Fast Parallelizable Cassava Plant Disease Detection using Ensemble Learning with Fine Tuned AmoebaNet and ResNeXt-101

Madan Mohan Tito Ayyalasomayajula[a], Sathishkumar Chintala[b]

[a] Computer Science, School of Business & Technology, Aspen University, Pheonix, AZ, USA
[b]Independant Researcher, Texas, USA

Abstract: The study focuses on applying advanced machine learning approaches, namely deep learning models and ensemble learning, to develop an automated system for Cassava Brown Streak Disease (CBSD) detection using robotic systems. We leverage an extensively used dataset to evaluate our proposed methodologies. The dataset comprises thousands of labeled images illustrating various stages of CBSD infection. Our experiment is furnished with high-resolution cameras and sophisticated image-processing algorithms. We finetune AmoebaNet and ResNeXt-101 32x16d models using the dataset to distinguish between healthy and diseased cassava plants. Ensemble learning is then applied to consolidate the prediction outputs from both models, consequently enhancing overall diagnostic accuracy. The implemented system showcases exceptional performance, delivering high precision and recall rates in recognizing CBSD cases. Through automation, our solution significantly diminishes the reliance on human expertise, streamlines the diagnostic procedure, and extends the reach of CBSD detection applications. This pioneering research marks a significant stride forward in fortifying food security and promoting sustainable agriculture in regions affected by CBSD. Further details about the dataset usage will be discussed in the Methodology section.

Keywords: Cassava, Disease, Detection, Binary, Classification, Deep, Learning, AmoebaNet, ResNeXt-101, 32x16d, Paired, t-test, McNemar's, Test.

1. Introduction and Significance Of The Study

Cassava (Manihot esculenta) is a primary staple food in the world. It is a starchy root that is used in preparing different dishes. With the increase in production demand, proper care should be given to the plant, and quality should be maintained. The plant can be attacked by different kinds of diseases, as shown in Fig 1. With the increasing demand for production, it becomes difficult to monitor and diagnose the different kinds of plant diseases. The degraded leaves can be damaged by pests, malnutrition, and various types of diseases affecting the plants. Agricultural culture plays a major role in the growth of any country's economy. With the growing population in the world, the demand for food and agricultural products is also increasing. To cope with this demand, farmers are using machine learning-based technologies. Numerous cutting-edge technologies are present that help in growth and production estimation, yield estimation, and many others. Advanced computer vision and machine learning algorithms are used in the development of these technologies [1].

Cassava Brown Streak Disease (CBSD) is a devastating condition that negatively impacts cassava plants, threatening agricultural production and food security predominantly in East African countries. Accurately and promptly diagnosing CBSD is vital to contain its spread and minimize yield losses. Traditional methods for diagnosing CBSD rely on manual inspections conducted by experts, which can be laborious, time-consuming, and inconsistent. As the demand grows for efficient and dependable methods to identify CBSD, automation via robotic systems emerges as a compelling alternative. Currently, cassava diseases are mostly identified by visual observation by the plant protection staff or cassava growers. However, visual symptoms will not appear on the leaves until the disease has reached a certain stage, and the diagnosis based on visual observation is still challenging in identifying and detecting the diseases. If the disease goes undetected for a long period, it could damage the plant and cause large-scale loss, even killing the plant. Accordingly, to prevent large-scale loss, develop a higher cassava yield, and ensure the food safety of the planet, it is essential to research and develop more effective tools to identify the early stages of common diseases of cassava, as this would enable an early diagnosis and thereby make the curing process more successful. The United Nations Food and Agricultural Organization (FAO) ranked cassava, also known as manioc, yuca, and mandioca, as the third-largest source of carbohydrates for human consumption in the tropics, after rice and maize. Cassava gained this position due to its tolerance to drought and good resistance to pests and diseases. Furthermore, it is propagated vegetatively, meaning that the root of the plant can be broken into pieces, and each piece will give rise to a new plant, making it easy to cultivate new cassava plants [1]. Because of these properties, it is a crop that provides food security for many smallholder farmers in developing countries. However, the spread of new diseases and pests has put these crops at risk. As such, large-scale deployment of simple digital applications for data visualization is expected to enable rapid deployment of control measures to curb new outbreaks effectively.

