

OPEN

A Machine Learning Approach to Growth Direction Finding for Automated Planting of Bulbous Plants

Brian G. Booth*, Jan Sijbers & Jan De Beenhouwer

In agricultural robotics, a unique challenge exists in the automated planting of bulbous plants: the estimation of the bulb's growth direction. To date, no existing work addresses this challenge. Therefore, we propose the first robotic vision framework for the estimation of a plant bulb's growth direction. The framework takes as input three x-ray images of the bulb and extracts shape, edge, and texture features from each image. These features are then fed into a machine learning regression algorithm in order to predict the 2D projection of the bulb's growth direction. Using the x-ray system's geometry, these 2D estimates are then mapped to the 3D world coordinate space, where a filtering on the estimate's variance is used to determine whether the estimate is reliable. We applied our algorithm on 27,200 x-ray simulations from *T. Apeldoorn* bulbs on a standard desktop workstation. Results indicate that our machine learning framework is fast enough to meet industry standards (<0.1 seconds per bulb) while providing acceptable accuracy (e.g. error < 30° in 98.40% of cases using an artificial 3-layer neural network). The high success rates of the proposed framework indicate that it is worthwhile to proceed with the development and testing of a physical prototype of a robotic bulb planting system.

In recent years, agricultural robotics has become an area of intense research as governments around the world target the economic potential of automation in this sector^{1,2}. These economic advantages stem not only from reduced labour costs, but also from more precision planting, harvesting, and advanced care of crops, leading to higher yields³. These reduced production costs can then help in reducing food costs for consumers. The flower growing industry have also shown interest in robotic agricultural techniques in an attempt to produce higher and more consistent yields⁴.

Robotic techniques have been applied to a variety of agricultural tasks, from seeding and transplanting^{5,6}, to weeding⁷, to harvesting⁸. In each of these tasks, agricultural robots need the ability to identify the objects they wish to manipulate and be able to perform the manipulation without damaging the object⁹. These needs have led to a variety of computer vision techniques for object detection^{10–12} and obstacle avoidance^{13,14}. The computer vision algorithms developed in this field vary widely depending on the plants and tasks in question, highlighting the need for specific solutions to different agricultural robotics problems.

One such example of an agricultural robotics task is the automated planting of bulbous plants like onions, garlic, and flower bulbs. This task shares similarities with seeding and transplanting, but differs from both in one key aspect: a plant bulb has a clear direction of growth that is unknown to the robot¹⁵. Given the desire to optimize plant height, crop yield, and crop uniformity, plant bulbs should be planted with their growth directions pointed vertically upward^{16–18}. Positioning the bulb in this orientation reduces the amount of energy the plant has to expend in order to sprout above the ground, resulting in larger, more consistent yields^{19–21}. This constraint on bulb orientation is not present with traditional seeding tasks where the seed can be inserted into the ground in any pose. This task also differs from transplanting where the plant has already been planted and, therefore, its growth direction is already determined. Effectively, the planting of bulbous plants constitutes its own unique task and requires its own computer vision algorithm, one that can identify a plant bulb's growth direction.

To the best of our knowledge, no algorithm currently exists that can consistently identify the growth direction of a plant bulb, despite the noted desire for such an algorithm to be created^{19,22}. Aksenov *et al.* recently proposed

Imec-Vision Lab, Department of Physics, University of Antwerp, B-2610, Antwerp, Belgium. *email: brian.booth@uantwerpen.be

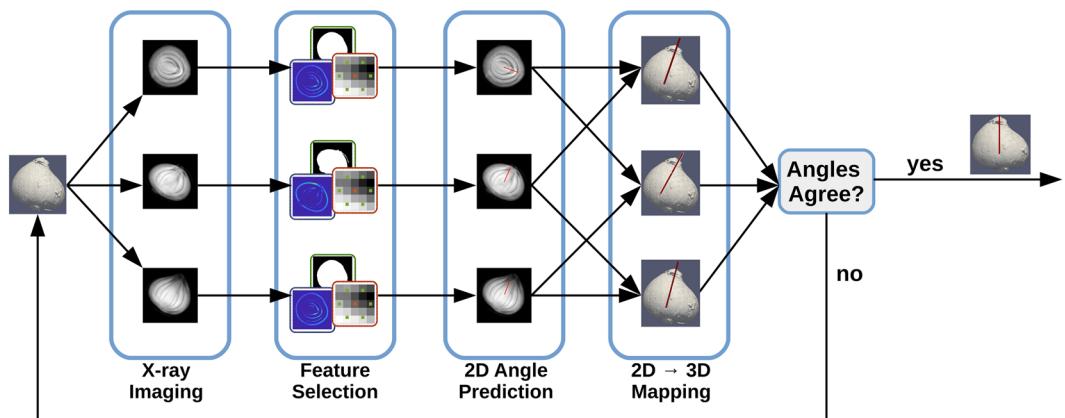


Figure 1. Flow chart of the proposed growth direction finding algorithm. Three x-ray images of the flower bulbs are processed to identify key shape, edge, and texture features. Machine learning algorithms then use these features in a regression to predict the growth direction within each 2D image. These 2D estimates are then paired and mapped to 3D to create three 3D estimates. If these estimates agree, then their average is outputted. Otherwise, the flower bulb is imaged again. See text for further details.

the first oriented planting robot that aims to grip the tip of a plant bulb's in order to position it appropriately during planting, however their approach only achieves a successful bulb planting 51% of the time¹⁶. These results indicate that a plant bulb's tip is not consistently a prominent feature. The bulb's shape, or the presence of attached bulblets, can make it difficult to identify the plant bulb's tip by shape alone. In an attempt to motivate research on this problem, a set of minimum criteria for such an algorithm have been published²². We specifically note that these criteria include (a) the planting of 10 plant bulbs per second, and (b) the planting of bulbs so that their growth directions are at least within 30 degrees of the upward vertical direction.

We hypothesize that multiple visual and morphological features are required from plant bulbs in order to estimate their growth directions in a way that satisfies the above criteria. Plant bulbs are known to have a complex internal structure that includes a root base, a central shaft from which the bulb sprouts, and shell-shaped scale surrounding the shaft that provide the plant with nutrients¹⁵. These internal structures provide a wealth of information on the bulb's growth direction and we aim to incorporate these structures into a growth direction estimation algorithm. X-ray imaging has the ability to visualize these information-rich internal structures, making it an appealing imaging technique for the growth direction estimation task. X-ray imaging has also been applied to similar agricultural tasks from seeding²³ to food quality assessment^{24–26}, suggesting that it can be safely and effectively used for our task of plant bulb growth direction finding.

