

Early Detection of Crop Diseases Using CNN Classification

Aryan Rajvanshi¹, Shiang-Wan Chin²

Received October 30, 2023

Accepted January 29, 2024

Electronic access February 29, 2024

Crop diseases stand as formidable adversaries to agricultural productivity and food security. Detecting these diseases at their inception is a critical prerequisite for effective management and intervention. However, the human eye's limitations often result in disease identification only after their detrimental effects have proliferated substantially within crops. In response, we introduce a bespoke Convolutional Neural Network (CNN) classification model, engineered to expedite and enhance the accuracy of plant disease recognition. Our study encompasses a comprehensive training and evaluation of the CNN model using diverse datasets encompassing apple, corn, and tomato crops, all procured from sources like Kaggle. Unlike traditional classification models that are trained and accustomed to specific datasets, our model is capable of classifying images with various lighting conditions, orientations and sizes which in turn can be used in various farm setups. Through a systematic process of training and validation, our CNN model consistently demonstrated testing accuracies of 97%-99% in categorizing diseases across all datasets. This accuracy is much higher than a majority of other CNN models used for crop diseases detection. The implications of our findings reverberate profoundly within the realm of agriculture and crop management. The integration of our AI-driven technology in crop disease detection not only addresses immediate agricultural concerns but also aligns with broader global challenges. The proactive identification of crop diseases through our innovative model contributes to sustainable agriculture by minimizing yield losses and promoting efficient resource utilization. This helps meet the needs of a growing population and reduces the impact of environmental stress on global food production. However, it is important to note that since this is a classification model, certain pictures with other distinct elements may yield low accuracies. This research underscores the instrumental role AI-driven technology can play in fortifying agricultural resilience and ensuring global food security.

Introduction

Agriculture stands as a pivotal economic and societal cornerstone. Over the past two centuries, the human population has grown sevenfold and the experts anticipate the addition of 2–3 billion more during the twenty-first century¹. With a projected global population of 10 billion by 2050, food production must increase by 60%². To meet this demand, reducing losses from pests and diseases, as well as food waste, is crucial. Currently, crop pests and diseases cause substantial global yield losses, including around 21.5% in wheat, 30.3% in rice, 22.6% in maize, 17.2% in potatoes, and 21.4% in soybeans³. Detecting crop diseases at an early stage to enhance crop quality and yield poses a challenge for human efforts. The multifarious origins of such diseases, ranging from climatic shifts to nutritional deficiencies and pest assaults, compound the complexity of their prompt identification⁴. Conventional crop disease detection involves labor-intensive manual techniques, including visual scrutiny and microscopic analysis. However, these methods are prone to inaccuracy due to varying human perceptions and interpretative errors, hindering consistent results⁴. Inherent limitations in human visual precision and the time-intensive nature of diagnosis result in inevitable errors. Traditionally, disease diagnosis relied on visual inspection by domain experts, plant

pathologists, and farmers⁵. However, relying on human visual precision for crop disease identification is impractical, especially in developing countries and small farms. It's a time-consuming and expert-dependent process, leading to additional costs for rare diseases. This method is not feasible for large farms and can result in biased decisions, highlighting the need for more efficient⁶. The contemporary focus has shifted towards automating disease detection using deep learning, outperforming traditional methodologies. AI-based crop disease detection surpasses traditional methods like molecular and serological techniques, offering faster and more accessible identification. Unlike expensive hyperspectral imaging, AI utilizes affordable digital cameras, providing a swift and accurate solution for widespread implementation, making it a superior choice for efficient and timely crop health monitoring⁶. Hence, a digital image-based automatic disease identification approach emerges as a practical and viable solution to address these challenges⁵.

A digital image-based automatic disease identification approach refers to the utilization of advanced computational techniques, particularly those derived from artificial intelligence and machine learning, to automatically detect and diagnose diseases in crops through the analysis of digital images. In this approach, sophisticated algorithms are trained on a diverse dataset of images showcasing both healthy and diseased plants.