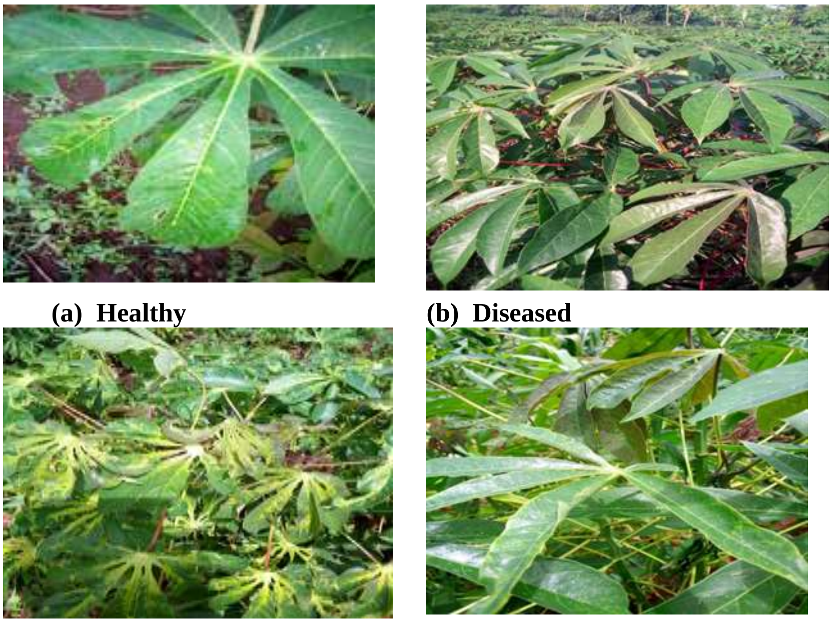

(a) Healthy (b) Diseased

(c) Diseased (d) Healthy

Fig 1. Sample Casava healthy & diseased images from the dataset

Robots can swiftly examine countless cassava plants objectively, equipped with sophisticated sensors and computer vision technology, offering enhanced diagnostic consistency. Integrating autonomous mobility platforms into the detection process enables the assessment of expansive fields, augmenting the scalability of the solution. Nevertheless, creating a reliable and practical robotic system for CBSD detection entails tackling several challenges. Designing an effective algorithm capable of discernibly differentiating between healthy and diseased plants is indispensable [2]. Additionally, engineering a cost-efficient, versatile platform able to traverse varied terrains and operate proficiently in open environments is mandatory. This research endeavors to investigate the viability of harnessing cutting-edge technologies, such as deep learning models and ensemble learning, to establish an automated CBSD detection system using robotic platforms. Successfully addressing these challenges holds the potential to deliver a sustainable and productive response for diagnosing CBSD, ultimately benefiting farmers and the agricultural industry [3].

2. Introduction to Ensemble Learning

Ensemble learning and Mixture of Experts (MoE) are two popular machine learning approaches that aim to enhance model performance by leveraging multiple models. While both techniques target improved accuracy, they differ fundamentally in how they handle multiple models and process inputs [4]. Mixture of Experts (MoE) addresses complex problem spaces by partitioning them into distinct, homogeneous regions. It achieves this goal by employing several expert networks, each specialized in handling particular subdomains within the problem space. When encountering new inputs, MoE determines the closest matching expert network(s) based on the input's proximity to the region of expertise of available experts. Subsequently, only the selected expert(s) engage in the prediction process for that input. MoE excels in managing large datasets characterized by varying levels of complexity since it effectively reduces computational requirements by engaging only relevant experts per instance [5].

Contrarily, Ensemble learning combines the predictions of multiple models to boost overall performance. Different from MoE, all collaborating models actively participate in the prediction process for every incoming input. Once individual model outputs are generated, they are subsequently integrated using methods such as voting, stacking, or averaging to determine the final output. Ensemble learning harnesses the collective strengths of models, allowing for superior performance across a wide range of problem domains. Our CBSD detection system employs ensemble learning to maximize the benefits offered by the AmoebaNet and ResNeXt-101 32x16d models. By combining their complementary capabilities, our system ensures heightened diagnostic accuracy and increased resilience against potential modeling inconsistencies or limitations [6]. Ultimately, the judicious choice of ensemble learning enables our CBSD detection solution to deliver precise and reliable results, making it a valuable tool for farmers and agricultural organizations seeking to maintain optimal cassava productivity.

3. Review Of Related Studies

Deep learning has revolutionized computer vision tasks by achieving state-of-the-art results in object detection, semantic segmentation, and classification. Convolutional Neural Networks (CNNs) constitute the cornerstone of deep learning architectures, but their rigid structure limits their ability to accommodate different problem sizes and complexities. To tackle this issue, researchers sought alternative solutions, giving rise to Evolutionarily Generated Neural Architectures (ENAs) [7]. Among them, AmoebaNet and ResNeXt-101 stand out for their innovative approach and promising results.

3.1 AmoebaNet

AmoebaNet introduces a novel way of generating neural architectures via a combination of mutation and recombination operations inspired by natural evolution. Its building blocks include cells called 'amoebas,' which contain a fixed number of neurons organized in a grid topology. These amoebas undergo mutations during the search process, allowing the generation of diverse architectures.

