Cost-effective, high-throughput phenotyping system for 3D reconstruction of fruit form

Mitchell J. Feldmann[1] **and Amy Tabb**[2]

[1]Department of Plant Sciences, University of California, Davis, Davis, California, United States of America
[2]USDA-ARS-AFRS, Kearneysville, West Virginia, United States of America

Reliable phenotyping methods that are simple to operate and inexpensive to deploy are critical for studying quantitative traits in plants. Traditional fruit shape phenotyping relies on human raters or 2D analyses to assess form, e.g., size and shape. Systems for 3D imaging using multi-view stereo have been implemented, but frequently rely on commercial software and/or specialized hardware, which can lead to limitations in accessibility and scalability. We present a complete system constructed of consumer-grade components for capturing, calibrating, and reconstructing the 3D form of small-to-moderate sized fruits and tubers. Data acquisition and image capture sessions are 9 seconds to capture 60 images. The initial prototype cost was $1600 USD. We measured accuracy by comparing reconstructed models of 3D printed ground truth objects to the original digital files of those same ground truth objects. The R^2 between length of the primary, secondary, and tertiary axes, volume, and surface area of the ground-truth object and the reconstructed models was > 0.97 and root-mean square error (RMSE) was $<$3mm for objects without locally concave regions. Measurements from 1mm and 2mm resolution reconstructions were consistent ($R^2 > 0.99$). Qualitative assessments were performed on 48 fruit and tubers, including 18 strawberries, 12 potatoes, 5 grapes, 7 peppers, and 4 Bosch and 2 red Anjou pears. Our proposed phenotyping system is fast, relatively low cost, and has demonstrated accuracy for certain shape classes, and could be used for the 3D analysis of fruit form.

Image Based Phenotyping | 3D Fruit Phenotyping | High-throughput Phenotyping | 3D Model Reconstruction | Camera Calibration | Computer Vision

Correspondence: ***mjfeldmann@ucdavis.edu***

Introduction

Fruit appearance is a key trait for many crops and can condition market viability of fruit products and the success of cultivars (1–3). Taken together, the shape and color, or appearance, of fresh fruit are often associated with quality as they reveal condition and impact consumer perception of taste (1). Fruit shape is a heritable, complex trait that is difficult to assess due to the complex nature of data acquisition, which can be both time consuming and computationally laborious (4–11). Fruit shape in agricultural studies have primarily been assessed subjectively, placing fruit into qualitative bins ranging from 'deformed' to 'uniform' or by using 2D geometric morphometrics (4, 7, 12–19).

Publicly available methods for 2D phenotyping of plants and plant organs have increased in recent decades to support high quality analysis of leaves, roots, shoots, stems, tubers, and fruits (6, 20–27). Computer vision has shown great potential to quantify external fruit quality and 2D imaging has been successfully implemented to measure the shape and size of fruits such as strawberries (4, 12), apples (5), carrot (6, 14), mangoes (28), and many others. More recently, methodologies for 3D reconstruction of plant organs have been developed with approaches that vary in speed, scale, cost, and accuracy; including laser scanners, x-ray computed tomography, and reconstruction from sequences of 2D images from digital cameras (8, 29–41). Methods that rely on sequences of 2D images are numerous and variable with their own complexities and nuances that provide different strengths and weaknesses(8, 27, 37, 40–44).

Modern technologies and analyses can be used to assess these physical characteristics and ultimately provide researchers with the tools necessary to support genetic inquiries and biological discoveries, expand what is known about modern germplasm, and enhance breeding practices in fruit and vegetable crops (4, 6, 8, 45–52). Multivariate and spatial statistics can be used to determine parameters that identify and quantify fruit defects (53), differentiate between marketable and non-marketable fruit (12, 50), and understand fruit phenotypes that impact markets requiring long shelf-life and sustained fruit quality through harvesting, handling, and shipping.

This paper describes a rapid (9 s), low-cost ($1,600), turntable-type system for 3D reconstruction of fruit and tubers. Fruit rotates on an automated pedestal while a remote-controlled digital camera acquires images, as shown in Figure 1. We use a multi-camera calibration method (54) to compute the calibration parameters of the camera at every time step. Fruit are segmented from non-fruit regions in the images. Finally, a reconstruction method using silhouettes as features (55) reconstructs the fruit or vegetable shape using the calibration and silhouette information (Figure 2).

Contributions. Our contributions to the state-of-the-art in fruit phenotyping and estimating 3D reconstructions are a high-throughput (9 second data acquisition), modular reconstruction system that can be used in lab or field settings (on a table) with high accuracy for objects that do not have local concavities. Our work is most similar to that of (40) and (8). In (40), a turntable system is used and the cameras rotate around the target object, rice inflorescences. Relative camera calibration parameters are estimated by detecting features and estimating matches from color checkerboard pages and a structure from motion approach generates point clouds. This approach works well for the target crop, but has an unknown scale factor that is not solved for. Consequently, the physical units such as a mm or cm are unknown, and comparison with another system of a different size may be difficult. In our system, we calibrate directly from patterns in the scene, so the physical size of the sample is estimated. The work of (8) uses a turntable system in the configuration that we do, where the fruit is rotating and the camera is stationary, to acquire images for 3D reconstruction. They use a commercial software package to reconstruct point clouds, and then process point clouds to extract fruit features. Our work differs from both of these in that we use silhouettes and directly use voxel representations instead of point clouds.

Core Ideas

- A low-cost 3D fruit phenotyping system is presented.
- Image capture using the proposed approach lasts for only 9 seconds.
- Accuracy is measured against 3D printed ground-truth objects.
- Camera calibration, background segmentation, and reconstruction does not rely on commercial software.
- An RMSE less that 3mm was obtained for ground truth objects without locally concave regions.

Materials and Methods

The 3D phenotyping system consists of multiple parts: the physical system for data acquisition and the algorithms for reconstructing shape from that data. Briefly, one or multiple digital cameras are mounted on an aluminum frame and remotely triggered at a frame rate of 7 frames per second (FPS) for 9 seconds while a stepper motor controlled by a microcontroller rotates an object. Captured images are calibrated using CALICO, a multi-camera calibration method (54) that relies on a combination of arUco and chArUco markers (56–58). The fruit foreground is segmented from the non-fruit background in each calibrated image. Segmented silhouettes of fruit are then used as features to reconstruct 3D models.

Hardware. The physical system is composed of an aluminium frame, digital cameras and camera control units, a USB barcode scanner, a microcontroller, and a microcomputer.

Frame. The frame's design is an inverted "T" shape structure with a 1.22m (4ft) horizontal base and a 1.22m vertical arm (Figure 1). The main structure is composed entirely of 80/20 t-slotted aluminum. We chose this material because it is lightweight, strong, and inexpensive. The t-slot design and the availability of different fasteners provides rigidity while remaining modular. The arm is connected to the base using a side-mount and hand brake. The side-mount and hand-brake combination means that the vertical arm can be positioned anywhere along the length of the horizontal base. The camera is mounted to the vertical arm using the same side-mount and hand brakes, again allowing it to be positioned continuously along the vertical span of the arm.

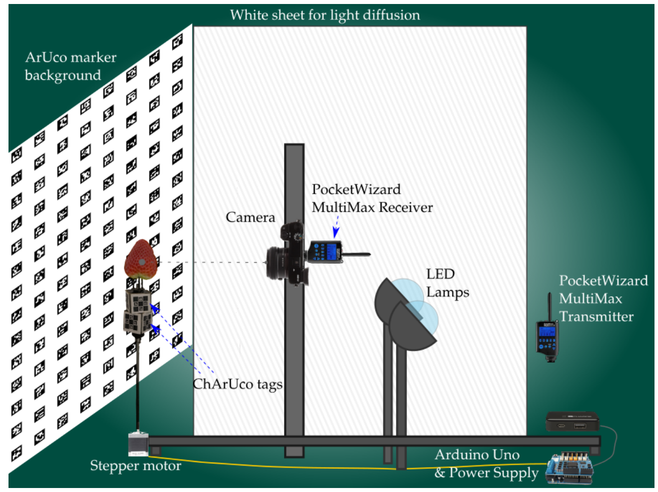

Fig. 1. Imaging system hardware. From left to right: the arUco marker backdrop; a stepper motor with a metal pedestal, chArUco tagged cubes, and a target object, a strawberry; the 80/20 t-slotted aluminum inverted T ($\perp$) frame; a digital camera is mounted on the vertical limb of the frame and attached to a PocketWizard MultiMax II radio transceiver; reverse facing LED light sources; Arduino microcontroller connected to stepper motor and power supply. **Best viewed in color**.

Cameras and controllers. We used one Sony α6000 mirrorless digital camera for this project. The camera was set to medium speed continuous image capture (7 frames per second), manual focus, and aperture priority mode with the aperture set to f/8. We controlled the camera with a PocketWizard MultiMax II transceiver unit. These units attach to the camera's multi-port and digitally control the camera's shutter button and can be programmed to "hold" the shutter button to allow for variable duration. We used 9 seconds of hold to match the rotation rate of our stepper motor. Two PocketWizard MultiMax II units are required: one unit to transmit a signal and one unit to receive a control signal for multiple cameras. With these, the camera is controlled from a single source which is triggered by the input of a barcode scanner.

