

Active Vision and Surface Reconstruction for 3D Plant Shoot Modelling

Jonathon A. Gibbs¹, Michael P. Pound¹, Andrew P. French, Darren M. Wells¹, Erik H. Murchie¹, and Tony P. Pridmore¹

Abstract—Plant phenotyping is the quantitative description of a plant’s physiological, biochemical, and anatomical status which can be used in trait selection and helps to provide mechanisms to link underlying genetics with yield. Here, an active vision-based pipeline is presented which aims to contribute to reducing the bottleneck associated with phenotyping of architectural traits. The pipeline provides a fully automated response to photometric data acquisition and the recovery of three-dimensional (3D) models of plants without the dependency of botanical expertise, whilst ensuring a non-intrusive and non-destructive approach. Access to complete and accurate 3D models of plants supports computation of a wide variety of structural measurements. An Active Vision Cell (AVC) consisting of a camera-mounted robot arm plus combined software interface and a novel surface reconstruction algorithm is proposed. This pipeline provides a robust, flexible, and accurate method for automating the 3D reconstruction of plants. The reconstruction algorithm can reduce noise and provides a promising and extendable framework for high throughput phenotyping, improving current state-of-the-art methods. Furthermore, the pipeline can be applied to any plant species or form due to the application of an active vision framework combined with the automatic selection of key parameters for surface reconstruction.

Index Terms—3D reconstruction, active vision, calibration, plant phenotyping

1 INTRODUCTION

UNDERSTANDING the physiological and molecular mechanisms underlying the growth of plants and crops is increasingly important in order to reach food security goals whilst achieving sustainability within agricultural systems. Therefore, methods are required to identify traits that translate into increased crop yield [1].

Phenomics, or phenotyping, is the measurement of traits which reflect plant growth, performance, composition and development. Measurements are often captured over a period of time, and are affected dynamic interactions with the genetic background (genotype) and the environment [2]. A comprehensive understanding of the possible range of plant phenotypes would aid breeding and genetic modification, facilitating the improvement e.g., of nutrient use and photosynthetic efficiency, thereby increasing crop yield and stability across diverse environments [3], [4]. Dimensions such as height and width, leaf area index (LAI), leaf area density (LAD) plus leaf angle and curvature are all important parameters that directly relate to the growth and light harvesting ability of plants. However, retrieving these measurements currently

constitutes a major bottleneck in assessing performance, or for the construction of quantitative models of plant development and pre-breeding applications. This, in part, is due to the drawbacks associated with manual measures of phenotypic parameters, which are time consuming and subject to inaccuracy [5]. Higher throughput methods may also be costly, inaccessible, or require specific expertise or appropriate analysis tools.

In recent years there has been a rapid increase in techniques aimed at extracting plant traits from two-dimensional (2D) images, to generate phenotypic information. However, many plant measurements such as area and leaf morphological properties can only be obtained from three-dimensional (3D) representations of plants. Consequently, the recovery of accurate 3D models of plants from 2D images is urgently required. High quality 3D models can provide a wide range of morphological and gross developmental data and can also be used to support simulations of plant function (e.g., [6], [7], [8]). While image-based modelling has made significant progress over the past decade, creating accurate representations of plants remains a challenging problem.

- J. A. Gibbs, M. P. Pound, A. P. French, and T. P. Pridmore are with the School of Computer Science, University of Nottingham, Nottingham NG7 2RD, United Kingdom. E-mail: {jonathon.gibbs1, michael.pound, andrew.french, tony.pridmore}@nottingham.ac.uk.
- D. M. Wells and E. H. Murchie are with the School of Biosciences, University of Nottingham, Nottingham NG7 2RD, United Kingdom. E-mail: {darren.wells, erik.murchie}@nottingham.ac.uk.

Manuscript received 21 Oct. 2018; accepted 20 Jan. 2019. Date of publication 25 Apr. 2019; date of current version 8 Dec. 2020.

(Corresponding author: Jonathon A. Gibbs.)

Recommended for acceptance by Y. SUN.

Digital Object Identifier no. 10.1109/TCBB.2019.2896908

1.1 Recovery of 3D Plant Models

The recovery of 3D descriptions of viewed objects from multiple images is a longstanding problem in computer vision. In recent years improvements to these descriptions, with respect to both quality and performance, have been made. However, most existing methods of representing objects in 3D have only been applied to relatively simple, predominantly convex objects; human heads and man-made artefacts such as buildings and vases.




Fig. 1. Overview of the reconstruction process on a Bromeliad (*Vriesea sp.*). (a) Sample image acquired by the active vision cell (AVC), (b) point cloud representation containing outliers, (c) volumetric proxy representation (PPR), (d) merged model, and (e) final 3D mesh model following surface reconstruction.

Crowded scenes, in which multiple, closely-packed objects are present, constitute a more difficult challenge and are troublesome to accurately represent. Crowded scenes generate high levels of occlusion - where part of the object is not visible from the current view - and parallax - the effect of the object appearing to differ when viewed at different angles, making accurate reconstruction harder than for simple convex objects. Plants are particularly challenging due to self-occlusion, the presence of many small, shiny surfaces that appear very similar, lack of texture for feature matching and difficulties when selecting camera placements. Moreover, plants are sessile organisms that adapt and acclimate to their fluctuating environment, from short-term changes such as the reorganisation of foliage to long-term growth patterns. As a result, plants have a complex structure which is able to change over time, making them difficult to model, particularly by standardized, fixed camera phenotyping platforms. Pipelines are required that can adapt to the expected wide variations in plant shape and/or size [9], [10], [11].

Many approaches to 3D modelling exist [12]. One particularly successful method is known as Multi View Stereo (MVS), which produces a series of 3D points, known as a point cloud, from an unordered set of 2D images taken around an object. However, image acquisition and selection are currently an insufficiently considered resource in MVS; this is particularly important when the target objects are complex and feature matching becomes more challenging. MVS uses feature matching between multiple overlapping views of an object or scene to determine the position of the object in 3D space. Seitz et al. [13] demonstrated that the quality and speed of 3D modelling depends significantly on the quantity and selection of input images, and that each image does not contribute evenly to the overall quality of the model. Without sufficient images, a faithful description of the plant cannot be produced, though also, without a prior representation of the object to compare to, it is difficult to evaluate whether sufficient images have been acquired. For crowded scenes, even when a large number of images are available, there is a high probability that some 3D data will be missing. An increase in the number of input images will also increase computational requirements in terms of both memory and time. In some instances, unnecessary input images can actually decrease the quality of the resulting model by introducing false points.

