

Automated identification of tomato leaf pathologies using deep learning via ResNet18 and a Tailored CNN Architecture

Yanto Supriyanto

STMIK IKMI Cirebon, Jl. Perjuangan No.10B, Karyamulya, Kec. Kesambi, Kota Cirebon, Jawa Barat 45135

Corresponding Author: yanto200445@gmail.com

Submitted: 3/7/2025

Revised: 29/7/2025

Accepted: 7/8/2025

ABSTRACT

Tomato leaves (*Solanum lycopersicum*) are susceptible to various diseases that significantly impact both yield and quality in agricultural production. To enhance the effectiveness of precision agriculture, deep learning-based image classification techniques have emerged as reliable tools for automatically detecting disease symptoms. In this study, two deep learning models are investigated for this purpose: a custom-built Convolutional Neural Network (CNN) and the ResNet18 architecture, which leverages transfer learning. The experimental workflow encompasses preprocessing of input images, data augmentation strategies, architectural development of the custom CNN, and the fine-tuning phase of the ResNet18 model. The evaluation was conducted on a validation dataset comprising 1,000 images evenly distributed across ten disease categories. Results show that the ResNet18 model attained a validation accuracy of 74%, whereas the custom CNN model achieved 55%. While the latter demonstrated lower predictive performance, it offers advantages in computational simplicity and execution speed. These findings suggest that transfer learning with ResNet18 is more suitable for complex, multi-class classification problems on limited datasets, whereas the lightweight CNN model may be better positioned for deployment on low-resource, edge-based agricultural systems.

Keywords: Deep learning; ResNet18; custom CN; tomato leaf disease; image classification

1. Introduction

Tomato (*Solanum lycopersicum*) is a horticultural crop with substantial economic significance and is extensively cultivated across various regions in Indonesia. Despite its agricultural importance, crop productivity is frequently hampered by leaf diseases, including bacterial spot, early and late blight, leaf mold, mosaic virus, septoria leaf spot, spider mite infestations, target spot, and yellow leaf curl virus [1]. These diseases not only manifest as visible damage on the foliage but also negatively impact both crop quality and yield. Traditionally, disease detection in crops has relied on manual inspection by farmers or agricultural advisors. However, this method is inherently subjective and susceptible to inconsistencies, as it depends heavily on individual skill and experience in identifying disease symptoms [2]. With the progression of precision agriculture, the use of artificial intelligence—especially deep learning—has gained traction as a promising solution. Among various AI techniques, Convolutional Neural Networks (CNNs) have been recognized for their ability to enhance diagnostic accuracy and efficiency [3]. One well-regarded model for image-based classification is ResNet18, which leverages residual learning to extract intricate features and mitigate the vanishing gradient problem that commonly arises in deep neural networks [4]. Alternatively, custom-designed CNNs offer greater adaptability to the specific nature of a dataset, although they often require intensive hyperparameter tuning to reach optimal performance levels [5]. Furthermore, employing strategies such as data augmentation and transfer learning has proven beneficial, especially when working with relatively small datasets [6]. In light of these considerations, this study explores and compares the performance of two deep learning



frameworks; ResNet18 utilizing transfer learning and a handcrafted CNN, for classifying various tomato leaf diseases. The models are quantitatively evaluated based on four key metrics: accuracy, precision, recall, and F1-score. The ultimate goal is to contribute to the development of a robust, image-driven diagnostic system that can support precision agriculture practices in Indonesia [7].

2. Method

This study followed a structured sequence of steps encompassing data acquisition, preprocessing of images, model development, and performance evaluation in classification tasks. Figure 1 primary aim was to examine and contrast two deep learning strategies namely, a manually constructed Convolutional Neural Network (CNN) and the ResNet18 architecture enhanced through transfer learning in classifying tomato leaf diseases. The dataset employed in this study was sourced from the publicly available Tomato Leaf Disease Dataset on Kaggle. It comprises ten disease categories, including bacterial spot, early blight, and yellow leaf curl virus. The dataset was partitioned into training and validation sets, with each class evenly represented. During the preprocessing phase, all images were resized to 224×224 pixels and standardized using ImageNet normalization parameters [8], ensuring input compatibility with ResNet18. Increase the model’s generalization capability and minimize overfitting, data augmentation techniques such as random rotations, horizontal flips, and other variations were applied to the training set. ResNet18 was utilized as a pretrained backbone within a transfer learning framework, where only the final classification layer was replaced and retrained according to the number of output classes. The earlier layers remained fixed to function as generic feature extractors [6]. Model training was carried out in Google Colab, taking advantage of the T4 GPU for improved computational speed [7], Both models were trained for 10 epochs using the Adam optimizer, a batch size of 32, a learning rate of 0.001, and CrossEntropyLoss as the loss function. Performance assessment, the models were evaluated using standard classification metrics: accuracy, precision, recall, and F1-score, alongside the use of a confusion matrix to further analyze class-wise prediction outcomes.