Microcomputers and stepper motor. The data acquisition process consists of rotating the fruit on the pedestal and acquiring images of that fruit. To automate this process, we used a Raspberry Pi 3 microcomputer as well as an Arduino Uno Rev3 microcontroller. To rotate the objects on the pedestal, we used a Nema 17 stepper motor controlled using an Arduino Uno Rev3 and an Arduino Rev3 motor shield. The pedestal is a thin metal rod approximately 20cm in length and 5mm in diameter. The Nema 17 stepper motor has 200 steps per rotation (1.8º per step). The motor is programmed to take 1 step every 45ms, which is a full rotation every 9 seconds.

Lighting. We used 4 LED lamps to illuminate the scene. These lights are all directed away from the object towards a reflective white sheet to reduce the intensity of the light on the scene. This enabled us to dramatically reduce, and in some instances eliminate, the glare on the surface of more reflective objects such as strawberries. The lights chosen do not have any temperature control and are likely not ideal for color accurate measurements.

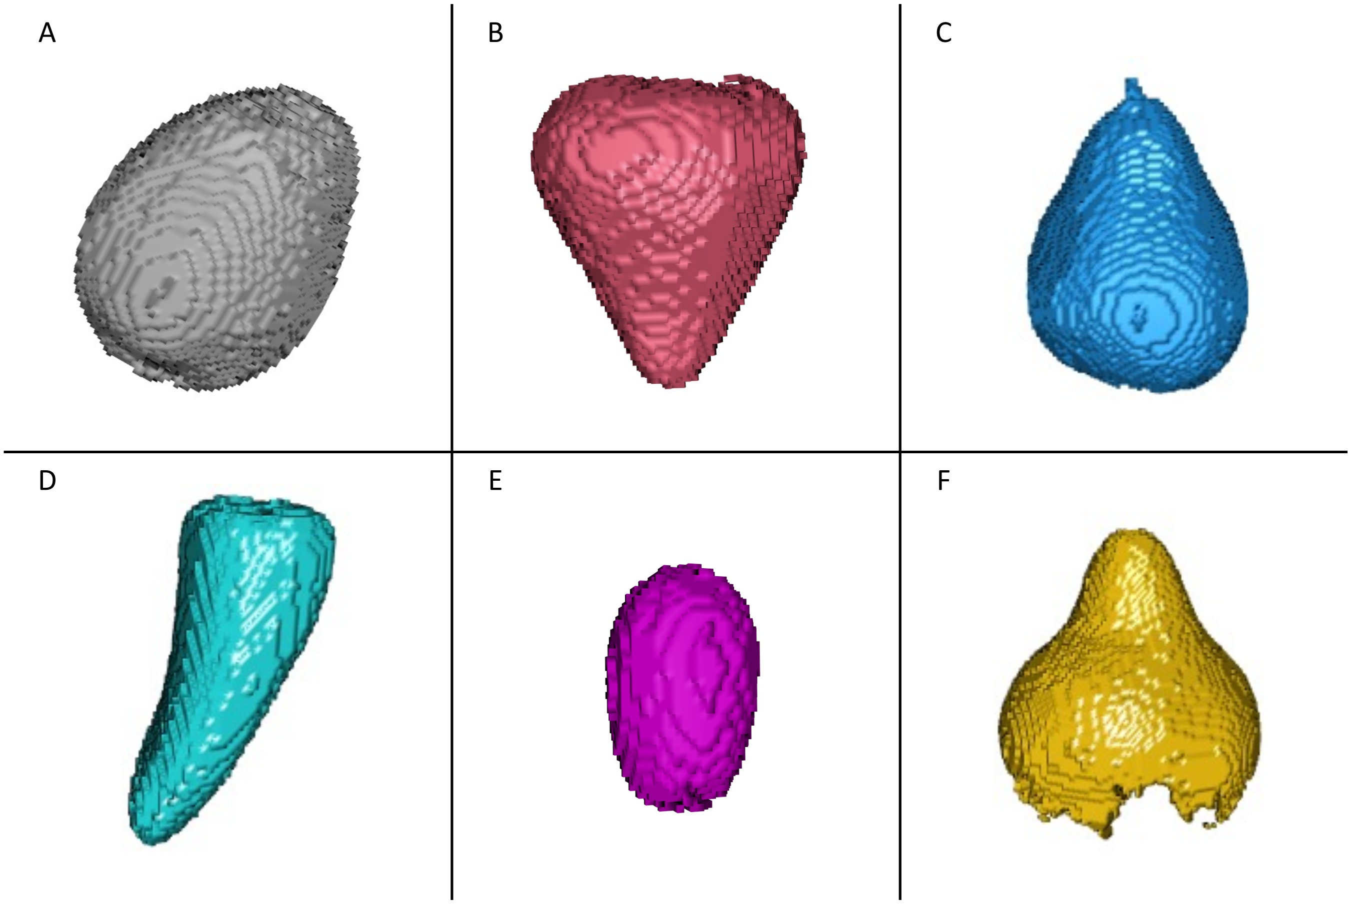

Fig. 2. Representative model reconstructions of the six types of fruit or tubers imaged in this study. The resulting 3D reconstructed models from the hardware and software presented in this study, including: (A) a baby yellow potato, (B) a strawberry, (C) a Bosch pear, (D) a sweet mini pepper, (E) a green table grape, and (F) an Anjou pear where the segmentation failed All models shown are at 1mm resolution. Images are not to scale. **Best viewed in color.**

Calibration Targets. Calibration is performed on image data that also contains the data for reconstruction. To accomplish one-step calibration and data acquisition, the workspace is prepared with calibration targets, which are shown in Figure 1. The fruit or tuber is mounted on the pedestal. A pair of offset cubes are mounted on the pedestal directly below the fruit or tuber (56). Each cube is 2.5cm $\times$ 2.5cm $\times$ 2.5cm with a small hole in the center for mounting onto the pedestal, and are rigidly attached.

On the cube faces without holes, chArUco markers are printed and attached to the cubes. The chArUco patterns are a 3×3 checkerboard with each square unit measuring 6.67mm$\times$6.67mm. The two cubes have eight faces with chArUco patterns on them, and multiple cube faces should be visible in any frame providing enough information such that the calibration method CALICO can compute camera poses.

We use a scene background, a 0.71m^2 (26in^2) aluminum panel, composed of arUco markers (57, 58). This type of background allows us to refine internal camera calibration parameters using the multi-camera calibration method, CALICO. Each arUco marker is 2.25cm^2 and adjacent markers are separated by 2.75cm of white space. Each image contains between 60-70 unobscured arUco markers, depending on the size of the object.

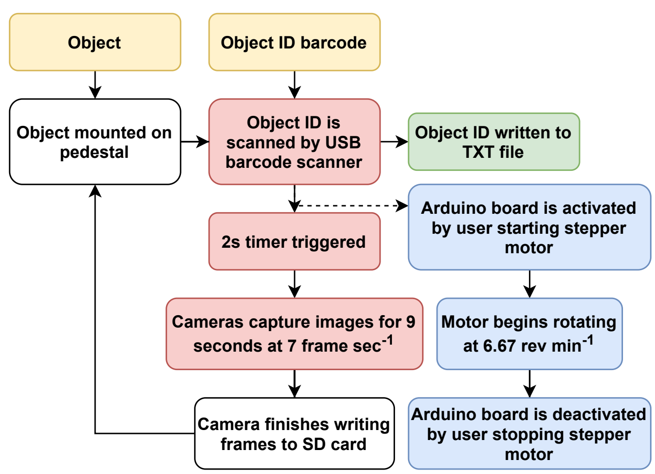

Fig. 3. Flow diagram of data acquisition strategy and steps. Input (yellow): A physical object with an associated barcode ID, e.g., QR code or data matrix. Staging (white): The object is placed on the pedestal. Camera Triggering (red): The object ID barcode is scanned using a USB barcode scanner attached to a Raspberri Pi, starting a 2 second timer before triggering the cameras for 9 s at 7 FPS. Intermediate output (green): The scanned barcode is written into a TXT file. Motor control (blue): During the 2 second timer, the Arduino board is activated by supplying power to the board initiating rotation of the object at 6.67 RPM (1 revolution every 9 seconds). Once the image capture is complete, the user deactivate the Arduino board, stopping the motor. Staging (white): The camera will need upwards of 15 seconds to finish writing the images to storage, during which time the next object can be staged on the pedestal, initiating the following session. During this process, the user is responsible for mounting the objects to the pedestal, scanning the barcode ID, and activating/deactivating the stepper motor. **Best viewed in color.**

Data Acquisition control. For every object, the Arduino board is first activated by the user, who starts the stepper motor's rotations. The cameras are triggered once the motor begins to rotate using a barcode scanner and a custom python script that sends a serial signal through the general-purpose input output (GPIO) of the Raspberry Pi to the transmitter unit. The cameras stop firing after 9 s and Arduino board is deactivated manually. Including swapping objects, the time for each sample is less than 25 s. In this time period, each camera captures approximately 60 frames during one complete rotation of the object. Following each session, the camera must be allowed to clear its on-board cache and to write the images to an SD card. Depending on the camera and the SD card, write speeds may vary. In our setup, this typically took about 10 s.

Reconstruction Pipeline. The 3D reconstruction pipeline consists of three stages: calibration, segmentation, and reconstruction. In this work, the stages consist of independent modules but they have been selected based on the assumptions of the reconstruction module.

We use a Shape from inconsistent Silhouette method from Tabb (2013) (55) that requires camera calibration information and uses silhouettes – or segmentations – of the target object to generate reconstructions. This reconstruction method tolerates calibrations with small camera calibration error and small image segmentation error. The methods used for calibration and segmentation are discussed in the Calibration and Segmentation sections, respectively.

