CrossMark

Validation of plant part measurements using a 3D reconstruction method suitable for high-throughput seedling phenotyping

Franck Golbach[1] · **Gert Kootstra**[1] · **Sanja Damjanovic**[1] · **Gerwoud Otten**[1] · **Rick van de Zedde**[1]

Abstract In plant phenotyping, there is a demand for high-throughput, non-destructive systems that can accurately analyse various plant traits by measuring features such as plant volume, leaf area, and stem length. Existing vision-based systems either focus on speed using 2D imaging, which is consequently inaccurate, or on accuracy using time-consuming 3D methods. In this paper, we present a computer-vision system for seedling phenotyping that combines best of both approaches by utilizing a fast three-dimensional (3D) reconstruction method. We developed image processing methods for the identification and segmentation of plant organs (stem and leaf) from the 3D plant model. Various measurements of plant features such as plant volume, leaf area, and stem length are estimated based on these plant segments. We evaluate the accuracy of our system by comparing the measurements of our methods with ground truth measurements obtained destructively by hand. The results indicate that the proposed system is very promising.

Keywords High-throughput phenotyping · Seedling phenotyping · 3D Plant model · Plant trait measurements

1 Introduction

Common morphological plant traits of interest include parameters such as main-stem height, size and inclination, petiole length and initiation angle, and leaf width, length, inclination, thickness, area, and biomass [1]. Until recently, most of these observations and their quantification, known also as plant phenotyping, relied on human assessments and measurements [2], being both costly and slow. Additionally, the exponentially growing amount of possibilities in all fields of plant sciences, such as genomics, and breeding make the application of automated methods in plant phenomics vital.

Plant phenotyping is the set of methodologies and protocols used to measure plant growth, architecture, and composition with a certain accuracy and precision at different scales of organization, from organs to canopies [3]. The plant phenotype includes complex plant traits that are assessed through measurement of the root morphology, biomass, leaf characteristics, fruit characteristics, yield-related traits, photosynthetic efficiency, and biotic and abiotic stress response [4]. Plant phenotyping has become a major field of research in plant breeding stimulated by the rapid development of plant genomics. Phenotyping is addressed by combining novel technologies such as non-invasive imaging, spectroscopy, image analysis, robotics and high-performance computing [1]. Plant measurements have been done on different scales, at the level of cells, organs, root systems, plants, and canopies, where different sensors are used for each scale [5]. Measurement are performed using two-dimensional (2D) images or 3D models. 3D geometrically accurate models of plants enable more accurate measurements and modelling of biological processes, such as, photosynthesis, yield prediction, and plant-growth modelling. Some non-invasive phenotyping systems make use of 2D hyperspectral imaging such as HyperART [6], or systems for measurement of structural parameters of plant canopies [7,8].

Visible-light imaging plays a significant role in measuring plant traits, such as, biomass, plant height, stem length, and leaf area. In current systems used in the field, a single

✉ Gert Kootstra
gert.kootstra@wur.nl

Franck Golbach
franck.golbach@wur.nl

[1] Wageningen UR Food & Biobased Research, Wageningen, The Netherlands

digital camera is mounted above the plant to generate a top view of the plant, sometimes accompanied with one or two additional cameras to generate side views, for the calculation of the shoot biomass or leaf area using the imaging software [9]. However, the estimation of the biomass of a plant from 2D images results in high inaccuracies, while a robust and accurate method is required for high-throughput phenotyping as proposed in [10]. The measurement accuracy of such methods highly depends on the relative position of the camera with respect to the plant. Accurate measurement of leaf area can be obtained if the camera observes a frontal view of the leaf. However, the actual observation may deviate significantly, resulting in a poor estimation of the true leaf area. In general, the measurement of leaf area from the 2D images is characterized by large standard deviations.

Phenotyping systems relying on 2D images are dominant in the literature [11–15]. This paragraph describes several 2D image-based systems together with their applications. The PHENOPSIS phenotyping system [12] uses 2D images to investigate the development of the Arabidopsis thaliana in different conditions. Phenoscope [16] is a 2D image-based phenotyping system monitoring rosette size and expansion rate during the vegetative stage, with automatic image analysis allowing manual correction. This system continuously rotates pots to minimize the influence of the camera perspective. In the initial stage of the growth of a seedling, 2D images can capture its development reliably, as the plant structure is still flat [17]. HTPheno [13] is an image analysis pipeline for high-throughput plant phenotyping, which provides the possibility to analyse colour images of plants that are taken in two different views (top view and side view) during a screening. Within the analysis, different phenotypical parameters for each plant, such as, height, width, and projected shoot area of the plants are calculated for the duration of the screening. HTPheno was applied to analyse two barley cultivars. Also PHIDIAS [18] is a system for high-throughput phenotyping, in this case successfully tested on Arabidopsis. In GROWSCREEN [15], leaf area and relative growth rate were measured based on images from a camera placed above an array of seedlings. The phenotyping system GlyPh [19] is a low-cost platform for phenotyping plant growth and water use, which allows the evaluation of plants growing in individual pots. In GlyPh, top- and side-view images of the plants are captured, to measure traits such as height, width, and projected leaf area.

A way to overcome the influence of the relative orientation of the camera and the leaf, in the phenotyping systems based on 2D imaging, on the measurement accuracy is to use additional height or depth information. These systems acquire a 2.5-dimensional (2.5D) model of plants. The depth can be measured by range cameras based on lasers (LIDAR) [20], by time-of-flight (ToF) cameras [8,21–23], or using stereoscopic cameras [7,24]. Other systems use projected-light cameras, such as Microsoft Kinect, to obtain 2.5D plant models as in [22].

Special class of the shoot-phenotyping systems are the systems which generate and use a 3D plant model. The fast development in computing power results in the development of techniques that process complete 3D models [25]. 3D plant models can be produced by several different techniques, such as, hand-held laser scanning [26], or using multi-view stereo or structure-from-motion to generate a 3D model [27,28]. However, not each 3D imaging technique is sufficiently fast to provide a high-throughput system. The speed of the system depends on the plant morphology and the challenges the imaging system needs to overcome to generate the 3D model. A promising new technique for generating 3D models is described in [27]; however, the time to recover the 3D model is significant and directly proportional to the plant morphology complexity. The 3D plant modelling from images presented in [28] is semi-automatic and uses 35–50 input images. In [25], 3D model is generated from 64 images using a commercially available 3D reconstruction method [29] based on an accurate shape-from-silhouette method. In [27], the multi-view stereo algorithm, [30], in conjunction with shape optimisation in the post-processing step, uses 36–64 input images to generate surface-based 3D model representations. Stereo-imaging and mesh processing-based systems, such as GROWSCREEN 3D [31], allowing more accurate measurements of leaf area, and extraction of additional volumetric data. Although those 3D systems result in very accurate 3D plant reconstructions and associated measurements, they are currently too slow to be used in high-throughput phenotyping system.

In this paper, we focus on high-throughput, non-invasive and non-destructive seedling phenotyping from a complete 3D plant model. We estimate the leaf area based on the 3D plant model, as estimated leaf area is one of the most often used traits in plant-phenotyping experiments. We have developed a system for fast 3D plant reconstruction from a small number of calibrated images using a shape-from-silhouette approach. Shape-from-silhouette methods were first introduced by [32] and further improved by, i.a., [33–35]. The method suits our needs because of its robustness and its potential for optimized implementation. The method results in a so-called visual hull of the object [36,37], represented by the set of occupied points in the voxel space. Based on such a 3D representation of a plant, we can accurately identify and segment the stem and the leaves from the 3D plant model, which allows us to accurately measure different plant features.

Using the 3D plant model, we accurately measure different plant features. We measure the leaf area by identifying the voxels that are on the surface boundary of the leaf. The length of the leaf is measured by taking the distance from the stipule (where the leaf connects to the stem), through the centre of

the leaf to the leaf tip (apex). Leaf width is estimated by taking the longest distance perpendicular to this line. We also measure stem length by tracing the stem from the plug up to the first leaf.

We will present our method in Sect. 2. The experimental setup and results are described in Sects. 3 and 4. We end the paper with a discussion on the method and future research in Sect. 5.

2 Methods

Our system for high-throughput phenotyping is illustrated in Fig. 1. The seedling is surrounded by different cameras, observing the object from different perspectives (Fig. 1a). The system can deal with a variable number of cameras. The quality of the 3D reconstruction generally improves when more cameras are used. However, as the increase in quality gets smaller with higher number of cameras and the computational time increases linearly with the number of cameras, a trade-off between accuracy and speed needs to be made. In our experiments, we used 10 cameras as an optimum of the trade-off. The silhouettes of the seedling in the acquired camera images (Fig. 1b) are used to reconstruct the object in 3D through a shape-from-silhouette method (Fig. 1c). Next,

the reconstruction of the whole plant is segmented in stem and leafs (Fig. 1d). Based on the whole plant reconstruction and the segmented stem and leafs, different quality features are calculated.

The 3D reconstruction method is described in more detail in Sect. 2.1. The leaf/stem segmentation is outlined in Sect. 2.2, and the different quality methods are described in Sects. 2.3 and 2.4.

2.1 High-throughput 3D reconstruction

We developed a shape-from-silhouette [38] method that calculates a 3D reconstruction of the sample from the silhouettes in all camera images. To perform the reconstruction, the precision position and orientation of all cameras with respect to the workspace need to be known, which are estimated through a calibration procedure. This procedure is outlined in Sect. 2.1.1, followed by a description of the fast 3D reconstruction method in Sect. 2.1.2 and a short discussion of the method in Sect. 2.1.3.

