

D1.3: Detailed description of “proof-of-concept” studies of dynamic phenotyping based on specific case-studies

*Tony Pridmore (UNOTT), Franck Golbach (DLO),
Llorenç Cabrera-Bosquet (INRAE), Xavier Draye
(UcL)*



This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731013. This publication reflects only the view of the author, and the European Commission cannot be held responsible for any use which may be made of the information contained therein.

Document information

EU Project N°	731013	Acronym	EPPN ²⁰²⁰
Full Title	European Plant Phenotyping Network 2020		
Project website	www.eppn2020.plant-phenotyping.eu		

Deliverable	N°	D1.3	Title	Detailed description of “proof-of-concept” studies of dynamic phenotyping based on specific case-studies
Work Package	N°	JRA1	Title	<i>Novel techniques and methods for environmental and plant measurements, including model-assisted phenotyping</i>

Date of delivery	Contractual	29/04/2020 (Month 36)	Actual	29/06/2020 (Month 38)
Dissemination level	X	PU Public, fully open, e.g. web		
		CO Confidential, restricted under conditions set out in Model Grant Agreement		
		CI Classified, information as referred to in Commission Decision 2001/844/EC.		

Authors (Partner)	Tony Pridmore (UNOTT), Xavier Draye (UcL), Franck Golbach (DLO), Llorenç Cabrera-Bosquet (INRAE)			
Responsible Author	Name	Prof Tony Pridmore	Email	tony.pridmore@nottingham.ac.uk

Version log			
Issue Date	Revision N°	Author	Change
29/6/2020	1	Tardieu	Edition by coordinator

Executive Summary

Objectives

The objectives of D1.3 are:

- to investigate and report on the potential of active imaging strategies to improve the data obtained from image-based plant phenotyping methods
- to present the results of proof of concept studies of dynamic phenotyping methods conducted within the consortium
- to propose methodologies for generalising these studies and identifying control strategies suitable for different types of high-throughput phenotyping installations.

Rationale: To date, work on high-throughput, image-based phenotyping has focussed on automation, the plants under investigation being brought before a camera, or cameras, by an automatic conveyor, images are captured, analysed and traits extracted. The same process is applied to each plant and image acquisition and analysis are independent: images are acquired from fixed camera locations and passed through a pre-determined analysis pipeline. That pipeline has no ability to control or influence the image acquisition process.

While this 'one size fits all' approach can be effective, it has limitations. Plants are complex and highly variable objects, and fixed camera systems cannot be relied upon to provide the most appropriate images in all situations. Dynamic phenotyping approaches owe more to robotics than automation. Here, image acquisition and analysis are integrated. Fixed cameras provide an initial survey and description of the plant, after which they or other cameras are automatically placed to acquire the images needed to perform the task at hand.

The hypotheses underlying this deliverable are that:

- i. the core technologies of robotics and plant image analysis are now sufficiently well-developed as to allow widespread use of dynamic phenotyping methods
- ii. the dynamic phenotyping approach has the potential to produce more accurate and consistent phenomic data than existing automated methods.

Main Results: This document first presents the case for dynamic phenotyping based on robotics as an alternative to existing methods based upon longer-established automation technologies. Section 1 outlines the background to the study and the hypotheses underlying the work done.

For a camera to be automatically positioned relative to a plant, the 3D position and orientation of the camera must be known and relatable to some plant feature(s). This can be achieved in two ways, either:

- i. a complete 3D model of the plant can be recovered and related to the 3D position of the camera or
- ii. biological knowledge can be embedded in image analysis methods which identify and recover the 3D position of only selected features of particular scientific interest.

Section 2 outlines work done on improved, high-resolution recovery of complete 3D plant models from multiple camera views, and on active 3D reconstruction. Two groups have taken complimentary approaches. DLO has developed a facility which produces a high-quality volumetric model of the viewed plant. This is effectively a 3D image: the space containing the plant is represented as a 3D grid of "voxels", with each being labelled as containing, or not containing, plant material. UNOTT also employs a volumetric model, but of lower quality, using it not to represent the plant but only to drive the camera. The images obtained are then combined to produce a surface-based representation of the plant. These activities are currently

being integrated, with the ability of the volumetric models produced in the Netherlands to better support active camera placement in Nottingham the focus of the work.