Calibration. The reconstruction method depends on camera calibration. Camera calibration usually means the internal camera parameters as well as the external pose (rotation and translation) of cameras relative to a world coordinate system. In the context of this work, by 'computing the calibration,' we mean determining the internal camera parameters as well as the external pose of the camera with respect to the calibration object at each image acquisition.

We use an existing method for multiple-camera calibration, CALICO (54), to compute the desired calibration parameters. To use CALICO in this context, chArUco tags have to be rigidly attached to the pedestal and multiple tags visible at each time instant, which is why the physical system is prepared as in the Hardware section. Some datasets had significant error in the camera pose because not enough chArUco tags were detected in each frame, so we extended the 'rotation' option of CALICO to also detect and use the arUco corners within the chArUco boards.

A successful calibration estimates internal camera calibration parameters, selects one of the chArUco boards as a world coordinate system, and estimates the relative pose of the camera at each image acquisition with respect to that world coordinate system. For instance, see Figure 4, which shows an example of an input image, (Figure 4(A)) and the chArUco markers below the pedestal. The reconstructed chArUco board poses, camera poses at each image acquisition, and a reconstructed strawberry are shown in Figure 4(C).

Segmentation. The reconstruction method depends on segmented images, where the fruit, tuber, or ground truth object is separated from the image background as in Figures 4A-4B and 6A-6B. The backgrounds consist of arUco tags and the pedestal with chArUco tags. With this background, consisting of dark and light intensities, we took an approach of modelling the actual intensities of the calibration tags per image as a Gaussian Mixture Model with two components, and then used background subtraction to determine the location of target objects.

First, the arUco tags are located in the image. Then, the dark and light regions of each tag are identified using Otsu's segmentation algorithm (59). The dark and light regions of all of the tags are used to estimate six Gaussian distributions (masks are shown in Figure 6C and 6D): $\mathcal{N}_{d,r}(\mu, \sigma^2)$, $\mathcal{N}_{d,g}(\mu, \sigma^2)$, $\mathcal{N}_{d,b}(\mu, \sigma^2)$, the distributions representing the dark intensities for red, green, and blue channels, and the same for all three channels of the light intensities.

Each image pixel x is evaluated against the distributions as in Equation 3. We use a typical background subtraction technique in that we subtract the mean and compare with a threshold; here the threshold is a constant multiplied by the standard deviation. The user provides constants k_d and k_l, and from the distributions computes Boolean values y_d and y_l for each pixel x. The segmentation result of whether the pixel represents the background (0) or not (1) is stored in $z = y_l \wedge y_d$.

$$y_d = \bigvee_{ch \in \{r,g,b\}} |\mu_{d,ch} - x_{ch}| > k_d \sigma_{d,ch} \tag{1}$$

$$y_l = \bigvee_{ch \in \{r,g,b\}} |\mu_{l,ch} - x_{ch}| > k_l \sigma_{d,ch} \tag{2}$$

$$z = y_l \wedge y_d \tag{3}$$

In our experiments, $k_d = 2.0$ and $k_l = 2.5$ for all tests.

Reconstruction. We used a Shape from Inconsistent Silhouette (SfIS) method (55) for 3D reconstruction of the plant organs and ground truth objects. With camera calibration and segmentation or silhouette provided, SfIS is a voxel-based method that searches for a labeling of voxels as occupied or empty such that the voxels match the provided segmentations. The match does not need to be exact, so some small camera calibration and segmentation errors can be present.

A key feature of the SfIS method is that it will not reconstruct concavities in 3D space. As examples of these types of shapes, the tetrahedron, sphere, and F ground truth objects (Figure 7A-7C) can all be reconstructed because they do not contain concavities, while the 6-sided spherical die cannot (Figure 7D). The stem or calyx region of an apple is also an example of a locally concave region on a surface. The reason that the SfIS method is not able to reconstruct locally concave regions if because of its dependence on segmentations as features.

We use the extension to SfIS of hierarchical search described in (60); the user specifies an initial voxel size, finds a solution

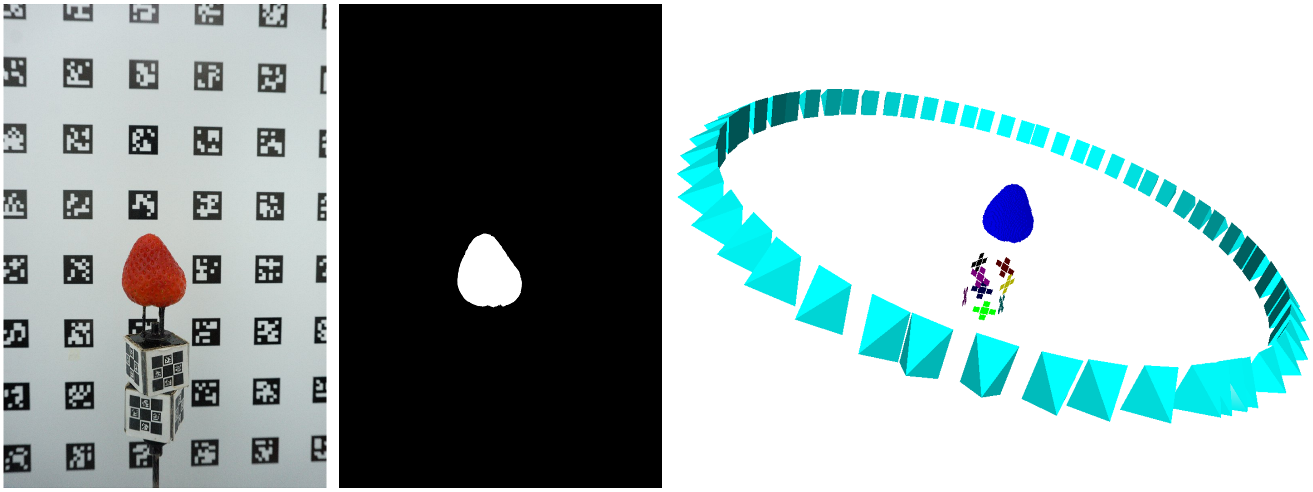

(A) Image from data acquisition, sample 53 **(B)** Segmented image **(C)** Reconstructed strawberry, calibration patterns (+ shaped forms), and estimated camera poses.

Fig. 4. An example of successful camera calibration results. Overview of image data and results. Images of fruit and calibration objects are captured while the pedestal rotates (a). Each image is segmented (b). The calibration and segmentation information is used to reconstruct the fruit shape, 1 mm voxel resolution shown in (c). A camera pose is represented as a pyramid, where the camera center is the tip of pyramid. **Best viewed in color.**

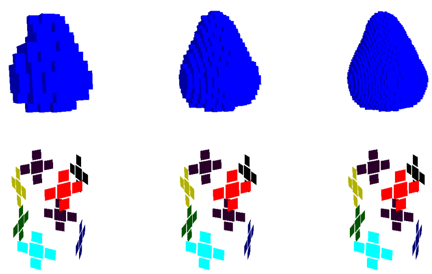

(A) 4 mm **(B)** 2 mm **(C)** 1 mm

Fig. 5. Reconstructions of a strawberry at 3 different resolutions. Strawberry reconstruction from Figure 4 during the hierarchical reconstruction process, with estimated calibration pattern positions below. The SfIS reconstruction method starts at a large voxel resolution (here, 4 mm), and refines the reconstruction at finer resolutions using the prior level's results. **Best viewed in color.**

with SfIS, divides the voxel size by eight and continues with search with SfIS, using the previous voxel size's result as an initial solution. In this work, we performed experiments on all of the samples with two different parameter sets. In the first, the initial voxel size is 4 mm, the number of voxel divisions is two, and the final voxel size is 2 mm. In the second set of experiments, the initial voxel size is 2 mm, the number of voxel divisions is two and the final voxel resolution is 1 mm. An additional parameter is the factor that the the input image is resized down, that value is 4 for both experiments. The initial image size is 6000 by 4000 pixels.

Experiments. We focus on two primary experiments. The first is to quantitatively measure and evaluate our 3D model reconstructions against objects with a known shape. These objects with a known shape are the ground truth samples, with 3D model files that are 3D printed. The second experiment is to qualitatively assess the system's ability to reconstruct various fruit models across different scales. We reconstructed the objects in both experiments at 1mm and 2mm resolution.

Ground Truth Samples. Fruit form, especially in 3 dimensions, can be difficult to quantify. To assess the accuracy of our system, we selected shapes for which we had 3D model files, printed those files, and then reconstructed the models from image data with the phenotyping system. Through this process we can characterize the performance of our method on reconstructing different shape types with durable objects versus individual fruit measurements, where the fruit decays quickly and the human-made measurement cannot be precisely replicated.

A motivation for using 3D printed objects is to have a way to quantitatively assess the performance of the phenotyping system, with a durable artifact that can be stored indefinitely and re-printed and/or scaled if needed. Since we have the originating 3D model file, we can compare the reconstruction and the ground truth object in ways that human-made measurements are unable to, by assessing differences in surface area and volume. This is in contrast to measurements on fruits or tubers that will not persist past a single session and may suffer from measurement error.

We identified 4 digital objects from Thingiverse (`https://www.thingiverse.com`) that had good representation of many different shapes that are both common and uncommon in 3D biological structures, such as fruit and tubers: convex regions, saddle regions, and locally concave regions, shown in Figure 7. We scaled these 4 objects prior to printing so that we would have different size representations. We 3D printed these 11 object×scale stereolithography (STL) format files (61–64) using a commercial-grade 3D printer. The 3D printed objects were then imaged in our system and reconstructed from the 2D images.

