TrimBot2020 Deliverable D2.7

Final evaluation and dissemination

Principal Author: Wageningen University (WU) &
Wageningen Research (WR)
Contributors: A.-Ludwigs-Univ. Freiburg (ALUF)
University of Edinburgh (UEDIN)
Dissemination: PU

Abstract. This report describes the bush trimming and rose stem clipping evaluation with the final manipulator and the final tools. For evaluation of the bush trimming performance a total of 29 individual bushes (spherical, cylindrical and cuboid shaped) were autonomously trimmed by the robot. Trimming of one object was performed from multiple vehicle positions around the target. The vehicle positioning was evaluated using a motion capture system. To compare the bush shape before and after trimming in the 3D space a Structure-from-Motion photogrammetry method was applied. This allowed to reconstruct precise 3D models of the bushes. In general the robot was not able to trim the bushes with the required accuracy. In some cases the fitted target mesh produced by the robot was erroneous. Aligning the target shape of one bush from multiple vehicle positions turned out to be challenging. The percentage of correctly trimmed points did not exceed 60%, many areas were trimmed too deep or not trimmed at all. This was in agreement with the qualitative results. In the manual scoring each bush had a final score < 3.1 on a scale from 1 (very bad) to 5 (very good). Despite these, the results show that the robot was able to fully autonomously trim bushes. The experimental evaluation of the vision part for rose stem clipping showed that the neural network is capable of segmenting the stems from the background. Scanning the rose bush is a simple yet effective method to capture the 3D structure of the plant. For the rose stem clipping 93% of reached targets were fully cut. This corresponds to 78% overall cutting success rate. This is sufficient to trim the bush completely from several positions around the bush while allowing additional attempts at previously uncut sites.

Deliverable due: Month 48

Contents

1 Introduction

This report describes the bush trimming and rose stem clipping evaluation with the final manipulator and the final tools of the developed Trimbot. The final system is described in Deliverable 2.6. Separate descriptions of the evaluation procedure, the results, and discussion and conclusions are given for topiary bushes (Section 2) and rose bushes (Section 3). Some general remarks and recommendations after evaluation are given in Section 4.

2 Bush trimming evaluation

This section explains how the final bush trimming evaluation was executed. A demonstration video clip of topiary bush trimming is available on the Trimbot2020 YouTube channel [5]. First, a description is given of the materials (bushes, platform, evaluation equipment) and methods (system control, experimental outline, evaluation method, data analysis) used to assess the trimming performance of the Bush Trimming platform. The result section then presents detailed results on all the experiments that have been carried out for the evaluation. The bush trimming evaluation chapter is concluded with a discussion and conclusion section.

2.1 Material and Methods

2.1.1 Test set of bushes

For evaluation of the Trimbot's trimming performance, spherical, cylindrical and cuboid[1] bushes were used. All bushes were obtained directly from the grower. The spherical bushes were previously trimmed to sphere-shape, and contained outgrowth from the last growing seasons. Cylinder and cuboid bushes were not available as-is from the grower, and thus created from spherical bushes that were pre-trimmed by the grower to resemble more or less the desired shape. Prior to the trimming experiments, extreme outgrowth and shape inconsistency were reduced by some rough pre-trimming to have shapes that could also be recognized and trimmed properly. In total, 10 spheres, 9 cylinders and 10 cuboids were tested. An overview of the bushes tested (before trimming) is shown in Figure 1.

2.1.2 Experimental hardware

This section describes the hardware that was used in the experiment. The hardware used for obtaining the quantitative evaluation of the trimming performance is described in section 2.1.6.

Trimbot platform For evaluation of the trimming performance, the duplicated platform 3b (blue version) as described in Deliverable 1.3 was used, together with the arm and bush trimming tool as described in Deliverable 2.6. As during initial testing, the visual approach was not working and this also failed sometimes during demo preparation. It was decided for the trimming evaluation to make use of a Velodyne Puck Lidar [4] as backup, together with the Lidar-based mockup components as mentioned in Deliverable 6.4. For accessing and processing the camera images and controlling the manipulator, a Razer Blade laptop was installed on the platform.

[1]The original plan was to use cube shaped bushes. However, the bushes obtained from the grower were relatively high, and cutting them back to a cube would results in many strong wooden branches at the top of the bush, with a clear risk of damaging the cutter. Thus, it was decided not to consider the top of these bushes, and only trim the sides, thereby making them cuboids.

Figure 1: Overview of bushes used in the trimming evaluation, before trimming. Left the cylinders, in the middle spheres and on the right cuboids.

Optitrack System To evaluate the the accuracy of vehicle positioning with respect to the bush and to have the ability to relate the actual pose of the bush with data acquired by the platform, an Optitrack motion capture system [2] was used. This system consisted of 6 Prime 13 cameras [3] placed on tripods around the test area and connected to a dedicated computer running the Motive software [1]. After calibrating this setup, the platform's position and orientation could be obtained by using 8 reference markers attached to the mobile platform. Figure 3 shows a photo of the platform with some of the markers visible. Similarly, the position of the bush was obtained by using 6 markers on the ring that was placed around the bush. An overview of the Optitrack setup in the garden is shown in Figure 2.

System control and logging Remote control of the Trimbot platform and its state machines was done via a dedicated laptop, which also stored the processing results, such as the images and arm poses from scanning and the calculated mesh and trimming trajectories. A second laptop was used to acquire data from the Optitrack system and register all experimental information, vehicle poses and general observations in an Excel-sheet to ensure all information for post-experiment evaluation. Furthermore, a GoPro Hero 7 was used to capture the experimental process.

2.1.3 Experimental settings

Trimming was performed on 4 poses around the bush, followed by a 5th pose where the platform was placed closer for trimming the top of the bush. In the experiment, the following settings were used for the hardware:

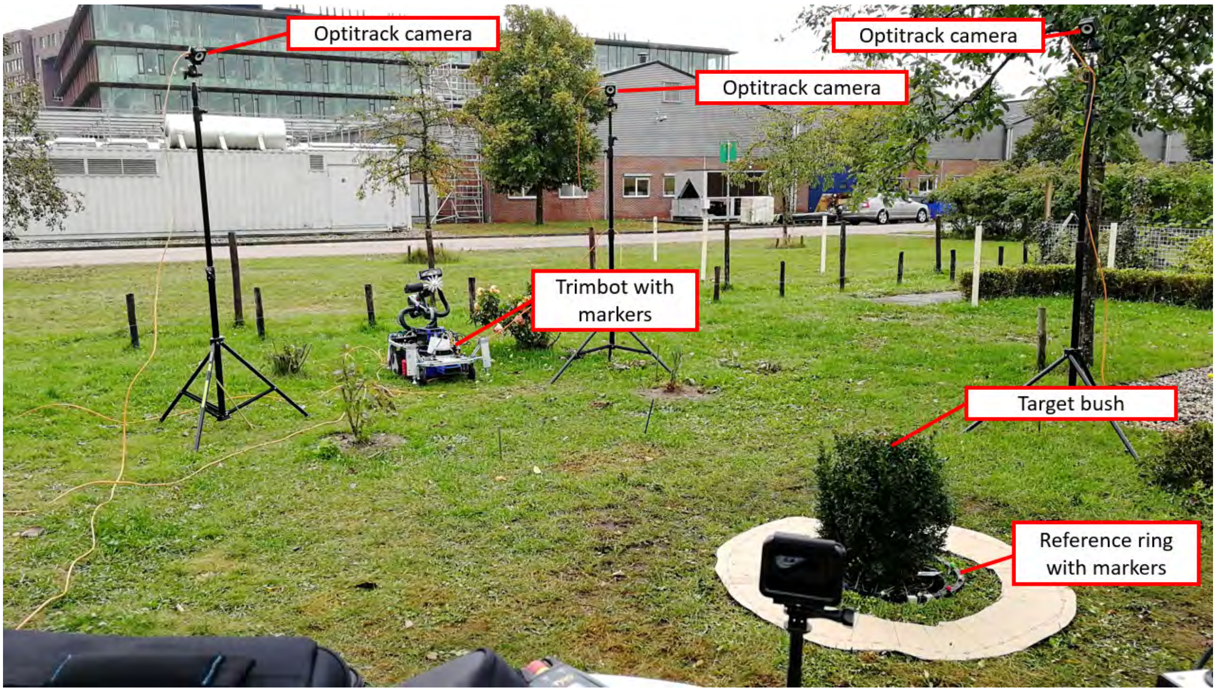

Figure 2: The Optitrack cameras are mounted on the tripods. This system was used for the evaluation to determine the position of the platform and bush.

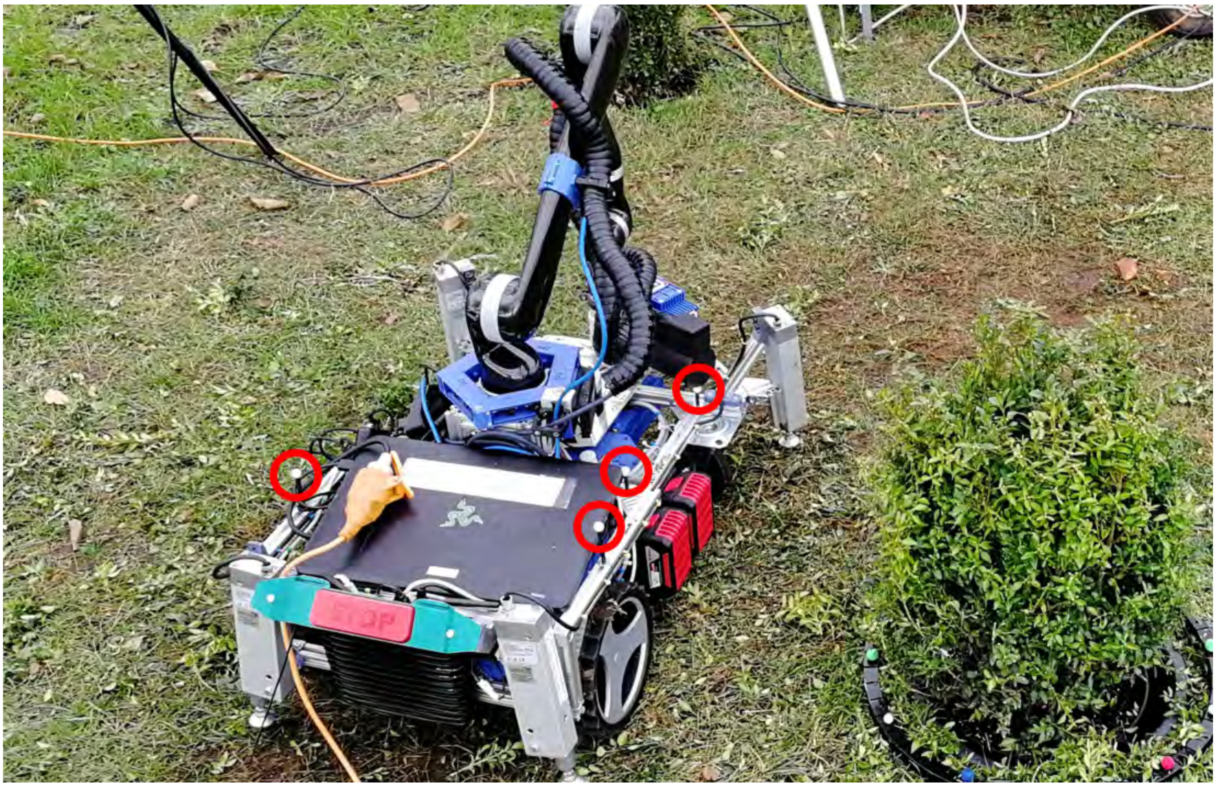

Figure 3: The platform with mounted Optitrack markers (highlighted by red circles) for measuring the platform pose.

- Target distance for approaching the bush was 0.65m between arm base and closest point of the bush point cloud.
- Joint speed for the manipulator was set to $0.1\pi\ s^{-1}$ during cutting and $0.2\pi\ s^{-1}$ during arm transitions.
- End-effector motor speed was set to 3000 rpm, which equals 20 rpm of the blade.

For the shape fitting algorithm by the robot, the following settings were used:

- For fitting spheres, a fixed diameter of 34 cm was used.
- For fitting cylinders, a diameter of 33 cm was used, and a height of 49.5 cm (a ratio of 3/2 for height/diameter).
- For the cuboids, a cube was fitted with an edge size of 30 cm (except for bush 4 -3641 and 5 - 3643, which used 34cm).
- For calculating the trimming path, the mesh diameter was shrunken by 2 cm (and height accordingly). The diameter of the shrunken fitted shape was therefore, 30 and 31 cm for the spheres and cylinders respectively. The cuboids had an edge size of 28 cm. Fitting was supposed to be done on a bush with outgrowth, while the planning should result in a trimmed bush, i.e. one shrunken with respect to the outgrowth.
- In the Iterative Closest Point (ICP) algorithm, the cutoff distance was set to 0.002 for spheres and 0.005 for other shapes.

For the coverage planning algorithm, the following settings were used

- Triangles grouping threshold of 0.5 for sphere sides and 0.3 for cuboid, cylinder, and sphere top. This dimensionless value is an upper limit for the fraction of the tool diameter length that is used to merge neighboring mesh triangles into a single patch. More details about the purpose of this quantity are in [6].
- Minimal arm reach was set to 0.7m, extended to computed distance from target center for more distant plants.

During the experiment, the following information was automatically stored:

- Images from arm cameras and arm configurations during scanning of the bush.
- The resulting fitted shape (point cloud).
- The target poses of the Trimbot platform.
- The planned and executed trimming trajectory.
- Results of evaluation actions during trimming.

During the experiments, the only setting that was changed between bushes was whether the trimming trajectory was repeated or not. For this, three options were used:

- Never repeat the trimming trajectory.
- Only repeat the trimming trajectory if the trimming error associated to the bush points affected by the cutting action was evaluated as too high (see Deliverable 2.6 for details about trimming result computation).
- Always repeat the trimming trajectory.

For each combination of repetition setting and bush shape, at least three bushes were trimmed.

