Data-Hungry Models: Deep Learning of Phenotypes in Crop Plants

Dr. Michael Pound

University of Nottingham

Introduction

- Deep learning for fast and accurate recovery of plant traits
  - Traits drive downstream tasks
- Case studies from my work and our lab:
  - Classification
  - Segmentation
  - Localisation
- Recent work in techniques for more challenging images
- Some thoughts about where we go from here over the next few years

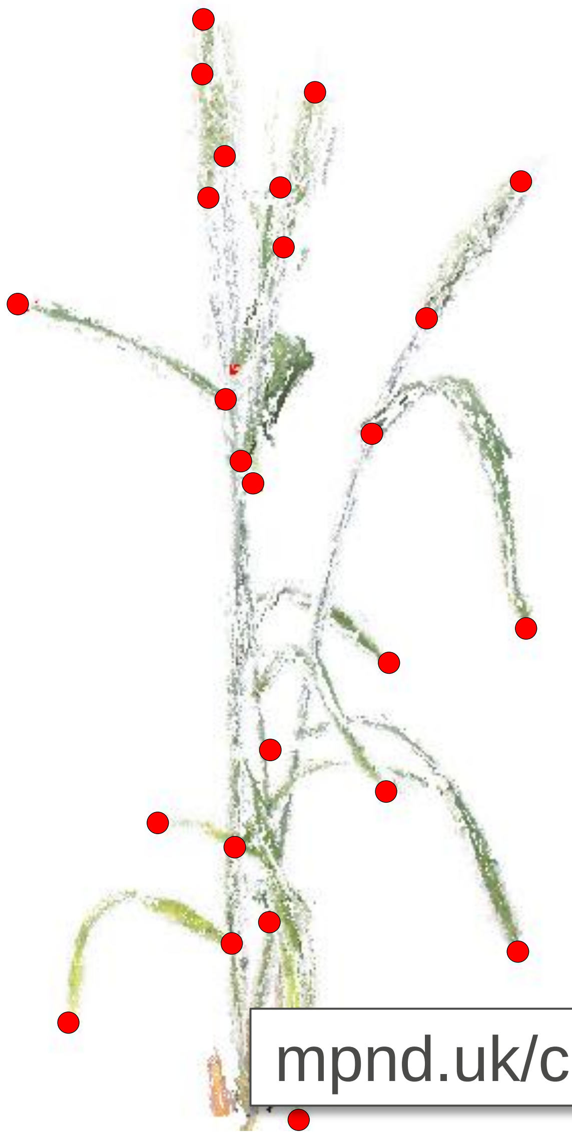

mpnd.uk/cis21

About Us

- We use image-analysis, machine learning and deep learning to drive research in Bioscience
- We focus on the development of new techniques and deep network architectures
- My research often focuses on segmentation and object localisation

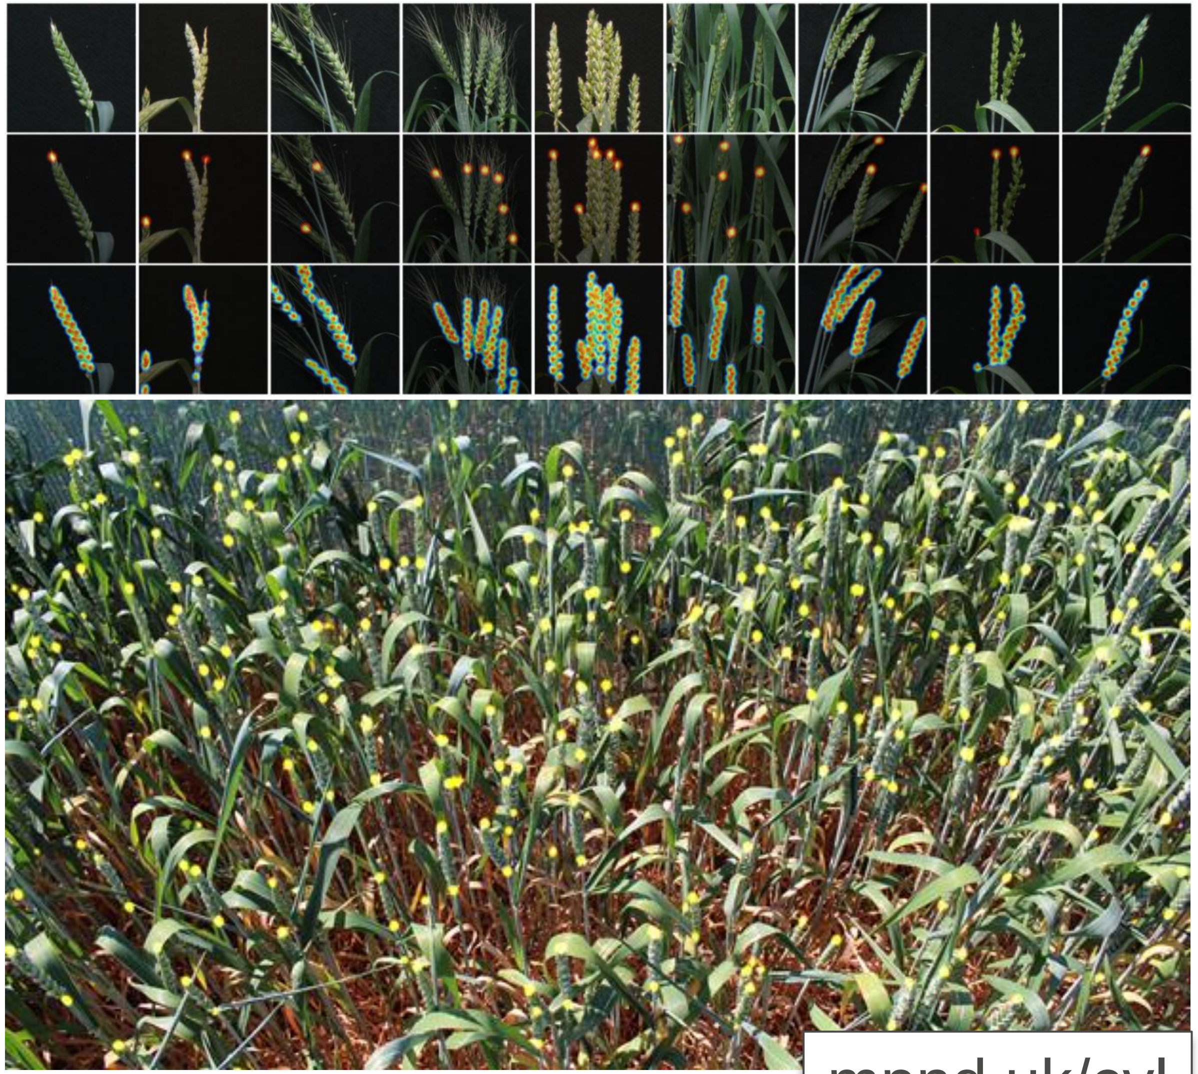

mpnd.uk/cvl

Capturing Plant Traits - An Example Problem

RootNav

- In 2014 we released RootNav
- Aims to quickly quantify 2D root images in a variety of growth conditions
- Accuracy is a primary focus

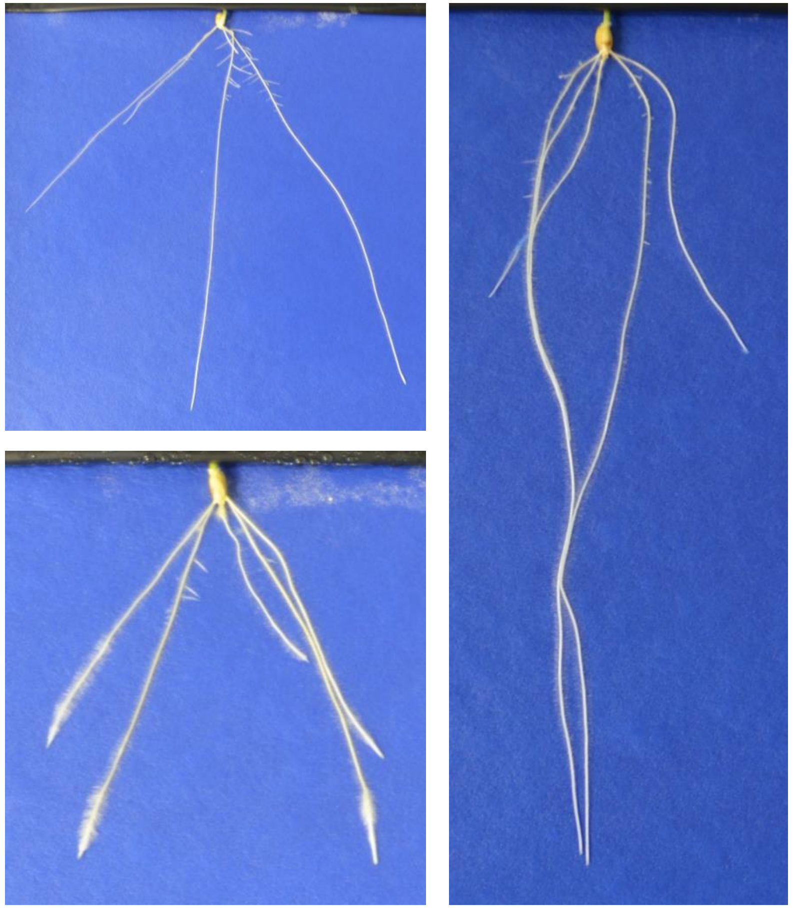

- High-throughput root traits can inform modelling approaches, as well as drive genomic studies such as GWAS and QTL

An Example Problem - RootNav

- RootNav is semi-automatic
- Usually users must identify root tips themselves

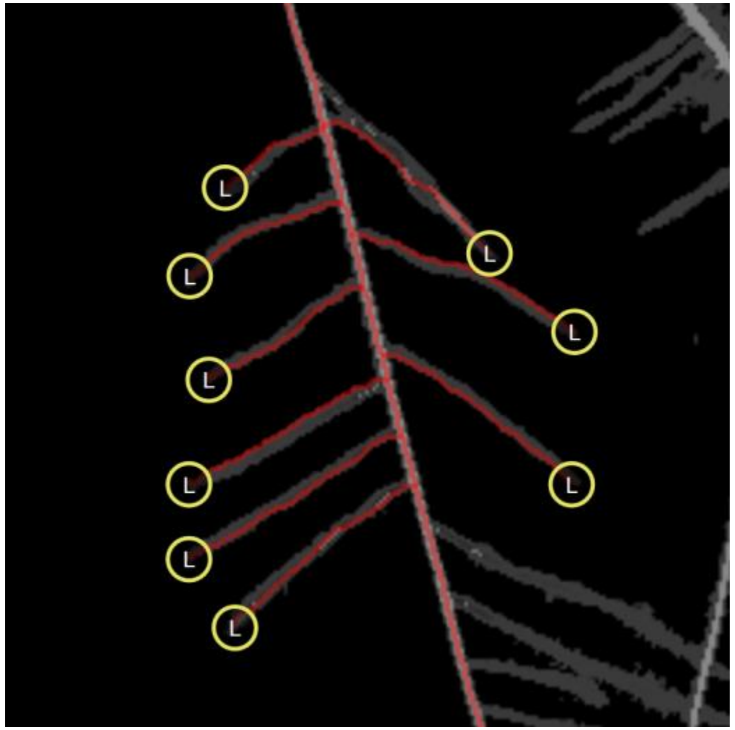

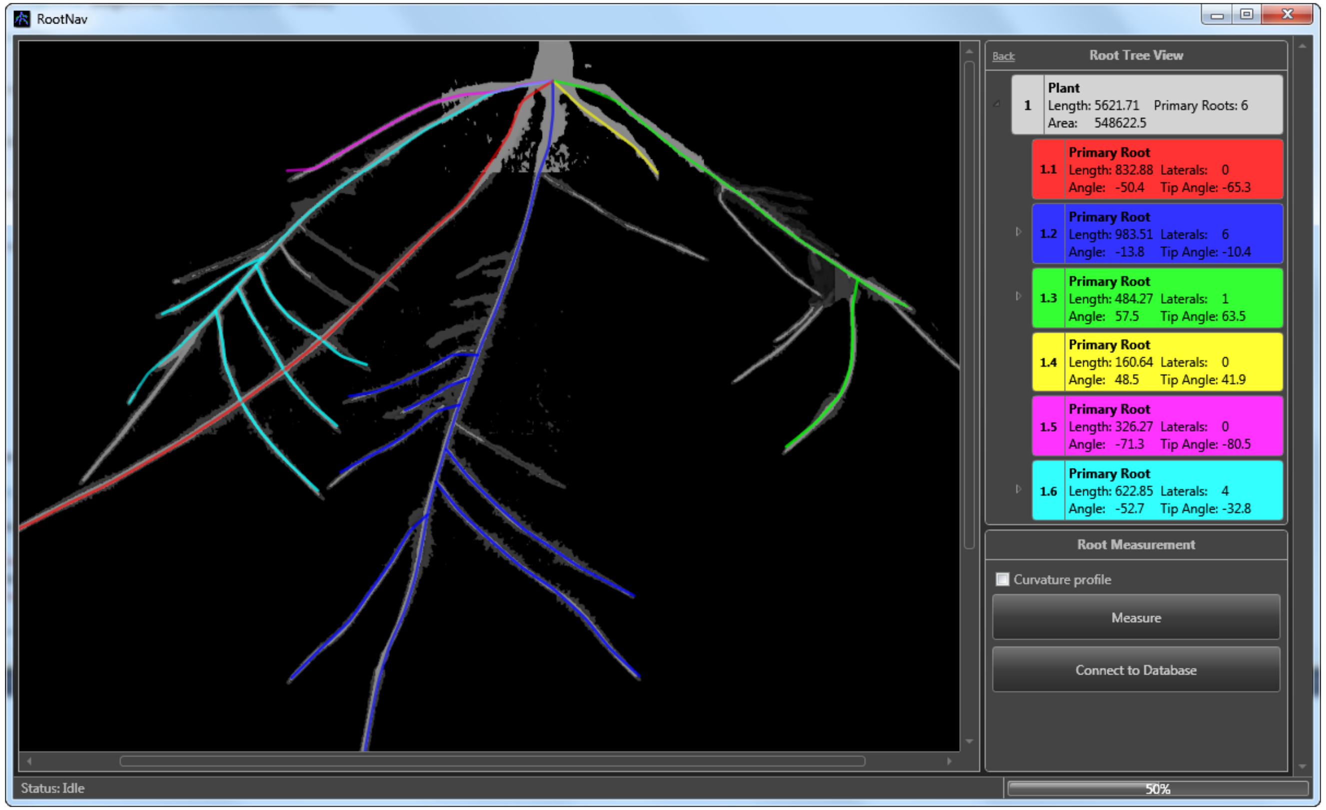

Representation Learning

- Deep learning has the potential to automate many of these detection processes
- Convolutional networks use convolution and pooling layers learn a feature representation
- A final neural network makes a decision based on these features