These algorithms learn to recognize intricate visual patterns and anomalies associated with various crop diseases. By processing and analyzing digital images of plants, these algorithms can accurately and swiftly identify the presence of diseases, enabling early intervention and effective disease management.

Deep learning, a subset of AI, has emerged as a contemporary technique for image processing and data analysis, demonstrating potential across various domains, including agriculture⁷. Deep learning plays a crucial role in agriculture seen in cases where it can accurately predict crop yields based on various factors like weather and soil⁸. This rise of machine learning (ML), driven by vast data and advanced computing, has found applications in diverse sectors, prompting the exploration of deep learning (DL), a subset focused on modeling complex abstractions through layered algorithms inspired by neural networks⁹. DL's expansion into artificial intelligence has been evident in natural language processing and image classification⁹. In the realm of crop disease detection, DL's prowess is transformative⁷. By leveraging its ability to process extensive datasets and identify subtle patterns, DL offers a solution for early disease detection⁷. Trained on images of healthy and diseased plants, DL algorithms can discern visual cues signifying diseases, facilitating timely intervention⁷. This shift from manual inspection to automated precision agriculture presents a significant stride in enhancing productivity and sustainability⁹. Integrating DL into crop disease detection showcases the fusion of technology and agriculture, holding promise for more efficient disease management⁹.

In the realm of deep learning, Convolutional Neural Networks (CNNs) have emerged as a distinct class of neural systems, predominantly tailored for visual analysis¹⁰. These networks consist of multiple layers of artificial neurons, which mimic the behavior of natural neurons by calculating weighted sums of inputs to produce activation values¹⁰. Upon feeding an image into a CNN, each layer progressively extracts fundamental features like edges or corners, gradually identifying more complex attributes such as objects¹⁰. This approach to feature recognition aligns with the inherent structure of visual data, making CNNs particularly effective in tasks such as image classification^{10,11}. Image classification, a fundamental task in computer vision, involves categorizing images into predefined classes¹¹. Deep learning, heavily reliant on CNNs, excels at automatically extracting informative features from images¹¹. This efficiency and accuracy of CNNs in deciphering intricate visual information positions them as the ideal choice for tasks like the early detection of crop diseases, where subtle visual indicators play a vital role^{10,11}.

Numerous research papers have demonstrated the effectiveness of CNNs in the realm of crop disease detection¹²⁻¹⁵. However, within this body of work, certain gaps and challenges emerge. Many studies have primarily focused on the development and validation of their CNN models using specific datasets,

often failing to ensure the generalizability of their models across diverse conditions and crops. This lack of cross-applicability limits the practicality of these models in real-world agricultural scenarios where crops and environmental conditions vary considerably. Furthermore, another prevalent issue is the insufficiency of image data in datasets. The scarcity of high-quality images of various diseases and environmental conditions hampers the ability of CNNs to learn robust and representative features. This limitation in data diversity may hinder the model's ability to generalize to real-world scenarios, as it may struggle to recognize and accurately classify new instances of diseases or environmental variations not well-represented in the training data.

In this study, a Convolutional Neural Network (CNN) model with custom dense layers was developed for accurate and efficient crop disease detection. The model made use of three primary convolutional layers and two dense layers to increase the efficiency of the model. This is different when compared to the earlier studies as instead of using a large number of convolutional layers, this model prioritizes speed so that farmers can be notified about their diseases as soon as possible. The model's efficacy was evaluated using three diverse datasets, each a combination of images primarily sourced from pictures taken by us in farm fields, and those from the PlantVillage dataset¹⁶: Apple, Corn, and Tomato. The pictures were taken at 16 different farm locations in Gujarat during different time periods (7am to 8am, 2pm to 3pm, 6pm to 7pm IST), 12 of which were used for training, allowing the training images to have various backgrounds and lightings. The labelled images were verified with an agriculturalist to ensure reliable training data. Through this, the trained CNN model exhibited a remarkable ability to generalize when tested with a separate apple dataset consisting of the pictures collected in the other four farms, achieving good accuracy, thereby underscoring its potential for cross-dataset applicability. Across the original three crops, the CNN model consistently demonstrated high accuracy during disease classification. Each crop dataset was structured with three distinct classes: Healthy plants and two categories for different diseases. The hypothesis driving this research was that the CNN model would excel in identifying crop diseases across various types, yielding impressive accuracy rates. This study seeks to contribute to the advancement of early disease detection in agriculture and enhance overall crop health management.