In this paper, we introduce – to our knowledge – the first algorithm for the estimation of plant bulb growth directions. An overview of the algorithm is presented in Fig. 1. The algorithm takes as input 3 x-ray projections of a plant bulb from which shape²⁷, edge^{28,29}, and texture features^{30,31} of the bulb are extracted. These features are then used as input to machine learning regression algorithms^{32–34} in order to estimate the 2D growth direction of each bulb in each x-ray projection. These 2D estimates of the bulb's growth direction, one from each of the three x-ray projections, are then paired, and by taking into account the x-ray source and detector geometry, three 3D estimates of the bulb's growth direction are computed³⁵. The mean of these 3D estimates is then used as the final growth direction estimate, while the variance of these estimates is used to gauge the algorithm's confidence in the final result³⁶. Estimates with low confidence are ignored and the corresponding plant bulb is imaged a second time.

Our algorithm shows similarities to others seen in x-ray luggage screening and defect detection in manufactured parts³⁷. In both contexts, object features are extracted from multiple x-ray images and are used to classify the image as either passing or failing inspection^{38–42}. The choice of image features is often dependent on the application, though a thorough review of image feature detectors notes that detectors that incorporate orientation and scale perform particularly well in general⁴³. Meanwhile, the choice of classifier is often dependent on the complexity of the classification and on the amount of data available to train the algorithm. In x-ray screening tasks, popular choices have included thresholding⁴², support vector machines^{39,40}, and artificial neural networks⁴¹. In some cases, image features are used to identify outliers and the presence of a single outlier leads to a failed x-ray inspection³⁸. Our research problem differs from the ones in these works as we do not have a classification problem but a regression problem: we aim to predict the growth direction from x-ray image features.

We evaluate the proposed algorithm on 81,600 x-ray projections randomly simulated from computed tomography (CT) scans of 68 *T. Apeldoorn* flower bulbs. These x-ray simulations are combined into groups of three to produce 27,200 inputs to our algorithm (400 per flower bulb). The algorithm's training and testing is then performed in a leave-one-plant-bulb-out fashion. The results of these tests are then compared to industry-established standards for this growth direction finding task²².

To generate the simulated x-ray projections, we assume they are obtained from the scenario shown in Fig. 2, one which is similar to that described by Thompson *et al.*⁴⁴. Effectively, we envision each plant bulb being transported along a conveyor belt in a uniformly-likely random pose. As the bulb travels along the belt, it will be

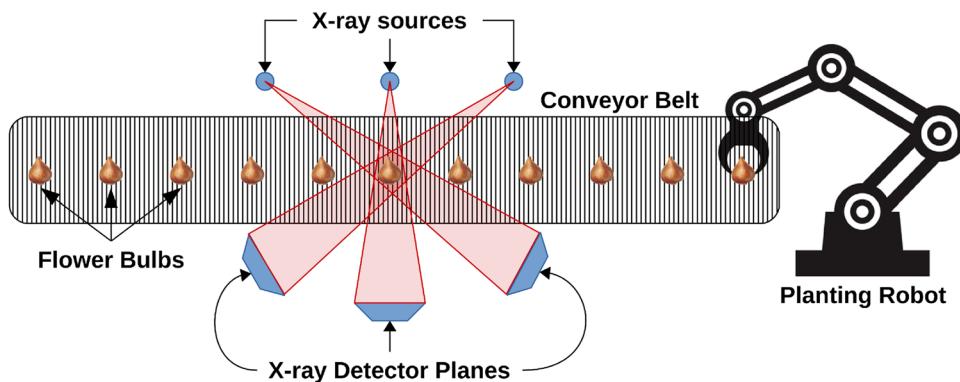


Figure 2. Schematic of the imaging setup simulated in this study: flower bulbs pass along a conveyor belt in the path of three x-ray projection systems before being planted by a robot. The x-ray systems are separated by 60° in order to give unique views of the same bulb.

imaged by three x-ray projection systems positioned at the sides of the belt, at the same height, and fanned out at angles 60° apart from each other. At the end of the conveyor belt will be a gripping robot that will grab the bulb, rotate it into position, and properly plant it in a growing tray^{9,22}. Our objective is to estimate the necessary rotation of the plant bulbs from the x-ray projections while the plant bulbs travel along the conveyor belt.

Results

Data collection & system implementation. Sixty-eight bulbs of the *T. Apeldoorn* flower were scanned in sets of 9 bulbs using the large field of view Hector CT scanner at Ghent University⁴⁵. The Hector CT scanner is equipped with a 240 kV X-ray tube (X-RAY WorX, Garbsen, Germany) and a $40 \times 40 \text{ cm}^2$ flat panel detector (PerkinElmer 1620, Waltham, Mass., USA) that are jointly combined on a high-precision rotation stage. The resolution of the resulting CT scans was $0.35 \times 0.35 \times 0.35 \text{ mm}^3$. During scanning, the plant bulbs were placed into styrofoam holders in order to keep them separated and stabilized during the imaging procedure. Each of the 68 flower bulbs were then manually segmented from the styrofoam background in the CT images using a combination of thresholding and morphological operations. Once segmented, a 3D mesh was generated by applying a contour filter to the bulb's segmented CT volume, and an expert manually annotated a vector along the growth direction twice for each bulb. The mean intra-expert annotation error was found to be 1.07° .

For each plant bulb, x-ray projections were generated by reprojecting the CT images using the ASTRA toolbox⁴⁶. The corresponding 3D annotated growth vector was also projected onto each generated x-ray image. The x-ray projections were generated to match the simulation environment shown in Fig. 2 with the bulb placed in a uniformly-chosen random orientation. Kernel support vector regression (SVR)³², kernel extreme learning machines (ELM)³⁴, and a 3-layer fully-connected artificial neural network (ANN)³³ were evaluated for the growth direction prediction. The proposed algorithm was implemented in MATLAB, version 2018a (The MathWorks, Natick, USA) running on a CentOS 7 operating system. All experiments were run on a desktop workstation with 16 Intel i7-5960X processor cores running at 3.00 GHz. We initially show results from the 3 machine learning algorithms' predictions of the 2D growth directions, one per digitally-reprojected radiograph. Subsequently, we display 3D growth direction results for the algorithm as a whole. Due to the randomness in the ELM algorithm, we show its aggregate results from 20 runs.

Machine learning regression. Figure 3a shows the histograms of prediction errors for the estimation of the 2D growth angles. As expected, the prediction errors cluster around zero, though the ANN and kernel ELM algorithms generally produced lower errors than kernel SVR. Kernel ELM produced a larger percentage of estimates with less than 5° error, but its remaining errors were more spread out than those of the ANN. This effect is most notable in the error range between $5\text{--}15^\circ$ where the ANN results outnumber kernel ELM. Figure 3b further shows a scatter plot of the prediction errors with respect to the ground truth 2D growth angles. Note that the prediction errors appear to be uniformly spread across all ground truth angles for all three algorithms. This suggests that the algorithms are unbiased to the ground truth growth angles. We note that the prediction errors do vary with respect to the flower bulb being tested, and that a small number of flower bulbs show significantly higher prediction errors than the majority of the dataset (see Supplementary Fig. S1).