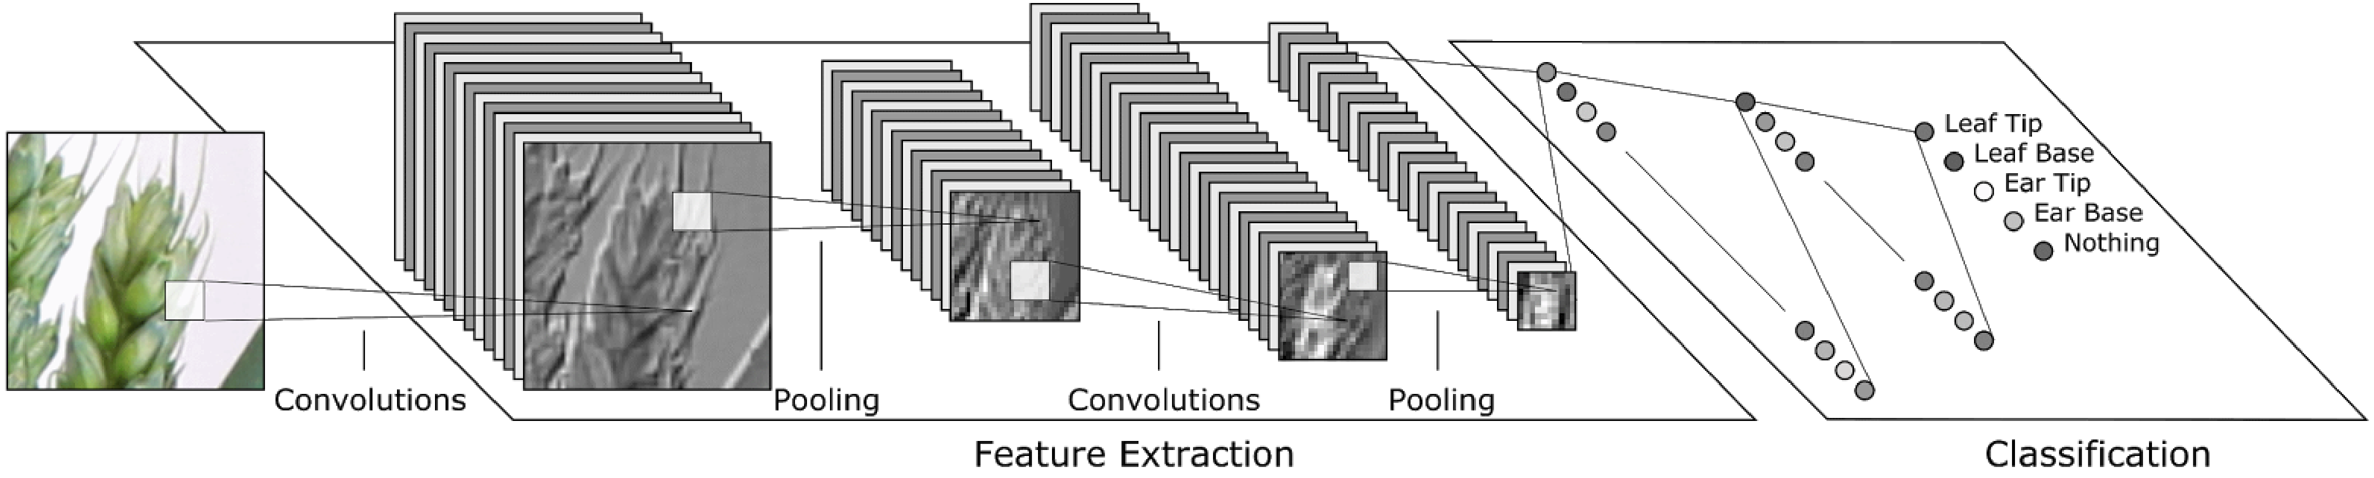

Shoot and Root Features

- Classification of shoot and root features

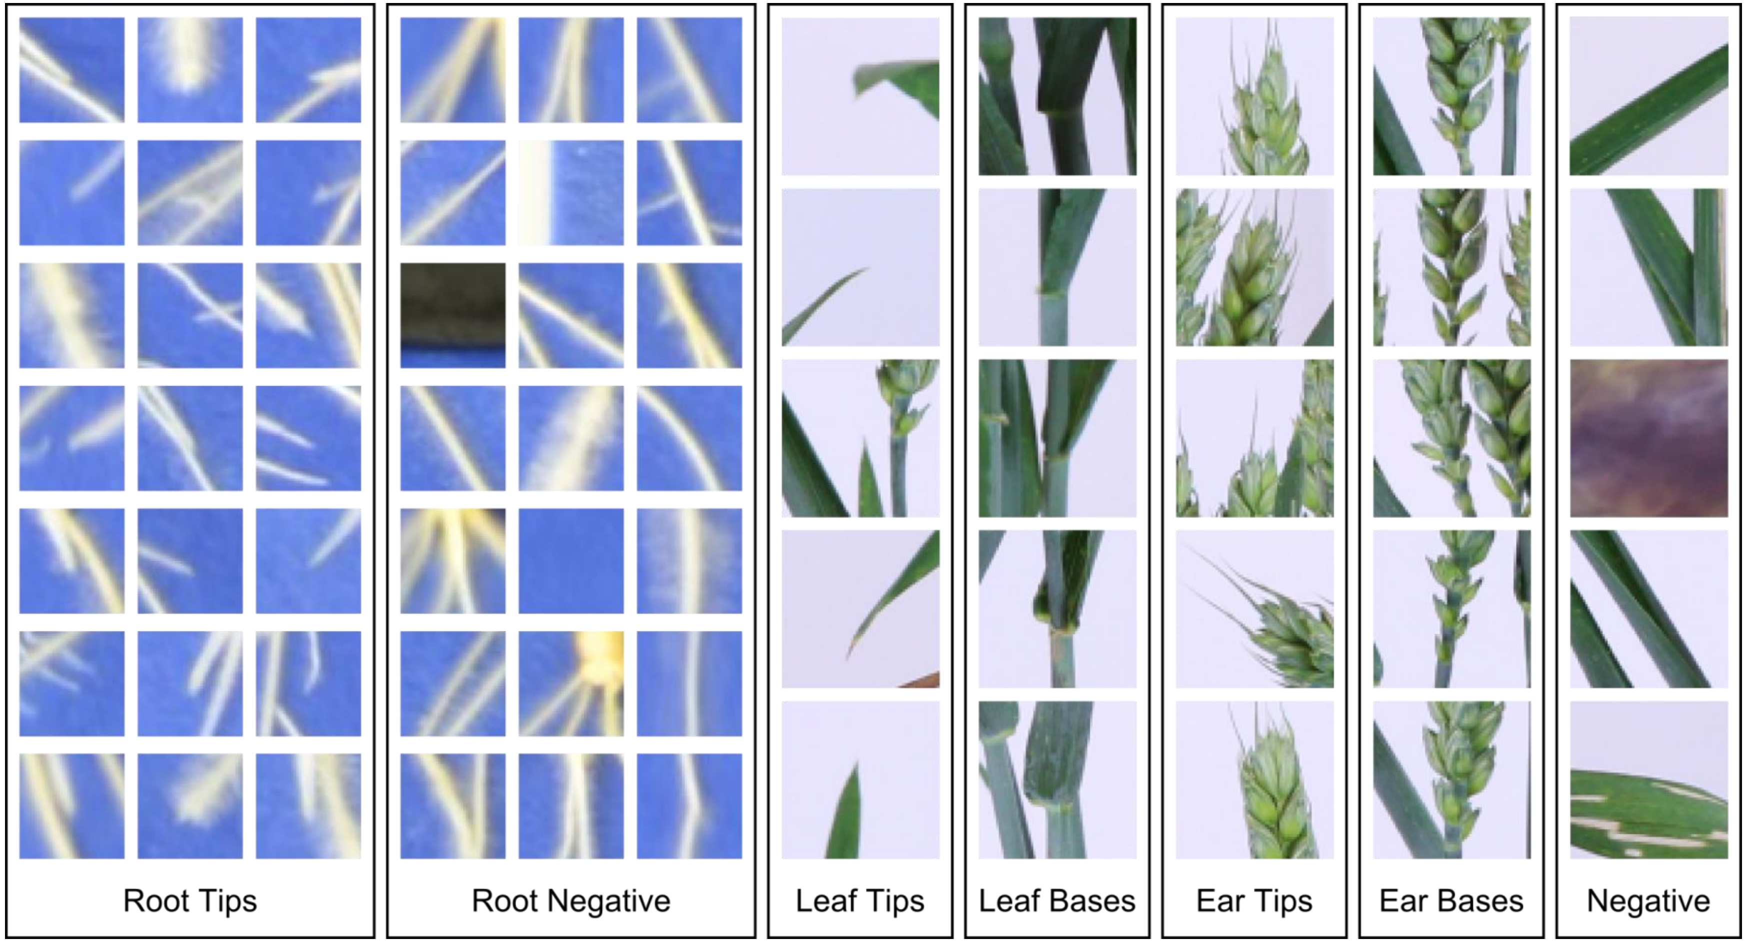

Shoot and Root Features

- Accuracy of the classifier is very high, with all classes >95%

Feature	Correctly Classified	Misclassified	Accuracy (%)
Root Tip	2904	73	97.5
Root Negative	5687	65	98.9
Total/Average	8591	138	98.4

Feature	Correctly Classified	Misclassified	Accuracy (%)
Leaf Tip	2225	113	95.2
Leaf Base	2299	52	97.8
Ear Tip	686	15	97.9
Ear Base	765	23	97.1
Shoot Negative	6110	136	97.8
Total/Average	12085	339	97.3

Localisation

- Use a sliding window approach
- We identified 12 of 14 root QTLs previously identified manually
- Useful features – but noisy

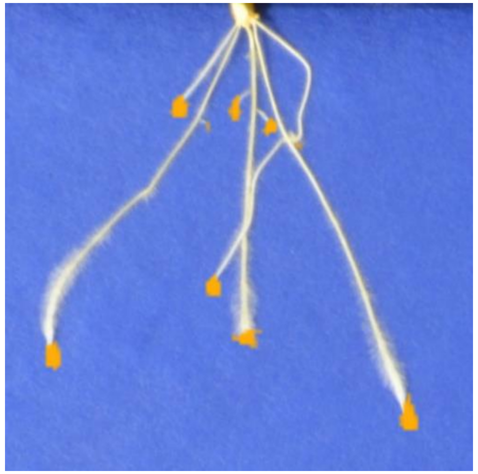

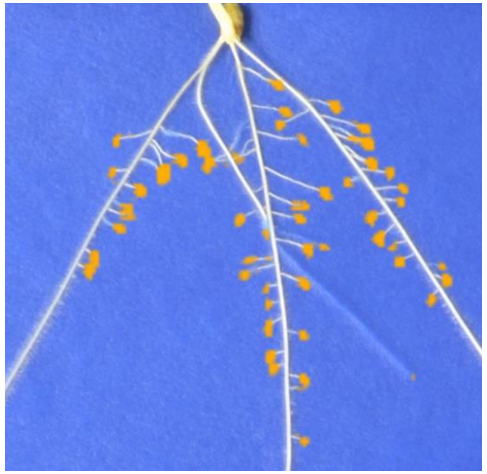

Deep machine learning provides state-of-the-art performance in image-based plant phenotyping. Gigascience 2017

mpnd.uk/giga

2D output - > segmentation, regression

- A sliding window does not enforce consistency between neighbouring pixels or regions
- Encoder-decoder networks are designed to return 2D output

