ORIGINAL RESEARCH **Open access**

Real-Time Root Crop Leaf Classification Using CNN Models on Raspberry Pi for Smart Agriculture

Ahmed Ben Ali[1*], Samira Gharbi[1], Nabil Jebali[2]

Abstract

This study describes the application of deep learning frameworks for distinguishing root crop foliage, including beetroot, potato, radish, and sweet potato. Using ResNet50 and DenseNet121, the research demonstrates classification on an extensive dataset comprising over 2,500 images sourced from multiple regions of Karnataka, India. Both frameworks delivered strong results, with ResNet50 attaining 99.60% accuracy and DenseNet121 reaching 97.60%. The models maintained high precision, recall, and F1 scores across all classes when running on the CPU. A notable milestone involved the practical deployment of these models on a Raspberry Pi 4B, where ResNet50 maintained 99.60% accuracy, and DenseNet121 achieved 96.81% accuracy on this constrained hardware. Field validation further confirmed practical utility, with the Raspberry Pi 4B assembly mounted onto a vehicle alongside a webcam to capture root crop leaves live, undergoing successful testing in a genuine farming environment. This established the system's suitability for instantaneous crop recognition. The findings underscore the potential of deep learning solutions within agricultural technology, delivering a trustworthy, efficient, and transportable method for categorizing root crop leaves. The unwavering high accuracy upheld across different computing platforms and real-life scenarios reveals the durability and flexibility of the constructed models.

Keywords Deep learning, Convolutional neural network, Agriculture automation, Raspberry Pi, Leaf classification

*Correspondence:
Ahmed Ben Ali
ahmed.benali@iresa.agrinet.tn
[1] Department of AI for Olive Fly and Citrus Pest Forecasting, Faculty of Agriculture, University of Tunis, Tunis, Tunisia
[2] Department of Smart Drip Irrigation Optimization, Faculty of Agriculture, University of Sfax, Sfax, Tunisia

Introduction

The expanding incorporation of Artificial Intelligence (AI) and Machine Learning (ML) into agriculture presents novel strategies for crop detection and oversight [1]. This investigation centers on fitting suitable deep learning models to the real-time differentiation of root crops—sweet potato, radish, beetroot, and potato. By automating the recognition workflow, the model promptly and accurately determines crop identity, reducing reliance on manual checks and improving operational efficiency.

In present-day agriculture, precise leaf classification is vital for equipping growers and agronomists to make informed decisions about crop vigor, harvest readiness, pest control, and yield optimization [1, 2]. Where standard crop identification relies on visual inspection by people—a process that is both tedious and prone to inaccuracies—this study proposes a more productive, precise, and expandable alternative grounded in deep learning methodologies.

Convolutional Neural Networks (CNNs) serve as the backbone for root crop leaf classification. CNNs excel at image recognition tasks because they can extract and propagate features to deeper network layers [3, 4]. A CNN comprises two

This is an open access article distributed under the terms of the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License (CC BY-NC-SA 4.0), which permits non-commercial use, distribution, reproduction, and adaptation in any medium or format, provided that appropriate credit is given to the original author(s) and source, a link to the license is provided, and any changes made are indicated. Adapted material must be distributed under the same license.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

essential components: feature extraction, which applies convolution and pooling operations to isolate meaningful image features, and classification, which produces output forecasts grounded in those features [5-8].

The arduous, slow pace of hand-sorting and identifying root crop foliage spotlights the pressing need for mechanization in farming, particularly to facilitate streamlined harvesting. This necessity prompted the present research. The investigation addresses this difficulty by developing a real-time root crop leaf identification scheme leveraging Convolutional Neural Networks (CNNs) running on a Raspberry Pi microcomputer. By marrying advanced deep learning approaches with affordable hardware, the project not only enables accurate crop recognition but also lays the foundation for deploying autonomous robotic platforms across cultivated land, amplifying output and operational capability in resource-tight contexts. This project uses ResNet50 and DenseNet121, two prominent deep learning architectures known for their performance in image classification [9-11]. The selection of ResNet50 and DenseNet121 stems from their proven ability to mitigate the vanishing gradient problem, promote efficient feature reuse, and maintain high accuracy even on devices with limited resources, rendering them especially suited for real-time farming applications [4, 9]. These models are trained on an exhaustive dataset of over 2500 real-time images, painstakingly gathered through visits to farming sites across diverse districts in Karnataka, India. Such field-driven data acquisition ensures the dataset reflects authentic and varied leaf features, thereby bolstering the models' discriminatory power. This direct engagement in data collection constitutes a meaningful scholarly contribution, providing a reliable foundation for rigorous model training and evaluation. The envisaged robotic carrier is designed to navigate autonomously across crop fields, specifically to classify root crop foliage. The carrier comes equipped with a Raspberry Pi microcontroller serving as the computing hub, along with an embedded webcam deliberately positioned for live image capture. The apparatus navigates field rows, snapping leaf images under changing perspectives and settings. These captures undergo on-the-spot analysis via cutting-edge deep learning routines for leaf sorting. This holistic strategy accentuates real-world relevance and readiness for field-level deployment.

Going beyond mere categorization, the project underwent trials directly in an agricultural environment, utilizing a robotic carrier [12]. It probes the performance of these models on CPUs and on hardware-scarce devices such as the Raspberry Pi [13-23], laying the groundwork for practical, in-field solutions [24, 25]. By blending top-tier deep learning algorithms with agronomic insights, this work aims to advance the rapidly evolving field of precision farming.

Mehdipour Ghazia *et al.* [3] evaluate three deep learning-driven object identification approaches—Faster R-CNN, SSD, and YOLOv3—operating on a Raspberry Pi to spot irregularities in melon foliage. The evaluation weighs detection precision, processing time, and resource usage to pinpoint the best algorithm for real-world deployment. Faster R-CNN emerged as the preferred candidate, credited to its elevated accuracy and mean average precision (mAP). SSD earned an endorsement for real-time spotting, striking a balance between speed and accuracy, whereas YOLOv3, despite broad acclaim, proved ill-suited for this task due to its low mAP. This wide-ranging assessment guides building robust deep learning-based object detection frameworks tailored to agricultural contexts on resource-frugal hardware such as the Raspberry Pi.

Tolgahan KARAHAN *et al.* [3] propose a scheme for categorizing grape disease images by combining a channel-wise attention (CA) module with streamlined CNNs. The team embedded Squeeze-and-Excitation (SE) blocks into ShuffleNet V1 and V2, lifting effectiveness while preserving a trim model footprint. Working from a collection of 4062 grape leaf images, the model achieved 99.14% accuracy, reducing memory requirements from 227.5 MB (AlexNet) to 4.2 MB. This lean model enables automated, human-free spotting of grape ailments, providing actionable guidance on pesticide deployment. Given its high accuracy, small size, and self-diagnostic capabilities, this technique has great potential for precision agriculture, equipping farmers to pinpoint and address diseases promptly using portable electronics.

Islam *et al.* [4] unveiled a web-driven crop disease forecasting platform built around the ResNet-50 model. The pipeline begins by gathering leaf snapshots from multiple plant types, such as potato, tomato, and pepper. Following preparatory adjustments, including rotation and scaling, the snapshots train a predictive engine. ResNet50, a 50-layer-deep CNN variant, outstrips VGG-16 and VGG-19 regarding both prediction accuracy and vanishing gradient handling. The

© 2023 The Author(s). | Published by Agro Intelligence Research Press

platform's purpose is to support cultivators with early-stage disease spotting, marking a meaningful leap forward in crop stewardship and agri-tech.

Leea *et al.* [10] conducted a linked investigation into how CNNs combined with transfer learning can be used to diagnose plant ailments. Swapping conventional convolutions for depthwise separable counterparts allowed the investigators to slash parameter counts and computational load. Pre-trained models like Inception, ResNet, and MobileNet were fine-tuned on the PlantVillage corpus, demonstrating transfer learning's ability to improve results on modest datasets. InceptionResNetV2 hit an accuracy mark of 99.11 %, harnessing complementary strengths from ResNet and Inception frameworks. This spotlights the capacity of CNNs to drive exact, hands-free plant disease detection.

Mahmudul Hassan *et al.* [11] focus on harnessing compact CNNs on a Raspberry Pi 4 to detect tomato leaf diseases. The group amplified a dataset of 18,160 tomato leaf images, partitioned into 10 classes. Through deep transfer learning, they trialed CNN designs including MobileNetV2, NasNetMobile, Xception, and MobileNetV3, carefully tuning operational parameters. A defining element was the Raspberry Pi, which empowered live disease identification. This effort demonstrates how slim CNN designs, combined with economical hardware, can serve agriculture through on-the-spot disease diagnosis.

Zhu *et al.* [12] this contribution delves into CNN-backed plant species recognition. The investigators trained a CNN to learn leaf characteristics directly from imagery, building a multilayered architecture. MATLAB handled pre-processing steps, such as noise stripping, and final categorization was performed via a Softmax classifier. The methodology outperformed earlier conventional techniques, demonstrating that deep learning approaches such as CNNs are superior for automating feature extraction and plant species sorting. The approach holds promise for real-time plant species identification, indicating substantial progress in botanical taxonomy.

Paryadi *et al.* [24] examine harnessing CNNs for flora identification, using convolutional and fully connected layers to extract features. A linear support vector machine (SVM) subsequently undertook the classification task. The inquiry established that CNN-centric methods outperform other methods in distinguishing various plant structures, including foliage, blooms, fruits, stalks, and limbs. These outcomes indicate that CNN-rooted schemes can eclipse older handcrafted-feature pipelines, marking a promising route for botanical sorting.

Zamir *et al.* [25] later examined a CNN model comprising convolutional, pooling, and fully connected layers for plant leaf categorization. A corpus of plant leaf photographs served as training material, supplemented with data augmentation procedures to broaden its richness. The performance yardsticks included F1-score, recall, accuracy, and precision. Although the model's accuracy was 86%, it faltered on alternative metrics such as precision, recall, and F1-score. Despite these uneven outcomes, the inquiry adds heft to the expanding literature affirming CNNs' capacity for plant leaf recognition and deep learning's broader role in agriculture [5, 25-29].

Bansal *et al.* [26] survey the application of transfer learning to plant disease spotting and benchmark assorted pretrained CNN models. DenseNet-121 achieved the highest F1 score (99.8%), demonstrating the effectiveness of transfer learning in this domain. ResNet-50, a 50-layer deep architecture, stood out for its skip-connection design, which helped counter vanishing gradients.

The collectives above have scrutinized diverse deep learning strategies, notably ResNet50 and DenseNet121, achieving substantial success in distinguishing crops beyond root vegetables. In our investigation, we emphasized subterranean crops such as beetroot, radish, potato, and sweet potato. We leveraged toolkits such as ResNet50 and DenseNet121, deploying them on hardware like the Raspberry Pi to enable live classification in farming conditions.

Materials and Methods

© 2023 The Author(s). | Published by Agro Intelligence Research Press

Figure 1 outlines the procedural workflow for training a deep learning system to discriminate among root crop varieties. The schematic encompasses every critical step, from initial data gathering and preparatory processing right through to model instruction, validation, evaluation, and final deployment. Adopting a disciplined, stepwise strategy, this project aims to develop a reliable apparatus for root crop identification, thereby advancing farming efficiency and ecological sustainability.

