https://doi.org/10.17221/67/2024-RAE

Towards interpretability: Assessment of residual networks for tomato leaf disease classification

Raphael Berdin, Rob Christian Caduyac*

Department of Electrical Engineering, University of the Philippines Los Baños, Laguna, Philippines

**Corresponding author: rmcaduyac1@up.edu.ph*

Citation: Berdin R., Caduyac R.C. (2025): Towards interpretability: Assessment of residual networks for tomato leaf disease classification. Res. Agr. Eng., 71: 1–9.

Abstract: The tomato occupies a prominent place in the Philippines' agricultural economy. However, tomato leaf diseases are challenges in tomato crop production leading to economic losses. Among the tomato leaf diseases, early blight and Septoria leaf spot are prevalent in the Philippines due to the climate. Thus, the accurate identification of diseases affecting tomato leaves is essential. Currently, a visual inspection is the primary method for diagnosing tomato leaf diseases which is time-consuming and inefficient. This study aims to develop a quantized Residual Network with convolutional 50 layer (ResNet-50) based model to classify tomato leaves as healthy or affected by Septoria leaf spot or early blight. Furthermore, to enhance the reliability of the models' classification, gradient-weighted class activation mapping (Grad-CAM) was implemented. In contrast with the visual inspection, a programmed system does not get tired and can provide consistent performance results. As a result, the original 32-bit floating point model attained an accuracy rate of 91.22%. The quantized 16-bit floating point model demonstrated comparable performance with 90.10% accuracy with a 50% reduction in the model size and inference time of 0.3942 seconds. The minimal accuracy loss of the 16-bit model relative to the 32-bit model is due to the post-training quantization. The reduction to 16-bit precision is significant for the future deployment of edge devices where resources are limited.

Keywords: tomato leaf disease classification; Grad-CAM; quantization; ResNet50

Tomato (*Solanum lycopersicum*) is extensively cultivated in the world, and it was reported that 187 million metric tonnes of tomatoes were produced worldwide in 2020 (Branthôme 2022). The tomato occupies a prominent place in the Philippines' agricultural economy. In fact, the Philippines is among the Asia-Pacific region's leading producers of tomatoes, with a production value of 6.15 billion pesos in 2021. It is also one of the most profitable crops in the Philippines and second most significant fruit vegetable in the country after the eggplant (Gorme et al. 2017).

The Philippines, due to favourable climate and soil conditions, has been a major producer of tomatoes. The country's tropical climate, with an average temperature of 21–24 °C, abundant rainfall, and fertile soil, is suitable for tomato cultivation. Tomatoes are typically cultivated year-round, but the main sowing season typically runs from the months of September to January in hilly locations, and from the months of November to February in lowland locations (Department of Agriculture 2017).

However, tomato plants are sensitive to a variety of diseases which limits tomato cultivation resulting in a significant economic loss. Among all tomato diseases, fungus-based diseases, such as early blight and Septoria leaf spot, are much more prevalent in the Philippines due to extended periods of leaf moisture and moderate temperatures, both of which are typical of humid climates (Maeda-Gutiérrez et al. 2020).

© The authors. This work is licensed under a Creative Commons Attribution-NonCommercial 4.0 International (CC BY-NC 4.0).

https://doi.org/10.17221/67/2024-RAE

Visual inspections, or the naked-eye method, are typically the initial step in detecting tomato plant diseases such as Septoria and early blight (Arnal Barbedo 2013). However, visual inspections are not always effective at detecting these diseases in their early stages since their symptoms are frequently confused with other diseases. It also emphasises that early blight and Septoria leaf spot can be difficult to distinguish visually, particularly in their early stages, and the incorrect identification can result in inappropriate management strategies. Hence, the automation of tomato leaf disease classification is a must.

Machine learning approaches have demonstrated the potential for automated tomato disease classification. Due to the rise of computing resources, deep learning algorithms are used since they do not require multiple pre-processing steps before the classification task. For tomato leaf disease classification, convolutional neural networks (CNNs) were commonly used since it can learn discriminative characteristics from pictures and generate predictions based on those (O'Shea and Nash 2015). Various studies (De Luna et al. 2018; Rangarajan et al. 2018; Zhang et al. 2018; Elhassouny and Smarandache 2019; Agarwal et al. 2020; Vinay et al. 2020; Attalah 2023; Baser et al. 2023) used CNN for tomato leaf disease classification, and they produced high accuracies. However, they did not consider the computational requirements of their algorithms, as well as their interpretability. For practical applications of tomato disease classifications, the computational requirement is a challenge when the algorithm will be deployed in resource-constrained devices (Alzubaidi et al. 2021). CNN models require significant memory bandwidth and computational power due to their numerous parameters, often surpassing edge device capabilities (Liu et al. 2019). Moreover, the algorithm's interpretability is a must especially when interventions will be undertaken. Interpretable algorithms enable knowledge-specific researchers to trust the detection process (Chopra and Wig 2021; Mahmud et al. 2024).