The point clouds generated via MVS are often integral to 3D modelling pipelines, where they are used as a basis to create a surface mesh representation. This is necessary for

automatic plant measurements or further applied modelling (e.g., [6], [7], [8], [14], [15], [16]). However, missing data at the point cloud stage can lead to mesh defects such as holes or overlapping datapoints [17] later on. Other imperfections may arise due to inherent noise in the point cloud itself, which alters the position of points, or incomplete sampling of the surface which may lead to holes, gaps and missing data. Moreover, the point cloud may include outliers; points lying away from the object surface that are not part of the object being modelled. These imperfections are likely to amplify inaccuracies arising during the surface reconstruction process; the process of fitting a surface to the point cloud to produce a 3D model. The complexity of the target object also leads to difficulties, affecting the ability to accurately detect and preserve the boundaries of complex shapes, and further increasing the processing time required for reconstruction.

In this work, a complete pipeline is proposed for the automatic recovery of 3D plant models. The pipeline can be split into two parts; 1. The automatic capture of 2D images of a given plant using an active vision cell (AVC) via manipulation of a camera mounted on a robot arm [18]. To do this we utilise an intermediate plant representation, which is interactively evaluated to ensure it has been adequately scanned, and automatically capture additional images where necessary. 2. The surface reconstruction of the model plant, achieved by utilising the 3D data obtained from Part 1 and merging the point cloud and volumetric data representations together to improve the accuracy of the 3D model. We then perform clustering on the merged data using a novel binning algorithm; and then surface reconstruction using triangulation. Level sets and a merging algorithm are applied to the triangulated clusters in order to refine the final surface representation. An overview of steps in the reconstruction process for real life plant is given in Fig. 1.

2 ACTIVE VISION

In controlled environment plant phenotyping, image acquisition is currently reliant on either manual capture [14] or static camera placements [19], [20], [21] that are unable to adapt to specific plant species or varieties. Consequently, there is an increase in labour, reduction in accuracy and often very expensive systems are designed only for a single plant form [22], [23], [24]. Here, an active vision approach is developed to reduce the manual requirements of image capture, ensuring that sufficient data is obtained, irrespective of plant species or form.



Fig. 2. The Active Vision Cell (AVC) hardware components: the LT360EX high precision turntable (left) and the Canon 650D mounted on the UR5 robot arm.

The proposed AVC [18] is composed of three primary hardware components; 1. A high precision turntable (LT360EX – Linear X Systems, Portland, USA) which is used to rotate a plant 360 degrees, providing accuracy of 0.1 degrees. 2. A robot arm, The Universal Robot 5 (UR5 – Universal Robots, Odense, Denmark), providing 6 degrees of freedom, which, in conjunction with the turntable, ensures the whole object can be viewed. 3. A digital-SLR camera (Canon 650D – Canon, Tokyo, Japan), for acquisition of RGB images (Fig. 2). The 650D was chosen due to affordability and high-resolution images of up to 18 megapixels, as well as the ability to mount it on a robot arm easily.

2.1 Calibration

To enable the manipulation of a viewpoint for active vision, the system must be accurately calibrated; that is, the location of the camera with respect to the world must be determined. This can be challenging, with the number of calibrations required dependent on the number of components in the system.

Within this cell, calibration of the camera and robot with respect to the world, robot base and end effector (the point in which a device is connected to a robotic arm) is performed. Calibration of this form is achieved by solving the linear equation $AX = YB$ (Fig. 3) in which $AXYB$ are transformation matrices consisting of rotation and translation, mapping a point in one space to one in another. B is the robot calibration from the base to the end effector, a transformation often known as forward kinematics. B is easily obtainable using Denavit-Hartenberg [25] parameters, a method for describing the structure of a serial link manipulator. A is the camera calibration and involves estimating the intrinsic and extrinsic parameters of the camera. The intrinsic parameters depend on the physical properties of the camera and lens. The extrinsic parameters represent the position and orientation of the camera within the world. Camera calibration is performed by capturing multiple images of a, planar target (e.g., a checkerboard) of known structure and dimensions.

Once these two calibrations are performed it is possible to calculate the position of the robot with respect to the world, Y , and the camera with respect to the end effector, X . The two unknowns are calculated from the linear equation using unit quaternions and a closed-form separable solution as proposed by Dornaika and Horaud [26]. Finally, the turntable is calibrated by performing multiple cell calibrations; in each, the centre of the world $(0, 0, 0)$ is obtained. Opposite rotations are connected, and line intersection is performed to calculate the centre of rotation.

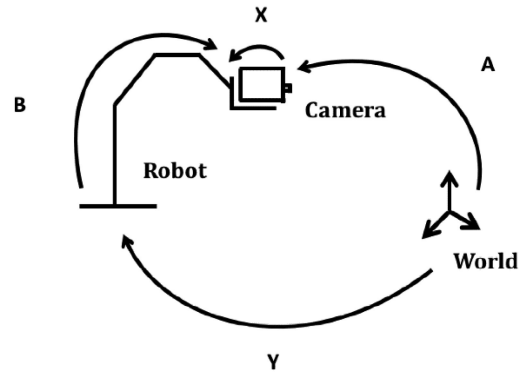


Fig. 3. The relationship between transformations of matrix equation $AX = YB$ for the calibration of the robot and camera.

Once the system is fully calibrated, the checkerboard calibration target can be removed from the environment. It remains possible to accurately obtain the position of the camera with respect to the world, through the full range of robot movement and turntable rotation, enabling 3D reconstruction.

For highly complex scenes that lack texture or distinct features, such as plants, feature matching approaches to calibration are unreliable. Thus, many phenotyping pipelines perform camera calibration online during reconstruction, requiring that a calibration target is at least partially visible in all images [14]. Consequently, the range of possible views are restricted, large objects may not fit, or may occlude much of the calibration target. The calibration method described above eliminates this constraint and can calibrate the camera for any position without the need for a calibration target, once all parameters are known.

2.2 Performing Active Vision

Given a calibrated cell, it is possible to capture images of any plant that is placed on the turntable. The process of selecting views by evaluating the environment to maximise some function is typically called next best view (NBV). Traditionally, NBV evaluates every possible viewpoint to determine which view to select. Within this pipeline, the proposed NBV algorithm differs from existing NBV algorithms in that it uses neither a depth sensor, nor a predefined set of images. Moreover, the proposed algorithm reduces the search space of viewpoints by using clusters of voxels, i.e., grouping them together, and an incremental view sphere, in which one 'optimal' position is evaluated and then expanded should it be inappropriate, as opposed to traditional methods which evaluate all views and all voxels. Redundant images are discarded to further improve efficiency. The proposed algorithm (Algorithm 1) can be broken down into three phases:

Phase 1: Select views that support fast convergence towards an initial volumetric representation, which broadly represents the object; this is evaluated in Phase 2 to determine quality and amount of coverage (i.e., the number of voxels that have been seen). An octree is used as an initial representation of the object due to the efficiency and ease of manipulation, particularly when performing ray tracing. The octree is initialised as a single voxel and, as views are captured, knowledge of the object is increased. The initial image set is captured from three height positions for 10 turntable rotations, resulting in 30 images.




Fig. 4. Incremental examination of the view sphere determines camera positions. The red dot in the center illustrates the optimal viewpoint defined by the normal of the cluster, the green is the first increment, and the yellow the second. This continues across the entire view sphere.