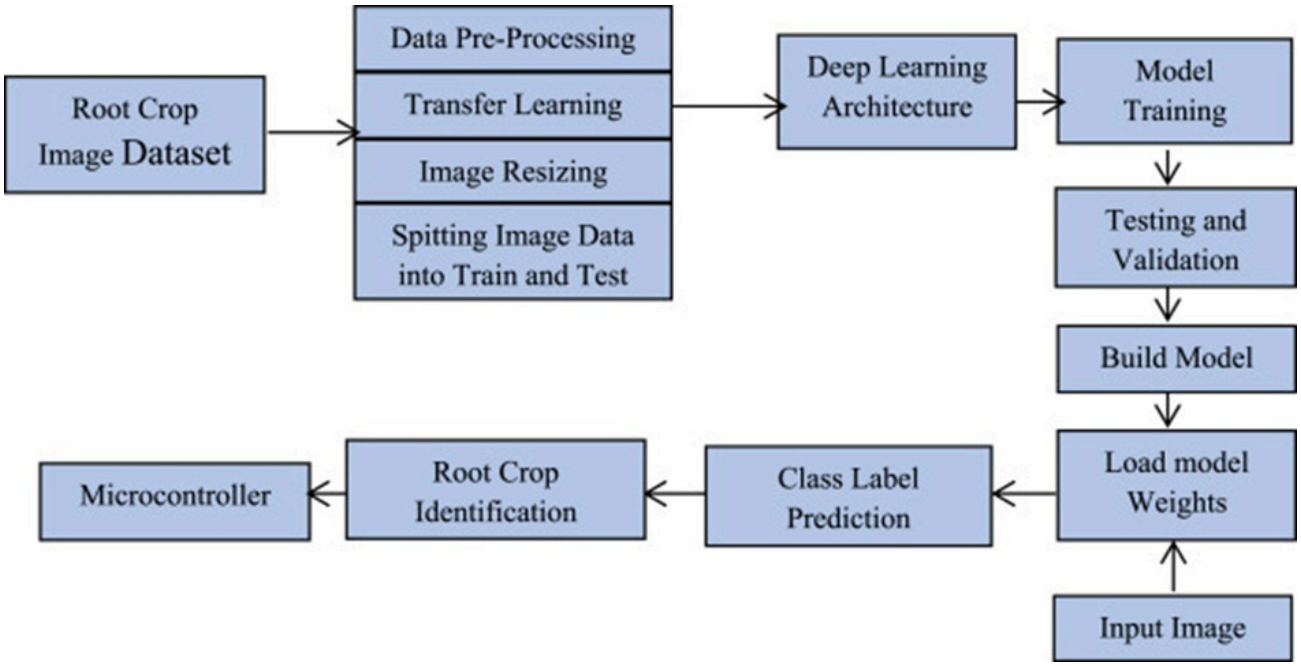

Figure 1. Proposed block diagram.

Data collection

Compiling the image corpus was far from straightforward, given the imperative to secure both breadth and authenticity in the sample pool. Undeterred by these hurdles, the present study drew upon an expansive repository of approximately 2500 naturally captured photographs—free of any artificial augmentation—harvested across multiple locales in Karnataka, India. Beetroot contributed roughly 770 frames, sourced from fields in and around Devarahalli, Gudlupet Taluk, within Chamarajanagar district, thus reflecting the local cropping profile. The radish portion, encompassing nearly 720 shots, originated from areas neighbouring Basavanahalli along Bannur Road in Mysuru district, shedding light on customary cultivation patterns there. Potato leaves furnished about 410 images, procured in Beluru, Hassan district, capturing the defining morphological traits of plants from that specific geography. Rounding out the assortment, sweet potato added close to 550 samples taken near HD Kote, Mysuru district. For training purposes, 70% of the entire collection was allocated; 20% for validation, and 10% for testing—a three-way split designed to underpin rigorous model appraisal. Such partitioning enables effective learning and trustworthy measurement of generalization performance.

Deep learning models

Among the most refined offshoots of machine learning sits deep learning, which has, over recent years, thoroughly reshaped the discipline of artificial intelligence. Deep learning rests upon the concept of artificial neural networks—elaborate, brain-inspired computational architectures that mirror the organisation and operation of biological cognition.

These neural constructs, assembled from stacks of interwoven computational units, possess the capacity—without being explicitly programmed—to discern subtle regularities and associations woven through massive datasets. The multiplicity

© 2023 The Author(s). | Published by Agro Intelligence Research Press

of concealed processing tiers inherent to these networks supplies the “deep” qualifier, empowering the system to autonomously distil layered, hierarchical representations directly from raw sensory input.

ResNet50

Falling under the broader ResNet umbrella, the ResNet50 instantiation leverages the ResNet50 skeleton and applies transfer learning by bootstrapping from weights trained on the ImageNet corpus. As shown in **Figure 2**, the 50 strata of the canonical generalised ResNet50 comprise convolutional operations, batch normalisation steps, and nonlinear activation functions, all partitioned into discrete residual blocks. Building on this backbone, a Global Average Pooling operation is introduced to collapse the spatial dimensionality. Thereafter, a cascade of dense layers—sized at 512, 256, 128, and 64 units—is interleaved with dropout regularisation intervals. The terminal layer invokes Softmax activation to perform the final category assignment. This hybrid arrangement allows the architecture to tailor itself to the designated classification objective while simultaneously leveraging the formidable representational-extraction capacity of ResNet50.

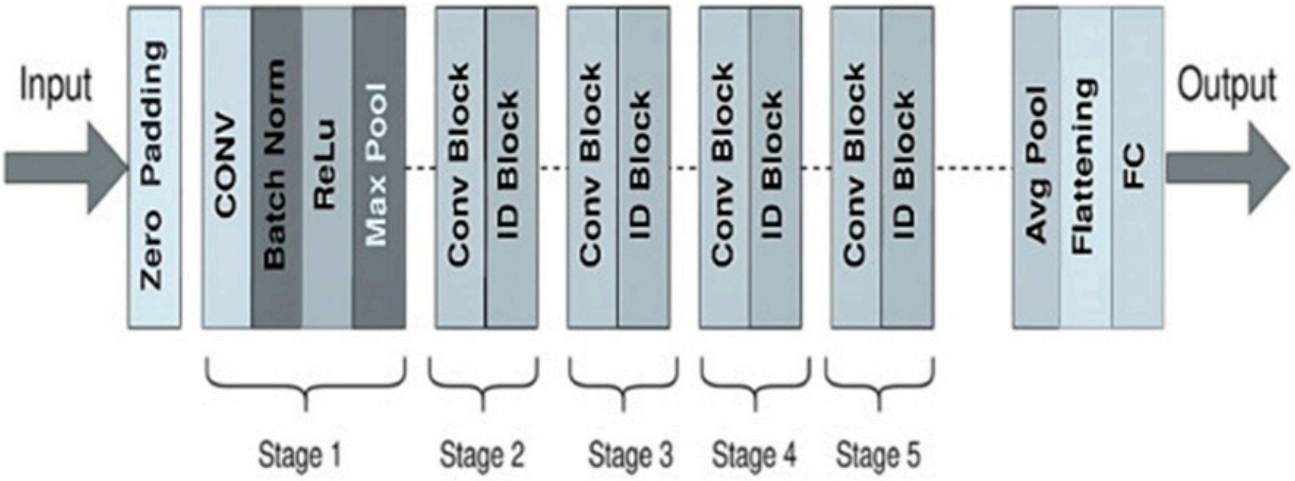

Figure 2. Generalized ResNet50 architecture.

The fifty layers populating ResNet50 break down roughly as follows:

- One opening convolutional layer
- One batch normalisation layer
- One max pooling layer
- Forty-five convolutional layers
- One final average pooling layer
- A single fully connected layer

Through their orchestrated arrangement, these components instantiate the hallmark residual learning paradigm characteristic of ResNet-type networks, counteracting gradient attenuation and thereby permitting the successful optimisation of remarkably deep topologies.

DenseNet121

© 2023 The Author(s). | Published by Agro Intelligence Research Press

Depicted in **Figure 3** is a highly performant convolutional neural network variant. The canonical DenseNet121 blueprint distinguishes itself through its dense connection scheme. Unlike conventional CNNs, DenseNet establishes direct feed-forward pathways between every pair of layers, an innovation that dampens vanishing gradient effects, amplifies feature propagation, promotes feature reuse, and dramatically reduces parameter overhead. In the specific embodiment realised here, the pipeline begins with an input receptor sized for 224×224-pixel RGB imagery. A pretrained DenseNet121 then acts as the feature extraction engine, mobilising its 121 layers' worth of dense blocks and transition modules to distil sophisticated patterns from the incoming photographs. Inside any given dense block, each convolutional unit inherits feature maps from all upstream layers and passes its own output forward to every downstream layer—a regime of total interconnectedness that maximises informational throughput. Following feature extraction, a GlobalAveragePooling2D operation is applied to the resultant 7 × 7 × 1024 activation maps, compressing the spatial extent while guarding against overfitting. The vectorised features subsequently traverse a chain of Dense modules, punctuated by Batch Normalisation and Dropout stages to enhance training robustness and curb over-specialisation. High-level abstractions are captured by a 256-unit Dense segment driven by ReLU activation, then subjected to a further round of Batch Normalisation plus Dropout. A dense concluding layer with Softmax activation, whose unit count equals the class count, finalises the model. This structural formulation proves admirably suited to a range of image recognition challenges, adeptly marrying supplemental refinement and regularisation layers with the representational head start conferred by a pretrained donor.

Layers	Output Size	DenseNet-121
Convolution	112 × 112	
Pooling	56 × 56	
Dense Block (1)	56 × 56	[1 × 1 conv, 3 × 3 conv] × 6
Transition Layer (1)	56 × 56	
	28 × 28	
Dense Block (2)	28 × 28	[1 × 1 conv, 3 × 3 conv] × 12
Transition Layer (2)	28 × 28	
	14 × 14	
Dense Block (3)	14 × 14	[1 × 1 conv, 3 × 3 conv] × 24
Transition Layer (3)	14 × 14	
	7 × 7	
Dense Block (4)	7 × 7	[1 × 1 conv, 3 × 3 conv] × 16
Classification Layer	1 × 1	

© 2023 The Author(s). | Published by Agro Intelligence Research Press

Figure 3. Generalized DenseNet121 architecture.

The constituents of DenseNet121 include:

· 1 initial convolution layer

· 1 max pooling layer

· 4 dense blocks, internally comprising a sum of 58 convolutional layers

· 2 convolutional layers (batch normalisation followed by convolution) per transition layer, across 3 transition layers

· 1 global average pooling layer

· 1 final classification layer

The above items tally to 66 layers (58 + 6 + 1 + 1). The remaining 55 layers, bringing the total to 121, consist of ReLU activations within the dense blocks, together with batch normalisation operations. These ancillary elements are indispensable for injecting nonlinearities and conditioning inputs, equipping the network to efficiently acquire complex patterns.

Pre-processing of the dataset

Preparing raw data is a foundational prerequisite for preparing a dataset for CNN-driven root crop discrimination. This preparatory stage imposes uniformity on the format and fidelity of incoming images, enabling model learning to proceed unhindered. Specific actions include:

· Resizing every sample to a fixed geometry, usually conforming to the dimensions expected by the chosen CNN variant, thereby guaranteeing consistent input dimensionality throughout the corpus. For the present work, the original heterogeneous set was pre-processed so that all images were uniformly scaled to 224×224 pixels, as evidenced in **Figure 4**.

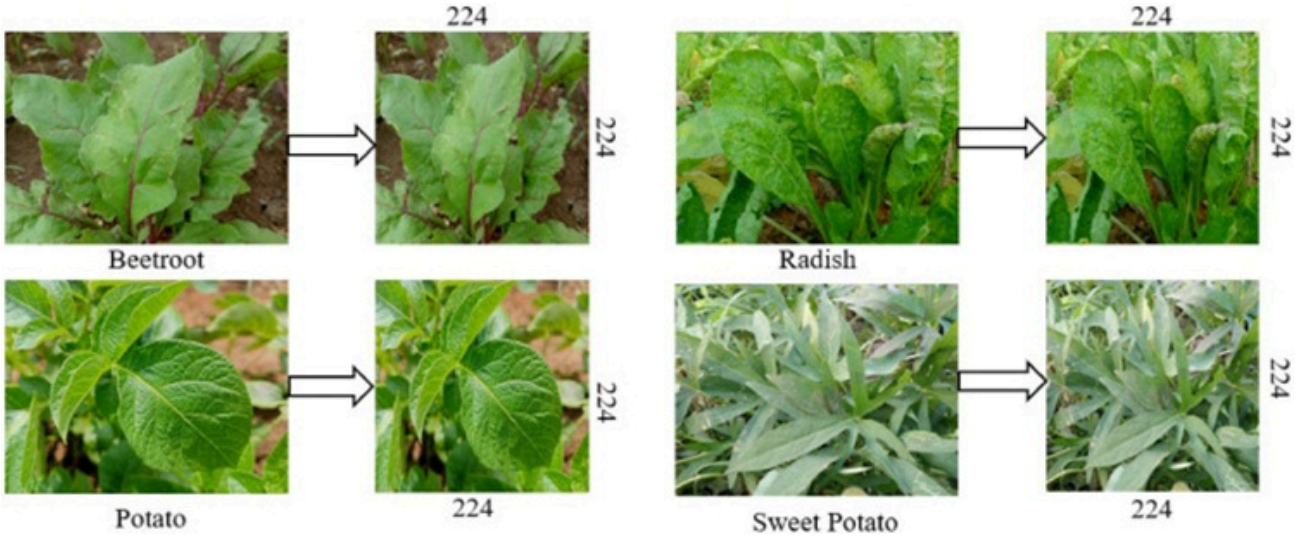

Figure 4. Image resizing.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

· Adjusting pixel intensities onto a standard interval, most commonly [0, 1], to temper the disrupting influence of variable illumination during model training.

$$\text{Normalized pixel}\,[i,j] = \frac{\text{original_pixel}[i,j]}{255} \quad (1)$$

Here, normalized_pixel[i,j] signifies the intensity value at position (i,j) after transformation; original_pixel[i,j] represents the raw intensity at that same coordinate; and i and j index the row and column locations, respectively (each spanning from 0 up to 223 for a 224×224 frame).