This study aims to develop a lightweight CNN model and assess the developed model through gradient-weighted class activation mapping (Grad-CAM). Specifically, residual networks with 50 layers (ResNet-50) will be optimised through grid search hyperparameter tuning. Post training quantization will be performed to the optimised ResNet-50, and assessment through Grad-CAM will be undertaken for the interpretability of the model. The research hypothesises that an accurate tomato leaf disease classification can be made using a lightweight and interpretable model.

MATERIAL AND METHODS

The main objective of the study is to develop a lightweight and interpretable ResNet-50 for tomato leaf disease classification. Figure 1 shows the overview of the methodology. It is divided into three parts namely, data pre-processing, deep learning training and evaluation, and assessment.

The training dataset is obtained from the Plant Village Dataset from Kaggle, an open-source data science and machine learning platform. The Plant Village Dataset is a reliable source utilised in the machine learning field and verified by agriculture experts worldwide (Plant Village 2024). In this study, only healthy, Septoria leaf spot, and early blight leaves were considered, and each class consists of 1 000 images. Figure 2 illustrates the sample images of the said class in the training dataset.

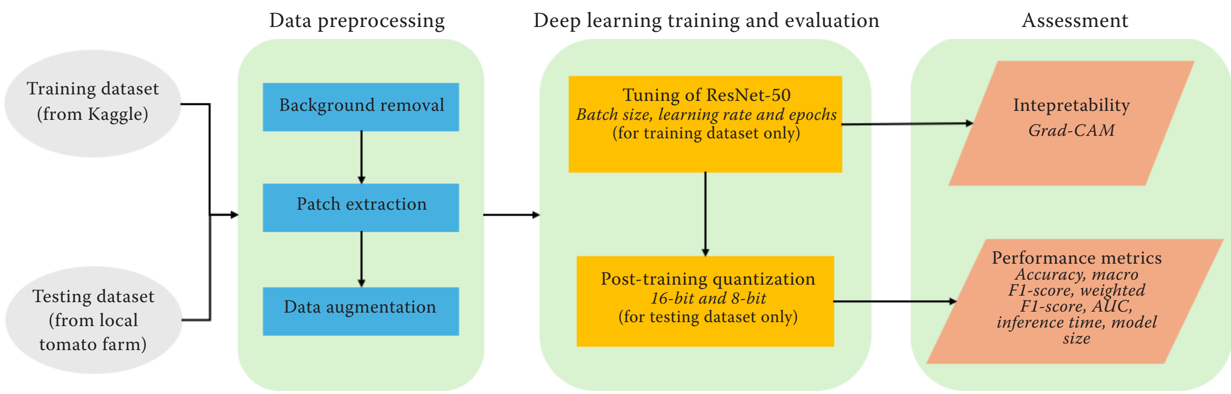

Figure 1. Methodology overview

https://doi.org/10.17221/67/2024-RAE

Figure 2. Left to right: early blight, Septoria leaf spot, healthy leaves of the training set

Meanwhile, the testing dataset was collected manually in Tiaong, Quezon, Philippines during the month of September 2023. It contains 85 early blight, 70 healthy, and 33 Septoria leaf spot leaves, which were 5 to 6 months old. Additionally, a licensed agriculturalist validated the testing set to ensure its credibility and verify the accuracy of the labels comprising the testing dataset. Figure 3 shows the sample image of each class, locally collected.

Data pre-processing. For interpretability purposes through Grad-CAM, all the images in each dataset were subjected to background removal and patch extraction to highlight the significant features of each class (Yebasse et al. 2021). Background removal is the process of replacing the background of an image to a white background. This allows the model to better capture the distinguishing characteristics of each class. Meanwhile, the process of patch extraction zoomed out all the images with a dimension of 90 × 90 pixels to highlight their unique regions. All the images were resized to a 256 × 256 dimension to tailor-fit with acceptable ResNet-50 input size. Their pixel values were normalised between 0 to 1 for faster solution convergence. Lastly, data augmentation was employed to improve the model's ability to generalise and avoid overfitting. The study employed rotation, flipping, and zooming, producing 6 000 and 188 training and testing images overall, respectively.