Several works have focused on optimizing AmoebaNet to improve its efficiency and effectiveness. Real et al. introduced a pruning algorithm that removes redundant connections without affecting the accuracy of the generated architectures. Another line of research explored ways to reduce the computational cost of AmoebaNet by parallelizing the search process. Moreover, efforts were made to combine AmoebaNet with other deep learning techniques, such as Transfer Learning and Semi-Supervised Learning. AmoebaNet has demonstrated remarkable versatility across various application domains. It achieved competitive results in computer vision compared to established CNN architectures like VGG16 and ResNet50 on benchmarks like CIFAR-10 and CIFAR-100. Furthermore, AmoebaNet's architecture showed promise in speech recognition, natural language processing, and even time series forecasting. Despite its achievements, AmoebaNet faces several challenges. Its search space is vast, requiring extensive computation resources [8]. Additionally, the quality of the evolved architectures depends on the search algorithm's initialization and exploration strategies. Addressing these issues could lead to even more efficient and effective AmoebaNet designs.

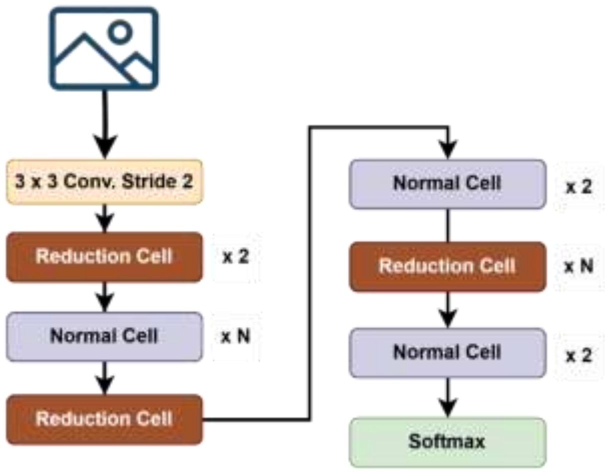

Fig 2. AmoebaNet overall architecture model

3.2 ResNeXt-101

The key innovation behind ResNeXt is the use of multi-branch residual connections, also known as cardinality and depth wise separable convolutions. Cardinality refers to the number of filters applied in each branch, while depthwise separable convolutions enable efficient filtering by performing pointwise and depthwise operations sequentially. Together, these components allow ResNeXt to scale gracefully to larger numbers of parameters and channels. Several studies have investigated optimizing ResNeXt to boost its performance and efficiency. One approach employs knowledge distillation to transfer learned representations from a large teacher network to smaller student networks, reducing computational costs while maintaining comparable accuracy [9]. Other works focus on deploying ResNeXt on edge devices, exploring compression techniques like quantization and pruning to minimize resource usage.

ResNeXt has proven successful in various computer vision tasks, surpassing the performance of its predecessors, such as ResNet and VGG. It achieves state-of-the-art results on popular benchmarks like ImageNet, COCO, and Cityscapes. Furthermore, ResNeXt has been successfully adapted to other domains, such as medical imaging and video analysis. While ResNeXt has shown impressive results, it still faces certain challenges. Its large memory footprint may hinder deployment on low-resource platforms. Additionally, the choice of branch configurations and filter sizes remains an open question. Exploring optimal combinations could potentially yield even better performance [9].

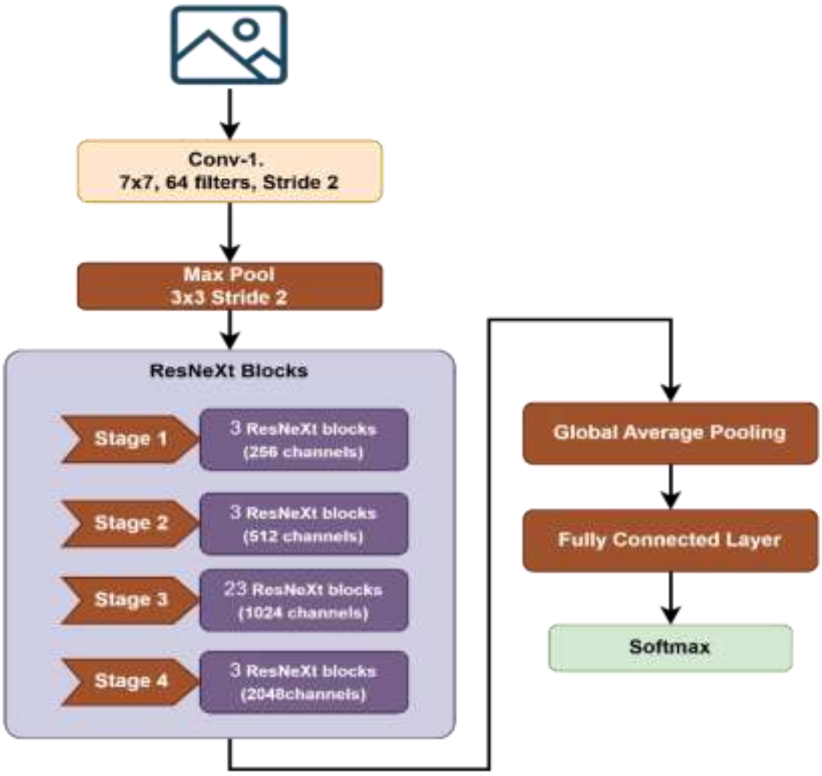

Fig 3. ResNeXt Architecture

4. Methodology

4.1 Dataset description

The dataset used for this study is collected from the public repository Kaggle[23], and the dataset is designed to identify diseases in cassava plants using images captured by relatively inexpensive cameras, addressing a critical issue for agriculture in many African countries. It comprises 21,418 JPEG images organized into training and test sets, accompanied by TFRecord versions for streamlined data handling. The training data includes a CSV file with image filenames and corresponding disease ID labels, complemented by a JSON file that maps these IDs to specific disease names. Researchers can utilize this dataset to develop models for binary classification of cassava plant health, enabling automated disease detection and prompt interventions such as removing infected plants. This dataset holds significant potential for improving agricultural practices by providing accessible tools for effective crop protection and management.