Results

Architecture of the Model

In this study, we used a Convolutional Neural Network (CNN) for the purpose of image classification and feature extraction. To evaluate the model's versatility and robustness, we subjected it to rigorous testing using Kaggle datasets encompassing tomato, corn, and apple crops. Central to our CNN model's construction

were several key techniques that collectively contributed to its effectiveness. The process of flattening, for instance, played a crucial role in transforming the complex image arrays into a simplified one-dimensional format. This not only streamlined the data for more efficient processing but also allowed the model to better capture essential features across different samples. Batch normalization, another integral technique, proved instrumental in standardizing intermediate outputs across the various layers of the network. Batch normalization enhanced both the speed and stability of the training process, ultimately leading to improved performance. Dense layers enabled the establishment of intricate connections between multiple neurons. This allowed the model to learn and recognize complex spatial hierarchies in the input images, enhancing its ability to differentiate between different crop diseases and healthy states. Moreover, the application of activation functions introduced non-linearity into the model, enabling it to capture intricate relationships and dependencies present in the data. This non-linearity is pivotal for distinguishing between diverse disease characteristics and ensuring the model's capacity to learn and generalize effectively.

Another key term, an epoch refers to a single iteration through the entire training dataset during the training phase of a model. During an epoch, the model processes each training example once, updating its internal parameters based on the calculated loss and the chosen optimization algorithm. The number of epochs used for training is a hyper parameter that determines how many times the model iterates through the entire dataset. A larger number of epochs can allow the model to better learn the underlying patterns in the data, but excessive epochs can lead to overfitting, where the model becomes too tailored to the training data and performs poorly on new, unseen data. Properly choosing the number of epochs is essential for achieving the right balance between learning and generalization. We used 6 epochs for tomato and apple as going beyond that caused the model to start overfitting, and results in a decrease in AUC and accuracy. In the case of corn, 8 epochs were used as training AUC consistently increased until the 8th epoch followed by continuous drops.

Additionally, we implemented a checkpoint mechanism to closely monitor the model's training progress. This technique ensured that optimal parameter configurations were saved at specific intervals, safeguarding against overfitting and equipping the model with enhanced generalization capabilities.

Testing the Apple dataset

We used three classes of apple leaves from the dataset. These were 'Healthy', 'Apple Scab', and 'Apple Cedar Rust.' (Figure 1) Apple Scab and Apple Cedar Rust are two diseases that infect leaves, and the healthy class contains uninfected pictures of apple leaves. All the data was split into three different sets: 80% of all the data was assigned as training data (2040 images), 10%

was assigned as validation data (255 images), and the remaining 10% was test data (255 images). We then used 1316 images, 504 images, 220 images for training; 164 images, 63 images, 28 images for validation; 165 images, 63 images, 27 images for testing for 'Healthy Apple Leaves', 'Apple Scab Leaves', and 'Apple Rust Leaves' respectively as shown in Table 1.



Fig. 1 (Left) Healthy Apple Leaf, (Middle) Apple Scab (Right) Apple Cedar Rust

Table 1 Apple dataset images for classes in training, validation and testing

Apple Leaf Images	Apple Healthy	Apple Scab	Apple Rust	Total
Training	1316	504	220	2040
Validation	164	63	28	255
Testing	165	63	27	255

We employed the categorical cross entropy loss function, a standard choice for multi-class classification tasks such as ours. This function quantifies the dissimilarity between the predicted class probabilities and the actual class labels for each input image. By minimizing this cross-entropy loss during training, our model learns to make accurate predictions across the various disease categories in our dataset.

