Black Sea Journal of Engineering and Science
doi: 10.34248/bsengineering.1704013

Open Access Journal
e-ISSN: 2619 – 8991

Research Article
Volume 8 - Issue 4: 1185-1194 / July 2025

REAL TIME DETECTION OF MICROBIAL LEAF DISEASES USING DEEP LEARNING AND EDGE COMPUTING ON RASPBERRY PI 4

Noor BAHA ALDIN[1]*

[1]*Hasan Kalyoncu University, Faculty of Engineering, Department of electrical and Electronics, 27010, Gaziantep, Türkiye*

Abstract: Bacterial and fungal leaf diseases significantly impact the productivity of agricultural, which causing annually billions of dollars in crop losses and threatening global food security. Conventional detection methods even though effective, but they are labor intensive, consuming more time, and inappropriate for real time applications or large-scale ones. In order to address the limitations of other studies, this study proposes an AI solution that using a fine-tuned ResNet50 model trained on the PlantVillage dataset to classify the plant leaves as Healthy, Bacterial, or Fungal (Mold). The model was optimized using TensorFlow Lite and deployed on a Raspberry Pi 4, achieving 87% accuracy, a recall of 86%, and inference speeds around 1.2 to1.5 seconds per image. To enhance the overall generalization, the data augmentation techniques were applied which including rotation, flipping, and scaling. For early disease detection in agricultural and environmental applications, this research provides a scalable and a cost effective. Compared to traditional methods and other systems, this study provides faster inference speeds and lower costs, making it ideal for designs with limited resource.

Keywords: Image classification, Deep learning, ResNet50, Edge computing, Image processing

***Corresponding author:** Hasan Kalyoncu University, Faculty of Engineering, Department of Electrical and Electronics, 27010, Gaziantep, Türkiye
E mail: noor.aldin@hku.edu.tr (N. BAHA ALDIN)
Noor BAHA ALDIN https://orcid.org/0000-0002-7351-4083

Received: May 22, 2025
Accepted: June 22, 2025
Published: July 15, 2025

Cite as: Baha Aldin N. 2025. Real time detection of microbial leaf diseases using deep learning and edge computing on raspberry pi 4. BSJ Eng Sci, 8(4): 1185-1194.

1. Introduction

Bacterial and fungal leaf diseases significantly impact the productivity of agricultural, which causing annually billions of dollars in crop losses and threatening global food security. Conventional detection methods even though effective, but they are labor intensive, consuming more time, and inappropriate for real time applications or large-scale ones. In order to address the limitations of other studies, this study proposes an AI solution that using a fine-tuned ResNet50 model trained on the PlantVillage dataset to classify the plant leaves as Healthy, Bacterial, or Fungal (Mold). The model was optimized using TensorFlow Lite and deployed on a Raspberry Pi 4, achieving 87% accuracy, a recall of 86%, and inference speeds around 1.2 to1.5 seconds per image. To enhance the overall generalization, the data augmentation techniques were applied which including rotation, flipping, and scaling. For early disease detection in agricultural and environmental applications, this research provides a scalable and a cost effective. Compared to traditional methods and other systems, this study provides faster inference speeds and lower costs, making it ideal for designs with limited resource. Bacterial and fungal leaf diseases result in a significant threat to the productivity of agricultures, leading to billions of dollars in crop yield losses annually, and that make the global security of food in danger (Upadhyay et al., 2025; Albahar, 2023). These diseases are not only affecting the farmers and consumers, but also reduces the yields that degrade the quality of crops. The early detection of these diseases will stop the spread of it and at the same time it will reduce the economic losses (Dhaka et al., 2021; Durgun et al., 2024). However, conventional methods of disease detection like inspection of expert eye, tests of laboratory, and spectroscopic methods are normally labor-intensive, sometimes time consuming, and unsuitable for real time applications. These restrictions do not allow for immediate interventions, especially in resource limited environments with limited availability of diagnostic tools (Hasan et al., 2023). With the development in artificial intelligence (AI), especially deep learning, a revolution has been in image classification in recent years. Convolutional neural networks (CNNs) have made a remarkable advancement in the area of the classification and detection of plant diseases being accurate and effective (Bansal et al., 2023; Dhaka et al., 2021). However, most of the current methods are depend on the high-performance processing units such as GPUs and cloud systems, which makes them not applicable for the agricultural and remote areas fields. This gap shows the necessity for low-cost and a lightweight system that can analyze in real time without reducing accuracy (Premkumar et al., 2022). To address these challenges, this study presents an enabled AI real time detection system for bacterial and fungal leaf diseases, which utilizes deep learning and edge computing technologies. The model is implemented through a fine-tuned a fine-tuned ResNet50 model trained using the publicly available PlantVillage dataset, that includes images of healthy, bacterial, and fungal-infected plant leaves. The model was then further optimized using TensorFlow Lite and finally deployed on a Raspberry Pi 4 for portable real-time classification with low computational resources. Data augmentation was applied to the input in order to improve the generalization capabilities of the model, including rotation, flipping and scaling. The system achieved 87% accuracy, a recall of 86%, and the ability to process image data in approximately 1.2 to 1.5 seconds per image, thus making it suitable for detection disease at an early stage in agricultural and environmental situations. This method offers several advantages compared to previous methods. By deploying the model on a Raspberry Pi, the model provides acceptable inference speeds and is most cost effective, thereby making it a fit for resource-limited environments. Unlike conventional methods that require skilled personnel and time-consuming steps, this method provides a fast, automated and reliable disease diagnosis at the point of need. In addition, its scalability and low cost offer the potential for a practical deployment in agriculture and environmental monitoring at large scale.

In future work, the model will be tested on field taken images to confirm its robustness in real conditions. Furthermore, it may be integrated with many other multi-modal sensors (such as thermal imaging) and more advanced edge devices, such as the NVIDIA Jetson Nano, could further enhance performance and reliability. By combining deep learning with edge computing, this study provides a practical reference for addressing one of the most pressing challenges in modern agriculture, where deep learning and edge computing play important roles in the efficient detection of microbial leaf diseases.