2.1.4 Experimental outline

At the start of each experimental day, all systems were placed in the garden: the Trimbot platform, the Optitrack system (including a calibration), the GoPro camera and the control laptops. Then, for each bush to be evaluated, the following steps were taken:

1. Execute manual pre-trimming, such that the bush shape is in roughly in agreement with the target shaped and outgrowth larger than ± 10 cm is removed before starting robot approach.
2. Place the bush at the evaluation spot. For cuboid bushes, make sure the orientation is correct, so the arm will start trimming at an edge of the bush.
3. Collect pre-trimming data for quantitative evaluation as described in section 2.1.6.
4. Measure and register dimensions (tape measure) and position of the bush (position of the reference ring measured with Optitrack).
5. Place the Trimbot platform at its starting pose, about 2m away from the bush.
6. Execute bush approaching manoeuvre and servo to the first platform trimming position (described in section 2.1.5).
7. Check the platform pose with respect to the bush (using a tape measure), and manually adjust the distance (and if necessary also orientation) if the platform is beyond the desired distance (0.7-0.8m between arm and bush center). Register the changes made and the final pose of the platform.
8. Execute trimming procedure according to the approach described in section 6 of Deliverable 2.6. At each new pose of the platform, check its position with respect to the bush according to step 7.
9. Collect post-trimming data for quantitative evaluation as described in section 2.1.6.
10. Store the bush for qualitative evaluation afterwards.

2.1.5 Target bush approach and visual servoing

Although the approaching of the target object and the positioning of the platform in the vicinity of the bush is not a core part of this work package, this information is still relevant to assess the actual trimming performance. As input for this evaluation, mainly direct observations and manual measurement of vehicle pose were uses, while Optitrack data was available but not used.

2.1.6 Quantitative evaluation of trimming result

Evaluation rig

The methodology used to evaluate the manipulator and tools performance in Deliverable 2.3 was based on a sequence of 2D silhouettes -from colour images- to compare the bush profiles before and after trimming. In this deliverable, the evaluation methodology was improved by using Structure-from-Motion (SfM) photogrammetry to reconstruct precise 3D models of bushes. With that, the comparison between the bush shape before and after trimming is carried out in the 3D space.

The photographic equipment used for data acquisition consisted of a single lens reflex (SLR) camera (Nikon Z6, Nikon Corp., Tokyo, Japan), with an Nikon Z 24-70 F/2.8 lens. The camera settings were configured in aperture-priority mode (± 11), ISO200, and focal length of 35mm. For each test bush a total of 250 images were taken before and after trimming. Images were taken from a horizontal distance of 0.5 m between the bush and cameras centres, and from 5 different heights: 20 cm, 30 cm, 40 cm, 50cm, and 60cm (Figure 4a). For each height, 50 images were acquired around the bush, corresponding to a 7.2 degrees rotation between two consecutive camera positions (Figure 4b).

To generate the 3D model of the test bushes from the acquired images, a multi-view SfM photogrammetry technique was employed by using the Agisoft Metashape software (v1.5, Agisoft LLC, St. Petersburg, Russia). Since the 3D reconstruction based on SfM is scale invariant, a set of known markers (depicted in Figure 4d) were used to scale the 3D model to a real-world scale. 6 of these markers were reflective balls, allowing to register the SfM model with the robot 3D vision system under the same coordinate system. This methodology was able to estimate the position of reference markers with less than a 5mm error (generally around 3 mm), of which the largest part originates from the manual placement of the markers in the image. Thus, the generated SfM 3D models were used to evaluate the bush shapes before and after trimming, with the errors in the point cloud being less than 1 mm.

Evaluation software

The software to evaluate the robot arm vision system and the bush trimming operation was developed in MATLAB. The software compares the bush shape before and after trimming, as well as the fitted shape defined by the robot arm vision system with respect to the actual target shape.

Figure 5 illustrates the evaluation software pipeline. The developed software processes 4 input files: the point cloud obtained with SfM photogrammetry before trimming the bush (3D model before trimming); the point cloud obtained with SfM photogrammetry after trimming the bush (3D model after trimming). the final poses of the robot arm (robot final poses) and the shapes estimated by the robot from the 5 trimming positions (robot fitted shapes).

From the point clouds obtained with SfM, the evaluation software detected automatically the floor plane, removed the points from the floor, and estimated the bush shape by fitting a target shape. The target shape can be a sphere, cylinder or cuboid. Each shape is fitted using the random sample consensus (RANSAC) algorithm (Figure 6b). This target shape fitting is only done on the non-trimmed point cloud. The Euclidean distance between each point and the target shape (δ_i) was computed (Figure 6c). All points presenting an Euclidean distance of

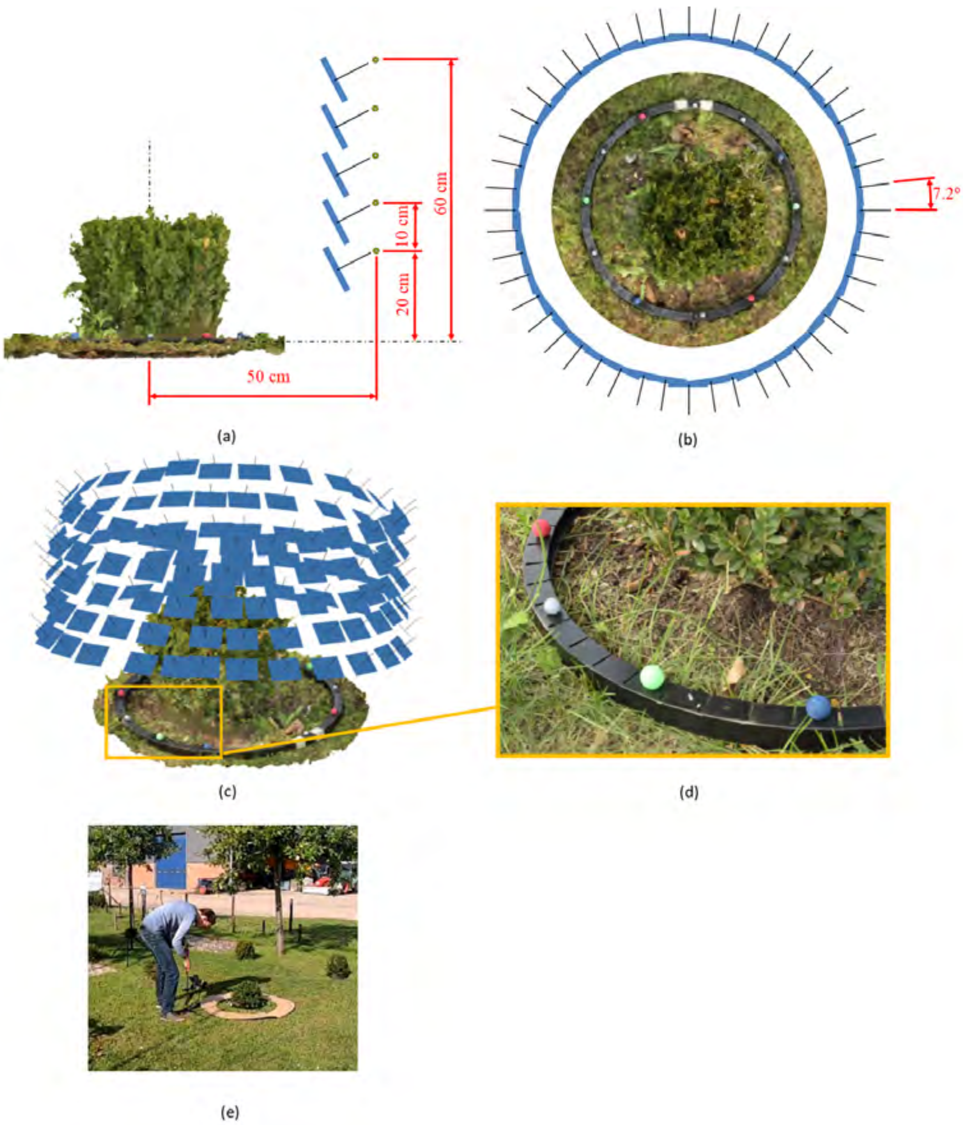

Figure 4: Photographic process layout. a) Elevation view of a bush showing the 5 camera height positions. b) Plan view showing the 50 camera positions around the bush. c) Isometric view showing all camera positions. d) Zoomed view of reference markers. e) Example in practice.

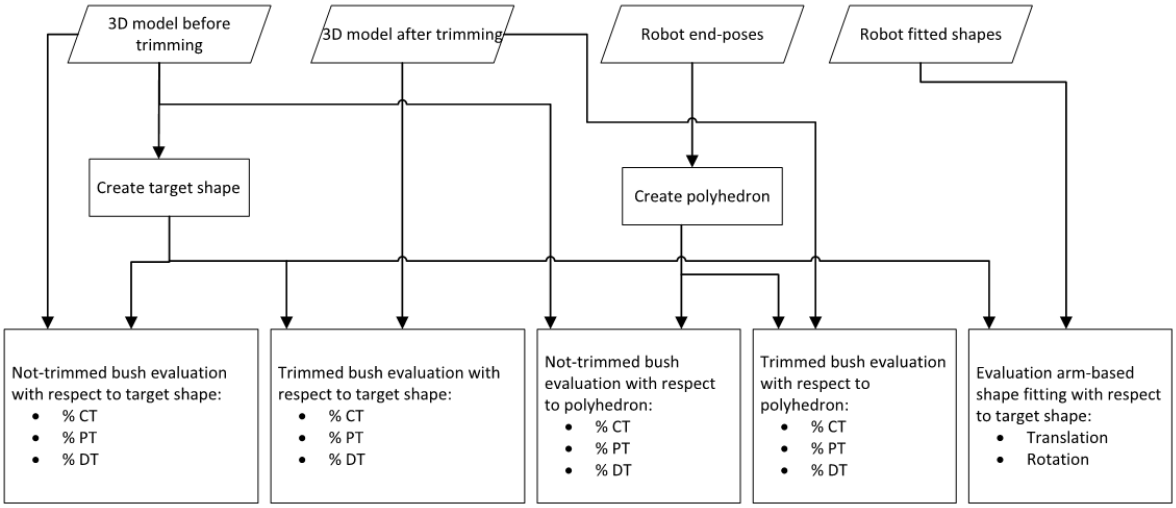

Figure 5: Evaluation software pipeline, which processes 4 input files: the point cloud obtained with SfM photogrammetry before trimming the bush (3D model before trimming), the point cloud obtained with SfM photogrammetry after trimming the bush (3D model after trimming), the final poses of the robot arm (robot final poses) and the shapes estimated by the robot from the 5 trimming positions (robot fitted shapes). The point cloud before and end-poses are used to calculate the target shape and polyhedron. Both are used in the last five blocks to calculate the performance of the robot: percentage correctly trimmed (%CT), poorly or undertrimmed (%PT) and deeply or overtrimmed (%DT) points.

$-0.02m < \delta_i < 0.02m$ were classified as a correctly trimmed points (P_{CT}). On the contrary, poorly or undertrimmed points with with $\delta_i > 0.02m$ (P_{PT}), while points with $\delta_i < -0.02m$ were classified as deeply or overtrimmed points (P_{DT}).

The evaluation software output provides a quantitative assessment in terms of five different analysis (Figure 5):

- Not-trimmed bush evaluation: Compares the 3D model before trimming with respect to the target shape.
- Trimmed bush evaluation: Compares the 3D model after trimming with respect to the target shape.
- Not-trimmed bush evaluation (polyhedron): Compares the 3D model before trimming with respect to a polyhedron. The polyhedron was generated by triangulating the robot arm end-poses.
- Trimmed bush evaluation (polyhedron): Compares the 3D model after trimming with respect to a polyhedron. The polyhedron was generated by triangulating the robot arm end-poses (as further detailed below).
- Robot shape evaluation: Compares the position and rotation of the five fitted shapes from the robot with the target shape.

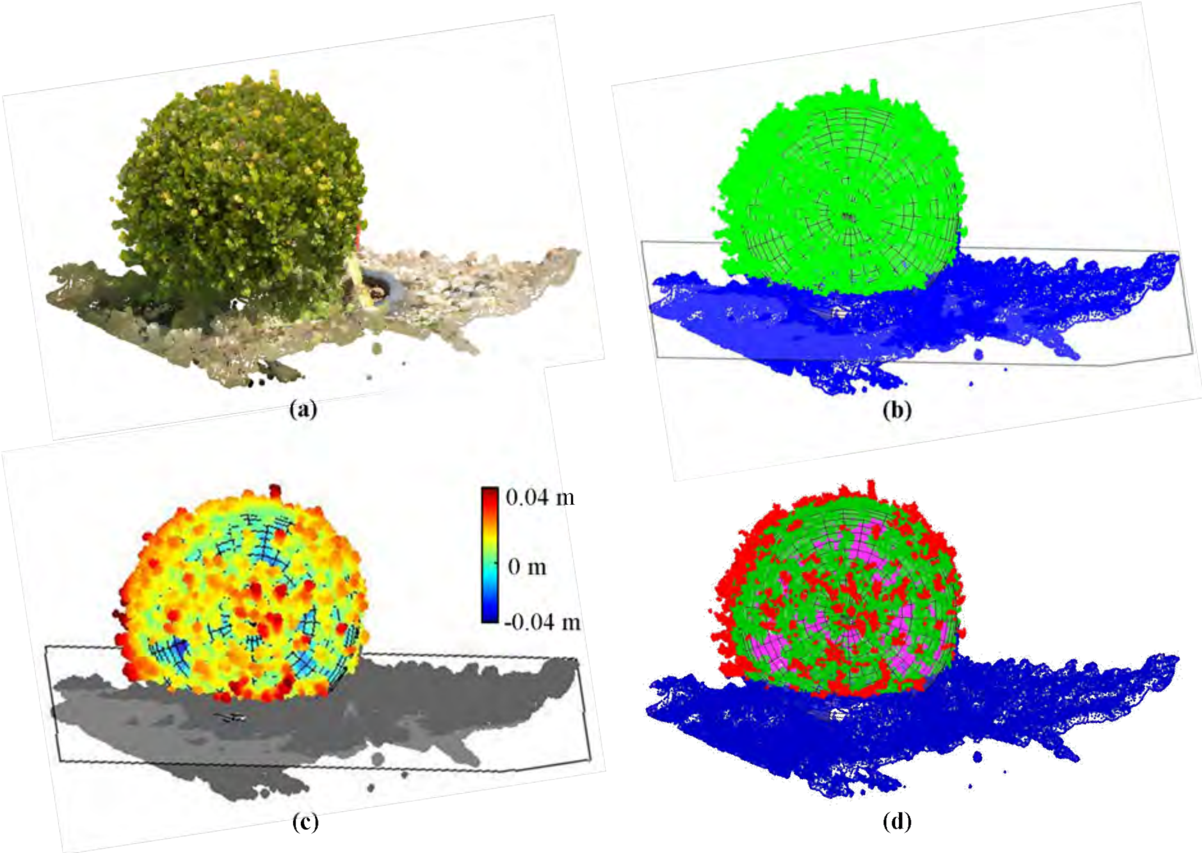

Figure 6: Floor detection and target shape fitting. a) bush 3D model. b) Segmented 3D model, showing the floor points (blue), bush points (green). c) Distance between each bush point and target shape. d) Segmented 3D model showing the floor points (blue), correctly trimmed points (green), poorly trimmed points (red), deeply trimmed points (magenta). In b), c) and d) also the target shape is visualized (grid of black lines).