Full simulation. Figure 4a shows, in colour, the percentage of flower bulb simulations that received a successful growth direction estimate (error $< 30^\circ$), and in grey, the percentage that received a poor growth direction estimate (error $\geq 30^\circ$). Results are shown cumulatively based on the number of times a bulb was run through the simulation. As expected from the 2D results, the ANN slightly outperformed the kernel ELM, while both algorithms significantly outperformed kernel SVR. After a maximum of three passes through the simulation, 98.40% of cases received a successful growth direction estimate from the ANN, while kernel ELM and kernel SVR achieved success rates of 96.76% and 60.68%, respectively.

Figure 4b further shows the histogram of estimation errors for the growth directions in the cases where an estimate was outputted by the system. Note that the errors have increased over those estimated in 2D (see Fig. 3a),

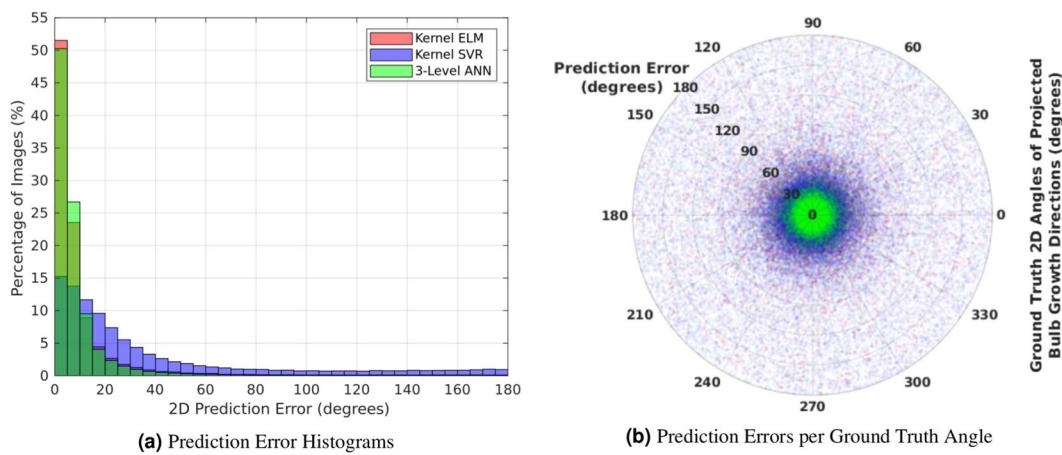


Figure 3. Prediction errors for the machine learning estimates of the 2D projections of the plant bulb growth directions, shown in both **(a)** histogram and **(b)** scatter plot form. The Kernel ELM and ANN algorithms outperform Kernel SVR, and the errors of all 3 algorithms seem independent of the ground truth angles. This suggests that the algorithms are working in an unbiased manner.

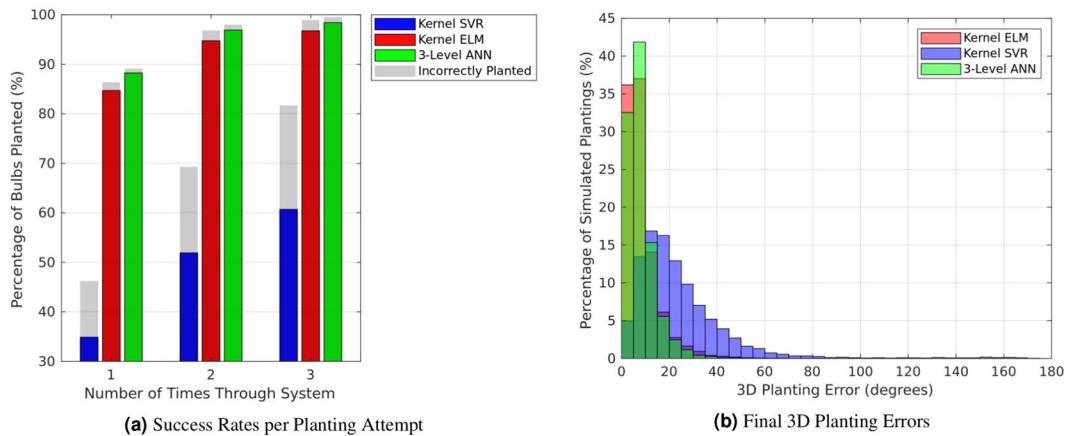


Figure 4. Final 3D growth direction planting errors for our simulation study, shown in both **(a)** success rate and **(b)** histogram form. After a maximum of 3 attempts to determine the flower bulb's growth direction, 98.4% of the simulated flower bulbs can obtain a successful growth direction estimate (error $< 30^\circ$) using the proposed 3-layer ANN.

but that the general trends persist: the kernel ELM produces accurate results (error $< 5^\circ$) for more cases than the other algorithms, but the ANN results generally cluster more tightly around zero. Both algorithms continue to outperform kernel SVR. We further observed that, for 66 of the 68 flower bulbs, the system provided a successful estimate over 90% of the time using the ANN. The remaining two bulbs, while obtaining relatively worse results, still obtained a successful estimate over 80% of the time (see Supplementary Fig. S2).

Finally, Fig. 5a shows the histogram of computation times for each of the 27,200 simulations performed in this study. In all cases and for all machine learning algorithms, an estimate was obtained in less than 0.1 seconds. The timings for each part of the system - the feature selection steps (shape, edge, and texture), the machine learning regressors, and everything else combined - is displayed in Fig. 5b. We observed that the edge features were the most computationally expensive part of the system, and that the ANN was the most computationally expensive machine learning algorithm. Nevertheless, all parts of the system are sufficiently fast to meet industry standards²².

Discussion

We have presented herein a machine learning framework for plant bulb growth direction estimation. The framework is based on shape, edge, and texture features collected from three non-collinear x-ray projections of the plant bulb. Using machine learning regressors and the geometric relationships between the x-ray projectors, we are able to obtain a 3D estimate of the plant bulb's growth direction from a triplet of 2D x-ray images. One of the key aspects of the proposed framework is the filtering of 3D estimates that do not internally agree with each other. This filtering stage was introduced to account for the fact that certain x-ray projections may have been acquired from an angle where the plant bulb's shape, edge, and texture features are uninformative. Figure 6b