The operation rescales intensities from the native 0–255 integer domain to a continuous 0–1 range. For a host of machine learning algorithms, neural networks in particular, this standardisation proves pivotal, as it homogenises the input distribution and frequently conduces to steadier, swifter convergence. The choice of 255 as the denominator reflects the convention that standard digital imagery allocates 8 bits per colour channel, yielding 2^8 = 256 distinguishable gradations, numbered 0 through 255.

Collectively, these pre-processing measures enhance the CNN's ability to capture class-discriminative features essential for precise crop identification while mitigating the confounding effects of raw input heterogeneity.

Data augmentation

A vital tactic for boosting the variability and adaptability of a dataset is data augmentation. By applying diverse manipulations to the source photographs—such as rotation, reflection, resizing, repositioning, and injecting random disturbances—synthetic training instances are generated. In effect, this expands the total volume of available data. Familiarising the model with a broader range of real-world perturbations equips it to better cope with fluctuations in brightness, camera angles, and spatial orientations.

As a regularisation technique, data augmentation curtails overfitting by deliberately adding measured variation to the training examples. When enriched samples are incorporated into the learning procedure, CNNs become more adept at learning salient representations and sharpen their precision in recognising root crop types.

Several established procedures are enlisted to enhance dataset variety:

· Image rotation: Entails revolving a picture by a set angular displacement. This teaches models to detect items irrespective of rotational stance.

· Flipping: Generating mirror counterparts along the horizontal or vertical axes. This infuses diversity in perceived object orientation.

· Cropping: Isolating selected segments of photographs. The technique can emulate close-up views or concentrate attention on specific details.

· Translation: Sliding images sideways or up and down. Models thereby learn to pinpoint objects independent of their placement within the visual field.

As shown in **Figure 5**, these modifications collectively enhance the model's generalisation and improve its output quality across a broader range of inputs.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

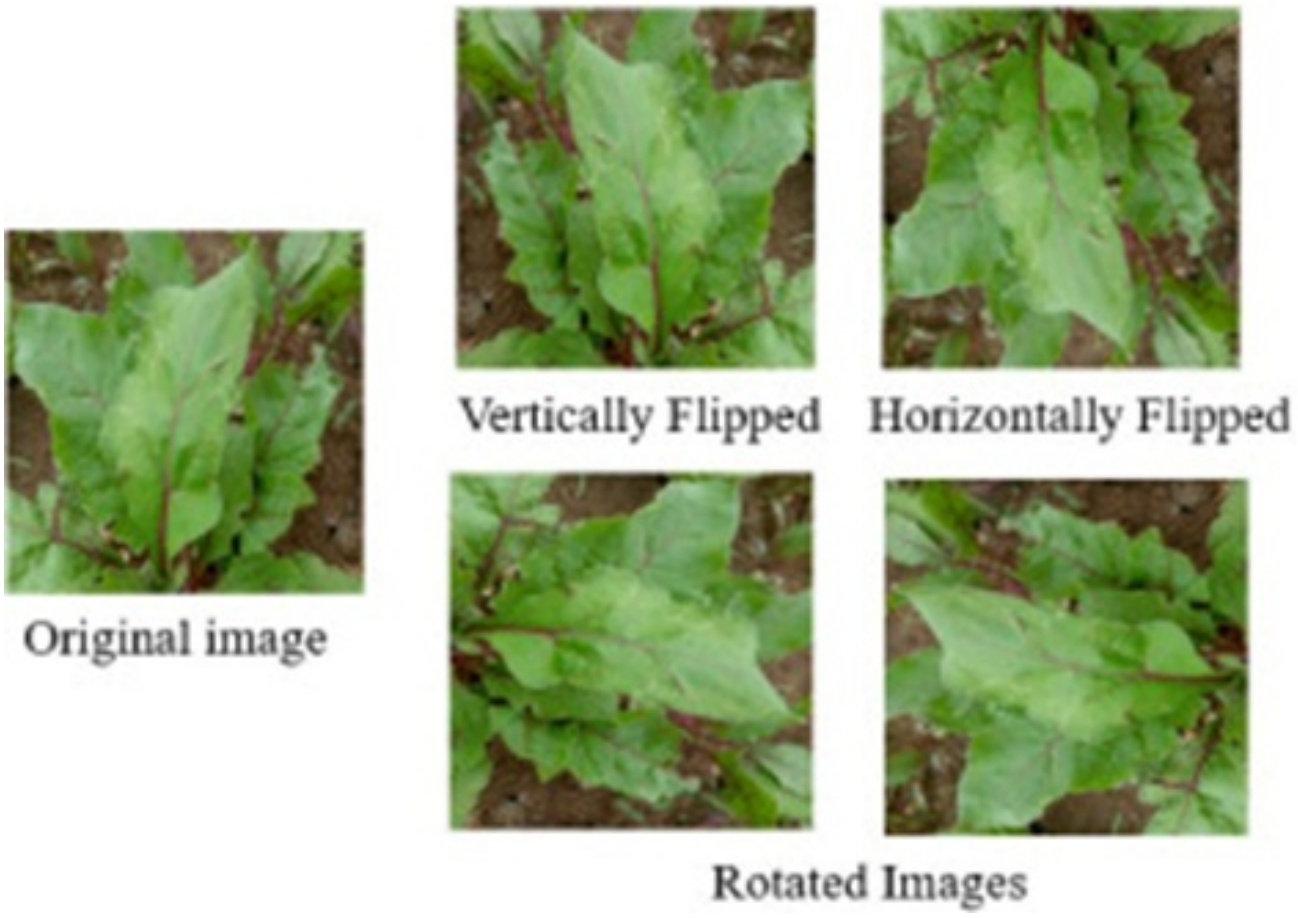

Figure 5. Data augmentation.

Training and validation

The assembled data is split into three portions: training, validation, and testing. Seventy percent goes toward training, twenty percent toward validation, and ten percent toward testing. This division ensures that most of the data drives the learning process, while distinct subsets are reserved for interim validation and final appraisal.

Training a CNN for root crop differentiation involves reducing a chosen loss function computed over the training partition by iteratively tuning model parameters. Over this stage, gradient descent combined with backpropagation steadily updates the network's weights and biases. By progressively extracting layered, hierarchical descriptors from the input images, the CNN eventually learns to discriminate among crop varieties. Validation entails evaluating the model on an independent subset that it has never encountered during training. The twin objectives of validation are to measure how well the model generalises beyond its training experience and to detect emerging issues such as overfitting. Performance indicators—accuracy, precision, recall, and F1 score—are derived by matching the model's outputs against the reference labels within the validation fold. These gauges reveal the system's proficiency at separating crop classes. Techniques such as cross-validation, where the dataset is split into multiple segments and validation is performed in successive rounds, can further strengthen the trustworthiness of validation findings.

Model evaluation

Testing refers to assessing the fully trained model's performance on the test split, an entirely separate dataset. To obtain an impartial measure of generalisation, this portion stays entirely hidden from the model during both training and validation.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

The confusion matrix is a vital tool for evaluating deep learning model performance, especially in classification frameworks [1]. It delivers a sweeping account of how dependably a model forecasts categorical results across a designated test set. By placing the model's predictions alongside the factual results, this matrix provides a clear visual summary of the model's correctness.

Fundamentally, the confusion matrix organises the model's outputs into four primary groupings:

· True positive (TP): Occurrences where the model correctly anticipated a positive outcome that coincided with the actual positive condition.

· True negative (TN): Cases in which the model rightly forecast a negative outcome matching the genuine negative state.

· False positive (FP): Designated a Type I error, this class includes instances where the model erroneously signalled a positive outcome when the truth was negative.

· False negative (FN): Also called a Type II error, this category encompasses situations in which the model incorrectly predicts a negative result despite the actual condition being positive.

a. Accuracy: **Eq. 2** formulates model accuracy as the ratio of total correct predictions to the overall number of predictions made. It offers a general indication of how frequently the model delivers correct judgments across all classes.

$$Accuracy = \frac{TP+TN}{TP+TN+FP+FN} \quad (2)$$

b. Precision: Precision reflects the share of the model's positive predictions that prove to be accurate. As laid out in **Eq. 3**, precision is obtained by dividing the count of true positive calls by the total of all positive predictions, both incorrect and correct alike.

$$Precision\,(p) = \frac{TP}{TP+FP} \quad (3)$$

c. Recall: Recall captures a model's ability to retrieve every relevant instance within a dataset. It measures how exhaustively the model identifies all positive examples and is defined as the ratio of true positives to the sum of true positives and false negatives **(Eq. 4)**.

$$Recall\,(r) = \frac{TP}{TP+FN} \quad (4)$$

d. F1 Score: **Eq. 5**, which defines the F1 Score, provides a unified metric that captures the model's overall effectiveness by balancing recall and precision. As the harmonic mean of recall and precision, it yields a more holistic appraisal than accuracy alone, particularly in datasets with uneven class distributions.

$$F1\ Score = \frac{2*Precision*Recall}{Precision+Recall} \quad (5)$$

© 2023 The Author(s). | Published by Agro Intelligence Research Press

To assess the model's accuracy, precision, recall, F1 score, and other pertinent metrics, the model's predictions are evaluated against ground-truth annotations in the test partition. These benchmarks reveal both the system's overall strength in root crop classification and its ability to handle novel, previously unseen instances.

Real-time testing

Live evaluation of the model was conducted using a Raspberry Pi 4B microprocessor fitted with 4GB of RAM. This unit was mounted to a vehicle, along with a Logitech C270 Digital HD 5MP 720p Webcam, to capture root crop images in real time. The Raspberry Pi 4B is a robust single-board computer suited to numerous edge computing undertakings. Powered by a quad-core ARM Cortex-A72 processor running at 1.5 GHz and backed by 4GB of LPDDR4 RAM, it delivers the computational power needed to host and execute machine learning models for instantaneous image recognition. The visual interface was relayed to the Raspberry Pi via Virtual Network Computing (VNC), which served as the means to control and observe the Raspberry Pi deployment.

Results and Discussion

The trajectory of model learning was followed via learning curves, which supply key insights into performance and generalisation behaviour. These plots graph training and validation accuracies as functions of epoch number.

As the epoch count increased, we observed training accuracy rise progressively, a sign that the model was becoming increasingly proficient at learning patterns embedded in the training corpus. Concurrently, validation accuracy was monitored, functioning as an indicator of how well the model extends its learning to new observations.

The interplay between these two trajectories yields meaningful intelligence about the model's learning dynamics. A progressively expanding rift where training accuracy climbs high while validation accuracy lags considerably lower may signal overfitting, whereas persistently depressed values on both curves could suggest underfitting. The desirable end state sees both traces meeting at an elevated accuracy plateau, thereby confirming that the model has extracted knowledge effectively from the training data while preserving strong generalisation capabilities.

ResNet50 model analysis

After the training cycle concluded, we subjected the ResNet50 model to a stringent examination on an independent test split to assess its real-world efficacy and extrapolation performance. The model returned remarkable fidelity, flawlessly recognising every crop sample within the test partition. This unblemished outcome on data never previously encountered further affirms the durability and potency of our customised ResNet50 architecture for this targeted crop differentiation challenge.