2.1.1 Calibration

To be able to perform the shape-from-silhouette method, the position of all cameras must be known with respect to the

(a) Multi-camera setup

(b) Image acquisition

(d) Leaf-stem segmentation
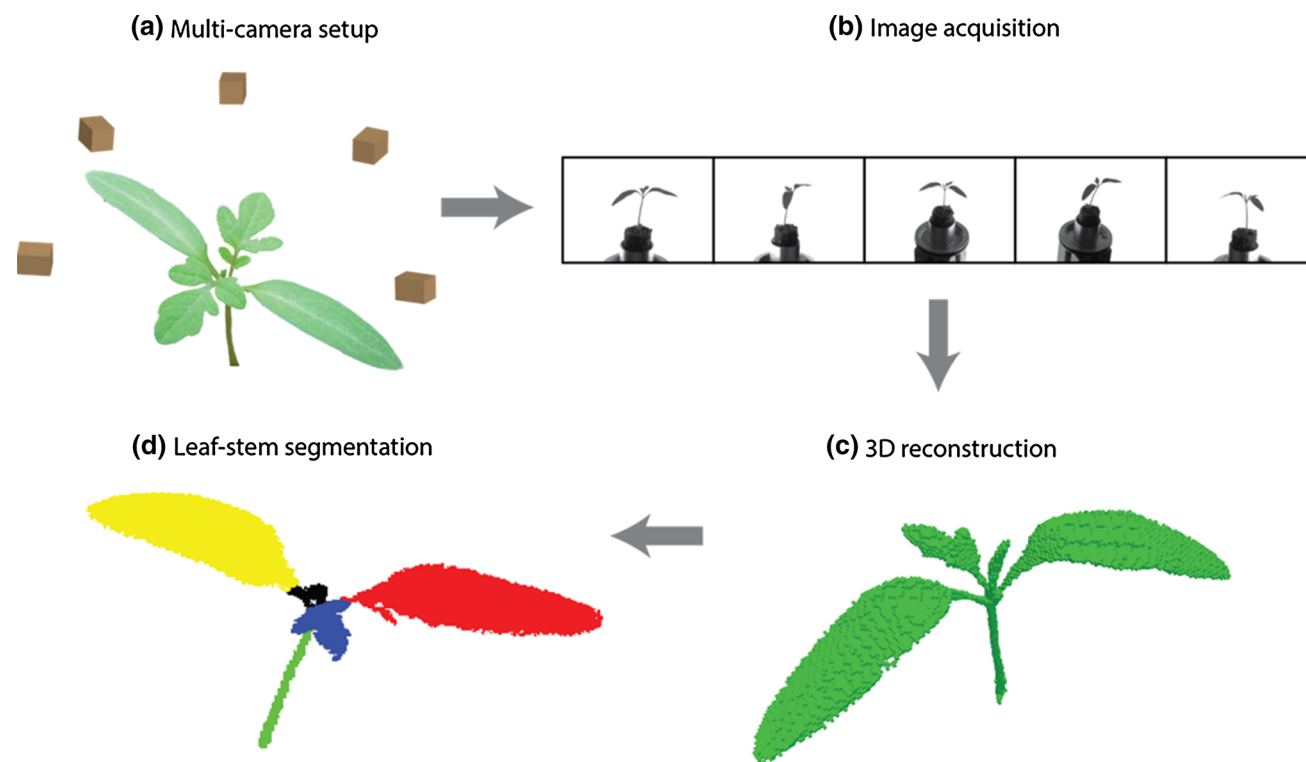
(c) 3D reconstruction

Fig. 1 An overview of 3D-reconstruction and segmentation pipeline. **a** The plant is observed by multiple cameras (10 used in the experiments). **b** Each camera provides an image from which the silhouette of the plant is determined. **c** Using a shape-from-silhouette method, the 3D shape of the plant is reconstructed. **d** The 3D model is segmented in stem (*green*) leafs (*coloured*) and unknown *black*)
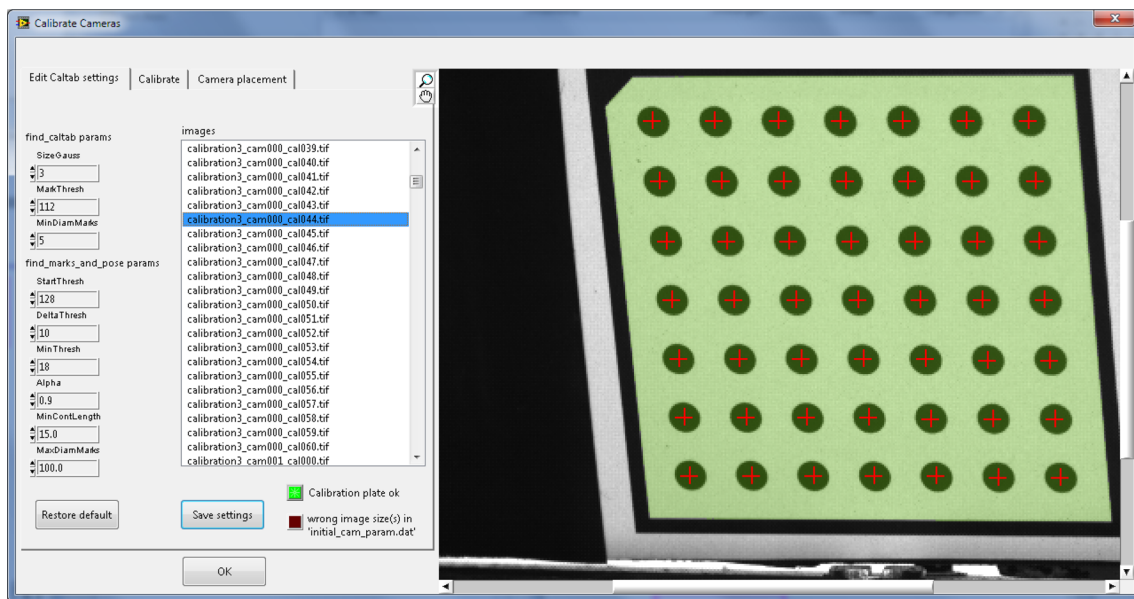
Fig. 2 The calibration procedure involves multiple views of the calibration plate by each camera; the plate is detected using Halcon software. For each camera, a 'calibration pathway' is calculated to one specific calibration plate position, which serves as a real-world reference. The pathway consists of a chain of affine transformations. The software determines all available pathways, and selects the chain with the smallest cumulative error

workspace. For every voxel in the 3D workspace, we need to know to which pixel(s) it maps on each of the cameras. To be able to calculate this mapping, we need to determine the internal and external parameters of the cameras. This procedure is known as multi-camera calibration [39]. The internal parameters describe the parameters of the camera and the lens, and include the focal distance of the lens, the size of a pixel on the sensor, the coordinates of the optical centre of the sensor and the radial distortion coefficient. The external parameters define the position and orientation of the camera in space with respect to a chosen origin, in our case the centre of the workspace.

To calibrate the system, a regular dot-pattern calibration plate is used. The plate is placed in a number of different orientations and positions and is synchronously observed by the cameras. In each camera image, a number of unique points can automatically be extracted. From the corresponding points between different cameras and by knowing the true dimensions of the plate, the internal and external parameters can be estimated. In general, the estimation improves when the calibration plate is placed in more poses, but typically around 20–25 observations are sufficient for a good calibration. We use the position and orientation of one of the calibration plates to determine the origin of the voxel space.

The whole procedure of multi-camera calibration was developed using Labview (Fig. 2), and is based on the single camera procedures provided by Halcon. Once both the internal and external camera parameters are determined for each camera, one specific plate position is chosen as a real-world reference. Through a chain of affine transformations, the correspondence of all camera positions to the real-world coordinates can be calculated. The software will optimize the chain by determining the smallest overall RMS error.

2.1.2 3D reconstruction

By applying the calibration, i.e. the projection matrices for each camera resulting from the external parameters and the corrected camera model, the mapping between the voxels in the 3D workspace and the pixels in the 2D camera images can be determined.

Knowing the projection of the voxels on pixels in the images, the object under observation can be reconstructed from the silhouettes in the camera images. The method for reconstruction is given in pseudo code in Fig. 3. All camera images are first segmented in foreground and background, resulting in binary silhouette images B_k. <**2.1**> This is done using a procedure known as background subtraction, where segmentation is obtained by subtracting an image containing only the background from an image containing the seedling. A pixel is part of the foreground when the Euclidian distance in RGB space between the two images is larger than a threshold value. The optimal value of this threshold is set manually for each camera to correct for small differences in apertures and lighting conditions.

```
for each camera image I_k
     B_k <- foreground_background_segmentation( I_k )
end

for each voxel v in the voxel space V
     V(v) <- occupied
end

for each voxel v in the voxel space
     for each camera k
          i <- get_corresponding_pixel( v, P_k )
          if B_k(i) = background
               V(v) <- empty
               break
          endif
     end
end
```

Fig. 3 Pseudo code of the shape-from-silhouette method

Next, all voxels in the voxel space V are initially set to 'occupied'. The occupancy of each voxel is then investigated by looking at the corresponding pixels in all camera images, which are determined using the camera parameters P_k. If all corresponding pixels are labelled as foreground, the voxel must be part of the object and remains occupied. Otherwise, the voxel is set to 'empty'.

Figure 4 shows three different views on a 3D reconstruction of a tomato seedling. The plant is reconstructed well. Stem and leaf shapes are clearly visible. Also smaller structures like the petioles are included in the model.

2.1.3 Discussion on the 3D reconstruction method

We intentionally chose the fastest and simplest implementation of space carving. This approach is known to have minor drawbacks: a voxel is eliminated from the reconstruction volume if one of the image pixels covered by a voxel shows background. Thus, as a voxel usually covers multiple pixels in the same image, it is not ensured that a voxel is kept if at least one corresponding pixel in every image contains foreground. In some cases, this may lead to discontinuities in the reconstruction. To avoid this, the finest structures need to have a diameter of at least $2\sqrt{3} \approx 3.5$ voxels for the worst-case scenario when the structures are oriented diago-

nally. In our experiments, we used a voxel resolution of 0.25 mm/voxel, whilst the diameter of the leaf stems was generally well above 1 mm. See also the description in Sect. 3.1.

Another known disadvantage of the shape-from-silhouette method is that only those parts of the object can be reconstructed well that are visible in the contour images, which means that occluded parts and concavities cannot be reconstructed. However, since plants, especially seedlings, are relatively open structures, all relevant parts are visible and this can hardly be called a drawback. Issues with occlusion are not unique for this method.

One of the strengths of the method is its high flexibility. The number of cameras, their viewpoints and their optics can be altered, requiring only a recalibration of the system, which is easy to perform. Also the dimensions of the workspace and the voxel resolution can easily be changed. The workspace can be adapted to the size and structure of the plant, which allows to work with very small workspaces of a few millimetres to workspaces that span several meters. There is, however, a practical limit to the number of voxels in the voxel space, as the method needs to store tables in memory with the relation of all voxels to the corresponding pixels in all images. The flexibility to easily adjust the size and resolution of the workspace allows the system to work with various types and sizes of plants, as long as the plant structure is relatively open, not to create too much occlusions and concavities.

<3.2> A major advantage of the space carving method is its ability to be used for in-line purposes at high speed using relative low-cost hardware, making it suitable for large-scale phenotyping experiments.