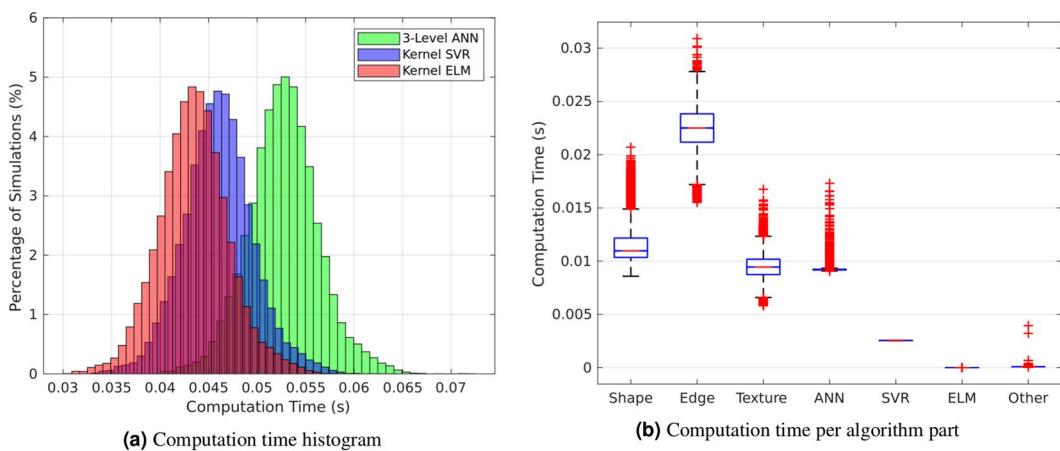


Figure 5. Computation time results for the growth direction estimation algorithm, shown both (a) cumulatively and (b) per step in the framework. Times are reported as the number of seconds per flower bulb. Note that each flower bulb receives an estimate within less than 0.1 seconds regardless of machine learning algorithm used. This allows us to achieve the industry criteria of 10 estimates per second in each case.

shows examples of such cases; they typically occur when the normal of the projection plane points in roughly the same direction as the bulb's growth direction. It is also for this reason that we chose not to have the machine learning algorithms estimate the 3D growth angle directly from the three x-ray projections at once: Since any one of the three projections could be uninformative, the algorithms would have the additional challenge determining which of the projection images to ignore.

With respect to the three machine learning regression algorithms we evaluated, we obtained the best overall results with the 3-layer ANN with kernel ELM running a close second. This outcome may be due to two factors. First, the kernel ELM sets some of its network's weights randomly as opposed to optimizing them. While this usually improves generalizability, it can also inflate errors due to the decreased flexibility in the regression function. Second, the additional layers in the ANN provide a function composition effect that can produce more powerful regressors than the single hidden layer network used in kernel ELM^{47,48}. Finally, both kernel ELM and the ANN far outperformed kernel SVR, suggesting that this regression problem is a non-linear one even after applying a Gaussian kernel to the inputted image features.

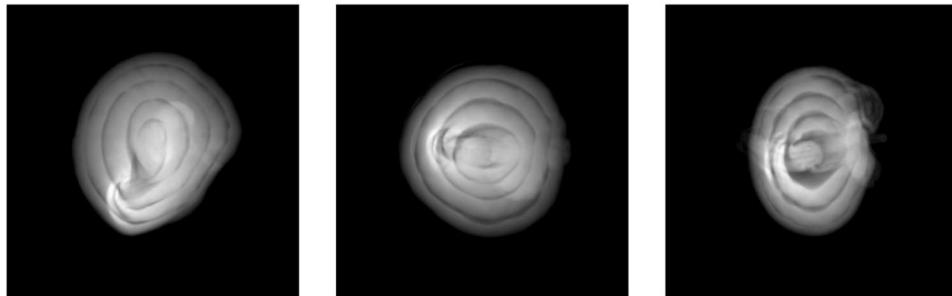
We further noticed that the quality of our algorithm's growth direction estimates depended on the plant bulb being imaged. As a result, we qualitatively examined the 3 plant bulbs that had the most successful growth direction estimates and the 3 bulbs with the least number of successful growth direction estimates. Those plant bulbs are shown in Fig. 6a,c, respectively. For the most successful bulbs, we note that the tip of the bulb is clearly visible and can be easily captured by the shape features. Also, the bulb is rather symmetric and has good contrast between its shells and the space between them, allowing for our edge features to easily identify the bulb's central shaft. This is not the case for the less successful bulbs, where a variety of visual challenges appear. These challenges include a bent central shaft (all three images in Fig. 6c), low contrast between the bulb's shells (right image in Fig. 6c), and asymmetries in the bulb's shape, in some cases due to the presence of bulblets along the side of the bulb (centre image in Fig. 6c, lower left corner of the bulb). These cases were rare in our database and it is possible that including additional bulbs like these in the training set could improve the performance of the proposed framework. Training and testing with more plant bulbs is something we intend to pursue as future work.

Overall, the proposed framework was able to achieve the industry-desired timing criteria (0.1 seconds) in all cases, and achieved the accuracy criteria (error < 30°) in up to 98.40% of cases using the ANN. These simulation results on 27,200 cases indicate that the proposed algorithm is reliable enough to proceed with evaluation as part of a physical robotic bulb planting prototype. This move from a simulation to a physical prototype may introduce additional challenges such as the inclusion of the conveyor belt within the image and additional background noise. We hypothesize that these additional challenges might be overcome with the addition of standard image processing techniques: the conveyor belt could be subtracted from the image using a template image acquired in the absence of a plant bulb⁴⁹, and additional noise could be addressed using image filtering techniques⁵⁰. These additional computational steps may increase the overall computation time. That being said, the use of unoptimized MATLAB code in this simulation study indicates that opportunities remain to improve the speed of the current algorithm in order to accommodate these additional image processing steps.

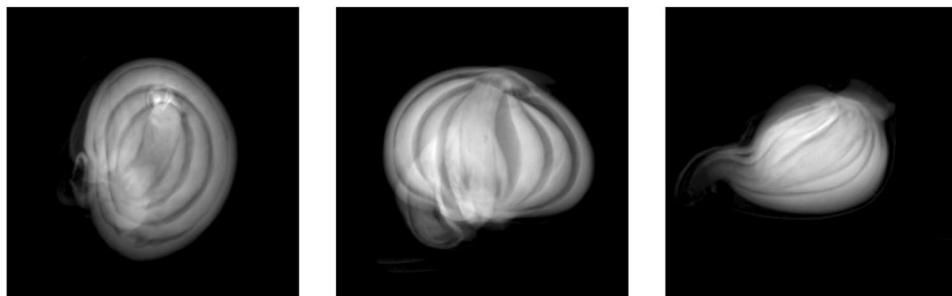
In conclusion, we have proposed that the automated planting of bulbous plants introduces the unique computer vision task of growth direction finding. To address this task, we have introduced the first algorithm for automated plant bulb growth direction finding. The algorithm makes use of machine learning regressors and visual features from x-ray projections of the bulbs in order to estimate the bulbs' projected 2D growth directions. These 2D estimates are then combined from image pairs and mapped to the 3D world coordinate system. Finally, estimates from multiple image pairs are compared to each other in order to determine the quality of the estimate, with poor quality estimates being discarded. Our results on a simulation of 27,200 cases resulted in successful estimates 98.40% of the time, suggesting that it is worthwhile to extend this algorithm to testing in a physical prototype.



