ISPRS Journal of Photogrammetry and Remote Sensing 104 (2015) 88–100

Contents lists available at ScienceDirect

ISPRS Journal of Photogrammetry and Remote Sensing

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

Octree-based region growing for point cloud segmentation

Anh-Vu Vo [a], Linh Truong-Hong [a], Debra F. Laefer [a,b,c,*], Michela Bertolotto [d]

[a] Urban Modelling Group, School of Civil, Structural and Environmental Engineering, University College Dublin, Ireland
[b] Earth Institute, University College Dublin, Ireland
[c] U3D Printing Hub, University College Dublin, Ireland
[d] School of Computer Science and Informatics, University College Dublin, Ireland

ARTICLE INFO

Article history:
Received 15 September 2014
Received in revised form 19 December 2014
Accepted 20 January 2015
Available online 24 March 2015

Keywords:
Segmentation
LiDAR
Octree
Voxelization
Point cloud
Building reconstruction

ABSTRACT

This paper introduces a novel, region-growing algorithm for the fast surface patch segmentation of three-dimensional point clouds of urban environments. The proposed algorithm is composed of two stages based on a coarse-to-fine concept. First, a region-growing step is performed on an octree-based voxelized representation of the input point cloud to extract major (coarse) segments. The output is then passed through a refinement process. As part of this, there are two competing factors related to voxel size selection. To balance the constraints, an adaptive octree is created in two stages. Empirical studies on real terrestrial and airborne laser scanning data for complex buildings and an urban setting show the proposed approach to be at least an order of magnitude faster when compared to a conventional region growing method and able to incorporate semantic-based feature criteria, while achieving precision, recall, and fitness scores of at least 75% and as much as 95%.

1. Introduction

Light Detection and Ranging (LiDAR), also known as laser scanning has been used as a cost-effective and reliable tool for documenting the urban environment. Laser scanning collects three-dimensional (3D) geometric points of object surfaces called point clouds. While point clouds are often used strictly for visualization or simple distance or angular measurement, their potential value is much greater, as seen in their application to road detection (Clode et al., 2004; Boyko and Funkhouser, 2011), disaster management (Laefer and Pradhan, 2006), and building reconstruction (Pu and Vosselman, 2009). Since buildings are complex, sections representing various surfaces can be identified individually, with subsequent aggregation for model reconstruction.

Segmentation is not a trivial task, as the point cloud datasets are unstructured and often massive. Additionally, many LiDAR point cloud processing tasks (including the investigated region growing algorithms) require hand-tuning. Consequently, such a process usually necessitates human intervention and can be quite time consuming. Improving speed of those algorithms even with a minor penalty of accuracy is useful in many applications. For example, in structural engineering, an offset of few centimeters can usually be considered as negligible. On the other hand, a faster program encourages users to try many different input thresholds and, therefore, increases the chance of selecting good input values, which could lead to even better results. Since data segmentation is computationally intensive, this paper proposes an efficient method for segmentation of building point data acquired from terrestrial or aerial laser scanning.

2. Related works

Segmentation plays an important role in building reconstruction from LiDAR data, which refers to the task of partitioning a 3D point cloud into subsets satisfying certain pre-selected criteria (Vosselman and Mass, 2010; Biosca and Lerma, 2008). Segmentation algorithms can be roughly classified as model fitting-based methods (e.g. Vosselman and Dijkman, 2001; Overby et al., 2004; Schnabel et al., 2007), region growing-based methods (e.g. Gorte, 2002; Tóvári and Pfeifer, 2005; Vieira and Shimada, 2005) and clustering feature based methods (e.g. Filin, 2002; Biosca and Lerma, 2008; Hofmann, 2003; Lari and Habib, 2014).

* Corresponding author at: School of Civil, Structural and Environmental Engineering, U3D Printing Hub, Earth Institute, University College Dublin, Newstead, Belfield, Dublin 4, Ireland. Tel.: +353 1 716 3226.

E-mail addresses: anh-vu.vo@ucdconnect.ie (A.-V. Vo), linh.truong-hong@ucd.ie (L. Truong-Hong), debra.laefer@ucd.ie (D.F. Laefer), michela.bertolotto@ucd.ie (M. Bertolotto).

http://dx.doi.org/10.1016/j.isprsjprs.2015.01.011
0924-2716/

2.1. Model fitting-based methods

As part of the model fitting-based category of approaches (sometimes referred to as the parameter-based approach), two widely employed algorithms are the Hough Transform (HT) (Ballard, 1981) and the Random sample consensus (RANSAC) approach as proposed by Fischler and Bolles (1981). The HT is used to detect planes (Vosselman et al., 2004), cylinders (Tarsha-Kurdi, 2007), and spheres (Rabbani and Heuvel, 2005). To accelerate the HT process and to increase the reliability of the results, Vosselman et al. (2004) determined the parameters through several separate steps. For example, plane identification employed two steps: (1) determination of the plane normal vector and (2) establishment of the distance from the plane to the origin.

