

# Fast Range Image Segmentation and Smoothing using Approximate Surface Reconstruction and Region Growing

Dirk Holz and Sven Behnke

**Abstract**—Decomposing sensory measurements into relevant parts is a fundamental prerequisite for solving complex tasks, e.g., in the field of mobile manipulation in domestic environments. In this paper, we present a fast approach to surface reconstruction in range images by means of approximate polygonal meshing. The obtained local surface information and neighborhoods are then used to 1) smooth the underlying measurements, and 2) segment the image into planar regions and other geometric primitives. An evaluation using publicly available data sets shows that our approach does not rank behind state-of-the-art algorithms while allowing to process range images at high frame rates.

## I. INTRODUCTION

As robots and autonomous systems move away from laboratory setups towards complex real-world scenarios, both the perception capabilities of these systems and their abilities to acquire and model semantic information must become more powerful. A key issue is the extraction of semantic information from sensory data and its decomposition into parts of interest that are relevant for the tasks of the robot. For mobile manipulation in domestic environments, for example, the perception of objects and their surroundings is a key prerequisite. A common approach [1] in 3D perception is to

- 1) detect horizontal support planes,
- 2) extract and cluster points on top of these planes, and
- 3) perform further processing, e.g., recognizing, classifying or tracking of the found clusters.

Here, one of the fundamental problems is to segment the 3D data into planes and other geometric primitives—or regions of local surface continuity in general.

In this paper, we address the problem of segmenting range images and organized point clouds in real-time on domestic service robots. The central idea of our approach is to approximately reconstruct the surface and segment the range image by growing regions using the resulting local mesh neighborhoods. By means of easily exchangeable components, our generalized region growing approach allows for different region models (e.g. planes) that are segmented in the data. We present models for segmenting planes, regions of local surface continuity, and simple geometric primitives at high frame rates (see Fig. 1).

Furthermore, we use the same mesh neighborhoods to 1) efficiently compute local surface normals and curvature estimates, as well as 2) smooth both the 3D measurements and the computed normals using a bilateral filter.

This research has been partially funded by the FP7 ICT-2007.2.2 project ECHORD (grant agreement 231143) experiment ActReMa.

D. Holz and S. Behnke are with the Autonomous Intelligent Systems Group, University of Bonn, Germany. Contact: dirk.holz@ieee.org.




Fig. 1. Example of surface reconstruction and (plane) segmentation on a RGB-D point cloud: (a) input cloud; (b) constructed mesh; (c) result of segmenting planes in the mesh (red points are assigned to multiple planes); (d) polygonalization as a collection of alpha shapes. Using the Microsoft Kinect camera at a resolution of  $160 \times 120$ , we can compute full polygonalizations with roughly 30Hz on a standard dual-core notebook.

This paper is organized as follows: After giving a brief overview on related work in range image and 3D plane segmentation in Section II, we present our approach in Section III and discuss how to detect geometric primitives using initial segmentations. We use multiple data sets for evaluating the efficiency and robustness of our approach and summarize the results in Section IV.

## II. RELATED WORK

Research on computer and robot vision produced a wide variety of approaches to range image segmentation—and plane segmentation in particular. Hoover et al. [2] compiled a survey and performed an evaluation of early work. Three different types of approaches can be distinguished according to the underlying working principle: methods using random sample consensus (RANSAC), 3D Hough transforms, and region growing. In the context of segmenting 3D laser range scans, another popular approach is to first detect lines in planar cuts, and to merge neighboring lines into local plane patches (see Vosselman et al. [3] for an overview).

### A. Segmentation based on Sample Consensus

RANSAC-based approaches try to find models for geometric primitives that best explain a set of points and the set of inliers supporting it. For segmenting a complete range image, Lee et al. [4] sequentially remove inliers from the original data set, and continue the segmentation with the residual points. Silva et al. [5] first identify connected regions and apply RANSAC region-wise. Gotardo et al. [6] compute an edge map for pre-segmentation and use a variant based on the M-estimator sample consensus (MSAC) to fit model parameters.