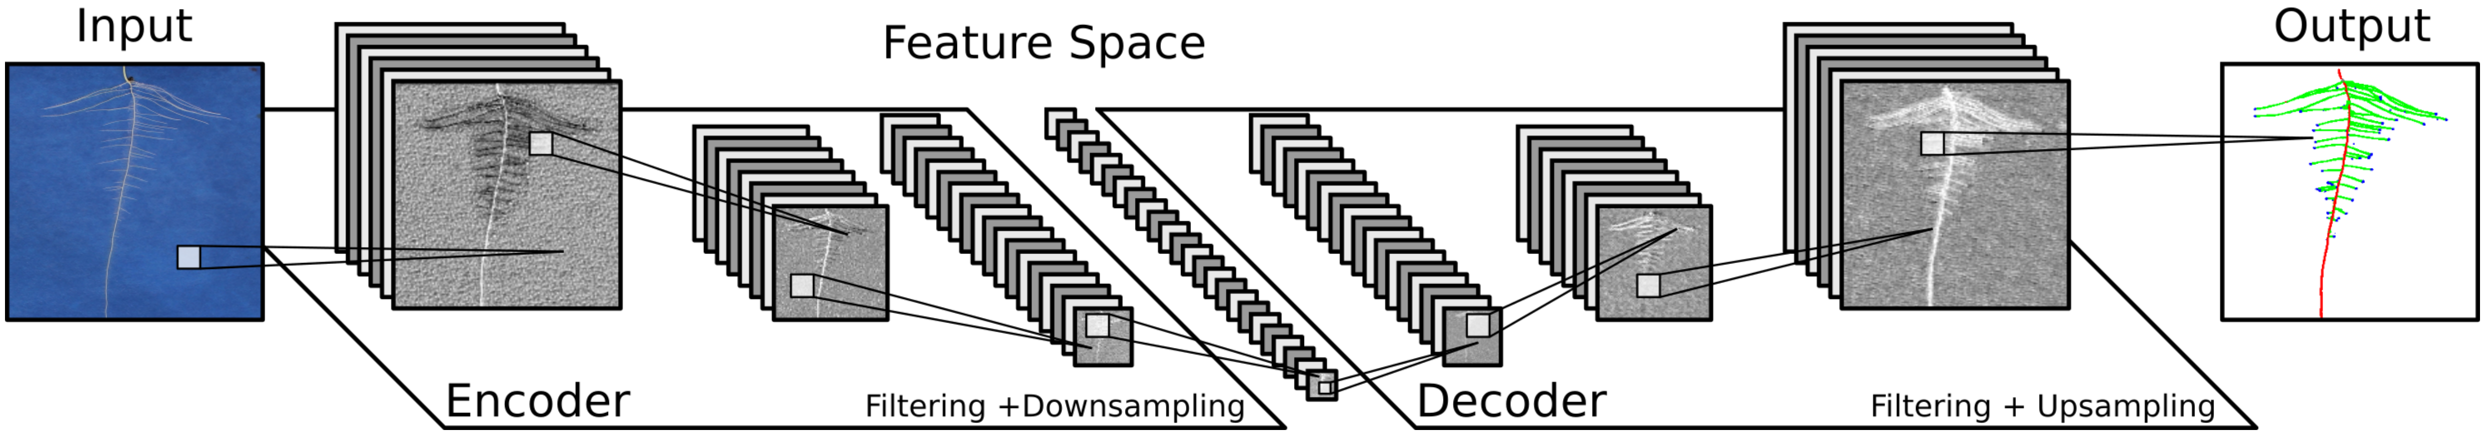

2D Segmentation

RootNav 2

- Automatic segmentation of roots, and localisation of tips and seed locations

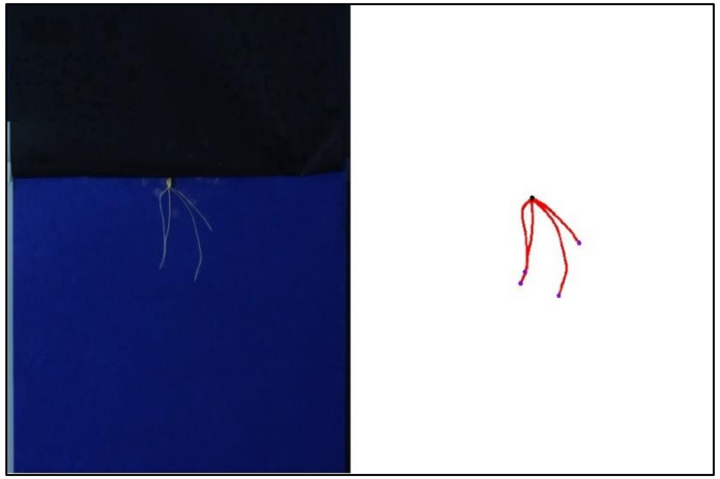
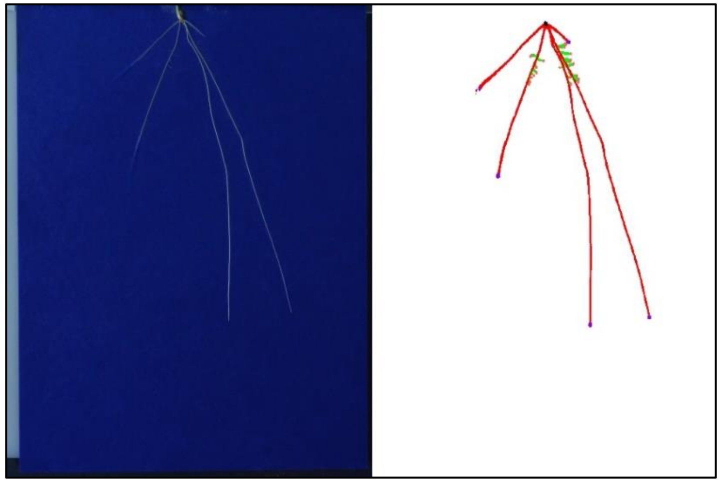
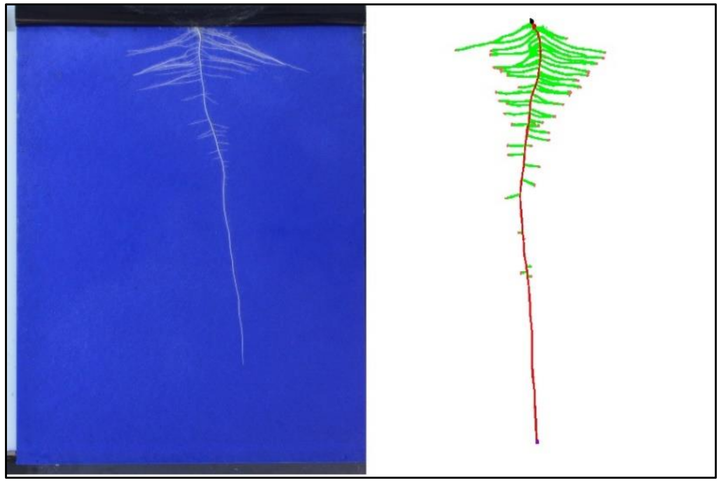
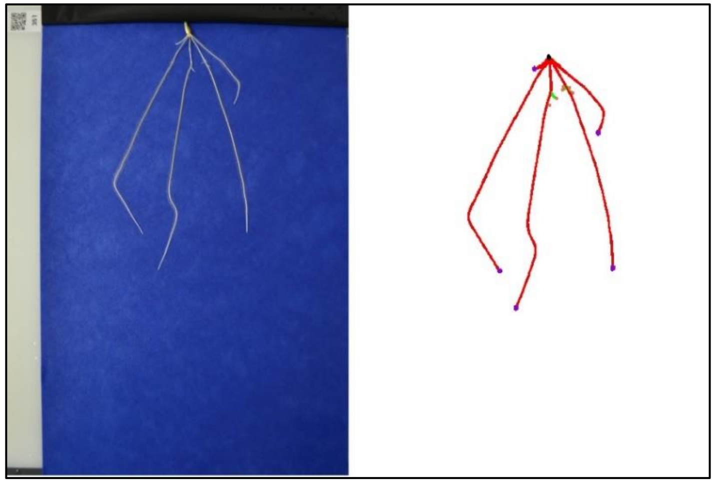
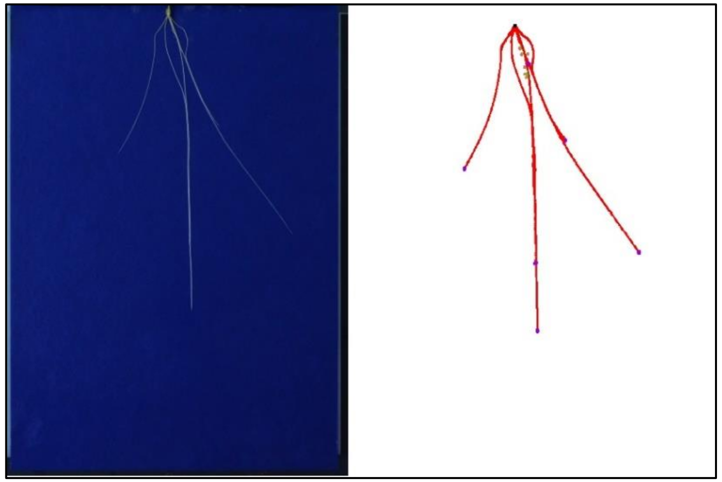

RootNav 2.0: Deep learning for automatic navigation of complex plant root architectures. GigaScience 2019

mpnd.uk/rn2

Transfer Learning

- Transfer learning to refine a network to new image types
- Our original dataset was 3000 images, this one is 200 images

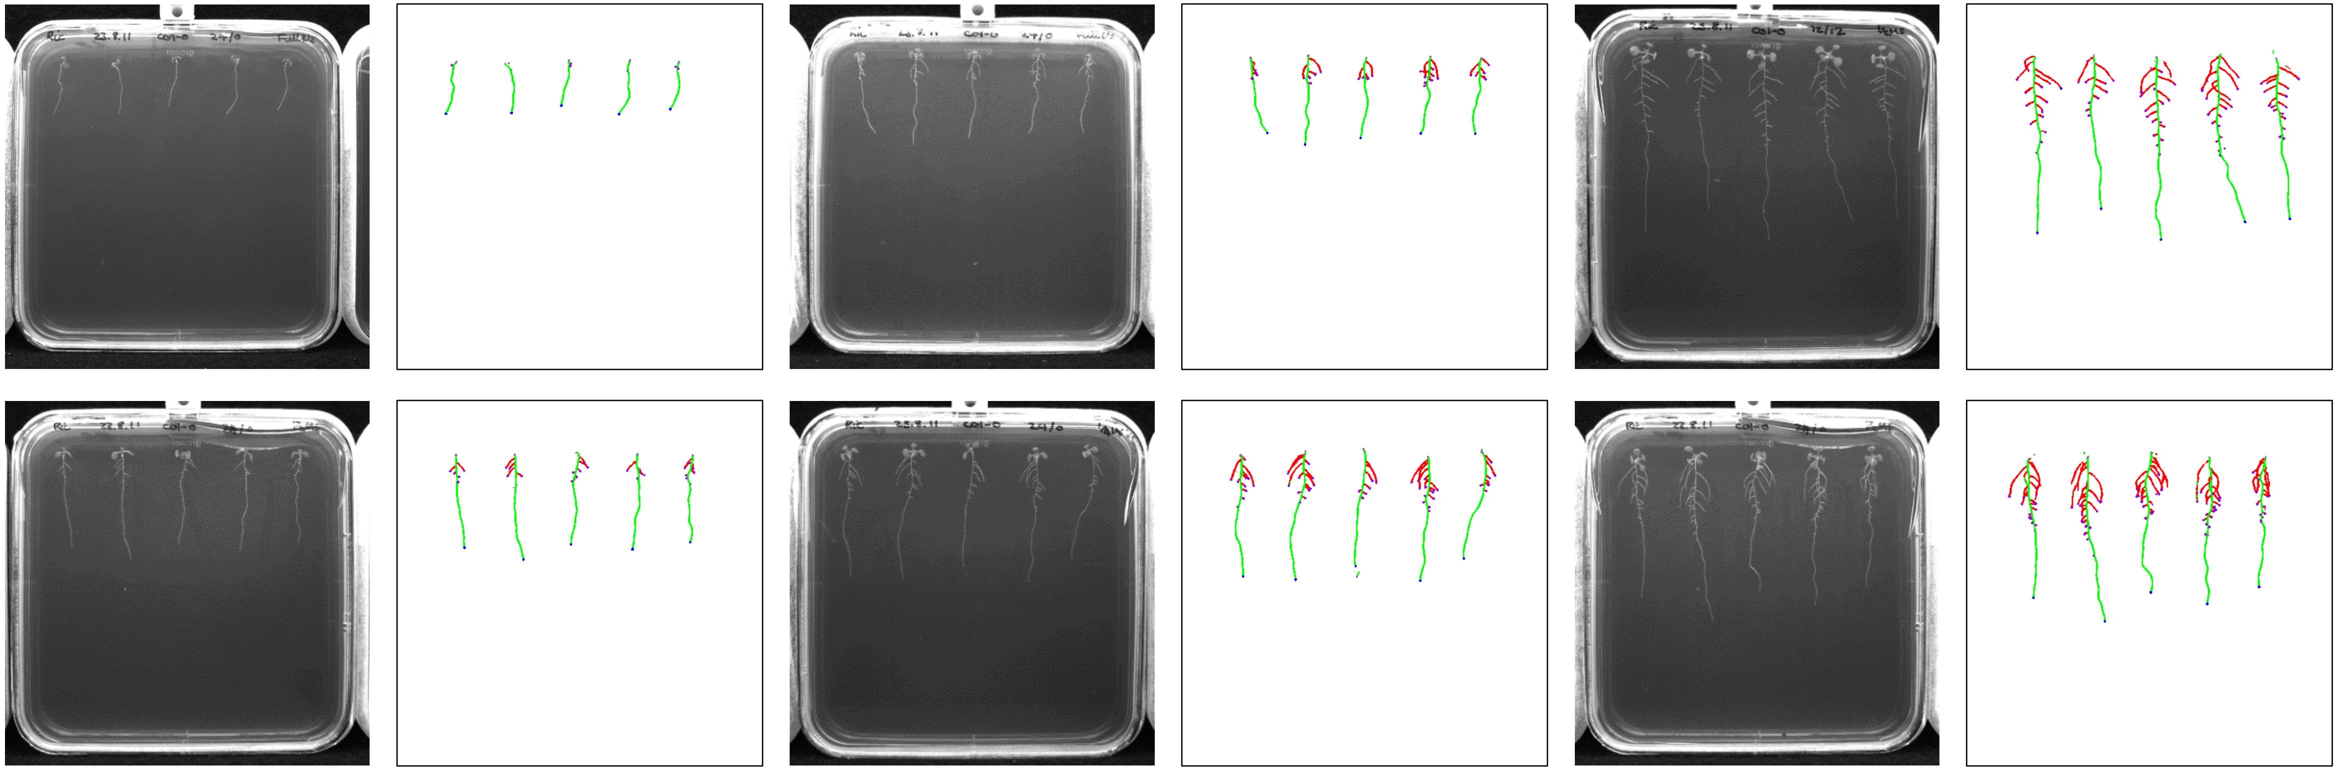

Multi-task Learning

- Heatmap regression to predict both the locations of spikes (spike tips) and the individual spikelets

