Computers & Graphics 35 (2011) 719–725

Contents lists available at ScienceDirect

Computers & Graphics

journal homepage: www.elsevier.com/locate/cag

Short Communication to SMI 2011

Identifying rectangles in laser range data for urban scene reconstruction

Thijs van Lankveld *, Marc van Kreveld, Remco Veltkamp

Utrecht University, Department of Information and Computing Sciences, Princetonplein 5, 3584CC Utrecht, The Netherlands

ARTICLE INFO

Article history:
Received 10 December 2010
Received in revised form
13 February 2011
Accepted 7 March 2011
Available online 21 March 2011

Keywords:
Boundary representations
Computational geometry
Object modeling
Geometric algorithms

ABSTRACT

In urban scenes, many of the surfaces are planar and bounded by simple shapes. In a laser scan of such a scene, these simple shapes can still be identified. We present a one-parameter algorithm that can identify point sets on a plane for which a rectangle is a fitting boundary. These rectangles have a guaranteed density: no large part of the rectangle is empty of points. We prove that our algorithm identifies all angles for which a rectangle fits the point set of size n in $O(n\log n)$ time. We evaluate our method experimentally on 13 urban data sets and we compare the rectangles found by our algorithm to the α-shape as a surface boundary.

1. Introduction

Automatic reconstruction of 3D geometric models is currently seeing an increase in demand. Applications like navigation, serious games for training, and urban planning require detailed and accurate models of very large data sets. Because of the size of the data sets and the constant changes to the scenes, the reconstruction process should be automated as much as possible.

Some data sources are directly useful for reconstruction of urban scenes. Photographs are easily obtained and ground-based and aerial views can quickly cover a large region. However, while photographs are very suitable for finding the edges of structures, they are less suitable for accurate positioning of the surfaces in the scene. By contrast, laser range scans enable accurate positioning, but do not have high quality color data.

In recent years, the resolution of laser range scanners has improved drastically and currently a single aerial scanning pass can produce a data set with a typical density of 50–100 data points per square meter [10]. Fig. 1 shows the data in one of the regions used in our experiments with the data points colored by surface. Some methods have successfully combined images and laser range scans [23], but to allow applications where no images are available we use only laser range scans.

Because there are many planar surfaces in urban scenes and few major directions, many of the surface boundaries have sharp, straight edges and right angles. In many urban data sets, rectangles are the most prevalent shapes, bounding 35% of the surfaces on average. Efficiently identifying these rectangles will solve a significant part of the geometry reconstruction problem.

We present an efficient algorithm that identifies the planes for which a rectangular boundary fits the data. We define "fitting" in a one-parameter coverage criterion that ensures a minimum local density across the whole rectangle. For each rectangular surface the algorithm produces the range of rotation angles for which a rectangle fits the point set. Experimental results show the quality of the algorithm in identifying "rectangular" point sets.

1.1. Problem

We aim to find the clusters of points in a plane for which a rectangle is the correct boundary. Intuitively, a rectangle fits as a boundary if no large parts are empty, i.e. have a low local density. We base a measure for this local density on the radius of an empty disk.

Definition 1. The δ-coverage region of a point set S in a plane is the union of disks in the plane with radius δ and center $c \in S$. A polygon $\mathcal{P}$ is δ-covered by point set S if and only if it is inside the δ-coverage region of S.

We use the δ-coverage region to determine the correctness of a boundary, because it has a clear geometric meaning. Besides indicating the location of the point set, this structure also is a way to handle any noise model if we can expect an error less than δ. The value of δ has another intuitive geometric meaning: the choice of δ is tied to the resolution of the laser scanner, as each disk inside the boundary should ideally contain multiple data points.

* Corresponding author.
E-mail addresses: thijsvl@cs.uu.nl (T. van Lankveld), marc@cs.uu.nl (M. van Kreveld), Remco.Veltkamp@cs.uu.nl (R. Veltkamp).

0097-8493/$ - see front matter © 2011 Elsevier Ltd. All rights reserved.
doi:10.1016/j.cag.2011.03.004

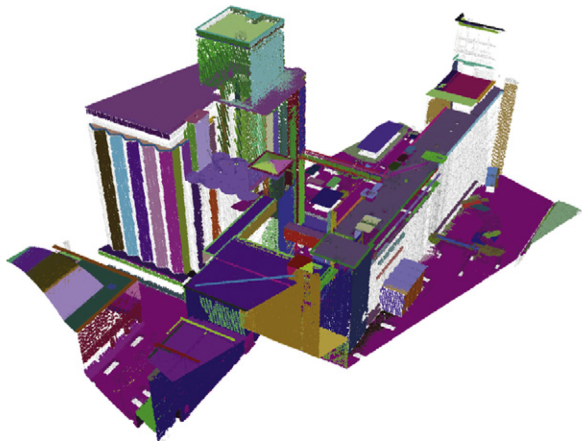

Fig. 1. One of the regions used in the experiments. Each cluster has its own color and remaining unclustered points are black. A detail of this scene is shown in Fig. 4. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Obviously, any rectangle bounding a point set must contain the convex hull $\mathcal{CH}$ of the set. Therefore, if the δ-coverage region does not contain $\mathcal{CH}$, no rectangle bounding the set can be δ-covered. If the δ-coverage region does contain $\mathcal{CH}$, there is a buffer in which a δ-covered rectangle may be placed. When reconstructing a complete scene the different rectangles should be connected along an edge. Therefore, we do not seek one optimal rectangle, but the class of all δ-covered rectangles from which an appropriate rectangle can be selected.

