Using Hierarchical EM to Extract Planes from 3D Range Scans

Rudolph Triebel and Wolfram Burgard
Department of Computer Science
University of Freiburg
George-Koehler-Allee 79, 79108 Freiburg, Germany
{*triebel,burgard*}*@informatik.uni-freiburg.de*

Frank Dellaert
College of Computing
Georgia Institute of Technology
801 Atlantic Drive, Atlanta, USA
dellaert@cc.gatech.edu

***Abstract*— Recently, the acquisition of three-dimensional maps has become more and more popular. This is motivated by the fact that robots act in the three-dimensional world and several tasks such as path planning or localizing objects can be carried out more reliable using three-dimensional representations. In this paper we consider the problem of extracting planes from three-dimensional range data. In contrast to previous approaches our algorithm uses a hierarchical variant of the popular Expectation Maximization (EM) algorithm [1] to simultaneously learn the main directions of the planar structures. These main directions are then used to correct the position and orientation of planes. In practical experiments carried out with real data and in simulations we demonstrate that our algorithm can accurately extract planes and their orientation from range data.**

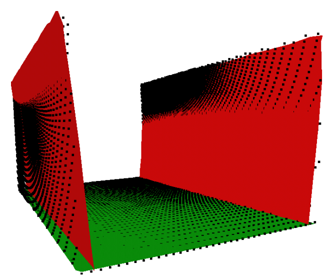

Fig. 1. Birds-eye view of a data set and its planar approximation obtained with the standard EM-based clustering approach. The angle between the normals of the two planes belonging to the parallel walls is 17.2 degrees.

I. Introduction

Whereas mobile robots act in the three-dimensional world, most of the research regarding spatial representations of the environment of mobile robots has focused on two-dimensional maps. The restriction to two-dimensional representations, however, is error-prone and has serious limitations. For example, the planning of paths can be incomplete if the three-dimensional world is mapped into two dimensions or even incorrect if not all obstacles are contained in the two-dimensional description. Additionally, two-dimensional representations do not support typical tasks like searching for objects. For example, without knowledge about the three-dimensional structure of a shelf, a robot cannot plan appropriate viewpoints to find an object in the shelf. Thus, two-dimensional maps are not sufficient in situations in which robots are deployed in real-world scenarios. On the other hand, three-dimensional models of buildings (exterior and interior) and man-made objects are envisioned to be useful in a wide area of applications, which goes far beyond robotics, like architecture, emergency planning, visualization etc. In all of these application domains, there is a need for methods that can automatically construct three-dimensional models.

The major disadvantage of three-dimensional representations lies in the amount of computational resources needed to store and to update them. In this paper we therefore study the problem of approximating three-dimensional range data by planes. Especially indoor environments typically contain many planar structures such as walls, floors, ceilings, tables etc. Whereas in the past several techniques have been applied to extract planes from range data, these approaches ignore that many planar structures in man-made environments are parallel. This can lead to errors in the resulting models, for example, when walls end in corners or when objects are placed on tables. As an example, consider the situation depicted in Figure 1. This figure shows a birds-eye view on a simulated data set with three orthogonal walls. It also shows the planes extracted with the standard approach for clustering linear structures with EM [2]. As can be seen, the resulting planes for the two parallel walls are not parallel. In this paper we present an approach to correct such errors in the approximation by using the hierarchical EM to simultaneously learn planes and their main directions from range data. In the maximization step of this algorithm we simultaneously maximize the parameters of the planes with respect to the data points and the main directions. As a result, the finally obtained planes are more accurate as those obtained with previous approaches, which do not consider main directions.