Section 3 outlines work done at INRAE on camera placement for maize silk analysis without access to a complete 3D model of the plant. Here multiple cameras are employed and calibrated using a technique related to that described in D1_1. Independent analysis of key views allows other cameras to be selected/rejected for use when examining the plant at hand and a separate x-y-z controllable camera to be placed to acquire a final image in which small, but biologically important features of the plant are clearly visible. This approach is currently being modified to a grape phenotyping task at INRAE, while a similarly structured approach is being taken to phenotyping aeroponically grown root architectures at UCL. Protocols and code for the INRAE maize system are already available. These represent the beginning of a library of scripts which will be extended in the final stages of EPPN²⁰²⁰.

Because they are front-of-science and need to be developed in day-to-day collaboration with scientists involved in their application, the activities described above have been carried out separately, but with frequent exchange, in different nodes of the consortium. In the final year of EPPN²⁰²⁰, the data and techniques produced will be brought together, transformed into software elements made available to and tested by the consortium of EPPN²⁰²⁰. Some elements are already available to the whole phenotyping community via public repositories. This will be the case at the end of the project for all other applications presented here.

Authors/Teams involved: Tony Pridmore (UNOTT), Llorenç Cabrera-Bosquet (INRAE), Franck Golbach (DLO), Xavier Draye (UCL).

Table of contents

Document Information	2
Executive Summary	3
Table of Contents	5
1. Introduction	6
2. 3D Reconstruction and Dynamic Phenotyping	8
2.1. 3D Reconstruction from Multiple Images	8
2.2. Improved 3D reconstruction by Space Carving (DLO)	9
2.3. Active 3D Reconstruction (UNOTT)	10
2.4. Next Steps (DLO, UNOTT)	12
3. Dynamic Phenotyping of Selected Plant Organs	12
3.1. Maize Silk Phenotyping (INRAE)	12
3.2. Next Steps (INRAE, UCL)	14
4. Conclusion	14

1. INTRODUCTION

The goal of high-throughput plant phenotyping is to provide accurate, objective measurements of the structure and/or function of large numbers of plants as efficiently as possible. This objective is at best difficult and often impossible to achieve using manual measurement tools or methods requiring high degrees of operator intervention. Automation is the key, and current large facilities are built upon concepts developed in automated manufacturing. Plants are often maintained on conveyors, travelling from one workstation to another as they undergo a series of fixed processes; RGB imaging, fluorescence imaging, laser scanning, etc. As in automated manufacturing, each process is clearly defined and all plants receive the same treatment.

In the standard phenotyping pipeline, characterised in Figure 1, the plant is placed in the field of view of some imaging device or sensor, raw data is acquired and analysed to provide trait data which is inserted in an appropriate information system. Here, image acquisition and analysis are independent. The analysis performed is tuned to the broad class of images that are expected, e.g. side views of a plant of an approximately known size against a known background, but not to the particular plant in view. This ‘one size fits all’ approach can clearly be effective, but it assumes a degree of uniformity in the objects considered - plants – which may not be present.

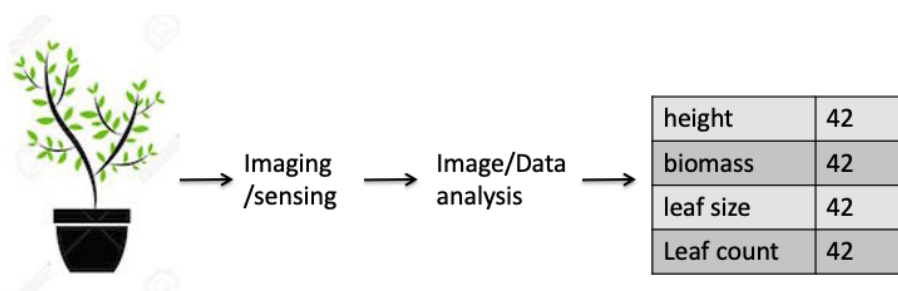


Figure 1. Overview of a standard plant phenotyping pipeline.