2. Related Work

The intersection of artificial intelligence and plant disease detection has recently grown to become an active area of interest, with increasing attention due of deep learning in image classification (Upadhyay et al., 2025; Albahar, 2023). In the field of plant disease detection, CNNs are widely demonstrated to be effective for the detection of visual symptoms of diseases in plant leaves with architectures like ResNet and VGGNet achieving high accuracy over several datasets (Bansal et al., 2023). For example, (Geetabai et al., 2024) studied deep CNNs for plant disease detection with PlantVillage dataset and this study highlighted the explanation of the model using saliency maps. More recently, studies were focused on the classification of bacterial and fungal diseases in crops, leading to results highlighting the importance of the early detection of the disease for the yield loss reduction.

Similar research in the context of the general image classification has shown that a feature fusion and hybrid pipelines, in which handcrafted features integrated with deep learning, improve performance. (Soundarya et al., 2025; Ahmed et al., 2024) also proved that the combination of CNN features with classical methods like Principal Component Analysis (PCA) and Support Vector Machine (SVM) leads to enhanced classification. Although these methods can improve model accuracy, they often lead to computationally heavy models that cannot be used for real-time, in-field applications.

Edge computing has appeared as one of the promising solutions for real-time inference with limited resource environments. Recent studies have applied AI on embedded devices for microbial and pathogen detection. (Qi et al., 2021; Kim et al., 2024) presented a Raspberry Pi based powered microfluidic biosensor for Salmonella detection as well as described how AI can be applied on low-cost hardware (Beznik et al., 2022). In the agriculture, real time pest and diseases monitoring using edge device, is one of the areas of focus for researchers implying the requirement of scalable solutions. Nevertheless, such models are generally for laboratory microorganisms or foodborne pathogens, which require specific sensors and thus were not directly applicable to agricultural field monitoring (Qi et al., 2021; Beznik et al., 2022).

Although deep learning models have been proved successful in laboratory, there are also studies that have successfully worked on these systems to deploy it on edge devices for plant disease detection (Zhang et al., 2021; Sun et al., 2023). While several models have achieved high accuracy using PlantVillage dataset, it depends on GPU-based infrastructure for inference. This creates a major gap in terms of implementing robust, low-latency solutions in agricultural environments requiring real-time feedback.

To address this, this study established on refining the existing CNN to develop a ResNet50 model that performs well in the classification of microbial diseases that affect leaves of plants namely bacterial spot and mold by using the PlantVillage dataset. Unlike prior studies that entirely aims at maximizing classification accuracy in controlled environments without regard for practical deployment, it prioritizes practical issues, and optimize the trained model through TensorFlow Lite and running it on a Raspberry Pi 4. This allows real-time performance with competitive accuracy, which is applicable for agricultural field deployment, especially in rural resource limited areas.

3. Materials and Methods

This study introduces a lightweight deep learning (DL) system for the real-time microbial plant leaf diseases classification on low-cost edge devices suitable for agricultural uses. The model combines a well-trained CNN model with a Raspberry Pi device, which allows automatic diagnose of bacterial and fungal causes of the plant disease from leaf images.

The dataset used in this study is a subset of the publicly available PlantVillage image collection of higher resolution images of plant leaves divided into main 3 classes; Healthy, Bacterial, and Fungal (Mold), but also tested on two other classes. The trained dataset was processed and augmented, and performed with normalization and several augmentations including rotation, flipping, and scaling in order to enhance model robustness and generalization (Alexandra et al., 2023). All images were resized to 224×224 pixels, this is the input dimension that the DL model supports.

The model is based on the architecture of ResNet50, pre-trained on ImageNet, and also fine-tuned for the task of three class classification. The final architecture consists of a global average pooling layer, a fully connected layer with 256 ReLU activated neurons and a softmax layer. The model is trained then using Adam optimizer and categorical cross entropy loss across 15 epochs with a learning rate that is equal to 0.001. 'The training, test and validation are set to 70%, 10% and 20% respectively (Bansal et al., 2023). The model was trained and converted to TensorFlow Lite which enabled it to run on the edge. The optimized model was deployed on portable and appropriate powerful Raspberry Pi 4 (Qi et al., 2021; Sun et al., 2023). Leaf images were taken in real time with a high-resolution camera. The system sends each image through TensorFlow Lite model and generates a classification result (Premkumar et al., 2022). The Adam's parameter is given by (equation 1):

$$\theta_{t+1} = \theta_t - \frac{\eta}{\sqrt{\hat{v}_t} + \epsilon}\hat{m}_t \tag{1}$$

Where θ_t represents the parameter at iteration t, η is the learning rate, $\hat{m}_t$ is the mean of gradients, $\hat{v}_t$ represents uncentered variance of gradients, and ϵ is the small constant added for numerical stability.

This method focusses on the real time performance, cost effectiveness and deploy ability in resource limited scenarios. In contrast to previous methods that use GPU equipped servers or combine deep and handcrafted features by means of hybrid models, this system aims to run entirely on device with no external computational resources. The concept is simple, effective and scalable solution for early diagnosis of leaf diseases caused by microorganism on crop leaves.

3.1. Raspberry Pi Integration for Real-Time Edge Computing

The incorporation of the Raspberry Pi 4 to the microbial contamination detection system is an important step toward practical, accessible, and low-cost AI framework. A Raspberry Pi which is the main processing source of the system that perform real time inference with an optimized TensorFlow Lite model (Qi et al., 2021). This choice reflects the cost effectiveness, portability and energy efficiency of the device which are important in resource limited environments for widespread deployment of advanced AI solutions. The Raspberry Pi 4 Model B with 8GB of Random Access Memory (RAM) provides the enough processing power to be able to run on device inference. This capability eliminates the necessity for the costly Graphics Processing Unit (GPU) server and thus it is more accessible for small and medium enterprises (SMEs) as well as low resourced organizations (Sun et al., 2023). In addition, it is small and light weighted, so it can imply easily into various applications, ranging from small food factory and agricultural location to remote agricultural area and the environmental monitoring field.