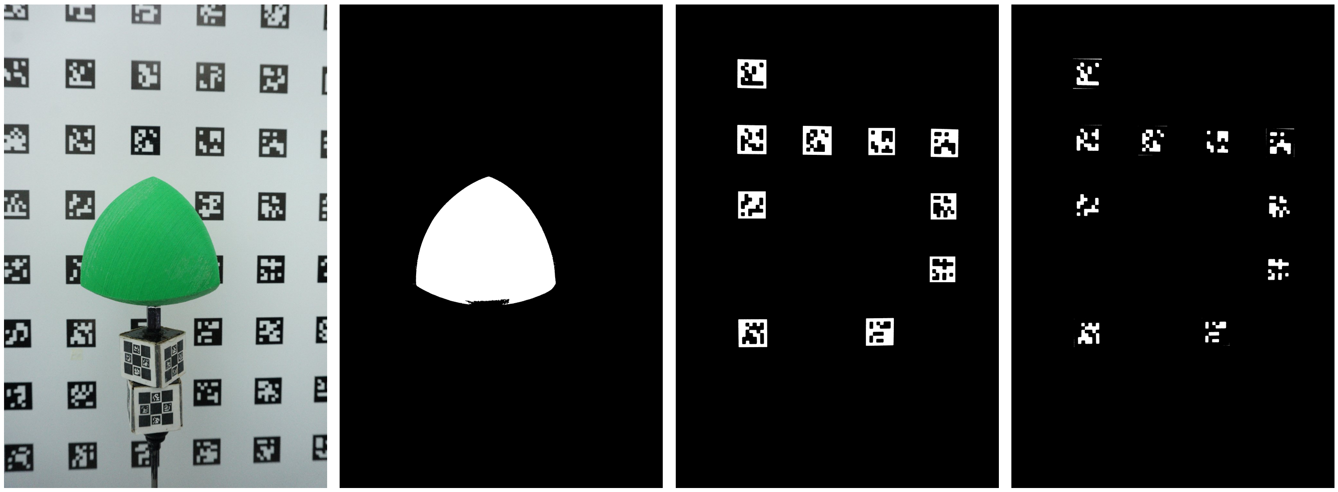

(A) Typical image from data acquisition **(B)** Desired segmented image **(C)** Tag dark intensities mask **(D)** Tag light intensities mask

Fig. 6. An example of successful background segmentation. Segmentation process from sample 1, one of the 3D printed objects. arUco tags are detected in individual images in the data acquisition step (a). Background subtraction generates the segmented image (b). Individual tags are segmented to separate the dark (c) and light (d) intensities; these regions are used to model the background. Notice a small segmentation error from shadow on the bottom of the tetrahedron. See text for more details. **Best viewed in color.**

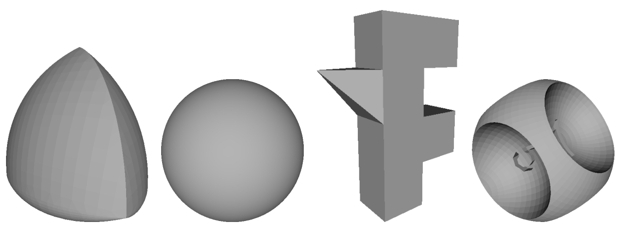

(A) Tetrahedron **(B)** Sphere **(C)** F-Object **(D)** 6-sided Spherical Die

Fig. 7. Ground truth objects. The four classes of ground truth objects used in this study. The 3D model files were used to print out physical copies, which were then imaged with our phenotyping system and reconstructed. 7A, 7B, and 7C can all be reconstructed with our system because they do not have locally concave regions, while the depressions in 7D cannot.

Quantitative Analysis. Once the ground truth objects are reconstructed, some postprocessing is done to align the reconstructed 3D model with the digital ground truth 3D model files. Specifically, we used R 4.1.0 (65) to perform quantitative comparisons between ground-truth objects and reconstructed models with the packages Morpho and Rvcg (66), rgl (67), Lithics3D (68), and mesheR (69). The reconstructed models are in Polygon (PLY) file format and were read using *Rvcg::vcgPlyRead()*. STL objects, e.g., the ground-truth objects, were read using *rgl::readSTL()* and converted to mesh3d using *rgl::as.mesh3d()*. The *mesheR::icp()* function was used to perform the iterative closest point algorithm between the reconstruction and the ground-truth triangular meshes, with 100 iterations and allowing for reflection.

For these quantitative analyses, we chose to measure the magnitude of the primary, secondary, and tertiary axes, e.g., X, Y, and Z, the surface area SA, and the volume Vol. the difference in magnitude between two models, δ_X, δ_Y, and δ_Z, are calculated as the difference between the magnitude of the first, second, and third axes of the reconstruction and ground-truth following ICP alignment:

$$\delta_X = (X_{R,max} - X_{R,min}) - (X_{G,max} - X_{G,min}) \quad (4)$$

$$\delta_Y = (Y_{R,max} - Y_{R,min}) - (Y_{G,max} - Y_{G,min}) \quad (5)$$

$$\delta_Z = (Z_{R,max} - Z_{R,min}) - (Z_{G,max} - Z_{G,min}) \quad (6)$$

where $X_{G,min}$ and $X_{G,max}$ are the minimum and maximum value of the first dimension of the ground-truth object G, respectively, and $X_{R,min}$ and $X_{R,max}$ are the minimum and maximum value of the first axis of the reconstructed object R, respectively. The *Morpho::meshDist()* function was used to calculate and visualise distances between 3D objects. The distance of the reconstructed model from the ground truth is summarized using root mean square error (RMSE). RMSE is calculated as:

$$RMSE = \left(\frac{1}{n} \sum_{i=1}^{n} \delta_i^2 \right)^{1/2} \quad (7)$$

where δ_i is the distance between i-th pair of n corresponding points on the surface of the reconstruction and ground-truth objects. The volume and surface area of models was extracted using *Lithics3D::mesh_volume()* and *Lithics3D::mesh_area()*, respectively The *rgl::shade3d()* function was used to visually compare 3D objects. All regressions were performed using *stats::lm()*.

Sample Collection. We purchased fresh fruit and produce from a local grocery store in Davis, CA, USA for qualitative assessment. In total we purchased, scanned, and reconstructed 48 objects; including 18 strawberries, 12 potatoes, 5 grapes, 7 peppers, and 4 Bosch and 2 red Anjuo pears. We want to test our approach for robustness, and so chose fruit and produce with different scales, colors, levels of glossiness, and other features.

Qualitative Comparisons. Reconstructed fruit were oriented using principal components analysis (PCA) with the *stats::prcomp()* function. PCA orientation of the 3D reconstructed models orients results in 3 axes corresponding to the primary (X), secondary (Y), and tertiary (Z) axes ordered by magnitude, e.g., X $\geq$ Y $\geq$ Z. PCA oriented models were visually inspected using *rgl::shade3d()* functions. Quantitative measurements (Volume, Surface Area, X, Y, and Z) were measured from the PCA oriented reconstructed models.

Results

Overall assessment of platform. Our combination of hardware and software was able to accurately reconstruct models of fruit (Figure 2) and ground-truth objects (Figures 8 and 9; Table 1). Image acquisition occurs in 9 s sessions and is buffered by approximately 15 seconds while the cameras finish writing photos to storage and the following object is staged on the pedestal (Figures 1 and 3). It is possible to achieve about two sessions per minute with current parameters and hardware.

The image calibration, segmentation, and reconstruction steps then proceed remotely following data organization and storage, which is an important consideration in practice. Using the 1 mm experiment to compute average run times, over all objects it took on average 27 seconds for calibration, 33.4 seconds for segmentation, and 413 seconds (6:53 minutes) for reconstruction, for an average total time of 7:53 minutes. All of these run times steps included load and write times of results and were generated on a workstation with one 12-core 2.70GHz processor and 192 GB of RAM.

The calibration, segmentation, and reconstruction steps are automated, and each of those steps are parallelized to some extent. Once the bounding box size was determined, the whole directory of samples was processed with a program that called each of the calibration, segmentation, and reconstruction steps, and was not supervised other than starting the process. Still, with new configurations or objects, or in case of failure, examining the output of each of the steps can indicate where there are problems, such as in the case of calibration or segmentation, on which the reconstruction step depends. For instance, in this hardware setup the calibrated camera poses should form a ring as in Figure 4C. Accurate segmentations may be problematic for some objects, such as in Figure 8(F), the segmentation had false negatives at the bottom of the pear. Consequently, that part of the fruit is not reconstructed.

