

Identification of Potato Virus Y in Potato Plants

using Deep Learning and GradCAM Verification

Roy Voetman

Supervisors: Willem Dijkstra, Klaas Dijkstra

Winter 2023

Introduction

- Infected plants could result in a large loss in crop production.
- Present-day research on classifying Potato Virus-Y (PVY) primarily employs hyperspectral data [1].
- This research
 - Adapt RGB based methodology for blight diseases from [2] to perform PVY classification using **VGG16** and **ResNet50**.
 - Applying Explainable AI techniques to validate if models focus on the same leaf characteristics as domain experts.

Materials and Methods

- Dataset of images from 12 healthy and 12 diseased plants.
- A total of 1106 leaves are annotated with polygons.
- Classifiers are trained on the individual leaves.

Abstract

Controlling potato plant diseases is crucial in agriculture because diseased plants can result in a significant loss in crop production. Farmers currently use manual visual inspections to classify potato plants that are diseased with the potato virus Y. This paper shows the feasibility of using RGB images to automate this process. A small dataset of potato plants pictures was acquired from an outdoor potato field and annotated with polygons to denote the location of the leaves. VGG16 and ResNet50, two convolutional neural networks, were trained to classify individual leaves. Additionally, experiments were conducted by changing the brightness of the images during training. This was done to counteract the different illumination levels present in our dataset. ResNet50's best performing experimental setup achieved an accuracy of 0.77, while VGG16 achieved an accuracy of 0.70. In general, classifiers with brightness augmentations outperformed other setups. Examining the class activation mappings of shallow layers revealed a focus on either the margins and midrib or the overall texture of the leaf. This corresponds to the characteristics used by domain experts who concentrate on colour and mosaic patterns on the foliage.



Figure 1: (a) Example of healthy foliage. (b) Example of PVY diseased foliage.

Materials and Methods

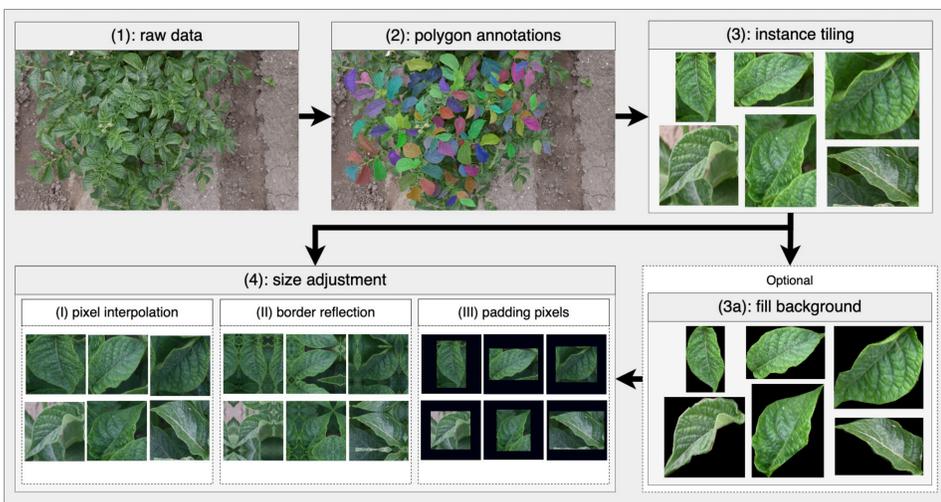


Figure 2: Data Processing Pipeline (1) Raw images (4000x6000) (2) Leaf annotations (3) Extracting the leaves. (3a) Optionally removing the background (4) Three possible algorithms to adjust the size are explored.

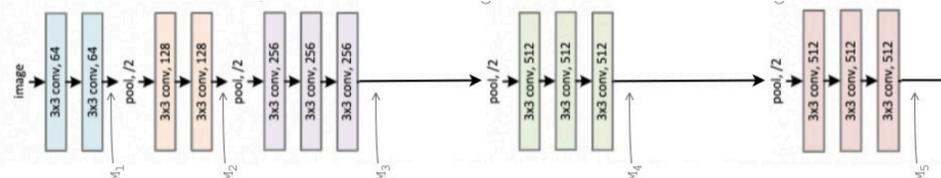


Figure 3. Class activation mappings are created at multiple points in the CNN architectures to evaluate fine-grained feature extraction capabilities. This figure shows the evaluation layers for the VGG16 architecture.

Experiments and Results

- Experiment 1: An ablation study was performed with ResNet50 and VGG16, a background removal procedure (Fig 2 step 3a), size adjustment methods (Fig 2 step 4), and brightness shift augmentations, for a total of 36 experimental setups.