Dynamic phenotyping methods differ in that they seek to adapt the sensing strategy to the plant at hand. Rather than take images from the same set of viewpoints, regardless of the size, shape and complexity of the object presented to them, they seek to acquire the images needed and ideally best suited to support those measurements. To achieve this requires a step away from automation and towards robotics; the key distinction being that while automated systems repeat the same operation multiple times, robotic systems sense their environment and take images accordingly. Figure 2. characterises a typical dynamic phenotyping pipeline.

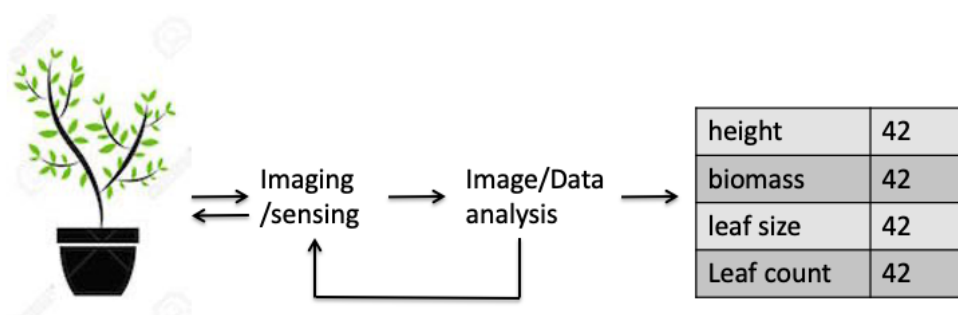


Figure 2. Dynamic plant phenotyping.

Here, image acquisition and analysis are integrated, in that a feedback loop exists between them. The process generally begins with an initial image, or set of images, gathered in the same way for each plant. These are not, however, the images used in the final analysis. Information recovered from the initial images is used to decide where camera(s) should be placed to best make the desired measurements. Further images are then acquired from those locations, analysed and, if necessary, the process repeated until sufficient images of sufficient quality are available to allow the desired traits to be recovered.

The dynamic phenotyping approach has the potential to:

- Reduce variations in trait values due to the (often arbitrary) positioning of the plant in front of the camera, leading to more robust and repeatable phenotyping
- Support more detailed analyses. The initial camera locations are often arranged to capture the whole plant, providing an overview image at comparatively low-resolution. These can then be used to move cameras closer to the plant to acquire higher-resolution images of smaller sections, perhaps individual organs. Data obtained from these close-up images is likely to be more fine-grained and accurate than that obtained from the more distant whole-plant views.
- Support multi-scale phenotyping. The dynamic approach offers the possibility of linking data obtained from whole-plant and selected close-up views to produce a richer representation in which detailed measurements of e.g. individual leaves are placed in the context of the surrounding plant architecture.

To obtain these benefits, however, the 3D position and orientation of the camera must be known and relatable to some plant feature(s). This can be achieved in two ways, either:

- i. an explicit 3D model of the plant must be recovered and related to the 3D position of the camera or
- ii. biological knowledge can be embedded in image analysis methods which identify and recover the 3D position of only selected features of particular scientific interest.

The work carried out in EPPN²⁰²⁰'s JRA 1.2 and described here comprises a number of case studies of the dynamic phenotyping methods which are now being drawn together. These span the two approaches to 3D information outlined above. Section 2 outlines work done at DLO on improving the 3D models available to drive dynamic approaches and at UNOTT on active approaches to 3D reconstruction. Section 3 focusses on work at INRAE and UCL which takes approach (ii), basing camera control on deeper analysis of 2D images acquired from known locations. Though the distinction is not a hard one, the work reported in section 2 is oriented towards reducing variation in phenotyping data, while that in section 3 is motivated by interest in multi-scale phenotyping. More detailed analysis is a goal shared across the consortium. Following description of the independent projects, plans for the remainder of the project are summarised, before conclusions are drawn in Section 4.