Figure 6 shows a selected set of test images along with the model's corresponding predictions. These concrete visual demonstrations offer palpable validation of the model's discriminative competence across diverse exemplars of beetroot, potato, radish, and sweet potato. The figure depicts the model's consistent success in pinpointing each crop variety.

© 2023 The Author(s). | Published by Agro Intelligence Research Press



Figure 6. Classified root crop image by ResNet50 Model.

The uniform correctness observed across all four categories in the test set mirrors the elevated metrics observed during training and validation. This suggests the model has not merely internalised the capacity to pick out the defining traits of each crop but has also grown exceptionally skilled at projecting those traits onto previously unobserved specimens.

Figure 7 plots the Accuracy and Loss profiles logged while the ResNet50 model was being trained and validated. The Accuracy trace shows that both validation and training accuracies rose steadily as epochs progressed, indicating uninterrupted assimilation and sound generalisation. The Loss trace reveals that as training progressed, validation and training losses shrank in lockstep, with no appreciable gap opening between them. This close correspondence indicates the model largely evaded significant overfitting.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

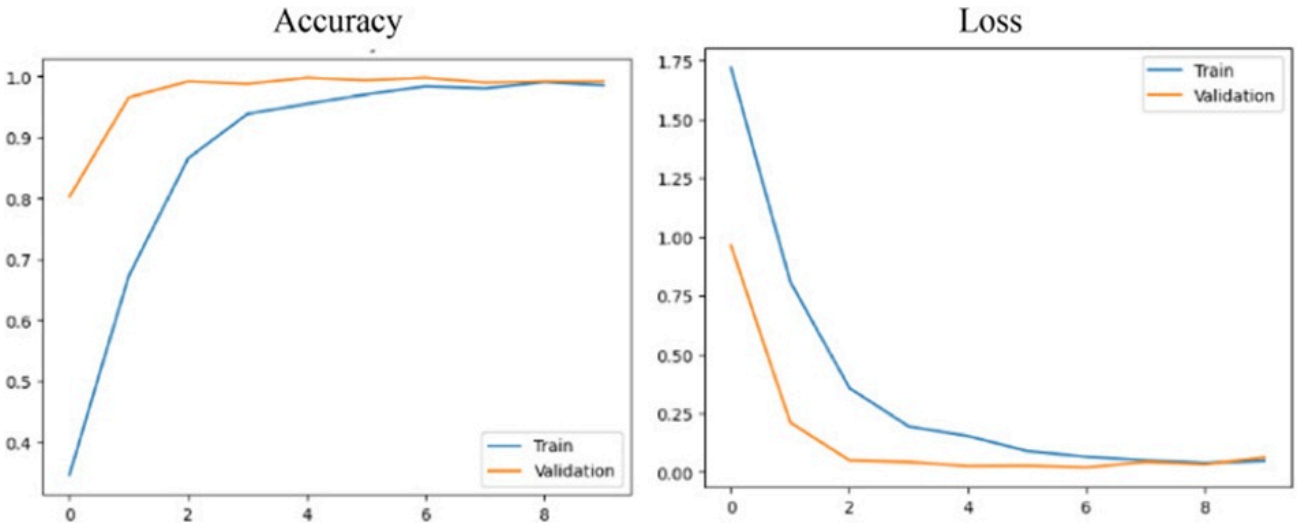

Figure 7. Accuracy and loss curve of the ResNet50 model.

The ResNet50 model posted striking performance statistics. It achieved a training accuracy of 99.73% and a corresponding loss of 2.75%. Crucially, test accuracy equalled 99.60 %, closely shadowing the training figure and signalling superb transfer to unseen data. The near-parity between training and test outcomes implies that the model has learned reliable, transportable feature representations from the training corpus.

A granular breakdown of our ResNet50 model's behaviour is conveyed through the confusion matrix displayed in **Figure 8**, which pertains to the model taught to discriminate among four root vegetable groupings: beetroot, potato, radish, and sweet potato. Analysing this matrix provides an exhaustive portrait of the model's predictions and enables a close inspection of its behaviour across every class.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

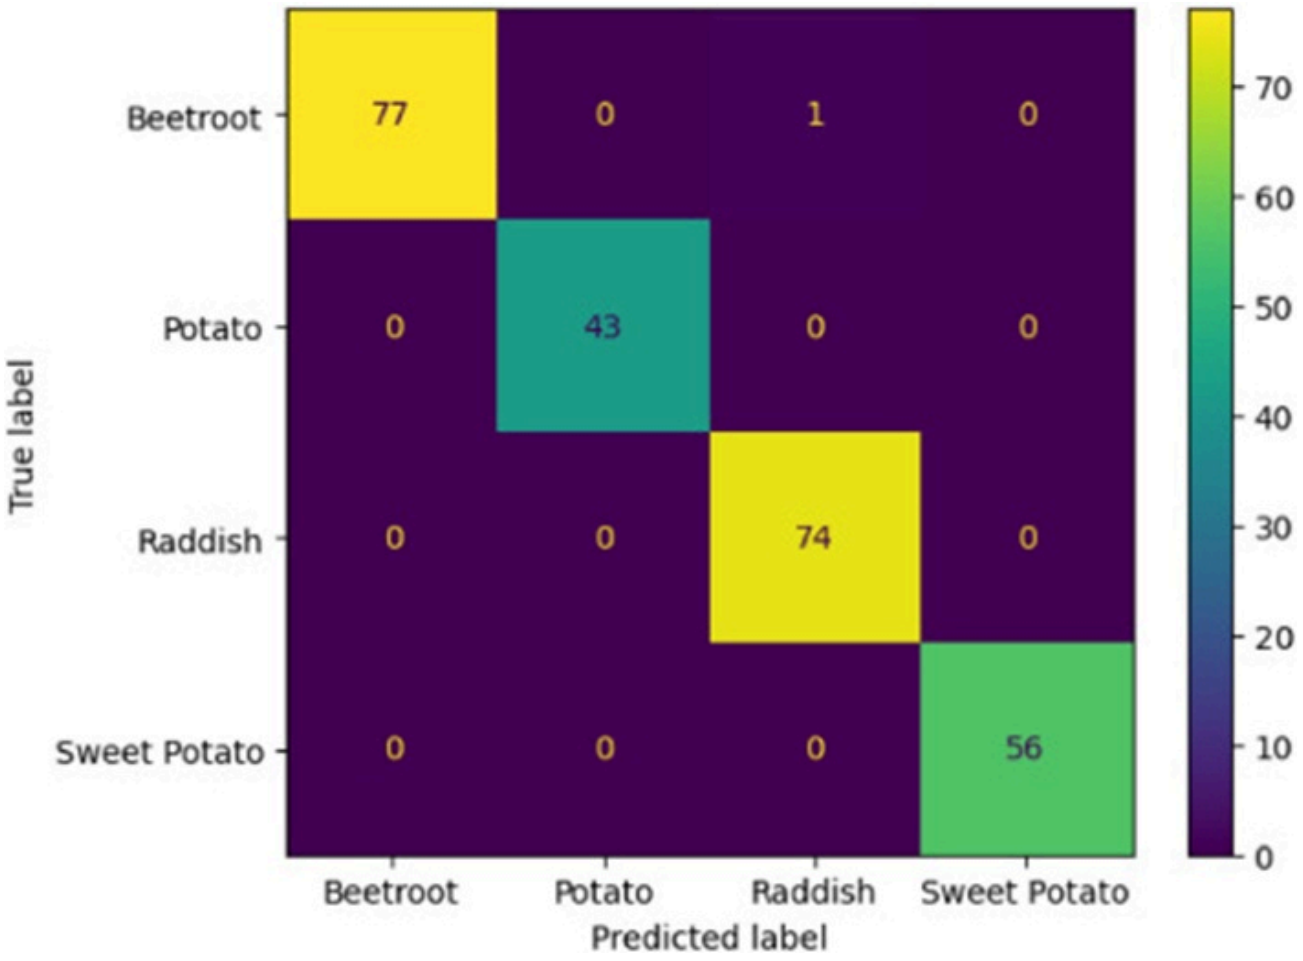

Figure 8. Confusion matrix of the ResNet50 model.

Table 1 consolidates the principal performance metrics broken down by class.

Table 1. Analysis of the confusion matrix of the ResNet50 model.

Root crop type	F1 score (%)	Recall (%)	Precision (%)	Accuracy (%)
Beetroot leaf	99.35	100	98.71	99.60
Potato leaf	100	100	100	100
Radish leaf	99.32	98.66	100	99.60
Sweet potato leaf	100	100	100	100

Figure 9 provides a side-by-side summary of overall performance for both the ResNet50 and DenseNet121 models.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

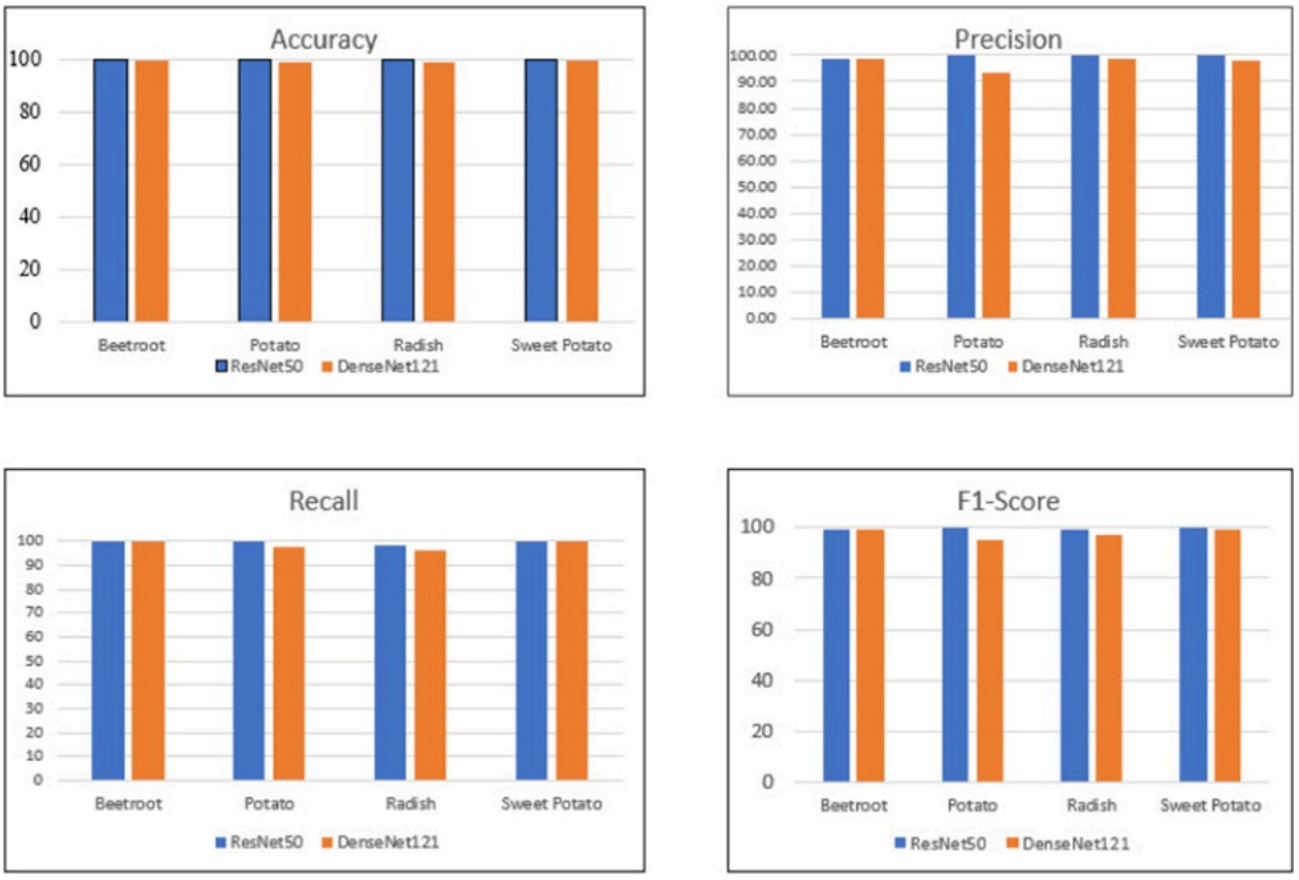

Figure 9. Accuracy, precision, recall, and F1 score of ResNet50 and DensNet121 models.