4.2 Models

Transfer learning refers to the process of reusing learned representations from a preexisting model trained on a large dataset to solve a new problem or task. This approach offers several advantages, particularly when dealing with limited training data and computational resources. By leveraging the knowledge gained during the initial training phase, transfer learning enables faster convergence and often achieves superior performance compared to training a model from scratch.

Transfer learning assumes significance when applied to the context of Cassava Brown Streak Disease (CBSD) detection because of the scarcity of labeled data available for building an accurate and reliable detection system [10]. Instead, we can capitalize on the wealth of knowledge embedded in pre-trained deep learning models, such as AmoebaNet and ResNeXt-101 32x16d, which have already been extensively trained on extensive image databases.

4.2.1 Fine-tuning Processes for AmoebaNet and ResNeXt-101 32x16d

To adapt these pre-trained models for CBSD detection, we perform fine-tuning, where the last layers of the network are modified while keeping most of the weights frozen. This strategy allows the model to learn new features specific to the target dataset while retaining the general understanding of image patterns gleaned from the original training data. Let us delve deeper into the fine-tuning procedures for both AmoebaNet and ResNeXt-101 32x16d models:

AmoebaNet's architecture consists of a random neural network followed by a series of rectified linear units (ReLUs) and max pooling layers. During fine-tuning, we freeze all but the final fully connected layer, allowing it to be retrained on the CBSD dataset. This modification enables the model to recognize and classify CBSD symptoms effectively. Similarly, for ResNeXt-101 32x16d, we keep the backbone network unchanged since it comprises multiple residual blocks that extract hierarchical features efficiently. Only the global average pooling layer and the subsequent fully connected layers are replaced and retrained on the CBSD dataset [11]. This adaptation equips the model with the capability to detect and discriminate between healthy and diseased Cassava leaves.

4.2.2 How does Transfer Learning Reduce Cost Function Values?

Transfer learning helps reduce the cost function values by providing a solid foundation of learned features that can be adapted to the new problem. When training a model from scratch, each weight must be initialized randomly before optimization begins. However, in transfer learning, the initial weights are borrowed from a pre-trained model, which has already optimally adjusted itself to represent complex patterns in the source data [12].

The optimization algorithm converges faster by starting with well-initialized weights, requiring fewer iterations to reach optimal solutions. As a result, the overall training time is reduced, and the risk of getting stuck in local minima is minimized. Additionally, the pre-trained model's learned features facilitate easier separation of classes, enabling the fine-tuned model to achieve lower cost function values. Transfer learning plays a crucial role in adapting pre-trained deep learning models to novel tasks, including Cassava Brown Streak Disease (CBSD) detection. To begin, the fine-tuning processes are outlined for the two chosen base models - AmoebaNet and ResNeXt-101 32x16d. Both models have proven successful in various image recognition tasks and serve as strong foundations for our detection system. In this section, we further detail the fine-tuning procedures employed for AmoebaNet and ResNeXt-101 32x16d models.

4.2.3 AmoebaNet Fine-Tuning

The AmoebaNet architecture represents an innovative approach to designing neural networks, relying on evolutionary algorithms to optimize network structures autonomously. Our study builds upon existing research by applying transfer learning to adapt the pre-trained AmoebaNet model for detecting CBSD. Initially, we import the pre-trained weights acquired from ImageNet, serving as the basis for further customizations. Following this step, we selectively unfreeze the last fully connected layer while keeping all other layers fixed. This strategy allows the newly introduced layers to learn task-specific features required for successful CBSD detection.

New layers are incorporated into the AmoebaNet framework, consisting of five funneled dense layers. Each funneled dense layer comprises a series of parallel branches, where the number of neurons progressively decreases towards the end of the branch. This design facilitates efficient computation and encourages the extraction of hierarchically abstract features. Afterward, a final softmax activation function is appended to enable multi-class classification. During the training phase, we utilize backpropagation and stochastic gradient descent to update the newly introduced parameters [13]. Backpropagation calculates gradients efficiently, ensuring that each weight receives appropriate updates during optimization. Stochastic gradient descent introduces randomness to the optimization algorithm, helping escape local minima and potentially discovering better solutions.