2. 3D Reconstruction and Dynamic Phenotyping

2.1. 3D Reconstruction from Multiple Images

Three-dimensional information can be recovered from multiple 2D images in two ways. In the first, points of interest are identified independently in each image and matched between views. If it can be established that feature A in image 1 depicts the same physical point as feature B in image 2 – the features here may be leaf tips, for example – then the 3D position of that physical point can be obtained by triangulation (Figure 3). The process requires accurately calibrated cameras, as the relative positions and orientations of the cameras used must be known, but once this information is obtained the transformation from multiple 2D features to one 3D feature is straightforward. 3D models produced in this way are formed of sets of 3D points lying on the visible surfaces of the object under consideration.

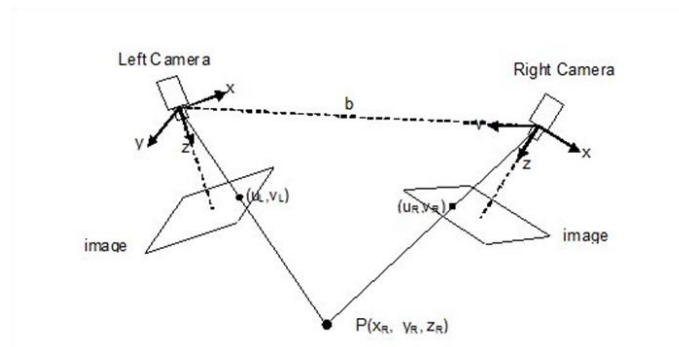


Figure 3. 3D reconstruction by point matching

The second approach, known as shape from silhouette or space carving, produces a very different, volumetric representation of the object being modelled. This is effectively a 3D image: the space containing the plant is represented as a 3D grid of “voxels”, with each being labelled as containing, or not containing, plant material. Space carving also requires the cameras involved to be calibrated, so that their relative positions and orientations are known, and proceeds by combining the results of analyses of individual images. These analyses, and the form of their combination are however different. Each image is segmented, to separate the object of interest (e.g. the plant) from its background. The result is a set of silhouettes of the plant, which are then projected into the voxel grid. Any voxel lying outside any projected silhouette is labelled as not containing object (plant) material (Figure 4).

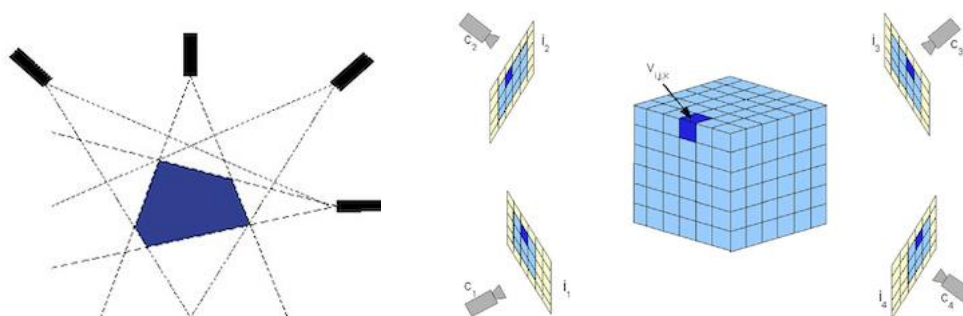


Figure 4. 3D reconstruction by space carving.

2.2. Improved 3D reconstruction by Space Carving (DLO)

Work on the recovery of 3D descriptions of plants at DLO is based upon space carving, and so produces a volumetric description. One of the challenges of the approach is achieving sufficiently high resolution at high throughput. The processing time required to produce a model depends mainly on the size, and so number, of the voxels making up the 3D representation. Increasing resolution significantly increases processing time; reducing resolution allows faster processing but risks producing a model which lacks fine details. This is particularly important when using the technique in plant phenotyping, as many important plant components are small. A specific concern is that stalks will be lost, and the connectivity of the plant will be reduced. Rather than produce a single plant volume, reduced resolution space carving quickly leads to a fragmented model from which traits are difficult, if not impossible, to extract.

