Understanding Deep Neural Networks For Regression In Leaf Counting

Andrei Dobrescu
A.Dobrescu@ed.ac.uk

Mario Valerio Giuffrida
v.giuffrida@ed.ac.uk
valeriogiuffrida.academy

Sotirios A Tsaftaris
S.Tsaftaris@ed.ac.uk
Tsaftaris.com

phenotiki.com

THE UNIVERSITY of EDINBURGH

Computer Vision Problems in Plant Phenotyping (CVPPP) 2019

CVPR LONG BEACH CALIFORNIA June 16-20, 2019

Motivations:

Neural networks are typically seen as **'black boxes'**, lacking a straightforward explanation of how a network achieves a prediction. The visualization of the inner mechanism of a deep network provides a **human level understanding** of how the deep learning models make decisions and what image representations they have learned. Investigating the 'black box' will increase the confidence in deep learning predictions and will permit the redesign of the architecture to improve performance.

Contributions:

- We focus on visualizing methods for direct **regression problems**.
- We demonstrate experimentally that **edges of leaf blades are the most influential for count**. The petiole or the central parts of leaves are not taken into account in the decision making.
- We show that the regression value is predicted mainly using the **foreground object (i.e. the plant).**
- Based on investigating intermediate layers we can **compress the networ**k, while not significantly impacting the performance.

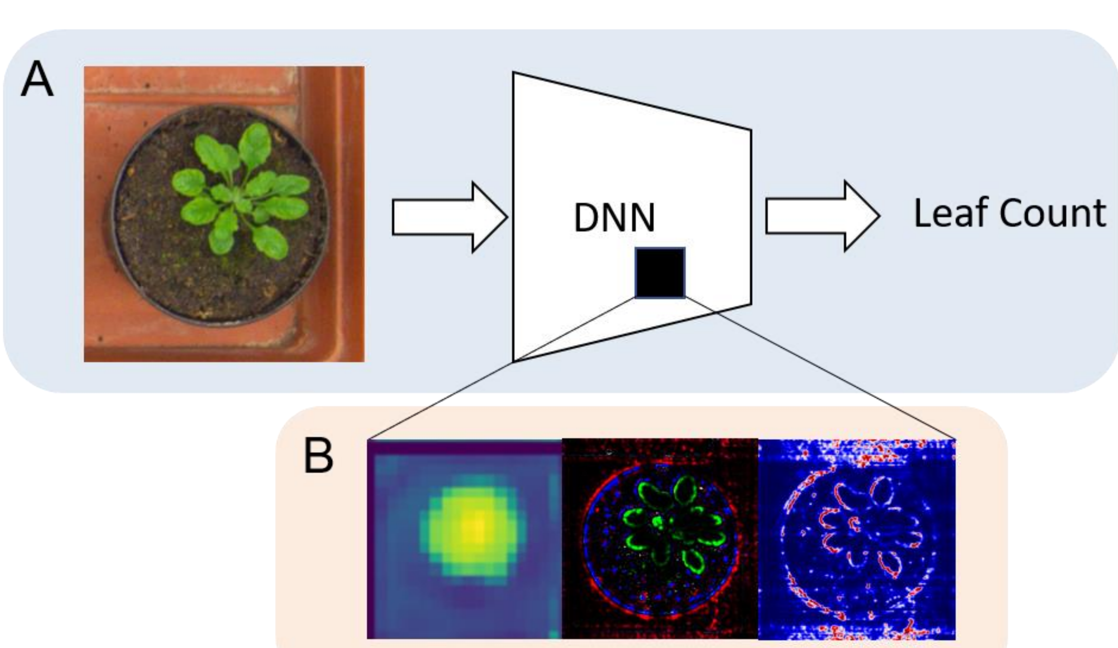

A: Common deep learning framework taking an image as input into a trained deep neural network (DNN) which outputs the leaf count. **B:** We investigate what elements of the input image contribute the most in computing a prediction and gain an understanding of the intermediate layers.