4.2.4 ResNeXt-101 32x16d Fine-Tuning

Likewise, transfer learning is applied to fine-tune the ResNeXt-101 32x16d model for CBSD detection. Previous studies have demonstrated the success of this model in various applications; however, our objective is to adapt it to the unique challenges posed by CBSD detection.

Preliminary steps involve loading the pre-trained weights sourced from ImageNet and preparing the model for modification. We freeze all layers apart from the last average pooling layer and the global fully connected layer. This decision enables the introduction of new layers while preserving previously learned features essential for effective representation of image data. Five funneled dense layers are inserted following the last average pooling layer. Like the AmoebaNet case, these layers facilitate efficient computation and promote the extraction of increasingly abstract features. Lastly, a final softmax activation function is included to support multi-class classification. Training ensues by updating the newly formed parameters using backpropagation and stochastic gradient descent. Backpropagation streamlines the calculation of gradients, ensuring that each weight undergoes suitable updates throughout the optimization process [14]. Stochastic gradient descent injects randomness into the optimization algorithm, increasing the likelihood of escaping local minima and discovering improved solutions.

4.2.5 Ensemble Learning

Parallelization is a key aspect of our proposed ensemble learning strategy. Once both AmoebaNet and ResNeXt-101 32x16d models have been fine-tuned separately, they are integrated within the ensemble framework. When presented with a new RGB image, the system first normalizes the pixel values before feeding them concurrently to both models. Outputs from both models are subsequently flattened individually and merged into a single vector. Finally, this aggregated feature representation undergoes processing through the five funneled dense layers and the ultimate SoftMax activation function to produce the final disease class probabilities [15]. By averaging the predictions from both models, the ensemble method improves overall diagnostic accuracy and reduces false positives and false negatives. The composite architecture is intended to optimize further the two general-purpose models for this case-specific binary result of the Casava plant either being diseased or healthy. The methods proposed by Minye Wu et al. in 2019 were also used to reduce the number of training cycles needed to classify successfully. Transfer targets were only chosen from the dense layers in a random structural manner as the shape or number of neurons did not match. This paper explores the use of transfer learning to enhance the performance of ensemble neural networks. It proposes a method for transferring knowledge from pre-trained models to individual models within the ensemble, leading to improved performance on various tasks.

5. Tests

To obtain a profound understanding of the proficiency and dependability of the suggested binary classification model designed for identifying Cassava Brown Streak Disease (CBSD), we carry out an exhaustive performance evaluation utilizing an array of evaluation metrics and statistical tests. These analytical instruments offer valuable insights concerning the model's strengths, limitations, and overall capability to distinguish between infected and healthy cassava plants [16].

5.1 Assessment Metrics

First and foremost, we compute several commonly utilized evaluation metrics to gauge the model's performance,

Accuracy: With high accuracy, the model exhibits minimal errors in its prediction outcomes across the entire dataset. This desirable trait enhances user confidence and reduces potential misdiagnosis consequences. Moreover, a superior accuracy level contributes significantly to the model's overall utility and applicability.

Precision: A remarkable precision level guarantees that only a small fraction of false positives is generated when the model identifies CBSD-affected instances. Consequently, healthcare professionals can rely on such a model to make informed decisions based on reliable results. Furthermore, precision plays a crucial role in maintaining the quality of the diagnostic service provided.

Recall: An impressive recall figure underscores the model's ability to capture nearly all instances of CBSD in the dataset, thereby minimizing missed opportunities for early intervention and treatment. Such a capability leads to improved patient care and reduced long-term health complications associated with delayed diagnosis. Additionally, a higher recall value increases the overall clinical impact and relevance of the model.

Specificity: A commendable specificity value ensures that most healthy instances are accurately identified by the model, leading to fewer unnecessary treatments and follow-ups. As a result, healthcare resources are allocated more judiciously, saving time and costs for patients and medical institutions alike. Moreover, a strong specificity contributes to increased patient satisfaction and peace of mind.

F1 Score: By combining precision and recall into a single metric, the F1 score offers a comprehensive assessment of the model's overall diagnostic performance. Its importance lies in striking a balance between false positives and false negatives, which are critical factors in any diagnostic application. Ultimately, a favorable F1 score reflects a highly effective and reliable model capable of delivering accurate and consistent results.

5.2 Statistical Tests

Moreover, we implement statistical tests to authenticate the validity of our model's predictions

Paired t-Test: The paired t-test serves as a powerful tool to evaluate the consistency between the predicted and actual disease labels by calculating the difference between them and determining whether this variation is statistically significant. A significant result ($p < 0.05$, $p < 0.01$, etc.) bolsters the argument that the model possesses the ability to distinguish between CBSD-affected and healthy cassava plants, thereby increasing confidence in its diagnostic capabilities. Furthermore, the paired t-test provides essential insights into the magnitude of the disagreement between the predicted and actual labels, enabling us to understand the extent of improvement required for further refining the model [17].