DLO's work in JRA1.2 has focused on

- i. increasing the resolution of the 3D acquisition process, so that even the finer structures remain connected
- ii. Improving and developing methods for automatic analysis of connectivity and segmentation

A new space carving facility (Figure 5) has been designed and implemented employing 15 high resolution colour cameras. The installation is able to model larger plants (40 x 40 x 70 cm) at ultrahigh resolution. The representations produced comprise a grid of 1400x800x800 voxels, describing the plant at a resolution of 0.5 mm, which is sufficient to maintain the connectivity of the plant.



Figure 5. The Wageningen large plant, space carving facility

Methods have also been developed for making this connectivity explicit, by skeletonising the voxel representation. This process locates points at the centre of plant components, e.g. stalks, and links them together to produce a simpler representation of the core structure, or skeleton, of the plant. Figure 6 shows an example of the data produced. To the left is one of the input images used; the central graphic shows the 3D volumetric representation, upon which the skeleton is overlaid in the rightmost illustration.

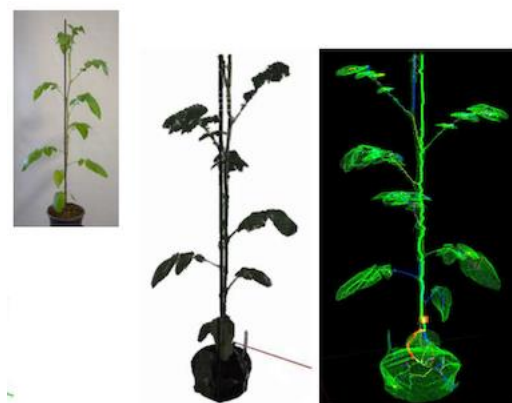


Figure 6. Results from the Wageningen space carving installation. Left: an input image, centre: the volumetric description, right: skeletonization.

Current work in Wageningen is directed towards further analysis of the volumetric representation. Techniques are being developed which can automatically identify plant organs, e.g. leaves, in the 3D data, exploiting the connectivity information made available in the skeleton. It is anticipated that this will provide more accurate and robust segmentation than previous algorithms which rely on voxel data alone. The skeleton will allow correct interpretation of the more complicated situations which can arise, e.g. when leaves touch, causing loops to appear in the plant volume. Figure 7 shows initial results.

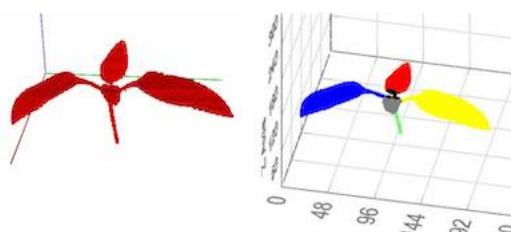


Figure 7. Initial plant segmentation results, on a small plant.

2.3. Active 3D Reconstruction (UNOTT)

Work on the recovery of 3D descriptions of plants at UNOTT is based upon point matching, and so produces a description of visible surfaces. Early work at UNOTT developed a point matching method in which an initial point cloud created by a general multi-view method was enhanced by techniques designed for specific use on plants. The points produced were clustered into sets likely to represent points on the same leaf, and a plane fitted through them. This very crude leaf description was then enhanced by a region growing process which operated in the leaf plane but referred to the input images. The approach showed improved performance over previous methods but required a large number of images to produce good quality representations. The complexity and variation present in plants means that it is impossible to predict a priori which views will be needed and a set of 40-60 were found to be necessary to give good coverage. There is, however, still no guarantee that the images needed to reconstruct a given plant will be available, and it is common for many of those captured to be found unnecessary.

The existing UNOTT system is the starting point for improvement made during EPPN²⁰²⁰, where the goal was to construct an active reconstruction system in which the images used to

create the 3D model are automatically selected in response to initial estimates of plant structure.



Figure 8. The UNOTT active vision cell.