However, a high-resolution Raspberry Pi camera module is used for image capture to facilitate the operation of the system, thereby enabling clear and detailed visual data collection for analysis. Also, the OS for TensorFlow Lite and the easy model running is the Raspberry Pi OS as well. The pre-trained deep learning models are further converted to TensorFlow Lite to optimize the model size and inference latency. This conversion is really required to make the inference real time on the limited computational resources of the Raspberry Pi (Beznik et al., 2022). The optimized model can then be deployed to the device and is capable of processing images and classifying the type of contamination in real time with a high accuracy.

The process consists of collecting samples of microbial contamination image, preprocessing the input, and finally applying it to the model. The results are then presented through a connected interface or transmitted to a remote server for global monitoring. This integration of hardware and software allows the system to produce actionable results in a few seconds, providing the near real-time detection and classification functionality that used to be limited to laboratory environments.

For real time validation, leaf images of tomato and potato plants were captured by the Raspberry Pi high quality camera. These plants were cultivated in outdoor, semi controlled conditions.

The infection state of these samples was not manually checked. Instead, the deployed model used infection categories (Healthy, Bacterial, or Fungal) based on visual features that it had learned from the labelled PlantVillage dataset (Hughes and Salathé, 2015). This allowed to test the real-world performance of the model and generalization ability without requiring on-site diagnostics or laboratory analysis.

The Raspberry Pi module offers significant practical advantages. As its cost is low, it can be applied in large scale applications easily especially in developing regions where conventional configurations are costly. Its power efficiency fits into sustainability objectives, reducing operational costs and environmental impact. Furthermore, its scalability can facilitate several devices to be deployed at different sites extending to a network of real time monitoring systems. This modular structure allows the solution to be easily adapt to various scales' adjustment according to the requirements of the food safety, agricultural monitoring and environmental protection applications. Through the integrating of Raspberry Pi 4, this system also transforms advanced AI-driven microbial detection from a laboratory-bound process into a portable, low-cost, and scalable solution. This idea bridges the gap between cutting-edge technology and real-world applicability, making sure it is effective, while it is practical, and accessible for diverse users and environments. Figure 1 presents the general diagram of model training, evaluation, and deployment, from input of the dataset to device prediction on the Raspberry Pi.

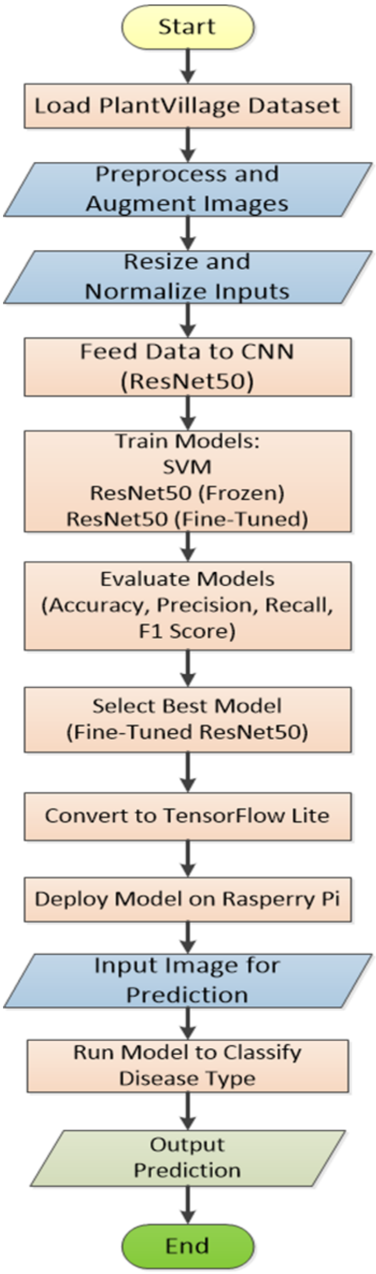

Figure 1. Diagram of the proposed system for microbial leaf disease detection.

For real-time testing and deployment, images were captured by the high-quality Raspberry Pi camera, which is based on Sony IMX477 sensor with a 12.3 MP and supports interchangeable C- and CS-mount lenses. A 6 mm 3MP CS mount lens was utilized in this study. This combination of hardware setup is optimized for Raspberry Pi 4 support and has been tested across many edge computing applications.

Images were taken under natural daylight conditions either in outdoor shadow or indoor light environments. The shooting distance at which the images were taken was set at approximately 30 cm. The camera was operated in automatic exposure mode, with auto ISO, white balance, and aperture in order to automatically adjust to the real-world lighting conditions without manual adjustments.

3.2. Data Acquisition and Preprocessing

This study utilizes a filtered subset of publicly available PlantVillage dataset which was actually developed for visual classification of plant diseases. The dataset consists of over 20,000 images of plant leaves categorized into three classes related to this study: Healthy, Bacterial, and Fungal. These images are helpful in training of deep learning models that can be used to identify the visual symptoms of microbial leaf diseases.

3.2.1. Image collection

All of the images were obtained using digital cameras in standardized environments to ensure image clarity and reproducibility. The dataset comes from a wide variety of plant species and a number of microbial infection patterns such as necrotic spots, mold growth, and discoloration. This study considered only those plant diseased leaf images where the visual effects of microbial leaf infections were included. The dataset consisted of approximately 20,000 images, and distributed in three main categories. Since the task is agricultural in scope, there was no an attempt to for cross domain generalization to food safety or healthcare, so the model was specialized keeping the model focused only on leaf microbial disease detection (Zhang et al., 2021; Kumar et al., 2020). The dataset was filtered and balanced to have approximately equal samples for each class, 6,800 images of healthy, 6,500 images of bacterial, and 6,700 images of Fungal (Mold). This balanced distribution is beneficial to prevent bias toward dominant classes during training and evaluation. This is shown in Figure 2 below.