Another efficient solution to segmenting even unorganized point clouds and detecting simple geometric primitives such as planes, cylinders and spheres has been proposed by Schnabel et al. [7]. They decompose unorganized point clouds using an octree subdivision and apply RANSAC only to subsets of the original point cloud.

In previous work [8], we adapted the perception scheme from Section I as well as the techniques from [1] and [7], and made them applicable to the measurements of time-of-flight (ToF) cameras. We presented techniques to cope with the specific error sources of the cameras, and to speed up processing by exploiting the image-like data organization. After detecting the most dominant plane, we applied the octree-based primitive detection of Schnabel et al. [7] only to already extracted and segmented points above that plane. In [9], we further speeded up the segmentation process by using integral images for computing local surface normals more efficiently, and using the index neighborhood underlying the 3D data to extract and track segments of points and object candidates, respectively. The overall approach is applicable in real time on a Microsoft Kinect camera and has been used for real-time object tracking and grasp planning [10].

### B. Hough-based Plane Segmentation

The Hough transform is the de-facto standard for finding lines and circles in 2D images. Various extensions to 3D exist that try to find, respectively, planes and maxima in histograms over the possible space of plane orientations and distances. For an overview, and an evaluation for Hough-based segmentation approaches, we refer to the works of Vosselman et al. [3] and Borrmann et al. [11].

RANSAC- and Hough-based segmentation share a common disadvantage. Points belonging to the same segment do not necessarily lie on connected components. Both approaches will merge plane segments if they share a common orientation and distance to the origin. In addition, Hough-based segmentation may suffer from discretization effects.

In [9], we present a fast plane segmentation approach that uses a similar parameter space as the Hough transform. We pre-cluster points and segment planes first in normal space and then, for each cluster, in distance space to obtain individual planes. We compensate for discretization effects by conducting a post-processing step in which neighboring segments are merged if their parameters do not considerably deviate. Still, unconnected planar patches may get merged into the same cluster.

### C. Segmentation using Region Growing

The idea of region growing-based segmentation is to exploit the image-like data structure. Hähnel et al. [12] connect neighboring points in 3D laser range scans to a mesh-like structure. The scans are then segmented recursively by merging connected patches that are likely to lie on the same planar surface. Poppinga et al. [13] apply the same approach to Time-of-Flight cameras and re-formulate the algorithm in an incremental fashion. They grow planar regions by adding neighboring points whose distance to the plane lies below a threshold. Centroid and covariance matrix for estimating the plane parameters are thereby incrementally updated.

Here, we follow a similar approach. Instead of incrementally computing the covariance matrix, however, we compute the normals for all points beforehand and simply average local surface normals to obtain an estimate of the plane's normal. That is, we only store and incrementally update the centroids in both Cartesian and normal space.

Other popular region growing-approaches to range image segmentation make use of local surface curvature. Regions are grown until points with a considerably larger curvature are reached. Just like Gotardo et al. [6] for RANSAC-based segmentation, Harati et al. [14], first compute an edge map to find connected regions of local surface continuity. Rabbani et al. [15] approximate local surface curvature by first fitting planar segments to local point neighborhoods and then computing, for each point, the distance to that plane. Recently, Cupec et al. [16] followed a similar approach. They first apply 2.5D Delaunay triangulation on a range image to obtain an initial triangular mesh and then use the maximum distance of an examined point to all triangles in a region to determine whether or not the point is added.

Here, we deduce a surface reconstruction directly from the image-like data structure, and use the local ring neighborhood around vertices to 1) efficiently compute local surface normals and curvature estimates, and 2) efficiently smooth the depth measurements using a bilateral filter. Our framework allows for using different types of models for region growing, including the approaches of Poppinga et al. [13], Harati et al. [14], Rabbani et al. [15] and Cupec et al. [16], as well as our fast approximation (see Section III).

