PLANAR 3D MODELING OF BUILDING INTERIORS FROM POINT CLOUD DATA

Victor Sanchez and Avideh Zakhor {victors,avz}@eecs.berkeley.edu
Video and Image Processing Lab, EECS Department, University of California, Berkeley

ABSTRACT

We present an automatic system for planar 3D modeling of building interiors from point cloud data generated by range scanners. This is motivated by the observation that most building interiors may be modeled as a collection of planes representing ceilings, floors, walls and staircases. Our proposed system, which employs model-fitting and RANSAC, is capable of detecting large-scale architectural structures, such as ceiling and floors, as well as small-scale architectural structures, such as staircases. We experimentally validate our system on a number of challenging point clouds of real architectural scenes.

***Index Terms*—** 3D indoor modeling, range scanners, point cloud, RANSAC, PCA, model fitting

1. INTRODUCTION

Three-dimensional (3D) modeling of architectural scenes from point cloud data generated by range scanners is an increasingly important research area with applications such as virtual and augmented reality, navigation, and virtual heritage conservation. While user-assisted methods have been proven effective, automatic methods for 3D interior modeling have received much interest in recent years. In this context, methods capable of modeling building interiors in terms of planar primitives are gaining popularity.

In contrast to building exteriors, the modeling of interiors poses two important challenges. First, a complexity challenge, as building interiors may contain several interconnected rooms and hallways. Second, a scalability challenge, given the prevalence of small-scale structures, such as walls and staircases, relative to the scale of the scene.

A number of methods for automatic planar 3D modeling of building interiors from point cloud data have been recently reported in the literature [1-2]. Although the method in [1] deals efficiently with architectural scenes comprising planar faces, it only produces a 3D model describing the exterior geometry of the building. In [2], the authors propose a method to model an architectural scene as a collection of cuboids, i.e., planar facets. The proposed method is robust to noisy and sparse data; however, it is restricted to modeling cuboids if at least five facets are detected as planes. Consequently, architectural structures with a limited number of facets, such as staircases, cannot be modeled.

This paper presents an automatic system for planar 3D modeling of building interiors from point cloud data by using a set of planar primitives. It is based on the observation that building interiors often feature highly structured floor plans comprising a strong planar arrangement. Therefore, building interiors may be modeled as a set of planar primitives representing four basic architectural structures: ceilings, floors, walls and staircases.

Our approach can be summarized as follows. First, we perform point-wise principal component analysis (PCA) on a point cloud representing a building interior. This is followed by a classification and segmentation process to identify and segment regions representing basic architectural structures. We then employ model-fitting on the segmented regions to detect ceilings, floors, walls, and staircases. We then generate a 3D model by assembling the detected architectural structures.

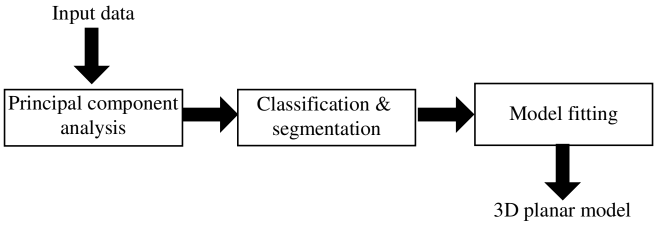

Fig. 1. Proposed planar 3D reconstruction system.

The rest of the paper is organized as follows. Section 2 describes in more detail the pipeline of our system. Section 3 presents and discusses the experimental results. Section 4 provides the concluding remarks.

2. PIPELINE OF THE PROPOSED SYSTEM

Our planar 3D modeling system comprises three steps, as shown in Fig. 1. The next sections describe, in more detail, these steps.

2.1. Principal Component Analysis

The input to a our system is a point cloud represented by the set $\mathbf{P} = \{p_i,\ i \in \{1,\dots,N\}\}$ in $\mathbb{R}^3$. We assume the point cloud data is generated by a mobile 3D acquisition system comprising a number of 2D range scanners, which are used for both localization and geometry capture [3-4]. Let c_i denote the 3D position of the range scanner for p_i. Note that c_i may be different for every input sample, as the acquisition system traverses the building interior.