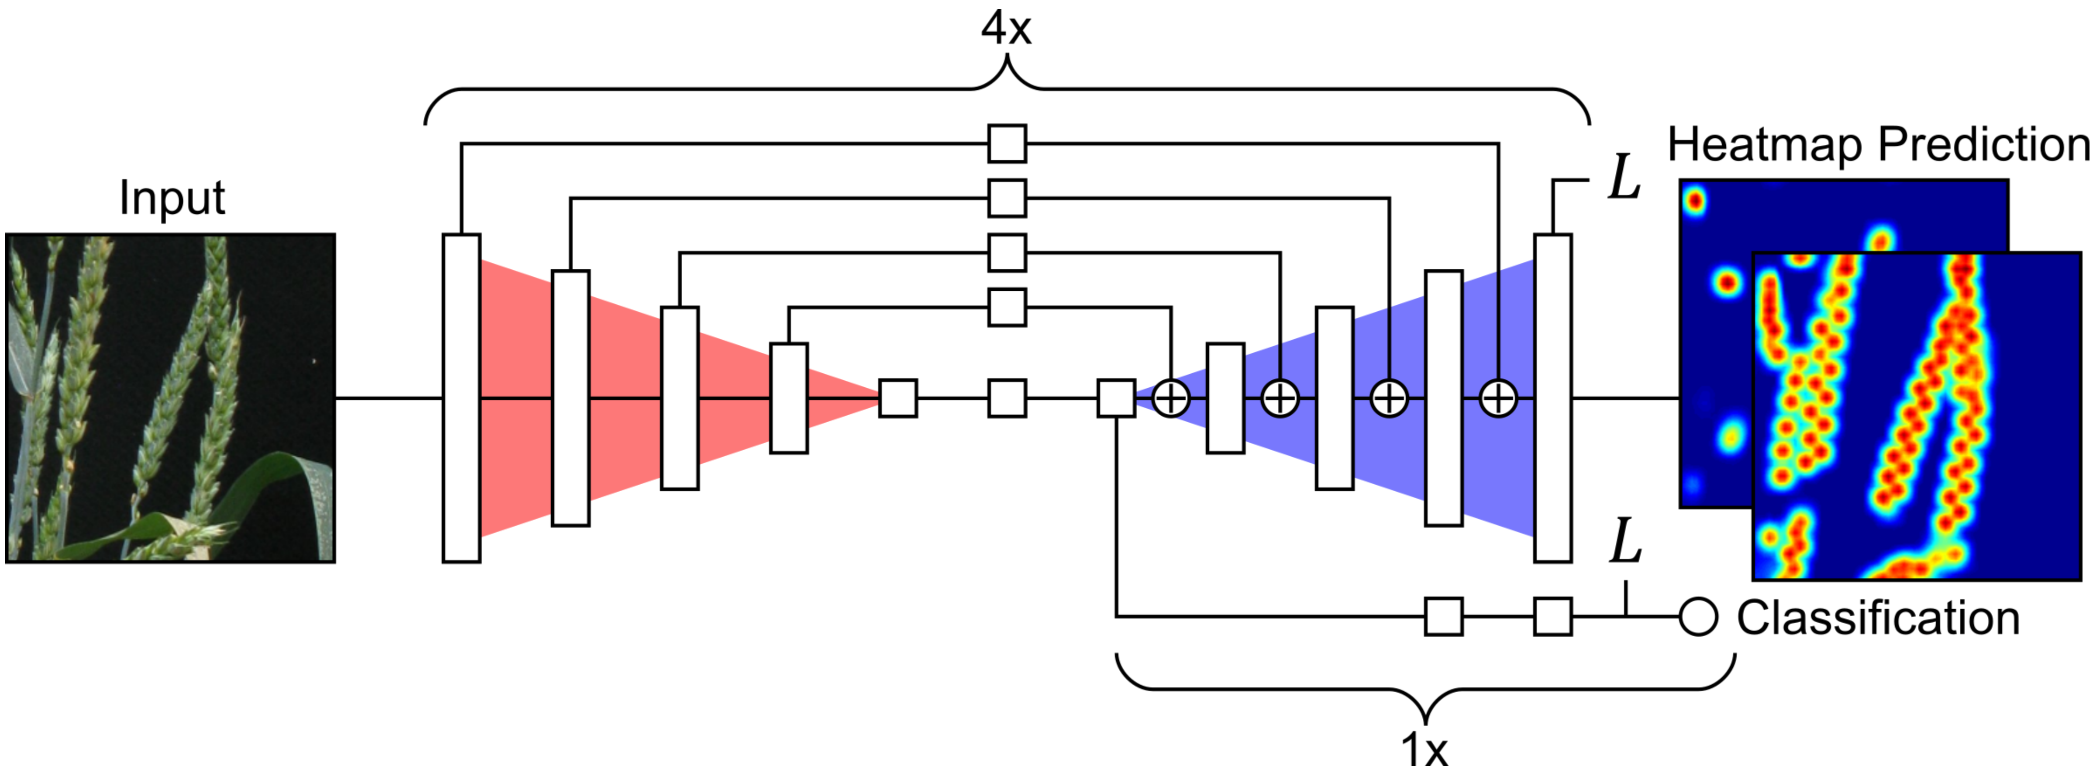

- Multi-task classification of awned phenotypes

Multi-task Learning

Spike counting: 95%

Spikelet counting: 99%

Awn identification: 99%

Deep Learning for Multi-task Plant Phenotyping. ICCVW 2017

mpnd.uk/iccv17

Other Applications

- On large images the approach can easily be run on smaller tiles

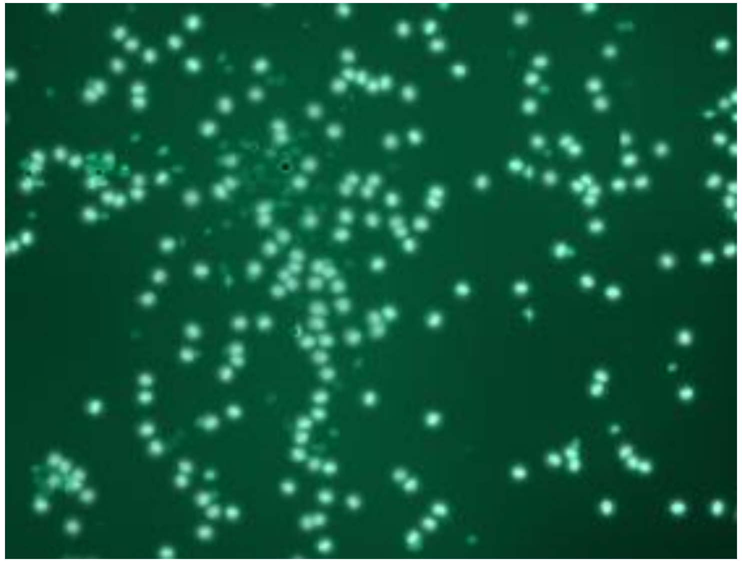

Harder Images

- Many standard techniques are based on common image datasets, more challenging data can pose a problem

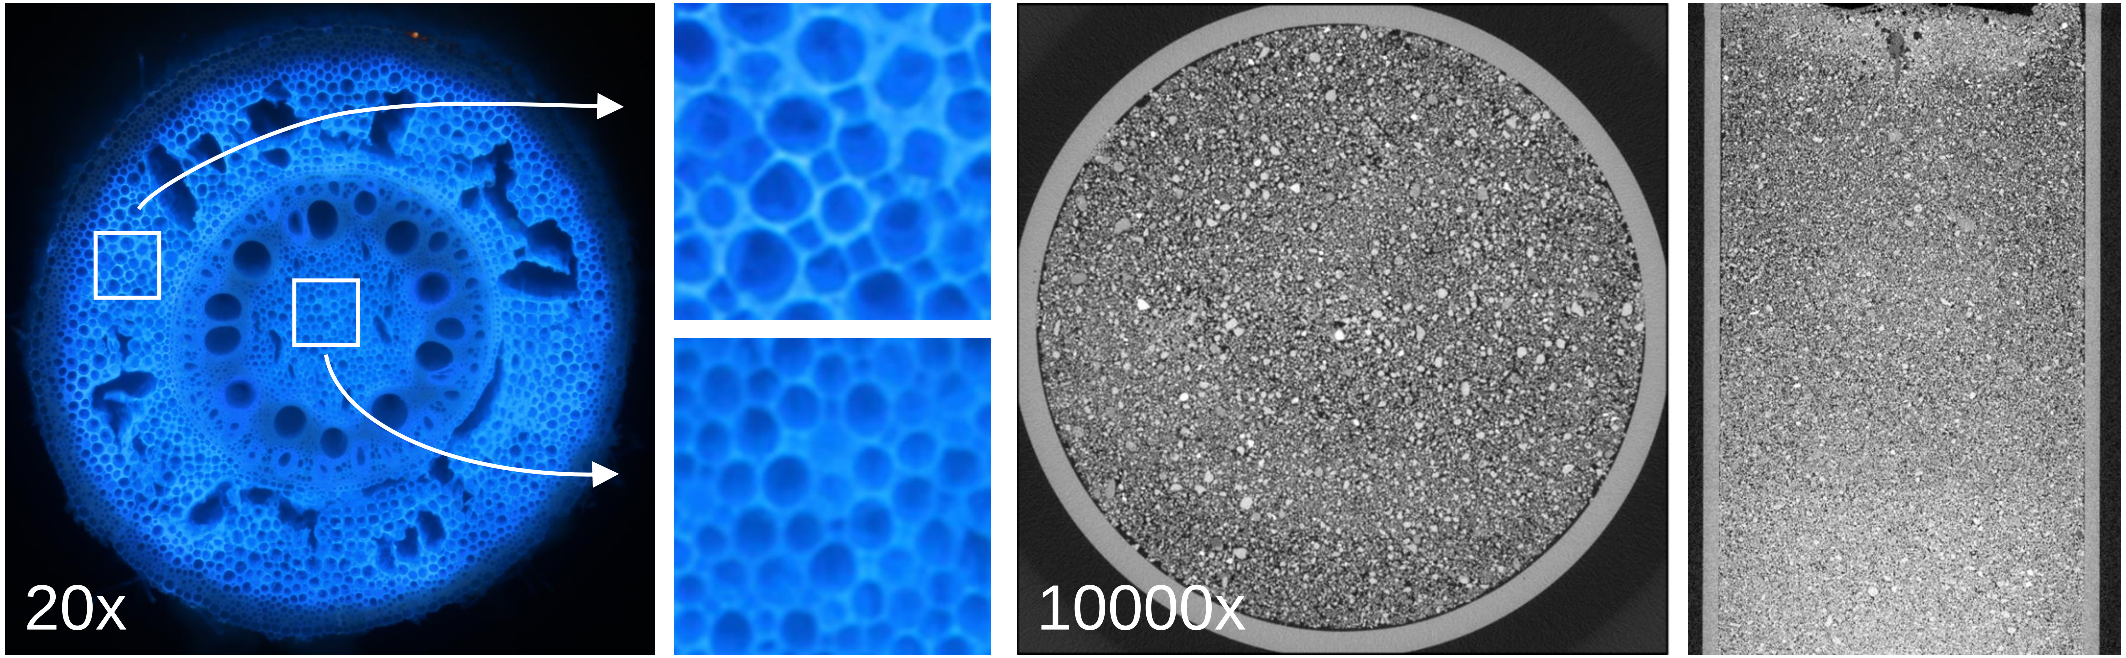

High-resolution Segmentation

- Common to split an image into smaller tiles, and process these sequentially
- This often loses context