The problem of constructing 3d-models of buildings (exterior and interior) and man-made objects has received considerable attention over the past few years. For example, Bajcsy *et al.* [3], Hakim and Boulanger [4], as well as Rous and colleagues [5] reconstruct three-dimensional structures from camera images. Furthermore, Sabe *et al.* use a variant of the Hough transform to extract planes from stereo images. Recently, several authors used 3d range scanners for the acquisition of volumetric models. For example, Stamos and Leordeanu [6] construct 3d-models by combining multiple views obtained with a 3d range scanner. Früh and

Zakhor [7] present a technique to learn accurate models of large-scale outdoor environments by combining laser, vision, and aerial images. Thrun *et al.* [8] use two 2d range scanners. The first is oriented horizontally whereas the second points toward the ceiling. By registering the horizontal scans the system generates accurate three-dimensional models. In a more recent work [9] several range scanners were used to learn models of underground mines. Nüchter and colleagues [10] developed a robot that is able to autonomously explore non-planar environments and to simultaneously acquire the three-dimensional model. Also several authors focused on the problem of extracting planar structures from range scans. For example, Hähnel *et al.* [11] use a region growing technique to identify planes. Recently, Liu *et al.* [2] as well as Martin and Thrun [12] applied the EM algorithm to cluster range scans into planes.

The approach described in this paper extends the EM-based techniques mentioned above by simultaneously estimating the planes and the clusters of their main direction and by incorporating the information about the major directions into the learning process. The approach borrows some ideas of the work of Coughlan and Yuille [13] as well as Schindler and Dellaert [14]. Both approaches estimate vanishing points in images and use this information for edge clustering. The paper presented here extends these techniques in two respects. We apply the hierarchical EM to three-dimensional range data and furthermore our algorithm is able to automatically estimate the number of main directions and planes from the data.

II. A Likelihood Model for Points, Planes, and Main Directions

Suppose we are given a set of N scan points $Z = \{\mathbf{z}_n\}$ with $\mathbf{z}_n \in \mathbb{R}^3,\ n = 1, \ldots, N$. Furthermore suppose there is a set Θ of M planes $\theta_m, m = 1, \ldots, M$ as well as a set Φ of K main directions $\phi_k, k = 1, \ldots, K$ of these planes. To correctly cluster the data points into planes and the planes into main directions, it is useful to introduce so-called correspondence variables α_{nm}, which take on a value of 1 if a data point z_n belongs to plane m and 0 otherwise, and similarly variables β_{mk} to indicate that a plane m belongs to main direction k. We collect these correspondence variables in the sets $A = \{\alpha_{nm}\}$ and $B = \{\beta_{nk}\}$. Our goal is to maximize the joint posterior $p(Z, \Theta, \Phi, A, B)$. Exploiting the independence between these variables this term can be rewritten as (see [15])

$$p(A, B, \Phi, \Theta, Z) \;\propto\; p(Z \mid A, \Theta) p(\Theta \mid B, \Phi) \tag{1}$$

$$\propto\; p(Z, A \mid \Theta) p(\Theta, B \mid \Phi). \tag{2}$$

It remains to describe how the individual terms $p(Z, A \mid \Theta)$ and $p(\Theta, B \mid \Phi)$ are computed. The first term specifies the likelihood of the data and the correspondence variables given the planes. Let us assume a plane is given as a tuple $\theta_m = (\mathbf{n}_m, d_m)$ where $\mathbf{n}_m$ is a unit normal vector and d_m is the Euclidean distance of the plane from the origin. For now, we assume the number M of planes is given. We define the distance $d_1(\mathbf{z}_n, \theta_m)$ of a scan point $\mathbf{z}_n$ from a plane θ_m as the Euclidean distance:

$$d_1(\mathbf{z}_n, \theta_m) \;=\; \mathbf{z}_n \cdot \mathbf{n}_m - d_m \tag{3}$$

If we assume Gaussian noise with variance ρ in the measurement processes we can write

$$p(z_n \mid \theta_m) \;=\; \frac{1}{\sqrt{2\pi\rho^2}} \exp\left\{-\frac{1}{2}\left(\frac{d_1(\mathbf{z}_n, \theta_m)}{\rho}\right)^2\right\}. \tag{4}$$

