International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences, Vol. XXXVIII, Part 5
Commission V Symposium, Newcastle upon Tyne, UK. 2010

AUTOMATIC 3D MODELLING OF INDOOR MANHATTAN-WORLD SCENES FROM LASER DATA

Angela Budroni, Jan Böhm

Institute for Photogrammetry, Universität Stuttgart, Geschwister-Scholl-Straße 24D, 70147 Stuttgart, Germany – (angela.budroni, jan.boehm)@ifp.uni-stuttgart.de

Commission V, WG V/3

KEY WORDS: 3D Model, Indoor, Reconstruction, Automatic, Point Cloud, Sweep

ABSTRACT:

We developed an automatic technique for the reconstruction of Manhattan-world interior scenes from point clouds. Our method mainly focuses on the reduction of human intervention in the modelling process, thus it aims to a full automation of the modelling phase with a consequent optimization of the efficiency of the laser-scanning project pipeline. We refer to a reconstruction complexity typical for indoor scenes. The work flow starts with a volume sweep reconstruction of the interior from the three-dimensional point cloud. As result of a discrete translational plane sweep, the input data is segmented into separate point sets including the floor, ceiling and wall points. Consequently, each point is assigned to a surface of the volume. The ground plan contours are extracted with a cell decomposition approach after partitioning the floor surface into rectangular cells of variable size. Only cells considered suitable are added to the ground shape and unified to define the ground plan. Along the ground plan contour, the walls are raised from the floor up to the ceiling level. Finally, the interior model is enhanced by the addition of built-in feature like doors.

1. INTRODUCTION

A growing attention concentrates upon the development of new automatic techniques for interactive building representations. A higher degree of automation in the reconstruction process may contribute to a fast and dynamic architectural documentation and facilitate the spread of oncoming CAD representations. A progression from 3D to 4D and BIM (Building Information Modelling) is expected, in which 3D geometric representations will only be a part of the BIM concept (U.S. GSA, 2010). Modern BIM design tools are flexible because they can be interpreted easily by computers by eliminating the redundancy typical of old CAD drawings. As described in (Eastman et al., 2008), the parametric aspect of BIM tools allows the objects to automatically update according to changing contexts. Thus, an automatic method to digitize existing architecture could be also integrated into a BIM machine.
The automatic reconstruction of building models has been mostly restricted to the determination of buildings' outer shapes. We developed an approach for indoor reconstruction from point clouds. The complexity of our 3D model is compatible with the highest level of details – LoD4 – required by the multi-scale representation CityGML (Kolbe et al., 2005). In particular we reconstruct a 3D model, which incorporates interior architectural details like rooms and interior doors. Also the accuracy of the final CAD model conforms to the CityGML requirements for the LoD4. The domain of our field of application is defined by the Manhattan-world assumption.

2. PREVIOUS WORK

The previous work, which has influenced our reconstruction approach, concerns the automatic and semi-automatic generation of three dimensional models from both point clouds and images. The convenience of such methods involves a wide range of situations depending on the type of scene to be reconstructed. For example, 3D models of ordinary city buildings may represent a truthful scenario for modern digital globes environments or they may support tools for urban planning and augmented reality applications. On the other hand, detailed models of historic architecture may contribute to cultural heritage preservation.
For the reasons above or just for record-keeping and fast visualization, current research continuously points on the development of new methods for 3D model generation. Noticeable work has been done for increasing the automation level of the reconstruction process, which is also one of the fundamental goals of our method. One possible step toward building modelling is feature extraction from terrestrial (and aerial) laser scanned data. According to (Pu and Vosselman, 2006), terrestrial laser data provides high density point clouds suitable for region segmentation. Constraints for segmentation are given by human knowledge about building features. The same approach is at the base of window extraction in (Pu and Vosselman, 2007). An overview of the different techniques for recognizing structures in point clouds is given in (Vosselman et al., 2004) with a special focus on the extraction of smooth and planar surfaces. (Boulaassal et al., 2007) evaluate a method to deliver planar façade segments by allowing automatic fitting of geometric primitives (segmentation) using RANSAC algorithm. The underlying principles of hybrid approaches that combine laser data and images for building reconstruction are reviewed in (Brenner, 2004).
Beside the point of automation, our method carries also the issue of reconstructing interior scenes. The difficulties inherent in the reconstruction of interiors are different from the problematics that come from the extraction of external shapes of buildings, thus they have to be treated differently. This is also the reason why model reconstruction and visualization of generic indoor scenarios is still a difficult task (Furukawa et al., 2009b). (Furukawa et al., 2009b) and (Furukawa et al., 2009a) address the special problem of reconstructing building interiors with a multi-view stereo approach for Manhattan-world scenes