Phase 2: The voxel representation is clustered, and rays are projected from each of the 30 cameras. From this a score for each cluster is determined based on the angle at which each image (the position which it was taken) sees the cluster, and the percentage that has been seen. Clusters that do not satisfy the evaluation function (i.e., those that are not seen enough) require more views. The new views are determined using an incremental view sphere in which the intersection of the normal from the cluster to the view sphere is used as the optimal position; this is evaluated by ray tracing to detect occlusions (Fig. 4). If the view is occluded, incremental expansion over the view sphere is performed until an acceptable view is found.

Phase 3: Redundant images are removed from the image set; these are any images that can be removed without decreasing the evaluation score (Phase 2) (Eq. 4.9), based on the number of times the cluster has been seen and the angles between the cameras that have seen it.

A more detailed specification of the process is given in Algorithm 1 and discussed in the remainder of this chapter. Algorithm 1 requires that the vision cell is calibrated and outputs a series of 2D images and corresponding matrices. A volumetric model is initialised, Algorithm 1 Line 1, and the robot is moved to the starting position in line with the turntable, Algorithm 1 Line 2. For each level (the position of camera along the vertical axis) a colour filter is applied to the cameras live stream and the robot arm is moved along the x and y axis until the object is encapsulated in the cameras field of view (FOV), Algorithm 1 Line 4, 10 images are captured at the current level rotating the turntable after each image and estimating camera matrices, Algorithm 1 Line 5-10. The volumetric model is then updated, Algorithm 1 Line 11, and the process is repeated for the next two levels. The next step involves active vision to obtain images of areas of the object that have not been viewed sufficiently. The volumetric model is clustered, and a view sphere is created, Algorithm 1 Line 13-15. The clusters are evaluated and if another view is necessary, the view to best see the cluster is determined based on the utility function, Algorithm 1 Line 17-21. If the model requires further evaluation, repeat the active vision process, Algorithm 1 Line 23. Once the

model is sufficiently scanned the volumetric resolution is increased and redundant images are removed, Algorithm 1 Line 24-25.

Algorithm 1. The NBV algorithm

Input: A calibrated active vision cell
Output: A series of Images, volumetric model, camera calibrations

- 1: $V \leftarrow \text{initialiseVoxel}()$
- 2: $Pos \leftarrow \text{determineStartPosition}()$
- 3: for $l \leftarrow 1$ to 3 do
- 4: $Pos \leftarrow \text{position}(l, V, \text{ColorFilter})$
- 5: $t \leftarrow \text{rotate}((l-1).12)$
- 6: for $n \leftarrow 1$ to 10 do
- 7: $I_n^l \leftarrow \text{Capture}(Pos)$
- 8: $M_n^l \leftarrow \text{camCalibration}(Pos)$
- 9: $t \leftarrow \text{rotate}(n)$
- 10: endfor
- 11: $V \leftarrow \text{update}(I^l, M^l)$
- 12: endfor
- 13: $p \leftarrow \text{points}(V)$
- 14: $C \leftarrow \text{cluster}(P)$
- 15: $S \leftarrow \text{viewSphere}(V)$
- 16: for $n \leftarrow 1$ to $\text{size}(C)$ do
- 17: $C_n^s \leftarrow \text{utilityFunction}(C, I, M)$
- 18: if $C_n^s < 1$ do
- 19: $Pos \leftarrow \text{incrementalViewSphere}(S, C_n)$
- 20: $V \leftarrow \text{update}(\text{capture}(Pos), \text{camCalibration}(Pos))$
- 21: endif
- 22: endfor
- 23: if !evaluate(V, C) goto line 16
- 24: $I \leftarrow \text{removeRedundant}()$
- 25: $V \leftarrow \text{increaseResolution}(V)$

The result of active vision is a set of 2D images that sufficiently capture the whole object, and their corresponding transformation matrices. We discuss the results of our AVC next, comparing it to more traditional image capture methods.

2.3 Evaluation

Evaluation of the effect of the AVC on surface reconstructions obtained by the method described in Section 3 was performed on six target plant species: Bromeliad (*Vriesea sp.*), *Aloe vera*, Cordyline (*Cordyline sp.*), *Brassica napus*, chilli (*Capsicum sp.*) and pumpkin (*Cucurbita pepo*). These plants were chosen based on their contrasting morphology (particularly size and leaf structure), demonstrating the generalizability of the approach.

For evaluation, X-ray μ CT images of target plants were obtained using a GE v|tome|x M scanner housed in the University of Nottingham's Hounsfield Facility, Sutton Bonington Campus. The v|tome|x M provides volumetric images with a voxel resolution of 5 - 150 μm and, more importantly, is not subject to the occlusion problems faced by visible light imaging; this therefore provides a gold-standard ground-truth dataset. Though some X-ray segmentation tasks are highly challenging, plant material and air are easily separable in the density data provided by μ CT and, following noise reduction, plant material can be identified by applying a user-defined threshold. From this, a complete 3D model of the target object is formed. The surface of each plant can be




Fig. 5. X-ray μ CT scans used as ground truth models. (a) Bromeliad (*Vriesea* sp.), (b) *Aloe vera*, (c) *Cordyline* (*Cordyline* sp.), (d) *Brassica napus*, (e) *Chilli* (*Capsicum* sp.), and (f) *Pumpkin* (*Cucurbita pepo*).

represented in standard triangular mesh format. This ground-truth data structure can be compared against point clouds and surface representations obtained from the proposed pipeline can be evaluated (Fig. 5).

Models used for the ground truth in this work were manually checked to ensure that the X-ray μ CT scanner was able to accurately reproduce the plant. It is worth noting that while the X-ray μ CT scanner produces accurate, highly detailed models suitable for evaluation, it is ill suited to most plant phenotyping tasks due to size restrictions, time requirements and a costly setup. Moreover, it is possible for the scanner to miss very thin areas of the plant such as stem material, resulting in an incomplete 3D model for some plant types.

2.4 Active Vision Results

The algorithm proposed here is evaluated against three other imaging methods:

One static; a single static camera is placed at the side of the plant at a fixed distance, with the plant placed on a turntable;

Two static uses two fixed cameras, one directly in line with the plant, the other placed slightly higher looking down, as this is a common approach in phenotyping systems, again with the plant on a turntable;

Arbitrary camera placement, which uses a series of random views taken from distinct positions, as is typically produced by manual imaging.

40 images were taken for each of the three imaging methods whilst the removal in redundant images in the AVC led to a 10–48 percent reduction in image set (i.e., between 21–36 images were required for the six plant species that were evaluated). Images are removed if their exclusion from the dataset does not decrease the utility function such that more views are necessary. Evaluation is based on the distance from points to the surface of the ground truth (i.e., the X-ray μ CT scan) and the number of points obtained per image (Table 1).