To further substantiate the model's suitability for practical deployment, we carried out live prediction trials by pairing a camera with the CPU. This configuration mimics field conditions where rapid, on-the-fly categorisation is necessary. The ResNet50 model preserved its elite performance in this fluid setting, reliably determining root crop identities directly from the live camera stream.

Figure 10 presents a real-time predicted root crop leaf image. This snapshot captures the model's output as it processes the live video feed, underscoring its ability to accurately distinguish between beetroot, potato, radish, and sweet potato foliage under live conditions. These instantaneous predictions exhibit a degree of accuracy that closely tracks the model's results on both the validation and test corpora.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

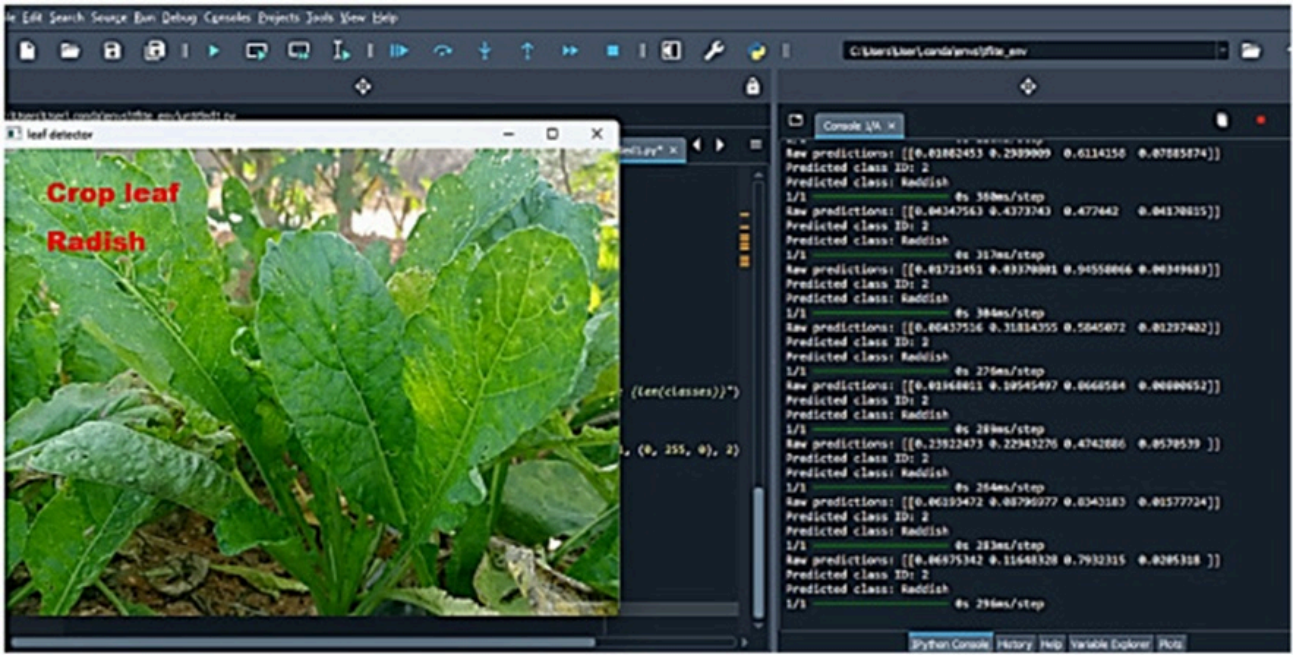

Figure 10. Real-Time root crop classification performance on CPU.

DensNet121 model analysis

Using an independent test collection, we appraised the practical performance and generalisation reach of our DenseNet121 model. **Figure 11** presents a curated set of predicted images sourced from this test pool, showcasing the model's classificatory outcomes across the four root crop leaf groupings: beetroot, potato, radish, and sweet potato.

© 2023 The Author(s). | Published by Agro Intelligence Research Press



Figure 11. Classified image of the root crop by DenseNet121 model.

These illustrative samples provide concrete proof of the model's sorting capability, demonstrating its ability to accurately distinguish the various root crop leaf types. The consistent performance across multiple samples indicates that the model has effectively acquired the knack for recognising the hallmark features of each crop variety.

The Accuracy and Loss trajectories for the DenseNet121 model during training and validation are shown in **Figure 12**. These plots provide meaningful insight into the model's learning behaviour and generalisation capacity.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

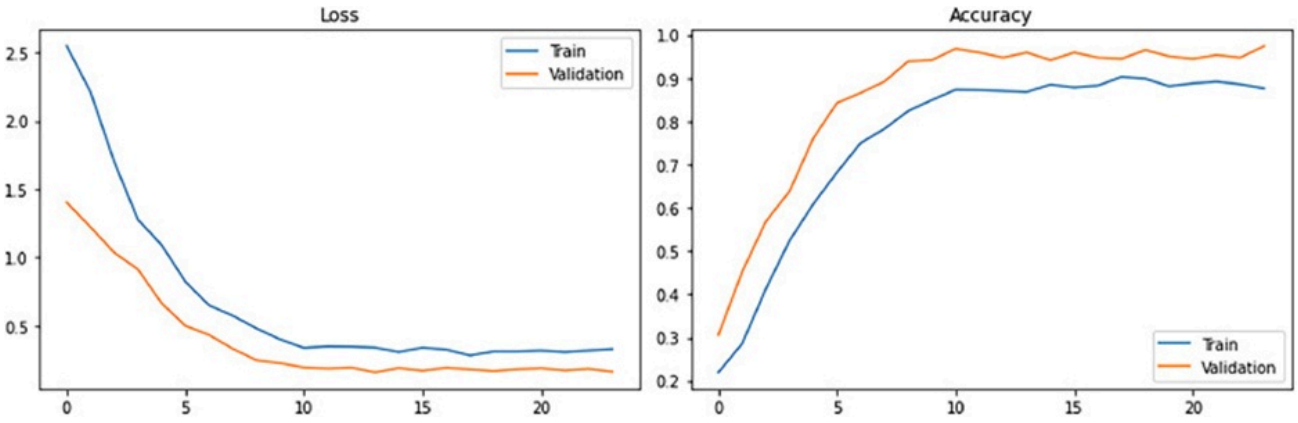

Figure 12. Accuracy and loss curve (Training and Validation) DenseNet121.

Uninterrupted learning throughout the training process is conveyed by the accuracy curve, which shows a continuous uptick in both training and validation accuracies across epochs. The Loss curve reflects this pattern with a parallel downtrend in both training and validation losses, lending further weight to the model's improving proficiency. Notably, the validation accuracy (97.60 %) surpasses the training accuracy (96.56%), a tell-tale indicator of healthy generalisation. This implies the model has not only absorbed the training material effectively but has likewise gained the capability to operate competently on novel inputs.

Figure 13 shows the confusion matrix for our DenseNet121 model, trained to separate four root crop foliage types: beetroot, potato, radish, and sweet potato. This matrix yields a thorough dissection of the model's outputs, allowing visualisation of its behaviour across all classes and facilitating detection of any recurring tendencies—correct or incorrect.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

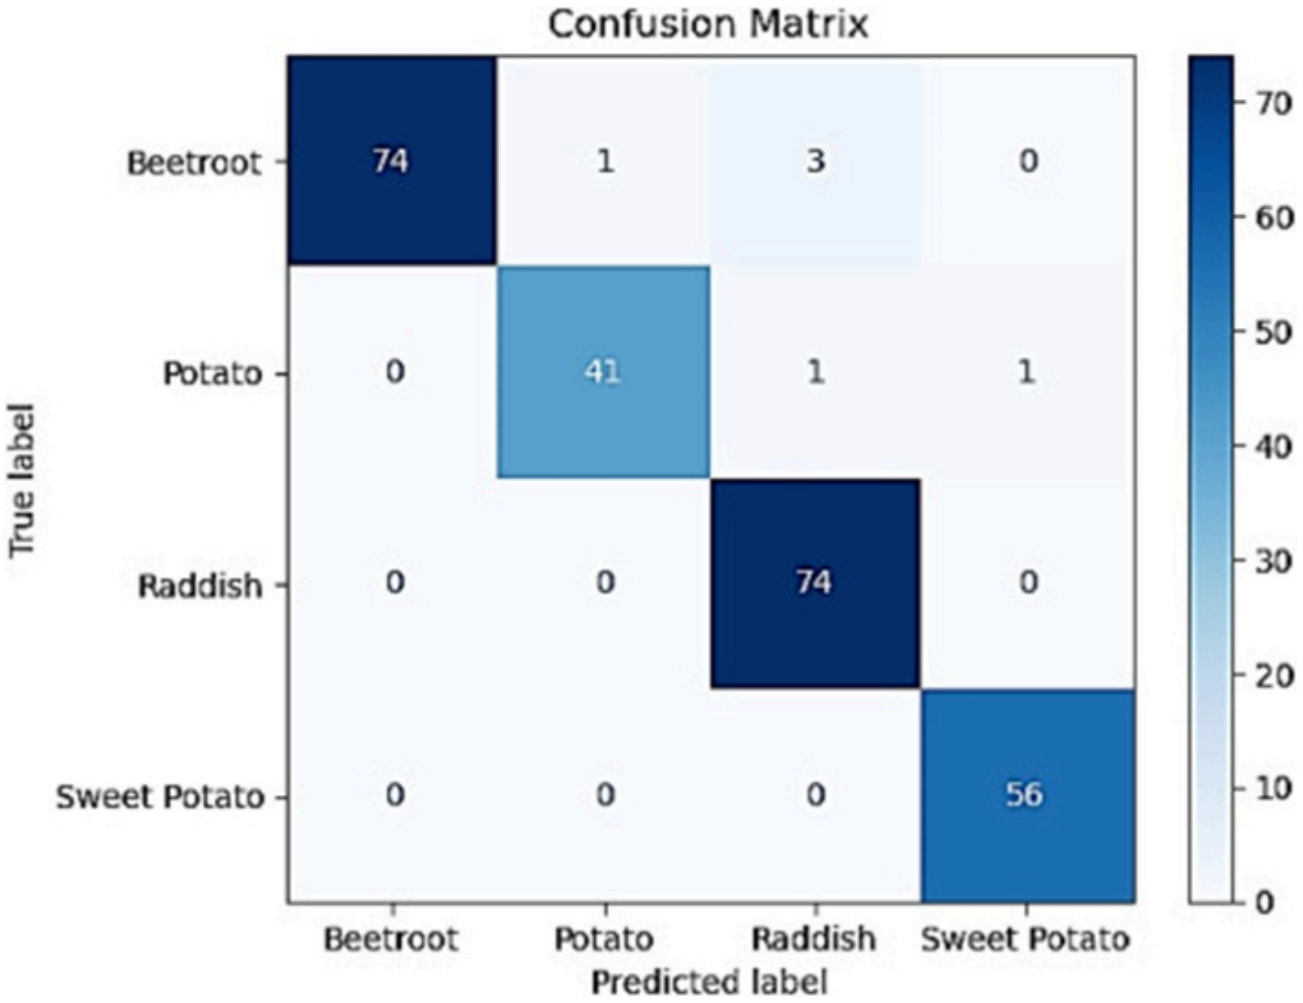

Figure 13. Confusion matrix of the DenseNet121 model.

Table 2 complements the confusion matrix by offering an exhaustive assessment of the DenseNet121 model's performance. This table presumably catalogues essential performance metrics—recall, precision, and F1-score—for each class.

Table 2. Analysis of the confusion matrix for the DenseNet121 model.

Crop Type	F1 Score (%)	Recall (%)	Precision (%)	Accuracy (%)
Beetroot Leaf	99.34	100	98.71	99.60
Potato Leaf	95.23	97.56	93.02	98.41
Radish Leaf	97.33	96.05	98.64	98.40
Sweet Potato Leaf	99.12	100	98.25	99.60

As shown in Table 2, the DenseNet121 model achieves admirable results across all four classes. Uniformly elevated precision, recall, and F1-scores attest that the model is not simply broadly accurate but also performs strongly across the full spectrum of crop types.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

To investigate the on-the-ground usefulness of our DenseNet121 model under dynamic, real-world conditions, we conducted real-time testing using a CPU-based setup. This evaluative step was designed to gauge the model's capacity to correctly and rapidly sort root crop leaves in a live setting, mimicking plausible real-world use.

