Artificial Cognition for Early Leaf Disease Detection using Vision Transformers

Huy-Tan Thai*, Nhu-Y Tran-Van†, Kim-Hung Le*
University of Information Technology, Vietnam National University Ho Chi Minh City
Ho Chi Minh, Vietnam
Email: *{tanth, hunglk}@uit.edu.vn
, †17521287@ms.uit.edu.vn

***Abstract*—There are many kinds of cassava leaf diseases firmly harm cassava yield, including four main types as followings: Cassava Bacterial Blight (CBB), Cassava Brown Streak Disease (CBSD), Cassava Green Mottle (CGM), and Cassava Mosaic Disease (CMD). In a traditional way, leaf diseases were diagnosed intuitively by farmers. This process is inefficient and unreliable. Several studies have recently relied on deep neural networks for identifying leaf diseases. In this research, we exploit the novel model named Vision Transformer (ViT) in place of a convolution neural network (CNN) for classifying cassava leaf diseases. Experimental results show that this model can obtain competitive accuracy at least 1% higher than popular CNN models (EfficientNet, Resnet50d) on Cassava Leaf Disease Dataset. These results also indicate the potential superiority of the ViT over established methods in analyzing leaf diseases. Next, we quantize the original model and successfully deploy it onto the Edge device named Raspberry Pi 4, which can be attached to a drone that allows farmers to automatically and efficiently detect infected leaves. This result has a significant capability for many future applications in smart agriculture.**

***Index Terms*—smart agriculture, leaf diseases, vision transformer, Raspberry Pi**

I. INTRODUCTION

In Africa, cassava is the second-largest source of carbohydrates. There are about 300 million Africans who consume cassava as primary nutrition every day. It is also a stable source of income for subsistence farmers due to the drought and pest-resistant characteristics. However, many plant diseases, especially leaf diseases, significantly affect cassava production quantity and quality. These diseases open a new challenge in the smart agriculture area, requiring considerable work for early and effectively detecting infected cassava leaves. Consequently, the profit of the cassava industry has been guaranteed.

With the development of computer vision in recent years, it is now possible to analyze leaf diseases without the farmers' observation, reducing time consumption, increasing accuracy, and protecting crops in time. Many studies are carried for this purpose, mainly divided into machine learning and deep learning-based approaches. However, the machine learning-based methods [1]–[4], generally require complex preprocessing and specific disease feature extraction before analyzing diseases. For solving these problems, many deep learning models are applied with high accuracy in plant disease detection, such as Faster Region-based Convolutional Neural Network (F-RCNN) [5], [6], SSD with Inception module, and Rainbow concatenation (INAR-SSD) [7], Deep Residual Dense Network [8], ResNet [9], EfficientNet [10]. Nevertheless, these deep models usually require substantial computational resources during the training phase. This limitation motivates the release of Vision Transformer [11] - a novel deep learning model which consumes significantly fewer arithmetical calculations to train than previous models. In Vision Transformer (VIT), a pure transformer encoder is directly utilized to sequences of image patches to accomplish image classification responsibilities. The experiments in [11] show that ViT achieves comparative results with the state-of-the-art CNNs while the computational cost is exceptionally reduced in the training phase. As far as we know, this paper is the first research that exploits the Vision Transformer (ViT) model for cassava leaf disease classification. Our experiments show excellent results compared to modern models such as EfficientNet, Resnet50 on Cassava Leaf Disease Dataset, published by Makerere University and National Crops Resources Research Institute. This is consistent with the findings of [11]–[13], which suggest that ViT can perform very well in image classification tasks. The main contributions of this paper are listed below:

- We exploit ViT for detecting infected leaves. We first use the ViT model pre-trained on ImageNet - 21k [14], published by Google Research Team. Then, we retrain this model using Cassava Leaf Disease Dataset [15] to improve the model's accuracy for analyzing infected leaves task. Subsequently, the model is quantized for reducing the size and accelerating the inference step.
- We successfully deploys the quantized model onto the Edge device named Raspberry Pi 4 Model B. This device can be attached to a fly-cam to form the Drone Pi for real-time detection of infected leaves. This success will pave the way for early leaf disease detection, an essential part of smart agriculture.