Visualization techniques

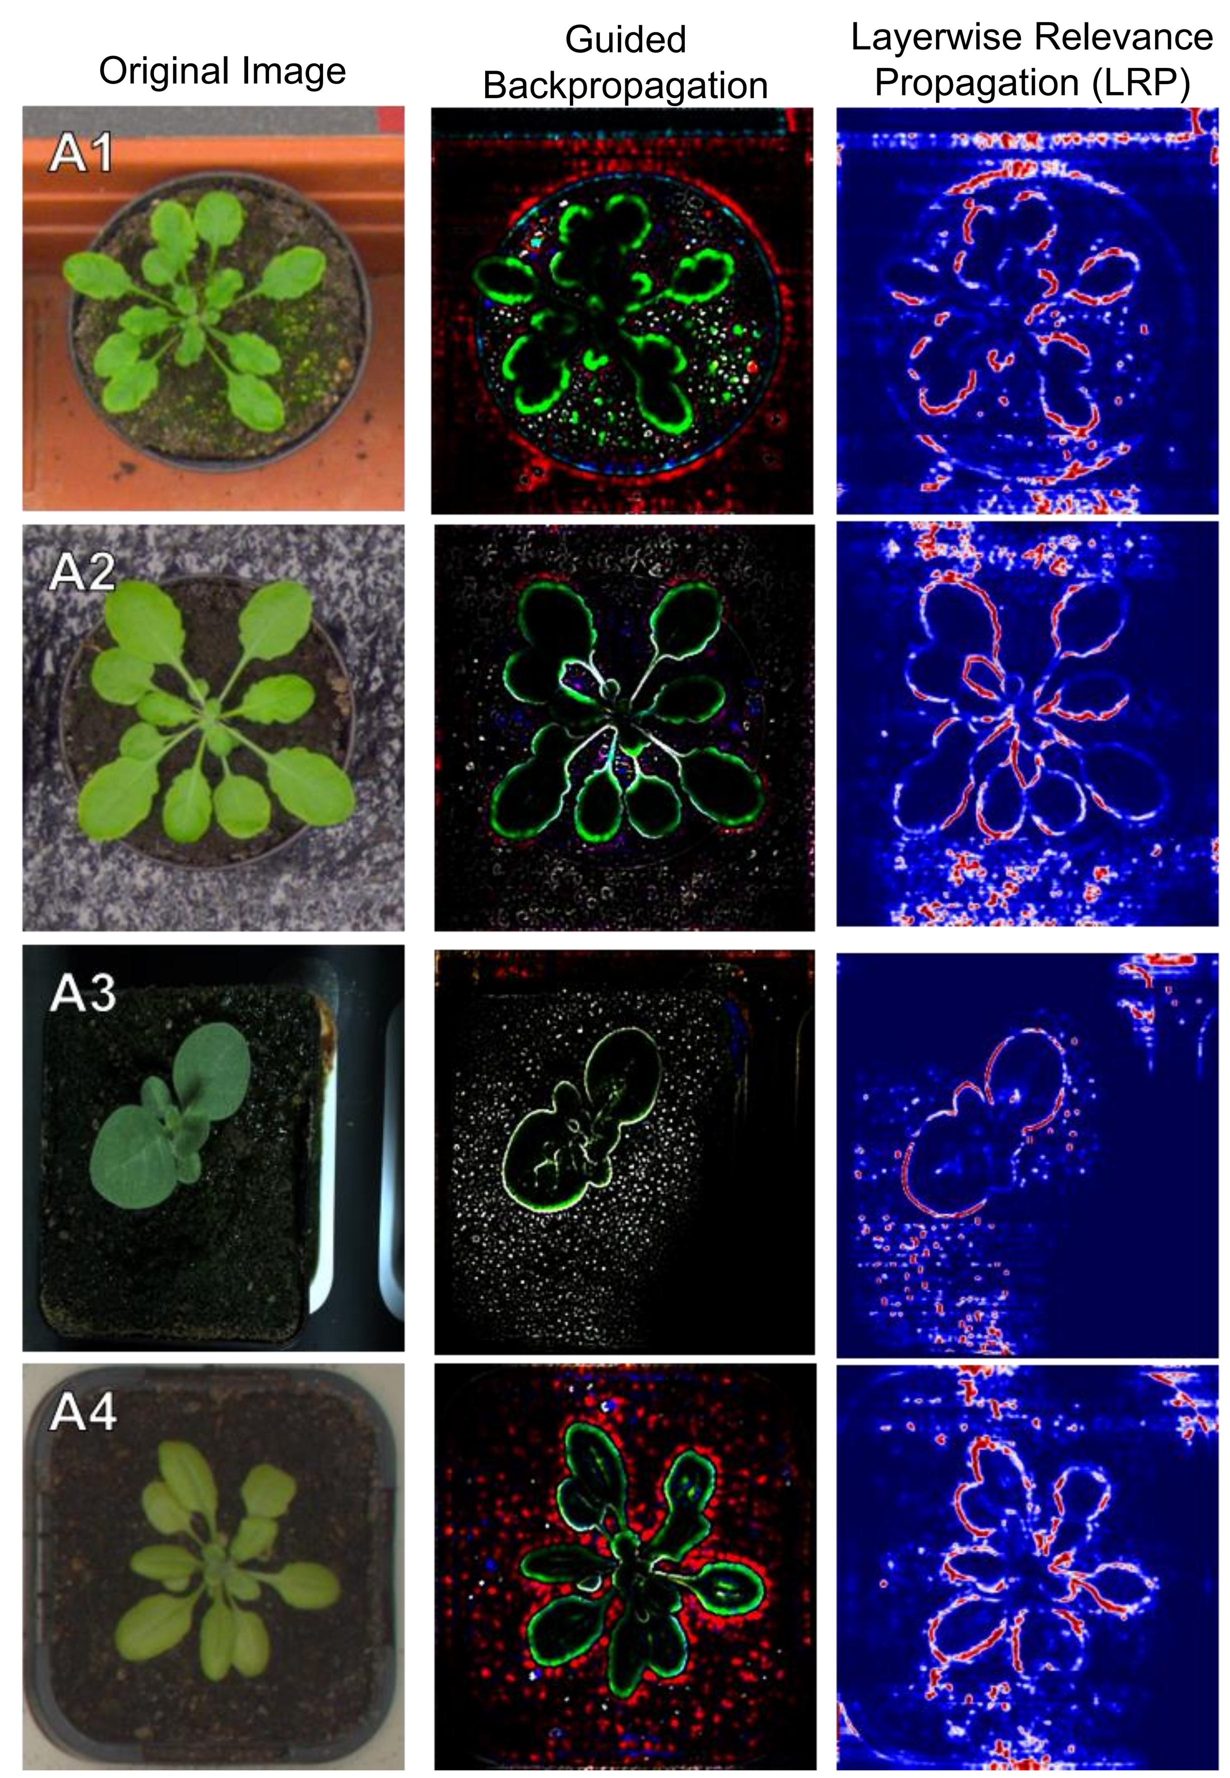

Guided Backpropagation:

A gradient-based visualization technique designed to highlight what parts of the input contribute to a given neuron in a neural network [1]. The method back propagates the gradient with relation to the input image while **masking negative values**.

Layerwise relevance propagation (LRP):

Each neuron receives a share of the network output and redistributes it to its predecessors in equal amount, until the input is reached [2]. The output of the LRP technique is a relevance heatmap highlighting **both positive and negative signals**.