When it comes to assessing the effectiveness of our model's performance, we turned to the AUC (Area Under the Curve) metric. Although accuracy is a commonly used evaluation metric, AUC offers a more comprehensive measure, particularly beneficial in scenarios involving imbalanced class distributions. AUC accounts for the entire range of possible classification thresholds, providing insight into the model's ability to distinguish between healthy plants and those affected by diseases. By focusing on AUC, we ensure a nuanced evaluation that accommodates the intricacies of our classification task and aligns with our objective of enhancing disease detection accuracy.

The model's training performance was notable, achieving a training AUC of 0.978 after 6 epochs. It exhibited a consistent upward trajectory in AUC values up to the 4th epoch, after which it reached a stable point. Similarly, the validation AUC displayed commendable progress, reaching 0.988 after 6 epochs. The validation AUC had an initial drop in AUC at

epoch 2, but increased consistently until epoch 4 followed by small changes till epoch 6 as seen in Figure 2. The initial dip in validation AUC at epoch 2 may stem from factors like the model adjusting to noise in early training data or exploring suboptimal weight configurations. As training continued, the model refined its parameters, recovering and stabilizing by discerning more meaningful patterns in the data. Impressively, the testing AUC performance reached a peak at 0.991 (Table 2).

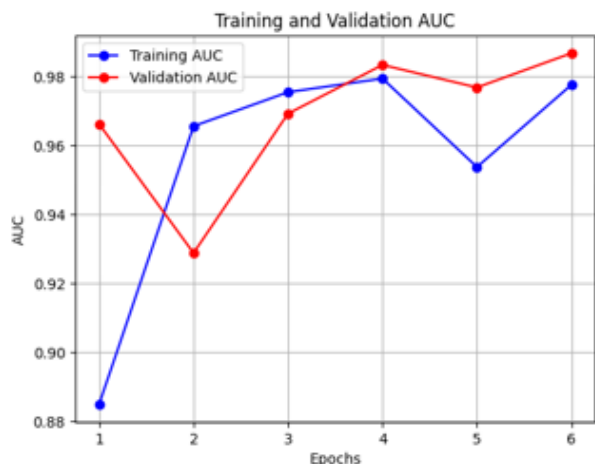


Fig. 2 Training Area Under the Curve (AUC) and validation AUC of the Apple dataset model

Table 2 AUC Results for all classes

AUC Results	Testing AUC	Validation AUC	Training AUC
Apple	0.982	0.988	0.991
Corn	0.988	0.999	0.996
Tomato	0.988	0.967	0.970

For a comprehensive evaluation of the model’s effectiveness, we also evaluated its accuracy. The training accuracy peaked at 0.927 after 6 epochs, demonstrating a steady improvement throughout the training process. In parallel, the validation accuracy peaked at 0.926 during the 5th epoch, with consistent growth leading up to this point. However, during the 6th epoch, a noticeable decline was observed in the model’s validation accuracy. Finally, the test accuracy recorded an impressive 0.946 (Table 3).

Testing the Corn dataset

We used three classes of corn leaves from the PlantVillage dataset. These were ‘Healthy’, ‘Corn Northern Blight’, and ‘Corn Common Rust.’ (Figure 3) Corn Northern Blight and Corn Common Rust are two diseases that infect corn leaves, and

the healthy class contains uninfected pictures of corn leaves. All the data was split into three different sets: 80% of all the data was assigned as training data (2672 images), 10% was assigned as validation data (333 images), and the remaining 10% was test data (334 images). We then used 930 images, 954 images, 788 images for training; 116 images, 119 images, 98 images for validation; 116 images, 119 images, 99 images for testing for ‘Healthy Corn Leaves’, ‘Corn Northern Blight’, and ‘Corn Common Rust’ respectively as shown in Table 4.




Fig. 3 (Left) Healthy Corn Leaf, (Middle) Corn common rust, (Right) Corn Northern Blight

The model achieved its highest training AUC performance of 0.988 after 8 epochs, displaying an ascending trend in AUC values until the 4th epoch, which then stabilized. The model’s validation AUC performance peaked at an impressive 0.999 after 6 epochs. In the testing phase, the model exhibited a strong AUC performance of 0.996 (Table 2).