The rest of the paper is organized as follows: Section II discusses related works using traditional machine learning and deep learning for leaf diseases detection, Section III presents our proposed approach, Section IV shows the result and discussion, followed by the conclusion in Section V.

II. RELATED WORKS

A. Machine learning approach

In the last decade, several traditional machine learning methodologies (e.g., K-Means clustering and support vector (SVM)) have been employed to identify and detect disease on plant leaves. The authors in [3] propose a model using image processing and SVM classifier for detecting diseases on cotton leaves. Yao et al. [4] use a similar approach to detect diseases on rice leave early and precisely. The presented SVM is capable of effectively classify various rice diseases (rice bacterial leaf blight, rice sheath blight, and rice blast) reported about 97% accuracy. Waghmare et al. [2] exploit a unique fractal to retrieve the segmented leaf texture and multiclass SVM for detecting common grape diseases. [1] introduce a novel image processing method and an artificial neural model to diagnose the disease of brinjal leaf. The authors apply K-Means clustering to extract appropriate features as the input of the model. However, the presented machine learning approaches require multiple preprocessing and extraction steps which are unsuitable for real-time detection or tasks requiring high performance.

B. Neural Network approach

Artificial neural network (ANN) exploits the combination of texture and color features to perform the leaf disease classification task. For example, in [16], [17], the authors use K-Means to extract image features before using these vectors for training the ANN classifier. Similarly, Sammany and Medhat [18] apply neural networks enhanced by genetic algorithms and support vector machines (SVM) to diagnose plant diseases. The study in [19] uses the Gabor filter to extract features and an ANN classifier to classify healthy and diseased samples. Recently, the authors in [20] utilize a hybrid metaheuristic feature selection and feed-forward neural network to identify mango leaves diseases. However, these ANN classifiers have low performance in some complex classification categories. Recent years have seen a rise in the number of deep learning studies, including image classification, object detection, and segmentation tasks [21]. With the continual advances in deep learning, many researchers aim to optimize deep neural networks for classifying crop diseases. For instance, the authors in [5] apply Faster Region-based Convolutional Neural Network (F-RCNN) to detect and recognize tomato plant leaf disease. Subsequently, Zhang et al. [6] improve F-RCNN by replacing VGG16 with a depth residual network resulting in 2.71% higher recognition accuracy compared with previous work. An optimized CNNs named INAR-SSD (SSD with Inception module and Rainbow concatenation) is proposed in [7] for real-time detecting apple leaf diseases. In [8], the authors attempt to restructure Deep Residual Dense Network for classifying Tomato leaf diseases within fewer parameters but higher accuracy than the original model.

III. PROPOSED APPROACH

A. System overview

As a result of recent advances in IoT, many industries, including agriculture, have been disrupted. IoT is an indispensable factor for improving agricultural quantity and quality at a lower cost. In the following years, the smart solution's usage powered by IoT will increase in the agriculture industry. As potential representatives, ground and aerial drones are widely utilized to evaluate crop health, crop monitoring, planting, crop spraying, and field analysis in smart agriculture. This paper proposes a system based on the Drone Pi, which is combined from a drone and a Raspberry Pi, for the early detection of infected leaves as shown in Fig.1. In this system, the Drone Pi is utilized to assess crop health via detect disease affecting leaves by the following steps:

- The Drone Pi's camera captures cassava leaf images, including the exact position of the spot in yield.
- A deep learning model named ViT, installed on the Drone Pi, is applied to classify and cluster infected leaves.
- The ViT classification's results, combine with the spot's position, are real-time send to the server via the 4G network to create a survey map about leave diseases of the cassava field. The farmers and rescue agency can instantaneously obtain this map via mobile devices and noticed the crop health problems beforehand, which prevents the high amount of loss or in some cases crop failure.

