2nd Call for Funding: Dr. Michael Pound – Computer Vision Laboratory, School of Computer Science, University of Nottingham

October 8, 2020 (updated October 8, 2020)

Project Title

Predicting plant root growth from time-series data using deep learning

Total Fund Requested – **£24,677.34**

This pilot project will study the potential of deep networks to predict the growth of plants by generating segmentation masks of root systems *into the future*. We will adapt an existing predictive network, PGGAN, into this new domain. The result will be an efficient root segmentation network that can perform predictive segmentation of what a root system will look like at a future time, based on a sequence of images showing h it has grown so far. This work holds the potential to speed up plant experimental cycles by predicting traits prior to their measurement; predicting phenotypes and traits even a few days early would allow experimenta cycles to run more quickly and efficiently.

Introduction: Mapping and predicting plant growth is a complicated and challenging task [1]. Traditional image processing and simulation methods have seen a great deal of use in plant root analysis, and more recent machine learning-based systems have shown extraordinary accuracy. Computer vision and machine learning methods have performed well in static plant analysis, and have shown they are able to learn complicated plant growth patterns [2-3]. A number of deep learning-based approaches are available to efficiently measure different plant root traits, powering genetic discovery [4-5]. Such approaches offer a great deal of enhanced performance compared with traditional image-based phenotyping approaches, and biologists are relying more and more on these outstanding results to capture complex features and structures of plants both above and below the ground.

Recent performance on static image analysis has yet to be extended to time series data. Predicting plant growth is essential for greenhouse growers, biologists and plant scientist to analyze the future growth patterns. It helps to understand how a specific plant species behaves under different environmental and various stresses (biotic and abiotic) factors. It also holds the potential to speed up plant experimental cycles by predicting traits prior to their measurement; predicting phenotypes and traits even a few days early would allow experimental cycles to run more quickly and efficiently.

Temporal analysis of plant growth has traditionally been performed through modelling. Several tools and approaches exist for modelling and simulation of plant growth [6-7]. Most of these tools are based on deterministic factors to map the plant growth, and offer 2D/3D deterministic plant simulations [8-9]. The models are capable of producing seemingly realistic root systems, but cannot easily be used for genetic discovery, as parameterising these models with species and genotypic information is challenging, and ultimately involves phenotyping studies. These works are also challenging to generalise; forecasting the growth of plants is challenging in the presence of differences between species and even individuals.

What is required is a computer vision solution that can capture root traits, but is also trained to predict future root traits based on current growth patterns. Such a system would be capable of accurately measuring root systems, while also speeding up the experimental cycle by predicting RSA and traits ahead of time.

This research will explore the design and development of the plant root phenotyping and growth analysis system for time-series data, and will embed the realistic (but uncertain) nature of plant growth forecasting into our existing root segmentation approaches. Using an existing image dataset of *Arabidopsis thaliana*, imaged over time, the project will explore the ability of deep networks to both capture plant root traits, and also predict these over time. The system will offer high-throughput root phenotyping capability, including Root System Architecture (RSA) segmentation, and direction and velocity of growth prediction.

This project will extend our existing work on RootNav 2 into temporal sequences, by incorporating a Generative Adversarial Network (GAN) [10] that maps plant growth patterns and forecasts plant growth.