Deep learning training and evaluation. A pre-trained ResNet-50 architecture was used as a classification algorithm for tomato leaf diseases. Figure 4 shows the architecture of ResNet-50. ResNet-50 is a 50-layer deep convolutional neural network that uses residual/skip connections to mitigate the vanishing gradient problem for training deeper networks (He et al. 2015). For the tuning of the pre-trained ResNet-50 model, a grid search was performed to identify the best hyperparameter from the preselected choices. The batch size, learning rate, and number of epochs were considered as the hyperparameter. Table 1 shows the hyperparameter choices of ResNet-50. Stratified 5-fold cross-validation was implemented in the training dataset, and the average accuracy is the basis for selecting the best hyperparameter combination.

Table 1. Hyperparameter choices for 50-layer residual networks

	Hyperparameter values
Batch size	32, 64, 128, 256
Epoch	25, 50, 75, 100, 125, 150
Learning rate	0.01, 0.001, 0.0001, 0.00001

Figure 3. Left to right: early blight, Septoria leaf spot, healthy leaves of the testing set

https://doi.org/10.17221/67/2024-RAE

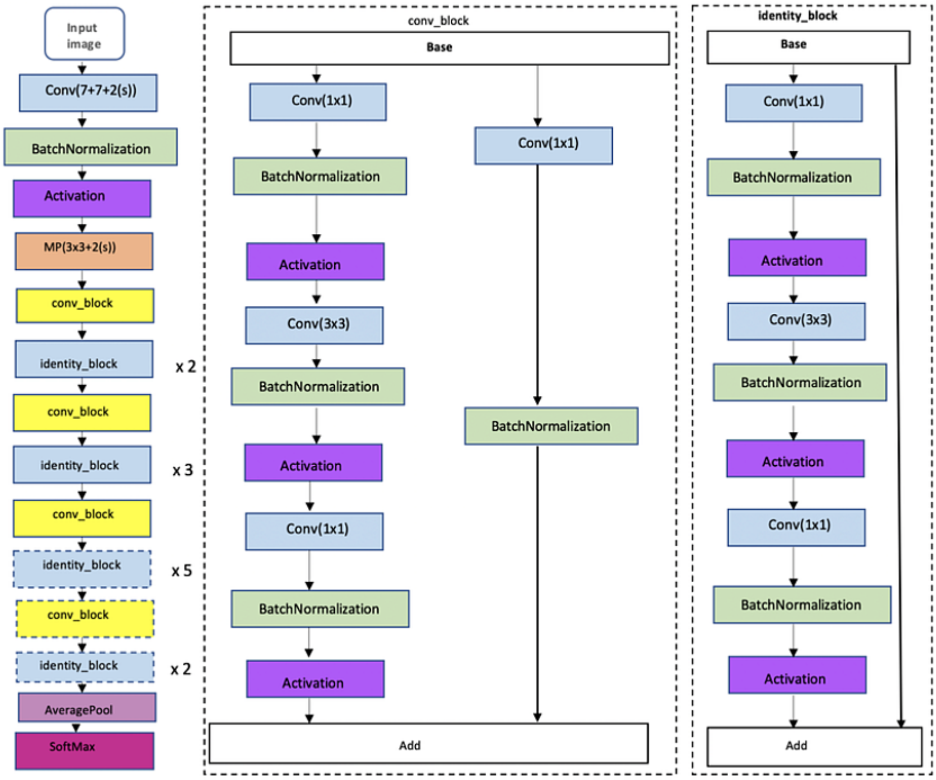

Figure 4. 50-layer residual network architecture (Mandal et al. 2021)

Table 2. Labels used in the qualitative evaluation of the (Grad-CAM) heatmaps (Molnar 2022)