We first perform a first-order approximation of the input data by estimating the normal vectors of the set $\mathbf{P}$. For each input point p_i of $\mathbf{P}$, we perform PCA on the ball neighborhood of radius σ centered on p_i, as shown in Fig. 2. PCA results in a set of eigenvalues, λ_{max}, λ_{mid}, λ_{min}, which encode the size and shape of the best fit ellipsoid over the ball neighborhood, and a set of eigenvectors, e_{max}, e_{mid}, e_{min}, which provide information about the principal directions of the fit ellipsoid. The eigenvector associated with the smallest eigenvalue constitutes an estimate of the normal vector, n_i, to the surface of the ball neighborhood centered on p_i. We assign n_i a consistent orientation by making sure that c_i is on its positive side, as shown in Fig. 2.

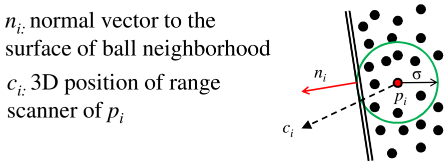
$n_{i:}$ normal vector to the surface of ball neighborhood
$c_{i:}$ 3D position of range scanner of p_i
n_i
σ
p_i
c_i

Fig. 2. Ball neighborhood of radius σ centered on point p_i.

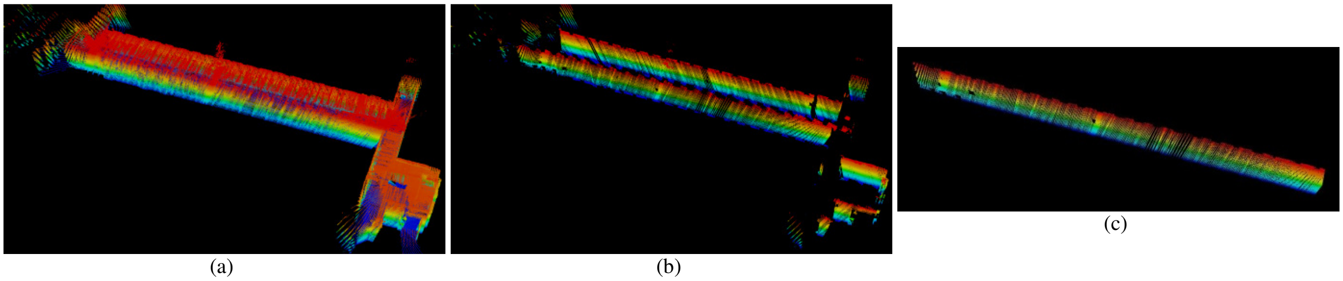
(a)
(b)
(c)

Fig. 3. Classification and segmentation of *wallsX* points of point cloud data representing a T-shaped hallway. (a) Original point cloud. (b) Subset *wallsX* after PCA and classification. (c) Segmented region representing a wall.

2.2. Classification and segmentation

Based on the results of PCA, we divide the input points into five disjoints subsets: *ceiling* points, *floor* points, *wallsX* points, *wallsY* points, and *remaining* points. Subset *wallsX* contains all the points whose normal vector n_i has a maximum component along the *x*-axis, whereas subset *wallsY* contains all the points whose normal vector n_i has a maximum component along the *y*-axis. Let $\theta(a,b)$ denote the minimum 3D angular distance between vectors *a* and *b*. Let us assume that points representing ceilings and floors lie on planes that are parallel to the *x*-*y* plane, and points representing walls lie on planes that are perpendicular to the *x*-*y* plane. Based on these assumptions, each input point p_i is classified into a single subset as follows:

$$p_i \in \begin{cases} \textit{ceiling} & \text{if } \theta\big(n_i,(0,0,-1)\big) < 15 \text{ deg.} \\ \textit{floor} & \text{if } \theta\big(n_i,(0,0,1)\big) < 15 \text{ deg.} \\ \textit{wallsX} & \text{if } \theta\big(n_i,(0,0,\pm1)\big) < 45 \text{ deg.} \\ \textit{wallsY} & \text{if } \theta\big(n_i,(0,0,\pm1)\big) < 45 \text{ deg.} \\ \textit{remaining} & \text{otherwise} \end{cases} \tag{1}$$

After dividing the input points into disjoints subsets, we grow a region through any pair of points in the same subset that are within an Euclidian distance d_{eucl}. We discard regions that contain fewer points than a threshold N_p, as they usually represent noisy regions. Note that d_{eucl} and N_p are user-specified parameters. Finally, we segment each found region from the disjoint subsets. Figure 3 shows the classification and segmentation of *wallsX* points of point cloud data representing a T-shaped hallway.