(a) Flower bulbs with the best estimates



(b) X-ray projections nearly aligned with the growth direction



(c) Flower bulbs with the worst estimates

Figure 6. Examples of x-ray projection images from flower bulbs with both (a) good and (b,c) bad growth direction predictions from our framework. Note that the framework performs well when the bulbs are reasonably symmetric with clear shells and a noticeable tip. Worse results are seen when the bulbs have a curved stem, low contrast between shells, or have growths on the outside of the bulbs. Additionally, x-ray projection images where the normal of the projection plane is nearly in line with the growth direction of the flower bulb also proved challenging for the framework. See text for further details.

Methods

An overview of our growth direction estimation algorithm was shown in Fig. 1 and consists of four main components: feature selection, 2D angle prediction, the 2D-to-3D mapping, and the filtering out of plant bulbs whose estimates disagree with each other. Each of these algorithmic components are presented in detail below.

Feature selection. The first step in the growth direction estimation algorithm is to extract a small set of features from the x-ray projection images that can simplify the regression step between the images and the corresponding growth directions. We identify three such features of interest below.

Shape features. The shape of the plant bulb is characterized by comparing it to the shape of an ellipse of the same area, positioned at the same location and orientation in the image as the bulb itself. In this way, we highlight the non-elliptical features of the bulb shape, particularly the bulb's tip, for the growth direction estimation.

We obtain the shape of the plant bulb by thresholding the image at the empirically-chosen threshold of 5% of the image's maximum intensity, resulting in a binary segmentation S_{bulb} . From this resulting segmentation, a principal component analysis is performed on the pixel locations within the plant bulb (i.e. $\{(x, y) | S_{bulb}(x, y) = 1\}$) in order to obtain the bulb's centroid (x_c, y_c) and major/minor axes of variation. These pixel locations are then

projected onto the two axes of variation to obtain the length and width of the bulb (a , b). This information is then used to define a comparable ellipse segmentation $S_{ellipse}$,

$$S_{ellipse}(x, y) = \begin{cases} 1 & \text{if } \frac{(x \cos(\alpha) + y \sin(\alpha) - x_c)^2}{a^2} + \frac{(x \sin(\alpha) - y \cos(\alpha) - y_c)^2}{b^2} \leq 1, \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

This ellipse segmentation is subtracted from the plant bulb segmentation in order to emphasize the non-elliptical elements of the plant bulb's shape: $S_{diff} = S - S_{ellipse}$.

Additionally, we desire to reduce the dimensionality of the shape information – and make it invariant to the bulb's position in the image – in order to aid the subsequent neural network regression. To achieve this goal, we bin the resulting segmentation differences into an angular histogram. The angular histogram h_{shape} is defined as being centred at the centroid of the plant bulb segmentation (x_c , y_c) and has each bin covering 10° . Finally, to make the histogram invariant to plant bulb size, we statistically normalize the histogram elements by their z-scores. The resulting normalized histogram h_{shape} , containing 36 entries, is then taken as the final set of shape features for the x-ray projection image. A summary of the shape feature definition is shown in Fig. 7a.

Edge features. Another notable visual feature in x-ray projections of plant bulbs is the pattern of shell-shaped structures that compose the majority of the bulb's mass. These structures are denoted as scales and contain nutrients for the plant¹⁵. These scales surround the shoot and the bulb's growth direction, which makes them a valuable image feature for the estimation of the bulb's growth direction.

To capture this visual feature, we apply the vesselness filter of Frangi *et al.* to the given x-ray projection image²⁸. Originally designed to highlight blood vessels in angiogram images, the Frangi vesselness filter works to highlight layered or tubular structures, structures similar in appearance to the shells in our x-ray projection images. The Frangi vesselness filter produces an edge image, V , defined as

$$V(x, y) = \begin{cases} \exp\left[-\frac{\lambda_2^2}{2\lambda_1^2\beta^2}\right]\left[1 - \exp\left[-\frac{\lambda_1^2\lambda_2^2}{2c^2}\right]\right] & \text{if } \lambda_1 < 0 \\ 0 & \text{otherwise} \end{cases}, \quad (2)$$

where $\beta = 0.5$ and $c = 15$ control the influence of the “blobness” and “structure” terms respectively. Both terms are based on the eigenvalues, $|\lambda_1| \leq |\lambda_2|$, of the image's Hessian matrix. The blobness term encourages one eigenvalue to be larger than the other, thereby highlighting areas where the image gradient is strong only in one direction. The structure term encourages the pair of eigenvalues to be large, thereby highlighting only areas where we see clear intensity changes. We ignore cases where $\lambda_1 > 0$ so as to only highlight high-intensity structures, which is how plant bulb shells appear in x-ray projection images. To capture shells of different thicknesses, we apply the Frangi filter at multiple scales by convolving the image with Gaussian filters of $\sigma = \{3, 5, 7\}$ prior to computing the Hessian^{28,29}. The scale that produces the greatest filtered response is then retained in the edge image V .

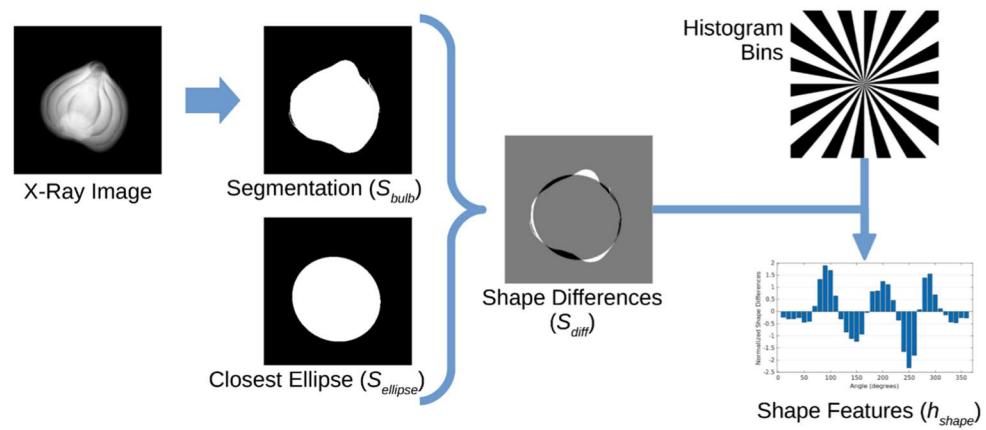
As with the shape features, we summarize this edge information in a statistically-normalized angular histogram. The resulting normalized histogram, h_{edge} contains 36 elements and is taken to be our final set of edge features. A summary of this feature detection technique is shown in Fig. 7b.

Texture features. While the previous feature detectors identify the shape and edges of the plant bulb, they fail to capture more complicated visual textures within an image. It has been shown that adding image texture features to a machine learning task can improve its performance⁵¹ and so we also follow that convention here. To capture the texture features of the plant bulbs, we employ local binary patterns (LBP)³¹.