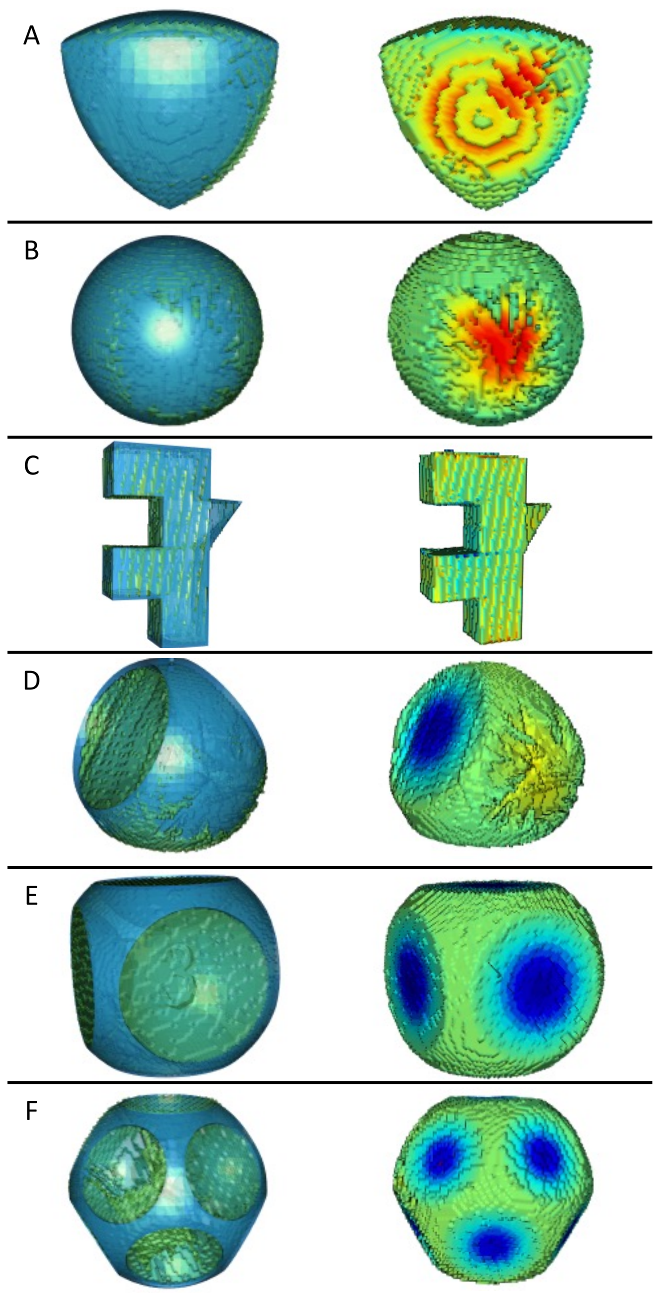

Fig. 8. Reconstructions of six ground-truth objects. (Left Column) Reconstructed model (green) overlaid by ground-truth (blue) following ICP alignment. (Right column) Heat map showing difference between reconstructed model and ground-truth object. Red represents regions where the ground-truth is larger than the model. Blue represents regions where the ground-truth is smaller than the model. Teal represents regions where there is no difference between ground-truth and model. (A) A representative tetrahedron (Tetra_1), (A) a representative sphere (Sphere_1) (A) the smaller "F" shaped object, (D) the three sided die, (E) the six sided die, and (F) the twelve sided die. Only 1mm resolution models are shown. **Best viewed in color.**

Quantitative assessment of ground truth samples. In general, our approach performed very well on the ground-truth examples (Table 1) and only failed in ways that are were expected given the assumptions and constraints of our system. [1] Major deviations between reconstruction and ground-truth in the major axis are typically small (maximum 4.74 mm) and RMSE for the entire surface is $\leq$ 2.70 mm, for the models without concavities. We found a strong correlation $R^2 \approx 0.99$ for most measurements between the reconstructed models and the ground-truth objects without concavities (Figure 9). The surface area (SA) R^2 was 0.979 for the models without concavities, and this R^2 value is the lowest value for the traits we examined on objects without concavities. We found that SA of the reconstructed models were upwardly biased relative to the ground-truth objects (110-

[1] These assumptions are discussed in the Materials and Methods section, 'Reconstruction' subsection.

Table 1. Accuracy metrics, including RMSE, difference in major axis length, and ratios of volume and surface area, from two experiments with eleven ground-truth objects. Differences in mm and ratios are reported between model and ground-truth object. Die_3, Die_6, and Die_12 have local concavities, while the other objects do not.

Object	Res[a]	RMSE[b]	δ^c_X	δ^d_Y	δ^e_Z	δ^f_{Vol}	δ^g_{SA}
Tetra_2	1	0.62	-0.12	-0.14	-0.14	0.97	1.20
	2	0.76	0.04	1.17	-1.17	0.96	1.18
Tetra_3	1	1.17	0.61	0.21	-0.93	0.96	1.24
	2	1.21	1.32	0.35	0.97	0.95	1.23
Tetra_4	1	1.62	-0.29	-0.83	0.05	0.96	1.26
	2	1.69	1.51	0.31	0.06	0.95	1.25
Sphere_2	1	0.97	0.47	-0.26	-0.53	0.96	1.23
	2	1.07	0.97	0.21	1.38	0.95	1.21
Sphere_3	1	1.54	-3.27	-1.01	-1.25	0.96	1.27
	2	1.57	-4.74	-1.07	-1.20	0.95	1.26
Sphere_4	1	2.61	-0.70	-1.70	-1.81	0.95	1.33
	2	2.70	0.33	-1.19	-0.90	0.94	1.33
F_2	1	0.84	0.73	-0.55	0.84	0.98	1.13
	2	1.03	1.10	0.75	1.17	0.98	1.12
F_3	1	0.66	4.66	0.82	1.93	1.01	1.16
	2	0.79	4.11	1.45	2.09	1.00	1.15
Die_3	1	7.05	-3.16	-9.64	5.68	1.17	1.13
	2	6.64	-2.70	-5.18	1.14	1.16	1.10
Die_6	1	6.31	8.71	2.45	6.52	1.29	1.03
	2	6.00	9.77	8.25	7.97	1.28	1.03
Die_12	1	9.34	0.28	-4.21	8.61	1.20	1.02
	2	8.70	0.27	-2.02	9.57	1.20	1.00

[a] Reconstruction resolution in mm.
[b] RMSE of the model surface against the ground-truth surface mm.
[c] Difference between the X axis length of the model and the ground-truth mm.
[d] Difference between the Y axis length of the model and the ground-truth mm.
[e] Difference between the Z axis length of the model and the ground-truth mm.
[f] Ratio of the volume ($\text{Vol}^{1/3}$) of the model over the volume of the ground-truth.
[g] Ratio of the surface area ($\text{SA}^{1/2}$) of the model over the SA of the ground-truth.

120%). This bias is most likely to do that fact that our models, which are made of voxels (Figure 5), have rough surfaces while the ground-truth objects are perfectly smooth (Figure 8). In general, the size measurement of these objects are very accurate, albeit imperfect, at both 1mm and 2mm resolutions (Figure 9).

We noticed minor segmentation false negative errors from shadows at the lower portion of the object in the images; in the reconstruction, these segmentation errors are realized as jagged portions where the printed object was attached to the pedestal, especially visible in Figure 8(A), (B), and (D). Reconstruction errors, resulting from small segmentation errors, do not have a large impact on the overall accuracy based on the metrics we assessed. However, large segmentation errors over multiple images will affect the reconstruction quality, such as in Figure 2F.

Our approach to reconstruction is unable to recover concavities, as demonstrated by the three spherical die examples (Figure 8D-F). The indexed faces of these models are sunken into the body of the model, resulting in multiple large depressions per die (Figure 7D). As is clearly shown in Figure 8D-F, our reconstructions are more similar to a 3D convex hull, yielding a flat surface over the large concavities in the true models. This is reflected by the rows corresponding to the three die in Table 1. In these cases, the Volume is 115-130% greater than the ground-truth model. In general this is not an issue for types of fruit that do not have concavities.

Qualitative assessment. We found that our platform and approach to reconstruction is both quantitatively accurate (Table 1; Figure 9), as well as visually accurate in most cases (Figures 2 and 8). For the peppers, grapes, strawberries, and potatoes, we found no systematic errors in reconstruction. However, the Anjou pears were troublesome to segment leading to the bottom half of the models being severely deformed. The reason for the segmentation error is the use of a general segmentation approach that worked without extensive tuning for the whole set of samples. However, if one were to have a large batch of objects with particular color features, fine tuning the user/session specific parameters for segmentation is important for yielding accurate models. Segmentation errors of this severe type appeared in 3 out of 59 objects that we imaged and the rest of the models appear to reflect the physical objects that were imaged.

Discussion

We have described a low-cost ($1,600 USD), high-throughput (9 s data acquisition), modular reconstruction system that can be used in lab settings or in the field on a table, with a fast data acquisition speed of 9 seconds per object. We will discuss several design decisions that lead to flexibility.

The use of consumer grade materials results in a relatively inexpensive system; multiple systems could be built and increase sample throughput during high-volume times of the year. This means that larger experiments can be executed enabling more robust studies. Our system is modular, allowing users with different interests to experiment and explore different cameras, sensors, or lights. This system is easily repaired and replaced if any damages are incurred by the hardware components.

The system has short session times and it only takes 9 seconds to acquire images on a single sample, regardless of the number of cameras. In fact, we found that were frequently rate limited by the write speed from the cameras on board cache to the SD card. More often than not, the next object was prepped around the same time that the cameras had cleared their on-board cache.

This system calibrates the camera from the image data acquired for the samples. The calibration is an absolute (as opposed to relative one, with an unknown scale factor), so the physical units of voxels are known.

Key assumptions and considerations. We highlight some key assumptions of the methodologies used in our system that are important for those considering it for a range of objects not treated in this paper.

