ADVANCING IMAGE CLASSIFICATION THROUGH PARAMETER-EFFICIENT FINE-TUNING: A STUDY ON LORA WITH PLANT DISEASE DETECTION DATASETS

Deeksha Aggarwal, Yash Mittal & Uttam Kumar
International Institute of Information Technology Bangalore (IIIT Bangalore)
Bangalore, India.
{deeksha.aggarwal003,yash.mittal,uttam}@iiitb.ac.in

ABSTRACT

Low-Rank Approximation (LoRA) has demonstrated remarkable efficiency in Large Language Models (LLMs), enabling the attainment of state-of-the-art results across various Natural Language Processing (NLP) tasks. This study shifts the focus to the application of LoRA, a low-rank approximation technique, within the realm of image classification for plant disease detection. The experiments unveil a notable reduction in trainable parameters to less than 1%. Notably, LoRA surpasses traditional fine-tuning methods, achieving a state-of-the-art accuracy of 73.83% on the label-deficient Pdoc dataset, 99.89% on the PlantVillage dataset and 81.81% on FieldPlant dataset. This research provides valuable insights at the intersection of PEFT and real-world plant disease detection, establishing a foundation for future explorations across diverse domains.

1 INTRODUCTION

As the demand for efficient and accurate image classification models continues to grow, the importance of focusing on parameter-efficient fine-tuning (PEFT) methods becomes increasingly crucial, particularly in scenarios characterized by limited computational resources and labeled data. This paper delves into the exploration of a low-rank approximation technique, specifically LoRA (Low-Rank Approximation), within the realm of image classification for plant disease detection.

The safeguarding of crops against plant diseases plays a pivotal role in meeting the rising demand for food quality and quantity (Strange & Scott, 2005). Timely disease detection in plants remains a challenging task for farmers who often have limited resources and alternatives, such as consulting fellow farmers or helplines, and must possess expertise in plant diseases. In this study, we investigate the potential use of computer vision for scalable and cost-effective (in terms of computation and storage) plant disease detection through PEFT.

PEFT involves the adaptation of a pre-trained model in machine learning to perform a specific task with minimal computational resources and data. The primary goal is to achieve optimal performance on a new task while leveraging knowledge gained from the pre-trained model. This becomes especially crucial in scenarios where computational resources or labeled data for the target task are limited. To the best of our knowledge, this study represents the first exploration of LoRA's application to plant disease detection, adding a novel perspective to the ongoing discourse in this domain.

2 DATA SET DESCRIPTION

The PlantDoc (Pdoc) (Singh et al., 2020), Plant Village Dataset (PVD) (Mohanty et al., 2016) and FieldPlant Moupojou et al. (2023) serve as datasets for plant disease detection, comprising visual plant images. The Pdoc and FieldPlant dataset features photos of plants captured in real-world settings, whereas the PVD is a collection of plant leaf images taken in a controlled environment. These three datasets encompass a total of 2,598, 54,303, and 5156 images, featuring different plant species categorized into 28, 38, and 26 healthy and unhealthy classes, respectively. PlantDoc dataset presents a real-world and challenging image classification scenario with limited labels per class.

Table 1: Test classification accuracy for Full FT: Fine-tuning all model parameters, FT: Fine-tuning selective model parameters, LoRA: Proposed method.

Data Set	Backbone	Full FT	FT	LoRA
PVD	Resnet50	99.12%	73.26%	**97.31%**
	ViT Base	99.44%	97.52	**99.89%**
Pdoc	Resnet50	32.20%	9.70%	**56.54%**
	ViT Base	56.11%	53.16	**73.83%**
FieldPlant	Resnet50	76.84%	60.02%	**65.03%**
	ViT Base	87.32%	80.47%	**81.81%**

3 METHODOLOGY

Traditional fine-tuning typically either updates all model parameters (Full FT) or freeze all and unfreeze and update selected model parameters (FT), incurring high computational and storage costs. In contrast, LoRA fully freezes the pretrained model weights and injects trainable rank decomposition matrices into selected layer of the model architecture, greatly reducing the number of trainable parameters for downstream tasks. During fine-tuning, only these matrices are trained, while the original model parameters remain unchanged. During inference, the update matrices are combined with the original model parameters to generate the conclusive classification result. See Hu et al. (2021) for more information on LoRA. This approach not only significantly decreases computational and storage expenses but also maintains or even improves model performance.