The quantitative parameters used for not-trimmed bush evaluation, trimmed bush evaluation and trimming evaluation were:

$$CT = \frac{P_{CT}}{P_{CT} + P_{PT} + P_{DT}} \tag{1}$$

$$PT = \frac{P_{PT}}{P_{CT} + P_{PT} + P_{DT}} \tag{2}$$

$$DT = \frac{P_{DT}}{P_{CT} + P_{PT} + P_{DT}} \tag{3}$$

where CT, PT, and DT are the percentage of the correctly trimmed, poorly (i.e. under) trimmed, and deeply (i.e. over) trimmed bush surface, respectively.

Trimming bush evaluation with respect to polyhedron

As the shape fitting is done separately for each platform pose, it is possible that the fitted shape is translated and rotated with respect to previous poses. Consequently, it can be interesting to

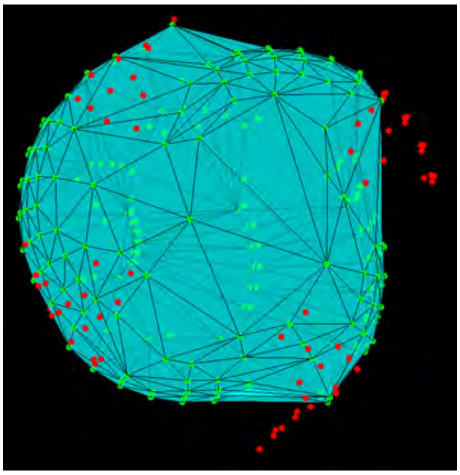

Figure 7: Creating a polyhedron from arm end poses. The red points are not triangulated since other green points were closer to the center.

compare the trimming result not only to the target shape, but also the (combination of) individual shapes that were fitted by the robot. In this case, it is expected that the amount of correctly trimmed surface is higher, since the actual target shape from trimming is used, and performance should thus better match this shape. This evaluation is done by assuming that the arm-end poses during trimming are located directly on the fitted target shapes. By triangulating all these robot arm end poses, a polyhedron can be created that resembles the desired shape after trimming. Only the inner points of the end poses are triangulated as shown in Figure 7. Similar to the target shape the CT, PT and DT points are calculated using an Euclidean distance of $-0.02m < \delta_i < 0.02m$.

2.1.7 Robot shape fitting evaluation

The robustness of the shape fitting method is first tested by scanning a bush at the same pose for 10 times with the robot at a static position under lab conditions. From this, 10 target shapes result, which are compared against each other to see the variation in center poses. The standard deviation of the translation between the 10 fitted shape gives information about the robustness of the shape fitting method, as ideally this should be close to zero (observations/results should be the same).

Additionally, to evaluate the shape fitting performance during the trimming experiments, the translation and rotation displacement of the robot fitted shapes with respect to the target shape were computed using the Iterative Closest Point (ICP) algorithm. ICP aligns the fitted shape with the target shape by minimizing the difference between corresponding points. The resulting transformation matrix to transform the fitted shape to the target shape gives information about the translation and rotation error.

2.1.8 Qualitative evaluation of trimming result

Qualitative evaluation of trimming performance (e.g. visual trimming result, smoothness of trimming, matching with desired shape) were obtained by manually scoring the trimmed bushes.

For this, the trimmed bushed were placed in a row (Figure 8), and 10 persons were asked to evaluate the trimming outcome based on aesthetics. For this, a standardized scoring form was used, as shown in Appendix A.1. Scoring was performed by 6 members of the Wageningen Trimbot Team, 2 colleagues from the Agro Food Robotics group, 1 boxwood grower and 1 professional gardener.

The following criteria were evaluated:

1. No deeply (i.e. over) trimmed points
2. No poorly (i.e. under) trimmed points
3. Bush trimmed into target shape
4. Symmetric result
5. Bush centred on stem
6. No large branches remaining
7. Smooth trimming result
8. Final mark

The data was normalized so each criterion ranged from 1 (very bad) to 5 (very good). Results were registered and summarized by averaging the scores of each person for each evaluation criterion over all bushes.

2.2 Results

2.2.1 General observations

In the experiment 29 bushes were trimmed: 10 spheres, 9 cylinders and 10 cuboids. Each experiment started at the same platform position Figure (9a). Approach of the bush is shown in *b*, while *c*, *d* and *e* show the subsequent trimming poses. A video of this procedure can be found in [5], while Figure 10 shows a photo of the Trimbot trimming a bush.

The trimming actions executed during the experiments generally had a clear impact on the bush, with significant trimming taking place that affected the bush shape. It was sometimes observed that the fitted shape was a bit misplaced (as seen in the RVIZ visualisation), or that the trimming manipulator seemed to be misplaced/rotated from the expected shape. This was most clear at the cubes, as here it is easier to identify if the trimmed and expected planes do match, whereas this was less apparent for the sphere and cylinder.

Furthermore, at some poses the position of the fitted shape produced by the robot showed a large translation/rotation offset with respect to the actual position of the bush. The reasons for this offset are unknown, but solved by restarting the scanning state, to re-run the bush scanning and shape fitting. Another problem observed in fitting was the influence of large branches, making it difficult to fit the shape accurately, especially if these approached the platform more closely. During trimming, branch positions and postures were affected by the trimmer. Due

Figure 8: Overview of the trimmed bushes for manual judgement of the trimming results. Upfront the cylinders, on the right the cuboids and in the rear the spheres.

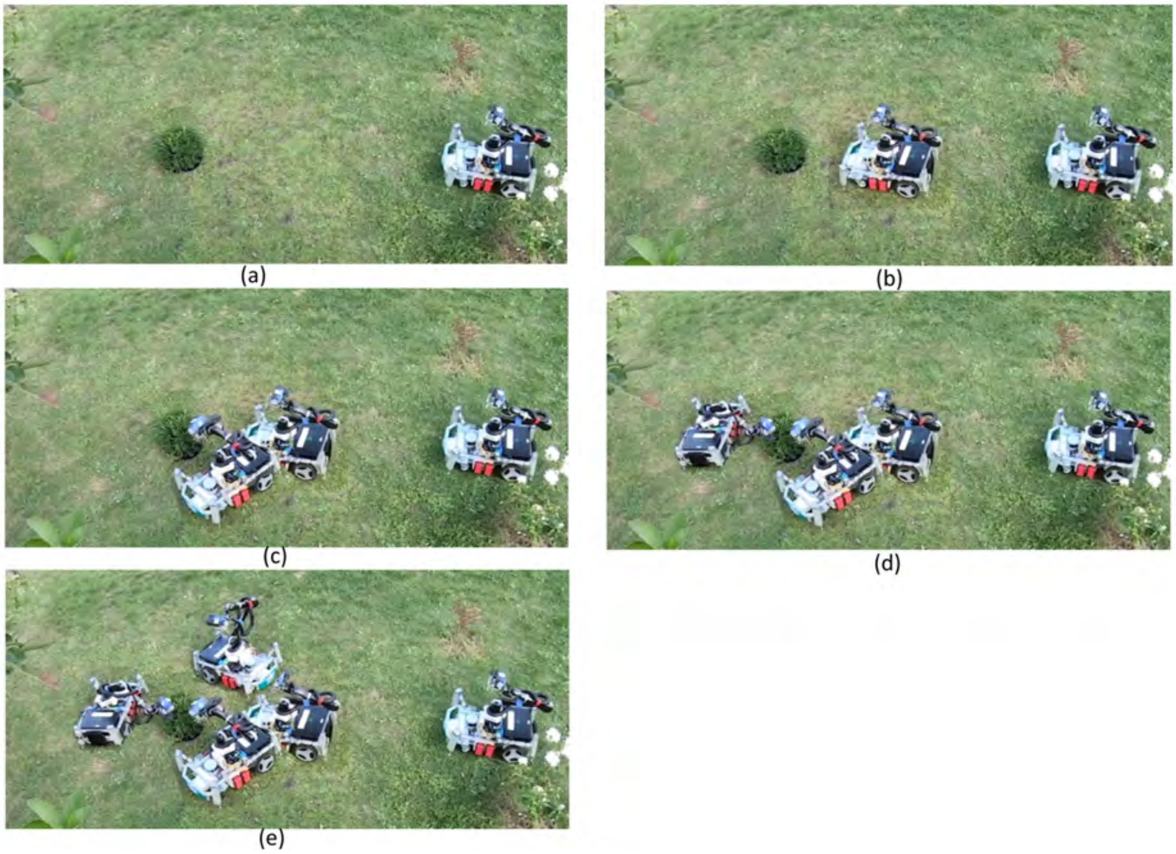

Figure 9: Showing the approaching and servoing of the platform around the bush, with start position (a), approach towards the bush (b) and trimming positions (c,d,e). Only three of the five trimming positions are shown.

to the flexibility of the bush, in general this had limited effect on the trimming capabilities. However, when the bushes frequently contained thicker branches (more then originally expected and harder to cut), the impact on manipulator and cutting performance was more pronounced. As result, the trimmer sometimes got stuck, requiring manual intervention by stopping the arm. This could also lead to displacements of the platform. Such behaviour was especially observed if trimming went rather deep into the bush, either due to a need to remove much outgrowth (like on the top of a cube), or as result of not-so-well matching shape fitting and trajectory planning. This not only influenced the cutting activity, but also the resulting bush shape, as can be seen in Figure 11.

At two moments, serious technical failures were observed. Halfway the experiments, the robot arm experienced a serious failure, with the motor of the second joint being broken. This was resolved by switching over the second arm, so the experiment could be continued. Almost at the end of the experiments, the internal mechanical connections in the platform got loosened, resulting in a collapse of the platform. This was resolved by fixing all mechanical connections.

2.2.2 Target bush approach and visual servoing

In general, the approaching and servoing methods did work, with the robot moving autonomously from pose to pose to trim the bush. A clear challenge, however, was in the actual accuracy of

Figure 10: Trimming a bush

Figure 11: Showing a sphere before (a) and after trimming (b). At some poses the robot trimmed too deeply. This can be observed in figure b since here the inner branches are visible.

the motions. In many cases, the robot was too far from the bush (and sometimes a bit too close), requiring a manual correction of the distance to the bush. In some cases, also a correction of orientation was needed, by rotating the platform such that the line arm-bush was perpendicular to the vehicle length direction.

For spheres, initial distances between robot arm and bush varied between 77 and 95 cm (mean of 84 cm), while after correction of the robot pose this was reduced to 73-83 cm (mean of 78cm). On the first trimming position, the robot was in all cases not on the correct pose, while on trimming positions two and three only two out of the 10 poses were initially correct. On trimming position four, performance was better, with only three times a correction of distance, although also three times an orientation correction was needed. On the 5th trimming position (used for top trimming, closer to the bush) a final distance in the range of 65-71 cm was used, which was achieved only once by visual servoing. For cylinders, the performance was similar to that of the spheres.

For cuboids, performance at pose one was similar, while at pose two and three, distances tend to get larger, with values between 90 and 105 cm occurring more often. However, this was less apparent in the average distance, which was 88 cm before and 75 cm after correction. No clear change was observed in the number of wrong orientations, although in one case the robot was found to move towards another bush. This indicates that there might be a relation between bush shape and approaching accuracy, with round surfaces performing better compared to edged surfaces, although this was not evaluated properly enough to draw real conclusions.

Furthermore, it should be noted that the corrections applied limited the observed errors, as for each servoing manoeuvre, the platform was more or less in a proper starting orientation. If the original pose had been used, it is likely that already at pose two, but for sure at pose three, the platform would have been too far off from the bush to allow reliable trimming.

2.2.3 Quantitative trimming results

Example point clouds before and after trimming are shown in Figure 12 for a sphere, cylinder and a cuboid (in top-down order). After trimming, each shape had less branches in the upper half of the bush, so there is a visible difference as result of trimming. Branches close to the ground could not be trimmed well by robot due to reachability and avoidance of floor collisions. An analysis of the performance for each shape is described below in more detail.

Spheres

For each sphere the percentage of correctly trimmed (CT), poorly trimmed (PT) and too deep trimmed (DT) point cloud points with respect to the target shape have been calculated. Table 1 shows an overview of the average performance for each trimming setting, with an increase in point cloud points in the correctly trimmed range. Also the amount of points where trimming went too deep increased, especially for the repeat and always repeat settings ($> 23\%$). A possible explanation for this might be that the fitted shape tended to be located too deep into the plant or being too small, with excessive trimming as result that is more pronounced if trimming motions are repeated. Such an effect is also visible for the poorly trimmed regions, where the decrease in points is larger if trimming motions are repeated, although there is no clear difference between the trials with result-based repetition and trials with systematically repeated segments. This may suggest that trimming result-based repetitions are useful too remove more outgrowth.

Another reasonable explanation for the errors in trimming might relate to the accuracy of the shape fitting that was performed. In Table 1 the average translation of the fitted shapes (which were used as input for trimming) with respect to the target shape in the evaluation is given, as 3D distance between both centers. As this average varies between 5.8 and 6.5 over the experiments, there is a clear offset of the fitted shape with respect to the desired target, which is then expected to have affected the trimming result as well. Thus, fitting errors might also have a large contribution to the final result, and are therefore evaluated in more detail in section 2.2.4.