Shape classes. Users who want to accurately represent locally concave shapes — shapes with egg-shaped depressions as demonstrated by the 3, 6, and 12-sided die in our calibration objects (Table 1; Figure 9; Figure 7D)— will need to substitute some portions of this system to recover such features. Shape from Silhouette is not able to recover locally concave regions. However, most of the types of objects we

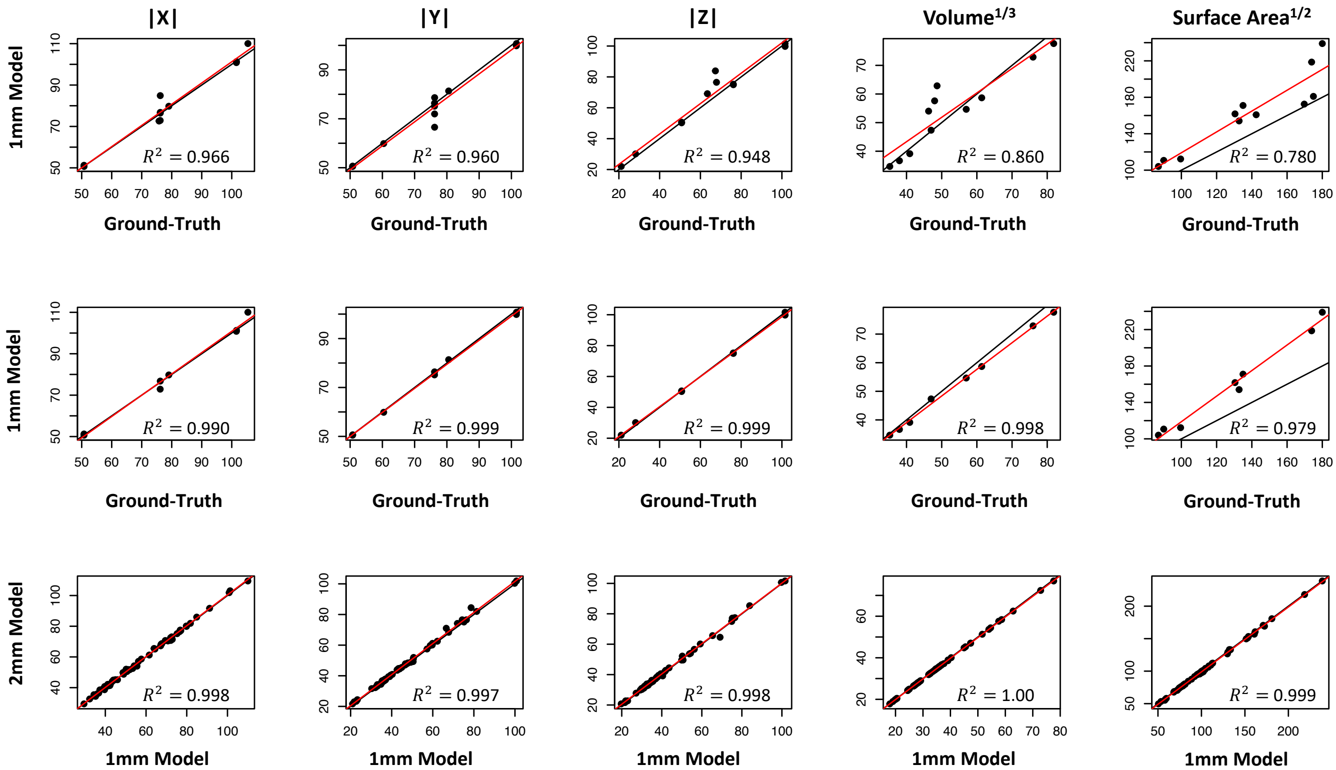

Fig. 9. Ground-truth calibration experiment results. *In silico* measurement of reconstructed 3D models and ground-truth objects. Measurements include length of primary (X), secondary (Y), and tertiary (Z) axes, the cube root of the volume (Vol$^{1/3}$), and the square root of the surface area (SA$^{1/2}$). (Top row) All 1mm reconstructed models against ground-truth objects including the three dice, (Middle row) 1mm reconstructed models against ground-truth objects excluding the three dice, and (Bottom row) all reconstructed fruit models in 1mm (x axis) and 2mm (y axis). All measurements are reported in mm. The adjusted R^2 from linear regression is reported in the plot. The solid black line is the identity line. The solid red line is the linear regression of y regressed onto x identity line.

envisioned imaging with this system, fruits and tubers, happen to be mostly non-concave.

Large or fragile objects. The stepper motor is non-continuous and takes "steps" to provide rotation which causes vibrations through the object. When objects are unbalanced, those vibrations can cause movement of the object out of the center of the scene between different frames. Users should pay close attention to lateral movement of the object during rotation. Similarly, if the larger fruit are of interest, some additional modifications will be required to stabilize the pedestal during rotation. If these issues are a concern, it may be beneficial to construct a multi camera network that surrounds the target (54) or a platform that enables the camera to move easily around a fixed target (40). The same is true for rice panicles or maize tassels, because they are not rigid body objects and the vibrations of the stepper motor are more likely to lead to changes in relative position of parts of the object between frames causing issues during reconstruction. These types of objects are better suited for systems where the cameras move relative to the object or in a multi-camera network (40, 54).

Measurement of chArUco markers. Third, accuracy is intimately tied to the measurement of the arUco and chArUco markers and any inaccuracies in those measurements will lead to systematic biases in the measurement of the 3D reconstruction. For example, if chArUco markers are declared to be 10mm, when they are really 20mm, all of the models will 2x smaller than the real object than they measured. Users should print all calibration targets (aruco or chArUco) and ground-truth samples with a high quality 2D or 3D printer to ensure sharp corners and well-defined edges. Further, users are encouraged to verify the proportions of the printed calibration targets with high-precision calipers prior to calibration.

Segmentation quality. Model quality is directly linked to segmentation quality (Figure 6A and 6B) as we have mentioned throughout this paper. If an object is only partially segmented, and this false negative error happens in multiple frames, part of the model may end up distorted or completely missing (Figure 2F). It is vital, as in any system, that users examine the quality of reconstructions prior to measurement and go back to calibration and segmentation outputs to identify the source of errors. In this work, we chose one set of segmentation parameters that performed reasonably well for all objects, but we recommend that users perform segmentation with parameters optimized for their research samples and imaging conditions.

Conclusions

In conclusion, we presented a phenotyping system for capturing, calibrating, and reconstructing 3D models of small-to-moderately sized fruit and tubers. The low-cost and reliance on consumer-grade materials makes it obtainable to almost any program; short session times allows researchers to increase the number of samples per hour, and high accuracy means that the digital representations will yield absolute measurements on objects that do not degrade over time, yielding a viable option for research and breeding programs interested in pursuing 3D fruit phenotyping.

DATA AVAILABILITY STATEMENT

The input data, including the images and configuration files, is available from Zenodo in (70), 10.5281/zenodo.5155765. All 3D model results, at 1 and 2 mm resolution, are at the same dataset source.
The ground truth objects were modified from objects at Thingiverse.com, and each had a different license. For this reason, there are four different datasets for the ground truth objects:

- Tetrahedra: 10.5281/zenodo.5153992 (61)
- Spheres: 10.5281/zenodo.5154029 (62),
- Sphere Dice: 10.5281/zenodo.5155690 (63).
- F-object: 10.5281/zenodo.5155743 (64)

ACKNOWLEDGEMENTS

The authors thank Scott Wolford at USDA-ARS-AFRS for 3D printing the ground truth objects.
Any mention of trade names or commercial products in this publication is solely for the purpose of providing scientific information and does not constitute recommendation or endorsement by the United States Department of Agriculture. USDA is an equal opportunity provider and employer.

FUNDING STATEMENT

This research was supported by awards from the University of California, Davis Henry A. Jastro Research Scholarship Award program. This research was supported by USDA-ARS Project 8080-21000-024-00-D.

AUTHOR CONTRIBUTIONS

Conceptualization: MJF, AT **Data curation:** MJF, AT **Formal Analysis:** MJF, AT **Funding Acquisition:** MJF, AT **Investigation:** MJF, AT **Methodology:** MJF, AT **Project administration:** MJF, AT **Resources:** MJF, AT **Software:** MJF, AT **Supervision:** MJF, AT **Validation:** MJF, AT **Visualization:** MJF, AT **Writing – original draft preparation:** MJF, AT **Writing – review & editing:** MJF, AT

COMPETING FINANCIAL INTERESTS

The authors declare no competing interests.