Label classification	Label name	Label acronym	Letter label
Main label	Class predicted correctly with high confidence	CPCHC	–
	Class predicted correctly with low confidence	CPCLC	–
	Class predicted incorrectly with high confidence	CPIHC	–
	Class predicted incorrectly with low confidence	CPILC	–
Sub label	All only spots are perfectly highlighted	AOSPH	A
Sub label	Some only spots are perfectly highlighted	SOSPH	B
Sub label	All only spots are highlighted	AOSH	C
Sub label	Some only spots are highlighted	SOSH	D
Sub label	Some spots perfectly are perfectly highlighted, but some spots are partially highlighted	SSPHSPH	E
Sub label	All spots are perfectly highlighted, but some insignificant areas are highlighted	ASPHIH	F
Sub label	All spots are majority highlighted, but some insignificant areas are highlighted	SSPHIH	G
Sub label	All spots are highlighted, but some insignificant areas are highlighted	ASHIH	H
Sub label	Some spots are highlighted, but some insignificant areas are highlighted	SSHIH	I
Sub label	Some spots are perfectly highlighted, but some spots are partially highlighted, and some insignificant areas are highlighted	SSPHSPHIH	J
Sub label	No spots are highlighted	NH	K

For label acronym description see Table 3

https://doi.org/10.17221/67/2024-RAE

Post training quantization (PTQ) was applied to the optimised ResNet-50 to make the optimised model lightweight enough for edge device deployments. PTQ is the process of reducing the precision of the weight and activation values inside the neural network model to lower-bit precision with consideration for performance degradation. Here, only 16-bit floating and 8-bit integer quantization were considered. The quantized ResNet-50 was evaluated using the testing dataset for generalisation purposes.

Assessment. To assess the quantized models in terms of its classification performance, the accuracy, macro F1-score, weighted F1-score, area under the curve (AUC), model size, and inference time were considered. The model size and inference time were considered for the assessment of the edge device compatibility and real-time processing speed. Ideally, a low model size and faster inference time were preferred. Meanwhile, other metrics were considered for the general classification performance. This is to ensure that the classification performance was not heavily degraded upon quantization. For interpretability, a visualisation tool named gradient-weighted class activation mapping (Grad-CAM) was used to create heatmaps from the ResNet-50 activation layer. Grad-CAM reveals the inner workings of the black-box of ResNet-50 and aids in debugging the model when overfitting. The qualitative assessment of Grad-CAM was analysed using the main and sub label, with each main level containing all the sub labels (Molnar 2022). Table 2 shows the various labels for the Grad-CAM heatmap analysis.

The labels are descriptive terms used to describe the visible symptoms in the images, excluding small spots due to the possibility of dirt. Insignificant leaf areas are those that do not correspond to a disease

Table 3. Description of each label found in Table 2 (Molnar 2022)

Label acronym	Description
CPCHC	A label that means that the model has correctly predicted the class of the image with a relatively high confidence.
CPCLC	A label that means that the model has correctly predicted the class of the image, but with a relatively low confidence.
CPIHC	A label that means that the model has incorrectly predicted the class of the image along with a relatively high confidence.
CPILC	A label that means that the model has incorrectly predicted the class of the image, but with a relatively low confidence.
AOSPH	Symptoms are the only ones highlighted, all the visible symptoms are highlighted, and a majority of the symptoms are highlighted.
SOSPH	Symptoms are the only ones highlighted and a majority of the symptoms are highlighted, but not all visible symptoms are highlighted.
AOSH	Symptoms are the only ones highlighted and all the visible symptoms are highlighted, but only some parts of a symptom are highlighted.
SOSH	Symptoms are the only ones highlighted, but only some parts of a symptom are highlighted and not all visible symptoms are highlighted.
SSPHSPH	Some visual symptoms have its majority highlighted and some symptoms only have some parts highlighted.
ASPHIH	All visible symptoms have their majority highlighted, but insignificant parts of the image are highlighted as well.
SSPHIH	Not all visible symptoms are highlighted, but the majority of a symptom is highlighted, and insignificant parts of the image are highlighted as well.
ASHIH	All visible symptoms are highlighted, but only some parts of a symptom are highlighted. Also, insignificant parts of the image are highlighted as well.
SSHIH	Some visible symptoms are highlighted and only some parts of a symptom are highlighted. Also, insignificant parts of the image are highlighted as well.
SSPHSPHIH	Some visual symptoms have their majority highlighted and some symptoms only have some parts highlighted. Also, insignificant parts of the image are highlighted as well.
NH	No visual symptoms were highlighted despite their presence on the leaf being analysed.

https://doi.org/10.17221/67/2024-RAE

Table 4. Optimal hyperparameters of the pretrained 50-layer residual networks

Hyperparameter	Value
Batch size	128
Epochs	50
Learning rate	0.001
Dropout rate	50

symptom or are not utilised by specialists. The high level of confidence is more than or equal to 85% and the low level of confidence is less than 85%. Table 3 contains the description of each label.