To achieve this, an active vision cell was constructed in which a camera was mounted in the gripper of a 6 degree of freedom robot arm capable of controlling camera position and orientation over a large working area and the plant placed on a nearby, computer-controlled turntable (Figure 8). A full calibration process recovers the relationships between camera, turntable and robot control parameters. A small set of images is captured from fixed positions to initiate the process. The plant is segmented from the background in each and the camera moved to centre the plant in the view. This creates a second set of images which are used to create an initial, low resolution 3D model of the plant via space carving.

The voxels making up this representation are clustered into small compact groups (Figure 9b), to reduce computation, and the visibility of each is assessed. For the point matching method to be successful, each area of the plant must be viewed by 2-4 images acquired from viewpoints satisfying constraints on their relative position and orientation. If any voxel clusters cannot be associated with images meeting these requirements a search is performed for viewpoints which would provide them. The robot then moves to those positions and captures the necessary images. Any redundant images are also discarded.

The UNOTT point matching algorithm is then applied to produce a set of 3D points lying on the plant surface. As a final filter, points lying outside the initial volumetric description are discarded; these are clearly the result of mismatched image features. A surface is then fitted over the remaining points to produce the final plant model. Figure 9 provides an overview of the process.

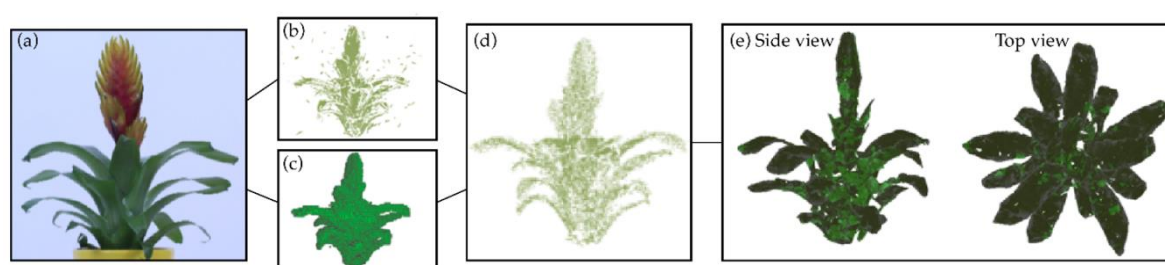


Figure 9. Overview of the reconstruction process on a Bromeliad (*Vriesea* sp.). (a) an image acquired by the active vision cell (b) point cloud representation containing outliers, (c) initial, volumetric representation, (d) merged model, (e) final 3D surface model

The active reconstruction system was evaluated by comparing the surface descriptions produced from colour images by the active cell with similar data structures obtained via X-ray computed technology in the University's Hounsfield Facility. A set of 6 plants with varying structures were scanned, segmented to separate plant material from the surrounding air and surfaces fitted. Note that while the Hounsfield's equipment produces high quality plant representations well-suited for use as ground truth, it is too expensive to be used for a task that can be achieved using more accessible technologies.

The active method was found to produce more accurate plant reconstructions than previous methods from fewer images, providing proof of concept of active 3D plant reconstruction.

2.4. Next Steps (DLO, UNOTT)

Because they are front-of-science and need to be developed close to application, the activities described above have been carried out separately, but with frequent exchange, in Wageningen and Nottingham. In the final year of EPPN²⁰²⁰ the data and techniques produced will be brought together and made available to the phenotyping community. Specifically:

- i. Data sets created by the two groups will be exchanged and the potential for further improving model connectivity by combining data structures will be assessed. The camera control strategies developed at UNOTT will be revisited and the potential for improved active reconstruction strategies based upon the organ segmentations provided by DLO's reconstruction and plant segmentation methods will be examined.
- ii. UNOTT's current system will be used, as a research tool, to identify imaging strategies suitable for different plant species/types. Two sets of plants will be modelled using the active cell, chosen for the differences in their physical structure. Attention will focus not on the quality of the models produced, but on the sets of viewpoints used to create them. These viewpoints will be plotted on a notional sphere surrounding the plant holder and smooth paths through those points sought. Identification of a path across the sphere which passes through the set of viewpoints used for a particular plant type will suggest a suitable viewing strategy. The models produced by a camera adopting that strategy will be assessed against ground truth and those produced by the active cell. Effective strategies could be used either to better initialize the active component for a given plant type, reducing the computation required, or be used in isolation in situations where active refinement of the strategy is not required/possible.