Figure 14 shows samples of live predictions generated by the DenseNet121 model running on a CPU. These captures show the model's output as it processes a live stream, highlighting its accuracy in distinguishing beetroot, potato, radish, and sweet potato foliage. The findings confirm that the model successfully upheld its elevated classification accuracy across all four groupings during this live exercise.

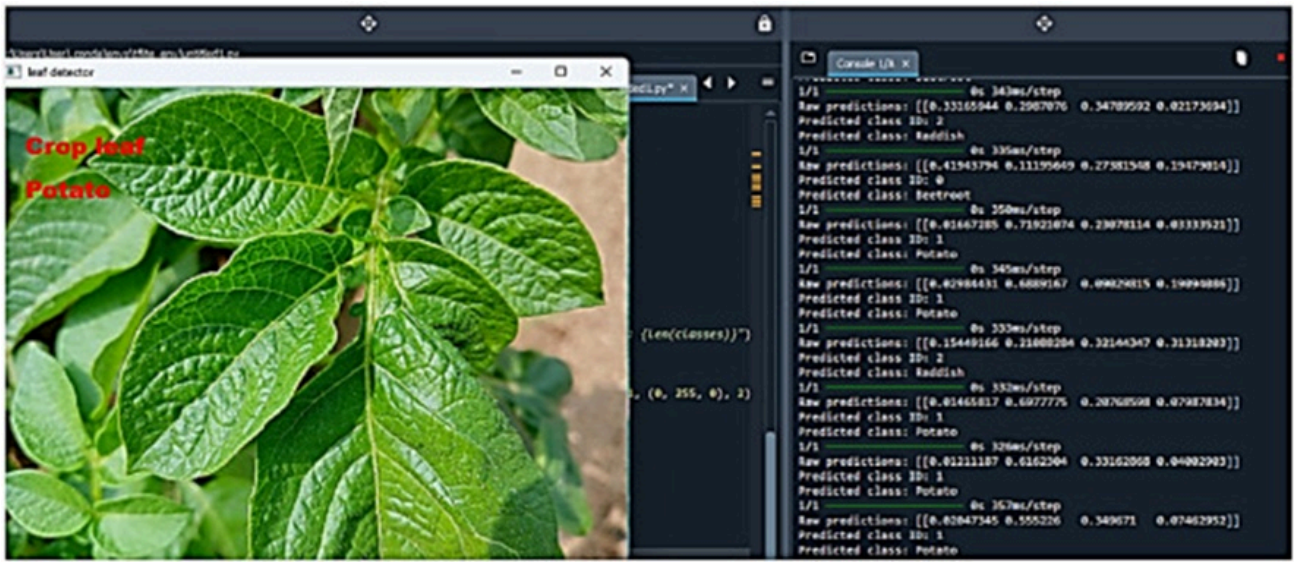

Figure 14. Real-time classification of root crop images by DenseNet121 Model on CPU.

The DenseNet121 model's ability to accurately label all root crop classes in real time while executing on a CPU arrangement is especially notable. This calibre of performance brings to light several pivotal strengths of our deployed model:

· Robustness: The model maintains high accuracy despite inherent variability in live capture, such as fluctuations in brightness, camera perspective, and background clutter.

· Efficiency: The model processes images at a sufficient speed to handle live input on a CPU, a prerequisite for real-time applications.

· Consistency: The shift from static test photographs to a dynamic live feed did not erode model performance, indicating strong generalisation.

· Versatility: The model's compelling real-time performance on a CPU indicates its suitability for deployment across a wide range of computing environments.

This fruitful live deployment significantly expands the potential applications of our DenseNet121 model. It validates the model's readiness for embedding within concrete agricultural technology platforms.

Model deployment on Raspberry Pi

The Raspberry Pi configuration, as visualised in **Figure 15**, uses a Raspberry Pi 4, the latest generation of this ubiquitous single-board computer. A Logitech C270 Digital HD 5MP 720p Webcam is directly attached to the Raspberry Pi 4B to

© 2023 The Author(s). | Published by Agro Intelligence Research Press

enable immediate image acquisition for leaf classification. The display was bridged to the Raspberry Pi through Virtual Network Computing (VNC), and the Raspberry Pi setup was steered and monitored wirelessly via this display. To optimise throughput on this hardware-constrained platform, the ResNet50 and DenseNet121 models were both converted to TensorFlow Lite, a format tailored for efficient inference on edge devices.

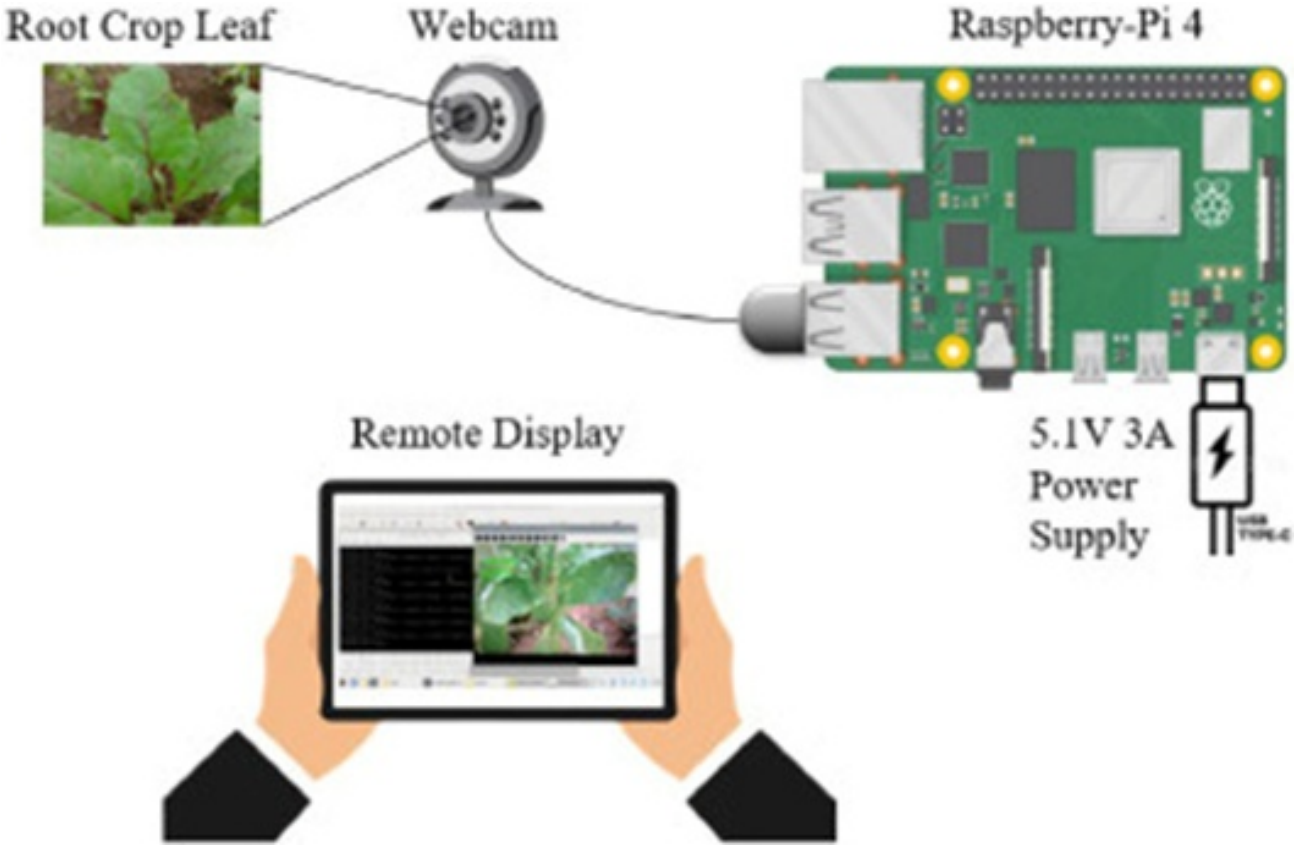

Figure 15. Raspberry Pi setup connection.

Both the ResNet50 and DenseNet121 models ran without issue on the Raspberry Pi, delivering performance practically indistinguishable from that observed on the considerably more muscular CPU. This noteworthy result spotlights the effectiveness of the TensorFlow Lite conversion pipeline in preserving model fidelity while dramatically cutting computational overhead.

Following the productive training, validation, and real-time CPU-centric testing of our ResNet50 and DenseNet121 models, we migrated them to a Raspberry Pi, a resource-constrained device well-suited for edge computing. This shift required translating the models from their original .keras and .h5 file representations into the slimmer .tflite (TensorFlow Lite) format, thereby tuning them for hardware with modest computational capacity.

The TensorFlow Lite conversion preserved the models' core capabilities while compressing their footprint and computational load. This measure is vital for eventual field rollouts in agricultural settings where access to high-powered computing resources cannot be assumed.

Figure 16 exhibits sample live predictions of root crop foliage generated by the ResNet50 model: **Figure 16(a)** displays the output image of the crop leaf, and **Figure 16(b)** presents the Raspberry Pi terminal, documenting the model's 99.60% accuracy in real-time discrimination of sweet potato, radish, beetroot, and potato leaves.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

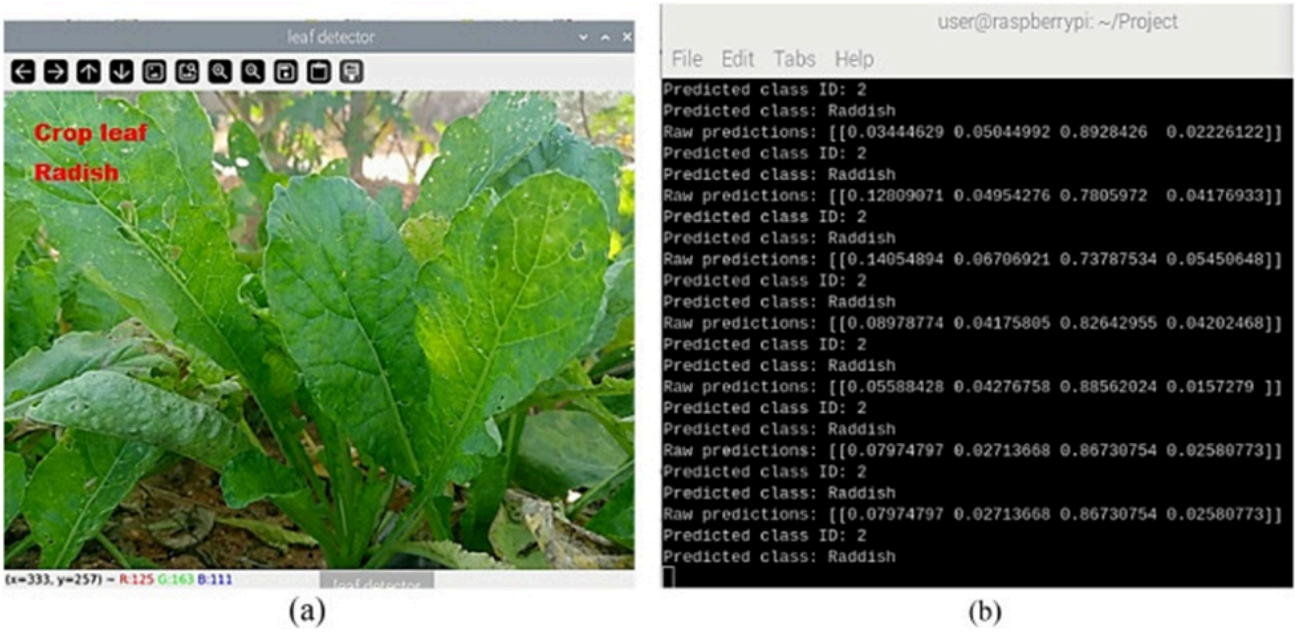

Figure 16. Real-time classification of root crop by ResNet50 on Raspberry Pi (a) Output image (b) Raspberry Pi terminal.

The DenseNet121 model achieved an impressive 96.81 accuracy in correctly identifying all root crop leaf types. **Figure 17** delivers visual proof of the model's successful porting and its exact classification capacity.