2.2 Stem and leaf segmentation

To be able to calculate the stem and leaf features, the complete 3D representation of the plant needs to be segmented into stems and leaves. We developed a segmentation method exploiting the structural layout of tomato seedlings (see Fig. 5 for a 2D illustration of the method). This algorithm is based on the breath-first flood-fill algorithm with a 26-connected neighbourhood, which iteratively fills a structure. In our case,

Fig. 4 Three different views on a 3D reconstruction of a tomato seedling

we start with the lowest point in the voxel representation, the bottom of the main stem of the plant (red square in figure). In every iteration of the flood-fill algorithm, the neighbouring points that are not yet filled are added and the iteration number is stored for these voxels (illustrated by decreasing brightness of the squares). As long as the main stem is filled, the added points in each generation are closely located in space. However, at the point that the first side branches and leaves appear, the spread of the newly added points increases. When this exceeds a given threshold, that iteration is labelled as the end of the main stem (yellow square). This threshold depends on the resolution of the voxel space and the char-
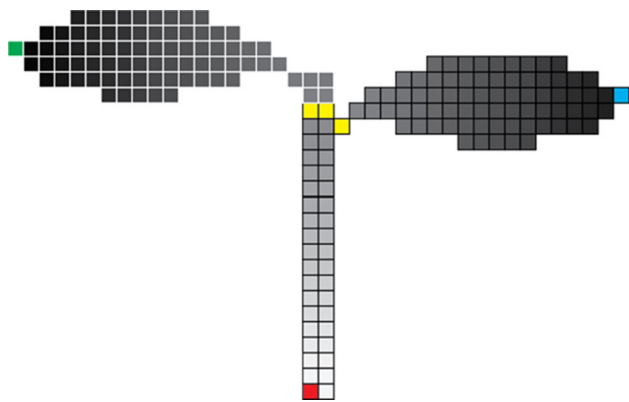
acteristics of the plant type. In our experiments, we used a threshold of 4 mm. When the flood fill is completed, we start a leaf segment with the last added point (green square), which will be one of the leaf tips, and subsequently backtrack the flood fill until the end point of the main stem is reached. In the process, all voxels are added to the leaf segment (squares white borders). We perform the same procedure from the next leaf tip (blue square), resulting in another leaf segment, and repeat this until all voxels have been labelled as either stem or leaf. If a leaf consists of a number of lobes, this algorithm separates the leaf into different segments. To correct this, segments that connect at a place other than the end of the main stem are merged.

2.3 Measuring leaf features

After stem and leafs are segmented, we calculate relevant phenotypic features of the leafs, specifically leaf length, leaf width, and leaf surface. These features are very predictive for the quality of the plant, as the size of the leafs play an important role in growth through the photosynthetic process.

2.3.1 Leaf length

We define the length of a leaf as the length of the midrib from the stipule (where the leaf connects to the stem) to the apex (the tip of the leaf). The leaf tip is determined by the segmentation method (see Fig. 5, blue and green point). It is the point on the surface of the leaf that is furthest from the stem end point. The stipule is determined in reverse order, as the point on the leaf surface that is furthest from the leaf tip. This method assumes an elongated leaf shape. Both points are marked in Fig. 6a by a purple cross. The Euclidian distance between these two points would be a very crude approximation of the leaf length and always an underestimation, as
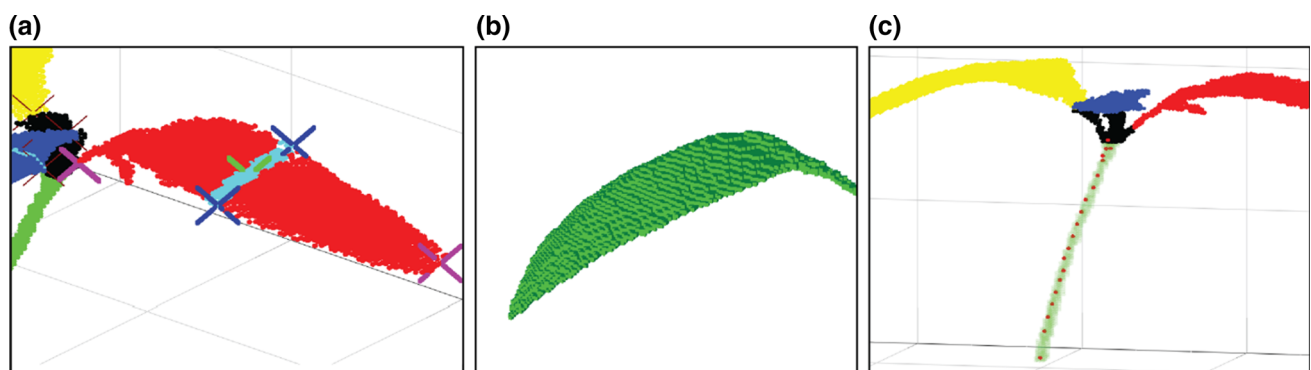

Fig. 5 Illustration of the stem–leaf segmentation algorithm. The methods starts filling the structure from the bottom (*red*). Subsequently, neighbouring points are filled and the iteration number is stored for these voxels (indicated by the gradient). As long as the neighbouring points are close together in space, we trace the stem. Once they spread out, we mark the end of the stem (*yellow*). When all voxels are filled, we start from the last added point, a leaf tip (*green*) and back-trace until the end of the stem. All back-traced voxels are added to the leaf segment. The same process is repeated for the other leaf tips (*blue*). This illustration gives a 2D example, however, in reality, it performs on a 3D reconstruction of the plant

Fig. 6 Measuring leaf and stem features. **a** Leaf length is approximated by the length of the 3D polyline running from the stipule (*purple cross*) to the centre of the leaf (*green cross*) to the tip of the leaf (*purple cross*). Leaf width is determined by the line between the two *blue crosses*. **b** The leaf surface is approximated by the leaf's surface voxels. **c** The stem length is approximated by the length of the polyline running through the stem-skeleton points (*red circles*)

a leaf is typically curved in 3D. Instead, we estimate the leaf length with an additional point in the middle of the leaf (green cross in Fig. 6a). To determine this midpoint, a band of points halfway between the begin and end point of the leaf is selected (marked light blue in Fig. 6a). The midpoint is set to the centroid of this band of points. The leaf length is then calculated as:

$$l^{\text{leaf}} = \left| \overrightarrow{\mathbf{m} - \mathbf{s}} \right| + \left| \overrightarrow{\mathbf{a} - \mathbf{m}} \right|,$$

where $\mathbf{s}$ is the stipule, $\mathbf{m}$ the midpoint and $\mathbf{a}$ the apex. $|\mathbf{x}|$ gives the length of the vector $\mathbf{x}$.

In a similar fashion, the length of the midrib in 3D can be determined more precisely by adding additional points between stipule and apex. However, we approximate the length by this three-point poly line, to keep computational costs low, as we aim to develop a high-throughput phenotyping system.

2.3.2 Leaf width

The algorithm for finding the width of the leaf aims to find the widest part of the leaf perpendicular to the axis through the begin and end points of the leaf, indicated by the purple crosses (Fig. 6a). The widest part is defined as the part were the Euclidian distance between the blue crosses is maximal. <**2.2**> To determine the position of the blue crosses, for all leaf points, the orthogonal projection on the line through the leaf's beginning and end point (purple crosses) is calculated. The distance between the purple crosses is divided into a number (20) of equidistant sections. Leaf points having their projection in a section are selected, indicated in blue. They form a band across the leaf. The outermost points in this band on the left, $\mathbf{m}^{\text{l}}$, and on the right, $\mathbf{m}^{\text{r}}$, are used to approximate the width of that section. This is repeated for all sections; the width of the leaf is defined as the maximum of these:

$$w^{\text{leaf}} = \max \left\{ \left| \overrightarrow{\mathbf{m}^{\text{l}} - \mathbf{m}^{\text{r}}} \right| \right\}$$

As with the leaf length this is an approximation that will result in a slight underestimation, because the Euclidian distance is used instead of the 'true' distance over the surface of the leaf.

2.3.3 Leaf area

Growth of a plant is for a large part based on photosynthesis. The amount of photosynthesis, and therefore the rate of growth, is related to the total leaf area of the plant, which is the sum of the surface areas of all leafs. Based on our 3D voxel representation of the leafs, we can accurately determine the leaf surface (see Fig. 6b).

A leaf is reconstructed by a set of voxels, where the thickness of a leaf is generally at least two voxels thick. To get the leaf area, we first determine the set of all surface points, that is, all occupied voxels that neighbour one or more non-occupied voxels. This set contains points on the top and on the bottom of the leaf. The leaf area is defined as the area of the top surface of the leaf. We obtain that value by:

$$a^{\text{leaf}} = \frac{1}{2} \left| S^{\text{leaf}} \right| \cdot r^2,$$

where S^{leaf} is the set of all surface voxels of the leaf, $|.|$ the size of the set, and r is the voxel resolution (mm per voxel). The surface of a voxel is thus approximated by the square of the voxel resolution.

We have chosen to work with this approximation because of its simplicity, but it should be noted that it is only fully correct for horizontally or vertically oriented surfaces. For tilted surfaces, the method will give an underestimation of the area, which is worst for a 45° angle, when the actual area will be underestimated by a factor $\sqrt{2}$. For the set of seedlings used in this experiment, this approximation worked rather well (see Sect. 4.4), but it is a limitation to be considered.

2.4 Measuring stem length

The stem length is another important quality feature of a seedling. Our measure of the stem length is based on the 3D stem segment resulting from the stem/leaf segmentation. The midline, or skeleton, of the stem is determined in 3D by a midline-tracking algorithm finding a number of points on the midline (see red dots in Fig. 6c). This is an iterative algorithm that starts from the lowest point, m_0. Each consecutive point on the midline is selected by searching the N nearest points connecting to the current point that have not yet been visited by the algorithm. From this set of N points, the centroid is calculated, which defines the new point on the midline, m_i. Starting from that new point, the algorithm iteratively continues until all points in the stem segment have been visited. The length of the stem is then determined as:

$$l^{\text{stem}} = \sum_{i=1}^{n} |m_i - m_{i-1}|, \quad M = \{m_0, \ldots, m_n\}$$

2.5 Processing time

Throughout the development, the aim has been to develop a system that is actually capable to act as a high-throughput phenotyping system, sufficient both in speed and accuracy. To achieve the necessary speed, approximations were used where needed.

Table 1 Processing times of all parts of the processing and analysis

Operation	Processing time (ms)
3D reconstruction[1]	20–60
Stem–leaf segmentation[2]	100–200, peaks to 1000
Leaf length[2]	5–10
Leaf width[2]	5–20
Leaf area[2]	10–20
Stem length[2]	5–10

[1] This code has been optimized for speed
[2] This code has not yet been optimized for speed
Most time-consuming part is the stem–leaf segmentation, which also shows the largest variation in processing times

Actual processing speed depends on the size of the voxel space, the number of cameras, and the size and complexity of the plant. All experiments were done using a voxel grid size of $240 \times 240 \times 300$ voxels, 10 cameras, and a PC with an i7 type processor (3.2 GHz). The implementation was done in C++ and Labview.