If we additionally incorporate the correspondence variables α_{nm} and under the assumption that their distribution is uniform (for details see [12]) we obtain

$$p(Z, A \mid \Theta) \propto \exp\left\{-\frac{1}{2} \sum_n \sum_m \alpha_{nm} \left(\frac{d_1(\mathbf{z}_n, \theta_m)}{\rho}\right)^2\right\} \tag{5}$$

The second term on the right side of Equation 1 specifies the likelihood of the planes and the correspondence variables given the main directions. To determine this quantity we proceed in the same way as above. A main direction for a set of planes is defined as a 3D unit vector. We define the distance of a plane θ_m to a main direction φ_k as

$$d_2(\theta_m, \varphi_k) \;=\; \sqrt{1 - (\mathbf{n}_m \cdot \varphi_k)^2}. \tag{6}$$

This corresponds to the sine of the angle between the normal vector $\mathbf{n}_n$ and the main direction φ_k. Again we assume that the planes belonging to a main direction are normally distributed with variance σ. Under the assumption that we know that plane θ_m belongs to main direction φ_k we can calculate its likelihood as

$$p(\theta_m \mid \varphi_k) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left\{-\frac{1}{2}\left(\frac{d_2(\theta_m, \varphi_k)}{\sigma}\right)^2\right\}. \tag{7}$$

In analogy to the derivation above we obtain

$$p(\Theta, B \mid \Phi) \propto \exp\left\{-\frac{1}{2} \sum_m \sum_k \beta_{mk} \left(\frac{d_2(\theta_m, \varphi_k,)}{\sigma}\right)^2\right\}. \tag{8}$$

This leads to this expression for the joint posterior:

$$p(A, B, \Phi, \Theta, Z) \;\propto\; \exp\left\{-\frac{1}{2} \sum_{nm} \alpha_{nm} \left(\frac{d_1(\mathbf{z}_n, \theta_m)}{\rho}\right)^2 - \frac{1}{2} \sum_{mk} \beta_{mk} \left(\frac{d_2(\theta_m, \varphi_k)}{\sigma}\right)^2\right\} \tag{9}$$

Our goal is to determine the model (Θ^*, Φ^*) that maximizes the likelihood of the data Z. Since, however, the values of the correspondence variables are unknown, we apply the EM algorithm which iteratively maximizes the expected log likelihood of the data and the model:

$$(\Theta^{[i+1]}, \Phi^{[i+1]}) = \underset{(\Theta,\Phi)}{\operatorname{argmax}}\, E_{AB}\left[\log p(A, B, \Phi, \Theta, Z) \mid \Theta^{[i]}, \Phi^{[i]}\right] \tag{10}$$

Inserting the expression for the posterior into this equation and exploiting linearity of the expectation we obtain

$$
\begin{aligned}
(\Theta^{[i+1]}, \Phi^{[i+1]}) = \underset{(\Theta,\Phi)}{\operatorname{argmax}} \\
-\frac{1}{\rho^2} \sum_{n=1}^{N} \sum_{m=1}^{M} E[\alpha_{nm} \mid \Theta^{[i]}](\mathbf{z}_n \cdot \mathbf{n}_m - d_m)^2 \\
-\frac{1}{\sigma^2} \sum_{m=1}^{M} \sum_{k=1}^{K} E[\beta_{mk} \mid \Phi^{[i]}](1 - (\mathbf{n}_m \cdot \varphi_k)^2) \quad (11)
\end{aligned}
$$

The expectations $E[\alpha_{nm} \mid \Theta]$ are computed as follows:

$$
\begin{aligned}
E[\alpha_{nm} \mid \Theta] &= p(\alpha_{nm} \mid z_n, \Theta) & (12) \\
&= \frac{p(z_n \mid \alpha_{nm}, \Theta) p(\alpha_{nm} \mid \Theta)}{p(z_n \mid \Theta)} & (13) \\
&= \frac{\exp\left\{-\frac{1}{2}\left(\frac{d_1(\mathbf{z}_n,\theta_m)}{\rho}\right)^2\right\}}{\sum_j \exp\left\{-\frac{1}{2}\left(\frac{d_1(\mathbf{z}_n,\theta_j)}{\rho}\right)^2\right\}} & (14)
\end{aligned}
$$