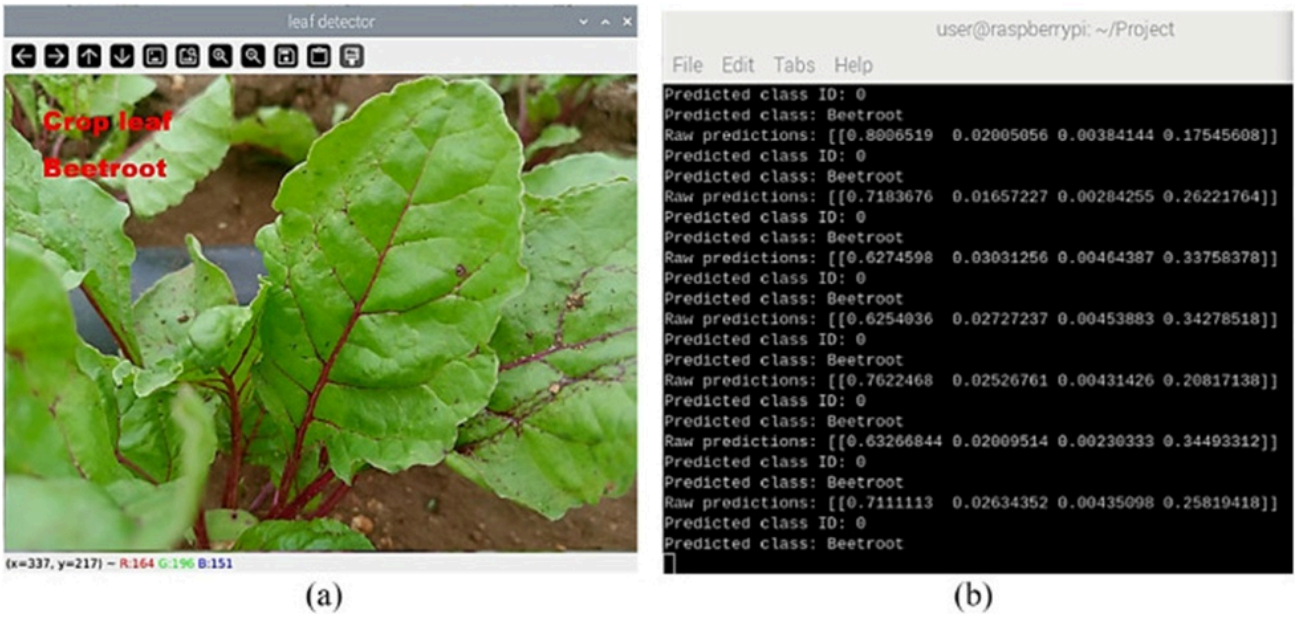

Figure 17. Real-time classification of root crop by DenseNet121 on Raspberry Pi (a) output image (b) Raspberry Pi terminal.

A field-level rollout using Raspberry Pi devices demonstrated how our models could operate in a real farming environment. **Figure 18** depicts the Raspberry Pi hardware arrangement placed on actual agricultural ground, demonstrating its role in live root crop leaf categorisation. **Figures 18a and 18b** show the model operating inside a radish plot, correctly labelling the crop as radish. In contrast, **Figures 18c and 18d** capture instances of the model in a sweet potato plot and accurately identify the crop as sweet potato.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

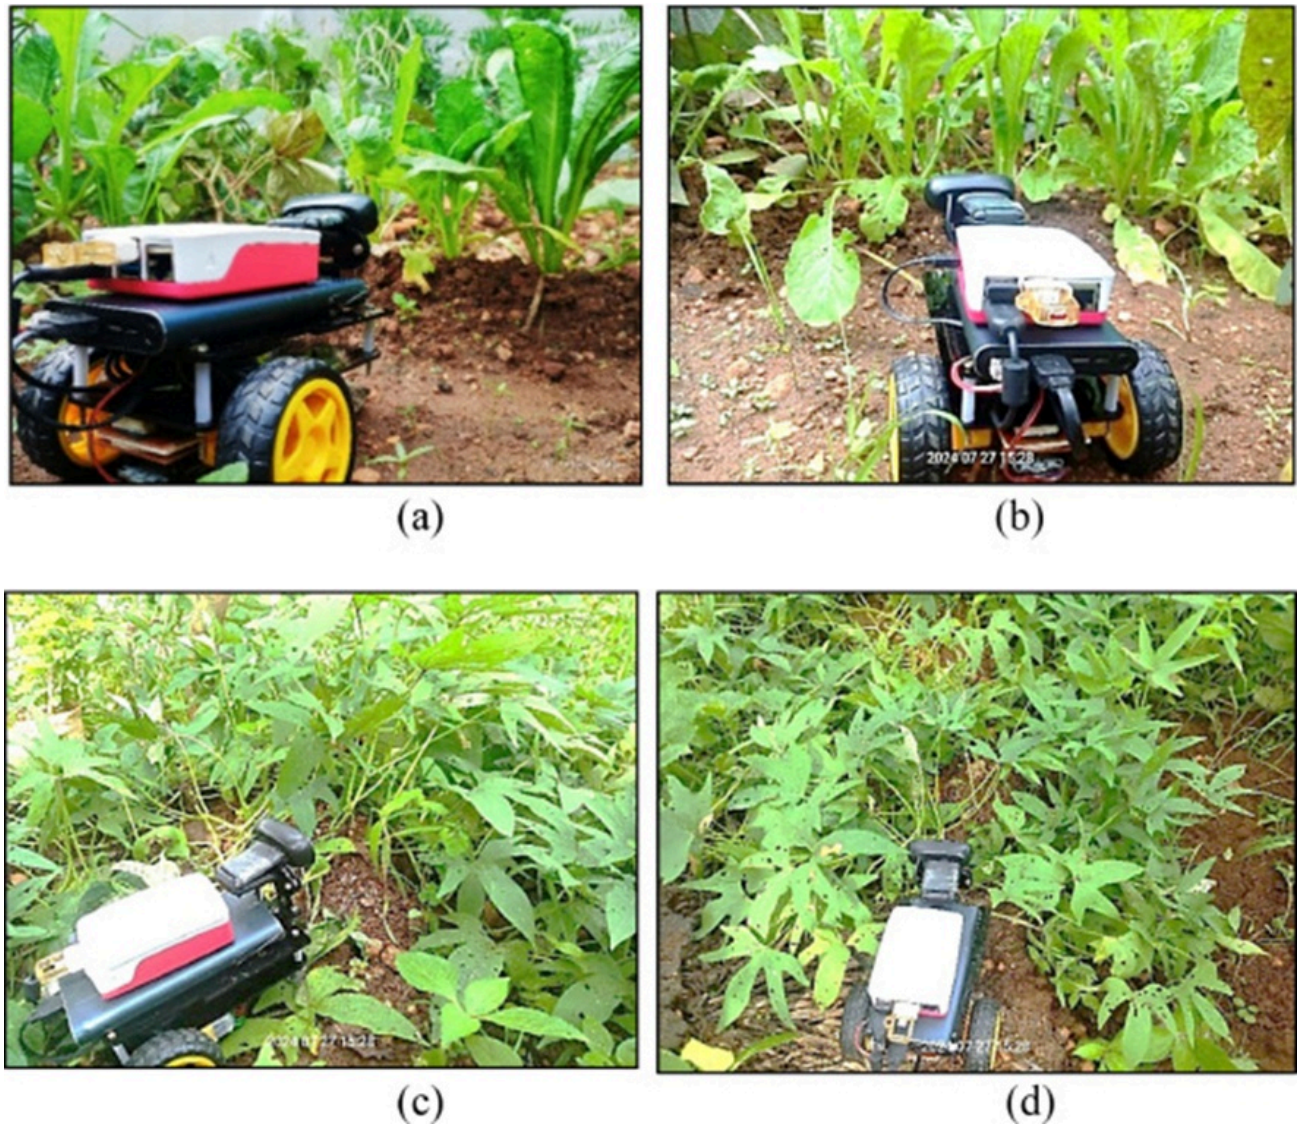

Figure 18. Raspberry Pi mounted on a vehicle to classify the root crop in agricultural land. (a) and (b) Radish classification in the agriculture field, (c) and (d) Sweet potato classification in the agriculture field.

This field-based trial marks a decisive stage in verifying the models' efficacy under authentic conditions. The Raspberry Pi 4, fitted with an attached webcam, was placed directly in the cultivated field to capture and process images of root crop leaves on-site. This arrangement allows instant, on-site leaf identification, delivering actionable real-time intelligence to growers and researchers.

The successful execution on Raspberry Pi brings several salient features of our models into relief:

· Portability: Both models can be ported and run effectively on resource-scarce devices.

· Efficiency: After format conversion, the models remain capable of live inference even on less potent processing hardware.

· Versatility/Scalability: Model behaviour remains consistent across dissimilar computing platforms, ranging from high-performance CPUs to compact single-board computers such as the Raspberry Pi.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

· Real-time processing: Growers receive on-the-spot leaf classification results, empowering swift decision-making.

· Cost-effectiveness: The use of affordable Raspberry Pi hardware makes this advanced technology accessible to a broad user base.

This deployment substantially extends the potential scope of our research, opening avenues for the construction of portable, low-cost, and energy-frugal crop identification systems operable directly at the field level. It heightens the concrete impact of our contributions to agricultural technology by furnishing a solution that can be readily instantiated across a wide array of farming environments.

Limitations and challenges

The operationalised model may run into constraints when faced with shifting environmental variables, such as fluctuating light levels or adverse weather, which may compromise its predictive reliability. On-the-ground difficulties faced by cultivators—including mixed crop stands, non-uniform field surfaces (e.g., uneven ground), and the platform's capacity to adapt to diverse landscapes—must be addressed. These hurdles can be softened by strengthening the model's resilience through training across a broader range of environmental settings. Equally, equipping farmers with adequate familiarity before operating the system will help ensure they can use it effectively and overcome any procedural obstacles they encounter.

Performance comparison on CPU and Raspberry Pi

Table 3 presents a thorough head-to-head breakdown of the hardware and software characteristics that distinguish the development CPU from the Raspberry Pi 4 used for rollout in this investigation. The development CPU, anchored by an Intel Core i5 processor (7th gen), 8GB DDR4 RAM, and generous storage (512GB SSD + 1 TB HDD), is a formidable workstation tailored for model training and algorithm development. It boots Windows 11 and uses Python 3.13 within the Anaconda ecosystem, TensorFlow 2.7.0, and coding environments Jupyter Notebook and Spyder.

Table 3. Comparison of CPU specifications with Raspberry Pi.

Parameter	Desktop computer configuration	Raspberry Pi 4 model B configuration
Central processing unit (CPU)	Intel Core i5 processor (7th generation)	Quad-core 64-bit ARM processor operating at 1.5 GHz
Memory (RAM)	8 GB DDR4 RAM	4 GB LPDDR4 memory operating at 3200 MHz
Data storage	A combination of 512 GB SSD and 1 TB HDD	32 GB microSD storage card
Operating platform	Windows 11, 64-bit edition	Raspbian operating system
Graphics unit (GPU)	Presumed integrated graphics processor	VideoCore VI graphics engine
Connectivity features	Wi-Fi 802.11n (2.4 GHz), Gigabit Ethernet, and Bluetooth 5.0 support	Gigabit Ethernet, Bluetooth 5.0, and dual-band Wi-Fi supporting both 2.4 GHz and 5 GHz
Power source	Conventional desktop PC power unit	5V/15A USB Type-C power adapter
Imaging device	5 MP camera module	Logitech C270 Digital HD webcam with 5 MP and 720p capability
Software environment	Python 3.13, Anaconda distribution, and TensorFlow 2.7.0	Python with TensorFlow Lite

© 2023 The Author(s). | Published by Agro Intelligence Research Press

Development tools	Jupyter Notebook and Spyder IDEs	Python-based integrated development environments (IDEs)

In contrast, the Raspberry Pi 4, intended for on-site deployment, is powered by a Quad-core 64-bit ARM processor clocked at 1.5 GHz, paired with 4GB LPDDR4–3200 SDRAM and a 32GB microSD card as its storage medium. It runs Raspbian OS and leverages TensorFlow Lite to streamline edge model inference. Both systems ship with networking capabilities, including Wi-Fi and Gigabit Ethernet, though the Raspberry Pi also supports dual-band Wi-Fi. The Raspberry Pi's diminutive form factor and frugal energy draw (5V/15A type C USB) make it a natural fit for portable, in-field deployment scenarios.