(Coughlan and Yuille, 1999). In (Delage et al., 2006) the first example on recovering an indoor model from a single image is described. So far, the task of interior reconstruction has mostly been associated with navigation problems and autonomous systems (Biber et al., 2005). On the other hand, solutions to military applications are shown in (Johnston and Zakhor, 2008), whose approach exploits a technique similar to our method.
The automatic reconstruction of interiors is an interesting topic among the architecture community, too. The convenience of indoor CAD models, which are easy to compute, can simplify the architectural design process by supporting a dynamic interaction between existing spaces and new structures. A challenging study about how to generate 3D models from architectural drawings is carried out in (Yin et al., 2009). The inputs of the method described in (Yin et al., 2009) are paper ground plans (images) that need to be converted into CAD files in order to automatically generate the model.

3. PROBLEM DOMAIN DEFINITION

Since the subject of indoor reconstruction may cover a wide range of different situations (from cultural heritage to industrial environments), we restrict our domain of application in order to develop a robust method that delivers optimal results in response to one specific issue. The reason lies in the automation of our reconstruction algorithm. An automatic tool in fact is supposed to work independently of external control, thus it requires to be specified into an operational domain, which defines its particular tasks. Referring to (Alexander et al., 1977), we create a language for the description of the interior building types that are the targets of our reconstruction algorithm. That means we are making hypotheses on the architectural characteristics of the building interior to be modelled. Our problem domain is determined by the following parameters:

Manhattan-world grid structure. As observed in (Coughlan and Yuille, 1999), most indoor and outdoor city scenes are based on a Cartesian coordinate system, which we can refer to as a Manhattan grid. The grid defines the positions of the walls.

Horizontal floor and ceiling are both orthogonal to the walls.

Lack of furniture or objects with an irregular geometry. We typically refer to public building hallways with only essential interior decoration elements.

Multiple doors connecting side rooms. Doors are expected to keep to standard measures.

On the basis of these requirements, which will be formalized in the following, we interpret the point cloud using simple rules of descriptive geometry. Therefore, we create a robust algorithm for processing the interior data and generating its 3D model. The coordinate system we refer to is Cartesian.

4. 3D MODELLING OF INTERIORS

Our solution to the problem of Manhattan-world interior reconstruction consists of two main subsequent steps (Figure 1). Minor issues are treated in parallel to give a contribution to the final model in terms of output verification (histogram visualization) and detail enhancement (door extraction). A third additional step may be added at the end of the pipeline to allow the application to handle different portions of the data by properly connecting them in a sequential fashion.
The steps we refer to are the sweep reconstruction of the indoor volume and the ground plan cell decomposition. The first is accomplished by a translational (linear) sweep to compute a volume enclosed by orthogonal faces. As a result, each face in the volume is segmented, too. In the second step, the hypothetical ground plan, represented by one of the previously segmented faces, is divided into rectangular cells whose effective correspondence with the real ground-plan depends on the number of points counted in the cell. The algorithm we propose processes a point cloud (input) in binary format and returns its 3D model as a CAD file.