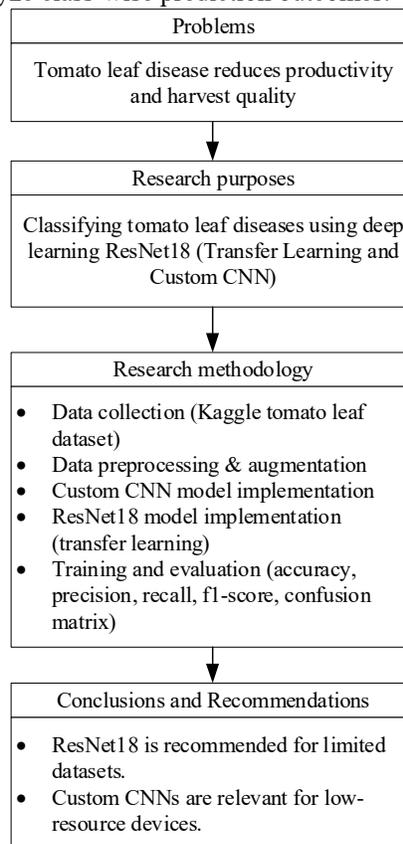


Figure 1. Research method flowchart for tomato leaf disease classification using CNN Custom and ResNet18



Dataset

The dataset utilized in this study originates from the Tomato Leaf Disease collection, created by publicly accessible on the Kaggle platform. It comprises 10 distinct disease categories, with each class containing 100 images for model training and another 100 for validation, yielding a total of 2,000 images with equal class distribution. All images were organized using a directory hierarchy that separates training and validation sets. Each image class is placed into dedicated subfolders, adhering to the PyTorch ImageFolder format. The raw images are provided in RGB color space and exhibit diverse native resolutions. During the preprocessing phase, they were resized to a uniform dimension of 224×224 pixels to meet the input specification requirements of both the ResNet18 and the custom CNN architectures.

Model architecture and data transformation

This research follows a structured methodology consisting of several phases, including data acquisition, image preprocessing, model training, and evaluation of classification performance. The primary objective is to compare two deep learning models: a manually constructed CNN (Custom CNN) and ResNet18, which incorporates transfer learning techniques. The dataset employed is the Tomato Leaf Disease collection retrieved from Kaggle [9], comprising 10 distinct disease categories, with 1,000 images used for training and another 1,000 for validation. The dataset was organized to fit the PyTorch ImageFolder structure, where each class is allocated to its own subfolder [9] [10]. During preprocessing, all images were resized to 224×224 pixels and normalized using ImageNet statistical values [11], to align with ResNet18 input expectations. For training, augmentation techniques such as $\pm 15^\circ$ random rotations and horizontal flips were employed to improve generalizability and reduce overfitting risks. The validation set, in contrast, was only resized and normalized [12].

```
Distribusi data train:
Tomato__Target_Spot: 100
Tomato__Mosaic_virus: 100
Tomato__Leaf_Mold: 100
Tomato__Early_blight: 100
Tomato__Septoria_leaf_spot: 100
Tomato__Yellow_Leaf_Curl_Virus: 100
Tomato__Spider_mites: 100
Tomato__Healthy: 100
Tomato__Late_blight: 100
Tomato__Bacterial_spot: 100

Distribusi data validation:
Tomato__Target_Spot: 100
Tomato__Spider_mites: 100
Tomato__Septoria_leaf_spot: 100
Tomato__Bacterial_spot: 100
Tomato__Leaf_Mold: 100
Tomato__Yellow_Leaf_Curl_Virus: 100
Tomato__Late_blight: 100
Tomato__Healthy: 100
Tomato__Mosaic_virus: 100
Tomato__Early_blight: 100
```

Figure 2. Dataset distribution

The Custom CNN was developed with an emphasis on simplicity and efficiency. It comprises three convolutional layers activated with ReLU and followed by max pooling, ending with a fully connected layer and a softmax output layer. Meanwhile, ResNet18 was applied as a pretrained model where only the final classification layer was fine-tuned, and the rest of the layers were kept static to act as feature extractors [13]. Both models were trained on Google Colab utilizing a T4 GPU environment [14], over the span of 10 epochs using the Adam optimizer, a batch size of 32, and a learning rate of 0.001. The loss function adopted was CrossEntropyLoss. Evaluation metrics include accuracy, precision, recall, and F1-score, along with a confusion matrix to examine prediction discrepancies across individual classes. A complete visual representation of the method is presented in Figure 1.

ResNet18 model implementation with transfer learning

In this study, ResNet18 is employed as a pretrained model to perform classification of tomato leaf diseases using a transfer learning approach. This model was chosen due to its efficiency in computation and its established performance in a wide range of image classification tasks, including those involving

plant pathology. The model is compatible with various hardware configurations and is run on a GPU when available, or on a CPU otherwise. ResNet18 is initialized with ImageNet-pretrained weights, allowing it to identify essential visual patterns such as shape boundaries, textures, and colors, an important feature for processing datasets of limited size. The implementation in this study, ResNet18 functions as a feature extractor, in which all convolutional layers are frozen, preserving the learned parameters. Only the final fully connected layer is adjusted and retrained to align with the ten classification labels in the dataset. If necessary, the entire model can be fine-tuned by unfreezing all layers, although this would increase the training time. Because the original ResNet18 is structured for 1,000-class classification, the final classification layer is replaced with a new one whose number of output units matches the number of disease categories in this task. Once model modification is complete, it is transferred to the appropriate device (GPU or CPU) for optimal execution during both training and evaluation, based on the available system resources. The ResNet18 architecture begins with an initial 7×7 convolutional layer, followed by batch normalization and a ReLU activation function. It proceeds through four residual blocks, with progressively increasing channel dimensions: 64, 128, 256, and 512. Each residual block incorporates skip connections that enable stable backpropagation of gradients during training. This design choice effectively addresses the vanishing gradient problem, ensuring that the network can continue learning as it deepens. Each block from layer1 to layer4 includes two convolutional layers and identity connections, facilitating smooth gradient flow during deep network optimization.