Similarly, the $E[\beta_{mk} \mid \Phi]$ are obtained as

$$
E[\beta_{mk} \mid \Phi] = \frac{\exp\left\{-\frac{1}{2}\left(\frac{d_2(\theta_m,\varphi_k)}{\sigma}\right)^2\right\}}{\sum_j \exp\left\{-\frac{1}{2}\left(\frac{d_2(\theta_j,\varphi_k)}{\sigma}\right)^2\right\}} \quad (15)
$$

In the M-step, we want to find new model parameters Θ and Φ so that the log likelihood function in Equation (11) is maximized. In our current implementation we apply the Fletcher-Reeves conjugate gradient algorithm to find a local maximum of the log likelihood function.

III. Estimating the Model Complexity

So far, we assumed that the number of planes M and the number of main directions K were given in advance. In practice, however, this is generally not the case. Instead, we need to estimate M and K – we will call this the *model complexity* – during the estimation process. In our implementation we apply the *Bayesian Information Criterion* [16], which is calculated as follows:

$$
BIC = -2L + (3M + 2K)\ln(N) \quad (16)
$$

In this equation, L is the log likelihood of the data given the current model where the factor -2 stems from the BIC formula. The term $3M + 2K$ corresponds to the number of free parameters (3 for each plane and 2 for each main direction). The goal is to find a model which has a minimal BIC value. As can be seen from Equation (16), a high model complexity results in a large BIC value and hence less complex models are preferred.

To minimize the BIC we constantly monitor its value. High BIC-values, for example, result from redundancy in the model. For example, it is possible that after convergence of the EM-algorithm two planes are equal. This can happen if the two planes are initialized too close to each other or if the data only supports a smaller number of planes. Such a case of redundancy can be detected applying the leave-one-out rule: after convergence of the EM we calculate the BIC for all possible models, in which one plane is left out. If there is a model that has a smaller BIC than the current one, then the plane, which has been left out to obtain this particular model, must be redundant and can safely be removed. The same strategy is applied to the main directions.

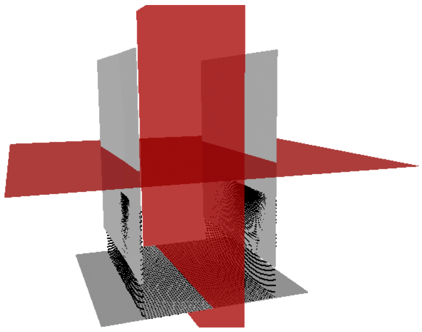

Fig. 2. Planes obtained by the initialization process for the data set shown in Figure 1. In this case our algorithm was initialized with five planes and five main-directions from which several were co-planar.

IV. Implementation Details

A. Initialization

Since the EM-algorithm can get stuck in local optima of the log-likelihood function, it needs to be initialized appropriately to converge to the global optimum, just like other optimization techniques, e.g., gradient-descent. In our case, the initialization is performed by sampling randomly from all scan points. During this sampling process each point is drawn with a probability proportional to the minimum distance to any of the already existing planes. Thus, when no planes are given each point is equally likely to be drawn. However in later initialization steps, the points that are badly explained by the current model, are more likely to be drawn. To initialize a plane for a point drawn we fit a plane to the points in its local neighborhood. A typical result of the plane initialization for the data depicted in Figure 1 is shown in Figure 2. Here, the algorithm was initialized with five planes and five main directions. Since the planes four and five and the main-directions three to five are co-planar to existing planes and main-directions, only three planes and two main-directions are visible.

B. Weighting Factors for Planes

When clustering the planes into main directions, each plane has a contribution to the resulting main direction according to its normal vector. This contribution is independent on the number of points that were used to calculate the plane in the plane clustering process. The problem that arises here, is that planes which are obtained from less data points (or with lower *support*) have the same influence as planes with a high support. This may result in wrong main directions if, for example, a plane resulting from spurious measurements, is clustered together with a wall.