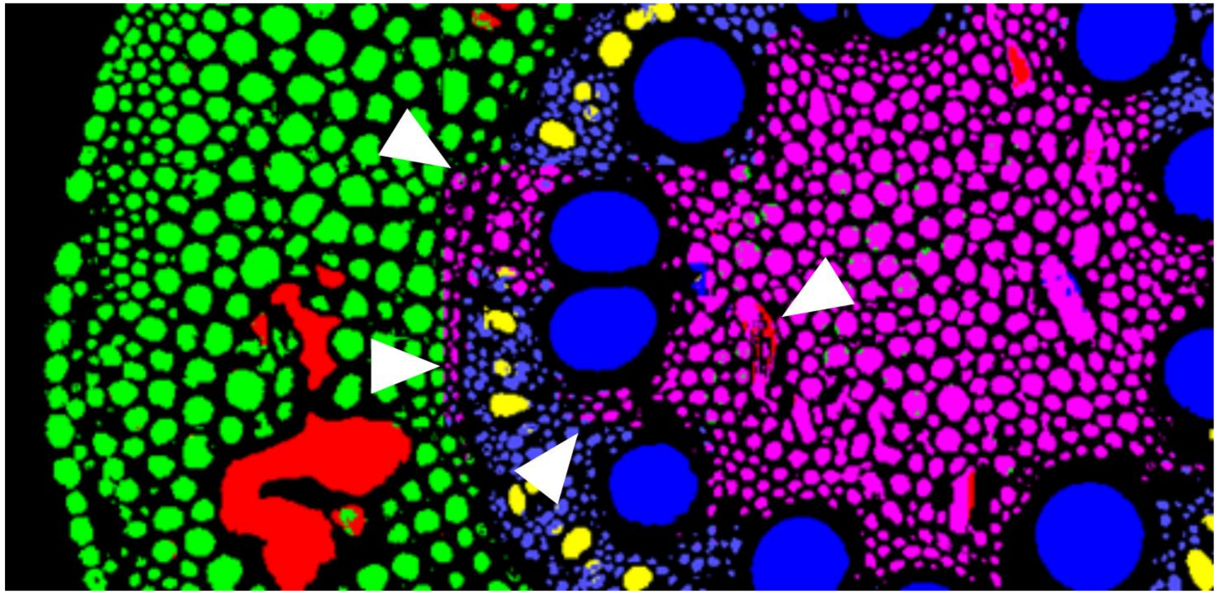

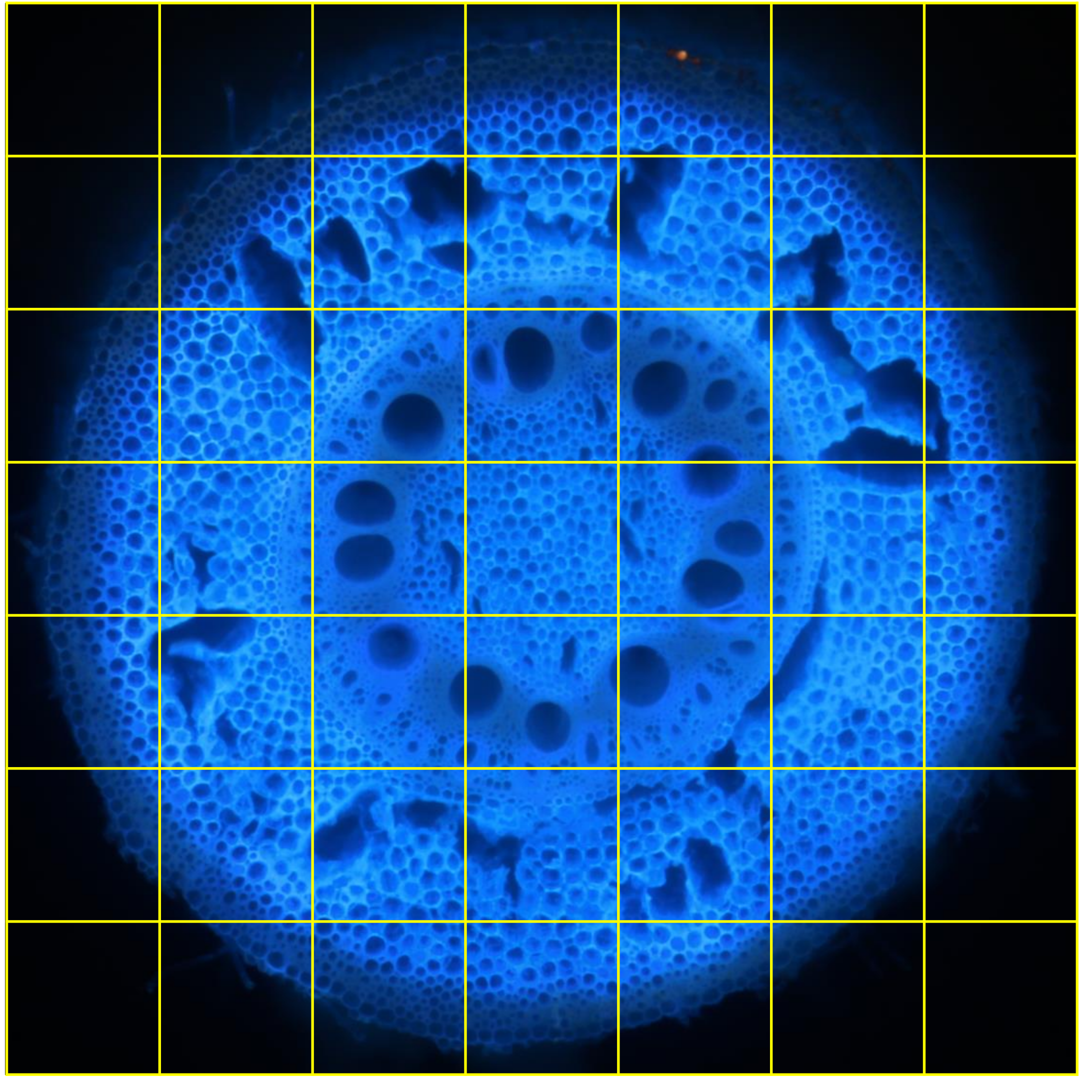

High-resolution Segmentation

- We save memory by splitting an image into small manageable tiles
- We restore context by sharing information between tiles
- These Crosstalk (XT) layers contain global filters and attention mechanisms