## III. FAST MESH CONSTRUCTION AND SEGMENTATION

In this section, we describe our approaches to approximate surface reconstruction and segmentation based on region growing. The overall processing pipeline is composed of the following components:

- 1) Deduce approximate mesh from image neighborhood.
- 2) Use mesh neighborhood to compute approximate local surface normals and curvature estimates.
- 3) Bilateral filtering to smooth both points and normals.
- 4) Segmentation based on region growing.

All components described in the above list use the same fixed neighborhoods. These neighborhoods come either directly from the mesh structure or can be pre-computed using a variety of search trees, if the input data is unstructured.




Fig. 2. Fast approximate meshing using a quad mesh (a) and different triangulations (b-d). Compared to the adaptive approach (b), triangulations using only left cuts (c) or only right cuts (d) can be obtained slightly faster.

### A. Exploiting Structure for Fast Approximate Meshing

The central idea of our surface reconstruction approximation is to deduce the desired mesh structure directly from the image-like organization of measurements. In fact, the algorithms presented in the following could easily be applied on local index neighborhoods in range images. However, an approximate mesh allows for 1) applying a wide variety of sophisticated algorithms from the field of computer graphics, as well as for storing certain edge weights, e.g., the difference vectors for integral image-based normal computation from our previous work [9].

We traverse a given range image  $R$  once and check for every point  $\mathbf{p}_i = R(u, v)$ :

- if  $R(u, v)$  and its neighbors  $R(u, v + 1)$ ,  $R(u + 1, v + 1)$ , and  $R(u + 1, v)$  (in the next row and the next column) are valid depth measurements, as well as
- if all edges between  $R(u, v)$  and these three neighbors are not occluded.

The first check is necessary because of the structure in the sensory data that we exploit. If the sensor cannot acquire a valid depth measurement for a certain pixel, it has to store an invalid one, in order to keep the structure organized.

The latter occlusion checks can be easily done by examining the difference vectors between  $\mathbf{p}_i$  and its three neighbors. If it falls into a common line of sight with the viewpoint from where the measurements have been taken, one of the underlying surfaces occludes the other. The condition for the validity of an edge between point  $\mathbf{p}_i$  and its neighbor  $\mathbf{p}_j$  can be formulated as

$$valid = ((|\mathbf{p}_i \cdot \mathbf{p}_j| \leq \cos \varepsilon_\theta) \wedge (\|\mathbf{p}_i - \mathbf{p}_j\|^2 \leq \varepsilon_d^2)), \quad (1)$$

where  $\varepsilon_\theta$  and  $\varepsilon_d$  denote maximum angular and length tolerances, respectively.

If all checks are passed,  $R(u, v)$  and its neighbors are used to extend the so far built mesh. Otherwise, holes arise. Referring to Fig. 2, we distinguish four types of meshes:

- 1) Quad meshes are formed by connecting pixel  $R(u, v)$  to  $R(u, v + 1)$ ,  $R(u + 1, v + 1)$ , and  $R(u + 1, v)$ .
- 2) Fixed left and right cut meshes are formed by cutting quads either from top right to bottom left (left cut) or from top left to bottom right (right cut).
- 3) Adaptive triangulation cuts the quad along the diagonal that has a smaller length.

For triangulations, a single invalid neighbor causes that only one triangle is added. After construction, we simplify the resulting mesh by removing all vertices that are not used in any polygon. An example triangulation is shown in Fig. 3(b).