Custom CNN model implementation

Besides employing ResNet18 with a transfer learning scheme, a custom Convolutional Neural Network was specifically developed to classify tomato leaf diseases. The model was constructed with a focus on simplicity and efficiency, aiming to accommodate the constraints and unique traits of the dataset used. The architecture comprises three primary convolutional blocks, each employing filter with a 3×3 kernel and a padding value of 1. Each convolutional layer is followed by ReLU activation and 2×2 max pooling, which serves to downsample spatial dimensions. The number of filters increases progressively; 32 filters in the first block, 64 in the second, and 128 in the third, allowing the network to learn hierarchical features from basic to more complex patterns. After the feature extraction stage, the output is flattened and fed into a fully connected layer consisting of 256 neurons. To reduce the risk of overfitting, a dropout layer with a rate of 0.5 is introduced before passing to the final classifier layer. The final output layer is a linear unit with 10 neurons, each representing one of the target disease classes, and is followed by a softmax activation function to convert raw outputs into class probabilities. Training, the model employs the Adam optimizer with the CrossEntropyLoss function. The training was conducted over 10 epochs, with a batch size of 32 and a learning rate of 0.001. To improve generalization, data augmentation techniques such as random resized cropping, horizontal flipping, and random rotation were applied during training. The model was trained on the Google Colab platform utilizing T4 GPU acceleration, which enhanced training efficiency. While the custom CNN has a simpler architecture compared to ResNet18, experimental results demonstrate that it can achieve competitive classification accuracy. With appropriate hyperparameter tuning and sufficient augmentation, this model offers a viable and computationally lightweight alternative for plant disease detection, especially in scenarios with limited hardware resources.

Model training process (Epochs)

The training process was performed over 10 epochs, with the objective of refining model parameters until reaching convergence on the training data. The training log indicated a progressive reduction in loss values across epochs, accompanied by an increase in prediction accuracy on both the training and validation sets. During the first epoch, the ResNet18 model utilizing transfer learning recorded a training loss of 1.7199, a training accuracy of 47.40%, and a validation accuracy of 55.50%. As training progressed, performance improved; by epoch five, the training loss dropped to 1.1553, with training and validation accuracies of 69.30% and 68.70%, respectively. By the end of the 10th epoch, the model achieved a training loss of 0.9147, a training accuracy of 78.80%, and a validation accuracy of 74.30%.

The sixth epoch onward, the validation accuracy stabilized in the range of 72–75%, suggesting that the model had begun to converge. Nevertheless, further improvements could still be attained through additional training epochs or refinement of hyperparameters. These observations are consistent with

findings in prior studies, which highlight that a declining loss curve combined with stable validation performance across epochs serves as a reliable indicator of model convergence. Enhancing the model's generalization capability and minimizing overfitting, particularly when applied to unseen or real-world test data can be addressed using techniques such as learning rate scheduling and early stopping [1]. Moreover, carefully configuring the hyperparameters plays a critical role in achieving optimal model performance. The full set of hyperparameter values applied to the ResNet18 and custom CNN configurations used in this study is detailed in Table 1.

Table 1. Hyperparameter training model ResNet18 and CNN

No	Hyperparameter	ResNet18 Transfer learning	CNN custom
1	Optimization Algoritm	adam	Adam
2	Learning Rate	0.001	0.001
3	Batch Size	32	32
4	Epoch	10	10
5	Loss Function	CrossEntropyLoss	CrossEntropyLoss
6	Scheduler	none	None
7	Dropout	-	0.5
8	Data Augmentation	Resize, Flip, Rotation, Normalize	Resize, Flip, Rotation, Normalize

3. Results and Discussion

The results demonstrate that the ResNet18 model, trained through transfer learning, achieves strong validation accuracy, reaching 74%. This suggests that the model possesses robust generalization capabilities, particularly when handling tomato leaf disease classification involving multiple categories. On the other hand, the Custom CNN designed with a simpler architecture delivers lower performance, attaining only 55% accuracy. Despite this, the model still performs relatively well on certain classes, such as Tomato Yellow Leaf Curl Virus and Mosaic Virus, which achieved F1-scores of 0.66 and 0.75, respectively. These classes are likely easier to identify due to their more distinctive visual patterns. Conversely, the lowest scores were recorded for Early Blight (F1-score: 0.50) and Target Spot (F1-score: 0.27), most likely because of visual overlap with other categories, which complicates feature extraction. Moreover, the limited ability of the Custom CNN to capture intricate features, along with potential shortcomings during preprocessing, may have further contributed to the drop in performance. Taken together, these findings emphasize the practical benefits of employing pretrained models like ResNet18, especially when working with small-scale datasets characterized by high inter-class visual similarity.