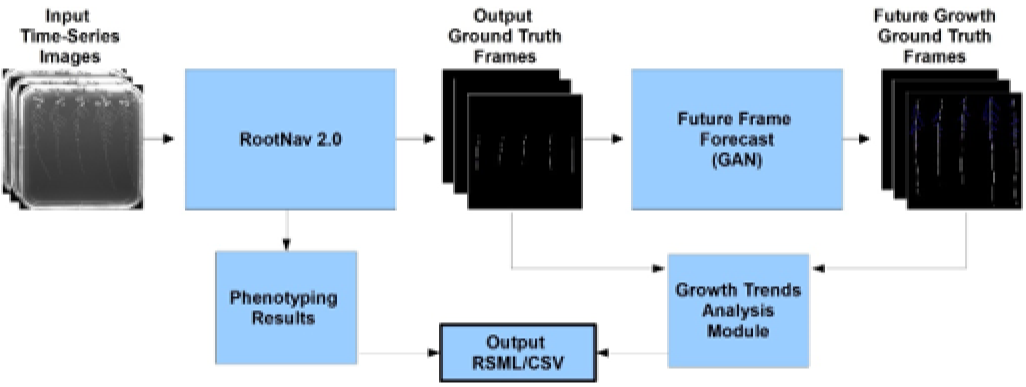

Figure 1 Proposed Architecture of Plant Root Phenotyping and Growth Analysis System for Time-Series Data: In this model, the Root-Nav 2.0 will extract phenotype details from root images and produce the GT frames which will be fed to a “future forecast” GAN module. This module will generate possible future growth patterns of a given set of inputs.

Figure 1 Proposed Architecture of Plant Root Phenotyping and Growth Analysis System for Time-Series Data: In this model, the Root-Nav 2.0 will extract phenotype details from root images and produce the GT frames which will be fed to a “future forecast” GAN module. This module will generate possible future growth patterns of a given set of inputs.

Proposed Methods: This project will explore the use of Generative Adversarial Networks (GANs) to capture both the spatial and temporal features of plant growth. GANs, introduced by Goodfellow et al. [10], are typically based around convolutional neural networks. GANs are composed of two networks that learn together, the generator and discriminator. Both networks compete against the other (hence “adversarial”), to produce new (synthetic) examples of data that can pose as real data. In this work the encoder “generator” will learn the patterns of growth for plant roots up to the current time, and try to predict those growth patterns by producing synthetic future frames, predicting the future growth of roots. The proposed approach is inspired by progressively growing GAN (PGGAN) [11]. PGGAN was initially designed to produce high-definition images from a set of random latent variables. The proposed network will be composed of a generative encoder model that will be trained to produce high definition future forecast images. The discriminator decoder will be trained to reduce the differences between real and synthetic (forecast) frames. This architecture will be available as an add-on to the existing RootNav 2.0 system.

The project will use images of plant growth captured over time. An existing segmentation network [4] will be used to extract detailed segmentation information for root systems, before a GAN is used to predict the growth of each plant over time. Our preliminary work shows this approach can produce realistic looking images of future frames, this project will also evaluate whether these are sufficiently accurate to produce biologically meaningful results. Images, whether captured or synthetic, must then be reliably quantified. Image noise is a key concern with prediction of real plant growth into the future, thus we will focus here on the prediction of segmentation masks rather than raw images in order to provide the cleanest and most reliably quantified output.

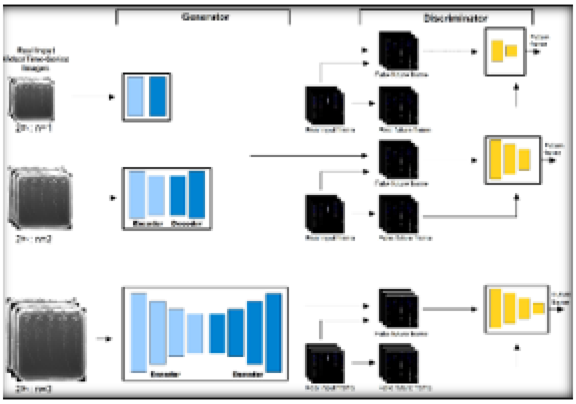

Figure 2 GAN generator and discriminator Modules designed to forecast future frames based on the early growth of plant roots.

Preliminary Data: As a demonstration of the efficacy of this GAN-based approach for this work, we trained GAN on an existing dataset of plant shoots captured over time. We used the publically available plant dataset by Uchiyama et al. [12]. This dataset, composed of Komatsuna plant (Japanese mustard spinach), comprises leaf growth captured from three high-resolution fixed cameras. Per leaf segmentation masks exist for training the network. The available resolution of images is 480×480, which we reduced to 128×128 for these initial experiments.