### B. Fast Computation of Surface Normals and Curvature

We compute the local surface normal  $\mathbf{n}_i$  for point  $\mathbf{p}_i$  as the weighted average of the plane normals of the faces surrounding  $\mathbf{p}_i$ . Using the cross product between the difference vectors of the bounding vertices to compute the face normals, and choosing the weights to be proportional to the area of triangles, removes the need of normalizing the face normals beforehand. Thus, we can obtain  $\mathbf{n}_i$  as:

$$\mathbf{n}_i = \frac{\sum_{j=0}^{N_T} (\mathbf{p}_{j,a} - \mathbf{p}_{j,b}) \times (\mathbf{p}_{j,a} - \mathbf{p}_{j,c})}{\|\sum_{j=0}^{N_i} (\mathbf{p}_{j,a} - \mathbf{p}_{j,b}) \times (\mathbf{p}_{j,a} - \mathbf{p}_{j,c})\|}, \quad (2)$$

where  $\mathbf{p}_{j,a}$ ,  $\mathbf{p}_{j,b}$  and  $\mathbf{p}_{j,c}$  form triangle  $j$ . In the actual implementation, we simply iterate over the faces, compute the difference vectors and their cross product, and add them to the normals of the involved points. Finally, we normalize all point normals at once. An example for computed local surface normals (color coded) can be seen in Fig. 3(b).

### C. Bilateral Filtering

Naturally, sensor measurements are affected by noise. Since this noise can hinder further processing like segmentation, we apply a bilateral filter for smoothing both the points and their normals while preserving edges in the sensed geometric structures. Again, instead of searching, we directly extract a point's neighborhood from the mesh. That is, we filter both a point  $\mathbf{p}_i$  and its normal  $\mathbf{n}_i$  over its 1-ring-neighborhood  $N_i$ , i.e., all points that are directly connected to  $\mathbf{p}_i$  by an edge in the mesh:

$$\mathbf{p}_i = \sum_{j \in N_i} w_{ij} \mathbf{p}_j / \sum_{j \in N_i} w_{ij}, \quad \mathbf{n}_i = \sum_{j \in N_i} w_{ij} \mathbf{n}_j / \sum_{j \in N_i} w_{ij}, \quad (3)$$

$$w_{ij} = \underbrace{e^{-\|\mathbf{p}_i - \mathbf{p}_j\|}}_{\text{distance term}} \underbrace{e^{-\|\mathbf{n}_i - \mathbf{n}_j\|}}_{\text{normal term}} \underbrace{e^{(I_i - I_j)/c_I}}_{\text{intensity term}}, \quad (4)$$

where the optional intensity term is only evaluated for colored point clouds and range images where also an intensity image is available. The normalization constant  $c_I$  is used to scale the intensity differences to lie in the interval  $[0, 1]$ . An example of filtering an input mesh can be seen in Fig. 3(c).

### D. Region Growing-based Segmentation

Despite the generalization over different neighborhood searches and region models, the implementation of our segmentation algorithm does not considerably deviate from other region growing algorithms in the literature. Given a set of seed points (and a priority queue of seeds),

Outer loop, until all points are processed:

- 1) we iteratively select the next seed point,
- 2) initialize the region model of interest, and
- 3) put the seed point onto the empty processing queue.

Inner loop, while the processing queue is not empty:

- 4) We take the next point from the processing queue,
- 5) check its compatibility with the region model,
- 6) and add it in case of compatibility.
- 7) We add the point's neighbors to the processing queue (again, only if they're compatible).




Fig. 3. Approximate surface reconstruction, bilateral filtering and segmentation on an example point cloud (a). The filtered mesh (c) is considerably smoother than the initially approximated mesh (b) and allows for cleanly segmenting regions of local surface continuity (d). The meshes in (b) and (c) are colored w.r.t. the vertices’ normal orientation (mapped to RGB space, object-space normal map). In (d) segments are colored randomly.

### E. Different Region Models for Segmentation

We have encapsulated the processing steps (2) initialization, (5) point compatibility, (6) model update, and (7) neighbor compatibility in exchangeable region models allowing to configure and control the behavior of the segmentation.

We have implemented several region models for approximate plane segmentation using local surface normals, as well as for segmenting regions of local surface continuity as a pre-segmentation for further processing. Both are detailed in the following. In addition, we added models that reproduce the behavior of segmentation algorithms found in the literature.