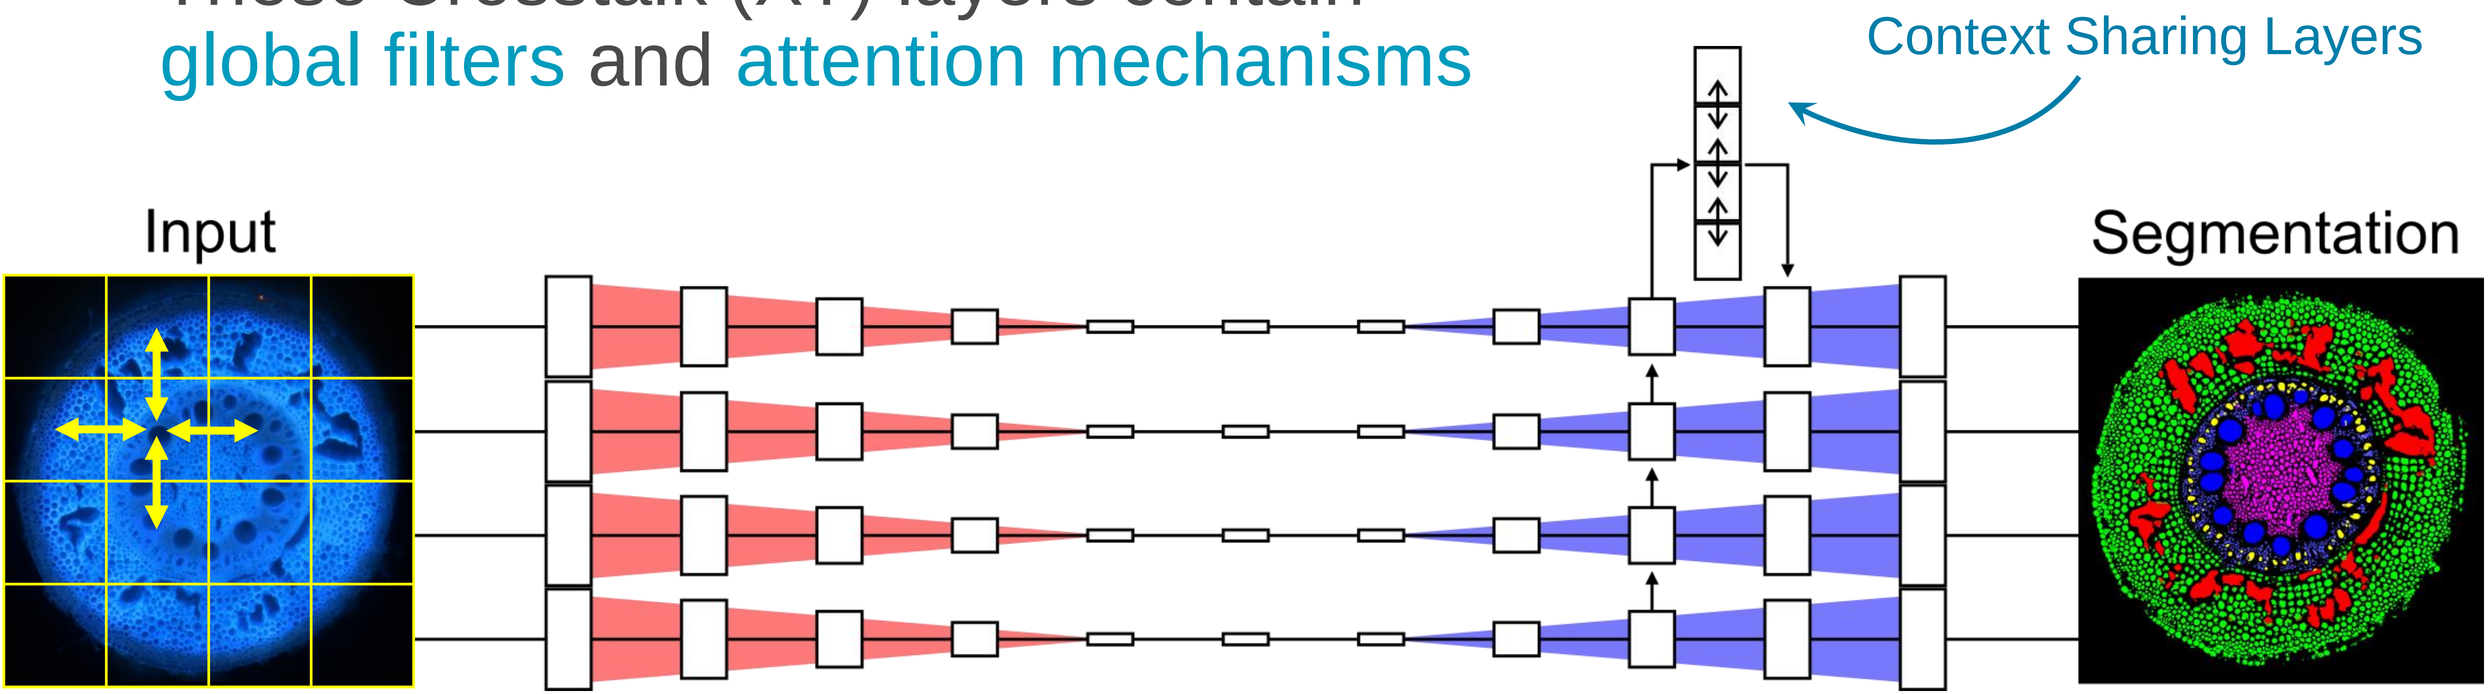

High-resolution Segmentation

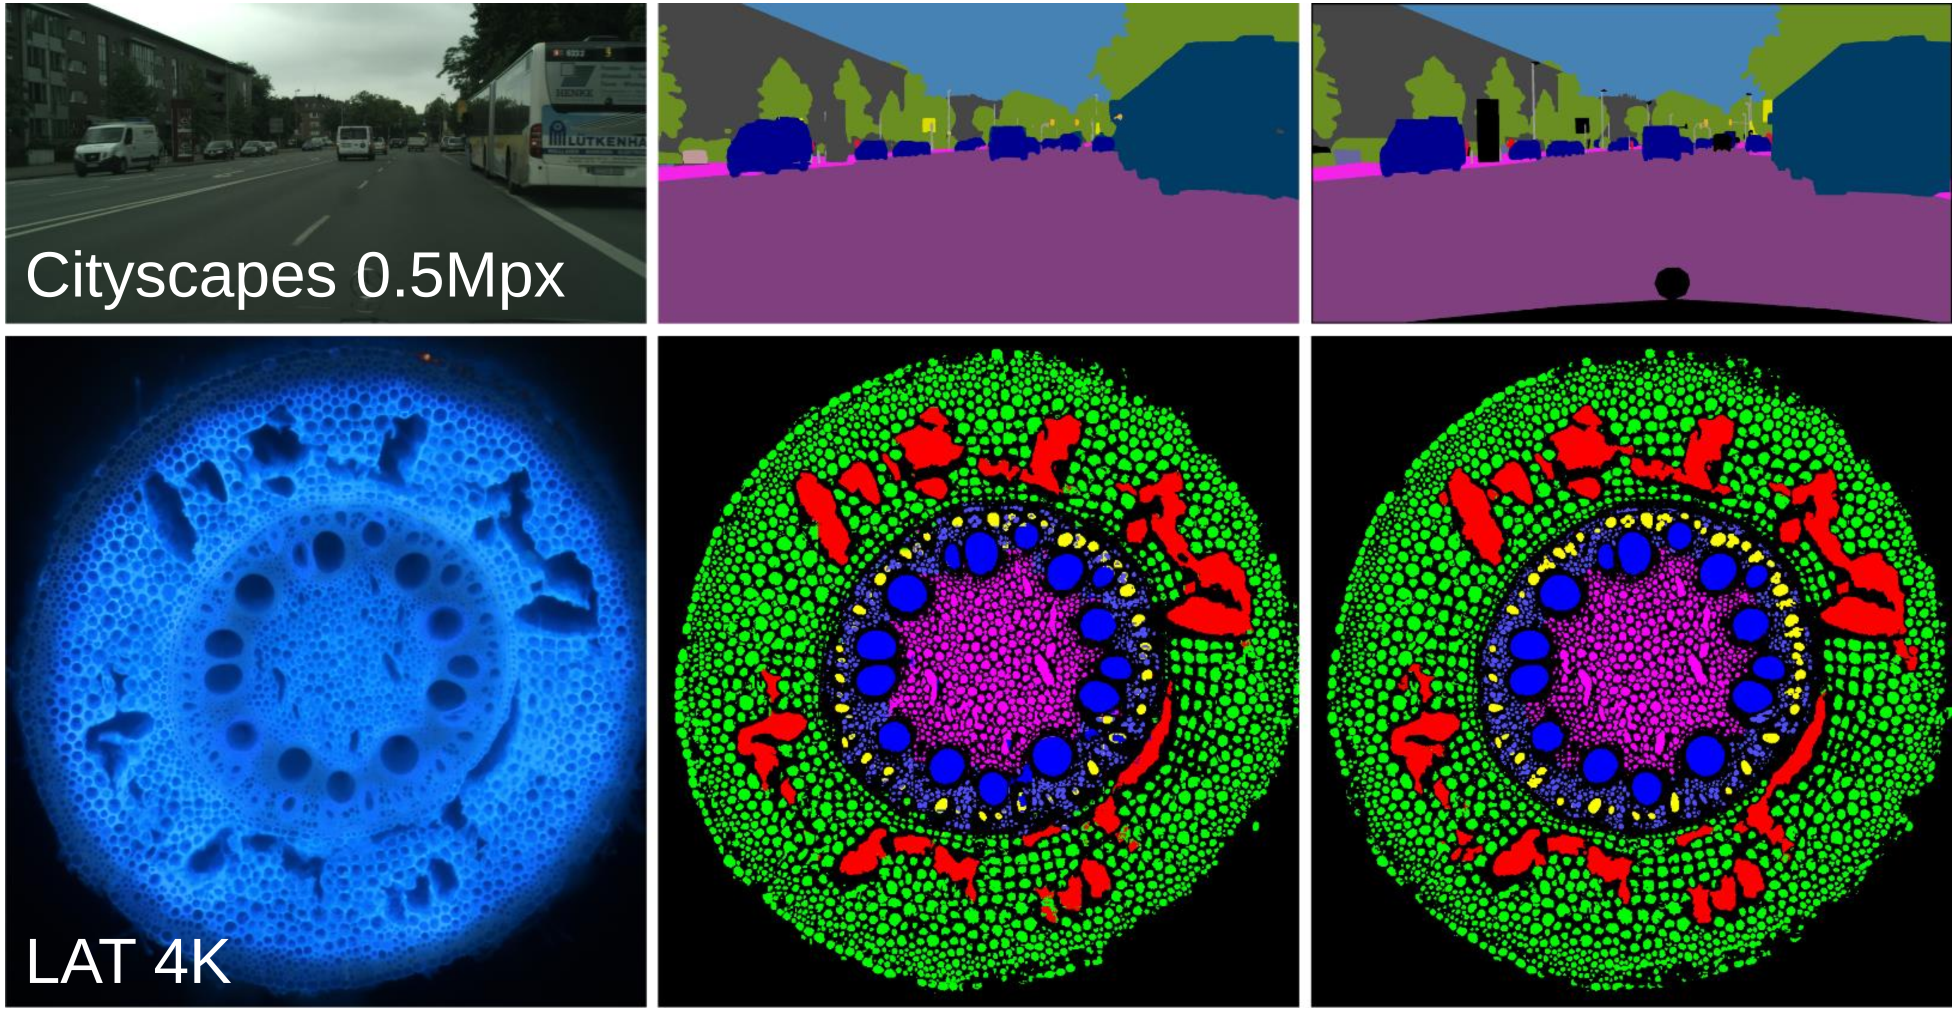

Source Results Groundtruth

36% Memory

27% Memory

High-resolution Segmentation

Standard network vs a tiled approach

DeepLab V3+	Accuracy Val (%)	Accuracy Test (%)	VRAM (GB) Train	VRAM (GB) Test
Standard	**76.2**	**75.6**	1.94 (1.0)	1.18 (1.0)
XT(ours)	75.0	74.1	0.70 (0.361)	1.00 (0.847)
Tiled(ours)	74.7	72.7	**0.69 (0.356)**	**0.95 (0.805)**

Tile size performance

Tile Size	Dice (%)	Resolution	Train VRAM (GB)
No tiling	**81.6**	1024×1024	8.43
512×512 XT	**81.6**	1024×1024	3.31
256×256 XT	79.8	1024×1024	2.52
128×128 XT	78.0	1024×1024	**2.32**
512×512 XT	**80.5**	2048×2048	7.70
256×256 XT	78.6	2048×2048	**6.41**
512×512 XT	-	3072×3072	-
256×256 XT	**78.0**	3072×3072	**11.9**

Higher Dimensions – Xray µCT

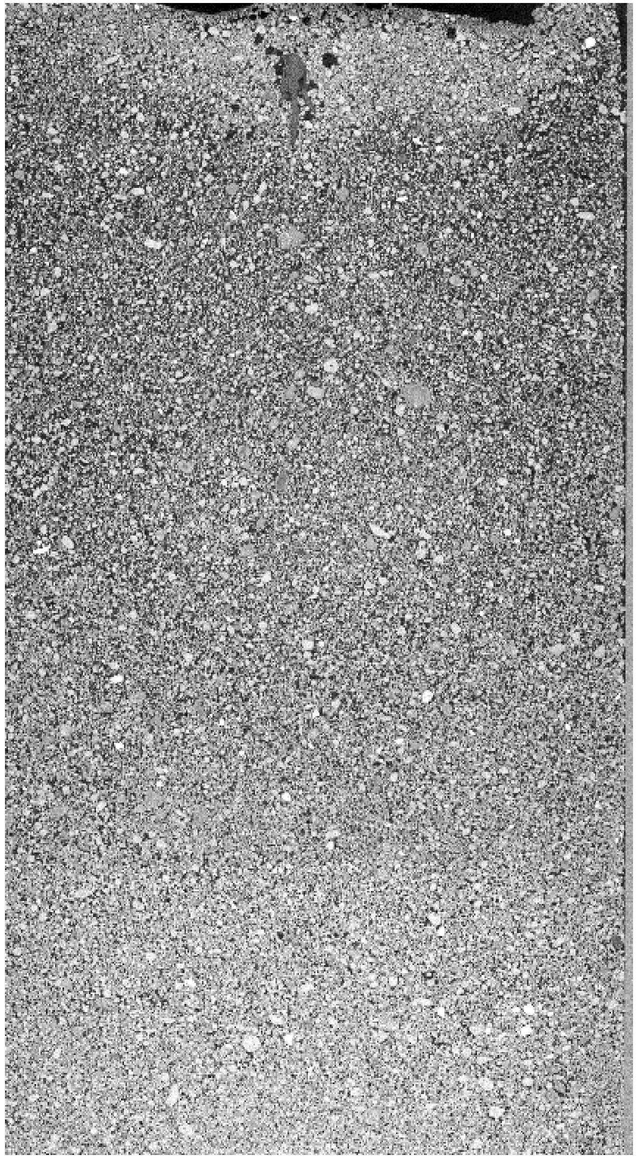
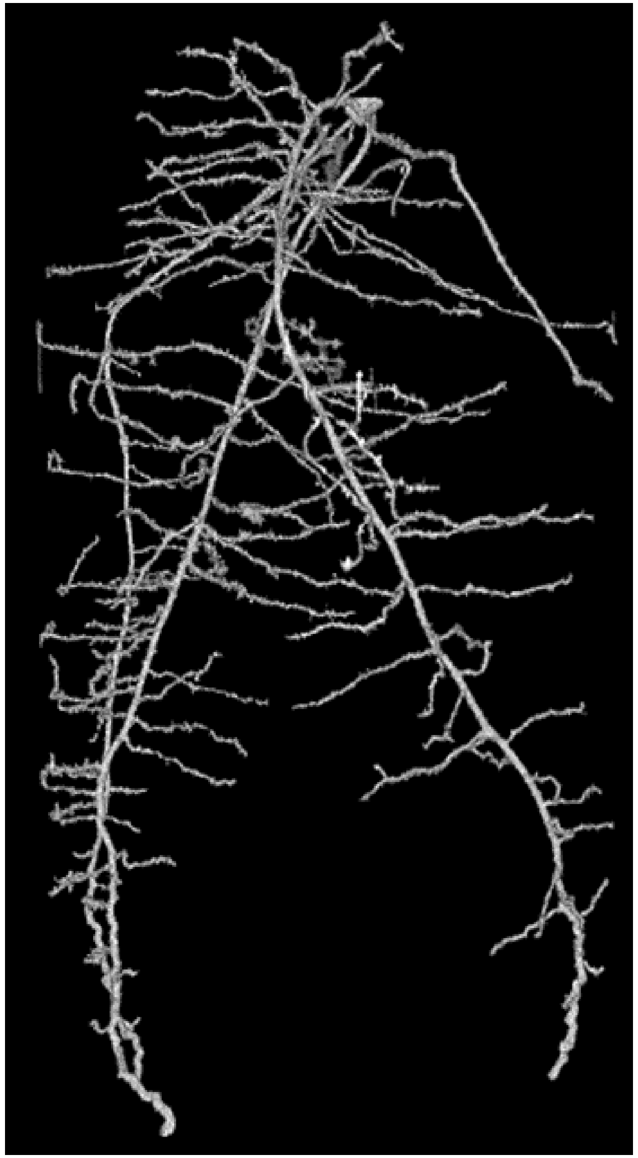
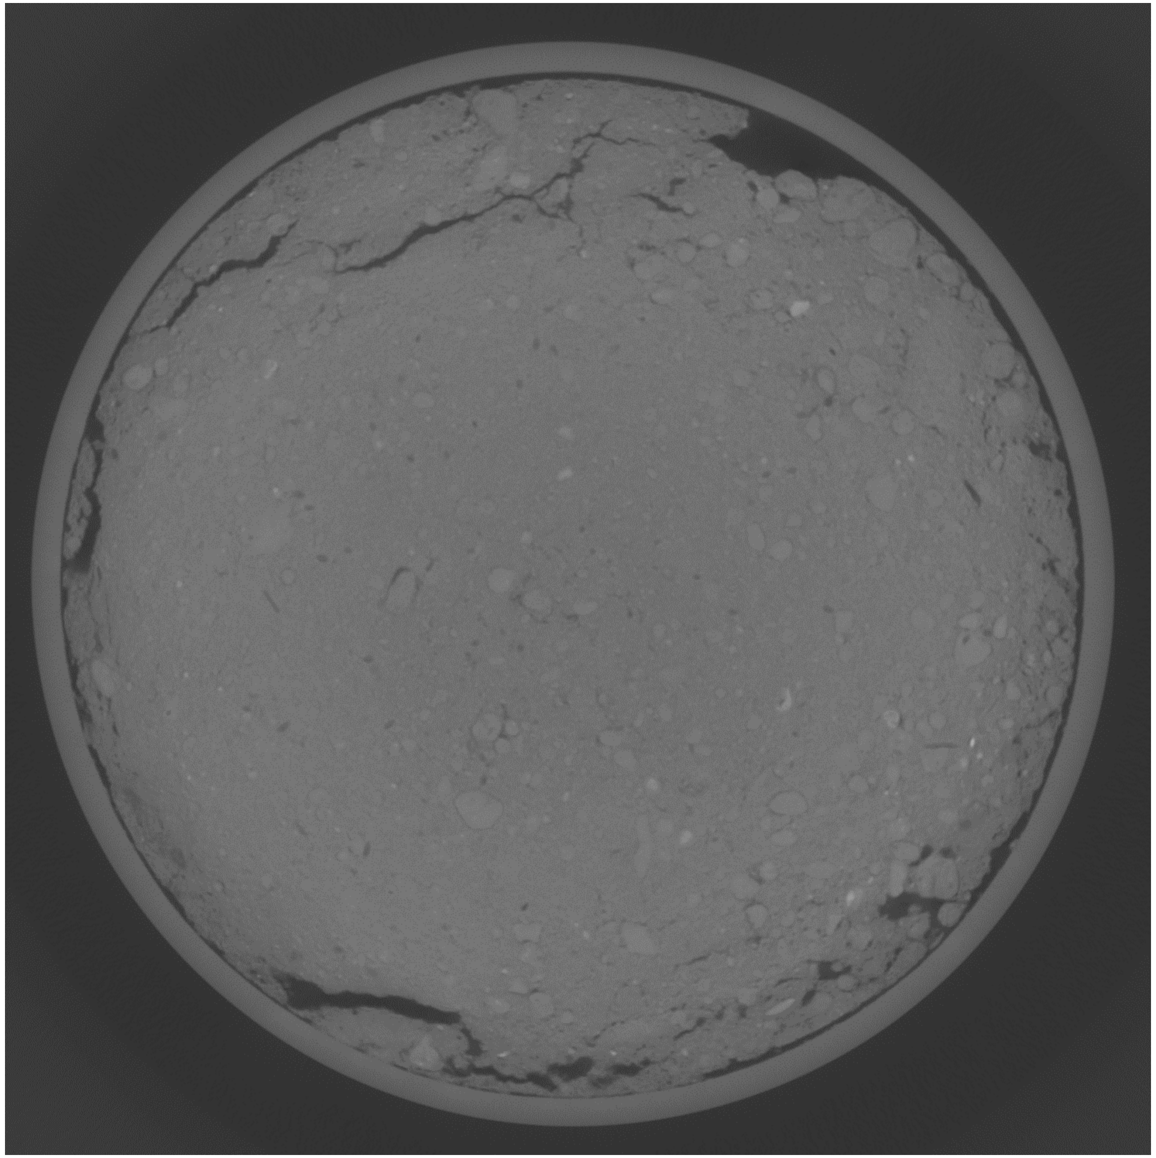