5. INDOOR VOLUME SWEEP RECONSTRUCTION

In computer graphics, a translational sweep is defined by sweeping a 2D area along a linear path normal to the plane of the area (Foley et al., 1996). The sweep computation creates a volume whose measure is given by the sweeping area times the length of the sweep. Depending on the type of primitive used to describe the sweeping area, the resulting volume takes on different shapes. The primitives that best satisfy our problem requirements are either a rectangle or a union of rectangles. In fact, according to the first two hypotheses made on the architectural characteristics of our data, the final indoor volume is expected to be a solid with faces, each being perpendicular or parallel to the others. However, our method does not require the extension of the sweeping area as prior knowledge as the contours (hence the shape) of the sweeping volume will be automatically determined in a following stage. Theoretically, this is done by intersecting perpendicular planes through the point cloud minimum and maximum coordinates along each of the three Cartesian axes.
Our algorithm consists in sweeping a generic plane with equation

$$ax + by + cz + d = 0 \tag{1}$$

along its normal vector

$$\mathbf{n}_s = (a, b, c). \tag{2}$$

The first issue is how to compute the coefficients that define the normal vector. To do so, we consider the surfaces that are part of the indoor volume to be reconstructed. The floor, the ceiling and some specific walls form the outer boundary of the indoor volume thus they are definitely part of the volume. Consequently, the trajectory of the sweep is given by the normal vectors to the plane of the floor (ceiling) and the walls.

Sweep along the Normal to the Plane of the Floor

Since the floor and ceiling (both horizontal) are parallel, sweeping a plane along the normal vector to the floor has the same meaning of sweeping a plane along the normal vector to the ceiling. In addition, we consider the laser scanner to be levelled so that the vertical axis of the coordinate system is aligned with the local vector of gravity. The coefficients that define the sweep trajectory are then determined by the components of the vertical direction of the coordinate system:

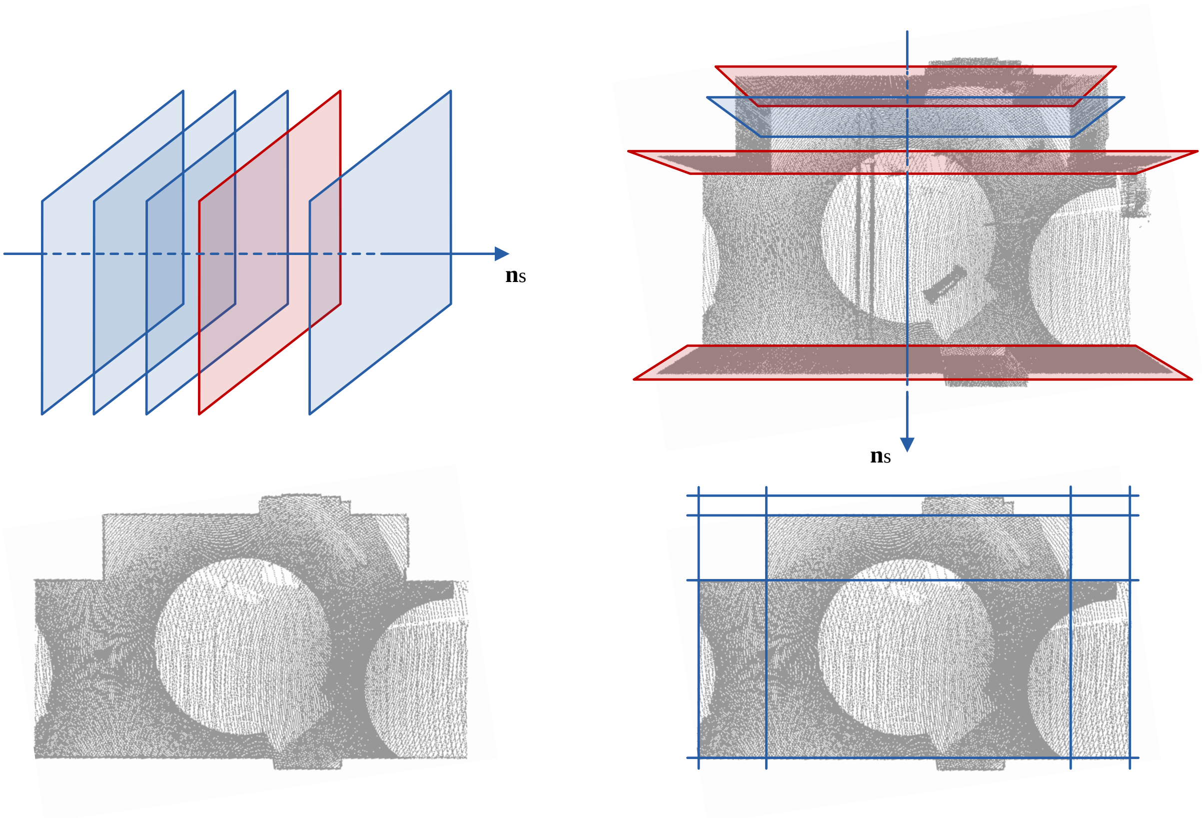

Figure 1: Two main steps for the solution to the problem of Manhattan-world interior reconstruction: Sweep reconstruction of the indoor volume (top) and ground-plan cell decomposition (bottom).

$$\mathbf{n}_v = (0, 0, 1). \quad (3)$$

Sweep along the Normal to the Plane of the Walls

Different walls of the same interior do not necessarily have the same orientation. According to the Manhattan-world assumption, the walls of our indoor volume can be positioned uniquely along two main orthogonal directions. Besides, the only constraint on the position of the walls with respect to the coordinate system is the parallelism to the vertical axis, which yields

$$c = 0$$

in equation 2. The remaining components of the sweep direction are expressed as trigonometric functions of an unknown angle α. Thus, the normal vector to the sweep plane is

$$\mathbf{n}_\alpha = (-cos\alpha, -sin\alpha, 0) \quad (4)$$

considering the orientation of one group of parallel walls and

$$\mathbf{n}_{\alpha+\pi/2} = (sin\alpha, -cos\alpha, 0) \quad (5)$$

for the perpendicular walls. In (Budroni and Böhm, 2009) we explain how the value of α is automatically computed.

Following from the above, three separate translational sweeps are needed for the complete reconstruction of the indoor volume. The further issue is the decision of a criterion to determine the length of each sweep as well as the occurrences of relevant surfaces along the sweep path. Along one sweep direction the length of the sweep corresponds to the difference between the maximum and minimum value of the point cloud coordinates. Potential wall – floor and ceiling – surfaces within the sweep domain are detected with an approach based on a hypothesis-and-test strategy. The discrete plane sweeping builds a hypothesis about the position of a surface at every step s_i. The physical existence of the current surface within the volume is tested by counting the number of 3D points available in the neighborhood of the surface. Only if the number of points exceeds a pre-determined threshold, the presence of a surface along that sweep at the i-th step is confirmed. Otherwise, no surface is detected and the following step s_{i+1} is tested.

5.1 Surface Segmentation

With the new information available after the sweep reconstruction, the point cloud can be segmented and each 3D point assigned to a surface of the volume. Every surface and its points are associated with one plane in space thus with one sweep direction – $\mathbf{n}_v$, $\mathbf{n}_\alpha$ or $\mathbf{n}_{\alpha+\pi/2}$ – and with a step position $s_i(\mathbf{n}_s)$ along that direction. For every direction, a histogram plots the frequencies of points that fall into each step interval. Points among the input data set are classified in three main groups: Those belonging to the floor, to the ceiling and to the vertical planes.

6. GROUND PLAN CELL DECOMPOSITION

Once the main surfaces that define the volume are determined by their positions along a sweep trajectory, the shape of the interior's ground plan is recovered by computing the union of rectangular ground cells. The ground plan represents the horizontal projection of the reconstructed volume thus its contours are given by the intersections of the vertical faces and the horizontal plane. To reduce the computational efforts, we consider the straight lines resulting from these intersections and resolve the problem into a two-dimensional domain.