1) *Approximate Plane Segmentation*: For plane segmentation, we initialize the centroid and the normal of the region model using the seed point and its normal (instead of computing them using an initial set of points as in the approach of Poppinga et al. [13]). In order to determine the compatibility of a point  $\mathbf{p}_i$  to the model, we simply check the angle between its normal  $\mathbf{n}_i$  and the normal in the plane model, as well as  $\mathbf{p}_i$ ’s distance to the plane. For updating the model w.r.t.  $\mathbf{p}_i$ , we incrementally update the plane’s centroid, but instead of incrementally updating a covariance matrix to derive a plane normal from it, we incrementally update its centroid in normal space. That is, by pre-computing the surface normals on the mesh neighborhood, and approximating the plane normal by averaging over point normals, we can reduce the number of computations considerably.

In order to obtain full polygonalizations as the one shown in Fig. 1(d), we 1) compute the convex or concave hulls (using alpha shapes) for all planar patches, and, if a triangulation is required, 2) decompose them again using ear clipping.

2) *Extracting Locally Smooth Surfaces*: For detecting geometric primitives, we need a rough pre-segmentation of the scene. This can easily be accomplished within the neighbor compatibility check of the models by, e.g., examining changes in the local surface curvature or comparing a point’s surface normal with either the region’s mean surface normal or the surface normal of the seed point. A typical result of applying a segmentation using the latter model is shown in Fig. 3(d). Points on the same physical (locally smooth) surface end up in the same segment.

### F. Detecting Geometric Primitives

For every locally smooth segment, we try to find the geometric primitive that best explains the underlying point set. Whereas we can directly compute a least squares plane fit to all the points in a segment, we use RANSAC to find




Fig. 4. Examples of detecting planes (yellow), cylinders (cyan) and spheres (magenta). Points and polygons belonging to multiple segments are colored red. All points are projected onto the found models.

the best sphere and cylinder model. Here, the computational efficiency of our approach comes from applying RANSAC only if the computed planar model does not fully explain the segment. Typical results of applying the rough pre-segmentation and primitive detection are shown in Fig. 4.

Since we still stick to the RANSAC-based primitive detection for spheres and cylinders, we focus the experimental evaluation to plane segments directly obtainable from region segmentation and extracted using our approximate model.

## IV. EXPERIMENTS AND RESULTS

We evaluate the correctness and efficiency of our approach using two publicly available data sets for which ground truth plane segmentations are available: the SegComp data set from Hoover et al. [2], and the recently published Kinect data set by Oehler et al. [17]. For the evaluation, we follow (and refer to) the scheme of [2]. Measured runtimes of the individual processing steps show that our approach can process range images at high frame rates—at a resolution of  $160 \times 120$  in real-time with 30Hz (see Table I). Compared the approach of Oehler et al. [17], we obtain better segmentation results (see Table II) while being a faster by a factor of 4 (they need  $>100\text{ms}$  for  $160 \times 120$ , and  $>2\text{s}$  for  $640 \times 480$ ).

### A. Estimating a Simple Isotropic Noise Model

For all the aforementioned region models, parameters like, for instance, the distance to the model and the deviation between surface normals play an important role. Whereas the latter can be neglected after applying the bilateral filter, the distance to the model is a parameter that is crucial for the quality of the segmentation. It resembles the amount of noise hindering a measurement to lie on the ideal model.

In order to obtain a rough estimate of the amount of noise at a given point, we use a simple isotropic noise

model. As suggested in [18], we assume Gaussian noise  $\mathcal{N}(0, \sigma^2)$  and use a simple quadratic polynomial of distance to determine  $\sigma$ , since noise in range sensors usually increases quadratically with the measured distance. Since the primary sensor used in our work is a Microsoft Kinect camera, we have computed a simple error model for this sensor. In 10 different scenes (ranging from scenes with only close range measurements to views of wide open space), we have collected 100 range images each. For each of the locations, we compute the mean and standard deviation per pixel and perform a least squares fit to find appropriate coefficients for the quadratic model; resulting in:

$$\sigma(d) = 0.00263d^2 - 0.00518d + 0.00752. \quad (5)$$

For the range images in the SegComp data set, that are slightly more affected by noise, we simply set  $2\sigma(d)$  as the maximally allowed deviation of a point to its region model.