Table 2. Classification performance evaluation

Model	Accuracy	presisi	Recall	F1-Score
<i>ResNet18</i>	74%	0.77	0.74	0.73
CNN Custom	55%	0.60	0.55	0.53

Table 2, the ResNet18 model consistently outperforms the Custom CNN across all evaluation metrics. These findings highlight the substantial advantage offered by leveraging pretrained models through transfer learning, particularly when working with limited datasets and high inter-class visual complexity.

Model accuracy

The ResNet18 model, implemented using a transfer learning technique, obtained a validation accuracy of 74%, demonstrating its capability in handling complex classification tasks. The highest recognition performance was observed in the Tomato Yellow Leaf Curl Virus category, which achieved an F1-score of 0.93. This can be attributed to its distinctive visual markers, such as curled leaf structures and yellow coloration. Additionally, high F1-scores were reported for the Bacterial Spot (0.83) and Healthy (0.82) categories. On the other hand, the Early Blight class recorded the lowest score at 0.38. This underperformance may stem from its visual resemblance to other diseases like Target Spot, as well as issues related to image preprocessing, illumination conditions, and camera capture angles. In contrast, the CNN Custom model yielded a lower validation accuracy of 55%. While overall accuracy was

reduced, the model retained the ability to detect visually unique disease patterns. However, it showed noticeable difficulty in distinguishing categories with overlapping symptoms, such as Target Spot (F1-score: 0.27) and Septoria Leaf Spot (F1-score: 0.38). These shortcomings are likely the result of architectural limitations in feature representation. In summary, ResNet18 exhibited superior performance across nearly all classes and evaluation metrics.

Confusion matrix

The confusion matrix serves as a key diagnostic tool, offering a detailed overview of the model's ability to differentiate between multiple classes in a classification setting. In this study, two confusion matrices were analyzed to compare the predictive behavior of the ResNet18 model (utilizing transfer learning) and the Custom CNN. Figure 5 illustrates the confusion matrix for the ResNet18 model, which demonstrates robust classification results across most disease categories. High prediction accuracy was observed for Tomato_Healthy (90 correct predictions), Tomato_Bacterial_Spot (75 correct), and Tomato_Yellow_Leaf_Curl_Virus (88 correct). Nonetheless, the model struggled to distinguish Early_Blight, often misclassifying it as Target_Spot or Late_Blight, revealing the difficulty of identifying diseases with overlapping visual symptoms. In contrast, Figure 6 shows the confusion matrix for the Custom CNN. This model achieved impressive accuracy for certain categories, such as Tomato_Mosaic_Virus and Tomato_Healthy, which reached 100% and 98%, respectively. However, its overall misclassification rate was higher than that of ResNet18, particularly in the Target_Spot category, which was frequently confused with Early_Blight. This may stem from limitations in the Custom CNN's capacity to capture intricate visual cues possibly due to its relatively shallow network depth and simpler design. Figure 3 a comparison of the two matrices highlights essential differences in model behavior: while ResNet18 exhibits more consistent and generalized performance across classes, Figure 4 the Custom CNN appears more effective for certain specific categories but lacks reliability in handling visually similar cases. These findings emphasize the necessity of thoughtful architecture selection and strategic data augmentation when designing robust plant disease recognition systems.

Detailed Classification Report:

	precision	recall	f1-score	support
Tomato__Bacterial_spot	0.93	0.75	0.83	100
Tomato__Early_blight	0.81	0.25	0.38	100
Tomato__Healthy	0.75	0.90	0.82	100
Tomato__Late_blight	0.81	0.71	0.76	100
Tomato__Leaf_Mold	0.73	0.77	0.75	100
Tomato__Mosaic_virus	0.64	0.90	0.75	100
Tomato__Septoria_leaf_spot	0.67	0.76	0.71	100
Tomato__Spider_mites	0.84	0.73	0.78	100
Tomato__Target_Spot	0.54	0.78	0.64	100
Tomato__Yellow_Leaf_Curl_Virus	0.98	0.88	0.93	100
accuracy			0.74	1000
macro avg	0.77	0.74	0.73	1000
weighted avg	0.77	0.74	0.73	1000

Figure 3. Per-class classification report using ResNet18 transfer learning model

Classification Report CNN Custom:

	precision	recall	f1-score	support
Tomato__Bacterial_spot	0.86	0.68	0.76	100
Tomato__Early_blight	0.47	0.54	0.50	100
Tomato__Healthy	0.40	0.98	0.57	100
Tomato__Late_blight	0.57	0.56	0.57	100
Tomato__Leaf_Mold	0.77	0.34	0.47	100
Tomato__Mosaic_virus	0.61	1.00	0.75	100
Tomato__Septoria_leaf_spot	0.53	0.30	0.38	100
Tomato__Spider_mites	0.35	0.37	0.36	100
Tomato__Target_Spot	0.58	0.18	0.27	100
Tomato__Yellow_Leaf_Curl_Virus	0.87	0.53	0.66	100