The approach behind our ground plan extraction is cell decomposition, a general form of space-partitioning modelling. According to the process of cell decomposition, we can represent any solid (or surface) as the sum or union of a set of cells into which is divided (Mortenson, 1997). This is well-suited for our problem since the full ground plan is not representable yet, but its cells are. A set of rectangular cells of variable size is mapped on the floor plane by intersecting two orthogonal sheaves of straight lines as shown in Figure 2. We refer to these lines as cut-lines. Each cut-line $l_i(\mathbf{n}_s)$ localizes the intersection of the vertical plane of one wall and the floor plane. The cut-lines are aligned either with the vector $\mathbf{n}_s = \mathbf{n}_\alpha$ in equation 4 or with $\mathbf{n}_s = \mathbf{n}_{\alpha+\pi/2}$ in equation 5. From the result of the cut-line intersection, we extract a distribution of points. The vertices of each cell are recovered from the point distribution by selecting neighbouring points according to a repetitive adjacency configuration. The data structure that describes the cells stores a sequence of points listed in anticlockwise order.

The cells – rectangles – are now available for the extraction of the ground plan of the interior. A union of floor cells is implemented after testing which cells are compatible with the floor. In fact, it is not guaranteed that every cell produced by the cut-line intersection lies within the ground plan contour. A check on the number of points inside the rectangles provides a separation of valid and irrelevant cells before assigning only the valid cells to the ground plan.

6.1 Valid Cell Specification

The criterion for a cell to be considered valid and accepted as a floor cell consists in a lower bound for the points inside the cell rectangle. Unfortunately, due to measurement limitations and to the discrete approximation of the sweep steps, cells that are clearly outside the ground plane contour may contain a considerable number of points anyway. That happens especially for cells adjacent to a valid cell. On the other hand, cells with a low sampling density may not be able to meet the general specification about the minimum number of points per cell rectangle. To solve these inconsistencies, which could drive us to an erroneous evaluation of the validity of the cells, we introduce in the requirements a dependency on the area of each rectangle. The lower bound for the points in the cell is compared to the ratio of the rectangle's area to a given global sampling density. In addition, we compute schrinked contours for every cell and count the number of points inside the new perimeter. By shrinking the cells, the influence of point overflow from neighbouring cells is avoided.

7. EXTRACTING DOORS

The extraction of details like doors adds completeness to the final 3D model and provides a more consistent solution to the indoor reconstruction problem. For our application it is necessary that the doors are open or receding since they are detected as linear gaps on the planar surfaces of the walls similar to (Becker and Haala, 2008). The doors are localized by processing a two-dimensional projection of the 3D data on the wall of interest (see Figure 3).

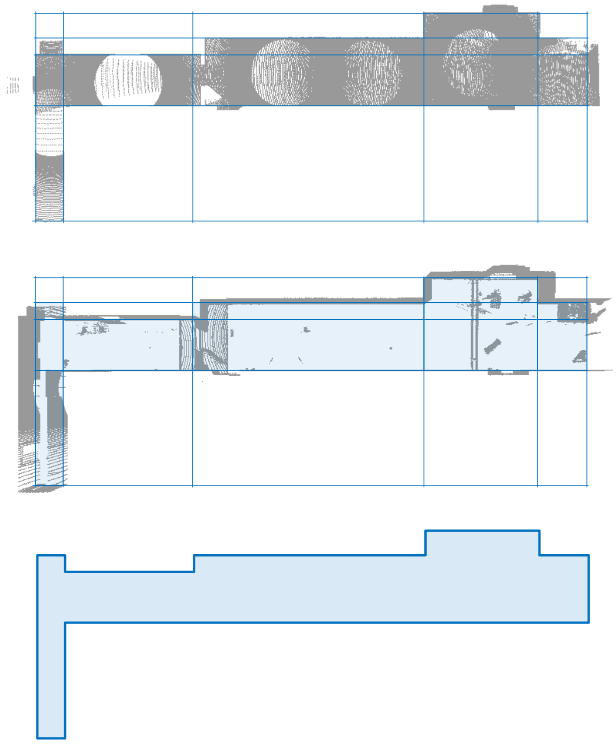

Figure 2: Cell decomposition of a ground plan: Intersection of orthogonal cut-lines (top), union of accepted cells (middle) and extraction of the ground plan contour (bottom).