### B. Plane Segmentation using the SegComp Data Sets

Typical results of applying the presented approximate plane segmentation on range images and organized point clouds can be seen in Figures 5 and 6. Considering our goal of obtaining a fast decomposition into dominant planes and other objects of interest, the obtained results are more than satisfying. Moreover, as can be seen in the detailed comparison in Table II, the proposed method does not considerably rank behind state-of-the-art range image segmentation approaches, while providing very efficient means to compute—within milliseconds—rough scene segmentations.

On the ABW data set [2], our approximate plane segmentation approach tends to over-segment the range image. This is caused by a special characteristic of the used camera that is not explicitly handled here resulting in inconsistent normal orientations. Besides over-segmented planar patches, our approach correctly detects 80% of the planes.

In the PERCEPTRON data set, no special sensor characteristics cause errors and our approach yields similar results as the work by Gotardo et al. [6]. Referring to Table II, over-segmentations are rare for our approach. Instead, a considerable amount of ground truth planes is not perceived. These planes are formed by only a small number of points and are neglected here due a minimum region cardinality that we use to eliminate very small planar patches.

We also evaluated our approach using the recently published Kinect data set by Oehler et al. [17]. This data set comprises two different sets of point clouds with corresponding ground truth segmentations – one for plane segmentation and one for cylinder segmentation. Naturally, we obtain a larger amount of over-segmentations here, since larger cylinders in the planes data set are labeled as noise while unconnected regions belonging to a single plane are labeled as one segment. However, visually inspecting the segmentation results reveals that all dominant planes are reliably segmented.

## V. CONCLUSIONS AND FUTURE WORK

We have presented a simple, yet efficient approach to segmenting range images and organized point clouds. Using

TABLE I  
MEASURED RUNTIMES\* FOR THE INDIVIDUAL PROCESSING STEPS

Resolution	160 × 120	320 × 240	640 × 480
Mesh Construction	11 ms	45 ms	188 ms
Computing Normals	2 ms	7 ms	33 ms
Smoothing/Filtering	10ms	38 ms	155 ms
(Plane) Segmentation	6ms	30ms	126ms
Overall Frequency	≈ 35Hz	≈ 8Hz	≈ 2Hz

\*Measured over 1000 runs on an Intel Core 2 DUO 2.26GHz (no parallelization), approximate plane segmentation (III-E.1) of a mesh constructed using adaptive triangulation (III-A).

an approximate polygonal reconstruction directly deduced from the image-like structure, we are able to efficiently compute point features such as local surface normals, and to smooth the measured data using a bilateral filter.

Experimental evaluation has shown that our approach does not considerably rank behind state-of-the-art range image segmentation techniques. However, regarding its computational complexity, it exploits several simplifications and approximations finally making the overall segmentation (and primitive detection) algorithm run within milliseconds on point clouds acquired by typical RGB-D cameras.

It remains a matter of future work to exploit the extracted information about segmented planar patches (and geometric primitives) in further processing steps like registration.

Implementations of all components presented in this paper are (or are going to be) publicly available within the open source Point Cloud Library PCL<sup>1</sup>.