Figure 4. Per-class classification report using custom CNN Model

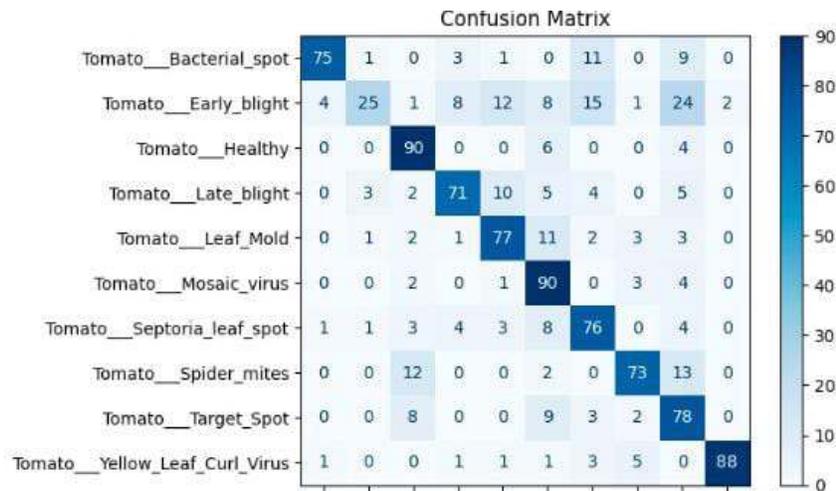


Figure 5. Confusion matrix visualization of ResNet18 model predictions on the tomato leaf disease validation data

Model prediction on test images (ResNet18)

Figure 7 illustrates the classification outcomes of the ResNet18 model when applied to several tomato leaf images sourced from the validation set. Overall, the model demonstrates a solid classification ability, correctly identifying most samples based on their ground-truth labels, with particularly strong performance in the Tomato_Bacterial_spot category. Green-labeled predictions in the figure indicate accurate classification and reflect the model's effectiveness in detecting unique visual traits, such as dark circular lesions with pronounced contrast.

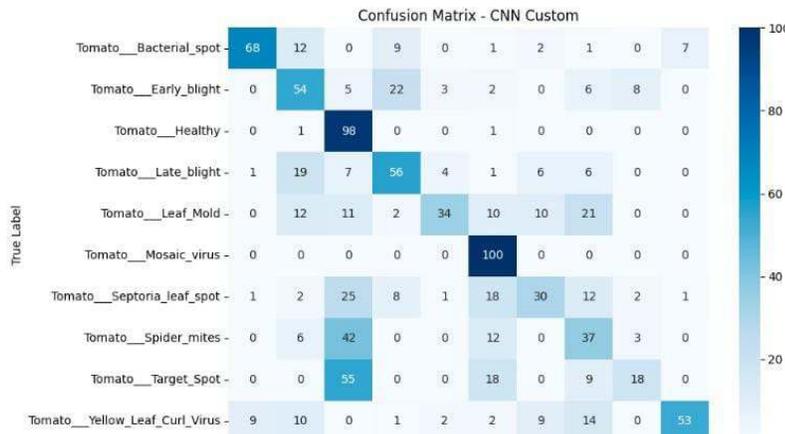


Figure 6. Confusion matrix visualization of custom CNN model predictions on the tomato leaf disease validation data.

Despite these strengths, some misclassification cases were observed. For instance, an image labeled as Tomato_Bacterial_spot was misidentified as Tomato_Early_blight, highlighted in red within the visualization. These inaccuracies likely stem from visual symptom overlap between disease types such as irregular dark spots—which can challenge the model, particularly under high intra-class variability. This observation is consistent with the claim [13], who noted that one major difficulty in tomato leaf disease identification lies in the visual similarity between disease classes. Factors such as image resolution, lighting conditions, and camera perspective can significantly influence model accuracy in real-world environments [15]. In summary, the visual evidence aligns with earlier quantitative results, reinforcing that while ResNet18 performs reliably on distinct disease classes, it still requires refinement for better differentiation of visually similar cases. Future work may focus on improving the network's

structural design, leveraging ensemble learning methods, or adopting multiview imaging strategies to enhance prediction reliability in practical scenarios.

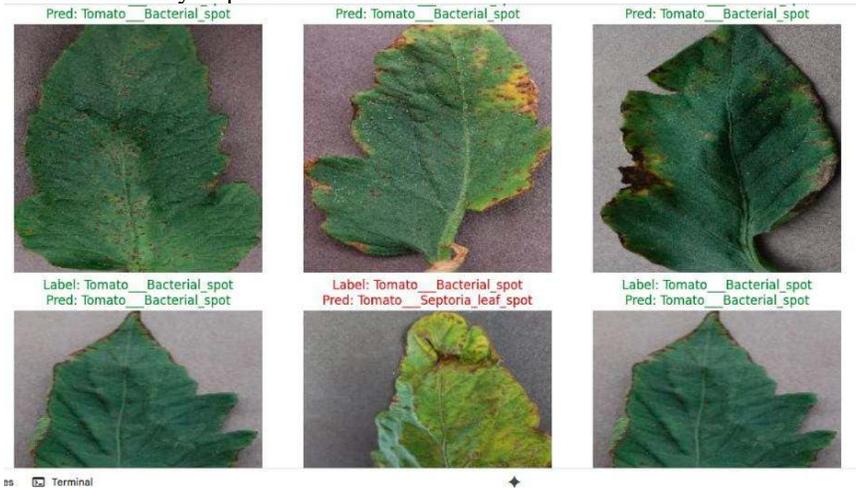


Figure 7. ResNet18 prediction on Bacterial_spot leaf images. Green text shows correct predictions; red text indicates misclassification as Early_blight due to symptom similarity.