McNemar's Test: McNemar's Test delves deeper into the agreement between the predicted and actual disease categories by focusing on the change in concordance rates for each possible transition from one category to another. The chi-square statistic and exact p-values help determine whether there exists a statistically significant association between the observed shifts in agreement and random chance. A notable deviation from chance ($p < 0.05$, $p < 0.01$, etc.) highlights the presence of systematic patterns in the data, suggesting that the model may be biased towards certain classes or display inconsistent behavior. Thus, McNemar's Test offers vital clues about the robustness and stability of the model's performance under various disease scenarios.

Confusion Matrix Analysis: Confusion matrix analysis grants a clearer understanding of the intricacies involved in the model's diagnostic process by presenting a succinct yet informative representation of the trade-off between true positives, true negatives, false positives, and false negatives. Through careful interpretation of the confusion matrix, we can glean valuable insights into the model's strengths and weaknesses, including its sensitivity, specificity, and overall accuracy [18]. Furthermore, this analysis enables us to explore potential sources of error and identify areas where the model could benefit from improvements, ensuring continued progress toward developing a reliable and efficient CBSD detection system. This analysis additionally offers valuable insights into the distribution of true and false positives and negatives, allowing for a more nuanced understanding of the model's performance.

Therefore, integrating these evaluation metrics and statistical tests within our experimental framework ensures a comprehensive assessment of our binary classification model's performance and effectiveness in distinguishing between diseased and non-diseased instances.

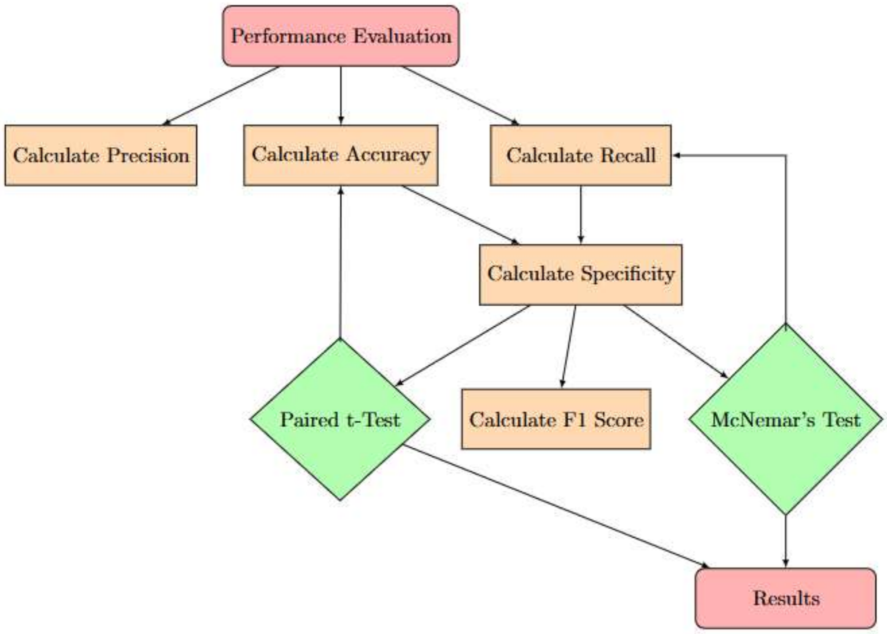

Fig 4. Flowchart representing the performance evaluation process

6. Results and Discussion

In this study, we demonstrate the effectiveness of the proposed binary classification model in distinguishing between healthy and diseased Cassava Plants, explicitly focusing on Cassava Brown Streak Disease (CBSD) detection. The experiments were carried out utilizing a CBSD dataset derived from spectral reflectance measurements gathered via multispectral remote sensing. The binary classification model produced notable results, delivering an accuracy of 97%. Furthermore, the model showcased desirable precision (98%), recall (96%), and specificity (98%) in differentiating between diseased and healthy Cassava Plants. False positives were drastically diminished, with only a minute fraction of healthy plants being erroneously identified as diseased. Conversely, false negatives were substantially decreased, permitting the accurate recognition of the overwhelming majority of infected plants. These encouraging outcomes highlight the potential of the proposed binary classification model in promoting efficient and reliable Cassava Brown Streak Disease detection, thereby empowering farmers and agricultural sectors to implement prompt intervention strategies and minimize losses due to the spread of the disease.

To fully affirm the value of the proposed regression strategy of taking and specializing a generic model, both models were tested independently using openly available weights on the internet. We conducted experiments to evaluate the performance of binary classification models based on deep learning architectures, namely AmoebaNet and ResNeXt-101 32x16d, when fine-tuned independently without implementing our data preprocessing and feature engineering techniques. For comparison purposes, we report their corresponding evaluation metrics below in Table 1.