In all cases AVC produced more points, with a lower mean distance to the ground truth, using a reduced image set. While static camera placement can produce a good model, success depends on the structure of the target plant, and efficiency is limited. By employing active vision, the amount of data obtained can be significantly increased despite requiring fewer images, thus reducing computational cost. While

TABLE 1
Evaluation of the Active Vision Cell

	Brom	Al	Cord	Brass	Chilli	Pump
Mean distance of points from the ground truth						
AVC	0.196	1.333	0.738	0.035	0.102	0.359
One St	0.357	1.452	0.757	0.201	0.238	1.122
Two St	0.344	1.691	0.864	0.087	0.184	1.210
Arb	0.269	1.896	1.028	0.168	0.254	0.698
Number of points per image						
AVC	7638	5705	3764	16634	10192	34953
One St	3527	3997	3071	2430	2832	17881
Two St	3885	4015	2355	3669	6186	12926
Arb	7638	4576	2004	4461	4976	21311

The mean distance of points is the average distance of all points in the point cloud to the ground truth X-Ray μ CT scan. A lower mean indicates higher accuracy. The number of points per image is the average number of points that each image generates where a higher number indicates higher quality images.

arbitrary camera placement can perform well, there are no guarantees on accuracy or repeatability.

The approach described here requires minimal user input. The AVC can adapt to objects of different shapes and sizes with varying levels of occlusion and complexity, with the only limitation to size being the reach of the robot arm (which could be overcome using a larger robot e.g., the Universal Robot 10; a larger version of the robot used here). The common difficulty for extremely dense scenes in which components, i.e., stems, cannot be separated, remains problematic. Static approaches often have cameras fixed in the environment making it extremely difficult and costly to adjust to different object sizes, which is crucial in plant phenotyping when trying to obtain growth information, or when creating a system capable of measuring multiple different species. The AVC is more accurate and requires fewer images than previous static imaging approaches, offering more flexibility than existing large-scale phenotyping systems by adapting to the natural variation of individual plants. As seen from Table 1 the AVC produces a set of points with a lower mean (root mean squared error) in all instances. The method proposed here is automatic with user input limited to changing the plant and clicking a single button to begin the process.

3 DATA MERGING AND CLUSTERING

The AVC produces several data forms, namely: a 3D point cloud generated by Patched Based Multi-View Stessreo (PMVS; Fig. 1b); a volumetric model (the Plant Proxy Representation- PPR; Fig. 1c); a series of camera calibration matrices; and a set of 2D images. It is important to note that the 2D images are not required for model merging (Section 3.1) or clustering (Section 3.2) but are used during level set processing in the surface reconstruction stage (Section 4). There are a number of challenges with using each of these data forms individually. For example, the point cloud often contains outliers; the PPR contains voxels that cannot be verified, that is, it is unknown as to whether they exist in the model; and the 2D data is not suitable alone due to occlusion and parallax (discussed previously), consequently limiting the ability to measure plant traits.



Fig. 6. Model merging. (a) The original point cloud data consisting of outliers, (b) the PPR representation containing unverified voxels, and (c) the merged model with the removal of outliers.

3.1 Model Merging

The accuracy and usefulness of data can be improved through merging the volumetric and point cloud representations. The model merging phase combines the point cloud and the PPR to produce a reduced, yet more faithful, point set which represents the object of interest (Fig. 1d). Fig. 6 illustrates the two starting datasets, point cloud and PPR, and the resulting merged model, from which the outliers have successfully been removed, yet the integrity of the object remains. The merged model produces a new point set which is used for the basis of cluster analysis and surface reconstruction (Fig. 1e).

Merging is performed by evaluating each voxel in the PPR, if it contains more than three points (as three are enough to form a single triangle) the voxel is considered to be part of the object. The points in each voxel that are part of the object are averaged to produce a single point. There are numerous advantages to model merging, as opposed to using either the point cloud or PPR independently. Namely; the ability to acquire connectivity information from the PPR (which is particularly useful in the surface reconstruction) where such information is not available in the point cloud, along with normals, and more accurate colour estimates which can be obtained from the point cloud. However, the biggest advantage is the ability to significantly reduce the size of the point set. For example, a complex 3D point cloud consisting of 300,000+ points can be reduced to a set of 3,000 points, whilst still retaining a faithful representation of the object.

3.2 Clustering

Following model merging, we perform clustering of the points to allow for efficient surface reconstruction in Section 4. A hybrid clustering algorithm has been devised, manipulating principal component analysis (PCA) and 1D data representations. The algorithm involves three key steps:

Step 1. Partition clusters based on the normal, using PCA to determine deviation from the plane – satisfying the normal constraint. The deviation of points from the plane (with respect to the Euclidean distance) determines whether the point set is flat, or close to flat, by evaluating λ^{Mid} . If the deviation from the plane is too great, the cluster is split. Traditional PCA algorithms split along the centre, or along some gap when evaluated using histograms. Here, a novel approach is used in which each 3D point is binned such that all points along λ^{Max} are projected to a single dimension. To determine a potential position to split the point set, an angle is calculated for each element in the array. The angle is the deviation between the current normal and the average sum of normals of values to the left and right of the current index. When the sum of normals to the left equal or are most similar to those on the right, the position to split is found (Fig. 7).




Fig. 7. Overview of the clustering algorithm. As you move along the array, the sum of differences between the normals (N) either side of position X is calculated. When the minimum difference is found such that A-B is close to 0, the position to split the points is assigned.

Step 2. When PCA (Step 1) is unable to improve the current point set, a modified automatic divisive hierarchical clustering (DIVFRP) [27] algorithm is applied to partition clusters if there are two distinct groups – satisfying the separation constraint. A dissimilarity measure based on the furthest reference points is used to overcome difficulties of clustering areas containing small distinct groups, which often occurs in complex scenes. DIVFRP aims to maximise the dissimilarity function in order to split points.

Step 3. Clusters are reclassified using the nearest centroid and proximal points – overcoming the outlier or incorrect patch constraint. In some instances, single points may be incorrectly placed. These are typically located at the most distal points of a clusters (i.e., the outer most points of a cluster). This can be due to noise or incorrectly estimated normals; arising from the point proximity during estimation. The reclassification step aims to move the distal points to another neighbouring cluster (usually one in close proximity), should the point be closer to the neighbouring cluster than the cluster it is currently assigned.