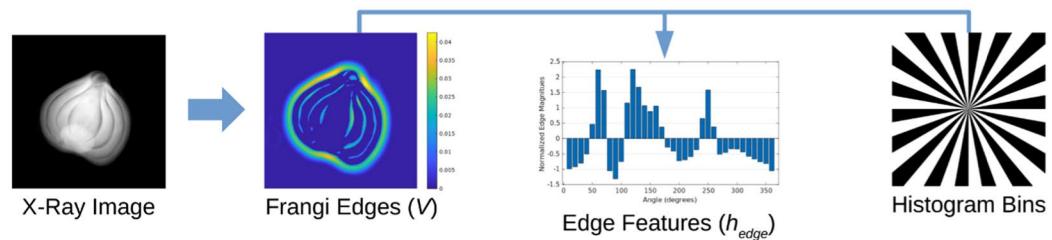
As the name suggests, LBP works at the level of individual pixels and generates a binary feature vector by comparing each pixel's intensity to its neighbours. A pixel's set of neighbours can be defined in many ways, but here we use the circular neighbourhood definition of Ojala *et al.*³⁰. Given a user-chosen radius r and number of neighbours k , a pixel's neighbours are defined as a set of k pixels uniformly spread around a circle of radius r centred at the chosen pixel (Fig. 7c). LBP then proceeds by comparing the intensity of a pixel to each of its neighbours. Wherever the intensity of the neighbour pixel is larger, a binary 1 is recorded for that neighbour. Otherwise, the neighbour is given the label 0. The result of these pixel comparisons are k binary digits which, when ordered in a clockwise manner, gives a k -digit binary number. Each binary number represents a different local pattern or texture. A summary of LBP is provided in Fig. 7c.

The LBP algorithm provides us with a binary number at each pixel, numbers which are then converted to decimal and binned into a histogram $h_{texture}$. In this work, we empirically set $r = 2$ and $k = 6$, which combined with the use of a global uniform pattern bin³⁰, results in a histogram of 33 entries.

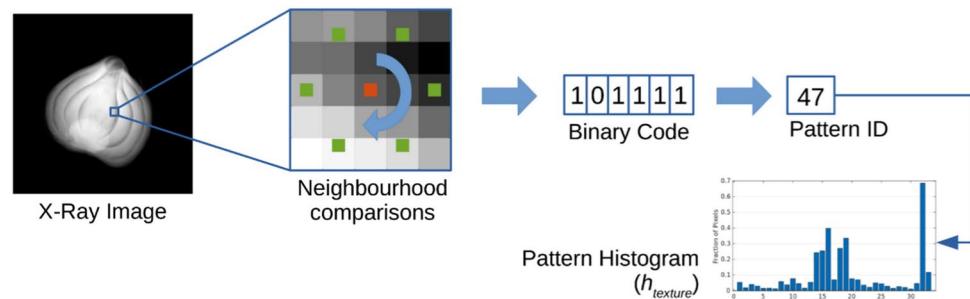
Machine learning regression. Using the combination of shape, edge, and texture features described earlier, we predict the 2D angle associated with the plant bulb's growth direction in the x-ray projection image. For this purpose, we evaluate three machine learning regression algorithms: kernel support vector regression (SVR)³², kernel extreme learning machines (ELM)³⁴, and a 3-layer fully-connected artificial neural network (ANN)³³. Gaussian kernels were used for the SVR and ELM algorithms with a kernel width of 500 being selected via a line search. The details of our ANN are presented in Table 1.



(a) Shape features are collected by comparing the flower bulb to an elliptical approximation. The difference in these two shapes are then grouped by angle and normalized to produce a shape feature set.



(b) Edge features are highlighted using a Frangi filter, then grouped by angle and normalized to produce a edge feature set.



(c) Texture features are highlighted using Local Binary Patterns.

Figure 7. Image features used for growth angle estimation include **(a)** shape, **(b)** edge, and **(c)** texture information. See text for further details.

A single input consisted of the 36 shape features h_{shape} , 36 edge features h_{edge} , and 33 texture features $h_{texture}$ described earlier. Each regressor's output is the 2D angle θ of the plant bulb's growth direction represented by its sine and cosine values. The sine and cosine of the angle are used to properly model the growth angle's periodicity⁵². The predicted growth angle is then retrieved as $\theta = \text{atan}2(\cos(\theta), \sin(\theta))$.

2D-3D geometric mapping. From the ANN, three 2D growth angles - θ_1 , θ_2 , θ_3 - are predicted, one for each of the x-ray projection images. At this point in the algorithm, it is unclear whether all three angle predictions are accurate. As a result, we choose to introduce redundancy into the algorithm by converting pairs of 2D growth angles into 3D estimates of the plant bulb growth direction. In this fashion, we obtain three 3D growth direction estimates, one for each pair of predicted 2D growth angles (θ_1 and θ_2 , θ_1 and θ_3 , θ_2 and θ_3).

To convert a pair of 2D growth angles into a single 3D estimate, we require information on the x-ray detector geometry. Let \mathbf{x}_i be the 3D world coordinate defining the origin of the detector plane for x-ray image I_i . Furthermore, let \mathbf{u}_i^x and \mathbf{u}_i^y be vectors defining, in 3D world coordinates, the x-axis and y-axis of that detector plane, respectively. Their cross product $\mathbf{n}_i = \mathbf{u}_i^x \times \mathbf{u}_i^y$ defines the normal of the detector plane in the direction of the x-ray source.

| Category | Item | Details |
|--------------|---------------------------|---|
| Structure | Input Layer | 1 input layer: 105 nodes (36 shape, 36 edge, 33 texture) |
| | Hidden Layer(s) | 3 hidden layers (60 nodes, 30 nodes, 15 nodes) |
| | Output Layer | 1 output layer: 2 nodes (sine & cosine of growth angle) |
| | Transfer Function | Hyperbolic Tangent |
| Optimization | Performance Function | Mean-Squared Error (with projection to unit circle) |
| | Backpropagation Algorithm | Scaled Conjugate Gradient (Hessian step size, σ : 5e-5) (Hessian regularization, λ : 5e-7) |
| | Maximum Number of Epochs | 1000 |
| | Gradient Tolerance | 1e-06 |

Table 1. Structure and parameters of the artificial neural network used to predict the 2D angle of a flower bulb growth direction from its corresponding x-ray image features.

Given this detector geometry and a 2D growth angle θ_i , we project outwards from the detector plane to identify the space to which the 3D plant bulb's growth direction can belong. First, we convert θ_i to a vector representation:

$$\mathbf{v}_{\theta,i} = \cos(\theta)\mathbf{u}_i^x + \sin(\theta)\mathbf{u}_i^y + \mathbf{x}_i. \quad (3)$$

In this formulation, we define the vector as originating at \mathbf{x}_i , though any other point on the detector plane could be used since only the direction of the growth angle is of importance, not its exact location.

Since the 2D growth angle θ_i is parallel to $\mathbf{v}_{\theta,i}$, we know that the corresponding 3D growth angle must also be parallel to a plane spanned by $\mathbf{v}_{\theta,i}$ and the normal of the detector plane \mathbf{n}_i . We define this plane as a solution space $S(i)$ and parameterise it by its normal $\mathbf{n}_{S(i)} = \mathbf{v}_{\theta,i} \times \mathbf{n}_i$. This process can similarly be done for the other x-ray projection.