B. Cassava Leaf Disease Detection

As the central heart of the system, our Cassava Leaf Disease Detection mechanism exploits the Vision Transformer model to identify infected leaves.

The Transformer, a transduction model based entirely on self-attention mechanisms, has played an essential role in various computer vision tasks. The authors in [22] propose this model as a simple network architecture but exceptionally effective in machine translation. Then, many studies attempt to apply Transformer to different natural language processing (NLP) tasks ranging from syntactic analysis to sematic analysis [23]–[29]. Based on NLP tasks' achievements, many researchers have recently integrated self-attention mechanisms and Transformer models with Convolution Neuron Networks to create a hybrid model to enhance vision tasks' performance [30]–[36].

In contrast, the authors in [11] propose a full version of the transformer model named Vision Transformer (ViT) that reaches state-of-the-art performance on the image classification task. ViT is pre-trained on large datasets (e.g., JFT300M) and then transferred to multiple mid-sized or small image recognition benchmarks. In this paper, we utilize this model for classifying on Cassava Leaf Disease Dataset as shown in Fig.2, including the following steps:

- Splitting the cassava leaf image into fixed-size patches. The dimension of each patch can be 16x16 or 32x32.

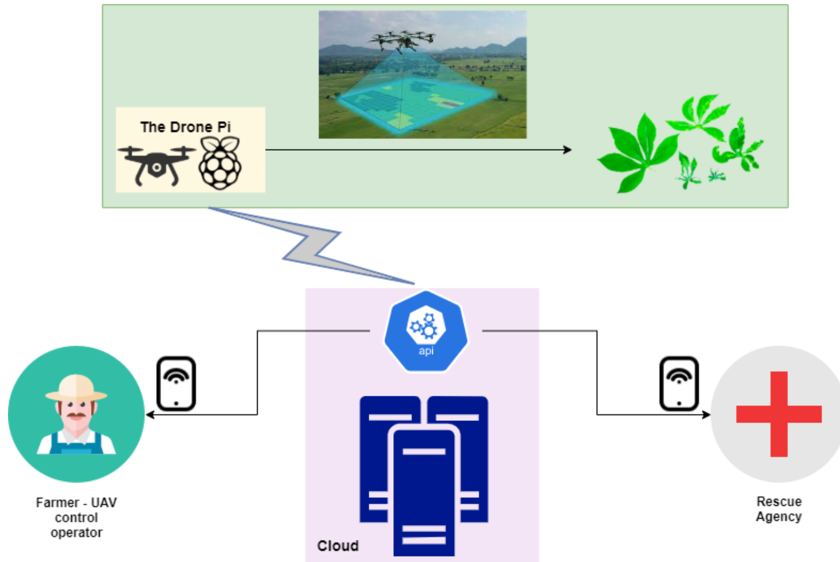

Fig. 1: Smart Agriculture Scenario for Cassava Plants

- All image patches are flattened and added with the positional embedding vector. A special token is also added at the start.
- All embedded patches are put into Transformer Encoder.
- The previous step's output is passed directly to Feed Forward Neural Network to classify and cluster images into five categories: CBB, CBSD, CGM, CMD, and Healthy.

Concretely, our Cassava Leaf Disease Detection mechanism includes two stages: training and inference stage, as shown in Fig.3. In the former stage, the pre-trained ViT model is retrained using Cassava Leaf Disease Dataset for classifying infected leaves. The output model is then quantized to reduce the size (from about 220MB to 84MB) for compatibility to the edge device with limited resources. Then, the quantized model is deployed onto the Drone Pi to perform classifying and clustering infected leaf images task in the inference stage.