Table 1 gives an indication of the time needed for all steps of the process.

3 Experimental setup

The main purpose of the experiments described in this section is to evaluate the performance of the 3D measurement system. We assess the accuracy and the usability of the image processing algorithms for phenotyping of individual plants. For each plant in the set, the following features were determined: the length, width and area of the leaves, and the length of the stem.

The seedling set consists of 541 tomato an bell pepper seedlings varying from 6 to 10 days of age. In this growth stadium, typical seedlings have two leaves. The majority of the plants in our set (474) had two leaves, 64 had three leaves, and 3 had four leaves. Seedling sizes varied from about 15 to 65 mm total height.
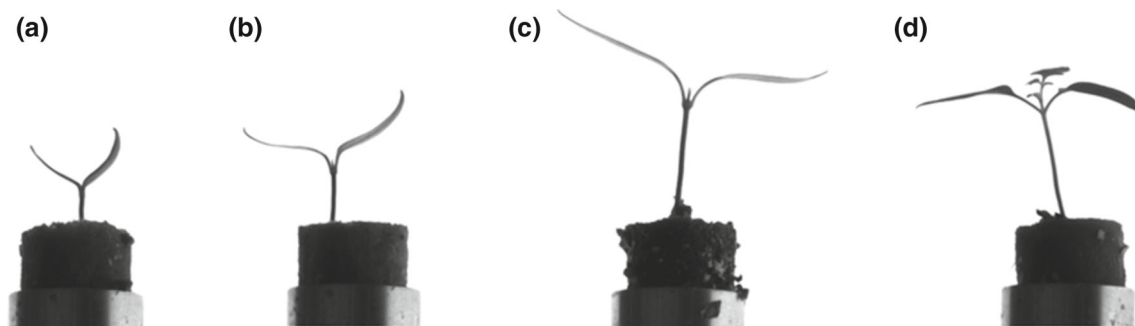
Figure 7 shows typical seedlings and their natural variation in size and shape.

We evaluate our non-destructive 3D acquisition system by comparing the results to ground truth measurements, taken by a calibrated 2D scanning of the separate plant organs: first, the plants were scanned in the 3D acquisition system, then, the leaves and stem of the seedling were cut off manually and placed on a flatbed scanner. Using straightforward 2D blob analysis tools, the required features were obtained and are used as ground truth.

3.1 The 3D acquisition system

The 3D acquisition setup consists of 10 Basler acA1300-30gm cameras with a resolution of 1280×960 pixels, mounted in two semi-circles around the semi-spherical background lightning (see Fig. 8). The cameras are placed at a distance of 900 mm from the seedling. Pentax lenses with a focal length of 25 mm were used. The cameras are triggered to ensure that images from all cameras are taken at exactly the same time. Shutter time was about 1 ms. The illumination is based on a backlight-illumination principle. Dimensions of the 3D voxel space were chosen to accommodate the range of seedling sizes. Our software allows for resizing, resampling, and shifting the 3D voxel space with respect to real-world coordinates afterwards, so the exact position of the seedling does not affect the results. Final resolution of the voxel space was set to 4 voxels/mm, using $240 \times 240 \times 300$ voxels (x, y, z), resulting in real-world dimensions of $60 \times 60 \times 75$ mm. The PC used has an i7 type processor (3.2 GHz) with 4 GB of RAM (restricted by the 32bit Windows operating system used).

3.2 Ground truth and 2D analysis method

To establish the ground truth, we manually cut individual leaves and the remaining stem and positioned them on a flatbed scanner set to 300 dpi (Fig. 9). Stems were manually straightened to minimize measurement error and leafs

(a)   **(b)**   **(c)**   **(d)**

Fig. 7 Four stages of development: two leaves, 20–45 mm (**a–c**); three leaves, approx. 40 mm (**d**). Images shown were taken with one of the cameras of the acquisition system

Fig. 8 The semi-spherical illumination with a seedling in position (**a**); The full cabinet with 10 cameras, illumination, with an opening for an optional conveyer belt (**b**)
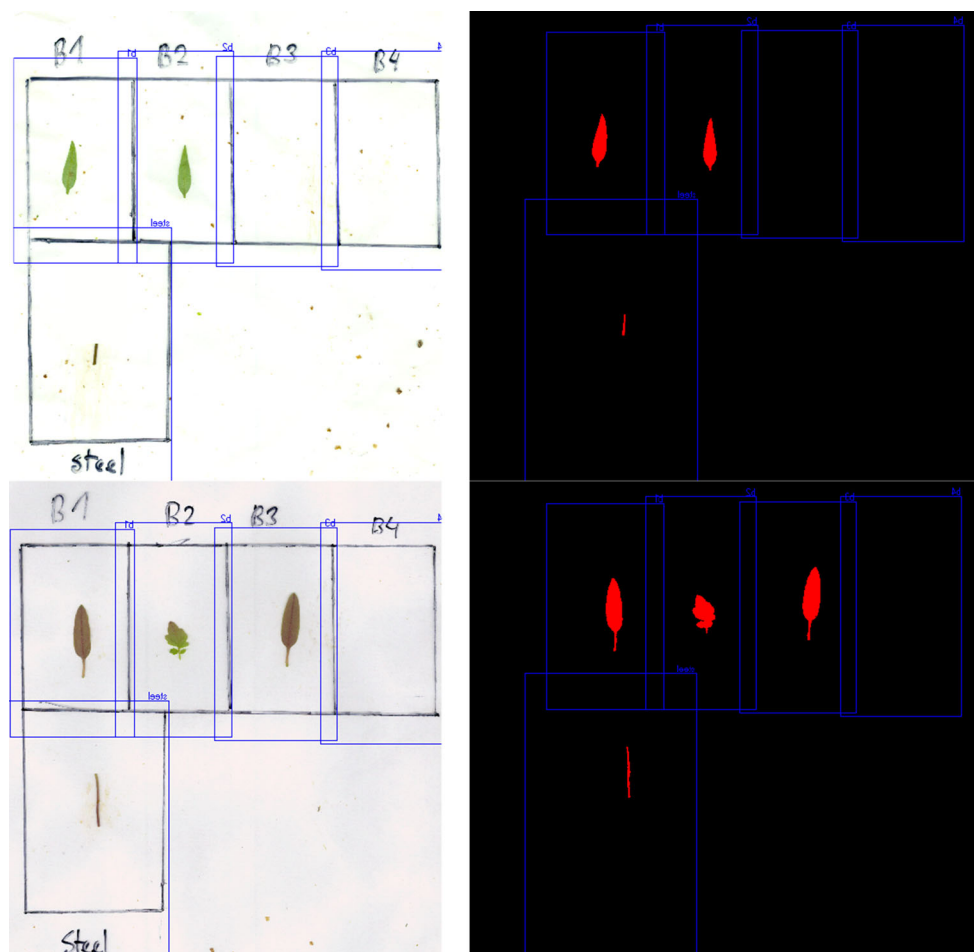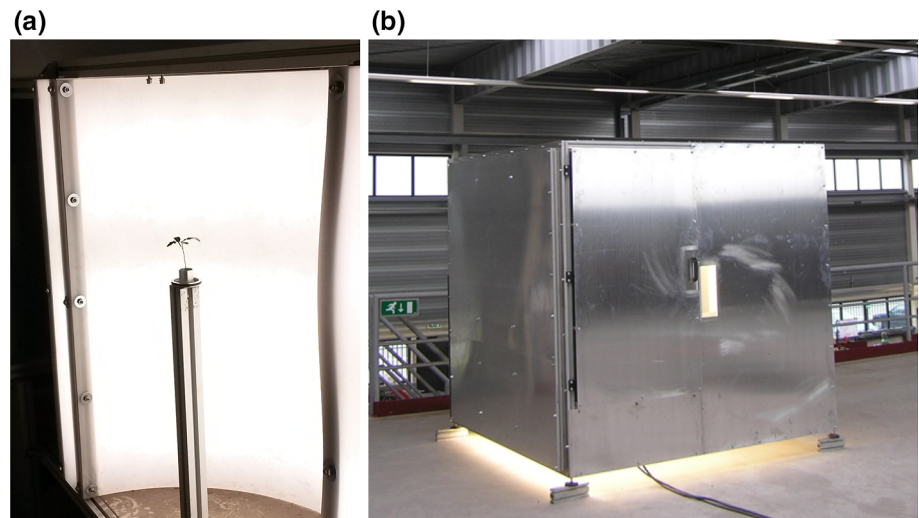

Fig. 9 Leafs were manually separated from the stem and positioned on a flatbed scanner (*left*); Basic image analysis methods (Labview IMAQ) were used to locate the individual parts and determine the features of interest (*right*)

were flattened as much as possible without tearing them. The resulting scans are used to calculate the ground truth measurements. A dedicated program was used to perform these measurements in a semi-automatic way. It is based on the

detection of individual plant parts in a set of user-definable ROIs. Plant parts are segmented from the background based on colour thresholding. The resulting binary objects are analysed using basic image processing tools as available in

Table 2 The overview of the measured plant features and their calculation methods

Feature	Calculation method
Length of the leaf	Longest Feret diameter
Width of the leaf	Widest section perpendicular to the longest Feret diameter
Surface of the leaf	Calibrated pixel surface
Length of the stem	Longest Feret diameter
Width of the stem	Calibrated surface divided by the length

the Labview IMAQ library: the leaf length is determined by finding the longest Feret diameter; the width of the leaf is defined as the widest part perpendicular to this length axis (see Fig. 10). A similar procedure is applied to the stem. The known resolution of the scanner (300 dpi) was used to transform the measurements to real-world units (mm and mm^2).

An overview of measurements performed is described Table 2.

We interpret these measurements as ground truth, but one must be aware of the fact that these measurements are not perfect, due to physical and computational reasons. Some of the leafs are curved in 3D space in such a way that a 2D projection on the flatbed scanner is not possible without folding or tearing the leaf, which obviously influences the measurements. Although carefully optimized, the segmentation method and the methods for doing the measurements also contain noise. This will contribute to the relative spread in the measured 3D feature values and act as a limiting factor in the overall apparent accuracy of the system.

4 Results

4.1 Stem length

Stem length was measured using the procedure described in Sect. 2.4. The total stem length was derived from the 3D voxel data by tracking the stem voxels from the plug up to the stipule (attachment point) of the first leaf, following the curvature of the stem (see Fig. 6c). This results are shown in Fig. 11.