Algorithm 2. Data Merging and Clustering

Input: A list P all points, V a volumetric model

Output: A list, C , of clustered points

```

1: for  $i \leftarrow 1$  to size( $P$ ) do
2:   for  $n \leftarrow 1$  to size( $V$ ) do
3:     if  $BB(P_i, V_n)$ 
4:        $V_n^p.add(P_i)$ 
5:     end
6:   end
7:  $V \leftarrow V$  where  $size(V^p) > 0$ 
8:  $C_0^p \leftarrow patch(V)$ 
9: while(true)
10:  for  $j \leftarrow 1$  to size( $C$ )
11:    if !criteria( $c_j$ )
12:       $C.add(split(c_j))$ 
13:       $C.remove(c_j)$ 
14:  if criteria = true  $\forall c_j \in C$ 
15:    return false
16: end
```

Algorithm 2 summarises the method. Algorithm 2 Line 3 determines whether some voxel is contained within v_n and if it is, adds point p_i to voxel v_n . $Size(\cdot)$ returns the length of the data structure. Algorithm 2 Line 8 determines the new point set (obtained from merging volumetric and point

TABLE 2
Evaluation of the Clustering Algorithm

Point Set	A	B	C	D	E	F
Correctly clustered points (%)						
Proposed	100	100	100	100	100	100
K-means	100	82	76	82	74	49
IPDDP	100	45	100	100	71	91
Spectral	100	100	100	100	100	100
Time taken to cluster (s)						
Proposed	0.183	0.374	2.512	2.988	7.008	2.762
K-means	0.005	0.001	0.005	0.008	0.007	0.006
IPDDP	0.113	0.121	2.26	1.928	3.106	1.564
Spectral	0.392	0.513	21.9	29.91	45.13	18.63

Correctly clustered points refers to the percentage of points that are correctly clustered for each of the methods. The time taken to cluster is the time taken to finish executing in seconds.

data). This modified point set, V , is used for the remainder of the work. Algorithm 2 Line 11 returns true if the cluster satisfies the constraints for a cluster to be correct and false otherwise. If the criteria function returns false the cluster is split, Algorithm 2 Line 12, which returns two new clusters and appends them to the end of the cluster list. Once all clusters meet the criteria, Algorithm 2 Line 14, clustering is complete; the criteria is defined as a series of constraints these are; points must be closer to the centroid of their own cluster and not any other (the outlier and incorrect patch constraint), points must have minimal deviation from the orthogonal plane (the normal constraint) to prevent the creation of long triangles, the loss of curvature and the incorrect triangulation of discrete components i.e., two leaves (the separation constraint).

3.3 Results

Evaluating clustering algorithms on complex 3D point sets (i.e., those with close and overlapping subsets) from scanned data is particularly challenging, the true number of clusters is often unknown, and identification of a ‘correctly’ clustered 3D point set is subject to opinion. Therefore, evaluation of the clustering algorithm is performed on data that can be visually inspected, which is obtained by segmenting a plant along the z -axis, and a correct number of clusters determined. Consequently, the number of actual clusters can be accurately identified through visual inspection. The clustering algorithm is evaluated against three existing algorithms on these segments: 1. K-Means clustering [28] iteratively groups points based on the nearest, 2. Iterative Principal Direction Divisive Partitioning (IPDDP) [29] clusters based on an embedding in a higher dimensional Euclidean space and 3. Spectral clustering [30] which clusters in fewer dimensions. These were chosen as they represent commonly used (K-means) and state of the art (Spectral) clustering methods. Each algorithm is evaluated with respect to the number of correctly clustered points and the time taken to cluster the point set (Table 2). The *Point Set* (Table 2) represents the point set which is used for evaluation, in which 6 are chosen varying in complexity as shown in Fig. 8.

For all segments evaluated, only the proposed method and spectral clustering were able to accurately detect the correct number of clusters and the correct number of points

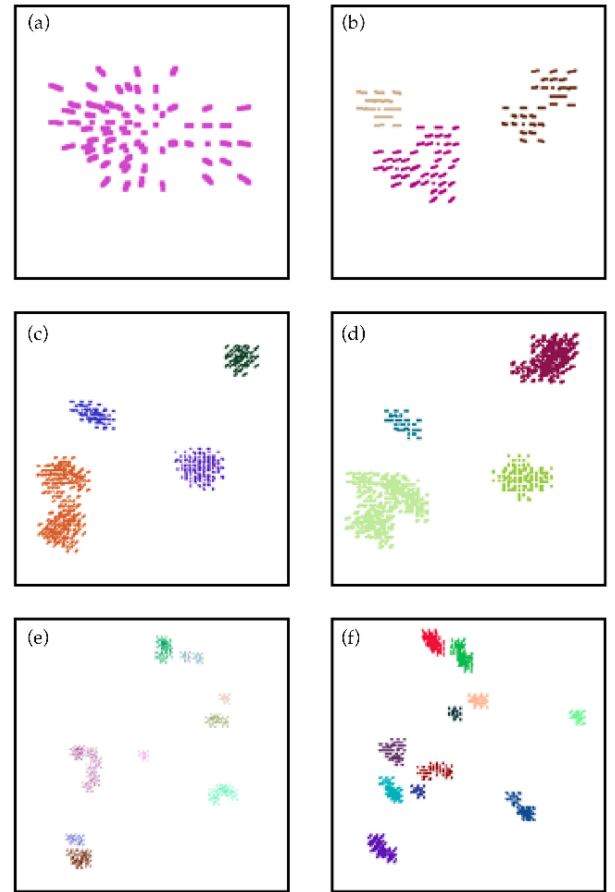


Fig. 8. Example point sets used for cluster analysis varying in complexity from a single cluster (a) up to 11 clusters (f) present.

within each cluster for all segments (Table 2 – Correctly clustered points), whereas K-means clustering and IPDDP often underestimated the number of clusters present, in one instance achieving only 45 percent accuracy. The proposed clustering algorithm can perform equally as well as the existing spectral clustering method, but with lower computational requirements (i.e., increase speed; Table 2 – Time taken to cluster), thus provides an improvement to current clustering methods in this instance.

4 SURFACE RECONSTRUCTION

A set of clustered points now represents the plant. The final stage of the modelling pipeline produces a surface representation of the object (Fig. 1e).

4.1 3D Modelling

Surface reconstruction is broken down into four phases, an overview is shown in Fig. 9 (see Algorithm 3).

Phase 1- Triangulation. Starting with the initial clustering (Fig. 9a), the points in each cluster are projected to an orthogonal regression plane through that cluster and are triangulated to produce a set of triangulated clusters, known from this point forward as cells (Fig. 9b).

Phase 2- Level Sets. Cell boundaries are expanded using level sets which can generate additional points plus maintain sharp features and edges (Fig. 9c). Directly applicable to this work, Pound et al. [14] used level sets during surface reconstruction to expand boundaries of clusters, as is done




Fig. 9. The surface reconstruction process. (a) Starting clusters, (b) triangulated clusters (cells), (c) level sets applied to cells, (d) smoothing of the cells, (e) cell merging, and (f) final smoothing step.

here. However, in Pound et al. [14], the resulting boundaries were not adjusted to faithfully represent curvature of the object as they were expanded on planar surfaces. In this work, this is overcome by re-projecting to 3D coordinates based on the surrounding point projection matrices maintaining the curvature of the original object.