2.3. Model fitting

After classification and segmentation, we fit a model to each segmented region. For regions segmented from the subsets *ceiling*, *floor*, *wallsX* and *wallsY,* we fit a planar model, which may result in a single planar primitive representing an architectural structure. For regions segmented from the subset *remaining,* we search for and possibly fit a staircase model, which may result in a set of planar primitives representing treads and risers of steps. The latter is based on the assumption that any staircase structure, if present in the data, should be contained in the subset *remaining*.

2.3.1. Planar model

We employ random sample consensus (RANSAC) to find the best fitted planar primitive in the least-square sense [5]. RANSAC randomly draws a set of three different points from the segmented region and computes a candidate planar primitive. A planar primitive consists of an oriented 3D plane and a set of inliers, which represent its spatial extent. A point p_i is an inlier of a planar primitive *T* if and only if the squared distance from p_i to *T* is less than a threshold *q*. The resulting candidate planar primitive is tested against all points of the segmented region to determine the number of inliers. After a given number of trials, the candidate planar primitive with the most inliers is selected and its spatial extent is computed.

It is expected that the spatial extent of most planar primitives describing walls may be approximated by four vertices, since walls are generally rectangular in shape. However, some walls are polygonal in shape and may require more than four vertices to approximate its spatial extent. Similarly, the spatial extent of most planar primitives describing ceilings and floors may require more than four vertices to be properly approximated, since these architectural structures are generally polygonal in shape.

In order to determine the spatial extent of walls, ceilings and floors, we first project all inliers on the corresponding oriented 3D plane. We then employ alpha-shapes to compute the vertices of the concave hull of the set of projected inliers [6]. Finally, we employ the split-and-merge algorithm to eliminate any redundant vertices, and determine the final vertices that approximate the spatial extent [7]. Fig. 4 shows the process to determine the spatial extent of a sample polygonal wall.

2.3.2. Staircase model

Our staircase model comprises six parameters as illustrated in Fig. 5. These six parameters are used to represent each step by a set of two planar primitives, corresponding to the tread and the riser. The spatial extent of each planar primitive is approximated by four vertices, since treads and risers are generally rectangular in shape.

The staircase model-fitting approach consists of two steps. The first step is a coarse search for valid staircase ramps, to limit the complexity of the second step, which further tests for the existence of steps. In the first step, we employ RANSAC to determine if a segmented region is a valid staircase ramp. A segmented region is a valid staircase ramp if it is possible to fit a planar primitive with

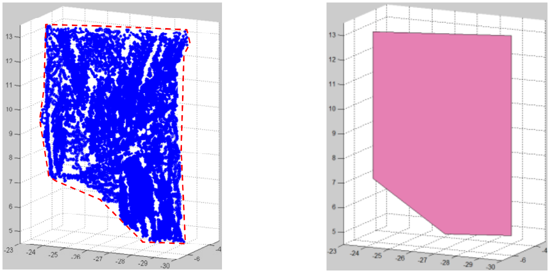

(a)
(b)

Fig. 4. Sample polygonal wall. (a) Projected inliers on the oriented 3D plane. Dotted red line denotes the concave hull calculated by alpha-shapes. (b) Final spatial extent after split-and-merge.

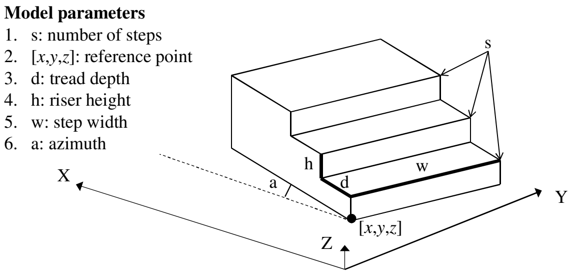

Fig. 5. The six parameters of the staircase model.