Given two 2D growth angles, θ_i and θ_j , we compute the normals of their solution spaces $\mathbf{n}_{S(i)}$ and $\mathbf{n}_{S(j)}$ respectively. We know that, since the 3D growth direction must lie both in $S(i)$ and $S(j)$, it must be orthogonal to both $\mathbf{n}_{S(i)}$ and $\mathbf{n}_{S(j)}$. Therefore, we define the 3D growth direction estimate, $(\theta_{i,j}, \phi_{i,j})$ as the azimuth and elevation angles defined by the cross product of these two normals:

$$\theta_{i,j} = \cos^{-1} \left(\frac{\mathbf{n}_{S(i),x}\mathbf{n}_{S(j),y} - \mathbf{n}_{S(i),y}\mathbf{n}_{S(j),x}}{\|\mathbf{n}_{S(i)} \times \mathbf{n}_{S(j)}\|_2} \right), \quad \phi_{i,j} = \tan^{-1} \left(\frac{\mathbf{n}_{S(i),z}\mathbf{n}_{S(j),x} - \mathbf{n}_{S(i),x}\mathbf{n}_{S(j),z}}{\mathbf{n}_{S(i),y}\mathbf{n}_{S(j),z} - \mathbf{n}_{S(i),z}\mathbf{n}_{S(j),y}} \right). \quad (4)$$

Filtering of poor estimates. Following the pairwise 3D growth direction estimation step, we have three growth direction estimates: $(\theta_{1,2}, \phi_{1,2})$, $(\theta_{1,3}, \phi_{1,3})$, and $(\theta_{2,3}, \phi_{2,3})$. We hypothesize that if these three estimates agree with each other, then they are likely to be accurate. To measure the agreement between these angles, we compute the angular distance between each pair of angles as

$$d(\theta_{i,j}, \phi_{i,j}, \theta_{a,b}, \phi_{a,b}) = \cos^{-1} [\sin(\theta_{i,j})\sin(\theta_{a,b}) + \cos(\theta_{i,j})\cos(\theta_{a,b})\cos(\phi_{i,j} - \phi_{a,b})] \quad (5)$$

These angular distances are then averaged across all three estimate pairs to obtain \bar{d} : a single measure of overall agreement between the growth direction estimates. A threshold τ is then applied to \bar{d} to determine if the estimates agree with each other. If $\bar{d} < \tau$, then the three growth directions are assumed to be accurate and the average of the three growth direction estimates is outputted; if $\bar{d} \geq \tau$, then the plant bulb is imaged again. In this work, we empirically set $\tau = 40^\circ$.

Received: 13 September 2019; Accepted: 20 December 2019;

Published online: 20 January 2020

References

- Duckett, T., Paerson, S., Blackmore, S. & Grieve, B. Agricultural robotics: The future of robotic agriculture. Tech. Rep., The United Kingdom Robotics and Autonomous Systems Network (UK-RAS) (2018).
- Pedersen, S. M., Fountas, S., Sorensen, C. G., Evert, F. K. V. & Blackmore, B. S. *Precision Agriculture: Technology and Economic Perspectives*, chap. Robotic Seeding: Economic Perspectives, 167–179 (Springer International, Cham, 2017).
- Roldán, J. J. et al. *Service Robots*, chap. Robots in Agriculture: State of Art and Practical Experiences (IntechOpen, 2017).
- Pekkeriet, E. J. & van Henten, E. J. Current developments of high-tech robotic and mechatronic systems in horticulture and challenges for the future. In Dorais, M. (ed.) *Processing of International Symposium on High Technology for Greenhouse Systems - GreenSys*, 85–94 (2009).
- Hu, J. et al. Dimensional synthesis and kinematics simulation of a high-speed plug seedling transplanting robot. *Comput. Electron. Agric.* **107**, 64–72 (2014).
- Iacomi, C. & Popescu, O. A new concept for seed precision planting. *Agric. Agric. Sci. Procedia* **6**, 38–43 (2015).
- Metha, P. Automation in agriculture: Agribot the next generation weed detection and herbicide sprayer - a review. *J. Basic Appl. Eng. Res.* **3**, 234–238 (2016).
- Shanshiri, R. R. et al. Research and development in agricultural robotics: A perspective of digital farming. *Int. J. Agric. Biol. Eng.* **11**, 1–14 (2018).