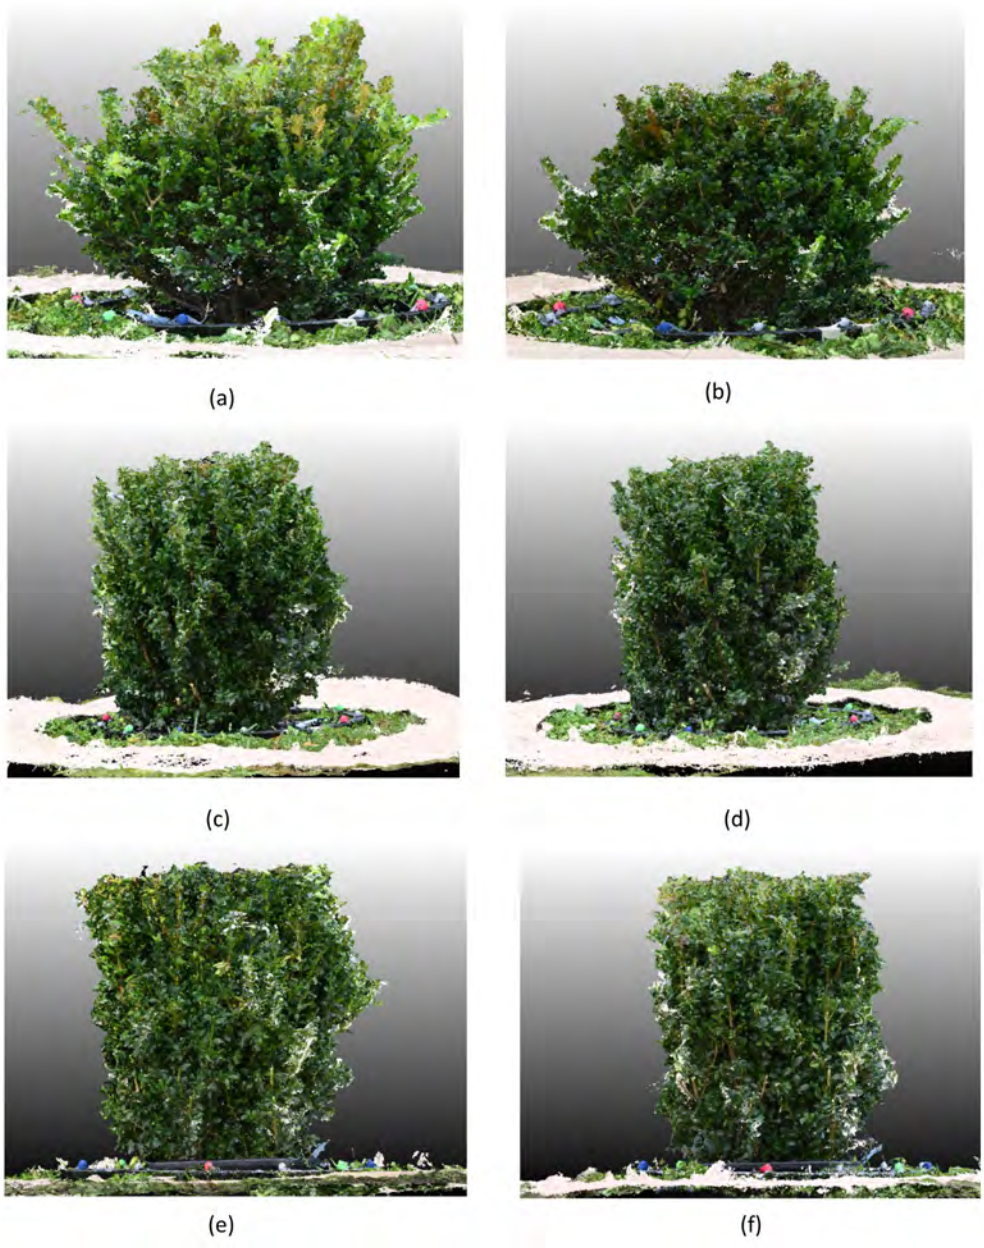

Figure 12: Point clouds before and after trimming. a) sphere before. b) sphere after. c) cylinder before. d) cylinder after. e) cuboid before. f) cuboid after.

In Figure 13, the distance between the points and target shape is visualized for sphere 2 (trimmed with decision based setting) as example of what results are found after trimming a bush. Here, colour indicates the distance, with a dark red colour meaning that the distance of the points towards the target shape is larger than 8 cm. The figures in the second row show a 2D projection of the distance in spherical coordinates. Before trimming the bush (left) has perceptibly more branches, after trimming the number of branches decreases but some large branches still exists, especially those closer to the ground plane. This is also seen from the distribution of points,

Table 1: Results of trimming experiments on spherical-shaped bushes with respect to target shape.

	No repetition		Repeat		Always repeat	
	Average	stdev	Average	stdev	Average	stdev
Translation error [cm]	6.5	3.4	5.8	2.1	6.1	3.1
Point cloud distribution Before						
% Correctly Trimmed (CT)	29	5	25	11	29	14
% Poorly Trimmed (PT)	68	5	71	12	68	14
% Deeply Trimmed (DT)	3	0	4	2	3	1
Point cloud distribution After						
% Correctly Trimmed (CT)	56	4	48	5	46	8
% Poorly Trimmed (PT)	36	4	25	11	23	8
% Deeply Trimmed (DT)	8	1	27	12	31	6

with before 9% CT and 90% PT, while after trimming this became 45% CT and 42% DT. Also some sides of the bush are trimmed too deeply, indicated by the dark blue points in the spherical coordinate plot. The histograms in Figure 14 show this as well, with a shift to the left after trimming. After trimming there is a lower ratio of points whose distance is > 10 cm, which means that large branches are removed. Nevertheless the number of inside points whose distance is higher than 5 cm (i.e. the points whose signed distance is lower than -5 cm) also increased after trimming. This indicates that the robot trimmed too deeply at some points, which is confirmed by the increase in percentage DT points after trimming (DT before 1%, DT after 13%). From Figure 13, especially the 3D view after trimming, it can also be seen that the trimming is not very smooth over the sphere, with an irregular distribution of points, and sometimes even too deeply and poorly trimmed areas next to each other.

Cylinders In Table 2, the translation error and average % of CT, PT and DT points are summarized for the cylinders. As expected the percentage of poorly trimmed points decreases after trimming in every experiment. More interesting is the increase in % of too deeply trimmed points, which roughly double in all cases. Moreover, the percentage of correctly trimmed points after trimming has hardly improved as a result of trimming, and even decreases in the always repeat experiment. From this, it seems that the bush was more cylinder-shaped before trimming than afterwards.

This can also be seen in Figure 15, which visualizes for cylinder 4 (trimmed with decision-based repetition) the distance between the points and target cylinder before and after trimming. In this example, the percentage CT points decreased after trimming (from 52% to 38%). This was mostly caused by the fact that the robot trimmed too deep, which is clearly visible in Figure 15 *f* where almost all points are negative (dark blue). From plot Figure 15 *a*, it becomes clear that this is more evident on the right side, while the left part tends to have more poorly trimmed points. Also should be noted that on the top of the cylinder, the points tend to be too deep before and after trimming, indicating the target cylinder might have been placed too low.

When looking at the histograms in Figure 16 several interesting observations can be made. For example, before trimming the points are rather well distributed around the target shape, although the distribution is somewhat wide with most points between -5 and +5 cm (trimming goal is to

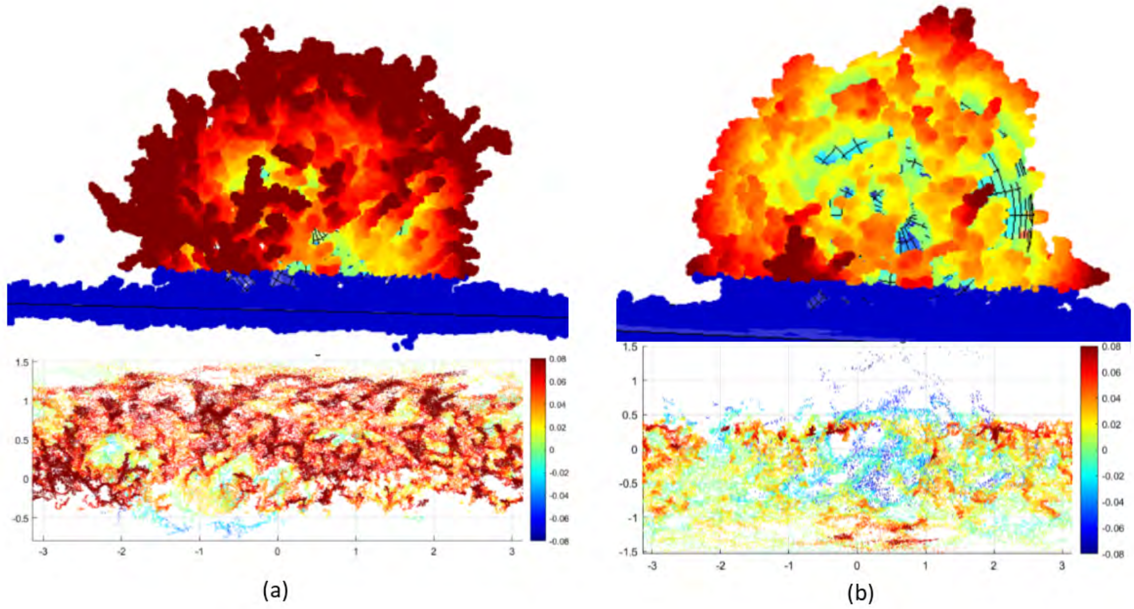

Figure 13: Distance between points and the target shape before (left) and and after (right) trimming. The top row shows a view on the 3D sphere, while the bottom row gives a 2D projection of the sphere. Green color represents 0 distance, orange to red positive distances (undertrimmed) up to 0.08m and light blue to dark blue color negative distances (overtrimmed) up to -0.08m.

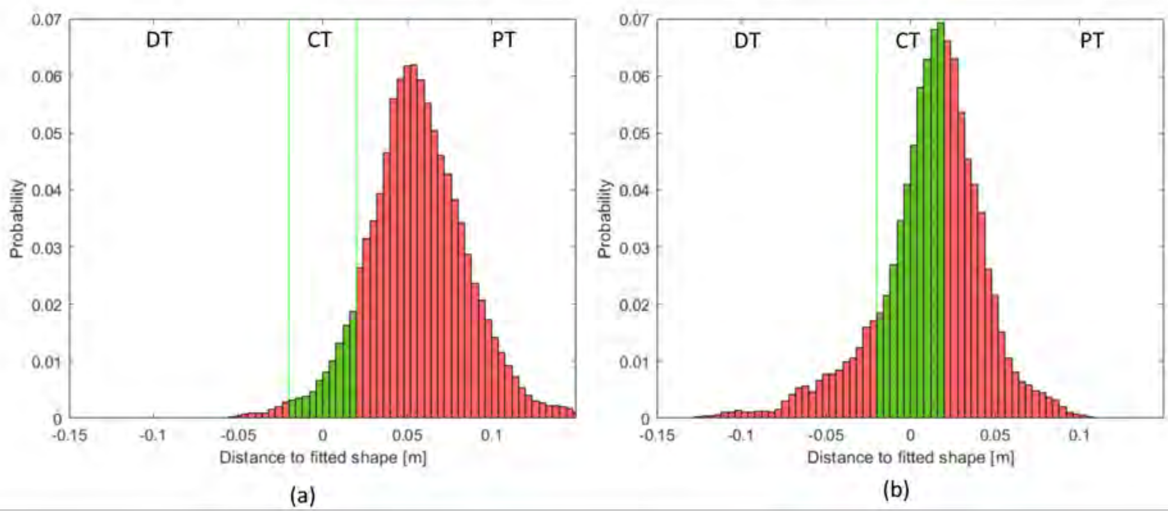

Figure 14: Histogram showing the distance from evaluation point to fitted shape. a) sphere before b) sphere after.

reduce this to -2 and +2 cm). After trimming, the distribution is changed such that the right tail became smaller (indicating proper removal of too long branches), but also that a shift towards the left that refers to bush parts trimmed too deep. Also, this might be related with a trimming target fitted too deep or a trimming plan aiming for a smaller bush or due to the translation error

Table 2: Results of trimming experiments on cylindrical-shaped bushes with respect to target shape.

	No repetition		Repeat		Always repeat	
	Average	stdev	Average	stdev	Average	stdev
Translation error [cm]	6.7	2.1	5.7	2.2	6.0	4.4
Before						
% Correctly Trimmed (CT)	38	4	43	7	35	2
% Poorly Trimmed (PT)	34	5	35	15	41	1
% Deeply Trimmed (DT)	27	2	22	8	24	2
After						
% Correctly Trimmed (CT)	40	5	47	7	31	6
% Poorly Trimmed (PT)	10	3	10	6	20	10
% Deeply Trimmed (DT)	51	9	43	12	49	5

during robot shape fitting.

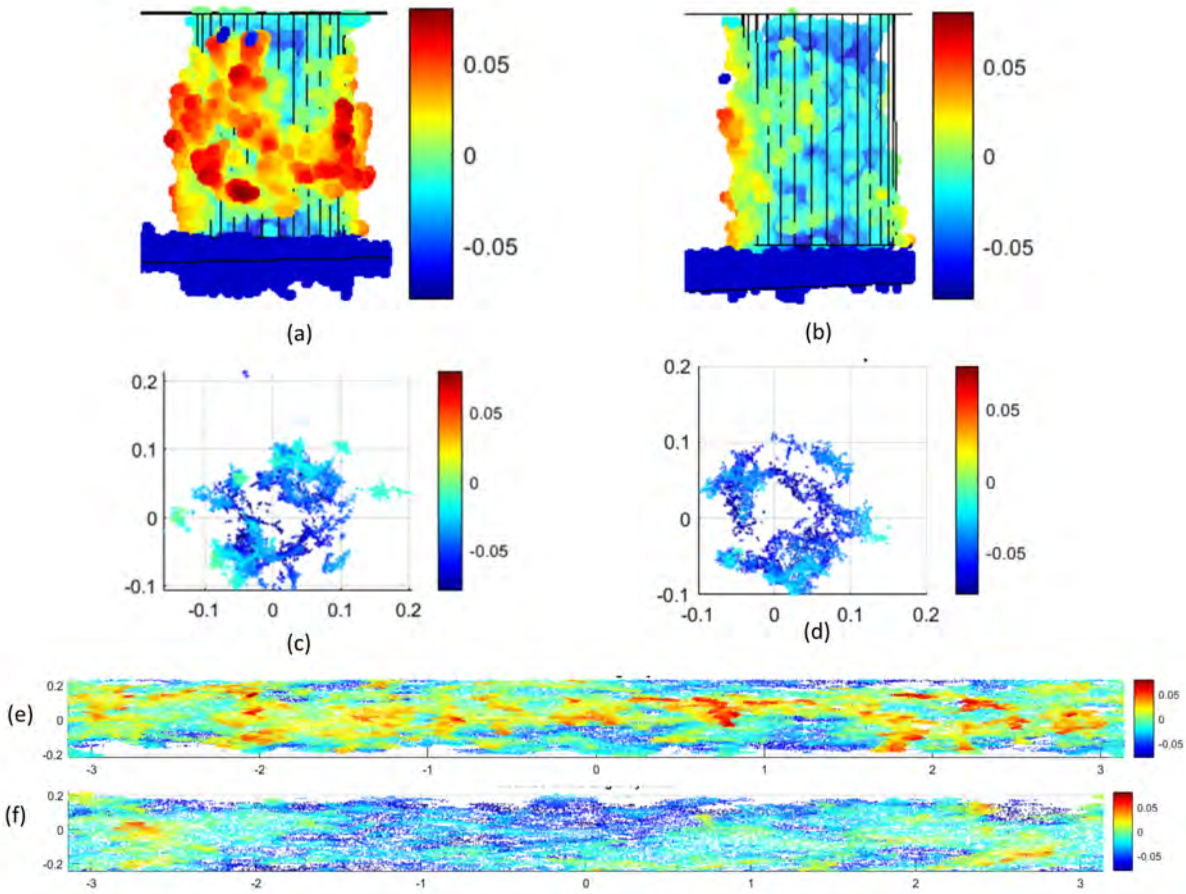

Figure 15: Distance between points and the target shape before and after trimming. Plots *a* and *b* show the distance to the target before and after trimming as side view, while plots *c* and *d* show this for the top plane. Plots *e* and *f* show the projection of the circumference of the cylinders. Colour codes as in Figure 13.