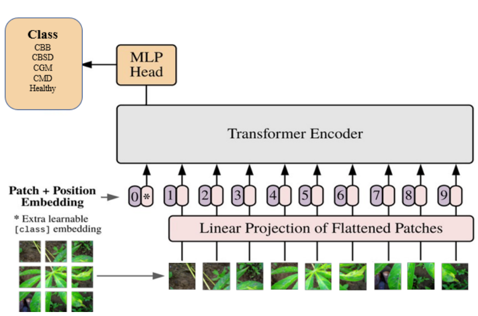

Fig. 2: ViT model for Cassava Leaves Disease Detection

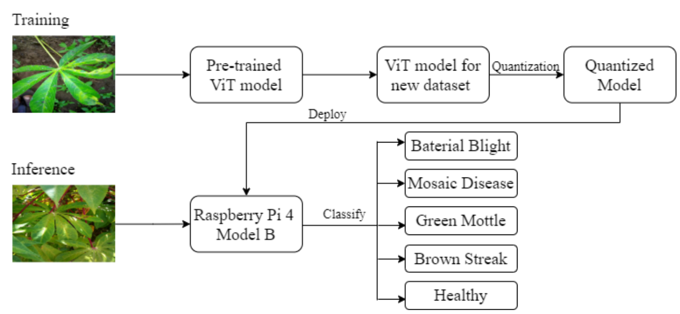

Fig. 3: Cassava Leaf Disease Detection mechanism

IV. RESULTS AND DISCUSSION

A. Dataset

This paper uses a dataset named Cassava Leaf Disease Dataset. This dataset includes 21397 images of cassava leaves divided into five categories: CBB, CBSD, CGM, CMD, and healthy leaves. Most images are taken from farmers and annotated by experts in Makerere University and National Crops Resources Research Institute. Fig. 4 depicts the representative of each category. Images are then resized to 224x224 as the requirement of the ViT model. To experiment, we use the k-fold cross-validation method with k = 5. The ratio of the training dataset and validation dataset is 4:1. The detail of training/validation images of each disease is shown in Table I.

B. Exprerimental enviroment and evaluation/measurement metrics

The model has been deployed and verified on Raspberry Pi 4 Model B with the specifications are shown in Table II. We use the F1 - score as the evaluation metric, which is the harmonic mean of precision and recall.

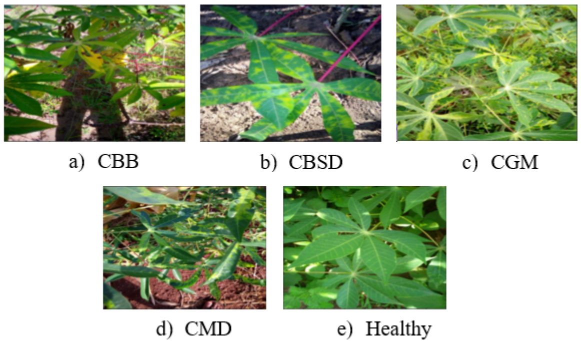

a) CBB b) CBSD c) CGM d) CMD e) Healthy

Fig. 4: Five common types of cassave leaf diseases

TABLE I: Cassava leaf disease dataset

Label	Disease	Training	Validation	Total Images
0	CBB	870	217	1087
1	CBSD	1752	437	2189
2	CGM	1908	478	2386
3	CMD	10526	2632	13158
4	Healthy	2061	516	2577
		17117	4280	21397

$$F1 = \frac{2 * (precision * recall)}{(precision + recall)} \tag{1}$$

$$precision = \frac{TP}{(TP + FP)} \tag{2}$$

$$recall = \frac{TP}{(TP + FN)} \tag{3}$$

TP (True Positive) represents the number of positive images classified correctly. FP (False Positive) denotes the number of negative images classified incorrectly. FN (False Negative) represents the number of undetected positive images. The model was trained on ImageNet - 21k by Google researchers. Then, we retrain and evaluate this model on the Cassava Leaf Disease Database.