RESULTS AND DISCUSSION

This section is subdivided into two parts, namely the performance of the quantized ResNet-50 and the qualitative assessment of ResNet-50. Each of them is discussed accordingly.

Performance of the quantized ResNet-50. The optimal hyperparameters of the pretrained ResNet-50 obtained using the training dataset is shown in Table 4. Here, grid search cross-validation was used, and the average validation accuracy was compared to various combinations of hyperparameters.

The performance of the original (32-bit) and quantized pretrained ResNet-50 were observed in the testing dataset through the accuracy, macro F1-score, weighted F1-score, and AUC. Meanwhile, the model size and inference time were also considered for the hardware performance of the models with the model size for the edge device deployment and inference time for real-time processing. 16-bit float and 8-bit integer were used for the quantization.

Figure 5 shows the comparison of the classification performance metrics of the 32-bit, 16-bit, and 8-bit pretrained ResNet-50. It was observed that a minimal degradation occurred on all the classification performance when converted from a 32-bit to 16-bit quantized model. However, significant degradation happened when converted to an 8-bit quantized model. For the 16-bit quantized model, there is a 90% chance that the model correctly predicts the rice leaf disease classification. A macro-F1 score of 88% indicated a balanced performance across all the considered rice leaf classes, especially the least occurring class Septoria leaf spot. Meanwhile, a weighted F1-score of 90% demonstrated the strong overall model performance, particularly for the more frequent class named early blight. Lastly, an AUC value of 91% reflected the model's capability to distinguish different diseases across various threshold settings.

For the hardware performance of the 32-bit, 16-bit, and 8-bit 50-Layer Residual Networks, Table 5 summarises the model size and inference time of each model. It was observed that a 32-bit model size was significantly reduced to two and four times when quantized to 16-bit and 8-bit, respectively. This was explainable since the reduction from a 32-bit floating point to a 16-bit float and an 8-bit

Table 5. Hardware performance of 32-bit, 16-bit, and 8-bit 50-Layer Residual Networks

ResNet-50	Model size (MB)	Inference time (s)
32-bit	93.5962	0.6085
16-bit	47.5045	0.3942
8-bit	23.3876	0.3281

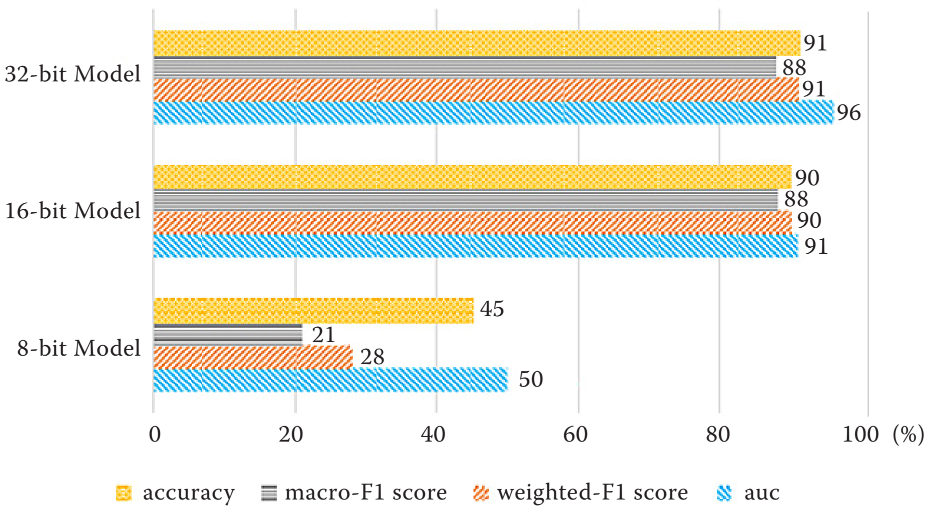

Figure 5. Classification performance of the 32-bit, 16-bit, and 8-bit 50-layer residual networks

https://doi.org/10.17221/67/2024-RAE

Table 6. Labels of the Grad-CAM images under the class predicted correctly with high confidence

Class	Grad-CAM labels										
	A	B	C	D	E	F	G	H	I	J	K
Early blight	0	0	0	0	1	3	0	10	16	123	1
Septoria	0	0	0	0	0	1	0	1	4	37	2
Healthy	0	0	1	4	3	9	0	19	25	82	0
Total	0	0	1	4	4	13	0	30	45	242	3

integer roughly compressed the bit counts to two and four times, respectively. Lastly, an almost 50% decrease in the inference time was observed when 16-bit and 8-bit quantized models are considered compared to the 32-bit model. While the reduction in the inference time is at a minimum in the context of processing one image, it is essential for real-time processing especially when dealing with a large number of image inputs to the model.