The actions will result in software elements first shared with the whole consortium for further test, and then made available to the whole phenotyping community via publication in an academic journal, whereas software elements will be available in a public repository.

3. Dynamic Phenotyping of Selected Plant Organs

3.1. Maize Silk Phenotyping (INRAE)

While 3D information is required to actively control camera placement, this need not take the form of a complete model of the plant. The point matching approach to 3D reconstruction can be applied to any image feature: in full 3D reconstruction methods it is applied to all those available, with the goal of producing as much 3D information as possible. An alternative approach is to recover and use the 3D locations of only those pixels that mark features of particular biological interest. This shifts the emphasis of the work away from more generic issues in 3D object modelling and towards the use of biological knowledge, and is the approach taken in the maize silk phenotyping system developed at INRAE.

Silk growth is a major trait for drought tolerance in maize, but one which is difficult to phenotype. First, maize plants have a distinctive, near planar architecture; images must be captured from viewing directions approximately normal to this plane to avoid high levels of occlusion which will prevent the silks from being visible at all. Second, silks are very small. Maize silks comprise hundreds of pollen-collecting filaments that must be imaged at close range if meaningful growth

data is to be obtained. The usual need for high throughput operation, however, makes close-range inspection of the entire plant impractical, even when a suitable view is available. To address these problems, INRAE has developed an active method capable of monitoring the silk growth of hundreds of maize plants every day.

First, the whole plant is imaged from multiple views: one top and 12 side views are acquired (Figure 10). The top view is segmented to separate plant material from background (Figure 10a) and a straight line fitted through the plant region. This line gives a good indication of the plant's major axis and allows some of the available side views to be rejected immediately, as not being close enough to perpendicular to that axis. Secondary linear regressions are performed on pixels lying far from the plant axis (shown blue in Figure 10b) and used to identify leaves which restrict the view of the remaining images (shown yellow in Figure 10c). Together, these processes allow usable images to be selected. In the example shown in Figure 10 this results in selection of the one acquired from the "90 degree" view. While not suitable for silk phenotyping, the selected image(s) are of sufficient quality as to allow the detection of maize ears.

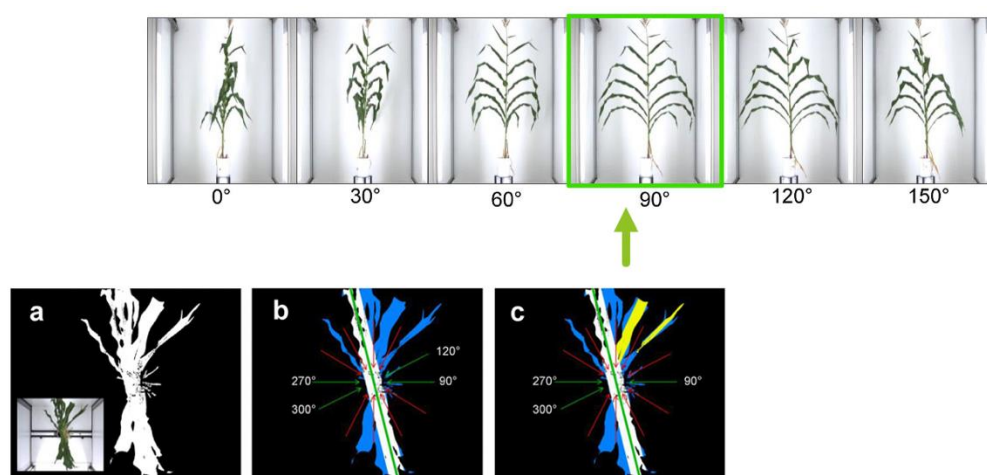


Figure 10. Side view selection in the INRAE maize silk phenotyping system. See text for details.