C. Results

Table III shows the result of the evaluation per category, including precision, recall, and F1-score. As can be seen, the highest score belongs to the classifying CMD task with a 96% F1-score. The second-highest F1-score belongs to the identification of CBSD task, followed by CGM, Healthy with 86%, 83%, and 79%, respectively. The lowest score belongs to the CBB disease classification with a 74.5% F1-score. This disease has the lowest score due to its significant within-class variations: first, symptoms are visible as translucent water-soaked spots; then, these translucent spots become dark green spots; finally, spots expand, nearby spots merge to create large brown patches. Meanwhile, the scores of other infected categories are higher than 83% due to the minor differences among lesions' patterns from the same category. In addition, the variety of lesions' appearances between categories is substantial. Therefore, these infected categories can be effectively classified by our approach.

TABLE II: Experimental enviroment

Equipment	Specifications
Name	Raspberry Pi 4 Model B
System	Ubuntu 18.04
Framework	torch 1.9.0, torch vision 0.9.0
Language	Python 3.7.3
CPU	Quad core Cortex-A72 (ARM v8) 64-bit SoC @1.5GHz
RAM	8GB LPDDR4-3200 SDRAM

Fig. 5a depicts the training and validation loss through 165 epochs. After 165 epochs, the training and validation loss scores decrease to an optimal value around 0.2 and 0.4, respectively, and converge. As shown in Fig. 5b, when the epoch value increases from 1 to 20, the training accuracy grows sharply from 81% to 88%, and validation accuracy rises from 77% to 84%. Afterward, these values fluctuate upward before reaching the highest value at 95% refer to the training accuracy, and 90% refer to the validation accuracy. From these results, it is clear that ViT model is fast convergent over 165 epochs with high accuracy.

Table IV illustrates the comparison with the state-of-the-art models on the cassava leaf disease dataset. The experimental results of these comparison models are taken from [37] through 165 epochs. Regarding to EfficientNet models, EfficientNetB3 achieves 88.1% weighted-average F1-score while EfficientNetB4 and EfficientNetB5 reach 88.3% and 88.7% weighted-average F1-score, respectively. Finally, EfficientNetB6 attains the highest score of Efficient models with an 89.1% weighted-average F1-score. Meanwhile, the Resnet50 model achieves an 89.2% weighted-average F1-score. Our system uses the Vision Transformer model, which obtains the overall highest score, 90.3% weighted-average F1-score. In addition, the results from all comparison models also show that the CBB classification score is the lowest through the experiments. Hence, it can be concluded that CBB is the most complex category in Cassava Leaf Disease Dataset. The experiment results highly indicate again that the ViT model can achieve competitive scores compared with state-of-the-art models in the classifying cassava leaf diseases task.

V. CONCLUSION

This paper builds a system that relies on the ViT model to identify infected leaves with full promising results. We then successfully quantize the model to reduce the model size by the factor of three and deploy it onto the Raspberry Pi 4 Model B. In the future, we intend to attach this device to a

TABLE III: Classification performances of each category

Category	Precision	Recall	F1-score
CBB	**79**	**70**	**74.5**
CBSD	87	85	86
CGM	81	85	83
CMD	**96**	**96**	**96**
Healthy	78	80	79
overall	90	90	90

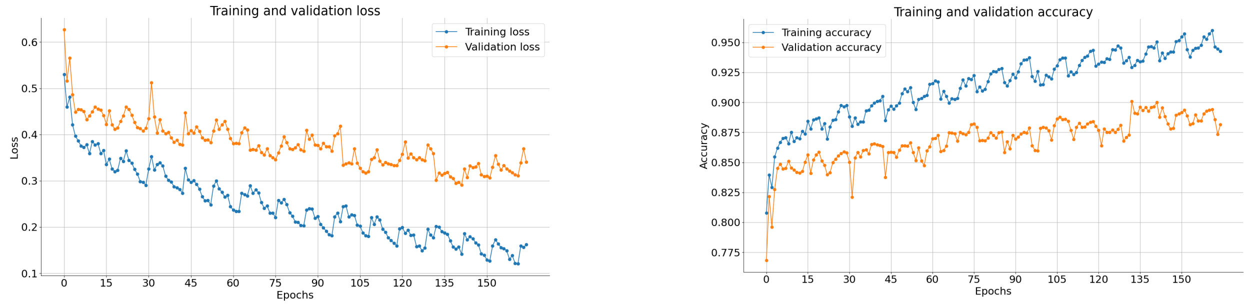

(a) Training, validation loss

(b) Training, validation accuracy

Fig. 5: Training, validation loss and accuracy through 165 epochs of our approach

drone to directly analyze and identify the infected leaves from the camera. This work will help detect early diseased plants, potentially protecting their yield before causing irreparable damage, making a small contribution to developing smart agriculture.

