Leveraging Multiple Datasets For Deep Leaf Counting

Andrei Dobrescu
A.Dobrescu@ed.ac.uk

Mario Valerio Giuffrida
valerio.giuffrida@imtlucca.it

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

The University of Edinburgh
IMT School for Advanced Studies Lucca
The Alan Turing Institute

Computer Vision Problems in Plant Phenotyping (CVPPP) Workshop 2017

Motivations and Contributions

- Number of leaves is a key phenotyping trait
- Plants have a dynamic and complex structure that challenges humans and state-of-the-art methods
- We present a deep learning method to predict leaf count as direct regression problem
- We exploit information available in different plant datasets while training the model on plant images only
- Our method achieved the 1st position in the CVPPP 2017 leaf counting challenge

Effect of Pooling more Data

Visualization of network learning when trained on all (Ac) versus just one dataset (A2) using a sliding window.

Ac

Only A2

The Deep Leaf Counter Using *Resnet 50* [H]

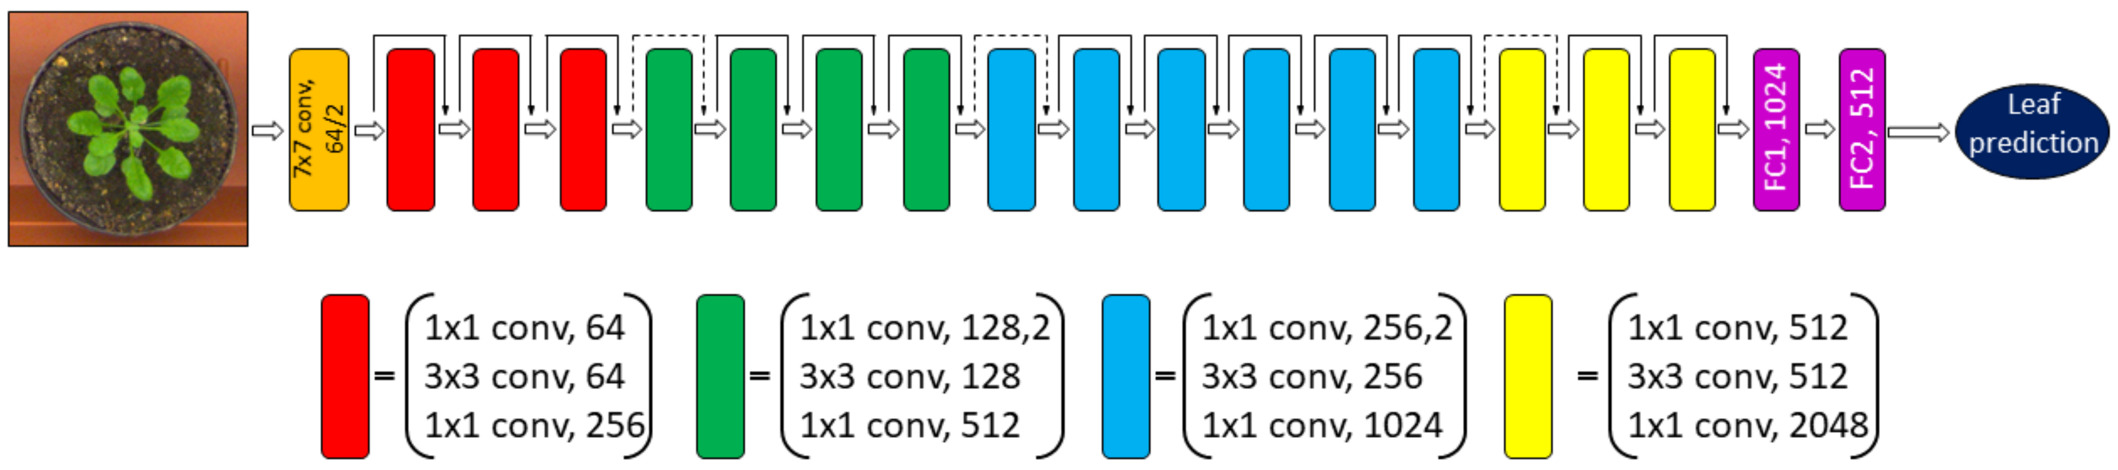

- We modified the ResNet, replacing the top layers with two FC layers and a single neuron with ReLU activation
- We perfomed tests to assess model training on one or more datasets.
- Combining datasets from different sources and species improves results more than just using affine data augmentation techniques.
- We split the training dataset four times with these proportions: 50% training, 25% validation and testing to emulate challenge ratios.
- The challenge results were obtained using an ensemble of four models fused by averaging their predictions.