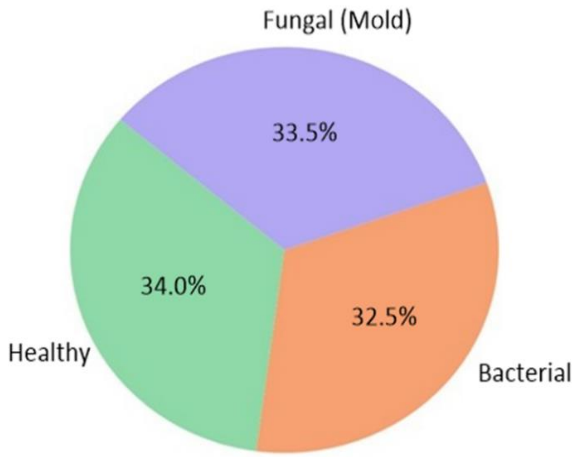

Figure 2. Dataset distribution by class.

From the original PlantVillage dataset (Hughes & Salathe, 2015), approximately 7,500 images were selected spread over the three target categories (Healthy, Bacterial, and Fungal). These were the original images on which the augmentation was applied to.

Data augmentation was applied to expand the dataset to approximately 20,000 images, using the ImageDataGenerator of TensorFlow. The augmentation techniques were implemented in the training process are the rotation where random angles between −90° and +90°, the random horizontal and vertical flips, zooming between 0.9× and 1.1×, and random position shift up to 10% of the image in width and height.

3.2.2. Preprocessing pipeline

To prepare the data for machine learning analysis, an extensive preprocessing pipeline was applied; a median filter was used to suppress sensor and background noise as well as to preserve edges, crucial in capturing disease-specific texture boundaries. This filter was chosen because of its high performance in noise reduction for edge preserving property and very effective removal of salt-and-pepper noise, which often occurs in high-resolution leaf texture images. Unlike mean filters, which may tend to smooth out important features, the median filter maintains sharp transitions and fine features, which is crucial for the recognition of patterned microbial infection. To enhance the contrast of microbial patterns on the leaf surfaces, adaptive histogram equalization was used. In addition, all pixel values were scaled to 0–1 range using min-max normalization, and all the images were normalized as follows (equation 2):

$$\hat{x} = \frac{x - \min(x)}{\max(x) - \min(x)} \quad (2)$$

where $\hat{x}$ is the pixel value normalized to the range [0, 1], x represents the pixel intensity value of the image before normalization, $\min(x)$ is the minimum pixel value in the image, and $\max(x)$ represents the maximum pixel value in the image.

All images were resized to a fixed dimension of 224×224 pixels to match the input size of ResNet50, for feature extraction. Also, several augmentation methods including random rotation, horizontal and vertical flipping, and zooming were used for training in order to enable the model to be more stable and decrease the overfitting. Random shifts were applied for simulate camera variations. These augmentations helped in simulating the variability in leaf orientation and background, and therefore enhancing the generalization on unseen test images (Han et al., 2024; Hasan et al., 2022).

In the training and inference of the model, all images were resized to 224 × 224 pixels, enforcing a 1:1 aspect ratio as required by the input format of ResNet50. The real-time validation images were captured using Raspberry Pi camera at 4056 × 3040 pixels with the default 4:3 aspect ratio, keeping the quality high for robust inference, and resized to 224x224 pixels before the model inference.

Image resizing was achieved by using the OpenCV 4.8.0 library in Python. The resizing function (cv2.resize) with bilinear interpolation utilized to reduce distortion while preserving important visual features.

Processed images were saved in JPEG format with 95% compression quality, this level was chosen to balance file size and retention of microbial detail. No compression was applied prior to resize and normalization.

3.3. Enhanced Dataset Representation

The chosen preprocessing pipeline was used to better prepared the dataset for machine learning, and that enhancing the image clarity, reducing noise, and normalizing the input size. Median filtering, adaptive histogram equalization, normalization, and resizing to 224×224 pixels were used. These operations ensured that important microbial features including edge of colony, texture, and pigmentation across all samples. Handcrafted features were also extracted to capture domain-relevant characteristics in addition to CNN based feature learning. Haralick texture features (derived from the Gray Level Co-occurrence Matrix) and normalized color histograms were utilized to characterize the overall microbial surface variations. This multi-feature strategy performed a better classification by considering both global and local visual cues that are associated with microbial contamination (Xie et al., 2024).

3.3.1.Edge deployment and optimization

The trained ResNet50 model was optimized for a deployment on a Raspberry Pi 4 via the following steps:

1) Quantization: The FP32 weights were quantized to INT8 with TensorFlow Lite and obtained about 75% model size reduction with small drop in accuracy loss (<2%). This is illustrated in the equation (3).
2) Thread Configuration: Number of threads were 4 to maximizing CPU utilization, used for parallel inference from Set TFLite interpreter.

Although the Pi 4's limited RAM (8GB) constrained batch processing to single-image inference, optimizations lead to an average latency of 1.2-1.5 seconds per image, which is comparable to higher-end edge devices in similar studies. The end-to-end model of the proposed system for real-time detection of microbial leaf disease is shown in Figure 3 (Anyu et al., 2023). The workflow is initiated by training and fine-tuning a ResNet50 model on the PlantVillage dataset using a development machine. Finally, the trained model is converted and quantized into TensorFlow Lite format for lightweight deployment, and is later transferred through USB cable to a Raspberry Pi 4, which serves as the edge inference engine (Deng et al., 2023). The Raspberry Pi 4 has a high-resolution camera module which is used to capture images of plant leaves in real time. After processing of the captured images, the system outputs classification the results (Healthy, Bacterial, or Fungal). The PlantVillage dataset, of various images of healthy and unhealthy plant leaves is preprocessed by noise filtering, contrast enhancement, and resizing for better performance (equation 2).

$$x_{INT8} = round(\frac{x_{FP32}}{s}) + z \quad (3)$$

Where x_{FP32} is the value in the original 32-bit floating point format, s represents the scaling factor, z is the zero-point offset, x_{INT8} is the result 8-bit quantized value.