The selected side-view image is segmented, and the plant region skeletonised. The stem is extracted from the skeleton and its width computed at regular intervals. Knowledge of the structure of maize plants and the shape of maize ears is embedded in an algorithm which then detects ears in the selected image. The 3D (x,y,z) position of the detected ear is computed from the top and selected side image(s), providing the 3D information needed to support active image acquisition.

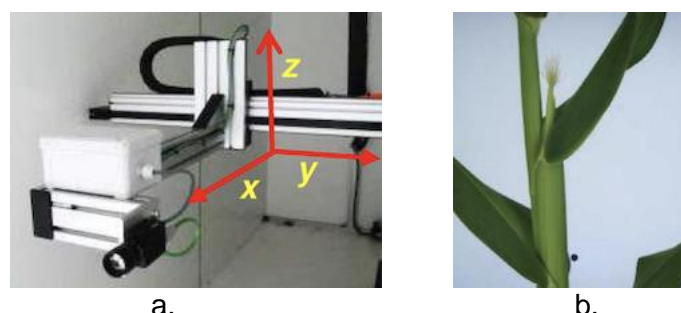


Figure 11. a. The INRAE active camera, b. sample image use in maize silk phenotyping.

At this point the plant is moved into a second imaging chamber and rotated so that the plane containing leaves (as previously identified from the top view) is perpendicular to the axis of a computer-controlled camera (Figure 11a). The 3D location of the detected ear then allows the camera to be positioned 30cm from the ear, providing an image (Figure 11b) from which data on silks can be extracted.

3.2. Next Steps (INRAE, UCL)

The image acquisition hardware and processing pipeline developed for maize silk analysis at INRAE have clear potential to be adapted to other phenotyping task. In the final year of EPPN²⁰²⁰ two such adaptations will be carried out:

- i. INRAE will create a grape phenotyping system based on their original hardware designs, image analysis and camera control code. Whole plant images will be acquired, and biological knowledge of the likely position and appearance of grape clusters embedded in software capable of their detection. Once detected, the 3D position of each cluster will be computed and a camera positioned to acquire a higher resolution image centred on the cluster which will be used to gather trait data.
- ii. UCL will complete development of a similar system, re-using components from INRAE where possible, for growth analysis of aeroponically-grown roots (Figure 12). Initial views will allow detection and 3D location of root tips, and an x,y,z-controllable camera will be positioned, at successive times over a period of growth, to assess root development.



Figure 12: Sample image from the UCL aerobics facility.

Code and designs for the INRAE maize silk system are already publicly available. The addition of similar materials from the grape and root tip work will initiate a library of scripts and tools for dynamic phenotyping of plant organs.

4. Conclusion

Automation is the key to high-throughput plant phenotyping, and current large facilities are built upon concepts developed in automated manufacturing: plants are often maintained on conveyors, travelling from one workstation to another as they each undergo the same series of fixed processes. Dynamic phenotyping takes a step away from automation and towards robotics, the key distinction being that while automated systems repeat the same operation multiple times, robotic systems sense their environment and take action accordingly. Rather than rely on fixed camera positions and image analysis pipelines, dynamic phenotyping methods seek to adapt the sensing strategy to the plant at hand and introduce a feedback loop between image acquisition. Dynamic phenotyping has the potential to increase the consistency of plant measurements, allow more detailed measurements to be made and support multi-scale phenotyping, but requires some 3D information to be available.

Work in JRA 1.2 of EPPN²⁰²⁰ to date has comprised a number of case studies which are now being brought together. Work on 3D reconstruction and maize silk phenotyping has demonstrated the feasibility of constructing practical plant analysis systems based on dynamic principles. In the final year of the project the techniques developed will be adapted to new phenotyping tasks, combined and consolidated. We shall seek new possible case studies that can be developed in other nodes of the project, on other species or organs. This will serve as a further test of the methods presented here. Some of the software tools are already available for the whole phenotyping community via academic journals and public repositories. This will be the case for all other methods presented here, at the end of the test procedures.