Regarding training accuracy, the model attained its peak performance of 0.987 after 8 epochs, demonstrating a consistent trend of improvement until the 6th epoch. Similarly, the model’s validation accuracy reached a notable 0.996 after 8 epochs. The model’s proficiency was also evident in the testing accuracy, achieving a commendable accuracy of 0.996 (Table 3).

Testing the Tomato dataset

We used three classes of tomato leaves from the PlantVillage dataset. These were ‘Healthy’, ‘Tomato Early Blight’, and ‘Tomato Late Blight.’ (Figure 4) Tomato Early Blight and Tomato Late Blight are two diseases that infect tomato leaves, and the healthy class contains uninfected pictures of tomato leaves. All the data was split into three different sets: 80% of all the data was assigned as training data (3600 images), 10% was assigned as validation data (450 images), and the remaining 10% was test data (450 images). We then used 1273 images, 800 images, 1527 images for training; 159 images, 100 images, 191 images for validation; 159 images, 100 images, 191 images for testing for ‘Healthy Tomato Leaves’, ‘Tomato Early Blight’, and ‘Tomato Late Blight’ respectively.

The model demonstrated robust performance in terms of AUC metrics. Specifically, the highest training AUC performance reached 0.988 after 6 epochs, with an upward trend consistently observed until the 6th epoch. The validation AUC performance

Table 3 Accuracy Results for all classes

Accuracy Results	Testing Accuracy	Validation Accuracy	Training Accuracy
Apple	0.927	0.926	0.946
Corn	0.987	0.996	0.996
Tomato	0.885	0.888	0.882

Corn Leaf Images	Corn Healthy	Common Rust	Northern Blight	Total
Training	930	954	788	2672
Validation	116	119	98	333
Testing	116	119	99	334

Table 4 Corn dataset images for classes in training, validation, and testing



Fig. 4 (Left) Healthy Tomato Leaf, (Middle) Tomato Early blight, (Right) Tomato Late Blight

peaked at 0.967 after 6 epochs. In testing, the model maintained a commendable AUC performance of 0.97 (Table 2).

Regarding accuracy metrics, the model's training accuracy reached its peak of 0.885 after 6 epochs, and this upward trend persisted throughout the training process. Correspondingly, the best validation accuracy was achieved at 0.888 after 6 epochs. During the testing phase, the model exhibited an accuracy of 0.882 (Table 3), reflecting its consistent and reliable performance across various evaluation stages.

Discussion

This study revolves around the development and evaluation of a Convolutional Neural Network (CNN) model aimed at image classification and feature extraction. Leveraging agricultural datasets sourced from Kaggle, our investigation delved into the model's performance across apples, corn, and tomatoes. The application of techniques like data flattening, batch normalization, dense layers, activation functions, and a checkpoint mechanism fortified our model. Notably, our assessment using the AUC metric corroborated the model's efficacy in effectively discerning between healthy and disease-affected crops, underscoring its potential for precise agricultural disease detection¹⁷.

With our foundational framework in place, a systematic eval-

uation of the model unfolded across diverse datasets. Commencing with the Apple dataset, housing categories like 'Healthy', 'Apple Scab', and 'Apple Cedar Rust', our analysis employed an 80-10-10 split for training, validation, and test sets, and the AUC metric served as our yardstick. Impressively, the test phase unveiled a robust AUC, spotlighting the model's acumen in distinguishing healthy from disease-stricken apple leaves. Transitioning to the Corn dataset, encompassing categories including 'Healthy', 'Corn Northern Blight', and 'Corn Common Rust', our model consistently attained remarkable AUC peaks during both training and validation phases, bolstering its ability to comprehend intricate disease-associated nuances. Similarly, the Tomato dataset, featuring 'Healthy', 'Tomato Early Blight', and 'Tomato Late Blight' classes, witnessed the model's consistent progression in performance.