Xray µCT

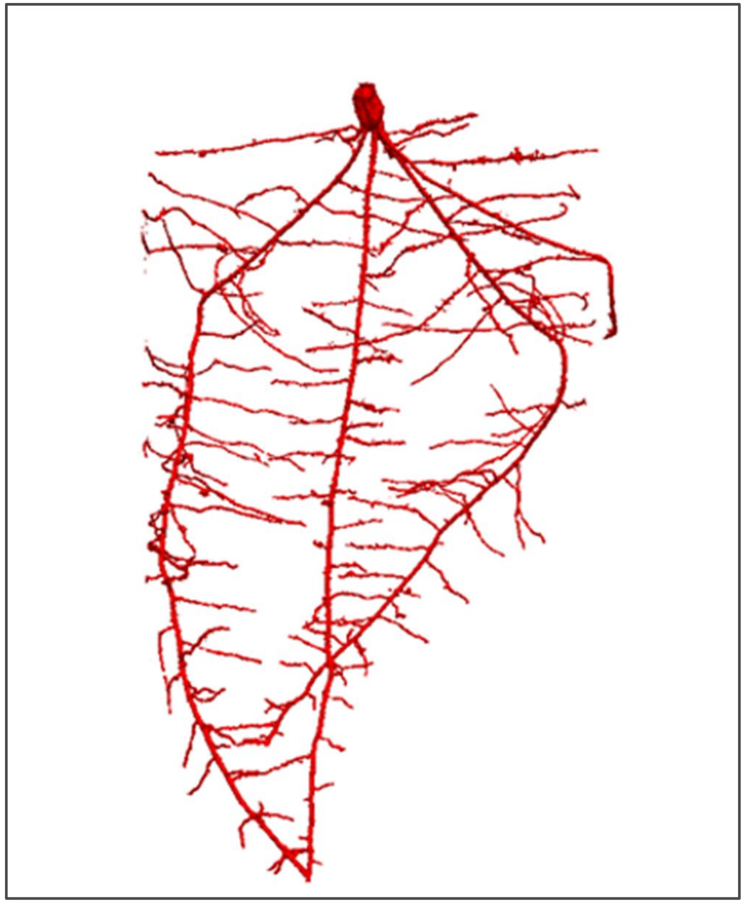

Ground Truth

16^3

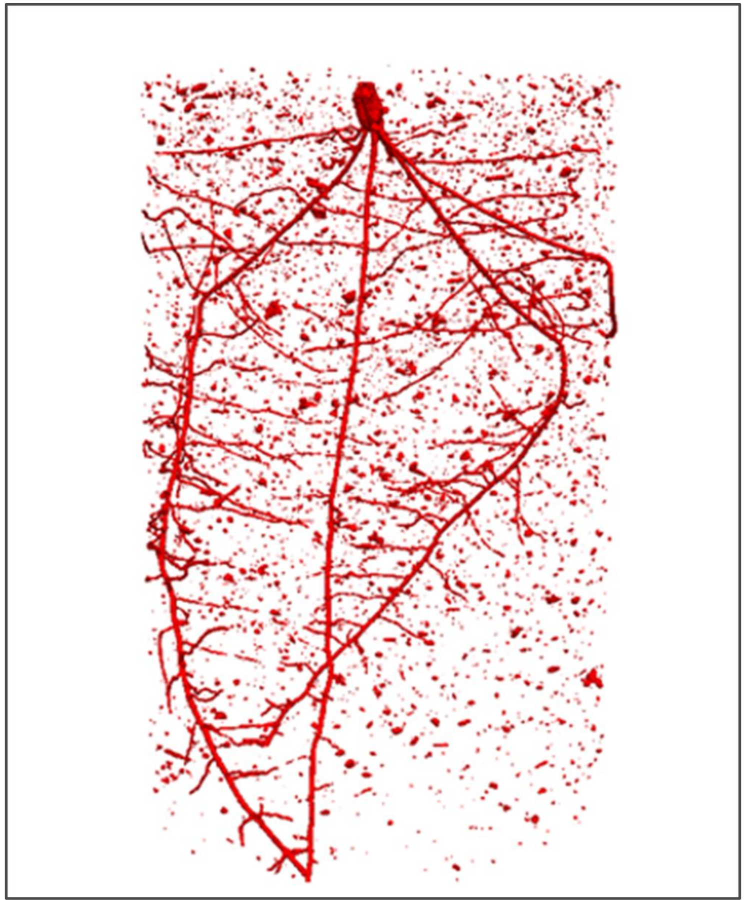

Fine Detail

128^3

Larger Receptive Field

Xray µCT

- The network jointly learns two input resolutions
- In essence this is another multi-task learning approach

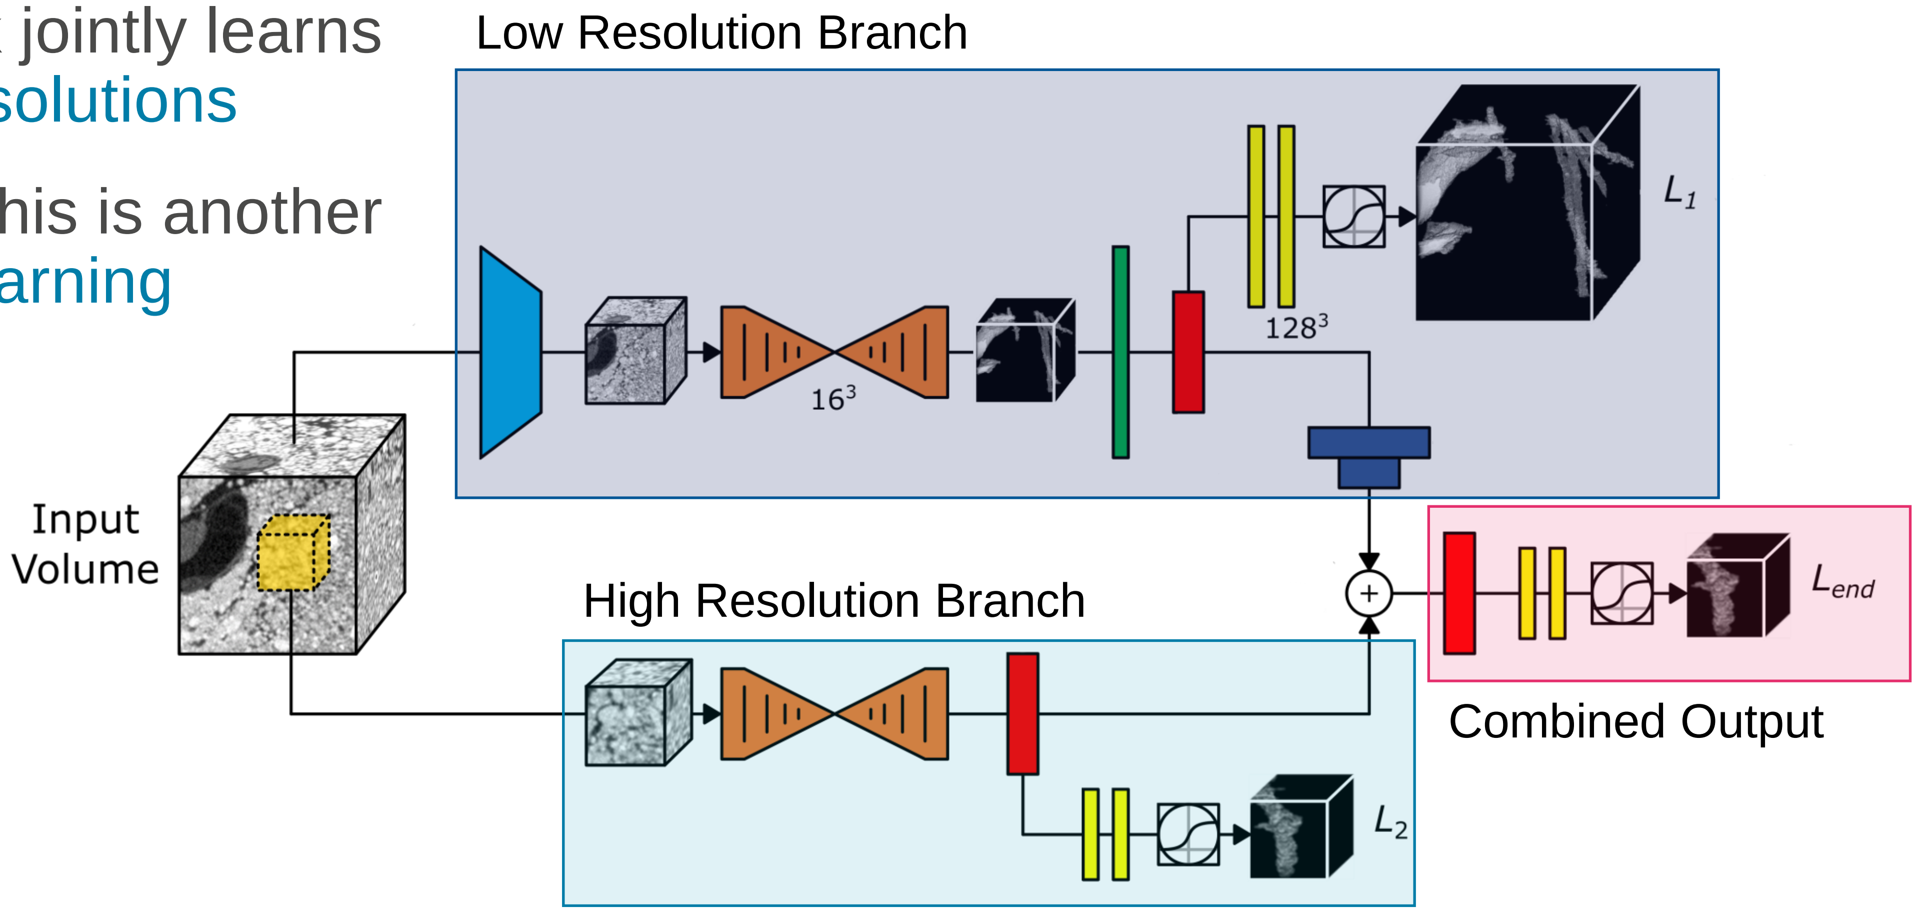

Three dimensional root CT segmentation using multi-resolution encoder-decoder networks. IEEE Transactions on Image Processing 2020

mpnd.uk/xray

Xray µCT

- Multi-resolution benefits from both views
- Hard negative mining further increases performance