Bibliography

1. Corrado Costa, Francesca Antonucci, Federico Pallottino, Jacopo Aguzzi, Da-Wen Sun, and Paolo Menesatti. Shape analysis of agricultural products: a review of recent research advances and potential application to computer vision. *Food and Bioprocess Technology*, 4 (5):673–692, 2011.
2. Amelia Gaston, Sonia Osorio, Béatrice Denoyes, and Christophe Rothan. Applying the solanaceae strategies to strawberry crop improvement. *Trends in plant science*, 25(2):130–140, 2020.
3. Megan M Mathey, Sonali Mookerjee, Kazim Gündüz, James F Hancock, Amy F Iezzoni, Lise L Mahoney, Thomas M Davis, Nahla V Bassil, Kim E Hummer, Philip J Stewart, Vance M Whitaker, Daniel J Sargent, Beatrice Denoyes, Iraida Amaya, Eric van de Weg, and Chad E Finn. Large-Scale Standardized Phenotyping of Strawberry in RosBREED. *Journal of the American Pomological Society*, page 12, 2013.
4. Mitchell J Feldmann, Michael A Hardigan, Randi A Famula, Cindy M Lopez, Amy Tabb, Glenn S Cole, and Steven J Knapp. Multi-dimensional machine learning approaches for fruit shape phenotyping in strawberry. *GigaScience*, 9(5):giaa030, 2020.
5. Zoë Migicovsky, Kyle M Gardner, Daniel Money, Jason Sawler, Joshua S Bloom, Peter Moffett, C Thomas Chao, Heidi Schwaninger, Gennaro Fazio, Gan-Yuan Zhong, et al. Genome to phenome mapping in apple using historical data. *The plant genome*, 9(2), 2016.
6. Sarah D. Turner, Shelby L. Ellison, Douglas A. Senalik, Philipp W. Simon, Edgar P. Spalding, and Nathan D. Miller. An Automated Image Analysis Pipeline Enables Genetic Studies of Shoot and Root Morphology in Carrot (Daucus carota L.). *Frontiers in Plant Science*, 9, November 2018. ISSN 1664-462X. doi: 10.3389/fpls.2018.01703.
7. Sofia Visa, Chunxue Cao, Brian McSpadden Gardener, and Esther van der Knaap. Modeling of tomato fruits into nine shape categories using elliptic fourier shape modeling and Bayesian classification of contour morphometric data. *Euphytica*, 200(3):429–439, December 2014. ISSN 0014-2336, 1573-5060. doi: 10.1007/s10681-014-1179-0.
8. Joe Q He, Richard J Harrison, and Bo Li. A novel 3d imaging system for strawberry phenotyping. *Plant Methods*, 13(1):1–8, 2017.
9. M Minervini, H Scharr, and S A Tsaftaris. Image analysis: The new bottleneck in plant phenotyping [applications corner]. *IEEE Signal Process. Mag.*, 32(4):126–131, July 2015.
10. Abbas Atefi, Yufeng Ge, Santosh Pitla, and James Schnable. Robotic technologies for high-throughput plant phenotyping: Contemporary reviews and future perspectives. *Frontiers in Plant Science*, 12:1082, 2021.
11. Allan G. White, Peter A. Alspach, Rosemary H. Weskett, and Lester R. Brewer. Heritability of fruit shape in pears. *Euphytica*, 112(1):1–7, Mar 2000. ISSN 1573-5060. doi: 10.1023/A:1003761118890.
12. T. Ishikawa, A. Hayashi, S. Nagamatsu, Y. Kyutoku, I. Dan, T. Wada, K. Oku, Y. Saeki, T. Uto, T. Tanabata, S. Isobe, and N. Kochi. Classification of Strawberry Fruit Shape by Machine Learning. *ISPRS - International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLII-2:463–470, May 2018. ISSN 2194-9034. doi: 10.5194/isprs-archives-XLII-2-463-2018.
13. Zoë Migicovsky, Kyle M. Gardner, Daniel Money, Jason Sawler, Joshua S. Bloom, Peter Moffett, C. Thomas Chao, Heidi Schwaninger, Gennaro Fazio, Gan-Yuan Zhong, and Sean Myles. Genome to Phenome Mapping in Apple Using Historical Data. *The Plant Genome*, 9(2):0, 2016. ISSN 1940-3372. doi: 10.3835/plantgenome2015.11.0113.
14. Graham W Horgan. The statistical analysis of plant part appearance—a review. *Computers and Electronics in Agriculture*, 31(2):169–190, 2001. doi: 10.1016/S0168-1699(00)00190-3.
15. Graham W Horgan, Mike Talbot, and Jonathan C Davey. Use of statistical image analysis to discriminate carrot cultivars. *Computers and Electronics in Agriculture*, 31(2):191–199, 2001. doi: 10.1016/S0168-1699(00)00191-5.
16. Shan Wu, Josh P. Clevenger, Liang Sun, Sofia Visa, Yuji Kamiya, Yusuke Jikumaru, Joshua Blakeslee, and Esther van der Knaap. The control of tomato fruit elongation orchestrated by sun, ovate and fs8.1 in a wild relative of tomato. *Plant Science*, 238:95–104, September 2015. ISSN 01689452. doi: 10.1016/j.plantsci.2015.05.019.
17. Shan Wu, Biyao Zhang, Neda Keyhaninejad, Gustavo R. Rodríguez, Hyun Jung Kim, Manohar Chakrabarti, Eudald Illa-Berenguer, Nathan K. Taitano, M. J Gonzalo, Aurora Díaz, Yupeng Pan, Courtney P. Leisner, Dennis Halterman, C. Robin Buell, Yiqun Weng, Shelley H. Jansky, Herman van Eck, Johan Willemsen, Antonio J. Monforte, Tea Meulia, and Esther van der Knaap. A common genetic mechanism underlies morphological diversity in fruits and other plant organs. *Nature Communications*, 9(1), December 2018. ISSN 2041-1723. doi: 10.1038/s41467-018-07216-8.
18. Antonio J. Monforte, Aurora Diaz, Ana Caño-Delgado, and Esther van der Knaap. The genetic basis of fruit morphology in horticultural crops: lessons from tomato and melon. *Journal of Experimental Botany*, 65(16):4625–4637, August 2013. ISSN 1460-2431, 0022-0957. doi: 10.1093/jxb/eru017.
19. Xiong Liao, Mengsi Li, Bin Liu, Miaoling Yan, Xiaomin Yu, Hailing Zi, Renyi Liu, and Chizuko Yamamuro. Interlinked regulatory loops of ABA catabolism and biosynthesis coordinate fruit growth and ripening in woodland strawberry. *Proceedings of the National Academy of Sciences*, 115(49):E11542–E11550, December 2018. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1812575115.
20. Anand Seethepalli, Larry M York, Hussien Almtarfi, Felix B Fritschi, and Alina Zare. A novel multi-perspective imaging platform (M-PIP) for phenotyping soybean root crowns in the field increases throughput and separation ability of genotype root properties. *bioRxiv*, page 309161, May 2018.
21. Anand Seethepalli, Haichao Guo, Xiuwei Liu, Marcus Griffiths, Hussien Almtarfi, Zenglu Li, Shuyu Liu, Alina Zare, Felix B Fritschi, Elison B Blancaflor, Xue-Feng Ma, and Larry M York. RhizoVision crown: An integrated hardware and software platform for root crown phenotyping. *Plant Phenomics*, 2020:3074916, February 2020.
22. Christopher N Topp, Anjali S Iyer-Pascuzzi, Jill T Anderson, Cheng-Ruei Lee, Paul R Zurek, Olga Symonova, Ying Zheng, Alexander Bucksch, Yuriy Mileyko, Taras Galkovskyi, Brad T Moore, John Harer, Herbert Edelsbrunner, Thomas Mitchell-Olds, Joshua S Weitz, and Philip N Benfey. 3D phenotyping and quantitative trait locus mapping identify core regions of the rice genome controlling root architecture. *Proc. Natl. Acad. Sci. U. S. A.*, 110(18): E1695–704, April 2013.
23. Mathilde Balduzzi, Brad M. Binder, Alexander Bucksch, Cynthia Chang, Lilan Hong, Anjali S. Iyer-Pascuzzi, Christophe Pradal, and Erin E. Sparks. Reshaping Plant Biology: Qualitative and Quantitative Descriptors for Plant Morphology. *Frontiers in Plant Science*, 08, February 2017. ISSN 1664-462X. doi: 10.3389/fpls.2017.00117.
24. Jirawat Salungyu, Suriphat Thaitad, Alexander Bucksch, Jitrana Kengkanna, and Patompong Johns Saengwilai. From lab to field: Open tools facilitating the translation of maize root traits. *Field Crops Research*, 255:107872, 2020.
25. Mao Li, Hong An, Ruthie Angelovici, Clement Bagaza, Albert Batushansky, Lynn Clark, Viktoriya Coneva, Michael J. Donoghue, Erika Edwards, Diego Fajardo, Hui Fang, Margaret H. Frank, Timothy Gallaher, Sarah Gebken, Theresa Hill, Shelley Jansky, Baljinder Kaur, Phillip C. Klahs, Laura L. Klein, Vasu Kuraparthy, Jason Londo, Zoë Migicovsky, Allison Miller, Rebekah Mohn, Sean Myles, Wagner C. Otoni, J. C. Pires, Edmond Rieffer, Sam Schmerler, Elizabeth Spriggs, Christopher N. Topp, Allen Van Deynze, Kuang Zhang, Linglong Zhu, Braden M. Zink, and Daniel H. Chitwood. Topological Data Analysis as a Morphometric Method: Using Persistent Homology to Demarcate a Leaf Morphospace. *Frontiers in Plant Science*, 9, April 2018. ISSN 1664-462X. doi: 10.3389/fpls.2018.00553.
26. Mao Li, Margaret H Frank, Viktoriya Coneva, Washington Mio, Daniel H Chitwood, and Christopher N Topp. The persistent homology mathematical framework provides enhanced genotype-to-phenotype associations for plant morphology. *Plant physiology*, 177(4):1382–1395, 2018.
27. Cedar Warman, Christopher M Sullivan, Justin Preece, Michaela E Buchanan, Zuzana Vejlupkova, Pankaj Jaiswal, and John E Fowler. A cost-effective maize ear phenotyping platform enables rapid categorization and quantification of kernels. *The Plant Journal*, 106(2): 566–579, 2021.
28. Sapan Naik, Bankim Patel, and Rashmi Pandey. Shape, size and maturity features extraction with fuzzy classifier for non-destructive mango (mangifera indica l., cv. kesar) grading. In *2015 IEEE Technological Innovation in ICT for Agriculture and Rural Development (TIAR)*, pages 1–7. IEEE, 2015.
29. Mathieu Gaillard, Chenyong Miao, James C Schnable, and Bedrich Benes. Voxel carving-based 3d reconstruction of sorghum identifies genetic determinants of light interception efficiency. *Plant direct*, 4(10):e00255, 2020.
30. Ayan Chaudhury, Frédéric Boudon, and Christophe Godin. 3d plant phenotyping: All you need is labelled point cloud data. In *European Conference on Computer Vision*, pages 244–260. Springer, 2020.
31. Alexander Feldman, Haozhou Wang, Yuya Fukano, Yoichiro Kato, Seishi Ninomiya, and Wei

Guo. Easydcp: An affordable, high-throughput tool to measure plant phenotypic traits in 3d. *Methods in Ecology and Evolution*, 2021.