The study's results are promising, notably in terms of disease detection accuracy, substantiated by impressive accuracy metrics. However, it's essential to acknowledge inherent limitations and potential sources of errors stemming from variations in data collection, image quality, lighting conditions, and camera angles. Moreover, the model's efficacy in handling novel diseases or variations not present in the training data warrants scrutiny. A notable assumption within the proposed methodology is that the labeled data is devoid of errors, a notion that may not hold true in real-world scenarios. The model's limitations may also become evident in situations where there are nuanced textures or subtle visual distinctions among different crop diseases. The CNN's reliance on visual features might result in misclassifications, especially when faced with uncommon diseases or environmental conditions that significantly differ from the training data.

To address these multifaceted challenges, a strategic approach is imperative. Resolving variations in data collection and image quality necessitates robust preprocessing techniques, such as image normalization and augmentation, to fortify the model's adaptability to diverse conditions¹⁸. Inclusion of diverse datasets spanning varying lighting conditions and angles

could potentially augment the model's generalization capabilities¹⁹. Regularly infusing the training dataset with new instances of diseases or variations can enhance the model's versatility in tackling novel cases²⁰. In response to the challenge posed by novel diseases, integrating transfer learning techniques can prove advantageous, pre-training the model on a wider spectrum of diseases before fine-tuning it on the specific dataset of interest to bolster its capacity to handle unseen variations²⁰. While suggesting future transfer learning, we chose not to incorporate it in this study for two main reasons. Firstly, the specific challenges of domain adaptation were a concern, as the source domain (pre-training data) might significantly differ from the target domain (crop disease dataset), potentially impacting model performance. Secondly, the selection of a suitable pre-trained model and the risk of overfitting during fine-tuning demanded careful consideration, and we opted for a more controlled approach in the initial phase of the study.

Additionally, exploring semi-supervised or weakly supervised learning methodologies can alleviate the model's dependence on perfectly labeled data²¹ by leveraging both labeled and unlabeled data, allowing the model to adapt to diverse patterns in real-world agricultural conditions. However, challenges include managing uncertainty in unlabeled data and maintaining a proper balance to avoid noise interference.

In the case of potential future experiments to build upon this study, replacing the current classification model with an object detection model could significantly enhance disease detection accuracy by not only identifying the disease but also precisely localizing the affected areas within the images. This level of detail can provide more accurate disease assessment, support targeted interventions, and accommodate scenarios where multiple instances of diseases are present in a single image²². Unlike the current classification model, an object detection model will not be hindered by multiple distinct elements in an image, and will instead be able to identify several diseases at a time. This could potentially lead to increase accuracy and efficiency.

An additional avenue explored could be the integration of drones in agriculture, as exemplified by studies such as²³ offering promising opportunities to enhance existing disease detection models. Unmanned aerial vehicles (UAVs) equipped with advanced cameras provide comprehensive crop imagery from different angles and altitudes, which can be seamlessly incorporated into current disease detection methodologies. This addition of aerial data can offer a more dynamic and holistic perspective on crop health, potentially improving the accuracy and timeliness of disease detection²⁴.

The developed CNN model exhibits remarkable scalability and practical applicability, paving the way for transformative impacts across agricultural landscapes. Its scalability is evident in its adaptability to an extensive array of crops, offering a versatile solution for disease detection. Beyond the initially evaluated apple, corn, and tomato datasets, this model can extend its disease

identification capabilities to a multitude of agricultural contexts. This model's adaptability is further enhanced by its potential integration with cutting-edge technologies like UAVs, enabling comprehensive crop health monitoring²³.

The significance of this study extends to reshaping agricultural practices through cutting-edge technology. The model's consistent performance across various crops underscores its potential in redefining early disease detection—a critical facet of modern agriculture. By enabling timely interventions, this approach can mitigate yield losses, optimize resource allocation, and ultimately contribute to global food security²⁵. The model's adaptability across diverse datasets further reinforces its practical utility in real-world scenarios. In setting a robust framework for disease detection, this study forms the basis for the development of sophisticated tools for sustainable crop disease management, advancing agricultural technology and offering a tangible solution to an imperative global challenge.