1.2. Overview

After examining related work in Section 2, we present our algorithm in Section 3. Our algorithm identifies rectangles that fit the data, and the way this is achieved is explained in Section 3.1. The algorithm scales well with the data size, which is proven in Section 3.2. We present the setup and results of our experiments in Section 4. We evaluate these results in Section 5 and discuss the broader implications in Section 6.

Our key contributions are:

- A novel approach for automatic urban reconstruction as the creation of a geometric model consisting of simple surfaces that fit the data. Our resulting geometry is simpler than smooth surfaces, and our paradigm is more general than building-model fitting, by allowing more building types.
- A new one-parameter algorithm that identifies clusters for which a rectangular boundary is appropriate and all angles for which a rectangle fits well. The algorithm has a time efficiency of $O(n\log n)$, where n is the size of the point set.
- An analysis of the parameter settings for which the algorithm works best. The same settings can be used in the α-shape [8], to which we compare our results.

2. Related work

Geometry reconstruction from laser range scans can be divided into interactive and automatic methods. While interactive methods like SmartBoxes [14] significantly speed up manual reconstruction, their reliance on a human operator makes them less appropriate for handling massive data sets within limited time.

Recent research into automatic geometry reconstruction from laser range scans has focused on smooth surfaces [2,11,20]. A likely reason for using smooth surfaces is that traditionally most high-density laser range scans were made using close range measurements of natural objects in a controlled environment. The Stanford Bunny, Dragon, and Happy Buddha [18] are well known examples of this type of objects.

Unlike these methods, we process urban scenes, which mainly consist of surfaces that are part of simple primitives like planes, spheres, and cylinders. The transitions between these surfaces are not expected to be smooth and the points are not expected to be exactly on the surface due to noise. Another important consideration is that urban scenes can contain a large number of outliers, generated by vegetation and a generally less controlled environment. Usually, smooth surface reconstruction methods have difficulty in identifying these artifacts and use heuristics to solve this problem.

Related research in geosciences has focused on fitting models of complete buildings to laser range data [4,17] or modeling only roof planes and vertical walls [24]. This way the scene can still be reconstructed from very sparse data if the possible shapes of the buildings are modeled a priori. However, this approach is limited by the number and complexity of the building models and is greatly influenced by vegetation. Unlike these methods, we use dense data sets and reconstruct each individual surface. Combining the surfaces will produce a geometric model equivalent to the predefined building models when this fits the data. In other cases, our building shapes are not limited to models determined a priori.

Some earlier methods for urban reconstruction classify data points into vegetation and buildings, and cluster the data set into a point set per surface [16,20]. However, limited effort has been invested in creating an explicit geometric boundary for these clusters. Schnabel et al. used the primitives to create an implicit model of the scene and made it explicit using marching cubes [15]. However, marching cubes output lacks the simplicity and elegance of the primitive shape geometry.

Various methods have been developed for computing the boundary of a set of unordered points in the plane. These methods can broadly be divided into two groups. The first group, including methods like the crust [3] and γ-neighborhood graph [22], assumes all points lie on the boundary. The second group, including the α-shape [8] and $\mathcal{A}$-shape [13], assumes the boundary contains all points within its interior. Our problem is most related to the second group, and our results are compared to the α-shape in Section 4.1.

3. Rectangular boundaries

Our algorithm determines whether a cluster of points in a plane can be bounded by a rectangle that does not have a large region void of the points. The problem of identifying all δ-covered rectangles can be reduced to identifying the angles of rotation for which there is a δ-covered rectangle. All other δ-covered rectangles will have edges parallel to an identified rectangle.

Because a δ-covered rectangle must lie within a buffer around the convex hull, at any given angle we choose the minimal area bounding rectangle. Our algorithm can easily be adapted to find other predefined convex shapes such as triangles, as shown in [21]. We chose rectangles because these occur most often in urban scenes. The algorithm may be adapted to identify non-convex shapes with predefined angles, like L-shapes, but at the cost of a decreased efficiency and simplicity.

The algorithm is presented in Section 3.1, but for completeness we first explain the preprocessing step used to cluster the data as well as the α-shape.

The input of our algorithm is a point set in a plane. As the laser range data sets consist of unordered points in 3-space measured