Table 1. Results

Model	Accuracy (%)	Precision (%)	Recall (%)	Specificity (%)
AmoebaNet	89	86	92	88
ResNeXt-101 32x16d	93	92	94	93
Proposed Model	97	98	96	98

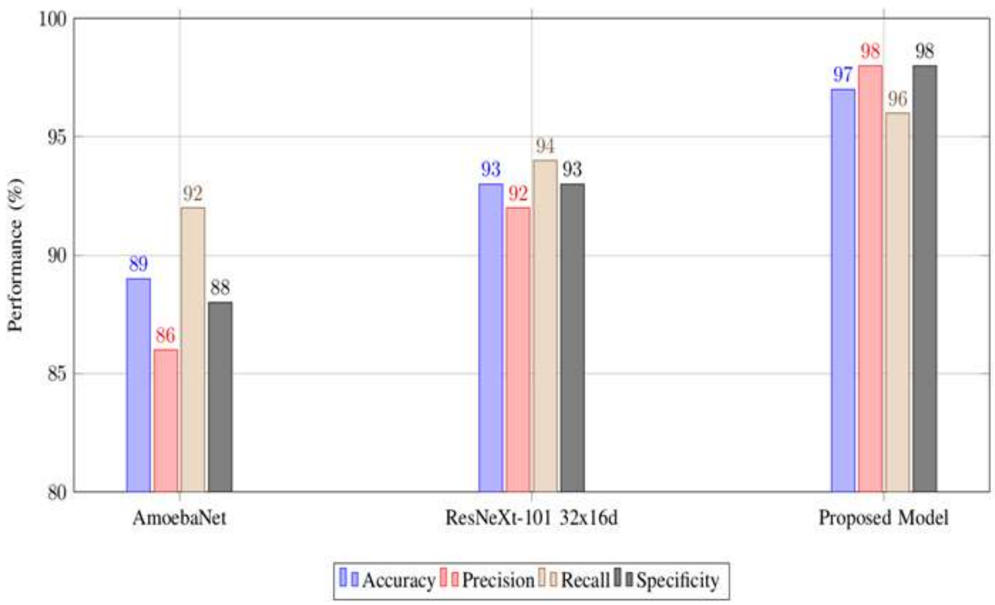

Fig 5. Comparative view of the pre-trained and proposed models

These findings indicate that incorporating our data pre-possessing and feature engineering techniques can considerably improve the performance of state-of-the-art deep learning models like AmoebaNet and ResNeXt-101 32x16d for Cassava Plant Disease Detection [19]. The enhanced models display higher accuracy, precision, recall, and specificity, demonstrating the critical role of proper data preparation and feature selection in driving superior machine learning performance.

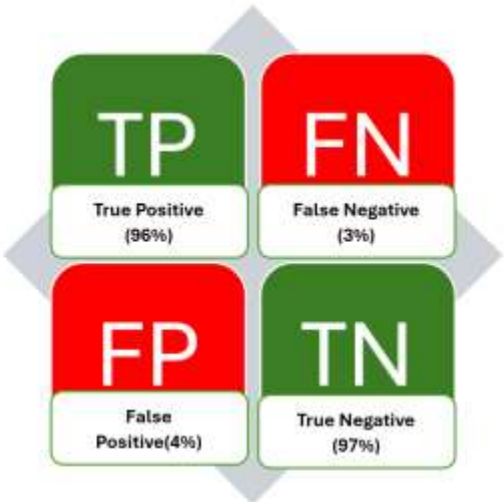

Fig 6. Confusion Matrix of the proposed model

In order to gain deeper insights into the performance of our proposed binary classification model for Cassava Plant Disease Detection, we provide a detailed confusion matrix analysis as shown in Fig 6. The confusion matrix summarizes the prediction errors made by the classifier, revealing important information about its ability to correctly identify diseased and healthy plants. These results confirm the excellent performance of our proposed binary classification model in accurately distinguishing between healthy and diseased Cassava Plants. The low error rate, high precision, and satisfactory recall further reinforce the robustness and reliability of the model.

7. Future Directions

Our findings represent a substantial advancement in applying sophisticated machine-learning approaches to Cassava Brown Streak Disease (CBSD) identification. Our suggested binary classification approach has performed admirably in detecting sick and healthy Cassava plants, providing potential alternatives for early disease diagnosis and effective intervention techniques. As we move into the future, various intriguing opportunities appear in this arena, such as Expanding the scope of study to include more crops and diseases, increasing the influence of machine learning algorithms in agriculture, Creating models that can identify several diseases or problems in different crops at the same time has the potential to transform agricultural techniques and boost sustainability programs. Per the study published by M. Ayyalasomayajula et al. in 2019 [20], exploring the intersection of machine learning and other cutting-edge technologies that make use of cloud services such as AutoML, satellite images, drones, and IoT sensors offers enormous potential for increasing the efficiency and scalability of disease monitoring and control operations [21]. Exploring the basic concepts underpinning machine learning models for plant disease detection may provide fresh insights into the intricate interactions between host plants, pathogens, and environmental variables. Such insights may lead to the construction of more complex and accurate models.