Materials and Methods

In this phase of the study, we harnessed data from the PlantVillage dataset, a comprehensive collection of agricultural images available on Kaggle, a prominent platform for data science collaboration. Kaggle serves as a hub for diverse datasets and resources, fostering robust research and analysis. The dataset comprised standardized high-resolution images, with each individual leaf image scaled to dimensions of 256 x 256 pixels. To optimize model performance, vital preprocessing steps were executed. Data normalization emerged as a critical initial procedure, involving the rescaling of pixel values within images to a standardized range of [0, 1].^{2,4} Normalization makes the model adaptable by standardizing feature scales, ensuring it focuses on relevant patterns in data. This consistency helps the model generalize effectively across diverse conditions in crop disease recognition, making accurate predictions in various agricultural scenarios. Additionally, data augmentation techniques were employed to enrich the diversity and quality of the training dataset. Notably, zoom and flip augmentations were introduced, introducing variability in leaf orientations and sizes. These augmentations significantly contribute to enhanced generalization and the reduction of overfitting risk. The model faced significant overfitting due to the insufficient amount of training data decreasing accuracy down to 78% in the case of apple, but data augmentation helped fix this issue.

CNN Architecture details

The employed Convolutional Neural Network (CNN) architecture and training configuration were pivotal components of this study's methodology for image classification. Implemented using the Keras library, the model's structure comprised a sequential arrangement of layers. The initial layer was a 2D con-

volutional layer with 32 filters and a ReLU activation function, designed to capture foundational image features. I opted for Convolutional (Conv2D) layers with Rectified Linear Unit (ReLU) activations to extract meaningful spatial features from the input images. This was followed by a max-pooling layer to down sample the feature maps. The Max-pooling layers reduce dimensionality and introduce translation invariance, which enhances generalization. Subsequently, two additional convolutional layers with 64 and 128 filters, respectively, were integrated, each paired with ReLU activation and max-pooling layers. A 'Flatten' layer converted the output to a one-dimensional vector for further processing. The model's interpretative capacity was enhanced through a dense layer comprising 128 units and a ReLU activation function. Lastly, a dense layer with a softmax activation function was employed to yield class probabilities for the three target categories. These layers take the high-level features learned by the convolutional layers and use them to make predictions. They introduce non-linearity and capture complex relationships in the data, making them essential for classifying images into different categories. The selection of filters and units in the Convolutional Neural Network (CNN) was a strategic decision, balancing model complexity and efficiency. Beginning with 32 filters in the initial layer captures basic features like edges, while the subsequent use of 64 and 128 filters in deeper layers allows the model to learn increasingly complex patterns. The dense layer with 128 units facilitates the understanding of intricate relationships within the learned features. These choices were made through experimentation and tuning, aiming for a balanced architecture that captures the nuances of the data without becoming overly complex. The model was compiled using the Adam optimizer and categorical cross-entropy loss function, which is particularly suitable for multiclass classification tasks. The Adam optimizer was chosen for its adaptive learning rate capabilities and effective handling of sparse gradients, contributing to enhanced performance and faster training convergence in our CNN algorithm for crop disease recognition. Other optimizers like Nadam and Adagrad were used but Adam gave the best results and hence was used. During training, AUC was tracked as a key metric, followed by a rerun using accuracy as the key metric. The 'ModelCheckpoint' callback mechanism was instituted to facilitate the preservation of the optimal model based on validation AUC/accuracy. The ModelCheckpoint strategy safeguards against overfitting by regularly saving the best-performing model weights during training. This ensures that the model generalizes well to new data and can revert to optimal states, mitigating the impact of overfitting.

Acknowledgements

I would like to acknowledge Mr. Shiang-Wan Chin for guiding me through the process of building the neural network model.