Individual Neurons Contribution

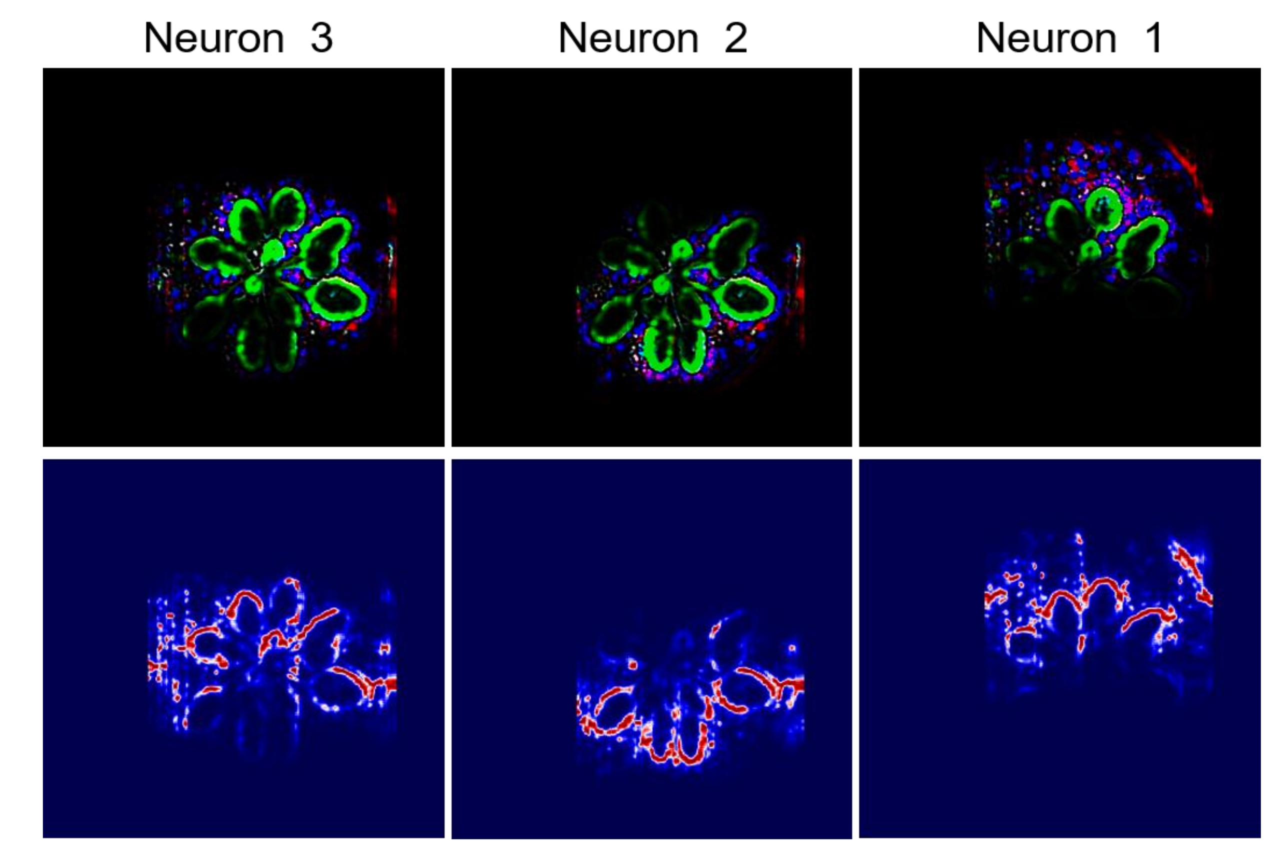

Each of these nodes acts as an 'attention map': **they get excited by different parts of the plant**, ignoring the rest of the image.

Compressing the network

Parameters	DiC	|DiC|	MSE	%
57M	0.11 (1.10)	0.74 (0.80)	1.20	41
36M	-0.21 (1.13)	0.79 (0.83)	1.33	41

We obtained a **reduction** of the total network parameters by **37%**, without significantly impacting accuracy.