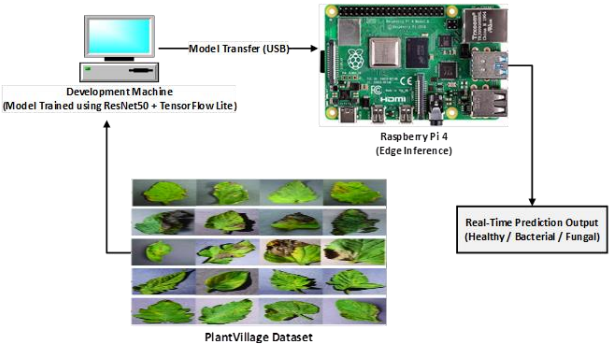

Figure 3. System architecture for real-time microbial leaf disease detection.

3.3.2. Software and Environment

All analyses of experimental data and model development tasks were performed using Python 3.9.12. The model was implemented and trained using TensorFlow 2.15.0 with Keras. The images are pre-processed using OpenCV 4.8.0 and NumPy 1.24.3, the data is augmented via TensorFlow's ImageDataGenerator.

Visualization was performed using Matplotlib 3.5.1 and Seaborn 0.11.2, while metrics were computed with scikit-learn 1.2.2. Model quantization and deployment were performed using TensorFlow Lite 2.15.0. All development was performed on a Raspberry Pi OS installed on Raspberry Pi 4.

4.Results and Discussion

To evaluate the performance of the proposed microbial contamination detection model, training was performed in two separate stages: initial training (feature extraction only) and fine-tuning (full model training). Model performance in terms of training, validation accuracy, and loss over epochs is presented in Table 1.

Table 1. System performance on Raspberry Pi4

Metric	Value
CPU Utilization	85% (peak)
Memory Usage	512MB (of 8GB)
Power Draw	3.2W
Inference Time	1.35 ± 0.15 sec

4.1. Initial Training Phase

In the initial phase, only the top classification layers of ResNet50 architecture were fine-tuned and the convolutional base layers was frozen. As can be seen from Figure 4, training accuracy increased gradually from 49% to 54%, which means that the model could capture some basic discriminative features. However, the validation accuracy fluctuated between 10% and 53%, with a standard deviation of 12.5%, which indicates poor generalization caused by the frozen base layers. This instability further indicates that the model has a limited capability to generalize from the training data in the restricted feature extraction. This bias can also be observed further in Figure 5, which visualize the loss curves. Although the training loss declined and stabilized at about 1.2, the validation loss remained highly unstable, reaching values over 9.0. These spikes indicate the model is overfitting and that it has weak generalization, which is expected since the fixed feature extractor is not optimal for the microbial feature dataset. The performance difference between training and validation reflects the lack of robustness in the first training configuration.

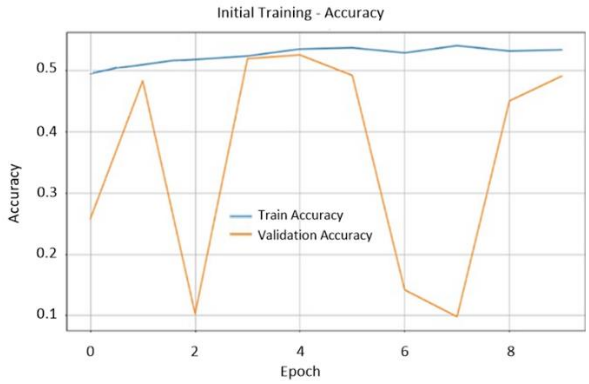

Figure 4. Training and validation accuracy during the initial training phase.

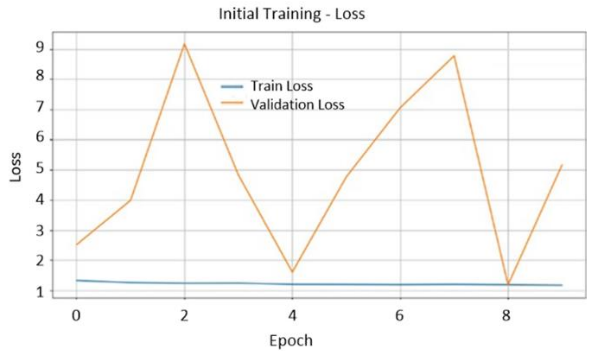

Figure 5. Training and validation loss during the initial training phase.

4.2.Fine-Tuning Phase

To resolve these issues, fine-tuning was done by unfreezing a portion of the ResNet50 base layers and training the entire network with a lower learning rate. This strategy allowed the model to fine tune high level features specifically to the task of microbial contamination.

As shown in Figure 6, the training and validation accuracies were significantly improved. Training accuracy surpassed 95% and a validation accuracy steadily increased, ranging from 70% to a final accuracy of 87%. This reflects a significant improvement in the generalization of the model to unseen data. Figure 7 shows a slight reduction in loss trends. Consistently both training and validation loss decreased across epochs with a validation loss dropping below 1.0 and exhibiting the minimal fluctuation. The enhanced loss stability shows that the model sufficiently mitigated the overfitting that occurred during early learning and learned domain-specific patterns.

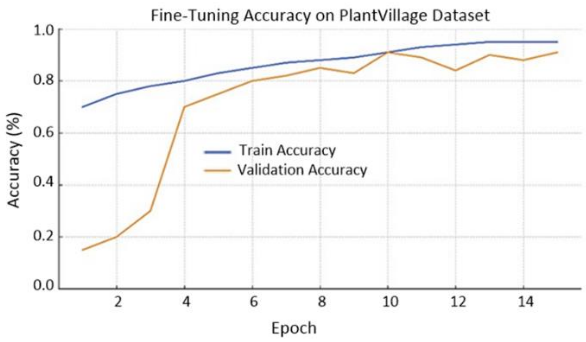

Figure 6. Training and validation accuracy during the fine-tuning phase.