References

- 1 N. Fedoroff, *Agriculture amp; Food Security*, **4**, year.
- 2 J. Ristaino, P. Anderson, D. Bebbler, K. Brauman, N. Cunniffe, N. Fedoroff, C. Finegold, K. Garrett, C. Gilligan, C. Jones, M. Martin, G. MacDonald, P. Neenan, A. Records, D. Schmale, L. Tateosian and Q. Wei, *Proceedings of the National Academy of Sciences*, **118**, year.
- 3 S. Savary, L. Willocquet, S. Pethybridge, P. Esker, N. McRoberts and A. Nelson, *Nature Ecology amp; Evolution*, **3**, 430–439.
- 4 A. Hidayah, S. A. Radzi, N. A. Razak, Y. Wira Hidayat Mohd Saad and A. Naja, *Journal of Robotics and Control (JRC)*, **3**, 790–99.
- 5 M. A. Haque, S. Marwaha, C. Deb, S. Nigam, A. Arora, K. Hooda, P. Soujanya, S. Aggarwal, B. Lall, M. Kumar, S. Islam, M. Panwar, P. Kumar and R. Agrawal, *Scientific Reports*, **12**, year.
- 6 H. Orchi, M. Sadik and M. Khaldoun, *Agriculture*, **12**, 9.
- 7 A. Kamilaris and F. Prenafeta-Boldú, *Computers and Electronics in Agriculture*, Elsevier BV, vol. 147, p. 70–90.
- 8 S. Lee, Y. Jeong, S. Son and B. Lee, *Sustainability*, **11**, year.
- 9 F. BAL and F. KAYAALP, *International Advanced Researches and Engineering Journal*, **5**, 309–323.
- 10 P. Pai, A. Bakshi, A. Kumar, B. Anand, D. Bhartiya, B. D. R and R., *Journal of Computing and Natural Science*, 46–54.
- 11 A. Elngar, M. Arafa, A. Fathy, B. Moustafa, O. Mahmoud, M. Shaban and N. Fawzy, *Journal of Cybersecurity and Information Management*, 18–50.
- 12 M. Jung, J. Song, A.-Y. Shin, B. Choi, S. Go, S.-Y. Kwon, J. Park, S. Park and Y.-M. Kim, *Scientific Reports*, **13**, year.
- 13 P. Khobragade, A. Shriwas, S. Shinde, A. Mane and A. Padole, 2022 International Conference on Smart Generation Computing, Communication and Networking (SMART GENCON). 2022 International Conference on Smart Generation Computing, Communication and Networking (SMART GENCON). IEEE.
- 14 P. B.S, N. M. A, R. M, S. S and S. R. B.S, *INTERNATIONAL JOURNAL OF ENGINEERING RESEARCH TECHNOLOGY (IJERT) NCCDS*.
- 15 2020 International Conference on Artificial Intelligence, Big Data, Computing and Data Communication Systems (icABCD). 2020 International Conference on Artificial Intelligence, Big Data, Computing and Data Communication Systems (icABCD). IEEE.
- 16 A. Ali, *Kaggle*.
- 17 A. Carrington, P. Fieguth, H. Qazi, A. Holzinger, H. Chen, F. Mayr and D. Manuel, *BMC Medical Informatics and Decision Making*, **20**, year.
- 18 P. Pradhan and R. Routray, *International Journal for Research in Applied Science and Engineering Technology*, **10**, 1511–1514.
- 19 E. Bart and S. Ullman, Proceedings of the British Machine Vision Conference 2004. British Machine Vision Conference 2004. British Machine Vision Association.
- 20 S. Ada, E. Ugur and H. Akin, Generalization in Transfer Learning (Version 2). arXiv.
- 21 R. Bruce, *A Bayesian Approach to Semi-Supervised Learning*.

-
- 22 M. Pratama, S. Kim, S. Ozawa, T. Ohkawa, Y. Chona, H. Tsuji and N. Murakami, 2020 International Joint Conference on Neural Networks (IJCNN). 2020 International Joint Conference on Neural Networks (IJCNN). IEEE.
- 23 M. Rajagopal and B. MS, *Artificial Intelligence based drone for early disease detection and precision pesticide management in cashew farming*, <https://doi.org/10.48550/ARXIV.2303.08556>, arXiv.
- 24 S. Kumar, D. S. Vemuru and D. A. Srinath, *Blue Eyes Intelligence Engineering and Sciences Engineering and Sciences Publication - BEIESP*, **9**, 931–935.
- 25 M. Masood, H. Saim, M. Taj and M. Awais, *Early Disease Diagnosis for Rice Crop*, arXiv.