Network activity before and after training

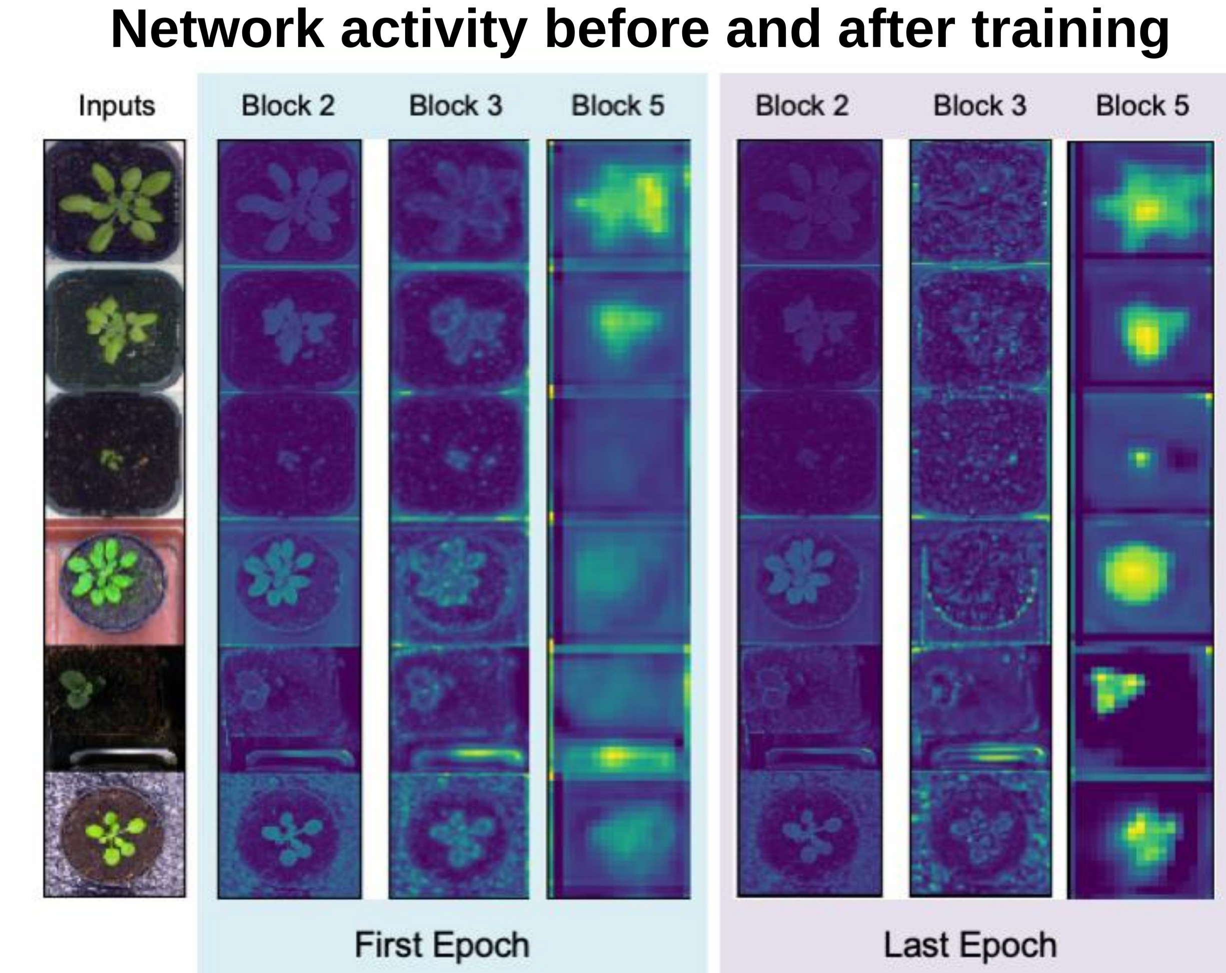

The average activations at start and end of training. The network learn to **focus on the plant** even with just the final leaf count as supervision.