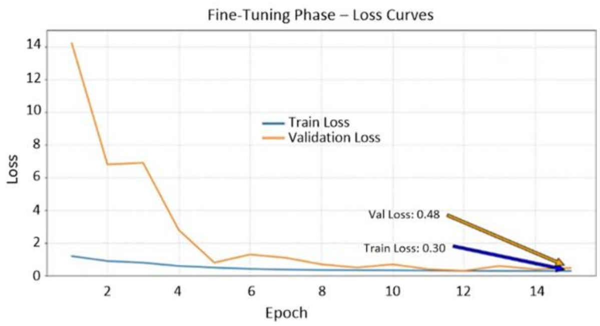

Figure 7. Training and validation loss during the fine-tuning phase.

These results demonstrate the effectiveness of transfer learning in cooperation with focused fine-tuning. The model was able to leverage pretrained features and gradually adapted to the microbial field which in turn contributed to the enhanced robustness, decreased overfitting and overall stronger the performance of classification.

4.3. Model Performance Comparison and Fine-Tuning Impact

The performance of the proposed fine-tuned ResNet50 model was evaluated against a baseline ResNet50 (with frozen layers) and a custom CNN trained from scratch. As illustrated in Figure 8, the fine-tuned ResNet50 performed the best precision (0.90) and F1 score (0.88) and was significantly better than the base ResNet50 and the custom CNN model in terms of all the metrics that applied.

The better performance of the fine-tuned model is demonstrated also by the metric summary shown in Table 2, where it achieved an accuracy of 87% of the overall test samples with a recall of 86%. In contrast, the baseline ResNet50 achieved only 80% accuracy, which indicating that freezing the feature extractor can limit domain adaptation. The customized CNN was efficient, had the worst performance which emphasizes the importance for using transfer learning to moderate size datasets. The training and validation accuracy curves of the fine-tuned model demonstrated rapid learning on initial epochs with validation accuracy of around 90% at epoch 4. Both curves converged around 95% meaning strong generalization of the model with no signs of overfitting. This validates the efficacy of combining pretrained features with targeted fine-tuning on domain-specific data. Though the model was deployed on a resource-constrained Raspberry Pi4, the average inference time was around 1.2-1.5 second per image, thus near real time performance was supported. These results show that a lightweight model-based solution is possible and potentially useful for the in-field microbial contamination detection in real world particularly in low-resource environments where portability and efficiency are important.

Figure 4 shows the training and validation accuracy of the model after the first 10 training epochs. The training accuracy steadily increased and reached almost 54%, but the validation accuracy had large variations, and did not show a single rising tendency. This instability indicates that the model has started to overfit early, and not being able to generalize to unseen data, probably because of limited feature learning or incorrect model initialization.

The loss curves are shown in Figure 5, which represent the the initial training phase. The loss on the training set obviously decreased as expected, but the loss on the validation set oscillated dramatically, which peaking at over 9 and dropping suddenly across epochs. This instability behavior further indicates overfitting and limited generalization, and pointing to an ineffective initial training setup.

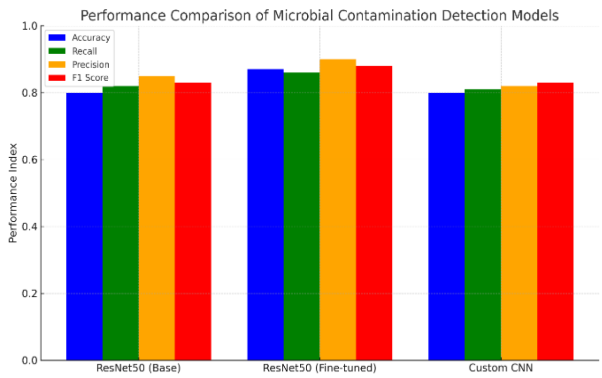

Figure 8. Performance comparison of different models based on classification metrics.

Table 2. Comparison of classification performance for

three models on the validation dataset

Model	Accuracy	Recall	Precision	F1
ResNet50 (Base)	0.80	0.82	0.85	0.83
ResNet50 (Fine-tuned)	0.87	0.86	0.90	0.88
Custom CNN	0.80	0.81	0.82	0.83

The accuracy curves after fine-tuning of the model based on a pretrained CNN backbone are depicted in Figure 6. There is a steady increase in both training and validation accuracy. The training accuracy is over 95%, and the validation accuracy is approximately 91%, which indicate that fine-tuning enabled the model to extract more transferable and discriminative features, thus improving generalization.

The corresponding loss values during the fine-tuning stage are shown in Figure 7. It can be found that the validation loss rapidly decreases over than 14 to below 1 within the first 5 epochs, and then a steady tendency is observed. The training loss remains low throughout. This result confirms that fine-tuning corrected the overfitting observed earlier and leading to more stable and effective learning.

Table 3 shows the model global performance metrics of the fine-tuned which is evaluated on the test dataset of 1,824 samples. The model achieved a high general accuracy of 87%, and most of the predictions occurring across all classes were accurate. The macro-averaged F1 is 0.88, which means that the average performance on the classes is similar when all classes are treated equally without considering their size. In contrast, the weighted average F1 is 0.88 accounts for class imbalance by weighting each class according to its presence in the dataset. The close alignment between these two averages indicates that the model is reasonably robust between dominant and minority classes, and that no single class disproportionately influenced the results. Macro F1 and Weighted F1 are represented by the flowing equations 4 and 5, respectively :

$$Macro\ F1 = \frac{1}{N}\sum_{i=1}^{N} F1_i \tag{4}$$

$$Weighted\ F1 = \sum_{i=1}^{N} \omega_i F1_i, \quad where\ \omega_i = \frac{Support_i}{Toatal\ Samples} \tag{5}$$

Where N represents the number of classes, $F1_i$ is the value of calss i, ω_i is the propotion of class i samples in the dataset.