Table 4 delivers a concise side-by-side appraisal of the two deep learning architectures, ResNet50 and DenseNet121, when running on a CPU versus a Raspberry Pi for the task of root crop leaf identification.

Table 4. Performance comparison on CPU and Raspberry Pi.

Model	Raspberry-Pi		CPU	
Empty Cell	**Time required in seconds**	**Accuracy (%)**	**Time required in seconds**	**Accuracy (%)**
ResNet50	0.3 s	99.60	0.313 s	99.60
DenseNet121	0.4 s	96.81	0.312 s	97.60

The interval each model needs to process a frame, from the moment of webcam capture, was carefully measured. The average per-image inference time was 0.3 and 0.4 seconds in both the Raspberry Pi environment and a conventional CPU-powered setup. ResNet50 demonstrates exemplary performance stability across both execution contexts. Running on the CPU, it delivers 99.60% accuracy. Once transferred to the Raspberry Pi, it maintains the same high-accuracy band, with any observed performance shift classified as negligible. This finding confirms that ResNet50 successfully transfers its strengths from a markedly more powerful CPU into the resource-sparse confines of the Raspberry Pi without incurring a meaningful accuracy penalty.

DenseNet121 delivers solid results on both platforms as well, though a slight accuracy drop is observed on the Raspberry Pi. On the CPU, it records an accuracy of 97.60%. When fielded on the Raspberry Pi, its accuracy settles at 96.81%, corresponding to a modest performance dip of 0.79%. This marginal-accuracy regression is notably impressive, especially given the substantial gulf in computational resources separating the two environments. The Raspberry Pi, hamstrung by comparatively meagre processing muscle and memory, nonetheless delivers results approaching those of a fully-fledged CPU. Such steadfastness in maintaining high accuracy across heterogeneous hardware regimes underscores the model's leanness and its promising potential for edge computing and on-farm agricultural technology solutions.

Testing was restricted exclusively to the 10% test partition, while models were validated against the 20% validation partition. Outcomes were derived from this image-based evaluation, with performance yardsticks including accuracy and classification latency documented for both the CPU and Raspberry Pi platforms. To supplement the image-based assessment, a field trial was conducted to evaluate the models' performance in a real-world agricultural landscape. This trial tested the models' robustness under routine daytime environmental conditions, including clear weather, moderately even terrain, ambient daylight, and typical humidity levels. The experimental configuration outlined above is illustrated in **Figure 18**.

Broadly, this comparison highlights the resilience and portability of both architectural designs. This shows that these sophisticated deep learning frameworks can be productively deployed on edge-class devices such as the Raspberry Pi for live, on-location crop discrimination with only a slight sacrifice in accuracy. This smooth migration from high-octane

© 2023 The Author(s). | Published by Agro Intelligence Research Press

computing environments to hardware-constrained platforms is a critical enabler for real-world agricultural deployments, where mobility and instantaneous processing are paramount.

Conclusion

The present study has fruitfully applied deep learning models to the task of sorting foliage across four root crop species: beetroot, potato, radish, and sweet potato. Capitalising on ResNet50 and DenseNet121 backbones, the investigation delivered outstanding capability supported by a sizeable repository of over 2500 images collected from numerous locations across Karnataka, India. Both architectures notched impressive accuracy figures, with ResNet50 reaching 99.6% and DenseNet121 reaching 97.6%, while steadfastly maintaining elevated precision, recall, and F1 scores across all crop classes. Critically, these models retained their robust showing under live classification conditions driven by CPU computation. A landmark accomplishment was the successful porting of these models to a Raspberry Pi 4 Model B, where ResNet50 achieved 99.6% accuracy, and DenseNet121 achieved 96.81% accuracy on this hardware-limited platform. This milestone received further corroboration through field-based trials, during which the Raspberry Pi assembly was mounted on a vehicle and operated successfully in an authentic agricultural environment. These outcomes illuminate the transformative promise that deep learning models harbour for agricultural technology, furnishing a trustworthy, efficient, and mobile solution for identifying root crop foliage. The steadfast high accuracy upheld across heterogeneous hardware platforms and genuine outdoor settings speaks to the durability and versatility of the engineered models.

Forward-looking research directions could emphasize broadening the image corpus to incorporate a greater assortment of root crop species and cultivation regimes, honing models for heightened efficiency on edge-bound hardware via tactics such as quantisation or network compression, launching multi-season longitudinal investigations to assess dependability and early anomaly flagging capabilities over complete crop cycles, and scrutinising the ecological impact and sustainability dimensions tied to deploying such technology across farming systems. Collectively, these pursuits would strengthen the models' capacity to generalise, reinforce their practical viability within resource-constrained settings, provide insight into performance over extended time horizons, and help ascertain that the technology contributes constructively to sustainable agricultural practice. Furthermore, upcoming efforts could explore the use of more muscular single-board computing platforms—such as the NVIDIA Jetson Nano, Odroid N2+, Rock Pi 4 Model C+, or ASUS Tinker Board 2—to boost the throughput of edge hardware, potentially clearing the way for more complex on-device computations and swifter field processing. Additionally, extending the system's repertoire to include functions such as disease identification, weed detection, and other sophisticated agricultural diagnostic capabilities would substantially enhance its utility for growers, making it a more comprehensive tool for crop stewardship and tactical decision-making. Such advancements would not merely elevate the platform's practical value but also likely broaden its applicability across a range of farming scenarios, thereby further propelling the evolution of intelligent agricultural methodologies.

Acknowledgements

None

Conflict of interest

None

Financial support

None

© 2023 The Author(s). | Published by Agro Intelligence Research Press

Ethics statement

None

Received: 21 Mar 2022 Revised: 26 May 2022 Accepted: 13 Jun 2022
Published online: 25 May 2023

Rights and permissions

Open Access The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at https://creativecommons.org/licenses/by-nc-sa/4.0/.

References

1. Harakannanavar SS, Rudagi JM, Puranikmath VI, Siddiqua A, Pramodhini R. Plant leaf disease detection using computer vision and machine learning algorithms. Glob Transit Proc. 2022:305–10.

2. Zhu X, Zhu M, Ren H. Method of plant leaf recognition based on improved deep convolutional neural network. Cogn Syst Res. 2018.

3. Mehdipour Ghazi M, Yanikoglu B, Aptoula E. Plant identification using deep neural networks via optimization of transfer learning parameters. Neurocomputing. 2020. https://doi.org/10.1016/j.neucom.2017.01.018.

4. Islam MM, Adil MAA, Talukder MA. DeepCrop: Deep learning-based crop disease prediction with web application. J Agric Food Res. 2023;14:100764.

5. Cao P, Wang T, Zhai L, Niu S, Liu L, Shi Y. Design of 6-DOF tomato picking lifting platform. Agriculture. 2022. https://doi.org/10.3390/agriculture12111945.

6. Buenconsejo LT, Linsangan NB. Classification of healthy and unhealthy abaca leaf using a Convolutional Neural Network (CNN). IEEE Access. 2021. https://doi.org/10.1109/HNICEM54116.2021.9732050.

7. Paul SK, Kewen XIA, Olanrewaju HS. A deep learning-based recognition technique for plant leaf classification. IEEE Access. 2021;9. https://doi.org/10.1109/ACCESS.2021.3131726.

8. Bisen D. Deep convolutional neural network based plant species recognition through features of leaf. Nature. 2020. https://doi.org/10.1007/s11042-020-10038-w.

9. Karahan T, Nabiyev V. Plant identification with convolutional neural networks and transfer learning. Cogn Syst Res. 2018. https://doi.org/10.5505/pajes.2020.84042.

10. Lee SH, Chan CS, Mayo SJ, Remagnino P. How deep learning extracts and learns leaf features for plant classification. Pattern Recognit. 2019. https://doi.org/10.1016/j.patcog.2017.05.015.

11. Hassan SKM, Maji AK. Identification of plant species using deep learning. Nature. 2021. https://doi.org/10.1007/978-981-15-7834-2_11.

12. Zhu H, Liu Q, Qi Y, Huang X, Jiang F, Zhang S. Plant identification based on very deep convolutional neural networks. Multimed Tools Appl. 2018;77(22):29779-97. https://doi.org/10.1007/s11042-017-5578-9.

13. Rahmat H, Wahjuni S, Rahmawan H. Performance analysis of deep learning-based object detectors on Raspberry Pi for detecting melon leaf abnormality. Int J Adv Sci Eng Inf Technol. 2022;12(2). https://doi.org/10.18517/ijaseit.12.2.13801.

14. Caballero J, Vergara F, Miranda R, Serracín J. Inference of recyclable objects with convolutional neural networks. arXiv. 2021.

15. Disabato S, Roveri M, Alippi C. Distributed deep convolutional neural networks for the Internet-of-Things. IEEE. 2021;7(8): 1239 - 52. https://doi.org/10.1109/TC.2021.3062227.

16. Ritzkal AA, Afrianto Y. Mask detection using framework tensorflow and pre-trained CNN model based on raspberry Pi. Iocsience J Mantik. 2020;4(3):1539–45.

17. Gupta V, Panchal VG, Singh V, Bansal D, Garg P. EmotionNet: resNeXt inspired CNN architecture for emotion analysis on Raspberry Pi. IEEE. 2021. https://doi.org/10.1109/RTEICT52294.2021.9573569.

18. Sarkar P, Raheman H. Development of a manually drawn single row onion set planter using a 2 DOF robotic arm. AgricEngInt CIGR. 2021;23(2). https://cigrjournal.org/index.php/Ejounral/article/view/6683/3619

19. Roshanianfard A, Du M, Nematzadeh S. A 4-DOF scara robotic arm for various farm applications: designing, kinematic modelling, and parameterization. Acta Technol Agric. 2021:61–6. https://doi.org/10.2478/ata-2021-0010.

20. Yoshida T, Onishi Y, Kawahara T, Fukao T. Automated harvesting by a dual-arm fruit harvesting robot. ROBOMECH. 2022;9:19. https://doi.org/10.1186/s40648-022-00233-9.

21. Altawil B, Can FC. Design and analysis of a four DoF robotic arm with two grippers used in agricultural operations. Int J Appl Math Electron Comput. 2023;11(2):079–87. https://doi.org/10.18100/ijamec.1217072.

22. Rahul K, Raheman H, Paradkar V. Design of a 4 DOF parallel robot arm and the firmware implementation on embedded system to transplant pot seedlings. Artif Intell Agric. 2020;4:172–83. https://doi.org/10.1016/j.aiia.2020.09.003.

23. Paradkar V, Raheman H, Rahul K. Development of a metering mechanism with serial robotic arm for handling paper pot seedlings in a vegetable transplanter. Artif Intell Agric. 2021;4:52–63. https://doi.org/10.1016/j.aiia.2021.02.001.

24. Paryadi C, Diqi M, Mulyani SH. Implementation of CNN for plant leaf classification. Int J Inform Comput. 2020;2(2).

25. Zamir M, Ali N, Naseem A, Frasteen AA, Zafar B, Assam M, et al. Face detection and recognition from images and videos based on CNN and Raspberry Pi. Computation. 2022;10(9):148. https://doi.org/10.3390/computation10090148.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

26. Bansal P, Kumar R, Kumar S. Disease detection in apple leaves using deep convolutional neural network. Agriculture. 2021;11:617. https://doi.org/10.3390/agriculture11070617.

27. Sahidan NF, Juha AK, Mohammad N, Ibrahim Z. Flower and leaf recognition for plant identification using convolutional neural network. Indones J Electr Eng Comput Sci. 2019;16(2):737–43. https://doi.org/10.11591/ijeecs.v16.i2.pp737-743.

28. Hu J, Chen Z, Yang M, Zhang R, Cui Y. A multiscale fusion convolutional neural network for plant leaf recognition. IEEE Signal Process Lett. 2018;25(6). https://doi.org/10.1109/LSP.2018.2809688.

29. Xenakis A, Papastergiou G, Gerogiannis VC, Stamoulis G. Applying a convolutional neural network in an IoT robotic system for plant disease diagnosis. IEEE Access. 2020. https://doi.org/10.1109/IISA50023.2020.9284356.

© 2023 The Author(s). | Published by Agro Intelligence Research Press