Additionally, the number of planes typically is very small compared to the number of data points. Thus, the influence of the main directions decreases the more data points are given. In practical experiments it turned out to be very useful to introduce weighting factors w_m for the planes which are dependent on the support of a given plane θ_m.

$$w_m = \sum_{n=1}^{N} E[\alpha_{nm} \mid \Theta] \quad (17)$$

In the EM we then use a modified distance function

$$d_2'(\theta_m, \varphi_k) = \sqrt{w_m} d_2(\theta_m, \varphi_k) \quad (18)$$

C. Sketch of the Algorithm

Our algorithm proceeds as follows:

1) Start with a fixed number of M_0 planes.
2) Initialize a main direction for each plane by taking the normal vector of that plane. This means the initial number of main directions K_0 equals M_0.
3) Apply EM until convergence
4) Drop one main direction as long as the *BIC* of the reduced model increases. This results in a new number of main directions $K_{i+1} \leq K_i$
5) Drop one plane as long as the *BIC* of the reduced model increases. This results in a new number of planes $M_{i+1} \leq M_i$
6) Select a new plane from the initialization queue and take its normal vector as a new main direction. Adding these to the model increases the complexity:

$$K_{i+1} := K_{i+1} + 1 \text{ and } M_{i+1} := M_{i+1} + 1$$

7) If no such plane can be found, stop. Otherwise go back to 3).

Note that it is not possible that K_i exceeds M_i in any time step i.

D. Post-Processing

So far, the goal of our algorithm was to find planes and main directions. In practice, however, we want to represent the environment as a set of *polygons*, because they indicate the faces of the objects in the environment. In general, a plane that is found in the data set contains more than one polygon. A typical situation is a wall that is "interrupted" by a doorway. In our implementation, we choose the following approach to extract polygons from planes:

- Determine all scan points that are close to the given plane (in our case: less than $0.1m$).
- Project all these points onto the plane.
- Perform a region growing on the points, where for each point all neighbors in a certain distance ϵ are added to the region. This is done efficiently using a 2D kd-tree (that is, a 2d-tree).
- Create an α-shape [17] from each region, where $\alpha = \epsilon$.

V. Experimental Results

The approach described above has been implemented and evaluated on real and simulated 3D data. Figure 3 contains the resulting planes obtained for the data set shown in Figure 1. The hierarchical EM is able to exploit the constraints introduced by the main directions of the planes and corrects the planes for the two parallel walls. In this example the angular error between the plane normals decreases from 17.2 degrees to 1.9 degrees.

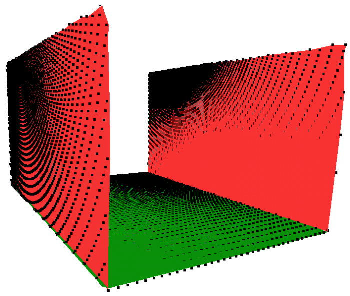

Fig. 3. Model obtained by our hierarchical EM. Compared to the model shown in Figure 1 the angle between the normals corresponding to the two parallel walls is 1.9 degrees.

TABLE I

Angles in degrees of the normals for the four planes in the newspaper rack in x, y and z direction.

Plane	plain EM			hierarchical EM		
	x	y	z	x	y	z
1	92.05	138.33	48.40	91.91	146.02	56.09
2	92.95	146.09	56.26	92.65	146.39	56.52
3	92.27	140.71	50.80	92.01	146.06	56.13
4	90.88	146.56	56.58	91.18	147.32	57.35