REFERENCES

[1] R Anand, S Veni, and J Aravinth. An application of image processing techniques for detection of diseases on brinjal leaves using k-means clustering method. In *2016 international conference on recent trends in information technology (ICRTIT)*, pages 1–6. IEEE, 2016.

[2] Harshal Waghmare, Radha Kokare, and Yogesh Dandawate. Detection and classification of diseases of grape plant using opposite colour local binary pattern feature and machine learning for automated decision support system. In *2016 3rd international conference on signal processing and integrated networks (SPIN)*, pages 513–518. IEEE, 2016.

[3] Namrata R Bhimte and VR Thool. Diseases detection of cotton leaf spot using image processing and svm classifier. In *2018 Second International Conference on Intelligent Computing and Control Systems (ICICCS)*, pages 340–344. IEEE, 2018.

[4] Qing Yao, Zexin Guan, Yingfeng Zhou, Jian Tang, Yang Hu, and Baojun Yang. Application of support vector machine for detecting rice diseases using shape and color texture features. In *2009 international conference on engineering computation*, pages 79–83. IEEE, 2009.

[5] Robert G. de Luna, Elmer P. Dadios, and Argel A. Bandala. Automated image capturing system for deep learning-based tomato plant leaf disease detection and recognition. In *TENCON 2018 - 2018 IEEE Region 10 Conference*, pages 1414–1419, 2018.

[6] Yang Zhang, Chenglong Song, and Dongwen Zhang. Deep learning-based object detection improvement for tomato disease. *IEEE Access*, 8:56607–56614, 2020.

[7] Peng Jiang, Yuehan Chen, Bin Liu, Dongjian He, and Chunquan Liang. Real-time detection of apple leaf diseases using deep learning approach based on improved convolutional neural networks. *IEEE Access*, 7:59069–59080, 2019.

[8] Changjian Zhou, Sihan Zhou, Jinge Xing, and Jia Song. Tomato leaf disease identification by restructured deep residual dense network. *IEEE Access*, 9:28822–28831, 2021.

[9] Ümit Atila, Murat Uçar, Kemal Akyol, and Emine Uçar. Plant leaf disease classification using efficientnet deep learning model. *Ecological Informatics*, 61:101182, 2021.

[10] V. Kumar, Hritik Arora, Harsh, and Jatin Sisodia. Resnet-based approach for detection and classification of plant leaf diseases. *2020 International Conference on Electronics and Sustainable Communication Systems (ICESC)*, pages 495–502, 2020.

[11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *CoRR*, abs/2010.11929, 2020.

[12] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, et al. Swin transformer: Hierarchical vision transformer using shifted windows. *CoRR*, abs/2103.14030, 2021.

TABLE IV: Comparison with state-of-the-art models on cassava leaf disease detection - 165 epochs

Model \ Category	**CBB**	**CBSD**	**CGM**	**CMD**	**Healthy**	**Overall**
Resnet50	66	83	82	96	74	89.2
EfficientNetB3	63	**80**	81	96	**71**	**88.1**
EfficientNetB4	**61**	82	81	96	75	88.3
EfficientNetB5	64	83	**80**	96	75	88.7
EfficientNetB6	66	82	82	96	75	89.1
OurModel	**74.5**	**86**	**83**	**96**	**79**	**90.3**

[13] Li Yuan, Yunpeng Chen, Tao Wang, Weihao Yu, Yujun Shi, Francis E. H. Tay, et al. Tokens-to-token vit: Training vision transformers from scratch on imagenet. *CoRR*, abs/2101.11986, 2021.