What is important for count?

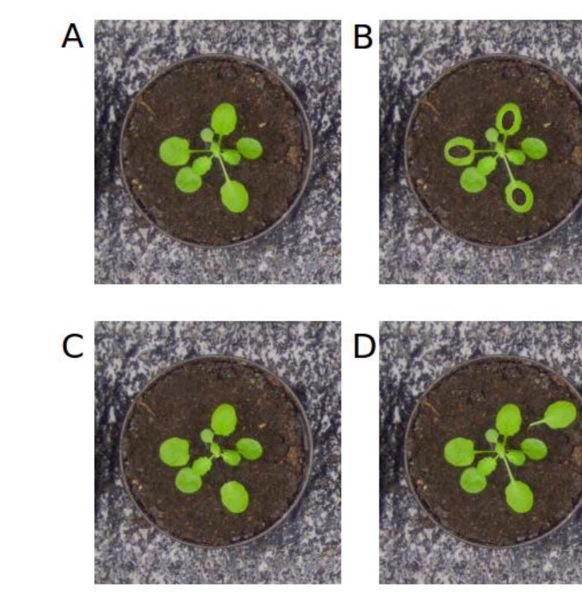

The leaf blade edge is the most important for computing count. **A:** The original image **B:** Leaves with deleted centres **C:** Leaves without petioles. **D:** Extra leaf was added.

The network predicts the **same leaf count** for images **B** and **C** as does for the original image.

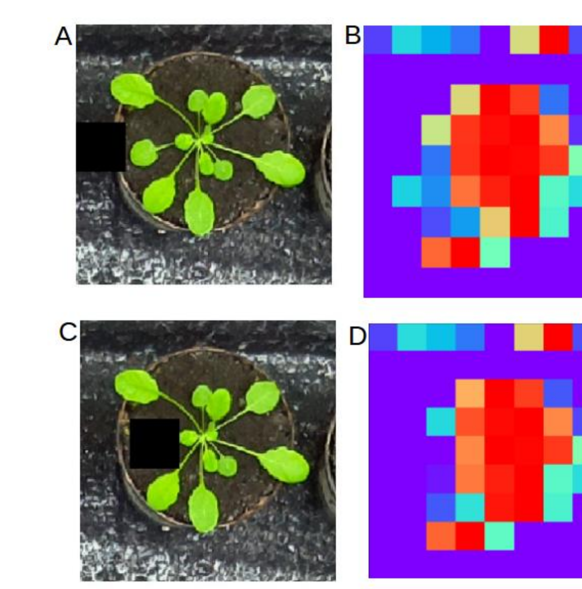

The heatmaps **B** and **D** are LRP at the final convolutional layer. In **D** the **black box** only has impact in the area where it occludes a leaf.

The band on the top is present in all images from the A1-A4 datasets.

Conclusions:

- We employed deep learning visualization techniques to better understand the decision contributing factors in the regression based plant phenotyping task of leaf counting.
- We experimentally determined that the blade edge is the most important part of the plant that contributes to the final leaf count, regardless of the background or species and scale.
- We show that the network reacts to occlusions of the input and does not store count information in areas not corresponding to the plant.
- We show how we can compress the network, while not significantly impacting the performance.

Acknowledgements:
This work was supported by the EPSRC DPT PhD fellowship EP/N509644/1. We acknowledge and thank NVIDIA for providing hardware essential for this work.

References:

- [1] J. Springenberg, A. Dosovitskiy, T. Brox, and M. Riedmiller.Striving for simplicity: The all convolutional net. In ICLR , 2015.
- [2] S. Bach, A. Binder, G. Montavon, F. Klauschen, K.-R.M¨uller, and W. Samek. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. PloS one, 10(7):e0130140, 2015.

EPSRC Engineering and Physical Sciences Research Council

NVIDIA