A. Experimental Evaluation on Real Data

The real data experiment was carried out with our mobile robot Zora shown in the left image of Figure 4. Zora is a B21R platform equipped with a 4DOF AMTEC manipulator which carries a SICK PLS range scanner. This setup allows our robot to flexibly scan complex scenes. The second image of Figure 4 shows a picture of a scene scanned with our robot. The third image of this picture depicts a typical data set obtained for this scene. The scan shown there consists of 21.479 points. To enhance visibility the data was smoothed and neighboring scan points were connected by triangles. The rightmost image of Figure 4 shows the result obtained with our hierarchical EM. The colors/grey-levels of the individual planes correspond to that of their main directions, which are also displayed. The final model consists of 7 planes and 3 main directions. The planes for the floor and the ceiling are only slightly corrected by the hierarchical EM. Whereas the plain EM

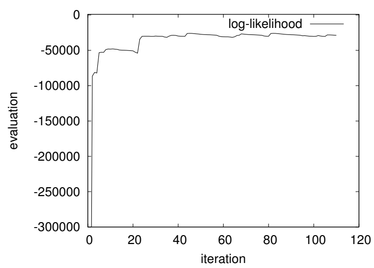

Fig. 5. Evolution of the log-likelihood during the estimation process.

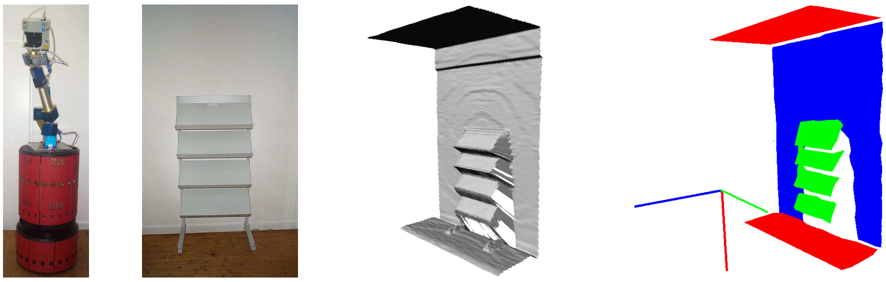

Fig. 4. Mobile robot Zora (left image), scene scanned by the robot (second image), resulting data (third image), and final model (right image).

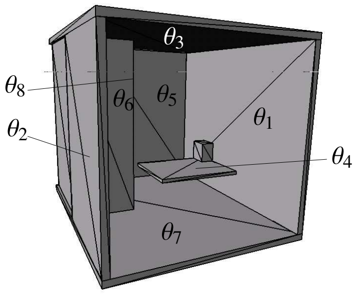

Fig. 6. Simulated 3d scene used to evaluate the quality of the resulting maps. On top of the horizontal plane is a small cube that introduces errors in the plane extraction.

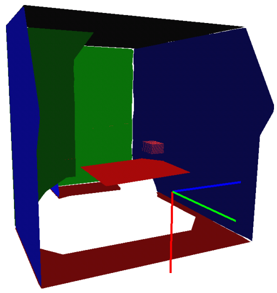

Fig. 7. Typical result obtained for the situation, in which the box was placed at the right rear corner of the small horizontal plane (see Figure 6).

TABLE II

Results for the simulated data set. The values are angular deviations from the ground truth in degrees.

Plane	plain EM			hierarchical EM		
	1	2	3	1	2	3
θ_1	0.220	0.234	0.227	0.155	0.208	0.224
θ_2	0.382	0.415	0.408	0.381	0.415	0.423
θ_3	0.479	0.512	0.492	0.289	0.067	0.190
θ_4	5.879	4.297	2.128	3.023	2.033	1.101
θ_5	0.105	0.070	0.083	1.524	0.318	0.121
θ_6	2.070	0.453	0.321	1.720	0.365	0.181
θ_7	0.672	0.940	0.889	0.244	0.608	0.624
θ_8	1.606	1.915	1.780	0.650	0.883	0.895

approach yields a deviation of 2 degrees for the ceiling and the floor, our algorithm generated planes with an angular distance of 1.7 degrees between the two planes. The most interesting part are the green/light grey planes for the newspaper rack. Note that their main direction is neither orthogonal to the main direction for the ceiling and the floor nor to that of the wall. Table I lists the individual angles of the normals for the four planes found for the rack in x, y, and z direction obtained with the plain EM approach and with our hierarchical EM. As can be seen from the numbers, our approach reduces the maximum deviation between the individual angles of the planes from 8 degrees to 2 degrees. Figure 5 plots the evolution of the log-likelihood during the hierarchical EM. Note that the log-likelihood does not always increase because of the introduction and removal of model components.