Algorithm 3. Surface reconstruction algorithm

Input: A list of clusters, C , An image set, I
Output: A surface representation, S

1. for $i \leftarrow 1$ to $\text{size}(C)$ do
2. $c_i^T \leftarrow \text{delaunayTriangulation}(c_i)$
3. $c_i^{LS} \leftarrow \text{levelSets}(c_i)$
4. $c_i^S \leftarrow \text{smooth}(c_i^T, c_i^{LS})$
5. end
6. $S = C$
7. while $\text{size}(C) > 1$
8. $C \leftarrow S$
9. $S \leftarrow \text{newSurface}()$
10. $C \leftarrow \text{detectBoundary}(C)$
11. for $j \leftarrow 1$ to $\text{size}(C)$
12. if merged (c_j) continue
13. $S.add(\text{cellMerge}(c_j))$
14. end
15. end
16. $S \leftarrow \text{smooth}(S)$

Phase 3- Mesh Modifications. The surfaces are smoothed using a Laplacian smoothing algorithm [31], long triangles are removed, and holes are filled (Fig. 9d).

Phase 4- Cell Merging. A merging algorithm is applied to the cells based on connectivity information obtained from the volumetric model and the Euclidean distance. Cells are merged together to produce a fully connected surface representation (Figs. 9e and 9f).

Algorithm 3 summarises the 4 phases described above. A series of clusters and an image set obtained using the AVC are used for surface reconstruction. For each of the clusters, triangulation is performed, Algorithm 3 Line 2, to produce a series of cells, followed by level sets to expand the boundaries, Algorithm 3 Line 3. The resulting cells are smoothed; first, the internal points and then the boundary points, Algorithm 3 Line 4. The second stage creates a fully connected surface. The boundary for each cell is detected, Algorithm 3 Line 10, cells




Fig. 10. The final surface reconstruction. (a) Bromeliad (*Vriesea sp.*), (b) *Aloe vera*, (c) *Cordyline (Cordyline sp.)*, (d) *Brassica napus*, (e) *Chilli (Capsicum sp.)*, and (f) *Pumpkin (Cucurbita pepo)*.

are then evaluated to determine if they are connected, and if so they are merged, Algorithm 3 Line 13.

4.2 Results

Example 3D models showing the final surface reconstruction of six plant species are given in Fig. 10.

The proposed surface reconstruction algorithm was evaluated against the reconstruction algorithm of Pound et al. [14] using the X-ray μ CT scans as a ground truth. Images were captured using the AVC presented in Section 2 and the same image set was provided to each reconstruction method (N.B. redundant images were removed by the AVC-based method, as the algorithm explicitly reduces redundancy, but not for the Pound et al. algorithm, so a larger set of images was available for the latter method). Evaluation was performed based on the mean distance and the percentage of the plant area represented relative to the ground truth model (Table 3).

For all plant species, the proposed surface reconstruction algorithm shows a reduced mean distance relative to the ground truth model, compared to the mesh produced using the canopy reconstruction algorithm [14]. Furthermore, for all but one tested plant species, a greater percentage of the plant is represented (calculated as mesh area) using the proposed method (Table 3). In all cases, an improved reconstruction is produced in terms of accurate surface representation, in which a higher percentage of plant area and lower mean (Table 3) constitute a more accurate representation. Moreover, in each of the cases, the number of

TABLE 3
Evaluation of the Surface Reconstruction Algorithm

	Brom	Al	Cord	Brass	Chilli	Pump
Mean distance of mesh from the ground truth						
Gibbs	1.18	1.91	1.78	0.32	0.53	1.35
Pound	3.37	5.54	6.85	0.48	0.55	1.50
Percentage of plant area represented (%)						
Gibbs	69.1	75.6	84.2	77.5	93.1	101.9
Pound	33.6	20.2	78.1	80.2	81.8	61.6

Mean distance is the distance of the surface reconstruction obtained from Gibbs and Pound compared to the ground truth model using euclidean distance, as root mean square error. *The percentage of plant area represented* is the total surface area of the reconstruction compared to that of the ground truth.

images used is at least 10 percent less due to the removal of redundant images in the AVC (Section 2).

The algorithm of Pound et al. is unable to differentiate between correct and incorrect points as it uses a green colour filter, though without applying a colour filter, no noise is removed and due to the complexity of the objects used in this work, numerous outliers in the point clouds are green due to erroneous matching in PMVS. This is overcome in the algorithm here as model merging, of the point cloud and PPR, can eliminate a large number of outliers, thus reducing the mean and standard deviation.

The algorithm of Pound et al. requires parameters that determine the size of the clusters (Phase 1); during this evaluation, a smaller size was chosen as a larger radius results in more erroneous reconstruction with respect to mean. Manually specifying parameters can be problematic, particularly when dealing with objects that vary in size and structure. Even plants of the same species and age require different parameters due to varying density and size. While some parameters are necessary to be specified by an operator, others should be calculated automatically by inspecting the dataset prior to processing. The proposed algorithm automatically calculates parameters by estimating, for example, the α radius and defining a clear set of cluster constraints should be satisfied by evaluating the data prior to reconstruction.

The algorithm proposed here significantly improves the quality and accuracy of the plant model, in most cases producing a fully connected mesh, as opposed to a set of patches; the output of the Canopy reconstruction algorithm. Consequently, it provides a more visually complete model and one that faithfully represents the target object. By performing cell merging (Phase 4), it is possible to cover a greater proportion of the object surface, while the algorithm of Pound et al. results in a series of triangulated clusters, each with a surrounding gap, consequently reducing the surface area and not faithfully representing the object. With respect to plant phenotyping, fully connected models are more desirable due to their realism. Moreover, connectivity of surfaces allows for more accurate modelling of light and photosynthesis dynamics, or of canopy movement [6], [7], [8], [15], [16], [32].

5 CONCLUSION AND DISCUSSION

In recent years, plant phenotyping has become increasingly popular, with an upsurge in affordable hardware and continued improvements in analysis tools. As the need for accurate high throughput plant phenotyping increases, so too will the demand for a robust reconstruction method.

Active vision (AV) systems have many advantages over traditional fixed-view and manual methods, particularly with the ability to overcome occlusions; one of the biggest challenges associated with 3D modelling. An AV system is able to intelligently manipulate the viewpoint such that it evades the part of the object causing the occlusion. Other advantages include the ability to overcome the limited field of view and develop image acquisition strategies with guaranteed accuracy and repeatability, which may not necessarily be possible with a manual approach. Moreover, manual methods require that some calibration target is visible in all images to calculate the required camera matrices. For crowded scenes or large target objects it is often problematic to ensure the calibration

target is visible in all views, however, AV overcomes this by estimating and maintaining the required camera matrices through a prior calibration step (Section 2.1).

The imaging and reconstruction pipeline presented here is well-suited for use in a large-scale phenotyping system. Image-based systems are highly desirable as a method of plant phenotyping; providing the information needed to calculate a number of key plant traits [33], [34], [35]. 3D reconstruction methods have previously been shown to accurately preserve and represent key physiological measurements including shoot traits such as leaf curling, shape and area or root traits such as morphology, geometry and topology [6], [7], [36]. Such features can be important yield determining traits for crops. Compared to other systems required for capturing plant structure (e.g., laser scanning such as LiDAR), image-based systems based on RGB cameras are cheap, flexible and can be used in multiple different settings.