[14] J. Deng, W. Dong, R. Socher, L. Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops (CVPR Workshops)*, pages 248–255, Los Alamitos, CA, USA, jun 2009. IEEE Computer Society.

[15] Makerere Artificial Intelligence (AI) Lab. Cassava leaf disease dataset. Available at https://www.kaggle.com/c/cassava-leaf-disease-classification/data.

[16] Aakanksha Rastogi, Ritika Arora, and Shanu Sharma. Leaf disease detection and grading using computer vision technology fuzzy logic. In *2015 2nd International Conference on Signal Processing and Integrated Networks (SPIN)*, pages 500–505, 2015.

[17] Usha Kumari, S. Prasad, and G. Mounika. Leaf disease detection: Feature extraction with k-means clustering and classification with ann. pages 1095–1098, 03 2019.

[18] Mohammed Sammany and T. Medhata. Dimensionality reduction using rough set approach for two neural networks-based applications. *International Conference on Rough Sets and Intelligent Systems Paradigms*, 2007.

[19] A. Kulkarni. Applying image processing technique to detect plant diseases. In *International Journal of Modern Engineering Research (Volume: 2, Issue: 5, Sept. 2012)*, page 3661–3664, 2012.

[20] Tan Nhat Pham, Ly Van Tran, and Son Vu Truong Dao. Early disease classification of mango leaves using feed-forward neural network and hybrid metaheuristic feature selection. *IEEE Access*, 8:189960–189973, 2020.

[21] Yann LeCun, Y. Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521:436–44, 05 2015.

[22] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, et al. Attention is all you need. *CoRR*, abs/1706.03762, 2017.

[23] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. *CoRR*, abs/1810.04805, 2018.

[24] Nikita Kitaev and Dan Klein. Constituency parsing with a self-attentive encoder. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2676–2686, Melbourne, Australia, July 2018. Association for Computational Linguistics.

[25] Alec Radford and Karthik Narasimhan. Improving language understanding by generative pre-training. 2018.

[26] Matthew E. Peters, Mark Neumann, Robert Logan, Roy Schwartz, Vidur Joshi, Sameer Singh, and Noah A. Smith. Knowledge enhanced contextual word representations. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 43–54, Hong Kong, China, November 2019. Association for Computational Linguistics.

[27] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, et al. Roberta: A robustly optimized BERT pretraining approach. *CoRR*, abs/1907.11692, 2019.

[28] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. Xlnet: Generalized autoregressive pretraining for language understanding. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.

[29] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, et al. Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc., 2020.

[30] Prajit Ramachandran, Niki Parmar, Ashish Vaswani, Irwan Bello, Anselm Levskaya, and Jonathon Shlens. Stand-alone self-attention in vision models. *Computer Vision and Pattern Recognition*, abs/1906.05909, 2019.

[31] Han Hu, Zheng Zhang, Zhenda Xie, and Stephen Lin. Local relation networks for image recognition. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3463–3472, 2019.

[32] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7794–7803, 2018.

[33] Yuhui Yuan and Jingdong Wang. Ocnet: Object context network for scene parsing. *CoRR*, abs/1809.00916, 2018.

[34] Jun Fu, Jing Liu, Haijie Tian, Yong Li, Yongjun Bao, Zhiwei Fang, and Hanqing Lu. Dual attention network for scene segmentation. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3141–3149, 2019.

[35] Yue Cao, Jiarui Xu, Stephen Lin, Fangyun Wei, and Han Hu. Gcnet: Non-local networks meet squeeze-excitation networks and beyond. In *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, pages 1971–1980, 2019.

[36] Minghao Yin, Zhuliang Yao, Yue Cao, Xiu Li, Zheng Zhang, Stephen Lin, and Han Hu. Disentangled non-local neural networks. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 191–207, Cham, 2020. Springer International Publishing.

[37] Dimitre Oliveira. Cassava leaf disease classification. Available at https://github.com/dimitreOliveira/Cassava-Leaf-Disease-Classification.