The estimations from our method correlate well with the ground truth measurements, with a correlation coefficient of 0.87. The root mean squared error of deviation (RMSED) is 4.3 mm and the average relative error is 21 %. Looking at the data, attention is drawn towards a series of 5 seedlings with estimated stem length equalling 0. In some cases, the automatic processing of the 3D data runs into a situation where no result is generated for the particular feature. For
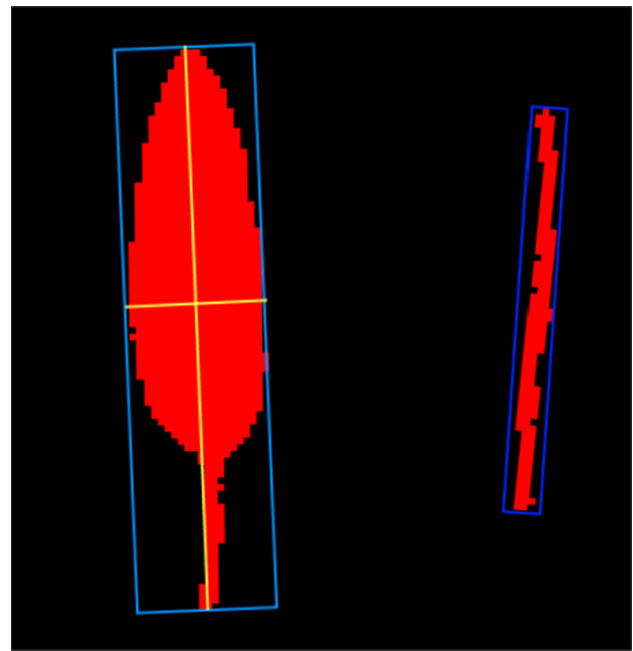

Fig. 10 Illustration of the measurement of the features using 2D scans of the individual plant organs. Leaf length is determined using the longest Feret diameter; leaf width is the widest part perpendicular to this (*left*). The length of the stem is determined using the length of the bounding box (*right*)

instance, in case of the encircled entry, we have the situation that the seedlings stem was not detected, as can be seen in Fig. 12. The stem is of unusual thickness, very short, and also bulging out on the lower side. This specific combination led obviously to a misclassification. In total, 530 out of 535 stems were segmented correctly by our method, which corresponds to 99 %.

The regression line has a slope of 0.79, which suggests that the length is systematically underestimated. In case of the stem length, this can be understood by looking closer to the way of working: in the given situation, there were some practical restrictions in the positioning of the lower end of the stem. A set of pre-set positions was used, always cutting of the lower end of the stem. In the future, we are looking to use an automatic separation algorithm for detaching the stem from the cup holding the seedling.

4.2 Leaf length

Leaf length was measured using the algorithm described in Sect. 2.3.1. The lengths were evaluated for all leafs individually; almost all seedlings have at least two leaves, a minority (64) had 3 leafs, and 3 seedlings had 4 leafs. Figures 13, 14 and 15 show the results of our 3D system compared to the 2D ground truth.

Our method performs well for leaf 1 and leaf 2 with correlations of, respectively, 0.91 and 0.90. The accuracy of the

Fig. 11 Stem length derived from the 3D data vs. ground truth from 2D flatbed scanner data. $N = 535$, RMSED = 4.3 mm, average relative error = 21 %, $R^2 = 0.75$; slope = 0.78, offset = 1.39
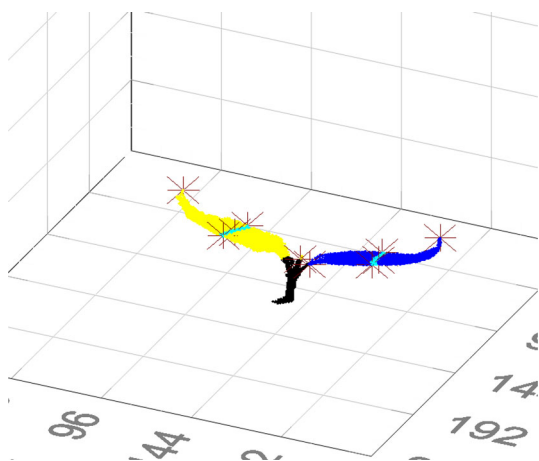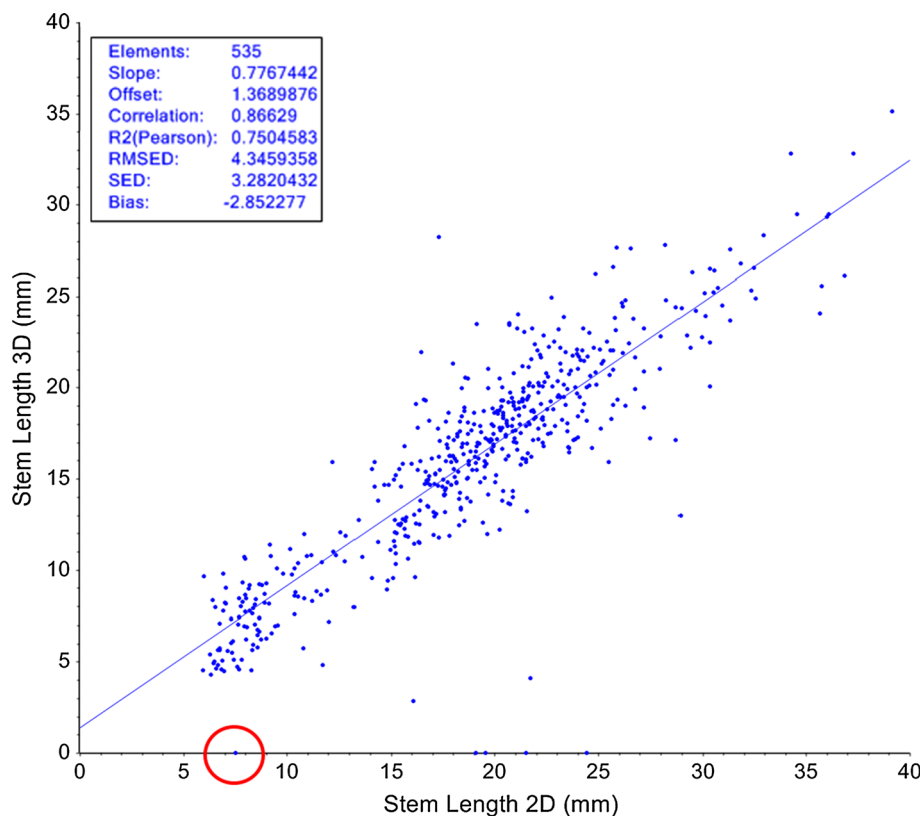

Fig. 12 The stem of the seedling is unclassified (*black*). The 'stem length' feature returns 0, 'no measurement'. These 'results' were not removed from the dataset, since we want to investigate applying the measurement system in a real-life phenotyping setting

leaf length estimations as measured by the RMSED is 4.3 and 3.6 mm for both leafs, and the average relative error is 18 and 16 %. Unfortunately, the results for the third leaf are not very good, as can be seen in Fig. 15 (correlation of 0.32, RMSED of 11, and average relative error of 55 %). Not only the spread is much larger, but there is a large number (16) of seedlings where the method gives no result (length = 0). Investigation shows that this is caused by a failure to correctly segment the leaf. The third leaf usually is the first true

leaf, which is more lobed than the cotyledons. In the first step of the segmentation method, the different lobes are often different segments. These are then merged in the second step. However, this sometimes fails, resulting in a small segment. If the length of such a segment is below a threshold of 3 mm <**3.1**>, the method returns 0.

4.3 Leaf width

<**2.3**> Leaf width was measured using the procedure described in Sect. 2.3.2. Results for leafs 1 and 2 are shown in Figs. 16 and 17:

The correlation of our method with the ground truth results is 0.83 for leaf 1 and 0.85 for leaf 2, with a RMSED of respectively 2.1 and 1.9 mm, and average relative error of 30 and 27 %. These errors are smaller than for the length of the leafs, but since the width of the leafs is about one-third of the length, the relative error is larger. This may not come as a surprise; the number of voxels in this direction is less, and hence, the accuracy.

4.4 Leaf area

Leaf area is based on the method described in Sect. 2.3.3. Results are shown in Figs. 18 and 19. The method has a correlation of 0.85, RMSED of 27, and average relative error of 22 % for the first leaf and correlation of 0.88, RMSED of 24, and average relative error of 19 % for the second leaf.

Fig. 13 Length of leaf 1
derived from the 3D data vs.
ground truth from 2D flatbed
scanner data. $N = 541$, RMSED
= 4.3 mm, average relative error
= 18 % $R^2 = 0.83$; slope = 0.89,
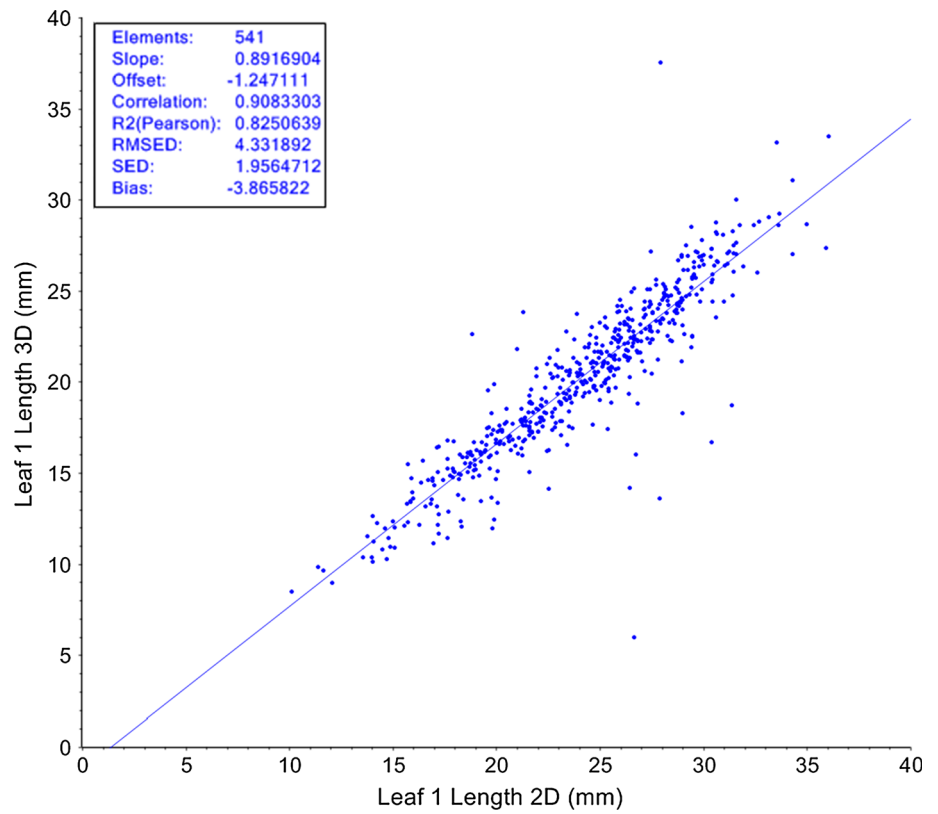offset = -1.25