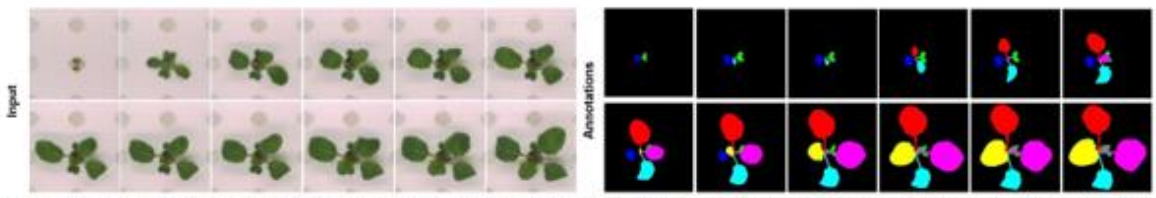

Figure 3 Example images from the datasets used during this work. This is a time-series dataset frame taken during the life-time of Komatsuna plant

Figure 3 Example images from the datasets used during this work. This is a time-series dataset frame taken during the life-time of Komatsuna plant.

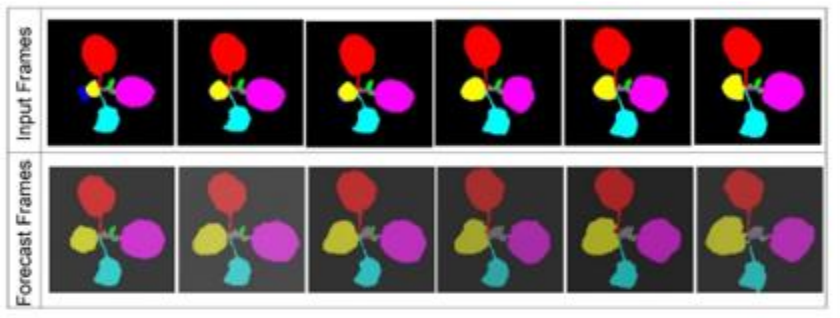

Figure 4 Sample Output from Trained GAN: The above column shows the input set of image frames. The below column shows the forecast frames from the trained GAN architecture. This GAN architecture trained to take six input frames to forecast six future growth frames.

This project will train a GAN to learn past plant growth trends and forecast future plant growth patterns based on those trends. In our initial experimentation of Komatsuna plant leaves, we used the annotated seg-mentation mask of leaves to train the GAN. Figure 4 shows the output results of our preliminary study, showing a high degree of similarity between the predicted frames and the ground truth. The original annotations may be seen in Figure 3. This project will adapt and retrain this technique for plant root images. Plant roots offer a more complex topological structure and continue to be an area of active study in image analysis, machine learning and plant phenotyping. All network designs and trained models will be released following the project, alongside the annotated dataset described below.

Experimental Design: We will use the Arabidopsis Thaliana root dataset of Wells et al. [13] during this project. Examples of these images may be seen in Figure 5. The dataset comprises 47 plant plates with each plate containing five plants. This time-series dataset was captured using near infrared and includes ten images of each plate imaged over a number of days. Different genetic accessions of Arabidopsis thaliana such as Col-0, Ler and Cvi were grown, however for this project we will focus on the general task of segmentation rather than identification. This dataset does not include segmentation masks, so we will produce initial masks using RootNav 2.0 [4], before alteration by an expert if necessary. The segmentation masks for each sequence will then be used to train the GAN for growth prediction. We will explore adaptations of the PGGAN network appropriate to this task, including an increase in output resolution beyond 128^2 pixels in order to promote higher accuracy of the eventual quantified root systems. We will evaluate the trained network against the known ground truth of images captured of later growth staged, to ensure that all generated images are biologically useful rather than simply looking plausible to a human observer.

Figure 5 Sample images from the experimental dataset.

Figure 5 Sample images from the experimental dataset.