## REFERENCES


- [1] R. B. Rusu, N. Blodow, Z. C. Marton, and M. Beetz, “Close-range scene segmentation and reconstruction of 3D point cloud maps for mobile manipulation in human environments,” in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, St. Louis, MO, USA, 2009, pp. 1–6.
- [2] A. Hoover, G. Jean-Baptiste, X. Jiang, P. J. Flynn, H. Bunke, D. B. Goldgof, K. Bowyer, D. W. Eggert, A. Fitzgibbon, and R. B. Fisher, “An experimental comparison of range image segmentation algorithms,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 18, pp. 673–689, 1996.
- [3] G. Vosselman, B. G. H. Gorte, G. Sithole, and T. Rabbani, “Recognising structure in laser scanner point clouds,” *Information Sciences*, vol. 46, no. 8, pp. 1–6, 2004.
- [4] K.-M. Lee, P. Meer, and R.-H. Park, “Robust adaptive segmentation of range images,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 20, pp. 200–205, 1998.
- [5] L. Silva, O. Bellon, and P. Gotardo, “A global-to-local approach for robust range image segmentation,” in *Proc. of the Int. Conference on Image Processing (ICIP)*, Rochester, NY, USA, 2002, pp. 773–776.
- [6] P. Gotardo, O. Bellon, and L. Silva, “Range image segmentation by surface extraction using an improved robust estimator,” in *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Madison, WI, USA, 2003, pp. 33–38.
- [7] R. Schnabel, R. Wahl, and R. Klein, “Efficient RANSAC for point-cloud shape detection,” *Computer Graphics Forum*, vol. 26, no. 2, pp. 214–226, 2007.
- [8] D. Holz, R. Schnabel, D. Droschel, J. Stückler, and S. Behnke, “Towards semantic scene analysis with time-of-flight cameras,” in *Proc. of the 14th RoboCup International Symposium*, Singapore, 2010.
- [9] D. Holz, S. Holzer, R. B. Rusu, and S. Behnke, “Real-time plane segmentation using RGB-D cameras,” in *Proc. of the 15th RoboCup International Symposium*, Istanbul, Turkey, 2011.

<sup>1</sup>The latest stable release of the Point Cloud Library PCL is available at <http://pointclouds.org>.



(a) 10 / 14, 4 misses (b) 18 / 24, 6 misses (c) 10 / 13, 2 over, 1 miss (d) 12 / 13, 1 over (e) 22 / 22

Fig. 5. Plane segmentation using the SegComp PERCEPTOR training data set (segments randomly colored). In total, 109 of 126 planes were correctly segmented (approx. 86.5%). Not correctly found are very small plane segments, e.g., the inner parts of the objects in (a) and (b). In addition, some planes are oversegmented due to noise, e.g., the support plane in (c). The estimated plane normals deviate from ground truth by roughly  $(2.5 \pm 1.6)^\circ$ .



(a) 10 / 11 (b) 9 / 10 (c) 9 / 9 (d) 5 / 8, 3 over (e) 5 / 9, 2 over, 2 miss

Fig. 6. Approximate plane segmentation using (a-c) the SegComp ABW data set [2], and (d-e) the data set by Oehler et al. [17]. ABW: On average, 12.2 planes out of 15.2 planes are correctly segmented, while roughly two planes per image are oversegmented. Kinect: for the data set from [17] over-segmentations are primarily caused by detecting planes on the surface of cylinders or in small regions not marked as belonging to planes in ground truth. Overall, roughly 58% of planes are correctly segmented (80% pixel overlap), while Oehler et al. only achieve 54.9% (at 51% pixel overlap, see [17]).