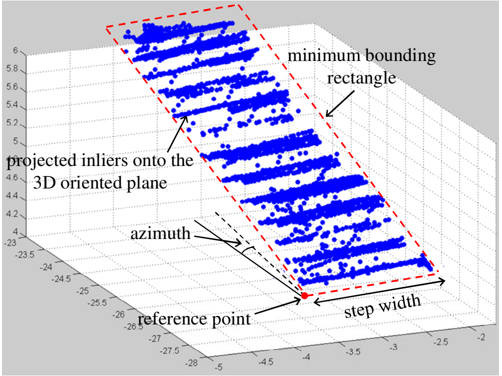

Fig. 6. Reference point, azimuth and step width of the staircase model.

an inclination between 25 deg. and 50 deg. with respect to the x-y plane, and with a spatial extent with a minimum width w_{min} = 0.5 m. and a minimum length l_{min} = 0.5 m. We determine the spatial extent of the planar primitive by projecting all inliers to the corresponding 3D oriented plane and finding the smallest bounding rectangle containing all projected inliers. Note that the values of w_{min}, l_{min}, and the inclination range of the fitted planar primitive represent the width, length and inclination of typical staircase ramps.

In the second step, we fit our staircase model to the segmented region only if a valid staircase ramp is detected in the first step. We use the smallest bounding rectangle containing all projected inliers, to determine the reference point, azimuth and step width, as shown in Fig. 6. The width of the bounding rectangle corresponds to the step width, while one of its vertices serves as the reference point. We project the bounding rectangle onto the x-y plane to determine the azimuth by calculating its orientation with respect to the x-axis.

In order to determine the tread depth, riser height and number of steps, we assume that the mobile acquisition system has a 2D pitch scanner capable of scanning the x-z or y-z plane, depending on whether the scanning path is on the x or y direction [3]. Using the pitch scanner data corresponding to the detected staircase ramp, we determine the beginning and end points of each riser and tread by fitting a model comprising two perpendicular lines, as shown in Fig. 7. The average Euclidean distance between the beginning and end points of all detected risers corresponds to the riser height, while the average Euclidian distance between the beginning and end points of all detected treads corresponds to the tread depth. The total number of steps is set to the total number of detected pairs of treads and risers. Fig. 8 shows a section of a sample *remaining*-subset with two valid staircase ramps and the corresponding fitted models.

3. EXPERIMENTAL RESULTS

Our system employs C++ libraries publicly available in the Point Cloud Library (PCL) [8]. We have tested our approach on three datasets, D1, D2, and D3, that capture 3D geometry of the interior of the building hosting the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley. Dataset D1 represents a T-shaped hallway, D2 three hallways connected by two staircases, and D3 a complex floor plan with interconnected rooms and hallways. We preprocess the point cloud data by applying voxel-grid filtering in order to reduce the noise. This filtering procedure divides the input data into voxels and approximates all the points within each voxel with their centroid. Table 1 lists the system parameter values and the statistics on the processed datasets, including processing times on a 2.13 Ghz Intel Xeon PC. The majority of processing time is devoted to point-wise PCA, which has a complexity of $O(N(\log N))$, where N is the number of points in the point cloud. Since dataset D1 is sparser than D2 and D3, for D1 we employ a larger d_{eucl} value and a lower N_p value during region growing. In all datasets, we discard any planar primitives with an area less than 0.5 m^2, as they usually represent noisy regions.

The final planar 3D models are shown in Fig. 9. Note that, in some places, the point cloud data are sparse and fairly noisy. Nevertheless, our system successfully generates simple models while discovering both, small-scale architectural structures, such as walls and staircases, and large-scale architectural structures, such as ceilings and floors. It is important to mention that, in the majority of the cases, a single planar primitive represents each detected wall,

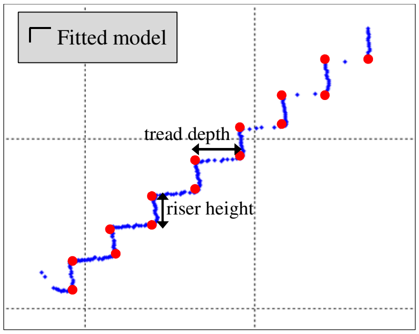

Fig. 7. Pitch scanner data of a sample staircase. Red dots represent the beginning and end points of risers and treads detected after model fitting.