Model prediction on test images (CNN Custom)

To assess its capability in classifying disease types based on visual cues, the CNN Custom model was evaluated using multiple images from the Tomato_Bacterial_spot category. Figure 8 illustrates the prediction outcomes, where each image is annotated with its actual label alongside the model's prediction. Correct classifications are highlighted in green, whereas incorrect ones are shown in red. The majority of the samples were accurately identified, showcasing the model's ability to detect distinctive features of the Bacterial_spot class such as the presence of dark, circular lesions on the leaf surface. This qualitative observation is consistent with earlier numerical results, where the same class achieved an F1-score of 0.76. However, one notable error occurred when an image belonging to the Bacterial_spot category was misclassified as Early_blight. This confusion likely stems from the overlapping visual characteristics shared between the two diseases, such as spot texture or variations in leaf coloration. Similar misclassification challenges have been discussed in earlier studies, including those [13], who emphasized that visual similarities in tomato leaf disease symptoms can hinder accurate automated recognition. Identified several influencing factors such as lighting inconsistencies, image quality, and differences in camera angles—that can impact model performance in visual recognition tasks al [15]. Visual interpretation methods like this are essential for complementing quantitative metrics, offering more comprehensive insights into the model's operational behavior under real-world imaging scenarios and helping identify aspects that warrant further refinement.

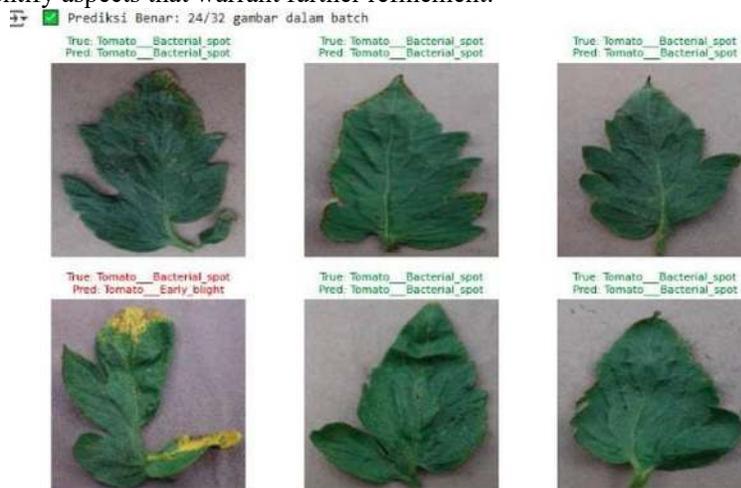


Figure 8. Example of CNN Custom model predictions on Tomato_Bacterial_spot class.

Green text indicates correct predictions, while red text indicates misclassifications. One image labeled `Bacterial_spot` was incorrectly predicted as `Early_blight`.

Comparison with previous studies

These findings reinforce that pretrained deep learning architectures such as ResNet18 generally outperform custom-designed models like CNNs tailored from scratch, particularly in handling multi-class classification problems, including those related to tomato leaf disease detection. In this study, ResNet18 yielded a validation accuracy of 74%, whereas the CNN Custom achieved only 55%. Similarly, in terms of average F1-score, ResNet18 reached 0.73, surpassing the CNN Custom's result of 0.53. This disparity highlights the advantages offered by pretrained networks like ResNet18, which benefit from prior exposure to large-scale datasets such as ImageNet, enabling the model to extract and generalize complex visual patterns more effectively through the use of residual learning blocks. Despite the performance gap, the custom CNN still proved capable of classifying certain disease categories with reasonable accuracy particularly those with distinct visual features like Tomato Yellow Leaf Curl Virus. However, it struggled with classes showing overlapping visual symptoms, such as Early Blight and Target Spot, which frequently led to misclassifications. This challenge is not unexpected, considering that the CNN Custom utilizes a more compact and shallow architecture, limiting its representational capacity for deeper feature extraction. This observation aligns with earlier studies suggesting that pretrained models are more adept at interpreting heterogeneous and complex imagery. Furthermore, the application of data augmentation techniques—such as image flipping and random rotation—used in this experiment, aligns with the findings by Wu et al., who emphasized the value of such transformations for improving model generalization. Although the CNN Custom's overall accuracy is lower, it presents a viable option for deployment in environments with constrained computational resources. Its training process is faster and less demanding in terms of hardware, making it suitable for edge computing applications or devices with limited processing power. Looking forward, future studies could investigate hybrid approaches—such as ensemble techniques or the integration of lightweight pretraining leveraging the strengths of both pretrained and custom models. Such approaches may lead to the development of classification frameworks that are not only robust and accurate but also computationally efficient and adaptable to real-world agricultural scenarios.