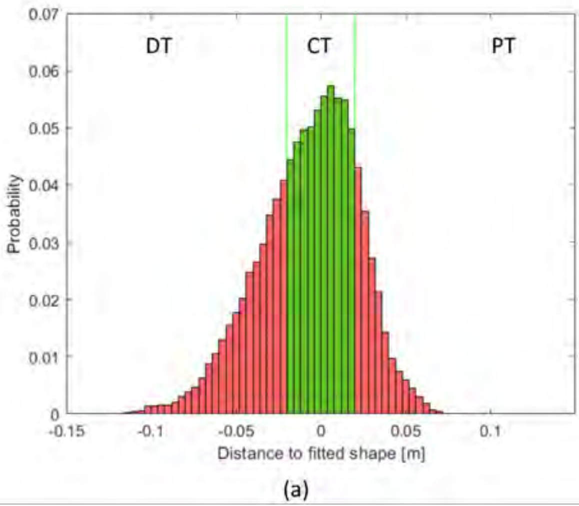

Figure 16: Histogram showing the distance from evaluation point to fitted shape. a) cylinder before. b) cylinder after

Cuboids

In Table 3, the translation error and average percentage of CT, PT and DT points are given for cuboids. The error of the top part is not included since the top part was not trimmed (see for explanation section 2.1). After trimming the percentage of correctly trimmed points increases in all experiments. This increase is caused by reducing the amount of poorly trimmed points. The always repeat experiments have the highest percentage of correctly trimmed points, but it should be taken into account that the input bushes of the always repeat experiment had by incident a higher percentage CT points before trimming compared to the other experiments. Similar to the spheres the no repetition experiments have the lowest average percentage DT points after trimming, at the cost of a higher percentage PT points. But due to the high standard deviation (8-15%) it is not possible to conclude there is a significant difference between the experiments.

Table 3: Results of trimming experiments on cuboid-shaped bushes with respect to target shape.

	No repetition		**Repeat**		**Always repeat**	
	Average	stdev	Average	stdev	Average	stdev
Translation error [cm]	5.1	1.8	4.3	1.2	5.7	0.9
Before						
% Correctly Trimmed (CT)	45	4	45	10	52	8
% Poorly Trimmed (PT)	42	9	39	15	32	11
% Deeply Trimmed (DT)	13	5	15	6	16	5
After						
% Correctly Trimmed (CT)	52	1	51	5	55	5
% Poorly Trimmed (PT)	29	8	23	15	15	4
% Deeply Trimmed (DT)	19	8	26	12	30	8

For explaining the results, cuboid 10 is used, which is trimmed using decision-based repetition. Here the amount of PT points decreased from 38% to 12% after trimming. The distance between the points and the target shape is visualized in Figure 17. Figure 17*c* visualizes the distance to each side of the cuboid before trimming, while Figure 17*d* does this after trimming. Some large branches are removed, but as shown in the first plot in *d* also many points are trimmed too deeply. Although this is less on the other sides of the bush, for the full bush still 39% of the points are trimmed too deep, whereas it was 18% before trimming. Still, on some places also too long branches remain, such as on side 3. As result, the cuboid has been trimmed, but not fully according to the target shape, which leaves it with rather irregular surfaces. When looking at the histogram (Figure 18), the expected reduction of the right tail as result of trimming is hardly present. Still, the whole distribution does shift to the left, with its peak at -0.02 after trimming. This indicates that more points are inside the target shape, as most points have a distance smaller than 0 to the target surface. From this, it can be concluded that in general there was trimmed too much, and that resulting surfaces were not very smooth.

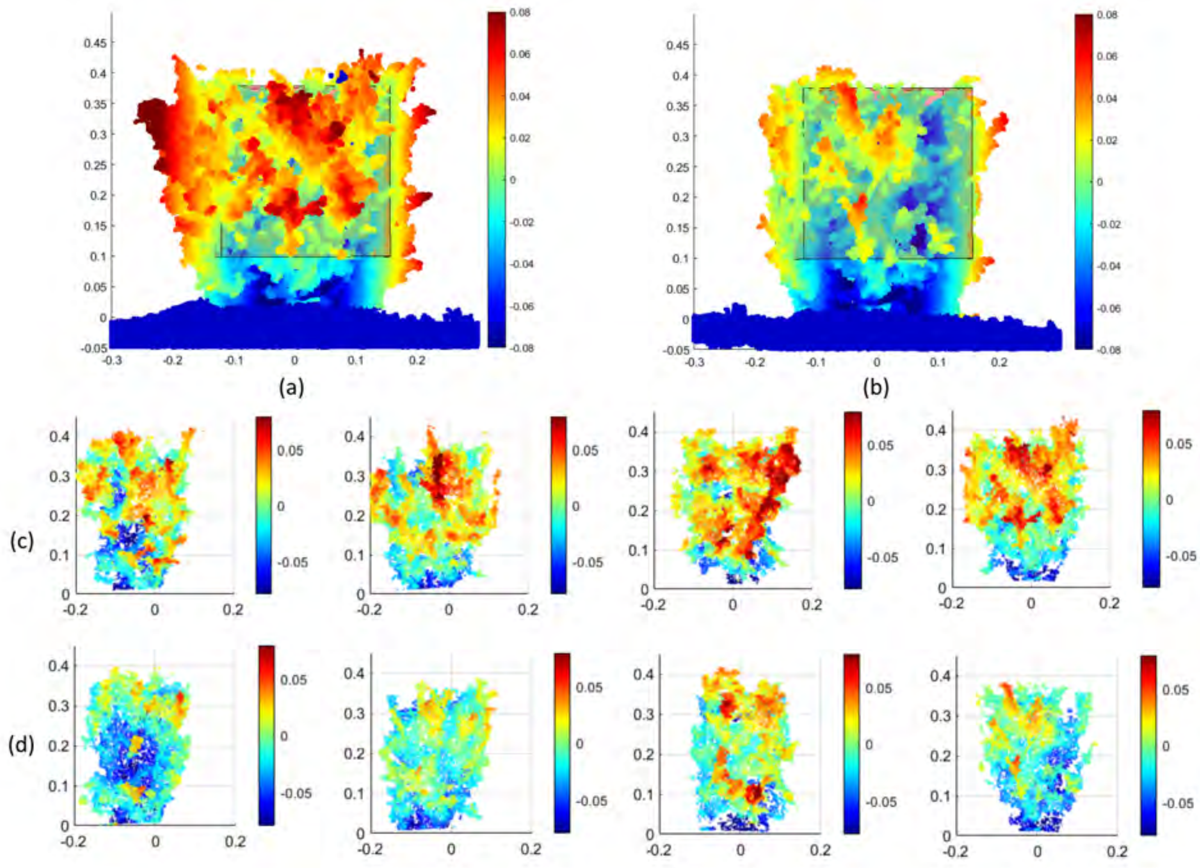

Figure 17: Distance between points and the target shape before and after trimming. a) cuboid before trimming b) cuboid after trimming c) Side planes of cuboid before trimming. d) Side planes of cuboid before trimming. Colour codes as in Figure 13.

Trimming bush evaluation with respect to polyhedron

CT, PT and DT were also calculated with respect to the polyhedron instead of the target shape. The results of this are summarized in Table 4 for the spherical bushes. In this table, the percentage PT points is ±12 %, whereas in Table 1 the percentage after trimming was at least 23 %, thus indicating an improved trimming accuracy. The possible source for this difference is visualized in Figure 19, where the red sphere indicates the target shape that was used previously.

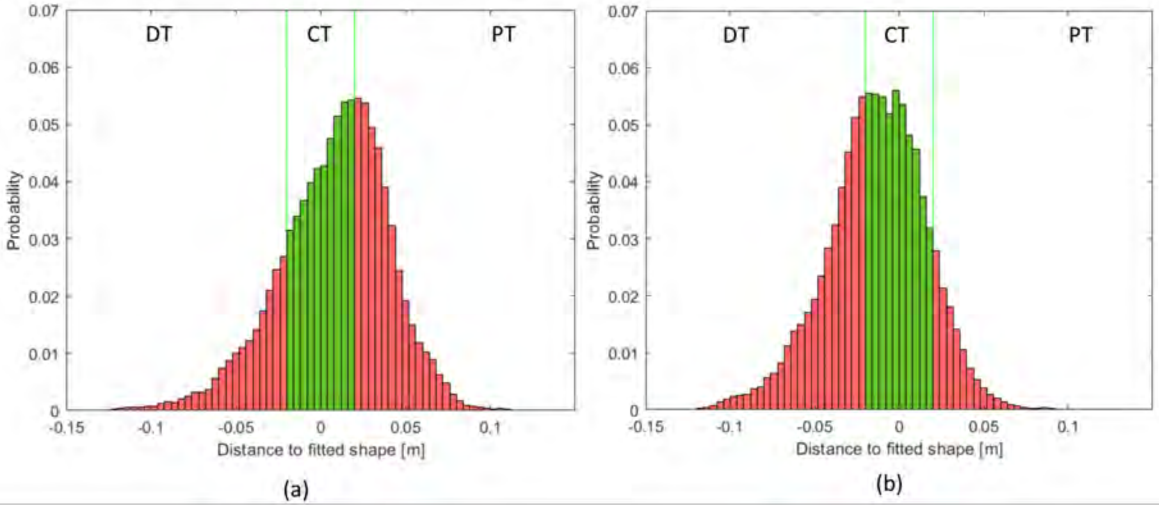

Figure 18: Histogram showing the distance from evaluation point to fitted shape. a) cuboid before. b) cuboid after

According to this figure, there is a clearly visible offset on the right side between the blue polyhedron and the red target shape. This means that the robot does not trim as closely to the target shape as desired. Consequently, the high percentage PT points with respect to the target can be explained by the offset between the arm end poses and target shape.

In Tables 5 and 6, the performance results with respect to the polyhedron are given for the cylinders and cuboids. In contrast to previous table (Table 4) the cylinder and especially the cuboids have a high percentage PT points after trimming. This can partially be explained by the fact that the end-poses were too deep to be trimmed. Another explanation is the movement of the robot arm. The manipulator slightly pushes the branches towards the inside. Consequently, less branches are trimmed than expected. When checking the videos recorded during the experiments, it can indeed be observed that during trimming sometimes the bush is moved substantially while limited trimming is performed.

Finally, for the cuboids all planes were trimmed twice, which means that the most inner plane is used to construct the polyhedron. In trimming, however, not all parts of this plane might be handled properly, such that always some outgrowth with respect to the polyhedron remains, and sometimes is even quite large.

The cylinders and cuboids did not have a high percentage PT with respect to the target shape (Table 2 and Table 3). The same cylinder as in Figure 15 is used to explain this phenomenon. In Figure 20, the arm end-poses are shown as large coloured points, while the red circle indicates the target cylinder. From this, it becomes clear that part of the arm positions are inside the target cylinder, with the manipulator trimming deeper than supposed and thus explaining the higher amount of deeply trimmed points (60% with respect to target shape). Furthermore, Figure 20 shows that the arm end poses from different platform poses (the different colours) do not match very well. For example between the green and yellow points there is a clear offset that likely created unsmooth surfaces. Also for the other poses this is the case. As the arm end poses are derived from the fitted shapes, these results indicate that either the fitted shapes are not accurate, or that the trajectory planning and execution does not properly follow the fitted shapes or that some post-movement registration replanning is needed. Since from previous experiments the latter seemed less likely, the accuracy of shape fitting is evaluated in more detail first.

Table 4: Results of trimming experiments for spherical-shaped bushes with respect to polyhedrons.

	No repetition		**Repeat**		**Always repeat**	
	Average	stdev	Average	stdev	Average	stdev
Before						
% Correctly Trimmed (CT)	45	5	41	5	33	4
% Poorly Trimmed (PT)	39	12	50	12	63	6
% Deeply Trimmed (DT)	16	8	10	8	3	2
After						
% Correctly Trimmed (CT)	57	4	49	8	58	5
% Poorly Trimmed (PT)	12	6	10	2	13	6
% Deeply Trimmed (DT)	32	10	42	6	29	11

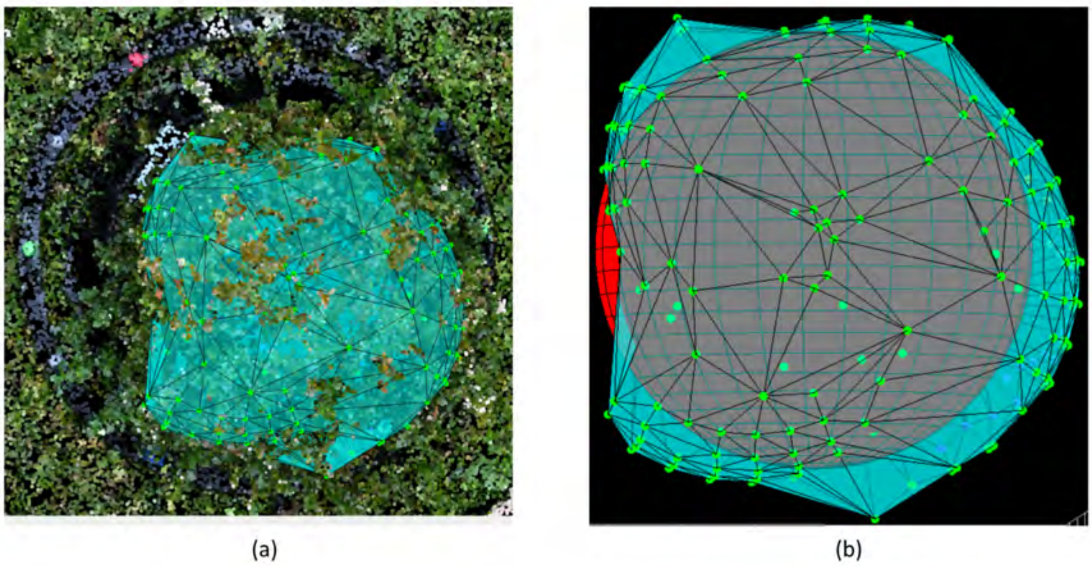

Figure 19: a) Figure a shows the trimmed bush with respect to the polyhedron (blue). The green points are the final poses of the robot arm. b) The red sphere is the target shape. 42% of the points were poorly trimmed, which was partly caused by the offset between the target shape (red) and polyhedron (blue).