Elements:	541
Slope:	0.8916904
Offset:	-1.247111
Correlation:	0.9083303
R2(Pearson):	0.8250639
RMSED:	4.331892
SED:	1.9564712
Bias:	-3.865822

Fig. 14 Length of leaf 2
derived from the 3D data vs.
ground truth from 2D flatbed
scanner data. $N = 511$, RMSED
= 3.6 mm, average relative error
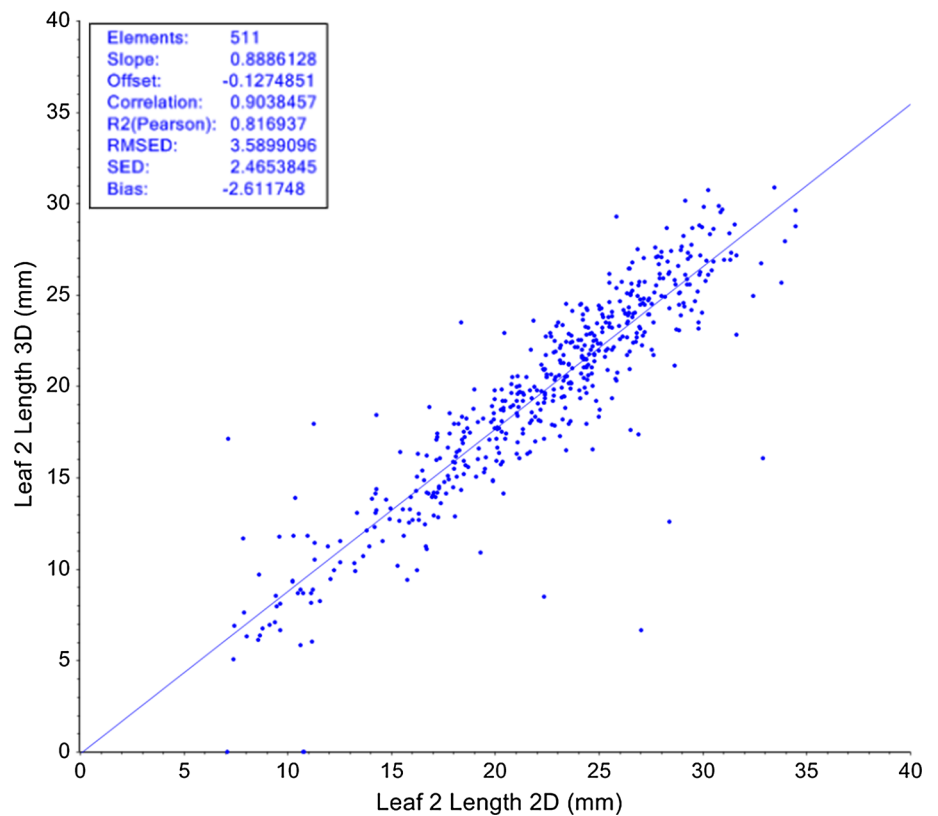= 16 %, $R^2 = 0.82$; slope =
0.89, offset = -0.13

Elements:	511
Slope:	0.8886128
Offset:	-0.1274851
Correlation:	0.9038457
R2(Pearson):	0.816937
RMSED:	3.5899096
SED:	2.4653845
Bias:	-2.611748

Fig. 15 Length of leaf 3 derived from the 3D data vs. ground truth from 2D flatbed scanner data. $N = 82$, RMSED = 11 mm, average relative error = 55 %, $R^2 = 0.10$. A number of cases returned no valid result for the leaf length of the third leaf. In some cases, this was due to incorrect segmentation of this leaf
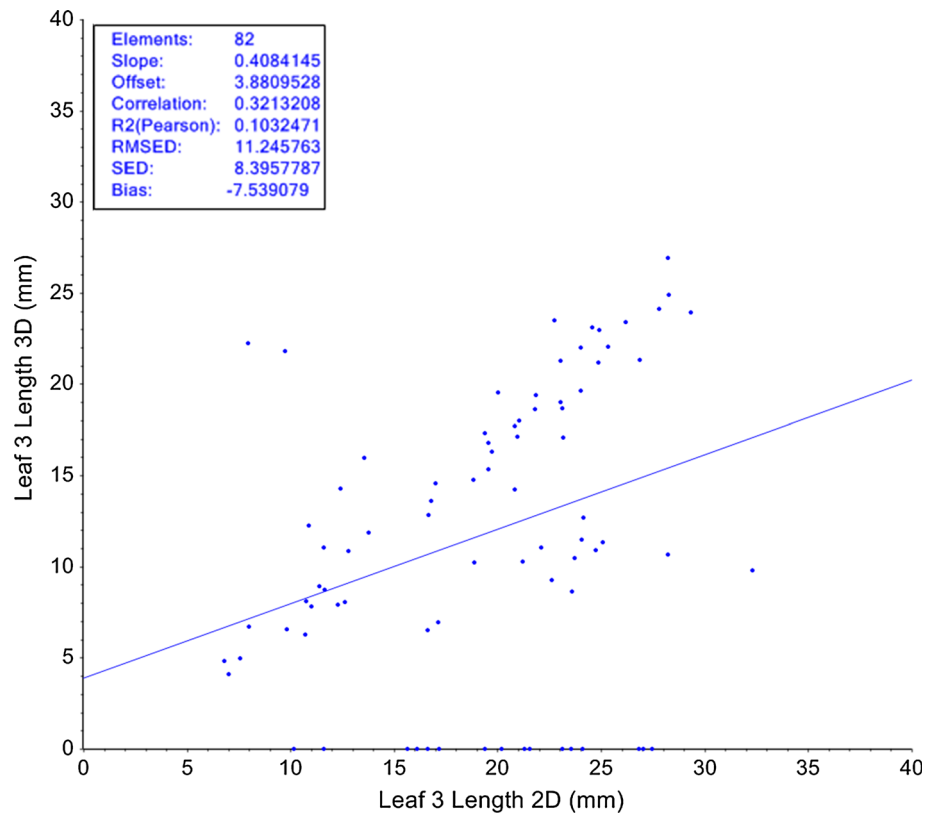
Elements:	82
Slope:	0.4084145
Offset:	3.8809528
Correlation:	0.3213208
R2(Pearson):	0.1032471
RMSED:	11.245763
SED:	8.3957787
Bias:	-7.539079

Fig. 16 Width of leaf 1 derived from the 3D data vs. ground truth from 2D flatbed scanner data. $N = 541$, RMSE = 2.1 mm, average relative error = 30 %, $R^2 = 0.69$; slope = 0.80, offset = -0.46
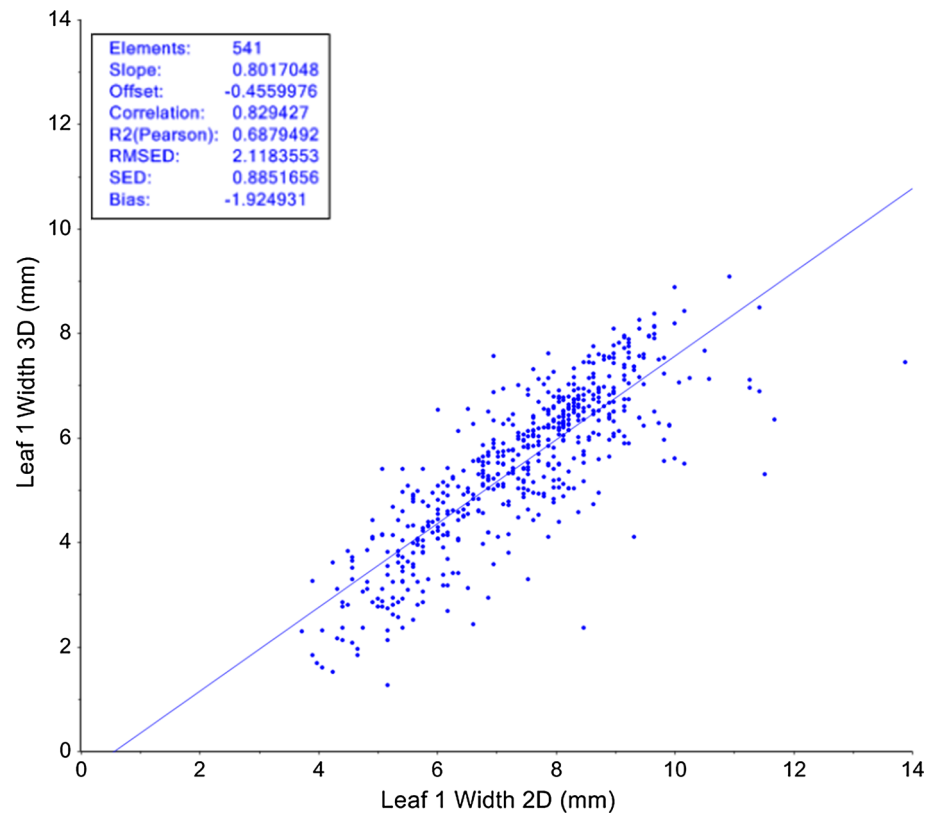
Elements:	541
Slope:	0.8017048
Offset:	-0.4559976
Correlation:	0.829427
R2(Pearson):	0.6879492
RMSED:	2.1183553
SED:	0.8851656
Bias:	-1.924931