16^3 128^3

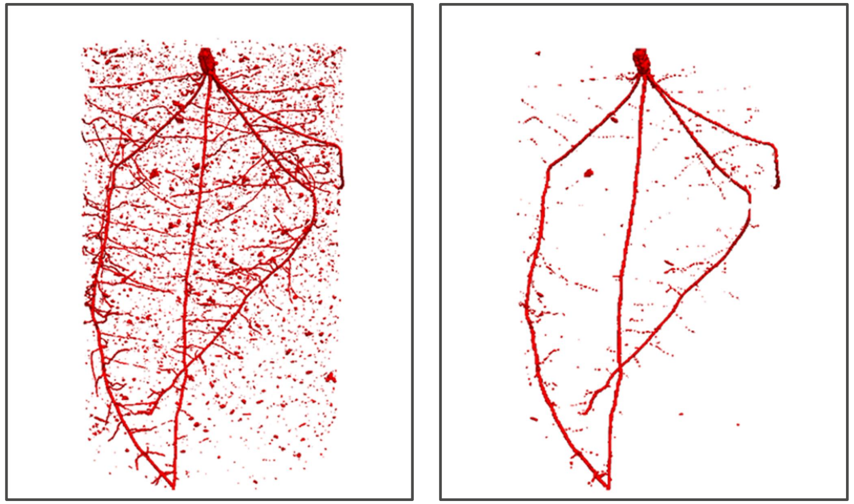

Fine Detail Larger Receptive Field

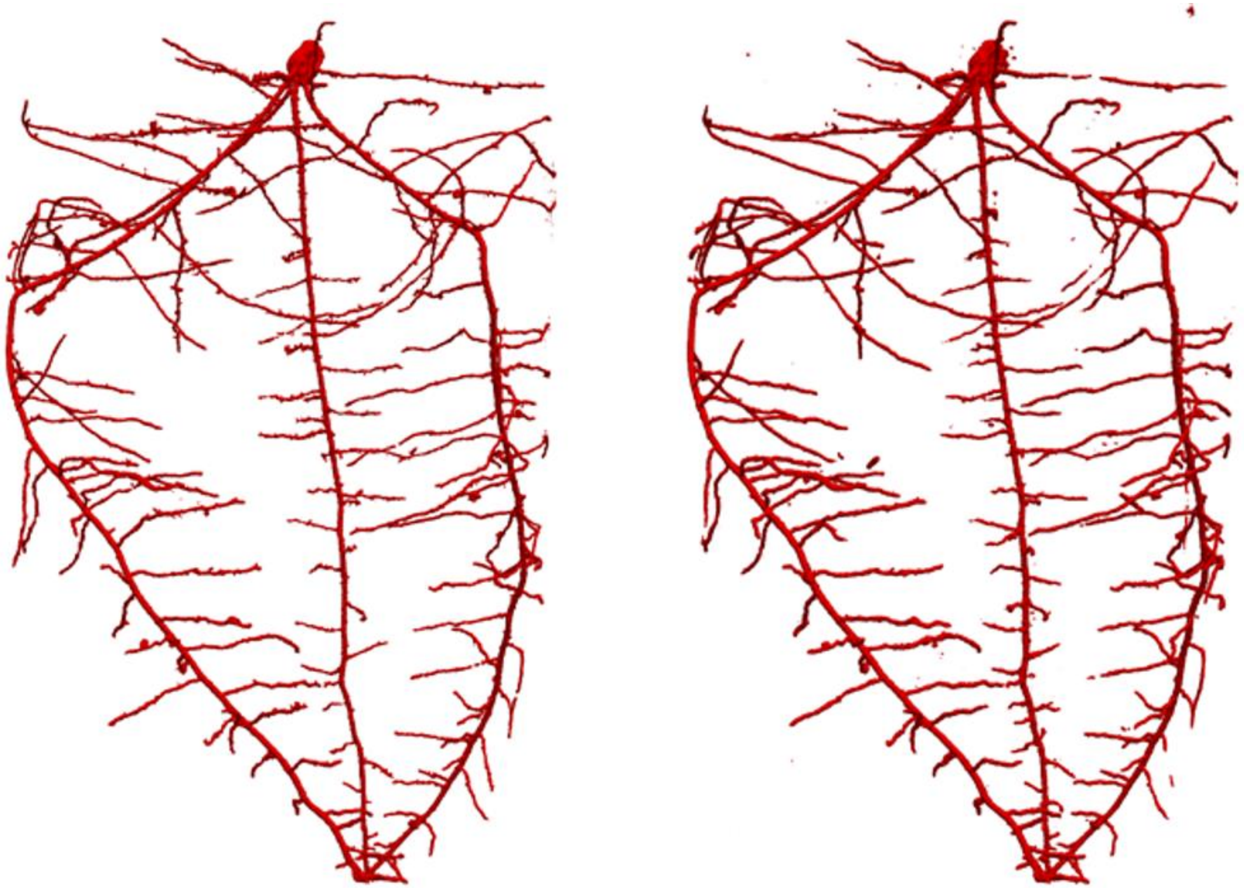

Ground Truth Hard Negative Mining

Looking Forwards

Current Challenges

- Modern computer vision and deep learning is not without its problems – particularly in challenging areas like Bioimaging

- Annotation is time consuming
- Bioimages are often much larger (spatial / dimensions) than CNNs are designed for
- Hardware cost is still a barrier for entry – at least for training
- The problem of generality

Looking Forwards

Data Annotation

- Many tools exist for fast data annotation but they tend to be quite general and slower for specific tasks
  - Unsupervised learning avoids annotation
  - Active learning reduces annotation time
  - Gaze-based solutions for easier annotation
  - Synthetic data for training

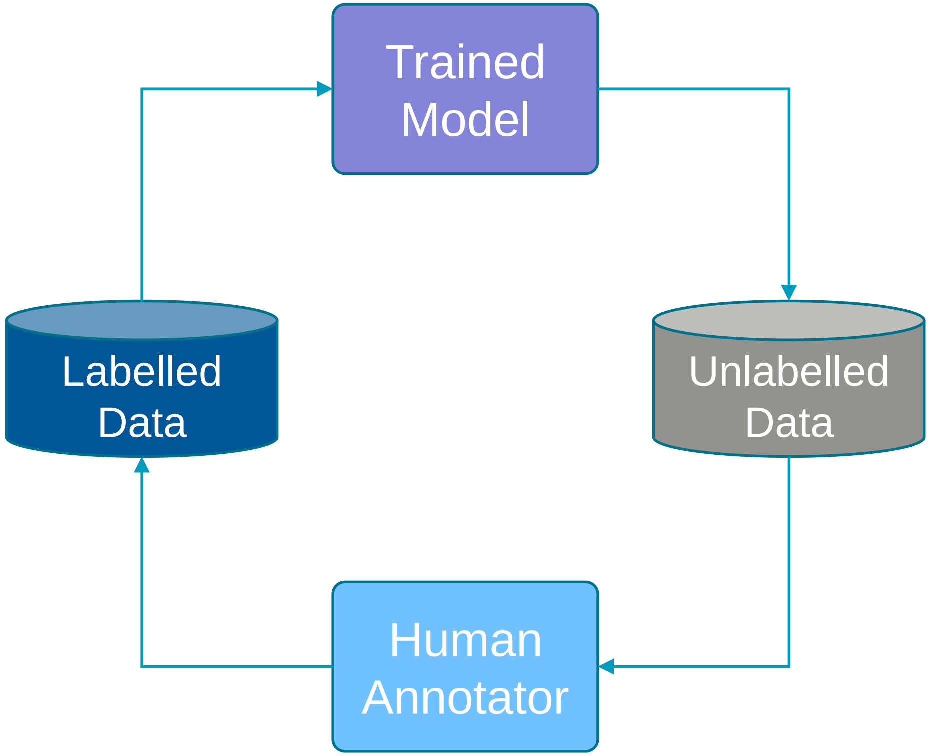

Looking Forwards

Data Annotation

- Many tools exist for fast data annotation but they tend to be quite general and slower for specific tasks
  - Unsupervised learning avoids annotation
  - Active learning reduces annotation time
  - Gaze-based solutions for easier annotation
  - Synthetic data for training

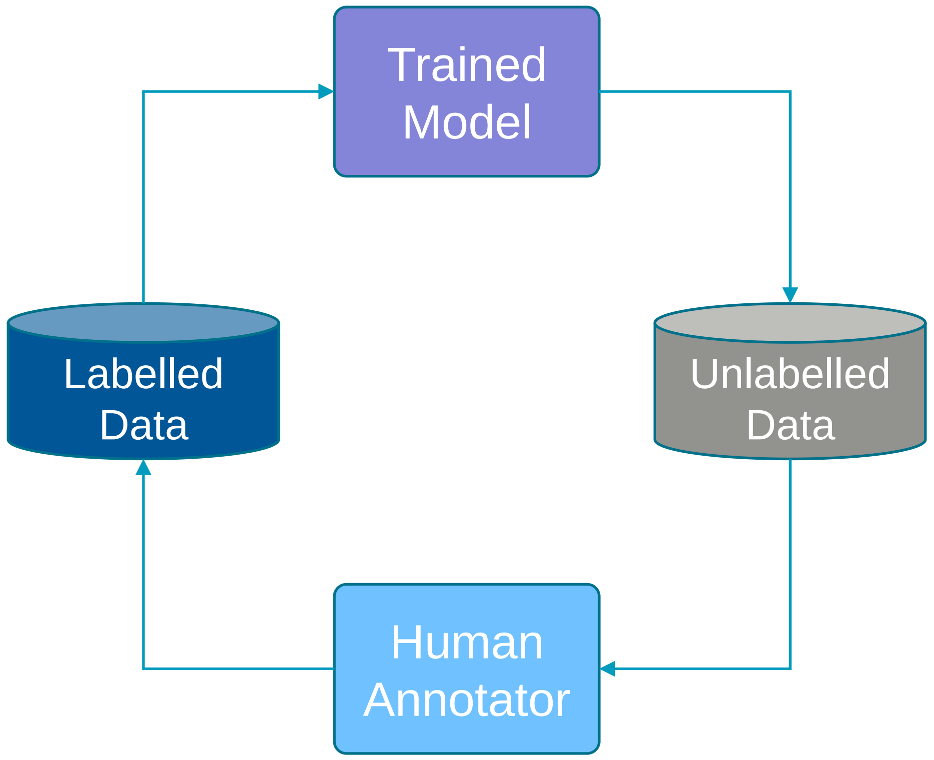

Looking Forwards

Data Annotation

- Many tools exist for fast data annotation but they tend to be quite general and slower for specific tasks
    - Unsupervised learning avoids annotation
    - Active learning reduces annotation time
    - Gaze-based solutions for easier annotation
    - Synthetic data for training

Looking Forwards

Data Annotation

- Many tools exist for fast data annotation but they tend to be quite general and slower for specific tasks
    - Unsupervised learning avoids annotation
    - Active learning reduces annotation time
    - Gaze-based solutions for easier annotation
    - Synthetic data for training

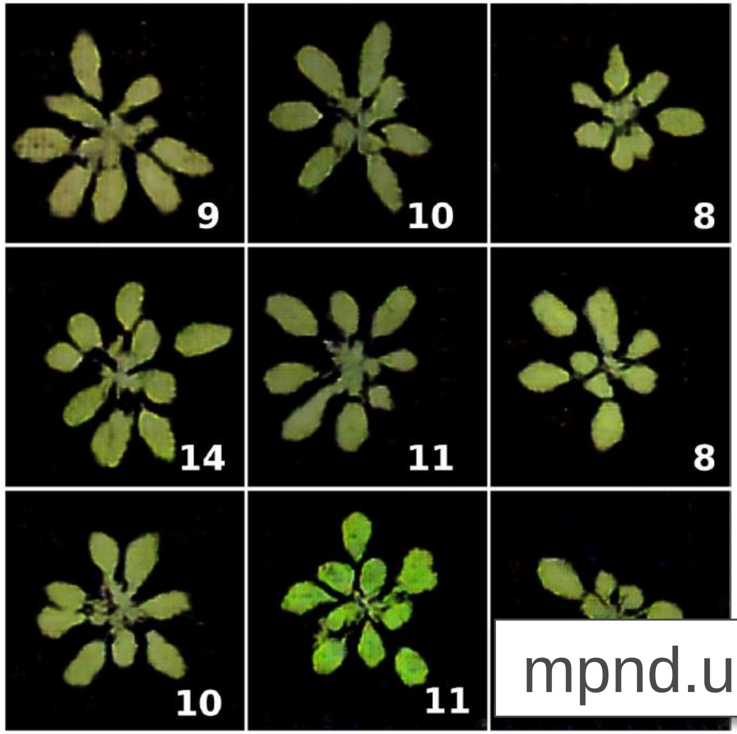

mpnd.uk/arigan

Giuffrida et al., ARIGAN: Synthetic Arabidopsis Plants using Generative Adversarial Networks. ICCVW 2017

Looking Forwards

Training on Challenging Data

- Bioscience has an abundance of hard data problems
- Larger images require more computational resources, and become a barrier for entry and reproducibility
- Problem areas are volumetric and multi/hyperspectral

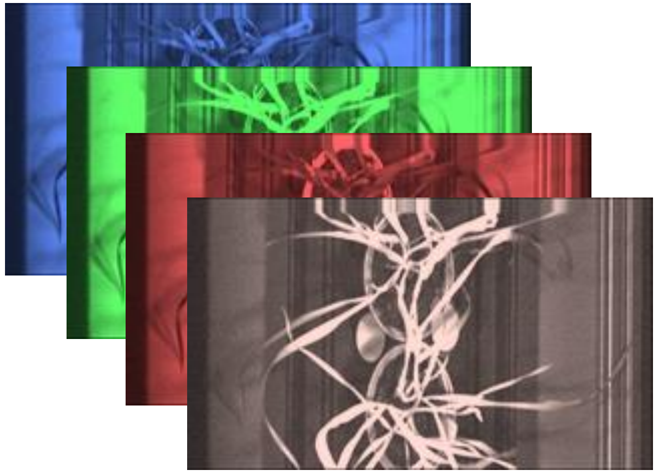

- Techniques exist to make trained networks efficient during inference
- Efficient training has seen very little work – most are case specific

Looking Forwards

Data and Model Sharing - Generalisability

- Deep learning still has a problem of generalisation
  - Models trained on one dataset will likely lose accuracy on another
  - Lab to field often requires retraining
- Large and varied datasets are important
  - Datasets such as the Global Wheat Challenge comprise varied images from around the world
- Providing access to trained models and code is a key part of this research

Conclusions

- AI is transforming the speed and accuracy of our plant phenotyping efforts
- Deep networks are powerful, but often the best performance comes from new architectures and techniques
- There now must be a renewed focus on collection and annotation of data, and model sharing

michael.pound@nottingham.ac.uk

- Acknowledgements

Tony Pridmore

Andrew French

Reza Soltaninejad

Eze Benson

Robail Yasrab

Aaron Jackson

Darren Wells

Jonathan Atkinson

Michael Wilson

Many thanks!