2.2.4 Evaluation of arm-based shape fitting

In previous paragraphs, the polyhedron figures indicated that the arm-based shape fitting pipeline might not be accurate. Thus, a more detailed analysis of the shape fitting accuracy was carried out. This was done in two parts: robustness and position accuracy.

The robustness of the shape fitting method is tested by scanning a bush from the same platform position 10 times. The average position (standard deviation between brackets) of center of the

Table 5: Results of trimming experiments for cylindrical-shaped bushes with respect to polyhedrons.

	No repetition		**Repeat**		**Always repeat**	
	Average	stdev	Average	stdev	Average	stdev
Before						
% Correctly Trimmed (CT)	39	4	43	1	34	4
% Poorly Trimmed (PT)	50	6	48	3	54	6
% Deeply Trimmed (DT)	11	2	9	2	12	2
After						
% Correctly Trimmed (CT)	53	1	59	3	45	12
% Poorly Trimmed (PT)	20	6	17	1	24	11
% Deeply Trimmed (DT)	28	5	24	3	30	1

Table 6: Results of trimming experiments for cuboid-shaped bushes with respect to polyhedrons.

	No repetition		**Repeat**		**Always repeat**	
	Average	stdev	Average	stdev	Average	stdev
Before						
% Correctly Trimmed (CT)	40	3	32	10	37	3
% Poorly Trimmed (PT)	50	2	62	12	55	5
% Deeply Trimmed (DT)	10	1	6	3	8	3
After						
% Correctly Trimmed (CT)	51	8	47	7	52	4
% Poorly Trimmed (PT)	35	7	44	9	32	2
% Deeply Trimmed (DT)	14	1	9	2	16	5

fitted shape with respect to the arm base is x=0.9 (0.3), y=-78.0 (0.5), z=10.4 (0.9) cm. By averaging the 3D displacement over all fitted shapes, it was found to be 0.9 cm on average with an SD of 0.4cm, which is lower than the displacement values reported in Tables 1 - 3. Since all points are on average within 1 cm of the centre, this was below the indicated threshold of 2cm, so it can be concluded that the shape fitting method is sufficiently robust.

The position of the fitted shape can be different at each pose because the bush is not homogeneous and the position of the robot with respect to the bush is not constant. Consequently, a translation of the fitted shape at each pose can result in a final bush without much similarity with the desired shape. To evaluate this, the five fitted shapes of sphere 1 are visualized in Figure 21(a). Each fitted sphere has a different colour. In the perfect case all fitted shapes would have been on the same position, but the figure shows that there is an offset between them. Figure 21b shows the 2D center poses (in the X-Y plane) of the all robot-fitted spheres, cylinders and cuboids With respect to the center of the target shape for each bush type.

From Figure 21(b), it shows that translation errors occur in both X and Y direction. For the sphere as example, the average translation found is 6.1cm (SD 2.9cm), which is large considering a bush radius of 15 cm as used in the experiments. No clear trend or pattern in the data could be observed, such that the actual distance between fitted shaped might be almost twice the deviation showed here. For example the red and green sphere in figure a have a

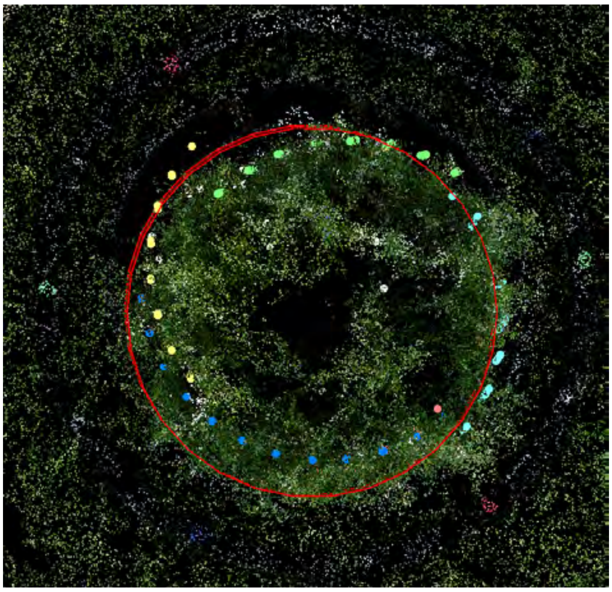

Figure 20: The red line is the target shape. The coloured dots are the final poses of the arm in the three different cutting positions of the mobile base. The blue, yellow and partly green are inside the model, which means the robot trimmed too deep compared to the target.

distance of -2 and +2 cm to the target. The distance between the centers of both the red and green one would then be 4 cm. From this, it might be concluded that the shape fitting pipeline is not able to fit a shape reliably at the same position for different platform poses.

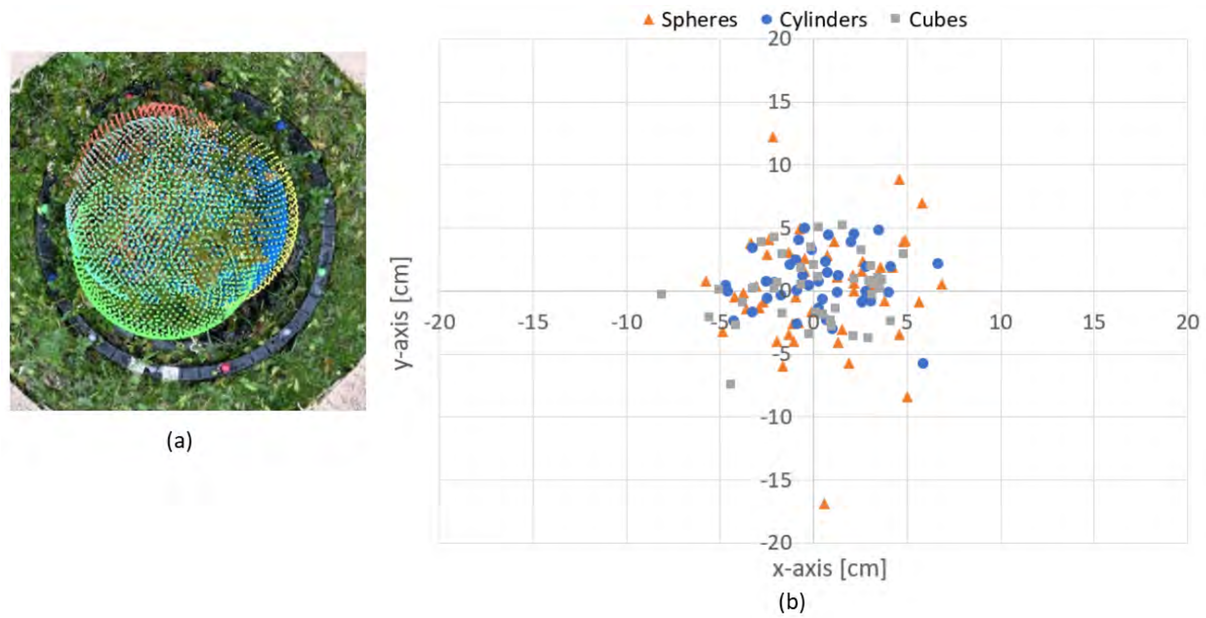

Figure 21: Translation of fitted shapes. a) Visualisation of translation. b) Translation of spheres (orange triangles), cylinders (blue points) and cuboids (gray rectangles) with respect to target shape.Note that the figures only show the translation in the horizontal plane.

Next to the displacement (as shown for the spheres), also the rotation of the fitted shape with respect to the target shape can have an important influence on the trimming accuracy, especially for objects that are not symmetric in all directions. For example, the cylinders had a maximum absolute tilt of 3.7 degrees. Since this maximum absolute rotation is relative small, no further analysis is done for the cylinders. For cuboids, the rotation error of the fitted shapes around the vertical axis is larger, as is visualized in Figure 22a for cuboid 10. This figure shows that the four fitted cubes are distinctly rotated with respect to each other. This observation triggered further investigation of this phenomena. The variation in the rotation is visualized with a boxplot in Figure 22b over 40 fitted shapes (originating from 10 bushes with 4 platform poses each). The average rotation error is indicated with a ”x” at 17.7 degrees. The variation is relatively high, as more than 25% of the fitted cuboids have an absolute rotation error $>$ 25 degrees. From this, it seems that the fitting method has clear problems to orient the cuboids consistently for each pose. From these results, it can be expected that trimming a cuboid according to these shapes will be rather destructive, as hardly any aspect of the original or desired shape will remain after trimming from multiple platform positions. However, the shape after trimming (as shown in Figure 22c, shows a much better shape then one would expect from combining the inner areas of the coloured squares in Figure 22a (the black line). This triggered further comparison of end-effector poses and robot fitted target shapes.

Warping errors The result of this comparison between arm poses and robot-fitted target shapes is shown in 23ab, where the coloured points should be on the fitted shapes. However, it shows that this is not the case, and the bush is trimmed using deviating end-effector poses. As this is something unexpected, further analysis was done to reveal the origin and occurrence of this phenomena. First of all, for the cuboid shown in Figure 22, this difference between the target shape and the actual arm poses turned out to be beneficial for the trimming result. The matching of actual trajectory and desired bush shape was better than what could be expected based on the fitted bush. Second, it was found that these errors between target shape and end-effector poses mainly occurred for shapes with a large translation or rotation error. This is for example visible in Figure 22b and c, displaying the consequence of a wrong shape fitting and warping trajectory. The red oval shows a high negative error, indicating that the trimming tool went too deep into the bush. When trying to explain this behaviour, a similar phenomenon was already observed in previous testing, but to a lesser extent (see also Deliverable 2.5), where further investigation of this issues was desired. This was carried out with respect to the general behaviour of the arm and the shape of the resulting path (showing the desired behaviour), but the distance to the actual target was not evaluated. Specifically for the case observed in Figure 22, part of the error might be explained from the warping method used during the experiments to speed up trajectory planning. Instead of calculating a new trajectory for each target, a set of pre-planned trajectories on a variation of target poses was used. From this, the best fitting trajectory was selected and warped to match the target shape (as described in Deliverable 2.6). This matching was done based on the absolute distance to the closest matching shape, without considering orientation of the shape. Such assumption was considered reseasonable, as the robot was expected to always be at the same orientation with respect to the shape (for the cuboids, at a 45° angle and facing an cuboid edge). The warping could then correct for minor differences ($<$2 cm) between the actual target shape and the pre-planned trajectory. However, for shapes where orientation does matter (such as cubes), ignoring the orientation in the selection method includes the risk of selecting of a wrong trajectory. This holds especially when the orientation of a fitted cuboid is very far from the expected 45° rotation with respect to the vehicle direction.

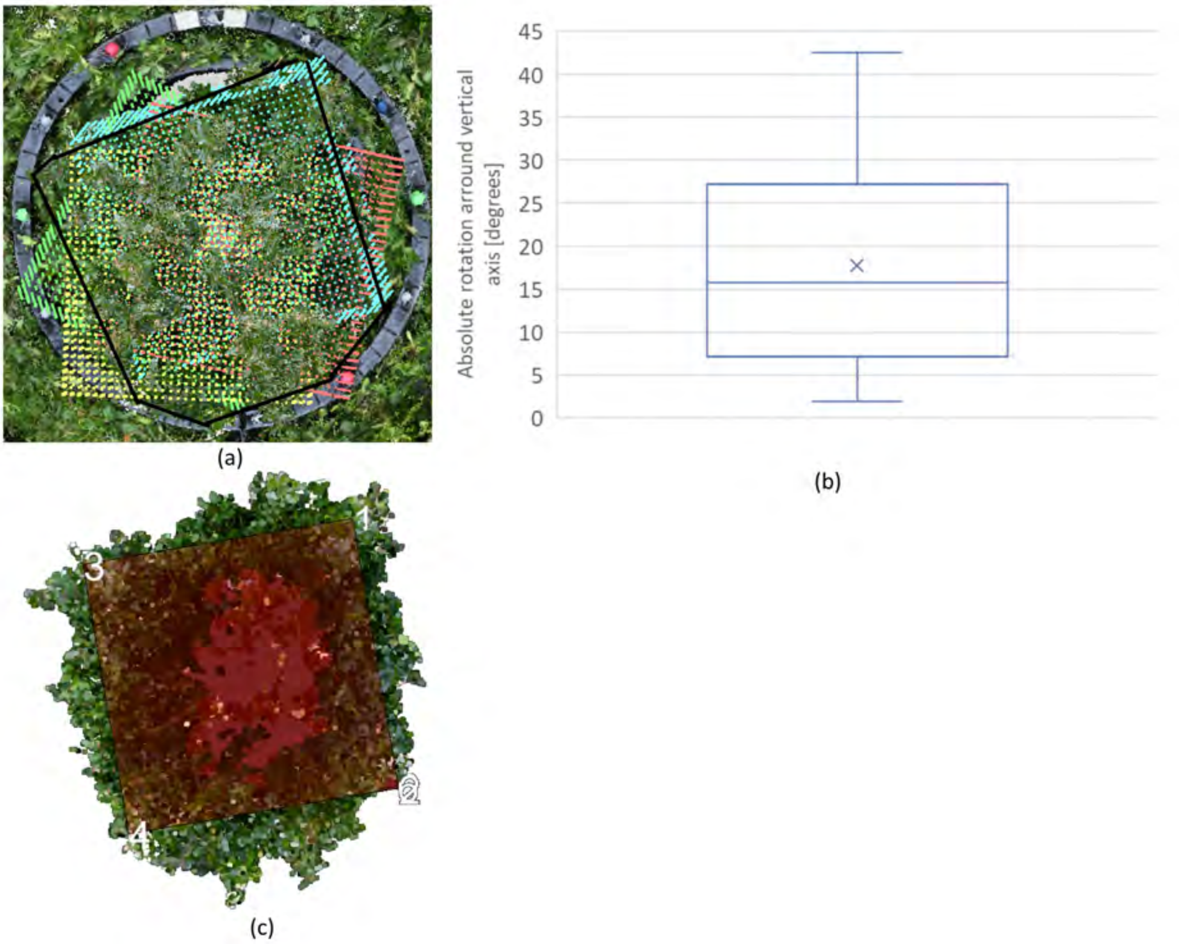

Figure 22: Rotational variation during shape fitting of a single cuboid from multiple platform poses. a) Fitted shapes are shown as dotted areas, with each colour indicating a different pose, and the solid line showing the reference shape. b) Boxplot of calculated rotations around the vertical axis. c) Top view of trimmed bush

In those circumstances, the robot end-poses were adapted taking into account deviations in the centroid position (again within the boundaries of the grid in the database), but could not correct for severe shifts in the cuboid rotation. This introduced a bias in the produced trajectories for cuboids, as it can be observed in Figure 23a, and turned out to be advantageous in producing acceptable trajectories in this particular case. Indeed, since the vehicle orientation was manually adjusted to the desired pose before scanning and trimming at each trimming position, a severe orientation shift with respect to the shapes in the database could only be a consequence of a fitting error. Therefore, a side effect of the employed warped knowledge base technique is that, assuming a properly oriented vehicle, a reasonable cutting behaviour can be achieved despite a wrong fitting result. Still, in future such cases should be noted to a supervisor to ensure the system is indeed behaving as expected.