Qualitative assessment of ResNet-50. Table 6 provides insights into the number of correct predictions falling under the main label class predicted correctly with high confidence (CPCHC). Note that a high confidence signifies an accuracy greater or equal to 85%. Within the CPCHC category, most images are classified under the sub-label J. This sublabel signifies that the model highlighted some visual symptoms perfectly and partially. It is worth noting that insignificant regions of the image are highlighted, potentially indicating areas of noise or non-disease-related features.

Figure 6 present sample Grad-CAM images of a tomato leaf affected by early blight under CPCHC. The heatmap revealed that the model successfully captures most of the diseased regions on the leaf. The intensity of the created heatmap defines the damage intensity of the tomato leaves. It was observed that the leaf was significantly damaged from the left side based on the heat intensity. However, it is worth noting that the highlighted regions appear coarse. This may be attributed to factors, such as the quality of the training dataset and discrepancies between the testing and training datasets. Despite the coarse highlighting, the model demonstrated high confidence in predicting the early blight class correctly. This suggests that the model has learned valuable features that are distinctive to tomato leaves exhibiting symptoms of early blight.

Figure 7 displays sample Grad-CAM images of a tomato leaf with Septoria leaf spot under the CPCHC category. The Grad-CAM image revealed that the model effectively highlights most of the diseased regions corresponding to Septoria leaf spot. It was observed that the leaf was significantly damaged everywhere based on the heat intensity. However, compared to the early blight class, the highlighted regions appeared coarser. This discrepancy may be attributed to the lower quality of the Septoria leaf spot (SLS) training set compared to the early blight training set. Additionally, discrepancies between the SLS training and testing sets such as the presence of different visual symptoms in the testing set not represented in the training set, may contribute to the coarser highlighting. Furthermore, the training set sourced internationally may not fully capture the visual symptoms prevalent in tomato leaves with SLS in the local Philippine environment. This leads to differences in the symptom expression. Despite

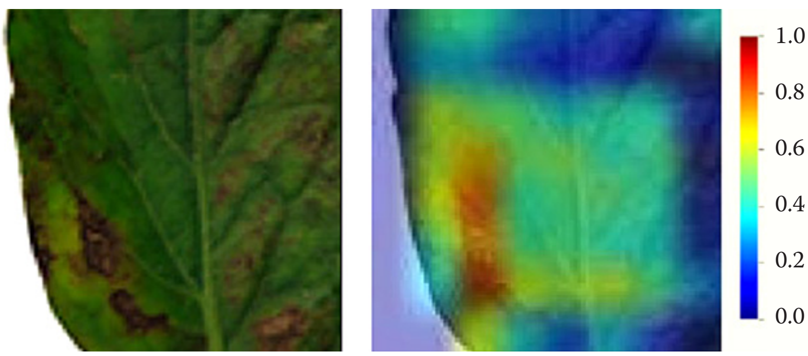

Figure 6. Sample Grad-CAM image of early blight under the class predicted correctly with high confidence

https://doi.org/10.17221/67/2024-RAE

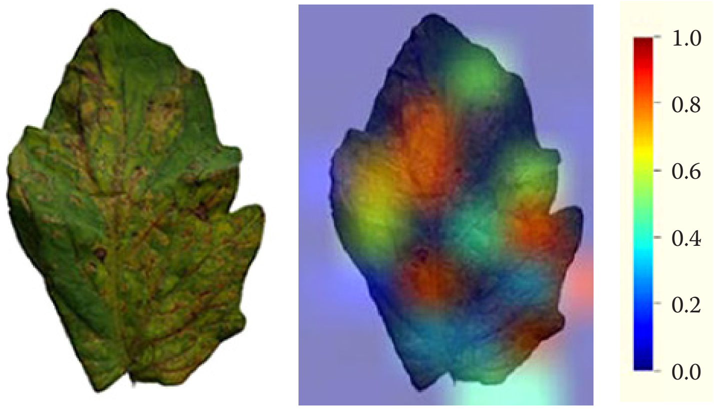

Figure 7. Sample Grad-CAM image of Septoria leaf spot under the class predicted correctly with high confidence

the coarser highlighting in the SLS class, the model still correctly predicted the SLS class. This featured its ability to generalise and identify distinctive features of tomato leaves with Septoria leaf spot.