9. Rodríguez, F., Moreno, J. C., Sánchez, J. A. & Berenguel, M. *Grasping in Robotics*, vol. 10 of *Mechanisms and Machine Science*, chap. Grasping in Agriculture: State-of-the-Art and Main Characteristics (Springer-Verlag, London, 2013).
10. Luo, L. *et al.* Vision-based extraction of spatial information in grape clusters for harvesting robots. *Biosyst. Eng.* **151**, 90–104 (2016).
11. Qiang, L., Jianrong, C., Bin, L., Lie, D. & Yajing, Z. Identification of fruit and branch in natural scenes for citrus harvesting robot using machine vision and support vector machine. *Int. J. Agric. Biol. Eng.* **7**, 115–121 (2014).
12. Rong, X., Huanyu, J. & Yibin, Y. Recognition of clustered tomatoes based on binocular stereo vision. *Comput. Electron. Agric.* **106**, 75–90 (2014).
13. Abdelmotaleb, I., Hegazy, R., Imara, Z. & Rezk, A. E.-D. Developement of an autonomous navigation agricultural robotic plantform based on machine vision. *Misr Journal Agric. Eng.* **32**, 1421–1450 (2015).
14. Bechar, A. & Vigneault, C. Agricultural robots for field operations. part 2: Operations and systems. *Biosyst. Eng.* **153**, 110–128 (2017).
15. Hertogh, A. A. D., Aung, L. H. & Benschop, M. *Horticultural Reviews*, chap. The Tulip: Botany, Usage, Growth, and Development (Wiley, 2011).
16. Aksenov, A. G., Izmaylov, A. L., Dorokhov, A. S. & Sibirev, A. V. Onion bulbs orientation during aligned planting of seed-onion using vibration-pneumatic planting device. *INMATEH-Agricultural Eng.* **55** (2018).
17. Hanks, G. R. Variation in the growth and development of narcissus in relation to meteorological and related factors. *J. Hortic. Sci.* **71**, 517–532 (1996).
18. Shropshire, F. M. *et al.* Significance of bulb polarity in survival of transplanted mitigation bulbs. *Bull. South. California Acad. Sci.* **115**, 112–125 (2016).
19. Castellanos, J. Z. *et al.* Garlic productivity and profi tability as affected by seed clove size, planting density and planting method. *HortScience* **39**, 1272–1277 (2004).
20. Nourai, A. H. Effects of planting methods and seed rates on yield, yield components, and quality of garlic (*Allium sativum* L.) in the Sudan. In *Proceedings of International Symposium on Alliums for the Tropics* **358**, 359–364 (1993).
21. Nazari, F., Farahmand, H., Khosh-Khui, M. & Salehi, H. Effects of two planting methods on vegetative and reproductive characteristics of tuberose (*Polianthes tuberosa* L.). *Adv. Nat. Appl. Sci.* **1**, 26–29 (2007).
22. Mohr, C. Bulb planting automation: General specifications. Tech. Rep., Vineland Research & Innovation Centre (2017).
23. Blunk, S. *et al.* Quantification of differences in germination behaviour of pelleted and coated sugar beet seeds using x-ray computed tomography (x-ray CT). *Biomed. Phys. & Eng. Express* **3** (2017).
24. Haff, R. P. & Toyofuku, N. X-ray detection of defects and contaminants in the food industry. *Sensors Instrumentation for Food Qual. Saf.* **2**, 262–273 (2008).
25. Janssens, E. *et al.* Neural network based x-ray tomography for fast inspection of apples on a conveyor belt system. In *Proceedings of IEEE International Conference on Image Processing (ICIP)*, 917–921 (2015).
26. Renu, R. & Chidanand, D. V. Internal quality classification of agricultural produce using non-destructive image processing technologies (soft x-ray). *Int. J. Latest Trends Eng. Technol.* **2**, 535–543 (2013).
27. Yang, M., Kpalma, K. & Ronzin, J. A survey of shape feature extraction techniques. *Pattern Recognit.* **41**, 43–90 (2008).
28. Frangi, A. F., Niessen, W. J., Vincken, K. L. & Viergever, M. A. Multiscale vessel enhancement filtering. In *Proceedings of Medical Image Computing and Computer-Assisted Interventions (MICCAI)*, 130–137 (Springer, Berlin, Heidelberg, 1998).
29. Nand, K. K., Abugharbieh, R., Booth, B. G. & Hamarneh, G. Detecting structure in diffusion tensor MR images. In *Proceedings of Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 90–97 (Springer, Berlin, Heidelberg, 2011).
30. Ojala, T., Pietikäinen, M. & Mäenpää, T. Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE Transactions on Pattern Analysis Mach. Intell.* **24**, 971–987 (2002).
31. Pietikäinen, M., Hadid, A., Zhao, G. & Ahonen, T. *Computer Vision Using Local Binary Patterns* (Springer-Verlag, London, 2011).
32. Amari, S. & Wu, S. Improving support vector machine classifiers by modifying kernel functions. *Neural Networks* **12**, 783–789 (1999).
33. Bishop, C. M. *Pattern Recognition and Machine Learning* (Springer, 2006).
34. Huang, G.-B., Zhou, H., Ding, X. & Zhang, R. Extreme learning machine for regression and multiclass classification. *IEEE Transactions on Syst. Man, Cybern. Part B (Cybernetics)* **42**, 513–529 (2012).
35. Carlbom, I. & Paciorek, J. Planar geometric projections and viewing transformations. *ACM Comput. Surv.* **10**, 465–502 (1978).
36. Athans, M., Ku, R. & Gershwin, S. B. The uncertainty threshold principle: Some fundamental limitations of optimal decision making under dynamic uncertainty. *IEEE Transactions on Autom. Control* **22**, 491–495 (1977).
37. Mery, D. *Computer Vision for X-Ray Testing*, chap. Applications in x-ray testing, 267–325 (Springer, 2015).
38. Mery, D. Inspection of complex objects using multiple-x-ray views. *IEEE Transactions on Mechatronics* **20**, 338–347 (2015).
39. Franzel, T., Schmidt, U. & Roth, S. Object detection in multi-view x-ray images. In Pinz, A., Pock, T., Bischof, H. & Leberl, F. (eds.) *Pattern Recognition*, 144–154 (Springer, Berlin, Heidelberg, 2012).
40. Ramirez, F. & Allende, H. Detection of flaws in aluminium castings: a comparative study between generative and discriminant approaches. *Insight-Non-Destructive Test. Cond. Monit.* **55**, 366–371 (2013).
41. Akcay, S., Kundegorski, M. E., Willcocks, C. G. & Breckon, T. P. Using deep convolutional neural network architectures for object classification and detection within x-ray baggage security imagery. *IEEE Transactions on Inf. Forensics Secur.* **13**, 2203–2215 (2018).
42. Shen, J. *et al.* X-ray inspection of TSV defects with self-organizing map network and Otsu algorithm. *Microelectron. Reliab.* **67**, 129–134 (2016).
43. Mikolajczyk, K. & Schmid, C. A performance evaluation of local descriptors. *IEEE Transactions on Pattern Analysis Mach. Intell.* **27**, 1615–1630 (2005).
44. Thompson, W. M., Lionheart, W. R. B., Moron, E. J., Cunningham, M. & Luggar, R. D. High speed imaging of dynamic processes with a switched source x-ray CT system. *Meas. Sci. Technol.* **26** (2015).
45. Masschaele, B. *et al.* HECTOR: A 240kv micro-CT setup optimized for research. *J. Physics: Conf. Ser.* **463**, 012012 (2013).
46. Van Aarle, W. *et al.* Fast and flexible x-ray tomography using the ASTRA toolbox. *Opt. Express* **24**, 25129–25147 (2016).
47. Lin, H. W., Tegmark, M. & Rolnick, D. Why does deep and cheap learning work so well? *J. Stat. Phys.* **168**, 1223–1247 (2017).
48. Rolnick, D. & Tegmark, M. The power of deeper networks for expressing natural functions. In *Proceedings of 6th International Conference on Learning Representations (ICLR)*, 14 (2018).
49. Alard, C. & Lupton, R. A method for optimal image subtraction. *The Astrophys. J.* **503**, 325–331 (1998).
50. Buades, A., Coll, B. & Morel, J.-M. A review of image denoising algorithms, with a new one. *Multiscale Model. Simul.* **4**, 490–530 (2005).
51. Mirmehdi, M., Xie, X. & Suril, J. (eds.) *Handbook of Texture Analysis* (Imperial College Press, 2008).
52. Hara, K., Vemulapalli, R. & Chellappa, R. Designing deep convolutional neural networks for continuous object orientation estimation. *arXiv preprint 1702.01499*, 10 (2017).

Acknowledgements

This research is part of the imec ICON project no. HBC.2016.0164 and has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement no. 746614. The Ghent University Centre for X-ray Tomography (UGCT) is acknowledged for the acquisition of the CT data of the flower bulbs. The special research fund of the Ghent University (BOF-UGent) is acknowledged for the financial support of the UGCT Center of Expertise (BOF.EXP.2017.0007).

Author contributions

J.D.B. and J.S. conceived the project. B.G.B. conducted the experiments and wrote the manuscript. J.D.B., B.G.B. and J.S. all analyzed the results and reviewed the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information is available for this paper at <https://doi.org/10.1038/s41598-019-57405-8>.

Correspondence and requests for materials should be addressed to B.G.B.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2020