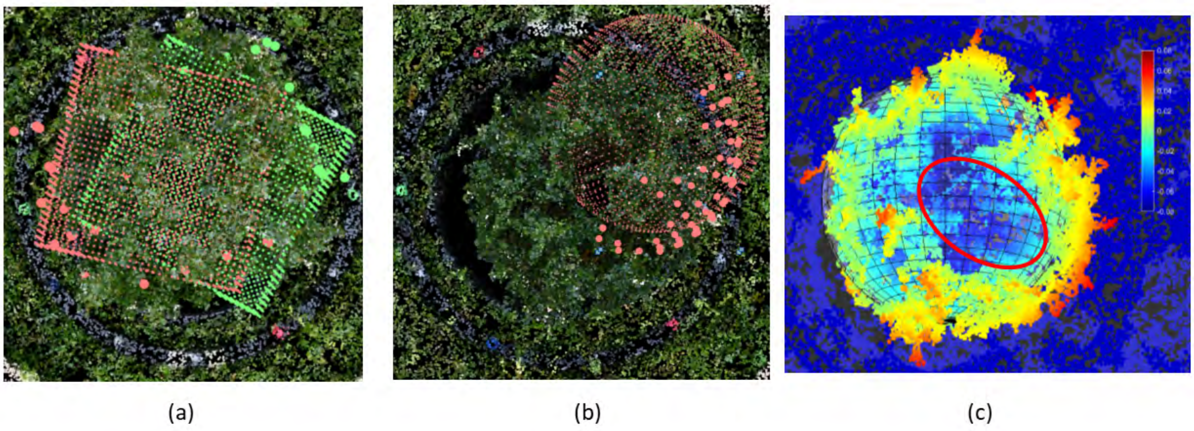

Figure 23: a) The coloured points are the end positions of the arm. The warping function does not seem to work properly because the points should be on the fitted shapes (red and green cuboid). b). The warping went also wrong for the spheres. c) The points visualize the distance to the target shape. At the red oval the robot trimmed too deep because of wrong shape fitting and warping error in Figure b.

2.2.5 Qualitative trimming results by manual bush scoring

The summarized results of the manual scoring are shown in Table 7 for the spheres, the cylinders and the cuboids. For each bush (with corresponding bush number) the average scores and standard deviations are shown for each evaluation criterion.

The last criteria (#9) is a calculated value average of criteria #1-#7. The average mark for the spheres is higher for the no repetition experiments than the repeat and always repeat experiments because some bushes were over-trimmed. This partly agrees with the results of the quantitative evaluation, as in Table 2 the no repetition spheres had the highest percentage of correctly trimmed (CT) points. Furthermore, on each shape the lowest average score is for the smoothness of the bush. It seems that each bush is not trimmed consistently, by having a combination of too deep and insufficiently trimmed patches.

All this data is also summarized in Figure 24a, which shows the single final marks given by the evaluators (criterion #8 in the tables) and Figure 24b, showing the average of the evaluation criteria given by the evaluators (criterion #9 in the tables). The x in both figures is the average of all bushes from that experiment (the combination of shape and treatment). The average scores over all criteria (from Figure 24b) are slightly higher than the values from the final scores given by the evaluators (Figure 24a). This matches the expectation, as an observer likely weights some criteria different from others.

In general, interpreting this data and individual criteria is difficult due to relatively large variation in the underlying data. For example, a standard deviation of 0.8 around a mean of 2.9 leads to a 95% confidence interval of 1.3-4.5, which is almost the full range of scores that could be given (1-5). However, it still provides some insights in the results obtained. For example, bushes can get completely different scoring on individual criteria by the evaluators, indicating that some aspects of the trimming on the bush might be rather OK, while others need clear improvement. When looking to the overall scores, this is also seen with no bushes having a score below 2.0 (indicating that at least something is preserved), while there is only one bush with a score higher than 3.5 (indicating that no bush was trimmed really well). Based on this analysis it is concluded that although some trimming is done, the performance must be improved in order to satisfy potential end-users of such a trimming robot.

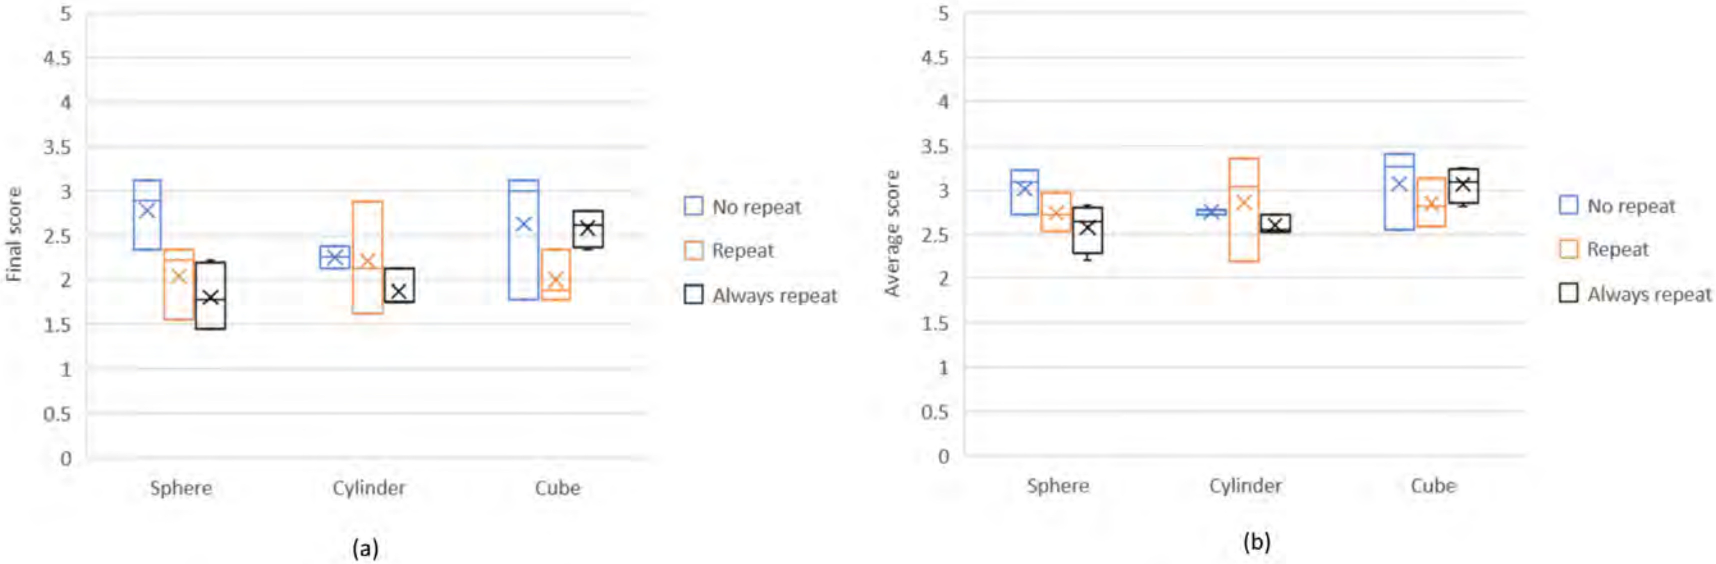

Figure 24: Figure a shows the final mark given by the evaluators, for each target shape and treatment. Figure b shows the average score calculated by averaging each criteria.

Table 7: The following three tables show the average rating and standard deviation for each scored bush. a) spheres, b) cylinders and c) cuboids.

(a)

Spheres		No repetition						Repeat						Always repeat							
		3636		3635		3639		3634		3637		3630		3631		3632		3633		3638	
		rating	st.dev.	rating	st.dev.	rating	st.dev.	rating	st.dev.	rating	st.dev.	rating	st.dev.	rating	st.dev.	rating	st.dev.	rating	st.dev.	rating	st.dev.
1	No deeply trimmed (DT) points	3.9	1.1	3.5	0.7	3.8	0.9	2.8	0.9	2.8	1.1	1.6	0.5	1.5	0.5	1.3	0.5	2.5	1.2	3.4	1.3
2	No poorly trimmed (PT) points	2.6	0.8	2.7	0.7	2.0	0.7	2.9	1.0	2.8	0.9	3.3	1.1	3.4	1.1	3.8	1.1	3.0	1.2	2.7	1.4
3	Bush trimmed into target shape	3.4	1.0	3.5	0.8	3.0	0.7	3.2	1.1	3.0	0.8	1.9	0.7	2.2	0.9	2.7	1.3	3.0	0.9	2.9	0.9
4	Symmetric result	3.6	0.7	3.3	0.8	2.4	1.0	3.0	0.9	2.7	1.2	2.2	1.1	1.7	0.7	1.8	0.9	2.3	0.8	2.5	0.8
5	Bush centred on stem	3.5	0.8	3.9	0.7	2.3	1.3	3.7	1.1	3.0	1.1	3.3	0.7	2.1	0.9	2.8	0.6	3.5	0.7	3.6	0.5
6	No large branches remaining	1.8	1.0	3.0	1.4	3.2	1.1	2.6	1.3	2.6	0.8	3.8	1.0	3.0	0.9	3.8	1.4	2.6	0.8	2.6	1.3
7	Smooth trimming result	2.8	0.4	2.7	0.7	2.3	0.8	2.6	0.7	2.1	0.8	1.6	0.7	1.5	0.5	1.6	0.5	2.2	0.6	2.1	0.7
8	Final mark	**2.9**	0.8	**3.1**	0.8	**2.3**	0.5	**2.3**	1.0	**2.2**	1.0	**1.6**	0.7	**1.4**	0.5	**1.4**	0.5	**2.1**	0.9	**2.2**	0.8
9	Average of criteria 1-7	**3.1**		**3.2**		**2.7**		**3.0**		**2.7**		**2.5**		**2.2**		**2.5**		**2.7**		**2.8**	

(b)

Cylinders		No repetition						Repeat						Always repeat					
		3652		3649		3653		3651		3655		3656		3650		3654		3657	
		rating	st.dev.	rating	st.dev.	rating	st.dev.	rating	st.dev.	rating	st.dev.	rating	st.dev.	rating	st.dev.	rating	st.dev.	rating	st.dev.
1	No deeply trimmed (DT) points	4.1	0.8	2.7	0.9	3.2	0.8	2.9	0.6	3.3	1.3	2.3	0.9	2.0	1.3	2.9	1.1	3.1	0.8
2	No poorly trimmed (PT) points	1.8	0.7	2.7	1.1	2.4	0.9	3.6	0.9	2.8	0.8	2.3	0.9	3.1	0.8	2.6	0.7	2.2	1.2
3	Bush trimmed into target shape	2.9	0.9	2.9	0.9	2.9	0.8	3.4	0.9	2.4	1.5	2.1	0.8	2.4	1.1	2.9	0.9	2.2	0.8
4	Symmetric result	2.8	0.7	2.8	1.2	2.4	0.5	3.6	0.5	3.1	0.9	1.7	0.5	1.9	0.8	2.6	1.1	2.2	1.1
5	Bush centred on stem	3.4	0.9	3.6	0.7	3.0	0.9	4.0	0.5	3.9	0.9	3.3	0.9	3.7	0.7	3.8	0.7	3.8	0.8
6	No large branches remaining	2.0	1.1	2.1	1.1	2.6	1.7	3.4	0.9	3.0	1.0	2.1	1.1	3.0	0.0	2.1	1.1	2.3	1.0
7	Smooth trimming result	2.4	0.7	2.8	0.8	2.4	0.5	2.6	0.7	2.7	1.0	1.4	0.5	1.8	0.7	2.2	1.0	1.8	0.7
8	Final mark	**2.3**	0.7	**2.1**	0.6	**2.4**	0.7	**2.9**	0.6	**2.1**	1.0	**1.6**	0.5	**1.8**	0.7	**2.1**	1.0	**1.8**	0.5
9	Average of criteria 1-7	**2.8**		**2.8**		**2.7**		**3.3**		**3.0**		**2.2**		**2.6**		**2.7**		**2.5**	

(c)

Cuboids		No repetition						Repeat						Always repeat							
		3647		3648		3668		3646		3658		3642		3640		3644		3641		3643	
		rating	st.dev.	rating	rating	rating	st.dev.	rating	st.dev.	rating	st.dev.	rating	st.dev.	rating	st.dev.	rating	st.dev.	rating	st.dev.	rating	st.dev.
1	No deeply trimmed (DT) points	3.8	1.0	4.0	0.8	2.4	1.0	3.3	1.2	2.6	1.1	4.0	0.9	2.9	1.0	3.3	1.1	3.6	1.1	3.1	1.3
2	No poorly trimmed (PT) points	2.5	0.8	2.1	0.9	2.4	1.0	2.6	0.7	3.3	1.2	2.1	0.9	3.0	0.9	2.6	0.8	2.3	0.7	3.0	0.9
3	Bush trimmed into target shape	3.5	0.5	3.3	0.9	2.3	1.0	3.5	0.8	2.8	0.9	2.1	0.7	2.7	0.8	3.1	1.0	3.2	0.6	3.6	0.5
4	Symmetric result	3.4	0.7	3.7	0.5	2.2	0.9	2.9	0.9	2.6	1.0	2.3	0.9	2.8	0.8	3.1	0.7	3.2	0.7	3.2	1.1
5	Bush centred on stem	4.1	0.3	3.9	0.7	2.8	1.0	3.5	0.8	3.6	0.8	2.8	0.8	3.2	0.8	3.8	0.8	3.9	0.6	3.6	0.8
6	No large branches remaining	3.4	1.6	2.8	0.6	3.6	1.0	3.4	0.8	3.0	1.3	2.8	1.1	3.0	0.9	2.6	0.8	3.4	0.8	3.2	0.6
7	Smooth trimming result	3.1	0.9	3.0	0.8	2.1	1.0	2.7	0.7	1.8	0.6	2.0	0.8	2.1	0.6	2.3	0.5	2.9	0.6	3.0	0.8
8	Final mark	**3.0**	0.5	**3.1**	0.8	**1.8**	0.7	**2.3**	0.5	**1.8**	0.7	**1.9**	0.6	**2.3**	0.7	**2.4**	0.5	**2.8**	0.7	**2.8**	0.7
9	Average of criteria 1-7	**3.4**		**3.3**		**2.5**		**3.1**		**2.8**		**2.6**		**2.8**		**3.0**		**3.2**		**3.2**	

2.3 Bush trimming Discussion & Conclusions

From the experimental evaluation, it seems that the robot is not able to trim the bushes with an accuracy that is sufficient for consumers. The results of the manual scoring showed that each bush had a final score smaller than 3.1 on a scale from 1 to 5. In addition, a boxwood grower stated that in their current state, the bushes trimmed by the robot are not saleable. The results of the manual scoring are in agreement with the quantitative results, since the percentage CT points did not exceed the 60% in Table 1, 2 and 3. This indicates that the remaining 40% is either trimmed too deep or not trimmed at all. Despite this, the results show that the robot is able to trim the bushes. Especially, in the sphere experiments an increase in the percentage of CT points took place. However, according to the manual scoring results it seems that a CT score of $>$50% is not yet good enough for a result accepted by the consumer.