Fig. 17 Width of leaf 2 derived from the 3D data vs. ground truth from 2D flatbed scanner data. $N = 541$, RMSE = 1.9 mm, average relative error = 27 %, $R^2 = 0.73$; slope = 0.96, offset = -1.44
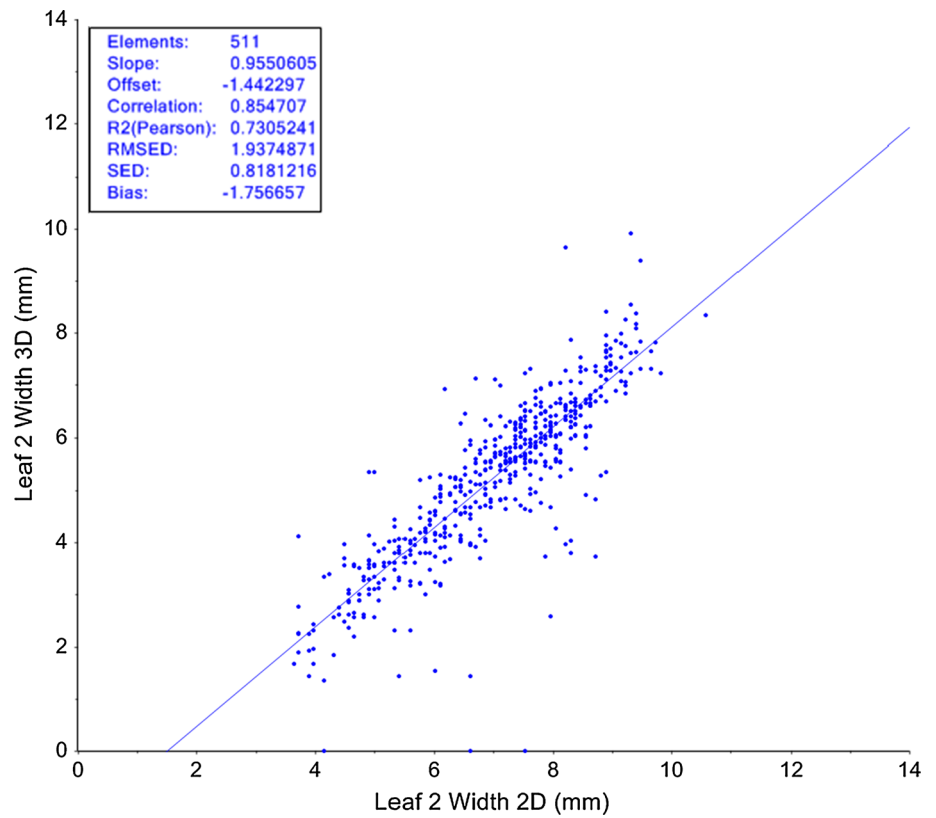
Elements:	511
Slope:	0.9550605
Offset:	-1.442297
Correlation:	0.854707
R2(Pearson):	0.7305241
RMSED:	1.9374871
SED:	0.8181216
Bias:	-1.756657

Fig. 18 Area of leaf 1 derived from the 3D data vs. ground truth from 2D flatbed scanner data. $N = 541$, RMSE = 28 mm^2, average relative error = 22 % $R^2 = 0.72$; slope = 1.05, offset = -14.8
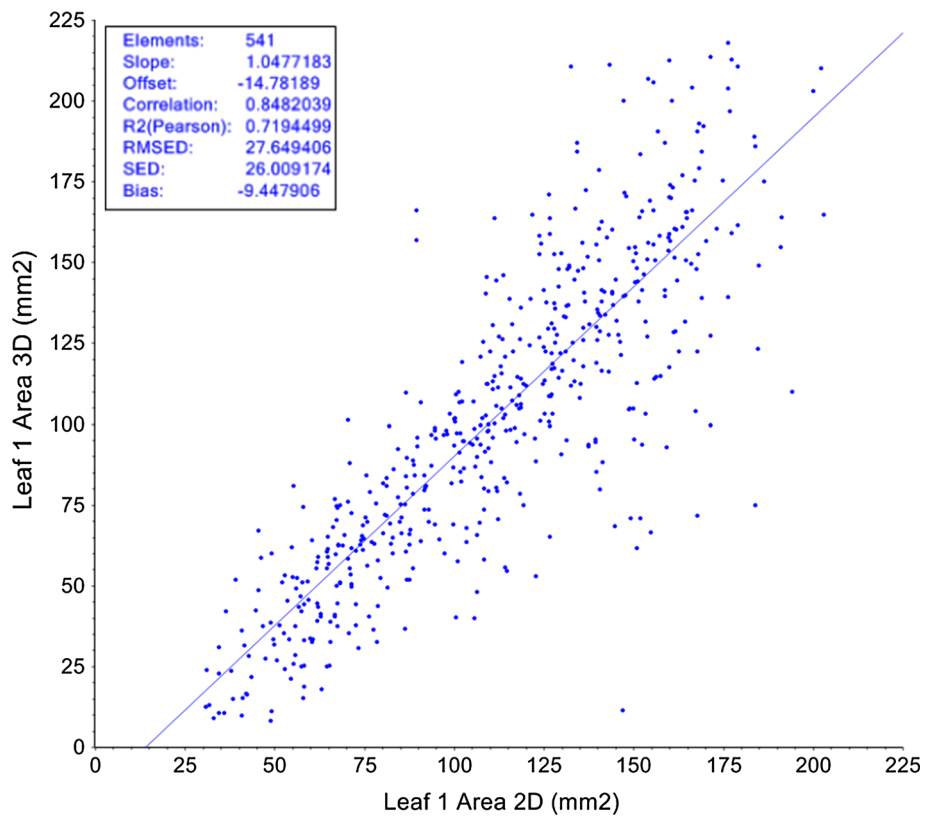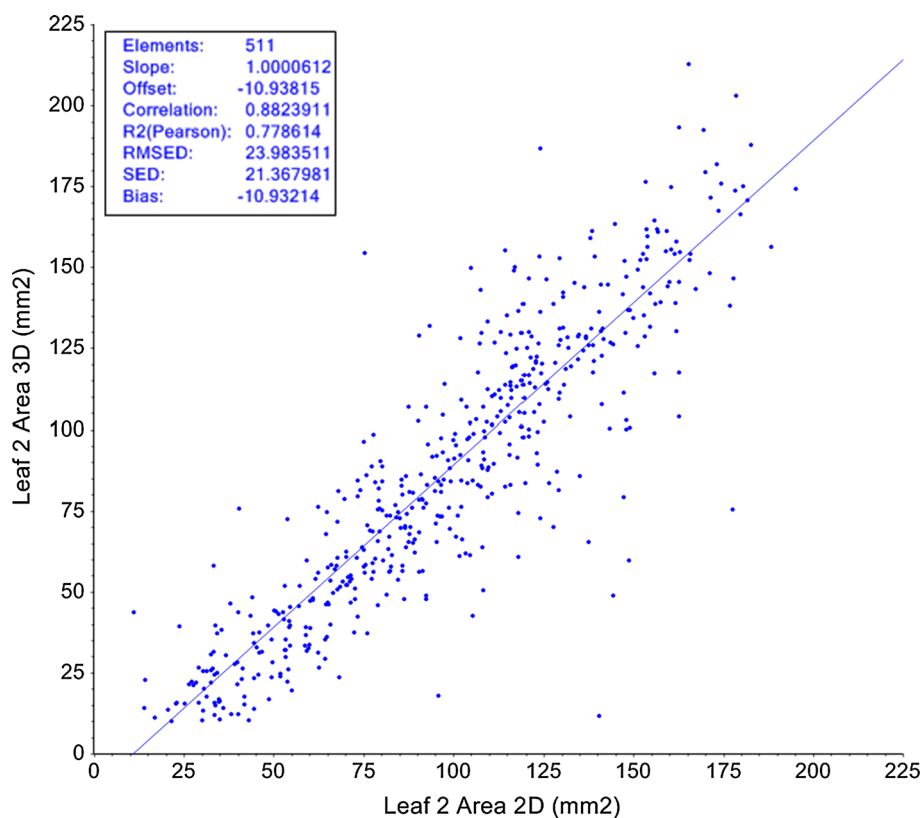
Elements:	541
Slope:	1.0477183
Offset:	-14.78189
Correlation:	0.8482039
R2(Pearson):	0.7194499
RMSED:	27.649406
SED:	26.009174
Bias:	-9.447906

Fig. 19 Area of leaf 2 derived from the 3D data vs. ground truth from 2D flatbed scanner data. $N = 511$, RMSE = 24 mm^2, average relative error = 19 %, $R^2 = 0.78$; slope = 1.00, offset = -10.9

5 Discussion

In this paper, we presented methods to measure specific features of a seedling based on a 3D reconstruction of the plant. Implemented features include length, width and area of individual leafs, and length of the stem. These are some of the basis phenotypic properties of a plant. The proposed method is optimized for speed, to be used in high-throughput phenotyping systems. The 3D reconstruction of the plant is created using a shape-from-silhouette method. The resulting 3D plant model is segmented in stem and leaf parts, from which stem and leaf parameters are calculated. Stem length, leaf length, and area can be determined with an average relative error of approx. 20 %; the error measuring the width of the leafs is close to 30 %.

A clear benefit of the proposed method is its fast processing time. Other existing methods for plant phenotyping using full 3D reconstruction from camera image are based on more complex 3D reconstruction methods, such as multiview stereo [27], structure-from-motion [28], space carving [25] or using laser-line scan devices [26]. These methods are more accurate, but with the cost of very long processing times, making them not applicable for high-throughput phenotyping (<1 s). We put much effort in optimizing the speed of our 3D reconstruction algorithm, which now runs in 20–60 ms, depending on the complexity of the plant, with a voxel space of $240 \times 240 \times 300$ voxels and using 10 cameras with a resolution of 1280×960 on a PC with a 3.2 GHz i7

processor. Other steps in our algorithm (stem/leaf segmentation and the calculation of stem and leaf parameters) have not yet been optimized for speed. Calculation of stem and leaf parameters has low complexity and is fast (25–60 ms), but the implementation of the stem/leaf segmentation is slow, taking 100–200 ms with incidental peaks up to 1000 ms. Depending on the structure of the leafs, the method's first step can result in many small segments, which then need to be merged in a second step. Moreover, the method has been implemented inefficiently in Labview. We are confident that we can greatly improve processing time with a more efficient implementation and adapting the method to better deal with irregularly shaped leafs.

To facilitate high-throughput plant phenotyping, we developed fast, relatively simple methods that approximate the relevant stem and leaf parameters. The leaf length is approximated by two line segments, which is an underestimation of the true length measured along the curved surface. This is reflected in Figs. 13 and 14 where the slope of the correlation between the 2D and 3D measurements is less than 1. Similarly, the leaf width is estimated by taking the Euclidian distance between two points instead of following the actual surface. Again, this results in an underestimation which can be noticed in Figs. 16 and 17. Finally, the method for measuring the leaf area assumes a mainly horizontally (or vertically) oriented leaf with an underestimation of the area when the leaf is tilted. Despite these limitations, we believe that the results of our system are very promising.

The camera setup has been developed such that from each camera perspective, the plant has a clear illuminated background, which assures robust foreground/background segmentation. We do not use a top view, avoiding the typical image-segmentation issues caused by pot, soil, moss and other clutter visible in the images.