4 EXPERIMENTS AND RESULTS

In our research, we utilized the low-rank approximation technique (LoRA) from the Hugging Face Transformers library for fine-tuning the ResNet-50 and ViT base models. LoRA significantly reduced the trainable parameters to 0.92% of the original model when the last block before the classifier layer of the resnet-50 model and attention blocks of the ViT model was selected to be updated. We compared LoRA with traditional fine-tuning (FT) methods, adjusting weights directly in two settings: full fine-tuning (Full FT) and fine-tuning only the selected parameters (FT). We conducted experiments using the Nvidia Quadro RTX 8000 GPU. Fine-tuning experiments were performed on models pre-trained on the ImageNet dataset, with a batch size of 128 and learning rates set to 0.001 for Full FT and FT and 5e-3 for LoRA, respectively and training epochs set to 100.

Table 1 presents a comparison of classification accuracy among three fine-tuning algorithms applied to the three datasets. Notably, LoRA demonstrates state-of-the-art (SOTA) classification accuracy compared to simple FT for all three datasets and attains the highest accuracy of 56.54% and 73.83% surpassing both full FT and FT for the Pdoc dataset.

5 CONCLUSION AND FUTURE WORK

This study explores image classification for plant disease detection using three open datasets: PlantDoc (Pdoc), Plant Village (PVD) and FieldPlant, emphasizing challenges in the real-world scenario. Leveraging PEFT with LoRA, we achieved a substantial reduction in trainable parameters to less than 1% of the original model. Through selective updates using low-rank “update matrices”, LoRA outperformed traditional fine-tuning, reaching a state-of-the-art accuracy on three challenging plant disease detection datasets. The results highlight the effectiveness of parameter-efficient low-rank adaptation techniques, showcasing not only high accuracy in resource-constrained environments but also superior performance when faced with limited labeled data for the target task such as Pdoc dataset. The study also contributes insights into the intersection of PEFT and real-world plant disease detection, opening avenues for future exploration across diverse domains and datasets. As future work, we plan to compare LoRA performance under cross domain fine-tuning with measurable catastrophic forgetting of the previous task to further establish our model’s state of the art results.

ACKNOWLEDGEMENTS

We are grateful to RRSC-South, National Remote Sensing Center (NRSC), Indian Space Research Organization (ISRO), Bengaluru for the research grant. IIIT Bangalore is acknowledged for the infrastructure support.

URM STATEMENT

The authors acknowledge that at least one key author of this work meets the URM criteria of ICLR 2024 Tiny Papers Track.

REFERENCES

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021.

Sharada P Mohanty, David P Hughes, and Marcel Salathé. Using deep learning for image-based plant disease detection. *Frontiers in plant science*, 7:1419, 2016.

Emmanuel Moupojou, Appolinaire Tagne, Florent Retraint, Anicet Tadonkemwa, Dongmo Wilfried, Hyppolite Tapamo, and Marcellin Nkenlifack. Fieldplant: A dataset of field plant images for plant disease detection and classification with deep learning. *IEEE Access*, 11:35398–35410, 2023.

Davinder Singh, Naman Jain, Pranjali Jain, Pratik Kayal, Sudhakar Kumawat, and Nipun Batra. Plantdoc: A dataset for visual plant disease detection. In *Proceedings of the 7th ACM IKDD CoDS and 25th COMAD*, pp. 249–253. 2020.

Richard N Strange and Peter R Scott. Plant disease: a threat to global food security. *Annu. Rev. Phytopathol.*, 43:83–116, 2005.

A APPENDIX

The code for this study can be found at `https://github.com/SCLIIITB/LoRAForPlants`