In the experiments three types of bushes (spheres, cylinders and cuboids) have been trimmed with three different settings (no repetition, decision based repetition and always repeat). It was expected that with more repetitions the percentage PT points would be lower after trimming than without repetitions. For the spheres and cuboids this is indeed true when using no repetitions (Table 1 and Table 3). This theory is however not applicable for the cylinders. These bushes had a higher percentage of PT points when using more repetitions. Also, the standard deviations of experiments were high. For example at the experiments with repeated trimming of the cuboids the standard deviation for the % PT is 15% (Table 3). Due to the limited amount of bushes tested and high standard deviations observed, it is not reliable to use statistical methods for comparison of approaches and to determine a best approach.

The bushes have been evaluated by reconstructing a point cloud from many images. For each point the distance between the target was determined and the percentage of CT, PT and DT points was calculated. In the calculation of the percentage DT points, an overestimation occurs, since each point being more than 2 cm within the target is considered as trimmed too deep. This assumption is not completely valid, since after removal of the outer leaves more inside branches will become visible and thereby increase the number of points in this area, although they might have been unaffected by trimming due to their natural location inside the bush. In the results, the impact of such overestimation seems small as the photogrammetry method requires multiple images from the same object to create a point, thereby reducing the impact of individual views.

In the evaluation, the trimming accuracy is calculated with respect to both the target shape and a polyhedron based on the robot fitted shapes. It was expected that the percentage of CT points after trimming would be much higher for the polyhedrons compared to the target shape. The results showed however that this is only the case for the cylinders. For cuboids, each side was trimmed from two poses, while the polyhedron was created from the inner one of these, the resulting shape was relative small. Consequently, all cuboids had a relative high percentage of PT points ($>$31%). For the spheres the small difference in the amount of CT points between the target shape and polyhedron is caused by the relative high amount of DT points ($>$29%) with respect to the polyhedron. Part of this can be explained by the robot trimming quite deep. Consequently, also large inner branches are visible in the point clouds, which increase the amount of DT points. In addition, the polyhedron is made by combining multiple poses. Due to translation errors between the fitted shape, the combination of these shapes, and thus also the resulting polyhedron, was not always very accurate. Still, this method proved useful to explain why the final shapes are not in agreement with the desired shapes as shown in Figure 19.

There are some improvements to be made to improve the accuracy of the robot. The first one is the approaching to the bush by the platform. As stated in section 2.1.5, the position of the robot was inaccurate in most cases and needed manual adjustment before trimming, as otherwise the robot was not able to trim the bush at all. Thus, improvements are required here to ensure that the platform is placed properly at the bush and the system can actually perform trimming as desired. Secondly, the results show that the shape fitting has a translation and rotation error for each pose (Figure 21 and 22). Consequently, some bushes like the cylinders were trimmed too deep due to wrong shape fitting (Figure 20). In other words, it seems that a part of the trimming errors observed are related to the shape fitting pipeline. A possible solution would be to trim the robot using an fitted shape from integration of multiple scans around the bush, but this puts even higher demands on the accuracy of platform localisation and positioning. Next to this, some small improvements are desired on the system control, such as proper reporting of deviations and errors.

In the current experiment, bushes were trimmed only once and the trimming system had to deal with significant outgrowth on the bush. In a future application, however, such a system could also exploit the advantages of applying more frequent trimming. Lawn mower robots, for example, perform their task more frequently compared to a human manually mowing the lawn, thereby taking advantage of the lower amount of material that has to be removed. Applying a similar approach could also benefit the trimming performance, as it becomes easier to recognize and trim a bush properly.

In this report the trimming performance of the robot was evaluated by creating a point cloud before and after trimming. Both point clouds are used to calculate the percentage of points correctly trimmed (CT), poorly trimmed (PT) or too deep trimmed (DT). The percentage of correctly CT points after trimming varied between 46-56% for the spheres, between 31-40% for the cylinders and between 51-55% for the cuboids. Each bush was also manually scored by 10 persons. The highest score was 3.1 (out of 5). Due to the fact that the highest score is relatively low and the %CT does not exceed the 60% it can be concluded that although Trimbot demonstrated the concept of automated topiary trimming, it is not yet able to trim either spheres or cylinders or cuboids with the required accuracy for practical application.

3 Rose stem clipping evaluation

3.1 Material and Methods

3.1.1 Test set of rose plants

The rose pruning pipeline was tested using several rose bushes. This includes rose stems placed inside a pot and a real bush in a garden. The thickness of the stems ranged between 0.6 to 1.0 cm. All the tests were performed outdoors in a garden, meaning that the system was tested in a uncontrolled environment. A sample of the plants used in the evaluation is shown in Figure 25.

3.1.2 Evaluation method

The evaluation of the vision module primarily focused on the detection rate, i.e. the ratio of bud-based cutting sites that were detected within 1 cm from the desired (GT) position along the stem. Internally, the module depends on stereo accuracy and segmentation accuracy, which are evaluated in detail in Deliverables 3.4 and 5.4 respectively.

The servoing evaluation measures the success rate of getting the stem inside the cutter and within 1 cm from the target position.

Finally, the stem cutting evaluation counts the number of cases when the upper part of the stem was separated from the lower part.

3.2 Results

A sample of the result of the rose pruning is shown in Figure 26. The process can be observed in a video online: Integrated Rose Bush Trimming (youtu.be/r9IHy5lH8YM).

3.2.1 Clipping sites from scanning

The results for the rose vision pipeline are given in *Deliverable 5.4 - Clipping site recognition*, from which we extract a brief performance overview below.

Branch segmentation. The pixel-wise accuracy of branch segmentation in input images is 82%.

Clipping sites detection. The accuracy from the multi-view scan is 90% of detected targets were within 1 cm from the true position, ie. 54 out of 60 locations were positively detected. There were no false positive detections.

3.2.2 Visual servoing and rotation maneuver

The quality of the visual servoing system was evaluated by letting the robot cutter navigate towards the cutting locations after the scanning. This result was compared against a blind navigation. The blind navigation consist on giving the location of the targets to the planner and navigate towards them without updating the target position. A navigation is considered successful if the robot reaches the target location and the target stem gets into the cutter. The total number of cutting locations found by the robot after scanning the bush were 54 out of 60 real cutting locations. Therefore, the evaluation of the visual servoing process was done using only the 54 locations found by the scanning process. Table 8 shows that blind navigation is not sufficient to drive the end-effector to the target location, mostly because the stems are moved by the wind which causes the cutting locations (targets) to change over time. This makes the end-effector usually end up on one side of the stem. On the other hand, our visual servoing approach is robust enough to make the robot navigate and reach the cutting points 94% of the time under dynamic conditions. The robot needs visual feedback of the current location of the target (in the stem) to navigate successfully to it. If only blind navigation is performed, the end-effector ends up on one of the side of the stem. This is caused mostly by interaction of the cutter with the bush in the final stage of stem approach, ie. the target branch moves because it is connected to another (closer) branch that is pushed away by the cutter. This effect is reduced by the visual servoing.

Stage	Cases	Rate
Test set	60	
Detected (Multi-view scan)	54	90 % of test set
Reached (Visual servoing)	51	94 % of detected
Reached (Blind navigation)	27	50 % of detected
Clipped (Fully cut)	47	93 % of reached 78 % of test set

Table 8: Rose pipeline results.

3.2.3 Stem clipping

The result of stem clipping is that 93% of reached targets were fully cut. This corresponds to 78% overall cutting success rate (see Tab. 8). A successful case is shown in Figure 26.

The failure cases resulted from the stem not getting fully inside of the cutter, then the stem was cut only partially or pushed away by the closing blade (leaving just a scratch). The shallow positioning was caused in some cases by the detector underestimating the distance to the stem, eg. when a leaf got directly between stem and camera. In other cases there were physical obstacles (other parts of the bush) that prevented the cutter from getting closer to the stem, often resulting in pushing away the whole bush or its part.

3.3 Rose stem clipping discussion and Conclusions

The experimental evaluation of the vision part shows that the neural network is capable of segmenting the stems of rose bushes from the background, even when the background and the stems have similar color. This result also demonstrates that the large dataset introduced in Deliverable 2.5 can indeed be used to successfully train a neural network to segment branches of different type of roses.

The proposed target localization approach, which consists of the combined process of stem detection, clustering and point cloud merging can successfully find the cutting points 90% of the times. These targets are found even when they are occluded by other stems or leaves. This approach also proves to be robust but fast enough to be used by the visual servoing to update the target location on the fly.

Scanning the rose bush in a square path is a simple yet effective to capture the structure of the plant. Different scanning methods can be done to improve the scanned bush model, like having different poses instead of a square shape or scanning the bush by navigating around it, however this would lead to further problems like localization and drifting. It can be argued as well that modeling the bush can help to improve the accuracy of the method, however it would lead to slower performance of the system.

Visual feedback is a key element to navigate in a garden where the wind can change the position of the stems, thus change the location of a target. The proposed visual servoing performs a good navigation with an accuracy of 94%. The combination of these steps results in a pipeline capable of finding cutting points in stems that are occluded by other stems or leaves and navigating towards them successfully in $\sim$ 12 secs with an average initial distance between the center of the cutting tool and a target stem of 0.6 m.

The combined success rate from a single view is 78%, which is sufficient to trim the bush completely from several positions around the bush while allowing additional attempts at previously uncut sites. Statistically, after the second attempt 5% of stems could be still uncut, which suggests that an average bush with 5-10 stems will be likely completely cut by then. This is what we also observed in practice.

Figure 25: A sample of the rose bushes used in the evaluation.

Figure 26: Rose bush after pruning. Figure 25 left shows the plant before being cut.

4 Recommendations and Outlook

Despite the limited results of the bush trimming performance, the evaluation showed that the developed robot was able to fully autonomously trim bushes and to clip rose stems. This is a major achievement of the project. For rose bushes, the performance was found acceptable when trimming a bush from multiple sites.

The developed method for the quantitative evaluation of the trimmed bushes is a very powerful instrument to identify the weak points of the system and can potentially also used for other projects that deal with non-rigid and natural objects.

Based on the evaluation experiments, the following improvements are proposed: Improving end-effector design by simplifying the concept and reducing tool weight and dimensions; Improving platform navigation using additional sensors and software. The updated navigation should be able to place the platform within a distance of 73-83 cm towards the bush to prevent manual pose correction; Merging fitted shapes from different platform poses to increase accuracy and reduce translation error between shapes; Improving arm path planning such that trajectories match the bush shape and if there is still an offset (warping error) between fitted shape and planned trajectory automatically restart scanning.

These suggested improvements are expected to enhance the robustness and accuracy, which will make it possible to better determine the influence of repetition settings. At the moment the inaccuracy in the shape fitting pipeline made it difficult to evaluate differences between the repetition settings. In addition, replacing the polyhedrons by the fitted shape makes it possible to determine if the platform trims well at the arm end poses.

References

[1] Motive optical motion capture software, 2019. URL: https://optitrack.com/products/motive/.

[2] Optitrack motion capture system, 2019. URL: https://optitrack.com/.

[3] Prime 13, 2019. URL: https://optitrack.com/products/prime-13/.

[4] Velodyne puck lidar, 2019. URL: https://velodynelidar.com/vlp-16.html.

[5] Trimbot2020 Consortium. Video of integrated topiary bush trimming in the garden, 2019. URL: https://www.youtube.com/watch?v=daUtzo1gew4.

[6] D. Kaljaca, B. Vroegindeweij, and E. van Henten. Coverage Trajectory Planning for a Bush Trimming Robot Arm. *Journal of Field Robotics*, pages 1–26, 2019. doi:10.1002/rob.21917.

A Appendices

A.1 Form for manual scoring

Naam:

Achtergrond:

	Type							Bol					
	Nummer		3629	3632	3633	3638	3636	3631	3630	3637	3634	3635	3639
Struik is te	Volledig onee	1											
diep gesnoeid	oneens	2											
/ heeft te diep		3											
gesnoeide	eens	4											
vlakken	Volledig eens	5											
Struik is te	Volledig onee	1											
ondiep	oneens	2											
gesnoeid /		3											
heeft ondiep	eens	4											
gesnoeide	Volledig eens	5											
Snoeiresultaat	Volledig onee	1											
komt overeen	oneens	2											
met gewenste		3											
vorm	eens	4											
	Volledig eens	5											
Snoeiresultaat	Volledig onee	1											
is symmetrisch	oneens	2											
		3											
	eens	4											
	Volledig eens	5											
Is de struik	Volledig onee	1											
netjes t.o.v de	oneens	2											
stam gesnoeid		3											
(stam in het	eens	4											
midden)	Volledig eens	5											
Hoeveel grote	Geen	1											
takjes steken	Enkele	2											
eruit	Veel	3											
Is het	Zeer onregelm	1											
snoeiresultaat	Onregelmatig,	2											
voldoende	gemiddeld reg	3											
regelmatig/vlo	Regelmatig/ne	4											
eiend	Zeer regelmatig/ zeer netjes	5											
Eind resultaat [1-5] 1=erg													