The cut-lines locate the wall positions along the sweep direction. The wall points are rotated by α to be aligned to the Cartesian axes, thus the cut-line directions $\mathbf{n}_\alpha$ and $\mathbf{n}_{\alpha+\pi/2}$ correspond to the X and Y axes. We refer to the new axis-aligned point coordinates as (x_α, y_α, z). Because of the Manhattan-world assumption, two separate cases are possible for the statement of the door extraction problem:

- Walls perpendicular to $\mathbf{n}_\alpha$. They have a fixed x_α coordinate for every cut-line $l_i(\mathbf{n}_\alpha)$. The linear gap is detected along the Y direction (Figure 4, left configuration).

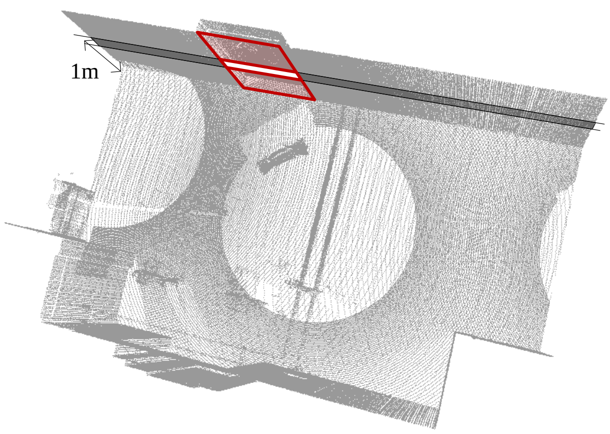

Figure 3: Two-dimensional solution for the localization of the doors: Point stripe extraction among the wall points.

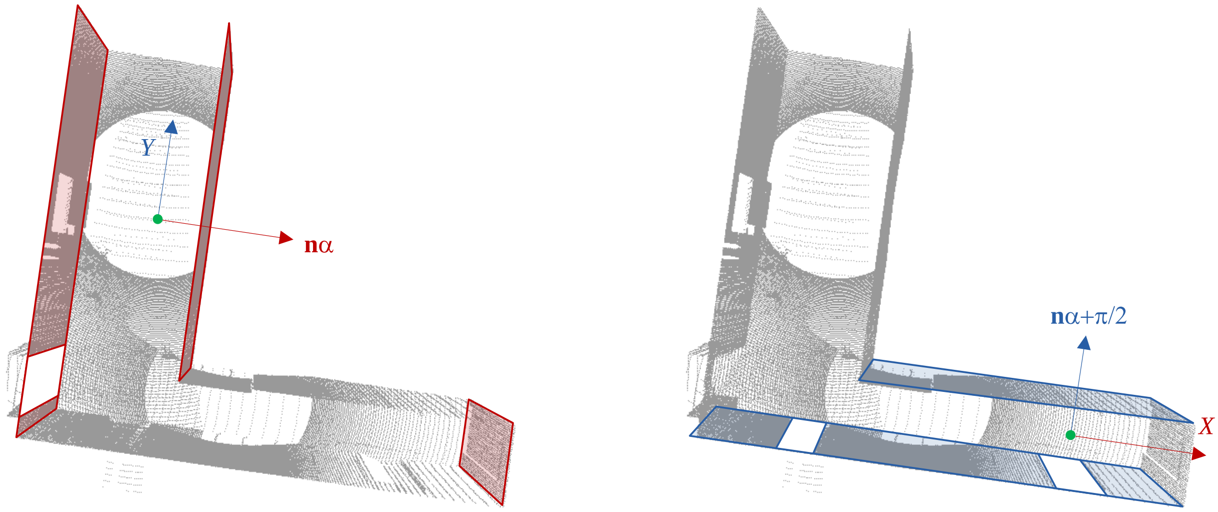

Figure 4: Detection of the doors along the two main cut-line directions. The linear gap is detected along the direction orthogonal to the sweep

- Walls perpendicular to $\mathbf{n}_{\alpha+\pi/2}$. They have a fixed y_α coordinate for every cut-line $l_i(\mathbf{n}_{\alpha+\pi/2})$. The linear gap is detected along the X direction (Figure 4, right configuration).