(A) Test results on training set A1

Train Sets	DiC	|DiC|	MSE	R^2	[%]
A1	-0.81(0.85)	0.94(0.70)	1.38	0.23	25
A1+A2	-0.06(1.03)	0.75(0.71)	1.06	0.76	41
A1+A4	-0.75(0.90)	0.88(0.78)	1.38	0.69	34
Ac	0.28(0.80)	0.53(0.66)	0.72	0.60	56

(B) Test results on training set A2

Train Sets	DiC	|DiC|	MSE	R^2	[%]
A2	-2.38(2.69)	2.38(2.69)	12.88	0.29	38
A1+A2	-0.56(2.06)	1.69(1.51)	6.31	0.65	25
A2+A4	-0.75(2.15)	1.75(1.45)	6.38	0.65	31
Ac	-0.38(1.11)	0.88(0.78)	1.38	0.92	38

(C) Test results on training set A3

Train Sets	DiC	|DiC|	MSE	R^2	[%]
A3	-0.57(1.50)	1.43(0.73)	2.57	0.46	14
Ac	0.71(1.03)	0.71(1.03)	1.57	0.47	57

(D) Test results on training set A4

Train Sets	DiC	|DiC|	MSE	R^2	[%]
A4	0.1(1.14)	0.91(0.85)	1.54	0.96	35
A1+A4	-0.01(1.06)	0.77(0.73)	1.12	0.97	39
A2+A4	0.05(1.04)	0.73(0.75)	1.10	0.97	43
Ac	0.12(0.99)	0.69(0.73)	1.01	0.97	46

CVPPP 2018 Challenge Results

	DiC				|DiC|					Agreement [%]			MSE		
	Ours	[G]	[RP]	[A]	Ours	[G]	[RP]	[R]	[A]	Ours	[G]	[A]	Ours	[G]	[A]
A1	**-0.39 (1.17)**	-0.79 (1.54)	0.2 (1.4)	-0.33 (1.38)	**0.88 (0.86)**	1.27 (1.15)	1.1 (0.9)	0.8 (1.0)	1.00 (1.00)	**33.3**	27.3	30.3	**1.48**	2.91	1.97
A2	**-0.78 (1.64)**	-2.44 (2.88)	-	-0.22 (1.86)	1.44 (1.01)	2.44 (2.88)	-	-	**1.56 (0.88)**	11.1	**44.4**	11.1	**3.00**	13.33	3.11
A3	**0.13 (1.55)**	-0.04 (1.93)	-	2.71 (4.58)	**1.09 (1.10)**	1.36 (1.37)	-	-	3.46 (4.04)	**30.4**	19.6	7.1	**2.38**	3.68	28.00
A4	**0.29 (1.10)**	-	-	0.23 (1.44)	**0.84 (0.76)**	-	-	-	1.08 (0.97)	**34.5**	-	29.2	**1.28**	-	2.11
A5	**0.25 (1.21)**	-	-	0.80 (2.77)	**0.90 (0.85)**	-	-	-	1.66 (2.36)	**33.2**	-	23.8	**1.53**	-	8.28
All	**0.19 (1.24)**	-	-	0.73 (2.72)	**0.91 (0.86)**	-	-	-	1.62 (2.30)	**32.9**	-	24.0	**1.56**	-	7.90

Conclusions

- Outperformed state-of-the-art results, using less supervision (e.g. no plant mask)
- Pooling data from different sources and setups helped to reduce overfitting
- The most discriminative features are focused in the plant region, as more training samples are given
- Averaging from different trained models reduces variance

References

[H] K. He, X. Zhang, S. Ren, J. Sun. Deep residual learning for image recognition. arXiv, 2015
[G] M. V. Giuffrida, M. Minervini, S. Tsaftaris. Learning to Count Leaves in Rosette Plants. CVPPP 2015
[RP] B. Romera-Paredes, P. H. S. Torr. Recurrent Instance Segmentation. ECCV 2016
[A] S. Aich, I. Stavness, Leaf Counting with Deep Convolutional and Deconvolutional Networks, CVPPP 2017 (arXiv preprint)
[R] M. Ren and R. S. Zemel. End-to-End Instance Segmentation and Counting with Recurrent Attention, arXiv 2016.

Acknowledgments

This work was supported by an EPSRC DTP PhD fellowship, number EP/N509644/. We acknowledge and thank NVIDIA for providing hardware essential for our work