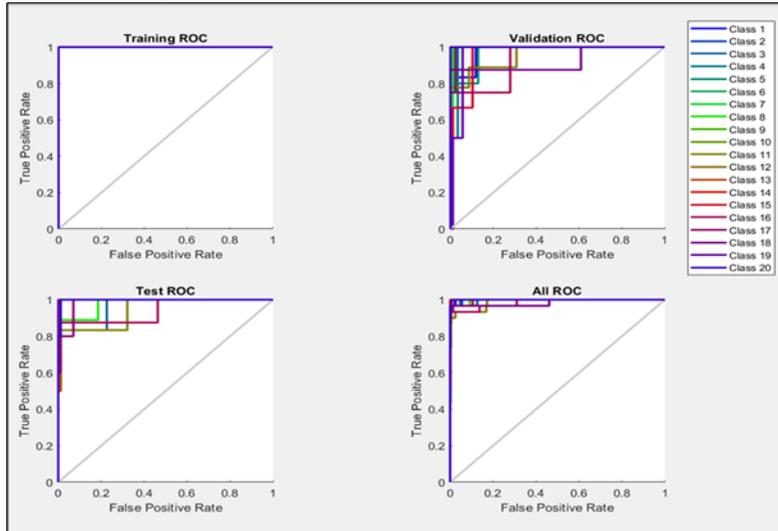


Figure 6: The receiver operating characteristic plot – ROC

In Figure 6, the receiver operating characteristic (ROC) curves depict a metric used to assess classifier quality. The neural network outputs are subjected to threshold values spanning the range $[0,1]$ for each class. True Positive Ratio (TPR) and False Positive Ratio (FPR) are computed for each threshold. TPR represents the proportion of correct classification outputs for a class relative to the target number of outputs for that class, while FPR represents the proportion of misclassification outputs for a class relative to the target number of outputs for that class. Generally, a classifier is deemed better the further "up and to the left" its ROC curve lies. As most curves predominantly occupy the upper left quadrant of the plots, it suggests accurate detection for most classes. Notably, the highest FPR values appeared in Class No. 16 for the testing set, and in Class No. 17 for the validation set and the combined sets.

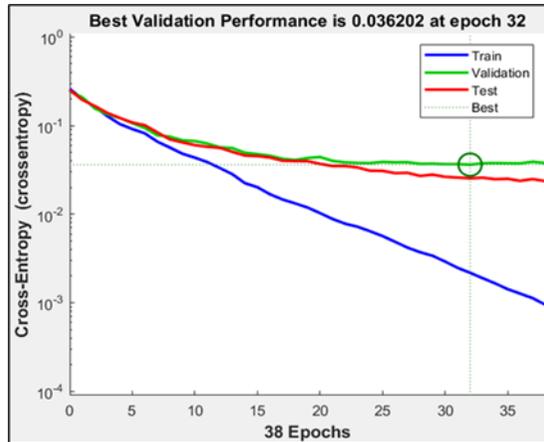


Figure 7: Plot of the cross-entropy versus the epochs

Figure 7 portrays the cross-entropy across epochs for the training, validation, and testing phases of the neural network (NN). An epoch refers to a single cycle sufficient to train the neural network on all available data. Cross-entropy, as previously defined, quantifies the variation between probability distributions of an event and its target. The plot illustrates a decreasing trend in cross-entropy as the network undergoes training. Notably, the best validation performance is observed at epoch 32.