Let us consider only the first configuration. The second has the same characteristics if a rotation of $\pi/2$ is applied (see Figure 5). Among the points associated with one wall plane, a horizontal stripe is selected. As shown in Figure 5, the stripe is centred at a height of one meter and contains points with coordinates (k, y_α, $1\pm\varepsilon$), where k is a constant value depending on the wall position along the sweep direction $\mathbf{n}_\alpha$. By sorting the coordinate y_α and subsequently detecting a gap of missing points between two neighbouring values, a door with the same width of the gap is added to the model.

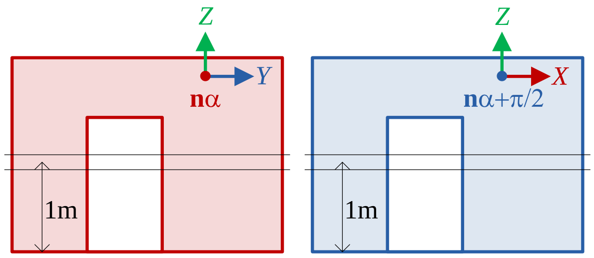

Figure 5: Door detection configuration for orthogonal directions.

We conclude from the test cases that the method is robust and the doors are located with an accuracy that is affected only from measurements errors (Figure 6). The only limitation of the door extraction method is the risk of false detections. Every convex or concave object deforming the wall profile may cause a gap in the data stripe on the plane of the wall. The width of the gap is set to 80 cm referring to the standard measure of a door (80 x 200 cm). For the same reason, the point stripe is selected at one meter.

8. CONCLUSIONS

We developed a robust method for reconstructing interior architecture (see Figure 7). Given a point cloud, the overall shape of the processed data is reconstructed as a three-dimensional model provided in CAD format. The data sets under test consist of point clouds of variable point density. The processing time grows with the number of points to be matched with the sweeping plane during the sweep reconstruction. The method was tested on a data set with over 2 million points with a resulting performance comparable to those in response to smaller data samples. We conclude that the algorithm robustness is not affected by huge dimension of the input data. Anyway, a further issue may be the connection of adjacent volumes after the single models are separately computed.
A limitation to the full optimization of the method is due to the Manhattan-world domain. Because of the Manhattan-world assumption, the primitives, which build the volume, are rectangular. If primitives with different shapes are used, structures with different characteristics may be reconstructed (for example, triangular primitives may reproduce a loft with an oblique roof).

Figure 6: Particular of the door extraction results.

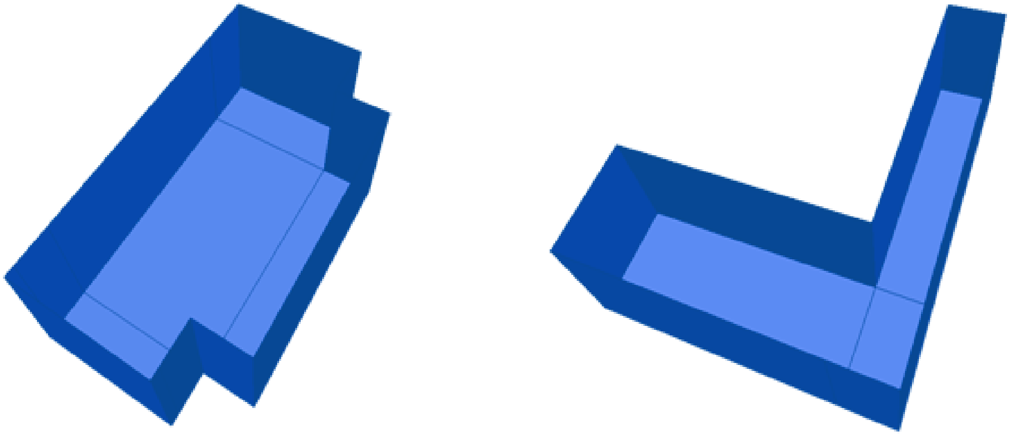

Figure 7: 3D Models. The models are displayed without the ceiling for better visualization. They are computed from point clouds with 800 thousand (left) and 650 thousand points (right).