Figure 8 illustrates sample Grad-CAM images of healthy tomato leaves under CPCHC. Notably, the Grad-CAM highlighted most of the regions of the tomato leaf which is a characteristically unique to the healthy class. However, some of the Grad-CAM images of a healthy tomato leaf highlighted the white background. Despite this, the model consistently classified the correct class indicating the robustness and a thorough understanding of the features of a healthy tomato leaf. This robust performance in identifying healthy leaves is evident, given the challenges posed by the background noise. Furthermore, it was observed that the healthy class within the training set exhibits higher quality compared to the early blight and Septoria leaf spot classes. This has offered valuable insights into model training. The clearer and less blurry images in the healthy class suggest that the model has had exposure to high quality data, enabling it to learn and generalise effectively. These healthy leaf features in the training set translate into finer Grad-CAM heatmaps, further validating the model's proficiency in classifying healthy tomato leaves.

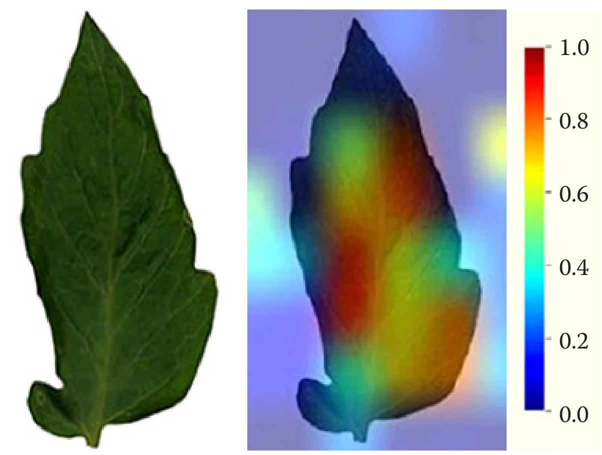

Figure 8. Sample Grad-CAM image of healthy leaves under the class predicted correctly with high confidence

CONCLUSION

This study developed quantized ResNet-50 architecture for classifying healthy tomato leaves, as well as those affected by early blight or Septoria leaf spot. The ResNet-50 architecture was fine-tuned using a dataset of labelled tomato leaf images. The training process involved data augmentation and optimisation techniques to enhance the model's generalisation capabilities. This ensured it could accurately distinguish between the various tomato leaf diseases.

The evaluation of Grad-CAM images served as a crucial tool for building confidence in the model's predictions. Grad- CAM was used to generate visual explanations of the model's decision-making process by highlighting regions in the images that contributed to the predictions. It was observed that despite highlighting small and insignificant regions, such as background noise and dirt, the model consistently predicted the correct class, underscoring its robustness. This qualitative analysis involved visually inspecting and interpreting numerous Grad-CAM heatmaps to ensure the model's reliability and interpretability.

In conclusion, the best quantized model in terms of performance metrics is the 16-bit model. It demonstrated a 50% reduction in the model size and had a faster inference time of 0.3942 seconds relative to the original 32-bit model. It also had a minimal performance loss, which makes it a suitable choice for the future deployment in environments with

https://doi.org/10.17221/67/2024-RAE

limited computational resources. This model provides an optimal balance between accuracy, model size, and inference speed which make it ideal for practical applications in the field.

REFERENCES

Agarwal M., Singh A., Arjaria S., Sinha A., Gupta S. (2020): ToLeD: Tomato leaf disease detection using convolutional neural network. Procedia Computer Science, 167: 293–301.

Alzubaidi L., Zhang J., Humaidi A.J., Al-Dujaili A., Duan Y., Al-Shamma O., Santamaria J., Fadhel M.A., Al-Amidie M., Farhan L. (2021): Review of deep learning: Concepts, CNN architectures, challenges, applications, future directions. Journal of Big Data, 8: 53.

Arnal Barbedo J.G. (2013): Digital image processing techniques for detecting, quantifying, and classifying plant diseases. SpringerOpen Journal Barbedo, 2: 660.