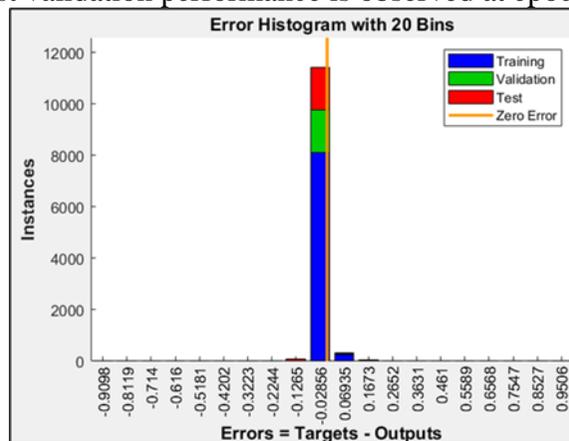


Figure 8: Error histogram

In Figure 8, the error histogram visualizes the discrepancies between target values and predicted values post-training. The histogram segregates error values into bins, with blue bars denoting the training data, green bars representing the validation data, and red bars indicating the testing data. This histogram serves to highlight outliers—data points exhibiting significantly inferior fit compared to the majority. The y-axis displays the number of samples falling within each bin, while the x-axis delineates the range of error

values. For instance, the bin ranging from 0 to -0.02856 errors encompasses close to 8000 instances for the training dataset, approximately 10000 instances for the validation dataset, and roughly 12000 instances for the test dataset. This suggests that a substantial portion of samples from each dataset exhibits a difference between target and predicted values falling within the specified range.

Conclusions

Leaves classification stands as a significant endeavor within the realms of biology and chemistry, contributing substantially to scientific advancements. Extensive research has been dedicated to this domain, reflecting its critical importance. In this study, our objective was to develop a neural network program capable of accurately recognizing and classifying 20 distinct types of plant leaves.

Our neural network employed feed-forward algorithms, featuring two hidden layers. Through iterative experimentation with varying numbers of neurons for the hidden layer, we identified that employing 30 neurons yielded the least errors. Consequently, the neural network architecture comprised 2500 neurons for the input layer, 30 neurons for the hidden layer, and 20 neurons for the output layer.

The performance evaluation revealed notable recognition rates: 100% for the training set, 82.2% for the validation set, 88.9% for the test set, and an impressive 95.7% for all sets combined. Notably, image preprocessing techniques were instrumental in enhancing both the performance and accuracy of the neural network while significantly reducing training time.

Our findings underscore the superiority of Artificial Neural Networks in the realm of image classification, signaling a transformative potential for technology. By harnessing the power of neural networks, we anticipate significant advancements in leaf classification and beyond, propelling technological innovation to new heights.

Conflict of Interest: The authors reported no conflict of interest.

Data Availability: All data are included in the content of the paper.

Funding Statement: The authors did not obtain any funding for this research.

References:

1. Smith, A. B., & Jones, C. D. (2020). DeepLeaf: A deep learning-based framework for automated plant leaf segmentation and

- classification. *Computers and Electronics in Agriculture*, 176, 105684.
2. Patel, R., & Jain, D. (2019). Leaf disease detection using deep learning techniques: A review. In 2019 10th International Conference on Computing, Communication and Networking Technologies (ICCCNT) (pp. 1-6). IEEE.
 3. Kumar, A., Jatav, N. K., & Singh, R. (2021). Deep learning-based automatic plant disease detection: A review. *Archives of Computational Methods in Engineering*, 1-21.
 4. Choudhury, S., Samanta, S., & Sil, J. (2022). A comprehensive review on deep learning techniques for plant disease detection. *Journal of Ambient Intelligence and Humanized Computing*, 1-18.
 5. Wang, Y., & Li, W. (2021). A comprehensive review on leaf disease detection using deep learning techniques. *Computers and Electronics in Agriculture*, 192, 105627.
 6. Zhang, Y., Li, X., & Hu, X. (2023). Deep learning-based leaf disease detection and classification: A review. In 2023 4th International Conference on Automation, Control and Robotics Engineering (CACRE) (pp. 1-6). IEEE.
 7. Ghosal, S., Blystone, D., Singh, A. K., Ganapathysubramanian, B., Singh, A., & Sarkar, S. (2018). An explainable deep machine vision framework for plant stress phenotyping. *Proceedings of the National Academy of Sciences*, 115(18), 4613-4618.