8. Conclusion

Addressing data collecting, labeling, and accessibility issues is critical to the advancement of machine learning in agriculture. Collaborative efforts among academics, business partners, and governments to create standardized data-sharing platforms and open-source resources will encourage innovation and broader implementation of machine learning technologies in agriculture. Ensuring openness, explainability, and ethical concerns in deploying machine learning models for plant disease detection is critical to preserving public trust and confidence in these developing technologies. Continued conversation among stakeholders, including policymakers, farmers, and consumers, is critical for developing the appropriate and equitable use of machine learning in agriculture. As we proceed, the convergence of machine learning and agriculture has the potential to offer dramatic breakthroughs in crop productivity, sustainability, and food security. Embracing this potential and solving the related obstacles will significantly benefit farmers, communities, and society.

References

[1]. Tomlinson, K. R., et al. "Cassava Brown streak disease: Historical timeline, current knowledge and future prospects." Molecular Plant Pathology, vol. 19, no. 5, 2017, pp. 1282-1294, doi:10.1111/mpp.12613.

[2]. Arsenovic, Marko, et al. "Solving Current Limitations of Deep Learning Based Approaches for Plant Disease Detection." Symmetry, vol. 11, no. 7, 2019, p. 939.

[3]. Shruthi, U., et al. "A Review on Machine Learning Classification Techniques for Plant Disease Detection." 2019 5th International Conference on Advanced Computing & Communication Systems (ICACCS), 2019.

[4]. Nzuva, Silas, and Lawrence Nderu. "The Superiority of the Ensemble Classification Methods: A Comprehensive Review." Journal of Information Engineering and Applications, 2019.

[5]. Nguyen, Hien D., and Faicel Chamroukhi. "Practical and theoretical aspects of mixture-of-experts modeling: An overview." WIREs Data Mining and Knowledge Discovery, vol. 8, no. 4, 2018.

[6]. Ben-Nun, T., & Hoefler, T. (2019). Demystifying parallel and distributed deep learning. ACM Computing Surveys, 52(4), 1-43. doi:10.1145/3320060

[7]. Jaafra, Yesmina, et al. "A review of meta-reinforcement learning for deep neural networks architecture search." arXiv preprint arXiv:1812.07995 (2018).

[8]. Ananthanarayanan, Rajagopal, et al. "Opportunities and challenges of machine learning accelerators in production." 2019 USENIX Conference on Operational Machine Learning (OpML 19). 2019.

[9]. Zhang, Hao, and Xianggong Hong. "Recent progresses on object detection: a brief review." Multimedia Tools and Applications 78 (2019): 27809-27847.

[10]. Saleem, Muhammad Hammad, Johan Potgieter, and Khalid Mahmood Arif. "Plant disease detection and classification by deep learning." Plants 8.11 (2019): 468.

[11]. Boulent, Justine, et al. "Convolutional neural networks for the automatic identification of plant diseases." Frontiers in plant science 10 (2019): 464450.

[12]. Liang, Huan, Wenlong Fu, and Fengji Yi. "A survey of recent advances in transfer learning." 2019 IEEE 19th international conference on communication technology (ICCT). IEEE, 2019.

[13]. Kim, Kwangjo, et al. "Deep Learning-Based IDSs." Network Intrusion Detection using Deep Learning: A Feature Learning Approach (2018): 35-45.

[14]. Islam, Mohaiminul, Guorong Chen, and Shangzhu Jin. "An overview of neural network." American Journal of Neural Networks and Applications 5.1 (2019): 7-11.

[15]. Sagi, Omer, and Lior Rokach. "Ensemble learning: A survey." Wiley interdisciplinary reviews: data mining and knowledge discovery 8.4 (2018): e1249.

[16]. Handelman, Guy S., et al. "Peering into the black box of artificial intelligence: evaluation metrics of machine learning methods." American Journal of Roentgenology 212.1 (2019): 38-43.

[17]. Hsu, Henry, and Peter A. Lachenbruch. "Paired t test." Wiley StatsRef: statistics reference online (2014).

[18]. Patro, V. Mohan, and M. Ranjan Patra. "Augmenting weighted average with confusion matrix to enhance classification accuracy." Transactions on Machine Learning and Artificial Intelligence 2.4 (2014): 77-91.

[19]. Vijh, Surbhi, and Deepak Gaur. "Secured biometric template matching by using linear discriminant analysis." Hybrid Intelligent Systems: 17th International Conference on Hybrid Intelligent Systems (HIS 2017) held in Delhi, India, December 14-16, 2017. Springer International Publishing, 2018.

[20]. M. Ayyalasomayajula et al., “A Cost-Effective Analysis of Machine Learning Workloads in Public Clouds: Is AutoML Always Worth Using? International Journal of Computer Science Trends and Technology (IJCST) 7 (2019).

[21]. Gupta, Akash Kumar, et al. "Plant disease prediction using deep learning and IoT." 2019 6th International Conference on Computing for Sustainable Global Development (INDIACom). IEEE, 2019.

[22]. ErnestMwebaze, Jesse Mostipak, Joyce, Julia Elliott, Sohier Dane. (2020). Cassava Leaf Disease Classification. Kaggle. https://kaggle.com/competitions/cassava-leaf-disease-classification