The issues associated with designing a high throughput phenotyping system remain a challenging problem, with many systems still reliant on static or manual image capture. The approach here provides a flexible framework from which future systems can be evolved. Within phenotyping systems, numerous sensors can be employed to gather information on plant growth and function. These include hyperspectral cameras, fluorescence and chlorophyll fluorescence cameras or sensors, near-infrared cameras and laser scanners. Different modules containing the different sensors can be combined within one system (e.g., [37]) to maximise the amount of information gathered, and thus the AVC and associated pipeline could form an image-based reconstruction module for a larger scale system. Such a system will be invaluable as the demand for increased productivity of crops continues to increase over the coming years. Applications for field grown crops may face further challenges from occlusion, illumination and real time movement.

Currently, the majority of the pipeline is automated, but user interaction is required in order to place the plant on the turntable, beginning the process with a single click and remove the plant following imaging. Multiple large-scale phenotyping systems have been created in which plants are supplied to the sensors via conveyor belts [19], [20], [21]. Alternatively, within the University of Nottingham's Hounsfield Facility, plants are supplied to the X-ray μ CT scanner via an automated laser guided vehicle and automatic FANUC robotic arm. Similar methods could be employed to this pipeline to fully automate 3D reconstruction of target plants. Future work here will look at the addition of other tools to the robot arm, such as a gripper, to manipulate the plant based on the 3D model.

ACKNOWLEDGMENTS

The authors would like to thank Dr. Alexandra Burgess for help with figures and writing, and Dr. Craig Sturrock for providing the X-ray CT imagery used here. This work was funded by the Engineering and Physical Sciences Research Council PhD Studentship Award 1499261 (to J.A.G.), the School of Biosciences, University of Nottingham and the Biotechnology and Biological Sciences Research Council [BB/R004633/1], "The 4-dimensional plant: enhanced mechanical canopy excitation for improved crop performance".

REFERENCES

- [1] M. A. Sutton, et al., *The European Nitrogen Assessment: Sources, Effects and Policy Perspectives*. Cambridge, UK: Cambridge University Press, 2011.
- [2] F. Tardieu, L. Cabrera-Bosquet, T. Pridmore, and M. Bennett, "Plant phenomics, from sensors to knowledge," *Current Biol.*, vol. 27, pp. R770–R783, 2017.
- [3] L. Quan, P. Tan, G. Zeng, L. Yuan, J. Wang, and S. B. Kang, "Image-based plant modeling," *ACM Trans. Graph.*, vol. 25, no. 3, Jul. 2006, Art. no. 599.
- [4] H. Wang, W. Zhang, G. Zhou, G. Yan, and N. Clinton, "Image-based 3D corn reconstruction for retrieval of geometrical structural parameters," *Int. J. Remote Sens.*, vol. 30, no. 20, pp. 5505–5513, Sep. 2009.
- [5] C. Preuksakarn, P. Ferraro, E. Nikinmaa, C. Godin, J. Durand, and E. Nikin-, "Reconstructing Plant Architecture from 3D Laser scanner data," in *Proc. 6th Int. Workshop. Functional-Struct. Plant Models*, 2010, pp. 12–17.
- [6] A. J. Burgess, R. Retkute, T. Herman, and E. H. Murchie, "Exploring relationships between canopy architecture, light distribution, and photosynthesis in contrasting rice genotypes using 3d canopy reconstruction," *Front. Plant Sci.*, vol. 8, May 2017, Art. no. 734.
- [7] A. J. Burgess, et al., "High-resolution three-dimensional structural data quantify the impact of photoinhibition on long-term carbon gain in wheat canopies in the field.," *Plant Physiol.*, vol. 169, no. 2, pp. 1192–204, Oct. 2015.
- [8] A. J. Burgess, R. Retkute, M. P. Pound, S. Mayes, and E. H. Murchie, "Image-based 3D canopy reconstruction to determine potential productivity in complex multi-species crop systems," *Ann. Bot.*, vol. 119, pp. 517–532, 2017.
- [9] L. Quan, P. Tan, G. Zeng, L. Yuan, J. Wang, and S. B. Kang, "Image-based plant modeling," in *ACM SIGGRAPH 2006 Papers (SIGGRAPH '06)*, ACM, New York, NY, USA, pp. 599–604, 2006, doi: <https://doi.org/10.1145/1179352.1141929>.
- [10] A. Reche-Martinez, I. Martin, and G. Drettakis, "Volumetric reconstruction and interactive rendering of trees from photographs," *ACM Trans. Graph.*, vol. 23, no. 3, pp. 720–727, 2004.
- [11] N. Ivanov, P. Boissard, M. Chapron, and B. Andrieu, "Computer stereo plotting for 3-D reconstruction of a maize canopy," *Agricultural and Forest Meteorology*, vol. 75, no. 1–3, pp. 85–102, Jun. 1995.
- [12] J. A. Gibbs, M. Pound, A. P. French, D. M. Wells, E. Murchie, and T. Pridmore, "Approaches to three-dimensional reconstruction of plant shoot topology and geometry," *Functional Plant Biol.*, vol. 44, no. 1, pp. 62–75, Jan. 2017.
- [13] S. M. Seitz, B. Curless, J. Diebel, D. Scharstein, and R. Szeliski, "A comparison and evaluation of multi-view stereo reconstruction algorithms," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, 2006, vol. 1, pp. 519–528.
- [14] M. P. Pound, A. P. French, E. H. Murchie, and T. P. Pridmore, "Automated recovery of three-dimensional models of plant shoots from multiple color images," *Plant Physiol.*, vol. 166, no. 4, pp. 1688–1698, Dec. 2014.
- [15] A. J. Burgess, et al., "The 4-dimensional plant: Effects of wind-induced canopy movement on light fluctuations and photosynthesis," *Front. Plant Sci.*, vol. 7, 2016, Art. no. 1392.
- [16] A. J. Townsend, et al., "Suboptimal acclimation of photosynthesis to light in wheat canopies," *Plant Physiol.*, vol. 176, no. 2, pp. 1233–1246, Feb. 2018.
- [17] T. Ju, "Fixing geometric errors on polygonal models: A survey," *J. Comput. Sci. Technol.*, vol. 24, no. 1, pp. 19–29, Mar. 2009.
- [18] J. A. Gibbs, M. P. Pound, A. P. French, D. M. Wells, E. H. Murchie, and T. P. Pridmore, "Plant phenotyping: An active vision cell for three-dimensional plant shoot reconstruction," *Plant Physiol.*, vol. 178, pp. 524–534, 2018.
- [19] "Photon Systems Instruments, Drasov, Czech Republic", <http://psi.cz/>
- [20] "LemnaTec, Aachen, Germany", <https://www.lemnatec.com/>
- [21] "Qubit Phenomics, Kingston Ontario, Canada", <https://qubitphenomics.com/>
- [22] D. Arend, et al., "Quantitative monitoring of Arabidopsis thaliana growth and development using high-throughput plant phenotyping," *Sci. Data*, vol. 3, Aug. 2016, Art. no. 160055.
- [23] S. Arvidsson, P. Pérez-Rodríguez, and B. Mueller-Roeber, "Methods a growth phenotyping pipeline for Arabidopsis thaliana integrating image analysis and rosette area modeling for robust quantification of genotype effects," *New Phytol.*, vol. 191, pp. 895–907, 2011.
- [24] A. Junker, et al., "Optimizing experimental procedures for quantitative evaluation of crop plant performance in high throughput phenotyping systems," *Frontiers Plant Sci.*, vol. 5, Jan. 2015, Art. no. 770.
- [25] J. Denavit, "A kinematic notation for lower-pair mechanisms based on matrices," *J. Appl. Mech.*, vol. 20, pp. 215–222, 1955.
- [26] F. Dornaika and R. Horaud, "Simultaneous robot-world and hand-eye calibration," *IEEE Trans. Robot. Autom.*, vol. 14, no. 4, pp. 617–622, Aug. 1998.
- [27] C. Zhong, D. Miao, R. Wang, and X. Zhou, "DIVFRP: An automatic divisive hierarchical clustering method based on the furthest reference points," *Pattern Recognit. Lett.*, 29, pp. 2067–2077, 2008.
- [28] D. MacKay, "An example inference task: Clustering," *Inf. Theory, Inference Learn. Algorithms*, Cambridge Univ. Press, vol. 20, pp. 284–292, 2003
- [29] S. K. Tasoulis and D. K. Tasoulis, "Improving principal direction divisive clustering," in *14th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining (KDD 2008)*, Workshop on Data Mining using Matrices and Tensors, Las Vegas, USA, Aug. 2008.
- [30] S. Reddy, "Image Segmentation by Using Linear Spectral Clustering," *J. Telecommun. Syst. Manage.*, vol. 5, no. 3, pp. 1–5, Oct. 2016.
- [31] L. R. Herrmann, "Laplacian-isoparametric grid generation scheme," *J. Eng. Mech. Div.*, vol. 102, pp. 749–756, 1976.
- [32] Q. Song, G. Zhang, and X.-G. Zhu, "Optimal crop canopy architecture to maximise canopy photosynthetic CO₂ uptake under elevated CO₂: A theoretical study using a mechanistic model of canopy photosynthesis," *Functional Plant Biol.*, vol. 40, no. 2, Mar. 2013, Art. no. 108.
- [33] D. Houle, D. R. Govindaraju, and S. Omholt, "Phenomics: The next challenge," *Nature Rev. Genetics*, vol. 11, pp. 855–866, 2010.
- [34] T. T. Santos and A. A. De Oliveira, "Image-based 3D digitizing for plant architecture analysis and phenotyping," in *Conf. graph. patterns images*, Ouro Preto, Workshop on industry applications.[SI]: Conference Publishing Series, vol. 25, 2012.
- [35] J. W. White, et al., "Field-based phenomics for plant genetics research," *Field Crops Res.*, vol. 133, pp. 101–112, 2012.
- [36] G. Lobet, L. Pages, and X. Draye, "A novel image-analysis toolbox enabling quantitative analysis of root system architecture," *Plant Physiol.*, vol. 157, no. 1, pp. 29–39, Sep. 2011.
- [37] N. Virlet, K. Sabermanesh, P. Sadeghi-Tehran, and M. J. Hawkesford, "Field scanalyzer: An automated robotic field phenotyping platform for detailed crop monitoring," *Functional Plant Biol.*, vol. 44, pp. 143–153, 2017.