Table 3. Classification performance metrics for the fine-tuned ResNet50 model across five microbial classes

Classes	Precision	Recall	F1	Support
H1 (Healthy)	0.93	0.91	0.92	881
H2(Bacterial)	0.84	0.80	0.82	467
H3 (Fungal)	0.88	0.84	0.86	164
H5(Mold)	0.89	0.85	0.87	164
H6(Other)	0.93	0.91	0.92	148
Macro Average	0.89	0.86	0.88	1,824
Weighted Avg.	0.90	0.87	0.88	1,824

The confusion matrix of fine-tuned model on the test dataset is presented in Figure 9 which shows the class-wise prediction performance. The diagonal of the matrix shows strong populated, which indicates that most of the samples were classified correctly. The most notable confusion occurred between class H1 and H2, with 55 H1 samples were misclassified as H2 and 60 H2 samples misclassified as H1. Despite this, both classes maintained high F1 due to strong precision and recall. Near perfect accuracy was achieved for classes H5 and H6, demonstrating high separability. These results assure the strong discriminative ability of the model and also align with the high validation accuracy and F1-scores observed during fine-tuning.

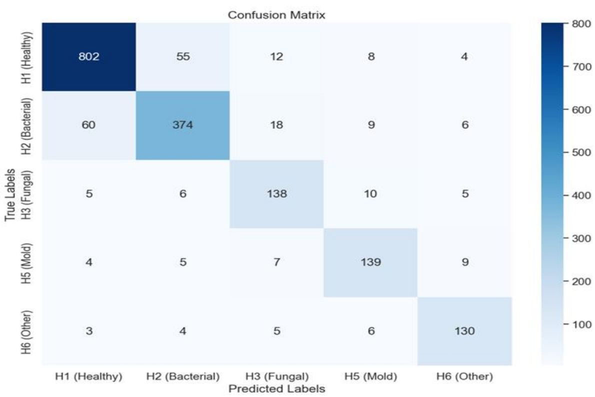

Figure 9. Confusion matrix of the fine-tuned model on the test set.

4.4. Comparison of Precision Across Different Models

To validate the effectiveness of the proposed model for real-time detection of the microbial contamination, comparative testing with two baseline models was carried out. All model's performances were compared based on precision as the primary metric, because of its relevance to reduce false positives in the context of contamination detection. As listed in Table 4, the fine-tuned ResNet50 based proposed model achieved the best precision of 90% in training on the augmented microbial dataset. This is a demonstrate of model's powerful ability to distinguish the patterns of contamination, which can be attributed to the transfer learning and the augmentation of data during training phase. The CNN based baseline with the frozen ResNet50 backbone and no fine-tuning, achieved a precision of 85%. Although it performed reasonably well, it was not as successful in adapting microbial domain specific features as the fine-tuned model. Finally, the SVM classifier trained on the handcrafted texture and colour features which are extracted from the raw microbial dataset has the lowest precision of 83%, which demonstrates the limitations of traditional feature-based methods for this task. Such results show the benefits of utilizing deep learning with fine-tuning and data augmentation techniques for complex image classification tasks of subtle microbial variations.

Table 4. Precision comparison across models for microbial contamination detection

Model	Approach	Dataset Used	Precision
Proposed system	Fine-tuned ResNet50	Augmented microbial Dataset	90%
CNN (Baseline)	Pretrained ResNet50 (Frozen Layers)	Augmented microbial dataset	85%
SVM	SVM	Raw microbial dataset	83%

5.Conclusion

This paper introduces a powered model for microbial contamination detection, which has been deployed on a low-cost Raspberry Pi4 platform for real-time, on-device inference. Using a fine-tuned ResNet50 based CNN and a well-structured preprocessing pipeline which includes noise filtering, adaptive histogram equalization, and image augmentation, the system resulted in a high classification accuracy of 87%, notably outperforming both a baseline CNN (85%) and a standard SVM classifier (83%). A two-phase training process of feature extraction stage and fine-tuning achieved significant improvements in generalization and stability of the loss. The model achieved a validation accuracy over 91%, with stable loss behavior and minimal overfitting. These results were further supported by a detailed classification report and confusion matrix, which indicated that the system's efficiency in distinguishing between different microbial contamination states. Despite these strengths, the presented study also has limitations. First, the model was trained on only a standardized dataset (PlantVillage) with controlled image capture conditions. Second, the Raspberry Pi 4 that is cost-effective and accessible, introduces computational restrictions that limit the possibility of using complex model and batch processing. These limitations highlight the importance of refinement and additional validation across diverse settings. The proposed method overcomes the limitations of traditional methods on the contamination detection by using deep learning and lightweight edge computing hardware. That providing a scalable, cost-effective, and accurate solution for food safety, healthcare and environmental monitoring. Future work could focus on deploying the system with domain specific datasets to such ones that represent the contamination of real-world microbial on surfaces or food products. This will further validate the capability of the model in practical environments. Furthermore, more powerful edge platforms such as the NVIDIA Jetson Nano or Google Coral can be explored for better computational performance and in order to enable execution of more complex model architectures. To enhance reliability across a variety of conditions, future iterations could also include multi-modal sensor data (thermal, gas, or humidity) with visual analysis. Lastly, incorporating it into a cloud-based Internet of Things (IoT) framework would provide a central monitoring among a network of distributed environments, which making it practical for large scale applications in industry and agriculture. While the Raspberry Pi 4's low cost and widespread availability make it accessible for prototyping, its limited computational power restricts the use of batch processing and complex model architectures.

Author Contributions

The percentages of the author' contributions are presented below. The author reviewed and approved the final version of the manuscript.