32. Helin Dutagaci, Pejman Rasti, Gilles Galopin, and David Rousseau. Rose-x: an annotated data set for evaluation of 3d plant organ segmentation methods. *Plant methods*, 16(1):1–14, 2020.
33. Simon Artzet, Tsu Wei Chen, Jérôme Chopard, Nicolas Brichet, Michael Mielewczik, Sarah Cohen-Boulakia, Llorenç Cabrera Bosquet, François Tardieu, Christian Fournier, and Christophe Pradal. Phenomenal: An automatic open source library for 3d shoot architecture reconstruction and analysis for image-based plant phenotyping. *bioRxiv*, 2020.
34. Shota Teramoto, Satoko Takayasu, Yuka Kitomi, Yumiko Arai-Sanoh, Takanari Tanabata, and Yusaku Uga. High-throughput three-dimensional visualization of root system architecture of rice using x-ray computed tomography. *Plant Methods*, 16:1–14, 2020.
35. Stefan Paulus, Jan Behmann, Anne-Katrin Mahlein, Lutz Plümer, and Heiner Kuhlmann. Low-cost 3d systems: suitable tools for plant phenotyping. *Sensors*, 14(2):3001–3018, 2014.
36. Lei Li, Qin Zhang, and Danfeng Huang. A review of imaging techniques for plant phenotyping. *Sensors*, 14(11):20078–20111, 2014.
37. Suxing Liu, Carlos Sherard Barrow, Meredith Hanlon, Jonathan P Lynch, and Alexander Bucksch. Dirt/3d: 3d phenotyping for maize (zea mays) root architecture in the field. *bioRxiv*, 2020.
38. Tyler Dowd, Samuel McInturf, Mao Li, and Christopher N Topp. Rated-m for mesocosm: allowing the multimodal analysis of mature root systems in 3d. *Emerging Topics in Life Sciences*, 5(2):249–260, 2021.
39. Ni Jiang, Eric Floro, Adam L Bray, Benjamin Laws, Keith E Duncan, and Christopher N Topp. Three-dimensional time-lapse analysis reveals multiscale relationships in maize root systems with contrasting architectures. *The Plant Cell*, 31(8):1708–1722, 2019.
40. Jaspreet Sandhu, Feiyu Zhu, Puneet Paul, Tian Gao, Balpreet K Dhatt, Yufeng Ge, Paul Staswick, Hongfeng Yu, and Harkamal Walia. Pi-plat: a high-resolution image-based 3d reconstruction method to estimate growth dynamics of rice inflorescence traits. *Plant Methods*, 15(1):1–13, 2019.
41. Weijuan Hu, Can Zhang, Yuqiang Jiang, Chenglong Huang, Qian Liu, Lizhong Xiong, Wanneng Yang, and Fan Chen. Nondestructive 3d image analysis pipeline to extract rice grain traits using x-ray computed tomography. *Plant Phenomics*, 2020, 2020.
42. Samantha Thi Porter, Morgan Roussel, and Marie Soressi. A simple photogrammetry rig for the reliable creation of 3d artifact models in the field: lithic examples from the early upper paleolithic sequence of les cottés (france). *Advances in Archaeological Practice*, 4 (1):71–86, 2016.
43. Yongjian Wang, Weiliang Wen, Sheng Wu, Chuanyu Wang, Zetao Yu, Xinyu Guo, and Chunjiang Zhao. Maize plant phenotyping: comparing 3d laser scanning, multi-view stereo reconstruction, and 3d digitizing estimates. *Remote Sensing*, 11(1):63, 2019.
44. Suxing Liu, Lucia M Acosta-Gamboa, Xiuzhen Huang, and Argelia Lorence. Novel low cost 3d surface model reconstruction system for plant phenotyping. *Journal of Imaging*, 3(3):39, 2017.
45. Christian Peter Klingenberg and Grant S McIntyre. Geometric morphometrics of developmental instability: analyzing patterns of fluctuating asymmetry with procrustes methods. *Evolution*, 52(5):1363–1375, 1998. doi: 10.2307/2411306.
46. Fred L Bookstein. Landmark methods for forms without landmarks: morphometrics of group differences in outline shape. *Medical image analysis*, 1(3):225–243, 1997.
47. Frank P Kuhl and Charles R Giardina. Elliptic fourier features of a closed contour. *Computer Graphics and Image Processing*, 18(3):236–258, 1982. doi: 10.1016/0146-664X(82)90034-X.
48. Daniel H. Chitwood. Imitation, Genetic Lineages, and Time Influenced the Morphological Evolution of the Violin. *PLoS ONE*, 9(10):e109229, October 2014. ISSN 1932-6203. doi: 10.1371/journal.pone.0109229.
49. Daniel H. Chitwood and Wagner C. Otoni. Morphometric analysis of Passiflora leaves: the relationship between landmarks of the vasculature and elliptical Fourier descriptors of the blade. *GigaScience*, 6(1), January 2017. ISSN 2047-217X. doi: 10.1093/gigascience/giw008.
50. Bo Li, Helen M Cockerton, Abigail W Johnson, Amanda Karlström, Eleftheria Stavridou, Greg Deakin, and Richard J Harrison. Defining strawberry shape uniformity using 3d imaging and genetic mapping. *Horticulture research*, 7(1):1–13, 2020.
51. Joseph L Gage, Elliot Richards, Nicholas Lepak, Nicholas Kaczmar, Chinmay Soman, Girish Chowdhary, Michael A Gore, and Edward S Buckler. In-field whole plant maize architecture characterized by latent space phenotyping. *bioRxiv*, page 763342, 2019.
52. Jordan Ubbens, Mikolaj Cieslak, Przemyslaw Prusinkiewicz, Isobel Parkin, Jana Ebersbach, and Ian Stavness. Latent space phenotyping: Automatic image-based phenotyping for treatment studies. *bioRxiv*, page 557678, 2019.
53. Bruce Wang, Timothy Sudijono, Henry Kirveslahti, Tingran Gao, Douglas M Boyer, Sayan Mukherjee, and Lorin Crawford. A statistical pipeline for identifying physical features that differentiate classes of 3d shapes. *The Annals of Applied Statistics*, 15(2):638–661, 2021.
54. Amy Tabb, Henry Medeiros, Mitchell J. Feldmann, and Thiago T. Santos. Calibration of asynchronous camera networks: Calico. *arXiv:1903.06811 [cs]*, March 2019. arXiv: 1903.06811.
55. A. Tabb. Shape from silhouette probability maps: Reconstruction of thin objects in the presence of silhouette extraction and calibration error. In *2013 IEEE Conference on Computer Vision and Pattern Recognition*, pages 161–168, June 2013. doi: 10.1109/CVPR.2013.28.
56. Gwon Hwan An, Siyeong Lee, Min-Woo Seo, Kugjin Yun, Won-Sik Cheong, and Suk-Ju Kang. Charuco board-based omnidirectional camera calibration method. *Electronics*, 7 (12):421, 2018.
57. Sergio Garrido-Jurado, Rafael Munoz-Salinas, Francisco José Madrid-Cuevas, and Rafael Medina-Carnicer. Generation of fiducial marker dictionaries using mixed integer linear programming. *Pattern Recognition*, 51:481–491, 2016.
58. Francisco J Romero-Ramirez, Rafael Muñoz-Salinas, and Rafael Medina-Carnicer. Speeded up detection of squared fiducial markers. *Image and vision Computing*, 76:38–47, 2018.
59. Nobuyuki Otsu. A threshold selection method from gray-level histograms. *IEEE Transactions on Systems, Man, and Cybernetics*, 9(1):62–66, 1979. doi: 10.1109/TSMC.1979.4310076.
60. Amy Tabb. *Shape from inconsistent silhouette: Reconstruction of objects in the presence of segmentation and camera calibration error - ProQuest*. PhD, Purdue University, West Lafayette, Indiana, USA, December 2014.
61. Mitchell J Feldmann and Amy Tabb. Ground truth 3d tetrahedra models, August 2021.
62. Mitchell J Feldmann and Amy Tabb. Ground truth 3d sphere models, August 2021.
63. Mitchell J Feldmann and Amy Tabb. Ground truth 3d sphere dice models, August 2021.
64. Mitchell J Feldmann and Amy Tabb. Ground truth 3d f-object, August 2021.
65. R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2019.
66. Stefan Schlager. Morpho and rvcg – shape analysis in R. In Guoyan Zheng, Shuo Li, and Gabor Szekely, editors, *Statistical Shape and Deformation Analysis*, pages 217–256. Academic Press, 2017. ISBN 9780128104934.
67. Duncan Murdoch and Daniel Adler. *rgl: 3D Visualization Using OpenGL*, 2021. R package version 0.106.8.
68. Cornel Pop. *Lithics3D: A toolbox for 3D analysis of archaeological lithics*, 2019. R package version 0.4.2.
69. Stefan Schlager. *mesheR: Meshing Operations on Triangular Meshes*, 2020. R package version 0.4.200213.
70. Mitchell J Feldmann and Amy Tabb. Data and Results for: Cost-effective, high- throughput fruit phenotyping system for three- dimensional reconstruction of fruit form, August 2021.