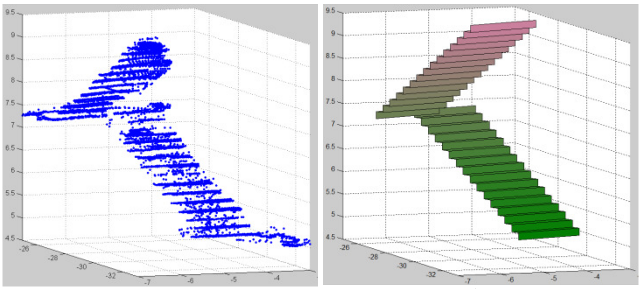

(a) (b)

Fig. 8. Section of (a) a sample *remaining*-subset comprising two valid staircase ramps and (b) the corresponding fitted models.

Table 1. System parameters and statistics on processed datasets

Dataset	Points	σ	q	d_{eucl}	N_p	Detected planes	Time
D1	546,563	0.05m	0.05m	0.15m	500	15	14.6 min
D2	4,683,693	0.05m	0.05m	0.10m	1000	201	145.4 min
D3	6,747,661	0.05m	0.05m	0.10m	1000	93	213.3 min

σ: PCA radius. q: threshold for RANSAC. d_{eucl}: Euclidian distance for region growing. N_p: minimum number of points per region.