B. Quantitative Evaluation

Additionally, we performed several simulation experiments to evaluate the quality of the resulting models compared to the ground truth. Figure 6 shows a simulated scene used for these experiments. This scene represents a room with five corner walls, the floor, and the ceiling. Additionally, it contains a horizontal plane with a box on top. The walls are parallel to the coordinate axes, so that their normal vectors are the standard basis vectors $(1,0,0)^T$, $(0,1,0)^T$, and $(0,0,1)^T$. In total, there were eight visible planes. These planes are enumerated from θ_1 to θ_8. To evaluate the capabilities of our algorithm we performed three different experiments in which we varied the position of the small box on the small horizontal plane. In the first case the box was placed in the right rear corner of the horizontal plane. In the second experiment the box was placed halfway between the center of the horizontal plane and its right rear corner. In the third situation the box was located in the center of the horizontal plane. A typical model obtained with our hierarchical EM applied to the third situation is depicted in Figure 7.

The performance of the plain EM algorithm and our algorithm on these simulated data is given in Table II. Each column contains for all three experiments the deviation in degrees for each of the eight normal vectors from its respective ground truth. Especially the table plane θ_4 is corrected by the hierarchical EM. However, other planes like θ_8 are corrected using the knowledge of the main directions. Note that the error in some planes, e.g., the walls θ_1 and θ_2 increases slightly. This is because the plane θ_8, which has the same main direction, has a higher deviation and therefore slightly increases the error of the planes with

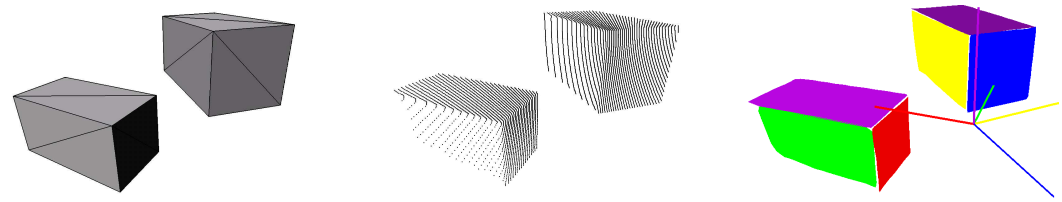

Fig. 8. Simulated scene with more than 3 main directions (left), simulated 3D-scan (center), and resulting planes and main directions (right)

the same main direction. In the final model obtained with our algorithm all three planes are almost parallel which indicates that the error is introduced by the constraints imposed by the corresponding main direction. The same holds for the planes θ_5 and θ_6.

C. Models with More Than Three Main Directions

In indoor scenes with mostly orthogonal or perpendicular planar structures such as offices, we only encounter three main directions. To illustrate that our algorithm can deal with more than just three main directions we performed an experiment with the simulated box world shown in the left image of Figure 8. The 3d data used as input to our algorithm is depicted in the middle image of Figure 8. Applied to this data set our algorithm found six planes and five main directions (see rightmost image of Figure 8). In this particular example, the hierarchical EM only slightly corrects the two top planes of the two boxes. All other planes were identical to the planes obtained by the non-hierarchical EM, since there was exactly one plane for each main direction.

VI. Conclusion

In this paper we presented a hierarchical approach to cluster 3d data points acquired with laser range scanners into planes. In contrast to previous approaches our algorithm uses a hierarchical model in which planes are also clustered into main directions. To find the model that maximizes the likelihood of the data we apply the EM algorithm. During the clustering process our approach simultaneously estimates the number of planes and the number of main directions.

