RDCNet: Instance segmentation with a minimalist recurrent residual network
– Supplementary Material –

Raphael Ortiz[1], Gustavo de Medeiros[1], Antoine H.F.M. Peters[1,2], Prisca Liberali[1,2], and Markus Rempfler[1]

[1] Friedrich Miescher Institute for Biomedical Research (FMI), Basel, Switzerland
[2] Faculty of Sciences, University of Basel, Basel, Switzerland

A Data augmentation

Table 1. Data augmentation parameters used for experiments with the different datasets. Note that empty fields indicate that the particular augmentation method was not used for that dataset.

Augmentation	Parameter	CVPPP2017	MoNuSeg	3D-ORG
Random axis flip	p_{flip}	0.5	0.5	0.5
	Axes	X, Y	X, Y	X, Y, Z
Random offset $\sim \mathcal{N}(\mu, \sigma)$ (IID per patch)	μ	0	0	0
	σ	0.2	0.2	0.2
Random noise $\sim \mathcal{N}(\mu, \sigma)$ (IID per pixel)	μ			0.05
	σ			0.3
Random HSV shift	Δ Hue	± 0.3	± 0.3	
	Saturation	$[0.8, 1.2]$	$[0.8, 1.2]$	
	Value	$[0.8, 1.2]$	$[0.8, 1.2]$	
Random Gaussian blur σ range	p_{active}	0.5	0.5	
	σ	[0.5, 3]	[0.5, 2]	
Random affine transform	zoom factor	[0.9, 1.1]	[0.9, 1.1]	
	shear	±5	±5	
	rotation angle	±10	±10	
Random warp	A	20	20	
Random clipping $\sim \mathcal{N}(\mu, \sigma)$	μ_{min}	-1	-1	
	μ_{max}	1	1	
	σ	0.3	0.3	

We detail the data augmentation used during training for all our experiments in Table 1. Random warp was generated by smoothing uniformly distributed pixel offset in range $[-A, A]$ with a gaussian kernel $\sigma = 2\ A$ where A is the max amplitude.

B Semi-convolutional embedding

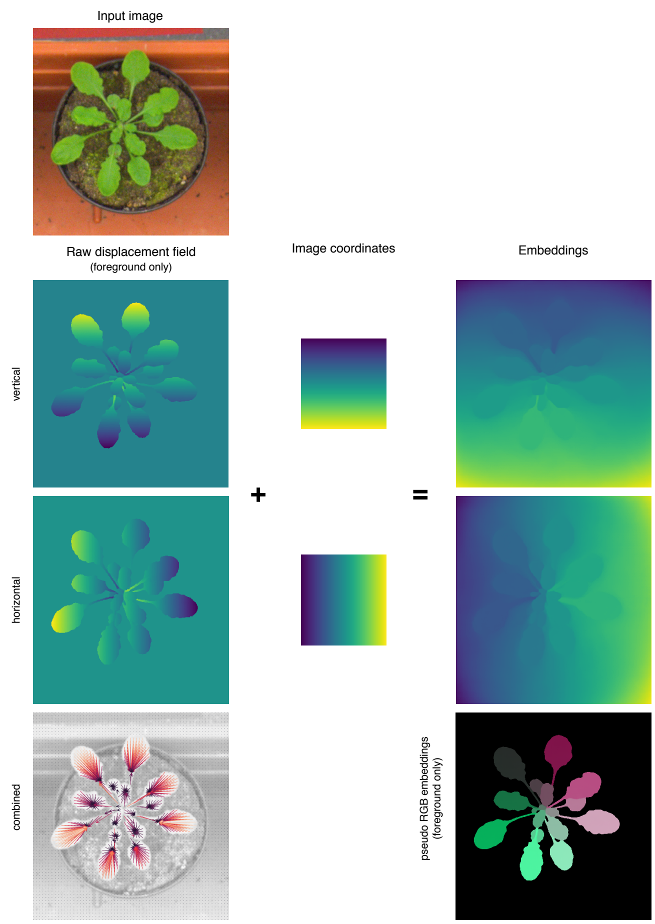

Fig. 1. Example of how so called semi-convolutional embeddings [1] are formed. Given the raw image on the top left, the network output a displacement field where all pixels of a given instance point to the same location. Note that this centre of attraction is not restricted to the center of mass or any specific point of the object. Adding the image coordinates at each position produces the final embeddings where each instance is ≈assigned a unique embedding. This works regardless of the size of the image as the effective field of the network does not need to cover the entire image.

C Centroids estimation during inference

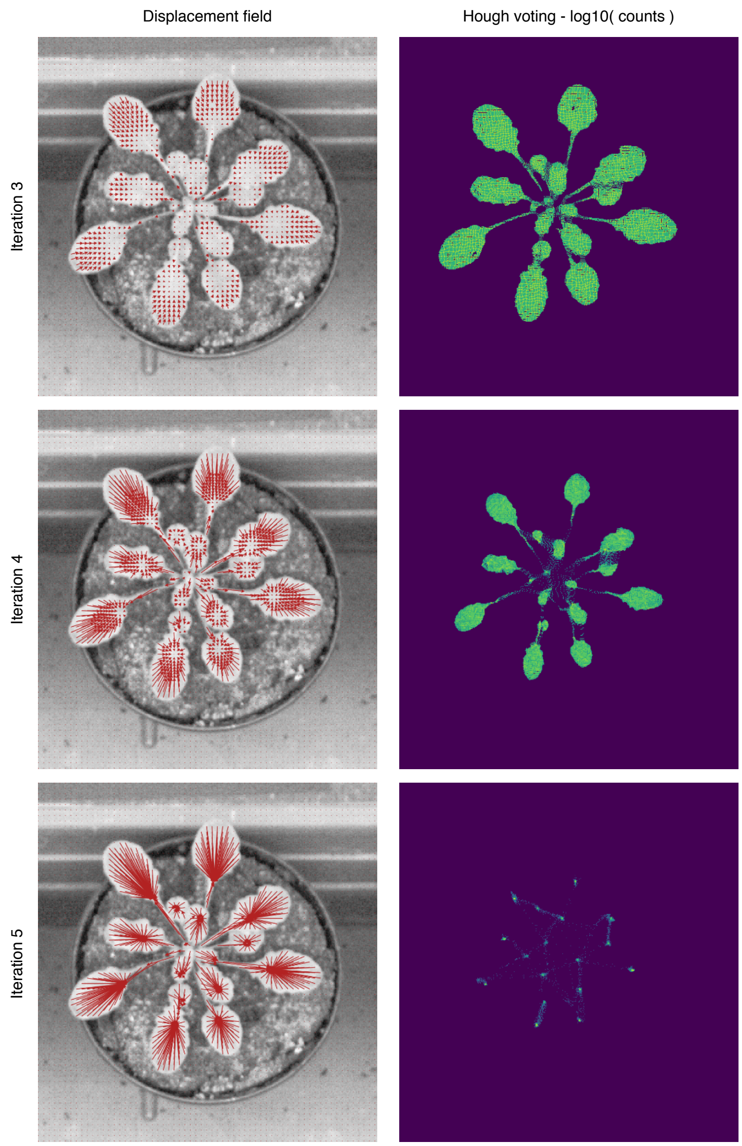

Fig. 2. Example illustrating post-processing of the network output to obtain the centres used for final embedding assignment. The model was trained with 5 iterations. Initially, each pixel is voting for himself and instances are gradually separated with each iteration until they can be identified with a trivial Hough voting step.

References

1. Novotny, D., Albanie, S., Larlus, D., Vedaldi, A.: Semi-convolutional operators for instance segmentation. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 86–102 (2018)