Attalah O. (2023): Tomato leaf disease classification via compact convolutional neural networks with transfer learning and feature selection. Horticulture, 9: 149.

Baser P., Saini J., Kotecha K. (2023): TomConv: An improved CNN model for diagnosis of diseases in tomato plant leaves. Procedia Computer Science, 218: 1825–1833.

Branthôme F.X. (2022): Worldwide (Total Fresh) Tomato Production Exceeds 187 Million Tonnes in 2020. Available at https://www.tomatonews.com/en/worldwide-total-fresh-tomato-production-exceeds-187-million-tonnes-in-2020_2_1565 (accessed July 10, 2024).

Chopra G., Whig P. (2021): Analysis of tomato leaf disease identification techniques. Journal of Computer Science and Engineering, 2: 98–103.

De Luna R.G., Dadios E.P., Bandala A.A. (2018): Automated image capturing system for deep learning-based tomato plant leaf disease. Detection and recognition. In: Proceedings of TENCON 2018 IEEE Region 10 Conference. Jeju, South Korea, Oct 28–31, 2018: 1414–1419.

Department of Agriculture (2017): Tomato Production Guide. Available at https://cagayanvalley.da.gov.ph/wp-content/uploads/2018/02/Tomato.pdf (accessed July 10, 2024).

Elhassouny A., Smarandache F. (2019): Smart mobile application to recognize tomato leaf diseases using convolutional neural networks. In: Proceedings of the 2019 International Conference of Computer Science and Renewable Energies. Agadir, Morocco, Jul 22–24, 2019: 1–4.

Gorme A.L., Gonzaga Z., Capuno O., Rom J., McDougall S., Goldwater A., Rogers G. (2017): Growth and yield of tomato (*Solanum lycopersicum*) as influenced by different soil organic amendments and types of cultivation. Annals of Tropical Research, 39: 116–128.

He K., Zhang X., Ren S., Sun J. (2015): Deep residual learning for image recognition. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, Nevada, June 27–30, 2016: 770–778.

Liu H., Yao D., Yang J., Li X. (2019): Lightweight convolutional neural network and its application in rolling bearing fault diagnosis under variable working conditions. Sensors, 19: 4827.

Maeda-Gutiérrez V., Galvan C.E., Zanella L.A., Celaya J.M., Galvan J.I., Gamboa H., Luna H., Magallanes R., Guerrero C.A., Olvera C.A. (2020): Comparison of convolutional neural network architectures for classification of tomato plant diseases. Applied Sciences, 10: 1245.

Mahmud T., Barua K., Barua A., Basnin N., Das S., Hossain M., Andersson K. (2024): Explainable AI for tomato leaf disease detection: Insights into model interpretability. In: Proceedings of the 26th International Conference on Computer and Information Technology. Cox's Bazar, Bangladesh, Dec 13–15, 2023: 1–6.

Mandal B., Okeukwu A., Theis Y. (2021): Masked Face Recognition using ResNet-50. Available at https://arxiv.org/abs/2104.08997 (accessed July 12, 2024).

Molnar C. (2022): Interpretable Machine Learning: A Guide for Making Black-box Models Explainable. Vol. 2. Available at christophm.github.io/interpretable-ml-book/ (accessed July 12, 2024).

O'Shea K., Nash R. (2015): An Introduction to Convolutional Neural Networks. Available at https://arxiv.org/abs/1511.08458 (accessed July 10, 2024).

Plant Village (2024): The Plant Village Solution. Available at https://Plantvillage.Psu.Edu/ (accessed July 10, 2024).

Rangarajan A., Purushotaman R., Ramesh A. (2018): Tomato crop leaf disease classification using pre-trained deep learning algorithm. Procedia Computer Science, 133: 1040–1047.

Vinay P., Kumar G., Aravind K. (2020): Classification of tomato plant leaf disease using neural networks. In: Proceedings of the International Conference on Innovative Computing and Communication. New Delhi, India, Feb 20–22, 2020: 1–5.

Yebasse M., Shimelis B., Warku H., Ko J., Cheoi K. J. (2021): Coffee disease visualization and classification. Plants, 10: 1257.

Zhang K., Wu Q., Liu A., Meng X. (2018): Can deep learning identify tomato leaf disease? Advances in Multimedia, 2018: 710865.

Received: July 29, 2024
Accepted: November 26, 2024
Published online: March 11, 2025