The RANSAC paradigm is used to extract shapes by randomly drawing minimal data points to construct candidate shape primitives. The candidate shapes are checked against all points in the data set to determine a value for the number of the points that represents the best fit (Fischler and Bolles, 1981). The method has been adapted to segment 3D point clouds. For example, using a standard RANSAC approach, Boulaassal et al. (2007) manually determined the optimal tolerance for extracting building façade planes from TLS data, while Schnabel et al. (2007) adapted RANSAC for precise and fast plane extraction. In the work by Schnabel et al. (2007), a point cloud with normal vectors was used to verify all sampling points before assigning a candidate shape. An octree was employed for efficiently extracting sampling points. That proposed method reduced running time from hours to seconds and enabled the processing of several million points in less than a minute. Similarly, Chen et al. (2014) improved the RANSAC algorithm through the localized sampling, to segment the polyhedral rooftop primitives and then through the application of a region growing based triangulated irregular network (TIN) to separate the coplanar primitives. The enhanced algorithm performed especially well with noisy data sets and was shown to be computationally efficient. In related work, Awwad et al. (2010) modified the RANSAC algorithm by using a clustering technique to divide the dataset into small clusters based on normal vectors of the points. To prevent spurious surfaces appearing in a group, the sum of the perpendicular distances between the points and a local surface was imposed as the condition to decide on whether the plane be accepted within the group or eliminated as being outside of the group.

As demonstrated above, the HT and RANSAC algorithms are well established as robust methods for segmenting 3D point clouds, with the RANSAC algorithm having the advantage of being able to do so for datasets with high noise and outliers. However, these algorithms have some disadvantages. First, since these algorithms only use point positions, many spurious planes that do not exist in reality may be generated. Second, the segmentation quality is sensitive to the point cloud characteristics (density, positional accuracy, and noise). Third, the algorithms perform poorly with large datasets or those with complex geometries. HT especially requires significant processing time and high memory consumption for large data sets, because all parameters must be stored. Furthermore, HT is very sensitive to the selection of surface parameters (Awwad et al., 2010; Tarsha-Kurdi, 2007).

2.2. Region growing-based methods

An alternative to model fitting methods are region growing based ones. The method introduced by Besl and Jain (1988) involved two stages: a coarse segmentation based on the mean and Gaussian curvature of each point and its sign, and a refinement using an iterative region growing based on a variable order bivariate surface fitting. This method was then adopted by others for 3D point cloud segmentation. For example, Gorte (2002) performed a region growing segmentation using a TIN as the seed surface and the angle and distance between the neighboring triangles for the growing. The seed region was used to merge the triangles. In contrast, Tóvári and Pfeifer (2005) used normal vectors, the distance of the neighboring points to the adjusting plane, and the distance between the current point and candidate points as the criteria for merging a point to a seed region that was randomly picked from the data set after manually filtering areas near edges. Nurunnabi et al. (2012) also used these criteria but with the seed point being that having the least curvature.

As the points in the interior region have been shown to be good seed points, Vieira and Shimada (2005) firstly removed points along sharp edges using a curvature threshold. Median filtering was then performed to reduce noise, and the remaining points were considered as seed points. Rabbani et al. (2006) proposed as an alternative the residual of a plane fitting to select the seed points followed by region growing using an estimated point normal and the residual. Similar to Besl and Jain (1988), Ning et al. (2009) proposed a two-step, region growing segmentation: rough segmentation to extract points on the same plane based on normal vectors followed by fine segmentation to extract architectural details based on the residual that is the distance from the point to the local shape. To improve efficiency and robustness of the region growing method, Deschaud and Goulette (2010) introduced the area of the local plane as a criterion for selecting the seed region and then employed an octree to search the neighboring points of those that should possibly be merged to the seed region.

As part of this approach to segmentation, region growing based on octrees or voxel grids have been introduced to improve efficiency. An example of this was the work done by Woo et al. (2002) to recursively divide a dataset into smaller voxels until the standard deviation of the voxels' normal vectors was less than the threshold, where the voxel's normal vector is the average of the point normal vectors. In this approach, the segmentation extracted edge-neighborhood points possessed by the voxels having the normal vector when the size was smaller than a predefined cell size threshold. The adjacent voxels were then merged to the leaf node, if the deviation of the voxels' normal vectors were less than the tolerance. Similarly, Wang and Tseng (2004) used the residual distance and an area of the data points within the voxels as the criteria for subdividing an initial bounding box. The voxels on the same layer having similar normal vectors were classified as being in the same group, where a normal vector of the voxel was computed from data points within the voxel. Subsequently, Wang and Tseng (2011) proposed splitting and merging algorithms to extract coplanar points from a connected voxel group. The region growing was then used to merge a group of coplanar points based on the angle variation of the fitting planes among them. In these works, the voxel size and how the normal vector of the voxel was computed strongly influenced the segmentation results.

While region growing-based methods are widely used in segmenting 3D point clouds as the methods are easily implemented, they are not particularly robust as has been shown experimentally (e.g. Woo et al., 2002; Biosca and Lerma, 2008; Teboul et al., 2010) in part because the segmentation quality strongly depends both on multiple criteria and the selection of seed points/regions, where no universally valid criterion exists (Biosca and Lerma, 2008; Teboul et al., 2010). Additionally, they also require extensive computing time when 3D point clouds are used (Woo et al., 2002; Boulaassal et al., 2007).

2.3. Clustering feature-based methods

Another major segmentation approach employs clustering of features. For example, Filin (2002) employed the formation of a feature space and a mode-seeking algorithm based on seven