Limitations

This study still encounters several challenges that warrant further refinement. Initially, although the dataset was designed to contain an equal number of images for each class, some instances particularly in certain categories were not properly processed due to issues during transformation or preprocessing. As a result, the final class distribution became unbalanced, which had a negative impact on the model's performance, especially for underrepresented categories like `Early_blight`, which recorded an F1-score of only 0.38. The limited data availability in this category hindered the model's ability to effectively learn its unique visual traits. Secondly, even though the architecture of the CNN Custom model has been thoroughly outlined, its experimental outcomes were not comprehensively incorporated into the earlier stages of analysis. A rigorous side-by-side evaluation of CNN Custom and ResNet18 is crucial to fully highlight the advantages and shortcomings of each model. Prior literature suggests that well-constructed CNN architectures can still deliver competitive results, especially when backed by proper parameter optimization. Third, the training duration was restricted to 10 epochs due to computational constraints on the Google Colab environment, which often experienced interruptions during longer sessions. This raised the possibility that the model failed to reach optimal convergence. Researchers such as Ahmed et al. recommend adopting incremental training strategies with dynamically adjusted epoch counts to ensure loss stabilization and model robustness. Fourth, hyperparameter optimization remains unexplored. Key components like learning rate scheduling, dropout configurations, and batch normalization have yet to be systematically investigated. Properly calibrating these elements could potentially enhance the model's generalization capability across a wider range of image types. Moreover, ensemble strategies have not yet been applied in this study, even though combining models—such as through soft voting or stacking between CNN Custom and ResNet18—may improve classification effectiveness, especially for classes that are harder to identify. Additionally, the current research lacks empirical testing on lightweight platforms or edge devices, such as

smartphones or agricultural IoT systems. As a result, the deployment feasibility of the compact CNN Custom model under real-world conditions remains unverified. Recognizing these constraints, future work should focus on enhancing the experimental design through dataset enrichment, extended training durations, advanced hyperparameter tuning, and integration of hybrid modeling techniques. These steps are expected to yield a classification framework that is not only more accurate but also adaptable and practical for real-world precision agriculture applications.

Deductive arguments and speculation

Drawing from the experimental outcomes and prior analysis, it is evident that employing a transfer learning framework using ResNet18 offers notable advantages, particularly in classifying tomato leaf diseases from digital images especially under limited data conditions. The model's validation accuracy of 74% and macro F1-score of 0.73 serve as strong indicators of its capability to consistently identify diverse visual features. This robustness stems largely from its residual block structure, which facilitates stable gradient propagation, allowing effective parameter optimization in deep networks without compromising performance. Despite this, the CNN Custom model also demonstrates value, especially for deployment on devices with restricted computational resources, such as smartphones or IoT platforms. Owing to its lower parameter count, it exhibits significantly better computational efficiency. Although this model's performance in the current study was suboptimal, findings from Ahmed et al. suggest that with careful design and fine-tuning, CNN Custom architectures can achieve accuracy levels comparable to those of pretrained networks, especially on more homogeneous datasets. A promising direction for future research involves the development of ensemble models that harness the respective strengths of both ResNet18 and CNN Custom. Such hybrid frameworks enable the combination of ResNet18's strong feature extraction with the fast inference capabilities of lightweight CNNs. Techniques like soft voting and stacking have shown efficacy in prior studies for enhancing multi-class classification in agricultural vision tasks. Further improvements may involve refining CNN Custom through advanced optimization methods, such as Neural Architecture Search (NAS) or metaheuristic strategies like Particle Swarm Optimization (PSO). These approaches allow for the design of models that maintain a practical equilibrium between predictive accuracy and computational efficiency. With continued refinement, these models are expected to become increasingly adaptive to real-world environments, thereby supporting the broader application of deep learning systems in precision agriculture, particularly within the Indonesian context.

4. Conclusion

This research conducted a comparative evaluation between two deep learning-based methods for identifying tomato leaf diseases through digital imagery: the custom CNN architecture and the pretrained ResNet18 model. The experimental analysis demonstrated that ResNet18 consistently delivered better results than the CNN Custom model, both in accuracy and overall evaluation metrics. It achieved a validation accuracy of 74%, reflecting its robustness in recognizing complex visual features characteristic of various plant diseases. In contrast, the CNN Custom model—built on a simpler structural design resulted in a lower validation accuracy of 55%. Despite this, it was still capable of correctly identifying several categories, such as Mosaic Virus and Yellow Leaf Curl Virus. However, its performance declined when faced with visually similar disease classes like Early Blight and Target Spot. From these findings, ResNet18 emerges as the more suitable option for tomato leaf disease detection. Its architectural advantages such as residual connections, the ability to extract rich features from pretrained ImageNet weights, and a consistent training process—contribute to superior generalization and predictive accuracy. Nonetheless, the compact and computationally efficient nature of the custom CNN model retains practical significance for real-world applications on edge computing devices with limited processing capabilities.