REFERENCES

Alexander, C., Ishikawa, S. and Silverstein, M., 1977. A Pattern Language. Oxford University Press.

Becker, S. and Haala, N., 2008. Integrated LIDAR and image processing of the modeling of building facades. In: PFG.

Biber, P., Fleck, S., Wand, M., Stanker, D. and Strasser, W., 2008. First Experience with a mobile platform for flexible 3D model acquisition in indoor and outdoor environments. In: SPIE Electronic Imaging Conference.

Boulaassal, H., Landes, T., Grussenmeyer, P. and Tarsha-Kurdi, F., 2007. Automatic segmentation of building facades using terrestrial laser data. In: ISPRS Workshop on Laser Scanning.

Brenner, C., 2004. Building reconstruction from images and laser scanning. International Journal of Applied Earth Observation and Geoinformation.

Budroni, A. and Böhm, J., 2009. Toward automatic reconstruction of interiors from laser data. In: IAPRS Workshop 3D-Arch on 3D Virtual Reconstruction and Visualization of Complex Architectures.

Coughlan, J. and Yuille, A., 1999. Manhattan world: compass direction from a single image by Bayesian Inference. In: IEEE International Conference on Computer Vision.

Delage, E., Lee, H. and Y. Ng, A., 2006. A dynamic Bayesian network model for autonomous 3D reconstruction from a single indoor image. In: IEEE Computer Society Conference CVPR, Vol. 2, pp. 2418-2428.

Eastman, C., Teicholz, P., Sacks, R. and Liston, K., 2008. BIM Handbook: A Guide to Building Information Modeling for Owners, Managers, Designers, Engineers and Contractors. John Wiley & Sons, Inc., New Jersey.

Foley, J. D., van Dam, A., Feiner, S. and Hughes, J., 1996. Computer Graphics: Principles and Practice. Addison-Wesley.

Furukawa, Y., Curless, B., Steitz, S. and Szeliski, R., 2009a. Manhattan-world stereo. In: IEEE Computer Society Conference CVPR.

Furukawa, Y., Curless, B., Steitz, S. and Szeliski, R., 2009b. Reconstructing building interior from images. In: IEEE International Conference on Computer Vision.

Johnston, M. and Zakhor, A., 2008. Estimating building floor-plans from exterior using laser scanners. In: SPIE Electronic Imaging Conference.

Kolbe, T., Gröger, G. and Plümer, L., 2005. Interoperable access to 3D city models. In: Int. Symposium on Geo-information for Disaster Management.

Mortenson, M., 1997. Geometric Modeling. John Wiley & Sons, Inc.

Pu, S. and Vosselman, G., 2006. Automatic extraction of building features from terrestrial laser scanning. In: ISPRS Image Engineering and Vision Metrology.

Pu, S. and Vosselman, G., 2007. Extracting windows from terrestrial laser scanning. In: ISPRS Workshop on Laser Scanning.

U. S. General Services Administration website, 2010. http://www.gsa.gov/Portal/gsa/ep/contentView.do?contentType =GSA_OVERVIEW&contentId=20917 (accessed 29 Apr. 2010)

Vosselman, G., Gorte, B., Sithole, G. and Rabbani, T., 2004. Recognising structure in laser scanner point clouds. In: International conference NATSCAN.

Yin, X., Wonka, P. and Razdan, A., 2009. Generating 3D building models from architectural drawings: a survey. In: IEEE Computer Graphics and Applications 29, pp.20-31.