Jonathon A. Gibbs received the BSc (hons) degree of the First Class in software systems from the University of Nottingham, in 2010, and the PhD degree from the University of Nottingham, in 2018. He previously held the position of manager of finance at Lockwood Publishing (2014) and currently works as a research fellow at the University of Nottingham. His research interests include computer vision, particularly 3D modelling, deep learning, hyper-heuristics, and active vision and robotics. In particular, automating extracting plant traits using robotics and 3D modelling.



Michael Pound is a Nottingham research fellow in the School of Computer Science, UoN, where he is a member of the Computer Vision Laboratory. His research focuses on applying computer vision and machine learning approaches to plant phenotyping problems. He has published a number of high-impact papers in *Plant Physiology*, *The Plant Cell*, and *GigaScience* along with associated tools for accurate plant phenotyping. Recently, his work has focused on accurate multi-view reconstruction of plant shoots, developing methods that underpin numerous publications on light modelling within crop canopies (CI, BBSRC project BB/R004633/1). He is also applying deep machine learning approaches to the detection of shoot and root features, and is currently developing high-throughput root analysis pipelines using deep machine learning (PI, BBSRC project BB/P026834/1). He regularly reviews for journals covering emerging technologies in plant phenotyping, including *Plant Methods*, *Functional Plant Biology*, *PLOS One*, and *Trends in Plant Sciences*.



Andrew P. French received the PhD degree in computer science, specializing in computer vision, in 2005. He is an associate professor in the Schools of Computer Science and Biosciences at the University of Nottingham. He works on the analysis of biological images, including developing novel approaches for the analysis and phenotyping of plants and crops, from cell scale to the field. Currently, he is pursuing deep machine learning approaches to allow phenotyping under challenging imaging conditions.



Darren M. Wells received the BSc (Hons) degree in biology from University College London, in 1997, and the PhD degree from the University of London, in 2001. He previously held research posts at Rothamsted Research and is currently a senior research fellow in plant and crop biophysics at the University of Nottingham. He has authored more than 50 peer-reviewed publications and 13 book chapters. His research addresses fundamental and applied problems in plant and crop science via phenotyping at the cellular, organ, and whole plant level.



Erik H. Murchie received the PhD degree from Sheffield University, in 1995. He is an associate professor in crop science at the University of Nottingham (UoN). He has worked on the physiology of sub-optimal crop photosynthesis since 1997, including collaboration with the International Rice Research Institute (IRRI) and the Maize and Wheat Improvement Centre (CIMMYT). He has focused particularly on photoprotection and understanding how such processes are integrated at the canopy level. He has more than 60 peer-reviewed publications in photosynthesis and photoprotection and ten book chapters. He is an editor for the *Annals of Botany* and *Frontiers in Plant Abiotic Stress*.



Tony P. Pridmore received the BSc (Hons) degree in computer science from the University of Warwick, in 1982, and the PhD degree in computer vision from the University of Sheffield, in 1987. He is a professor of computer science with the School of Computer Science, University of Nottingham, where he leads the Computer Vision Laboratory. He has worked on image-based plant phenotyping since 2008. He is a director of the UK plant phenotyping network PhenomUK, and (joint) editor in chief of *Plant Methods*.

▷ For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/csdl.