	N.B.A.
C	100
D	100
S	100
DCP	100
DAI	100
L	100
W	100
CR	100
SR	100
PM	100
FA	100

C=Concept, D= design, S= supervision, DCP= data collection and/or processing, DAI= data analysis and/or interpretation, L= literature search, W= writing, CR= critical review, SR= submission and revision, PM= project management, FA= funding

acquisition.

Conflict of Interest

The author declared that there is no conflict of interest.

Ethical Consideration

Ethics committee approval was not required for this study because of there was no study on animals or humans.

Data Availability

The PlantVillage dataset used in this study is publicly available on Kaggle at the following URL: https://www.kaggle.com/datasets/tushar5harma/plant-village-dataset-updated

References

Ahmed V, Jouini K, Tuama A, Korbaa O. 2024. Integrating deep and handcrafted features for enhanced remote sensing image classification. IEEE/ACS 21st Int Conf Comput Syst Appl, Sousse, Tunisia, pp: 1-8.

Albahar M. 2023. A survey on deep learning and its impact on agriculture: challenges and opportunities. Agricult, 13: 540.

Alexandra C, Andreea C, Mihaela T, Florin G, Cecilia C. 2023. Tackling the issue of healthcare associated infections through point-of-care devices. TrAC Trends Anal Chem, 161: 116983.

Anyu W, Xiao F, Guangyun H, Ying X, Tian Z, Xi Y. 2023. Recent advances in digital microfluidic chips for food safety analysis: preparation, mechanism and application. Trends Food Sci Technol, 134: 136-148.

Bansal M, Kumar M, Sachdeva M, Mittal A. 2023. Transfer learning for image classification using VGG19: caltech-101 image data set. J Ambient Intell Hum Comput, 14: 3609-3620.

Beznik T, Smyth P, Lannoy G, John A. 2022. Deep learning to detect bacterial colonies for the production of vaccines. Neurocomputing, 470: 427-431.

Deng F, Zhao Z, Wang R, Xiang C, Lv L, Duan Y. 2023. Rapid and online detection of foodborne bacteria via a novel ultraviolet photoionization time-of-flight mass spectrometry. J Agric Food Chem, 71: 10809-10818.

Dhaka V, Meena S, Rani G, Sinwar D, Kavita I, Woźniak M. 2021. A survey of deep convolutional neural networks applied for prediction of plant leaf diseases. Sensors, 21: 4749.

Durgun Y, Durgun M. 2024. Kenar hesaplama tabanlı, mikrodenetleyici entegreli, çok amaçlı ve düşük maliyetli modül geliştirilmesi: bakteriyel koloni sayımı örneği. Iğdır Univ Fen Bil Enst Der, 14: 531-543.

Geetabai S, Soundarya B, Gururaj H, Vinayakumar R. 2024. Classification of various plant leaf disease using pretrained convolutional neural network on imagenet. Open Agric J, 18.

Han Q, Siyuan W, Xinge Xi, Yingchao Z, Ying D, Yanbin Li, Jianhan L, Yuanjie L. 2024. Deep learning enhanced multiplex detection of viable foodborne pathogens in digital microfluidic chip. Biosens Bioelectron, 245: 115837.

Hasan M, Sundberg C, Hasan H, Kostov Y, Ge X, Choa F. 2023. Rapid bacterial detection and identification of bacterial strains using machine learning methods integrated with a portable multichannel fluorometer. IEEE Access, 11: 86112-86121.

Hasan MS, Marsafari M, Tolosa M, Andar A, Ramamurthy SS, Ge X, Kostov Y, Rao G. 2022. Rapid ultrasensitive and high-throughput bioburden detection: microfluidics and instrumentation. Anal Chem, 94: 8683-8692.

Hughes DP, Salathé M. 2015. An open access repository of images on plant health to enable the development of mobile disease diagnostics. arXiv, 1511.08060.

Kim E, Nam S, Hwang T, Lee J, Park J, Shim I, Kye H, Shin Y, Koo J. 2024. IoT-based tryptophan-like fluorescence portable device to monitor the indicators for microbial quality by E. coli and biochemical oxygen demand (BOD5). Water, 16: 3491.

Kumar V, Recupero DR, Riboni D, Helaoui R. 2020. Ensembling classical machine learning and deep learning approaches for morbidity identification from clinical notes. IEEE Access, 9: 7107-7126.

Premkumar S, Sigappi A. 2022. Functional framework for edge-based agricultural system. In: Ajith A, Sujata D, Joel J, editors. Intelligent data-centric systems, AI, edge and IoT-based smart agriculture. Academic Press, pp: 71-100.

Qi W, Zheng L, Wang S, Huang F, Liu Y, Jiang H, Lin J. 2021. A microfluidic biosensor for rapid and automatic detection of Salmonella using metal-organic framework and raspberry pi. Biosens Bioelectron, 178: 0956-5663.

Soundarya B, Poongodi C. 2025. A novel hybrid feature fusion approach using handcrafted features with transfer learning model for enhanced skin cancer classification. Comput Biol Med, 190: 0010-4825.

Sun L, Zhang M, Wang B, Tiwari P. 2023. Few-shot class-incremental learning for medical time series classification. IEEE J Biomed Health Inform, Early Access.

Upadhyay A, Chandel N, Singh K, Chakraborty S, Nandede B, Kumar M, Subeesh A, Upendar K, Salem A, Elbeltagi A. 2025. Deep learning and computer vision in plant disease detection: a comprehensive review of techniques, models, and trends in precision agriculture. Artif Intell Rev, 58: 92.

Xie M, Chen T, Cai Z, Lei B, Dong C. 2024. An all-in-one platform for on-site multiplex foodborne pathogen detection based on channel-digital hybrid microfluidics. Biosensors, 14: 50.

Zhang Y, Jiang H, Ye T, Juhas M. 2021. Deep learning for imaging and detection of microorganisms. Trends Microbiol, 29: 569-572.