We hope to further improve the results in the near future. First, the results show some limitations in the process of stem and leaf segmentation. In case of very fine structures, the method does not always find the optimal position of stem–leaf transition. We expect to improve this by applying a skeletisation algorithm to the seedling, which will give us a more accurate description of the plant connectivity and topology. This will not only allow us to do a better segmentation, but also to improve the measurements of, for instance, stem length. Next, to improve leaf measurements, we will describe the leaf by a fitted surface curved in 3D. Doing so will not only increase the accuracy of the leaf measurements, but also reduce the sensitivity to noise. It will furthermore add the possibility of measuring new features, like the shape of the leaf's contour. Finally, we are working on a way of modelling more complex plants. One of the major challenges is the way to deal with apparent loops in the model, i.e. situations where it is not clear whether plant parts are really connected, or just touching each other. To solve such ambiguities, we will incorporate expert knowledge in the system, so that the most likely representation can be found.

For phenotyping purposes, improving accuracy is not the only important issue. The ability to measure characteristics of individual plant parts is also a promising advantage of a 3D-based system. We hope that the system presented here can contribute to this demand.

References

1. Furbank, R.T., Tester, M.: Phenomics—technologies to relieve the phenotyping bottleneck. Trends Plant Sci. **16**(12), 635–644 (2011)
2. Kolukisaoglu, Ü., Thurow, K.: Future and frontiers of automated screening in plant sciences. Plant Sci. **178**(6), 476–484 (2010)
3. Fiorani, F., Schurr, U.: Future scenarios for plant phenotyping. Annu. Rev. Plant Biol. **64**, 267–291 (2013)
4. Li, L., Zhang, Q., Huang, D.: A review of imaging techniques for plant phenotyping. Sensors **14**(11), 20078–20111 (2014)
5. Rousseau, D., et al.: Multiscale imaging of plants: current approaches and challenges. Plant Methods **11**(1), 6 (2015)
6. Bergstrasser, S., et al.: HyperART: non-invasive quantification of leaf traits using hyperspectral absorption-reflectance-transmittance imaging. Plant Methods **11**(1), 1 (2015)
7. Biskup, B., et al.: A stereo imaging system for measuring structural parameters of plant canopies. Plant Cell Environ. **30**(10), 1299–1308 (2007)
8. van der Heijden, G., et al.: SPICY: towards automated phenotyping of large pepper plants in the greenhouse. Funct. Plant Biol. **39**, 870–877 (2012)
9. Eberius, M., Lima-Guerra, J.: High-throughput plant phenotyping—data acquisition, transformation, and analysis. In: Bioinformatics: Tools and Applications, pp. 259–278. Springer, New York (2009)
10. Golzarian, M., et al.: Accurate inference of shoot biomass from high-throughput images of cereal plants. Plant Methods **7**(1), 2 (2011)
11. Imaging robots. http://www.psb.ugent.be/infrastructure/391-image-robots
12. Granier, C., et al.: PHENOPSIS, an automated platform for reproducible phenotyping of plant responses to soil water deficit in *Arabidopsis thaliana* permitted the identification of an accession with low sensitivity to soil water deficit. New Phytol. **169**(3), 623–635 (2006)
13. Hartmann, A., et al.: HTPheno: an image analysis pipeline for high-throughput plant phenotyping. BMC Bioinf. **12**(148), 1–9 (2011)
14. Jansen, M., et al.: Simultaneous phenotyping of leaf growth and chlorophyll fluorescence via GROWSCREEN FLUORO allows detection of stress tolerance in *Arabidopsis thaliana* and other rosette plants. Funct. Plant Biol. **36**(11), 902–914 (2009)
15. Walter, A., et al.: Dynamics of seedling growth acclimation towards altered light conditions can be quantified via GROWSCREEN: a setup and procedure designed for rapid optical phenotyping of different plant species. New Phytol. **174**(2), 447–455 (2007)
16. Tisné, S., et al.: Phenoscope: an automated large-scale phenotyping platform offering high spatial homogeneity. Plant J. **74**(3), 534–44 (2013)
17. Subramanian, R., Spalding, E., Ferrier, N.: A high throughput robot system for machine vision based plant phenotype studies. Mach. Vision Appl. **24**(3), 619–636 (2013)
18. Minervini, M., Abdelsamea, M.M., Tsaftaris, S.A.: Image-based plant phenotyping with incremental learning and active contours. Ecol. Inf. **23** (2014)
19. Pereyra-Irujo, G.A., et al.: GlyPh: a low-cost platform for phenotyping plant growth and water use. Funct. Plant Biol. **39**, 905–913 (2012)
20. Dornbusch, T., et al.: Measuring the diurnal pattern of leaf hyponasty and growth in Arabidopsis—a novel phenotyping approach using laser scanning. Funct. Plant Biol. **39**, 860–869 (2012)
21. Alenya, G., Dellen, B., Torras, C.: 3D modelling of leaves from color and ToF data for robotized plant measuring. In: 2011 IEEE International Conference on Robotics and Automation (ICRA) (2011)
22. Chéné, Y., et al.: On the use of depth camera for 3D phenotyping of entire plants. Comput. Electron. Agric. **82**, 122–127 (2012)
23. Klose, R., Penlington, J., Ruckelshausen, A.: Usability study of 3D time-of-flight cameras for automatic plant phenotyping. Bornimer Agrartechnische Berichte **69**, 93–105 (2009)
24. Andersen, H.J., Reng, L., Kirk, K.: Geometric plant properties by relaxed stereo vision using simulated annealing. Comput. Electron. Agric. **49**, 219–232 (2005)

25. Paproki, A., et al.: A novel mesh processing based technique for 3D plant analysis. BMC Plant Biol. **12**(63) (2012)
26. Paulus, S., et al.: Low-cost 3D systems: suitable tools for plant phenotyping. Sensors **14**, 3001–3018 (2014)
27. Pound, M.P., et al.: Automated recovery of three-dimensional models of plant shoots from multiple color images. Plant Physiol. **166**(4), 1688–1698 (2014)
28. Quan, L., et al.: Image-based plant modeling. ACM Trans. Graph. **25**(3), 599–604 (2006)
29. Baumberg, A., Lyons, A., Taylor, R.: 3D S.O.M.—a commercial software solution to 3D scanning. Graph. Models **67**(6), 476–495 (2005)
30. Furukawa, Y., Ponce, J.: Accurate, dense, and robust multiview stereopsis. IEEE Trans. Pattern Anal. Mach. Intell. **32**, 1362–1376 (2010)
31. Biskup, B., et al.: Diel growth cycle of isolated leaf discs analyzed with a novel, high-throughput three-dimensional imaging method is identical to that of intact leaves. Plant Physiol. **149**(3), 1452–1461 (2009)
32. Baumgart, B.G.: Geometric Modeling for Computer Vision. PhD thesis, Stanford (1974)
33. Kim, Y., Aggarwal, J.: Rectangular parallelepiped coding: a volumetric representation of three dimensional objects. IEEE J. Robot. Autom. **2**, pp. 127–134 (1986)
34. Martin, W., Aggarwal, J.: Volumetric descriptions of objects from multiple views. IEEE Trans. Pattern Anal. Mach. Intell. **5**(2), 150–174 (1983)
35. Szeliski, R.: Rapid octree construction from image sequences. Comput. Vision Graph. Image Process. Image Underst. **58**(1), 23–32 (1993)
36. Laurentini, A.: The visual hull: a new tool for contour-based image understanding. In: Proceedings of the Seventh Scandinavian Conference on Image Analysis (1991)
37. Laurentini, A.: The visual hull concept for Silhouette-based image understanding. IEEE Trans. Pattern Anal. Mach. Intell. **16**(2), 150–162 (1994)
38. Koenderink, N.J.J.P., et al.: MARVIN: High speed 3D imaging for seedling classification. In: Seventh European Conference on Precision Agriculture, pp. 279 –286 (2009)
39. Kurillo, G., Zeyu, L., Bajcsy, R.: Wide-area external multi-camera calibration using vision graphs and virtual calibration object. In: IEEE Second ACM/IEEE International Conference on Distributed Smart Cameras: Stanford (2008)

Dr. Gert Kootstra is a researcher and scientific coordinator in Computer Vision and Robotics at the Wageningen UR. He received his PhD in Artificial Intelligence from the University of Groningen, The Netherlands, in 2009. During 2009–2012, he was a postdoctoral researcher at the Royal Institute of Technology (KTH) in Stockholm, focussing on machine vision for robotic grasping. His current research deals with visual quality inspection and robotic handling in the agro-food domain.

Sanja Damjanovic received the PhD degree in computer vision from University of Twente, The Netherlands, in 2012. In 2001 and 2005, respectively, she received the Dipl.-Ing. and Master of Science in electrical engineering from the University of Belgrade, Serbia. From 2001 to 2005, she was a researcher in the Department of Telecommunications, Mihajlo Pupin Institute, Belgrade, Serbia. From October 2005 to August 2006, she was a DAAD research fellow at Digital Signal Processing Group (DISPO), Ruhr-University, Bochum. From September 2006 to August 2011, she was a researcher at Signals and Systems Group, University of Twente, working on the research topic of stereo matching. From 2011 to 2015, she was a scientist at Computer Vision group at Food and Biobased Research at Wageningen UR, The Netherlands, in the fields of computer vision and signal processing applied to agrotechnology and consumer science. Sanja is currently a lecturer and researcher in the field of computer vision and signal processing with Saxion University of Applied Sciences, Enschede, The Netherlands.

Franck Golbach is a researcher and project manager at Wageningen University and Research Centre. He holds a Master of Applied Physics degree from Utrecht University. He specialized in computer vision within the agricultural domain. For some years, he focused on the analysis of 3D imaging data, aiming at quality control, high-speed sorting, and developing 3D vision solutions for phenotyping purposes. Current interests involve accurate and calibrated color measurements on curved surfaces.

Gerwoud Otten MSc., studied electrical engineering at Twente University from 1993 to 1997. From 1998 he is a researcher in Computer Vision and Robotics at Wageningen UR. Gerwoud's current research deals with 3D quality inspection of seedlings and robotic plant manipulation.

Rick van de Zedde is a senior researcher/ business developer for Computer Vision at the Wageningen UR Food & Biobased Research Institute in the Netherlands, where he has worked since 2004. His background is in Artificial Intelligence with a focus on imaging and robotics. In 2002 he graduated with an MSc in Artificial Intelligence from the University of Groningen. Since 2006 he has been a coordinator of GreenVision, the centre of expertise for image processing in the agrofood industry (http://greenvision.wur.nl).