	ResNet50		VGG16	
	Accuracy	F2	Accuracy	F2
No Value Shift				
Padding pixels + background removal	0.70	0.68	Pixel interpolation + background removal	0.65 0.78
Value Shift [-50, 50]	Accuracy	F2	Value Shift [-50, 50]	Accuracy F2
Pixel interpolation	0.77	0.82	Pixel interpolation + background removal	0.70 0.80

Table 1. The evaluation metrics for the best performing per-leaf classifiers on the test set

- Experiment 2: Aggregate individual leaf results to determine per-plant performance.

ResNet50	Targets: 0	0	0	1	1	1
Pixel interpolation	0	0.01	1	0.54	1	0.60
VGG16	Targets: 0	0	0	1	1	1
Pixel interpolation	0	0.06	1	0.77	1	0.76
				1	0.62	1
					1	0.59
						1
						0.56

Table 2. The evaluation metrics for the best performing per-plant classifiers on the test set

Experiments and Results

- Experiment 3: Evaluated feature extraction capabilities with class activation mappings (Fig 4) and visualisations of the embeddings (Fig 5) from the last convolutional layer.

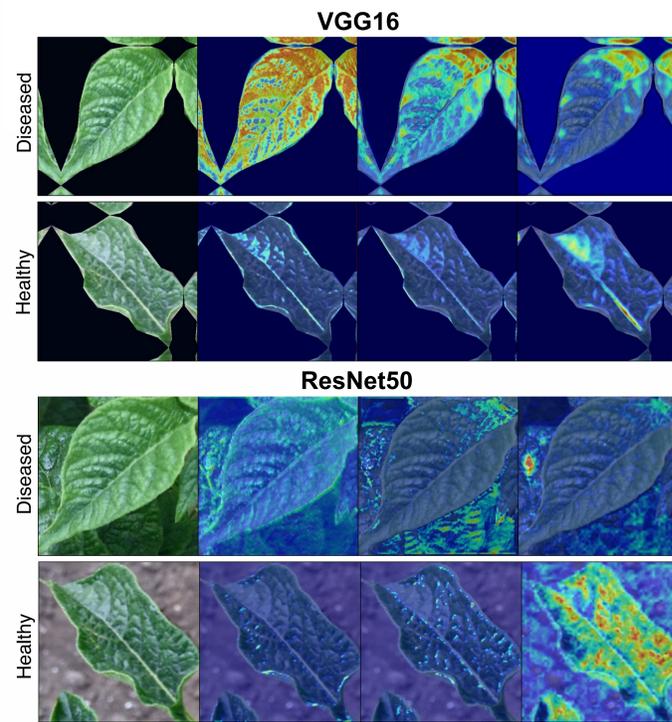


Figure 4. GradCAM outputs from the ResNet50 model with pixel interpolation and VGG model with border reflect + background removal. CAMs are computed at different layers (displayed from shallow to deep).

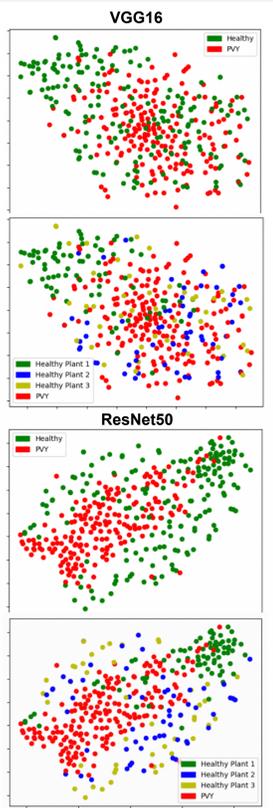


Figure 5. t-SNE distributions of the embeddings. The first plot shows a colour per target while the second plots specifies the associated (healthy) plant.

Conclusions

- It is feasible to classify PVY diseased plants with purely RGB data.
- ResNet50 models yield overall better performance than VGG16 models.
- Shifting brightness levels during training positively influences performance.
- While CNNs seem to focus on leaf texture details, a characteristic also used by domain experts, they are very dependent on the illumination level in the image.
- The performance will likely improve if the lighting conditions are always the same, or when the lighting conditions are balanced for all target classes.
- The labeling process could be improved by providing a label for each individual leaf.
 - We found that there is was a mismatch between ground truth labels when asking a second domain expert to classify individual leaves instead of entire plants.

References

- [1] G. Polder, P. M. Blok, H. A. C. De Villiers, J. M. Van der Wolf, and J. Kamp. Potato virus y detection in seed potatoes using deep learning on hyperspectral images. *Frontiers in plant science*, 10:209, 2019.
- [2] X. Li, Y. Zhou, J. Liu, L. Wang, J. Zhang, and X. Fan. The Detection Method of Potato Foliage Diseases in Complex Background Based on Instance Segmentation and Semantic Segmentation. *Frontiers in Plant Science*, 13, 7 2022.



computer vision & data science

Acknowledgements

This project is a collaboration between the Computer Vision & Data Science professorship and Smart Agri Technology B.V.