TABLE II  
DETAILED BENCHMARKING RESULTS ON THE SEGCOMP DATA SETS FOR PLANE SEGMENTATION

Approach	Regions in ground truth	correctly detected	orientation deviation	over-segmented	under-segmented	missed (not detected)	noise (non-existent)
SegComp ABW data set (30 test images) by Hoover et al. [2], assuming 80% pixel overlap as in [6]							
USF [6]	15.2	12.7 (83.5%)	1.6°	0.2	0.1	2.1	1.2
WSU [6]	15.2	9.7 (63.8%)	1.6°	0.5	0.2	4.5	2.2
UB [6]	15.2	12.8 (84.2%)	1.3°	0.5	0.1	1.7	2.1
UE [6]	15.2	13.4 (88.1%)	1.6°	0.4	0.2	1.1	0.8
OU [6]	15.2	9.8 (64.4%)	–	0.2	0.4	4.4	3.2
PPU [6]	15.2	6.8 (44.7%)	–	0.1	2.1	3.4	2.0
UA [6]	15.2	4.9 (32.2%)	–	0.3	2.2	3.6	3.2
UFPR [6]	15.2	13.0 (85.5%)	1.5°	0.5	0.1	1.6	1.4
Oehler et al. [17]	15.2	11.1 (73.0%)	1.4°	0.2	0.7	2.2	0.8
Ours	15.2	12.2 (80.1%)	1.9°	1.8	0.1	0.9	1.3
SegComp PERCEPTOR data set (30 test images) by Hoover et al. [2], assuming 80% pixel overlap as in [6]							
USF [6]	14.6	8.9 (60.9%)	2.7	0.4	0.0	5.3	3.6
WSU [6]	14.6	5.9 (40.4%)	3.3	0.5	0.6	6.7	4.8
UB [6]	14.6	9.6 (65.7%)	3.1	0.6	0.1	4.2	2.8
UE [6]	14.6	10.0 (68.4%)	2.6	0.2	0.3	3.8	2.1
UFPR [6]	14.6	11.0 (75.3%)	2.5	0.3	0.1	3.0	2.5
Oehler et al. [17]	14.6	7.4 (50.1%)	5.2	0.3	0.4	6.2	3.9
Ours	14.6	11.0 (75.3%)	2.6	0.4	0.2	2.7	0.3

- [10] J. Stückler, R. Steffens, D. Holz, and S. Behnke, “Real-time 3D perception and efficient grasp planning for everyday manipulation tasks,” in *Proc. of the European Conference on Mobile Robots (ECMR)*, Örebro, Sweden, 2011, pp. 177–182.
- [11] D. Borrmann, J. Elseberg, K. Lingemann, and A. Nuechter, “The 3D Hough transform for plane detection in point clouds: A review and a new accumulator design,” *3D Research*, vol. 2, pp. 1–13, 2011.
- [12] D. Hähnel, W. Burgard, and S. Thrun, “Learning compact 3D models of indoor and outdoor environments with a mobile robot,” *Robotics and Autonomous Systems*, vol. 44, no. 1, pp. 15–27, 2003.
- [13] J. Poppinga, N. Vaskevicius, A. Birk, and K. Pathak, “Fast plane detection and polygonalization in noisy 3D range images,” in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Nice, France, 2008, pp. 3378–3383.
- [14] A. Harati, S. Gächter, and R. Siegwart, “Fast range image segmentation for indoor 3D-SLAM,” in *Proc. of the IFAC Symposium on Intelligent Autonomous Vehicles (IAV)*, Toulouse, France, 2006.
- [15] T. Rabbani, F. van den Heuvel, and G. Vosselman, “Segmentation of point clouds using smoothness constraint,” *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. 36, pp. 248–253, 2006.
- [16] R. Cupec, E. K. Nyarko, and D. Filko, “Fast 2.5D mesh segmentation to approximately convex surfaces,” in *Proc. of the European Conference on Mobile Robots (ECMR)*, Örebro, Sweden, 2011, pp. 49–54.
- [17] B. Oehler, J. Stückler, J. Welle, D. Schulz, and S. Behnke, “Efficient multi-resolution plane segmentation of 3D point clouds,” in *Proc. of the International Conference on Intelligent Robotics and Applications (ICIRA)*, Aachen, Germany, 2011, pp. 145–156.
- [18] D. Anderson, H. Herman, and A. Kelly, “Experimental characterization of commercial flash lidar devices,” in *International Conference of Sensing and Technology*, Palmerston North, New Zealand, 2005.