5. References

- [1] M. Ahmad, M. Abdullah, H. Moon, and D. Han, "Plant Disease Detection in Imbalanced Datasets Using Efficient Convolutional Neural Networks With Stepwise Transfer Learning," *Ieee Access*, vol. 9, pp. 140565–140580, 2021, doi: 10.1109/access.2021.3119655. <https://doi.org/10.1109/ACCESS.2021.3119655>

- [2] K. Roy *et al.*, “Detection of Tomato Leaf Diseases for Agro-Based Industries Using Novel PCA DeepNet,” *Ieee Access*, vol. 11, pp. 14983–15001, 2023, doi: 10.1109/access.2023.3244499. <https://doi.org/10.1109/ACCESS.2023.3244499>
- [3] D. Li, Y. Pan, S. Mao, M. Deng, and H. Shen, “Deep Residual Networks With a Flask-Like Channel Structure,” *Ieee Access*, vol. 12, pp. 11711–11722, 2024, doi: 10.1109/access.2023.3347331. <https://doi.org/10.1109/ACCESS.2023.3347331>
- [4] A. Tabbakh and S. S. Barpanda, “A Deep Features Extraction Model Based on the Transfer Learning Model and Vision Transformer ‘TLMViT’ for Plant Disease Classification,” *Ieee Access*, vol. 11, pp. 45377–45392, 2023, doi: 10.1109/access.2023.3273317. <https://doi.org/10.1109/ACCESS.2023.3273317>
- [5] S. Barburiceanu, S. Meza, B. Orza, R. Malutan, and R. Terebeş, “Convolutional Neural Networks for Texture Feature Extraction. Applications to Leaf Disease Classification in Precision Agriculture,” *Ieee Access*, vol. 9, pp. 160085–160103, 2021, doi: 10.1109/access.2021.3131002. <https://doi.org/10.1109/ACCESS.2021.3131002>
- [6] E. Moupojou *et al.*, “FieldPlant: A Dataset of Field Plant Images for Plant Disease Detection and Classification With Deep Learning,” *Ieee Access*, vol. 11, pp. 35398–35410, 2023, doi: 10.1109/access.2023.3263042. <https://doi.org/10.1109/ACCESS.2023.3263042>
- [7] E. Özbilge, M. K. Ulukök, Ö. Toygar, and E. Ozbilge, “Tomato Disease Recognition Using a Compact Convolutional Neural Network,” *Ieee Access*, vol. 10, pp. 77213–77224, 2022, doi: 10.1109/access.2022.3192428. <https://doi.org/10.1109/ACCESS.2022.3192428>
- [8] X. Wang and W. Cao, “Bit-Plane and Correlation Spatial Attention Modules for Plant Disease Classification,” *Ieee Access*, vol. 11, pp. 93852–93863, 2023, doi: 10.1109/access.2023.3309925. <https://doi.org/10.1109/ACCESS.2023.3309925>
- [9] S. T. Y. Ramadan *et al.*, “Improving Wheat Leaf Disease Classification: Evaluating Augmentation Strategies and CNN-Based Models With Limited Dataset,” *Ieee Access*, vol. 12, pp. 69853–69874, 2024, doi: 10.1109/access.2024.3397570. <https://doi.org/10.1109/ACCESS.2024.3397570>
- [10] L. Li, S. Zhang, and B. Wang, “Plant Disease Detection and Classification by Deep Learning—A Review,” *Ieee Access*, vol. 9, pp. 56683–56698, 2021, doi: 10.1109/access.2021.3069646. <https://doi.org/10.1109/ACCESS.2021.3069646>
- [11] M. Umar, S. Altaf, S. Ahmad, H. A. Mahmoud, A. S. N. Mohamed, and R. Ayub, “Precision Agriculture Through Deep Learning: Tomato Plant Multiple Diseases Recognition With CNN and Improved YOLOv7,” *Ieee Access*, vol. 12, pp. 49167–49183, 2024, doi: 10.1109/access.2024.3383154. <https://doi.org/10.1109/ACCESS.2024.3383154>
- [12] C. Zhou, S. Zhou, J. Xing, and J. Song, “Tomato Leaf Disease Identification by Restructured Deep Residual Dense Network,” *Ieee Access*, vol. 9, pp. 28822–28831, 2021, doi: 10.1109/access.2021.3058947. <https://doi.org/10.1109/ACCESS.2021.3058947>
- [13] R. S. R. Singh and R. K. Sanodiya, “Zero-Shot Transfer Learning Framework for Plant Leaf Disease Classification,” *Ieee Access*, vol. 11, pp. 143861–143880, 2023, doi: 10.1109/access.2023.3343759. <https://doi.org/10.1109/ACCESS.2023.3343759>
- [14] H. I. Peyal *et al.*, “Plant Disease Classifier: Detection of Dual-Crop Diseases Using Lightweight 2D CNN Architecture,” *Ieee Access*, vol. 11, pp. 110627–110643, 2023, doi: 10.1109/access.2023.3320686. <https://doi.org/10.1109/ACCESS.2023.3320686>
- [15] S. Ahmed, Md. B. Hasan, T. Ahmed, Md. R. K. Sony, and Md. H. Kabir, “Less Is More: Lighter and Faster Deep Neural Architecture for Tomato Leaf Disease Classification,” *Ieee Access*, vol. 10, pp. 68868–68884, 2022, doi: 10.1109/access.2022.3187203. <https://doi.org/10.1109/ACCESS.2022.3187203>