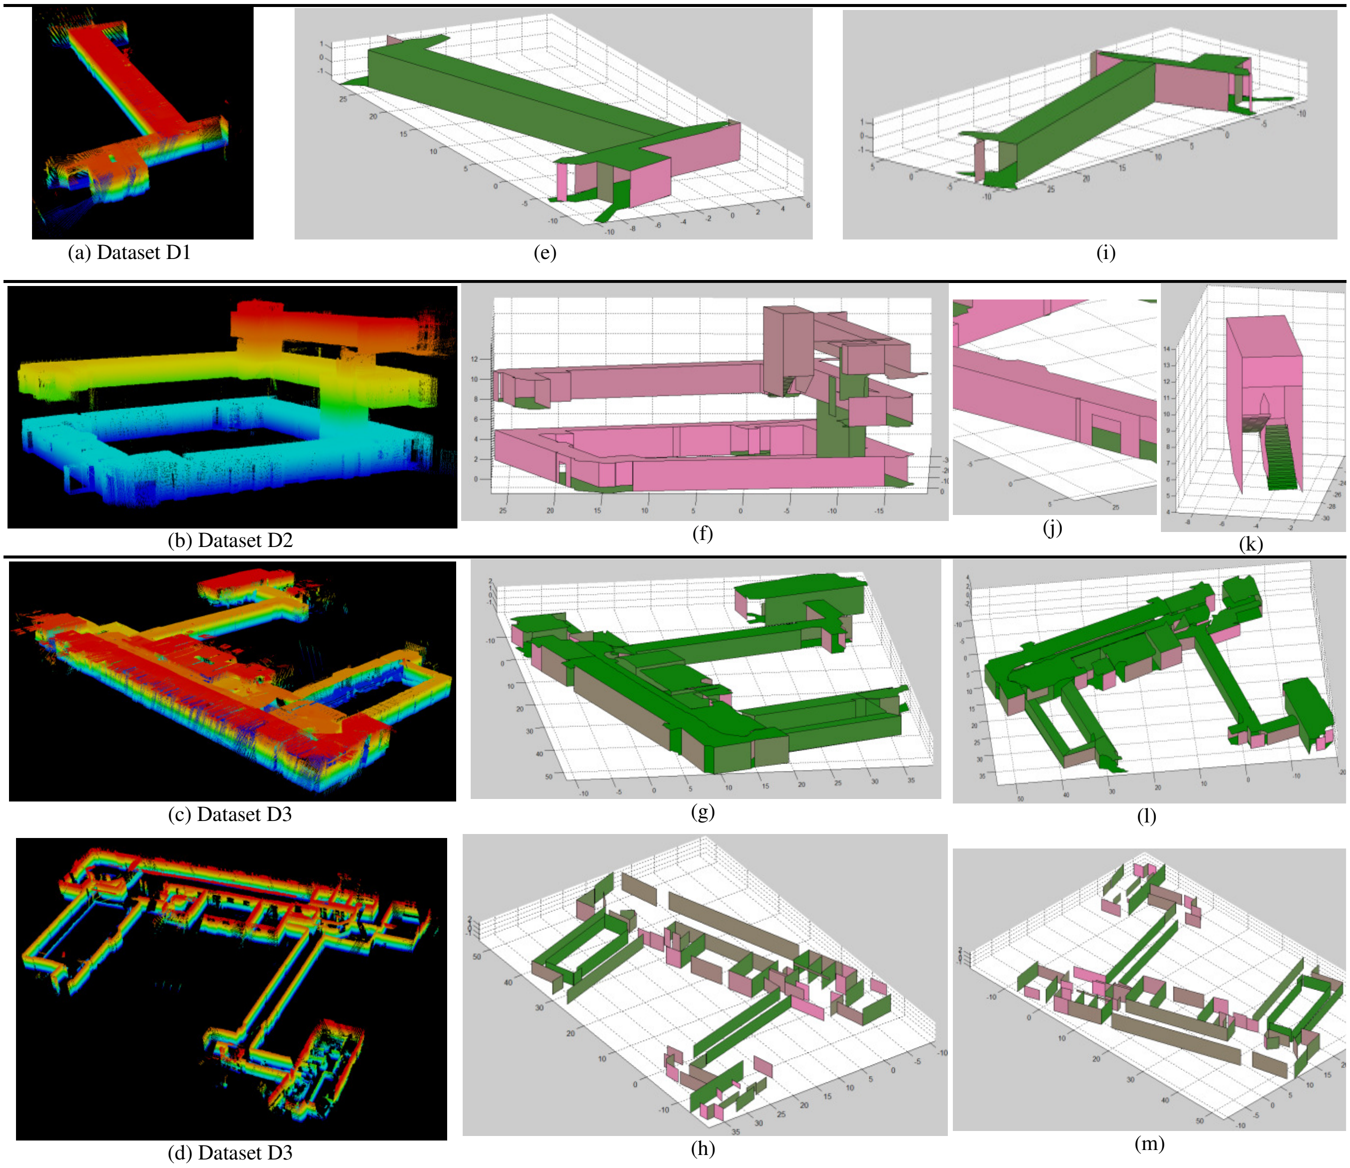
(a) Dataset D1

(e)

(i)

(b) Dataset D2

(f)

(j)

(k)

(c) Dataset D3

(g)

(l)

(d) Dataset D3

(h)

(m)

Fig. 9. Planar 3D models of datasets. (a)-(c) Input point cloud data. (d) View of dataset D3 with ceilings and floors removed to facilitate visualization of interconnected rooms and hallways. (e)-(m) Different views of the corresponding planar 3D models.

ceiling and floor. Note in Fig. 9(j) the benefits of using alpha-shapes and split-and-merge to approximate the spatial extent of polygonal walls, in this case, a wall with an open door.

4. CONCLUSIONS

We have presented an automatic system for planar 3D modeling of building interiors from point cloud data produced by range scanners. The system is capable of generating simple models in terms of a set of planar primitives. We have demonstrated the benefits of our system on challenging point clouds representing building interiors with various staircases, interconnected rooms and hallways. Future work involves modeling non-planar walls with curvature.

5. REFERENCES

[1] J. Chen and B. Chen, “Architectural Modeling from Sparsely Scanned Range Data,” *Int. J. Comput. Vis.*, vol. 78, pp. 223-236, 2008

[2] P. Jenke, B. Huhle, and W. Straßer, “Statistical Reconstruction of Indoor Scenes,” in *Proc. of WSCG 2009*, Plzen-Bory, Czech Republic, 2009

[3] G. Chen, J. Kua, S. Shum, N. Naikal, M. Carlberg, and A. Zakhor, “Indoor Localization Algorithms for a Human-Operated Backpack System,” *in Proc. Int. Symp. on 3DPVT*, 2010

[4] T. Liu, M. Carlberg, G. Chen, J. Chen, J. Kua, and A. Zakhor, “Indoor Localization and Visualization using a Human Operated Backpack System,” *in Int. Conf. on Indoor Positioning and Indoor Navigation*, 2010

[5] M. Fischler and R. Bolles, “Random Sample Consensus: a Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography,” *Commun. ACM*, vol. 24, no. 6, pp. 381-395, 1981

[6] H. Edelsbrunner, D.. Kirkpatrick, and R. Seidel, “On the Shape of a Set of Points in the Plane,” *IEEE Trans. Inform. Theory*, vol. 29, pp. 551-559

[7] U. Ramer, “An Iterative Procedure for Polygonal Approximation of Plane Curves,” *Comp. Graphics and Image Processing*, pp. 244-256, 1972

[8] Available online: http://pointclouds.org