The approach has been implemented and validated on real data and in simulation runs. The results demonstrate that the additional constraints imposed by the main directions in the hierarchical model allow to more reliably determine the planar approximations. The advantages of this are two-fold. First, the orientations of the planes are more accurate and second, we expect that points corresponding to objects close to planar structures can more reliable be separated in later segmentation steps.

Acknowledgment

This work has partly been supported by the German Research Foundation under contract number SFB/TR8 and by the EC under contract number FP6-004250-CoSy.

References

[1] A. Dempster, N. Laird, and D. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *Journal of the Royal Statistical Society*, vol. Series B, vol. 39, no. 1, pp. 1–38, 1977.

[2] Y. Liu, R. Emery, D. Chakrabarti, W. Burgard, and S. Thrun, "Using EM to learn 3D models with mobile robots," in *Proceedings of the International Conference on Machine Learning (ICML)*, 2001.

[3] R. Bajcsy, G. Kamberova, and L. Nocera, "3D reconstruction of environments for virtual reconstruction," in *Proc. of the 4th IEEE Workshop on Applications of Computer Vision*, 2000.

[4] S. Hakim, P. Boulanger, and F. Blais, "A mobile system for indoors 3-d mapping and positioning," in *Proc. of the 4th Conference on Optical 3-D Measurement Techniques*, 1997.

[5] M. Rous, A. Matsikis, F. Broicher, and K.-F. Kraiss, "Erzeugung eines planaren 3D-Modells aus Kameraaufnahmen zur Anwendung in der mobilen Robotik," in *Proc. Fachgespräche Autonome Mobile Systeme*, 2000. In German.

[6] I. Stamos and M. Leordeanu, "Automated feature-based range registration of urban scenes of large scale," in *Proc. of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2003.

[7] C. Früh and A. Zakhor, "An automated method for large-scale, ground-based city model acquisition," *International Journal of Computer Vision*, vol. 60, pp. 5–24, October 2004.

[8] S. Thrun, W. Burgard, and D. Fox, "A real-time algorithm for mobile robot mapping with applications to multi-robot and 3D mapping," in *Proc. of the IEEE International Conference on Robotics & Automation (ICRA)*, 2000.

[9] S. Thrun, D. Hähnel, D. Ferguson, M. Montemerlo, R. Triebel, W. Burgard, C. Baker, Z. Omohundro, S. Thayer, and W. Whittaker, "A system for volumetric robotic mapping of abandoned mines," in *Proc. of the IEEE International Conference on Robotics & Automation (ICRA)*, 2003.

[10] A. Nüchter, H. Surmann, and J. Hertzberg, "Planning robot motion for 3d digitalization of indoor environments," in *Proc. of the 11th International Conference on Advanced Robotics (ICAR)*, 2003.

[11] 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, pp. 15–27, 2003.

[12] C. Martin and S. Thrun, "Online acquisition of compact volumetric maps with mobile robots," in *IEEE International Conference on Robotics and Automation (ICRA)*, (Washington, DC), ICRA, 2002.

[13] J. Coughlan and A. Yuille, "Manhattan world: Orientation and outlier detection by bayesian inference," *Neural Computation*, vol. 15, no. 5, 2003.

[14] G. Schindler and F. Dellaert, "Atlanta world: An expectation maximization framework for simultaneous low-level edge grouping and camera calibration in complex man-made environments," in *CVPR*, 2004.

[15] D. Anguelov, R. Biswas, D. Koller, B. Limketkai, S. Sanner, and S. Thrun, "Learning hierarchical object maps of non-stationary environments with mobile robots," in *Proceedings of the 17th Annual Conference on Uncertainty in AI (UAI)*, 2002.

[16] G. Schwarz, "Estimating the dimension of a model," *Annals of Statistics*, vol. 6, pp. 461–464, 1978.

[17] F. Bernardini and C. Bajaj, "Sampling and reconstructing manifolds using alpha-shapes," tech. rep., Departement of Computer Sciences, Purdue University, 1997.