ISPRS Journal of Photogrammetry and Remote Sensing 113 (2016) 59–74

Contents lists available at ScienceDirect

ISPRS Journal of Photogrammetry and Remote Sensing

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

Octree-based segmentation for terrestrial LiDAR point cloud data in industrial applications

Yun-Ting Su [a,*], James Bethel [a], Shuowen Hu [b]

[a] Geomatics Engineering, School of Civil Engineering, Purdue University, 550 Stadium Mall Drive, West Lafayette, IN 47907, United States
[b] U.S. Army Research Laboratory, 2800 Powder Mill Road, Adelphi, MD 20783, United States

ARTICLE INFO

Article history:
Received 18 April 2015
Received in revised form 25 December 2015
Accepted 4 January 2016
Available online 19 January 2016

Keywords:
Terrestrial LiDAR
Point cloud segmentation
Split and merge
Piping systems
Cylinder detection

ABSTRACT

Automated and efficient algorithms to perform segmentation of terrestrial LiDAR data is critical for exploitation of 3D point clouds, where the ultimate goal is CAD modeling of the segmented data. In this work, a novel segmentation technique is proposed, starting with octree decomposition to recursively divide the scene into octants or voxels, followed by a novel split and merge framework that uses graph theory and a series of connectivity analyses to intelligently merge components into larger connected components. The connectivity analysis, based on a combination of proximity, orientation, and curvature connectivity criteria, is designed for the segmentation of pipes, vessels, and walls from terrestrial LiDAR data of piping systems at industrial sites, such as oil refineries, chemical plants, and steel mills. The proposed segmentation method is exercised on two terrestrial LiDAR datasets of a steel mill and a chemical plant, demonstrating its ability to correctly reassemble and segregate features of interest.

1. Introduction

Light detection and ranging (LiDAR) originated in the early 1960s, with its first applications in meteorology. As LiDAR technology matured, improving in accuracy, speed, and spatial resolution, it has been adopted for use in numerous applications, extending from topographic mapping with airborne LiDAR platforms to surveying of urban sites with terrestrial LiDAR systems, and more recently, to corridor mapping with mobile LiDAR systems/devices mounted on land vehicles. Static terrestrial LiDAR has become an especially effective tool for surveying, in some cases replacing traditional techniques such as electronic total stations and GPS methods. With terrestrial LiDAR, surveyors can scan an entire site at a standoff distance without requiring an individual to occupy the site (Lee, 2011). Current generation LiDAR scanners also have very fine spatial resolution, providing precise 3D point cloud data with millimeter accuracy (Soulard and Bogle, 2011). Typically, a site is surveyed from multiple locations to obtain more complete coverage, with the resulting scans aligned and combined to form a single high density point cloud dataset. Consequently, LiDAR point clouds usually contain hundreds of thousands to tens of millions of individual points, depending on the size of the site being surveyed. Furthermore, LiDAR data is nonuniformly sampled, as scanned points do not lie on a uniform spatial grid. The ultimate goal in most applications involving LiDAR is to process this high density point cloud data and reconstruct a 3D computer-aided design (CAD) model of the scene, first by segmenting the point cloud data into appropriate segments and then recognizing primitives from the segments to generate 3D models.

This work focuses on efficient segmentation of terrestrial LiDAR data of piping systems in industrial sites (e.g. chemical plants, oil refineries, and steel mills). Given an entire nonuniformly sampled point cloud of a scene, the objective is to perform segmentation and extract individual segments, determining which segments are likely to be pipes, vessels, or walls. We propose a robust octree-based split and merge segmentation algorithm that can efficiently process large LiDAR data. After initially splitting the dataset into octants (referred to as voxels in this work) based on point density using octree decomposition, the points within each voxel are further split into spatially unconnected components using graph theory based analysis. Following splitting, the merging process uses a series of connectivity criteria (proximity, orientation, and curvature) to intelligent merge components together. The novel split and merge procedures are the key contributions of this work. This proposed segmentation algorithm is a bottom-up approach that is scalable and parallelizable.

The organization of the manuscript is as follows. Section 2 describes prior, related work on the segmentation of terrestrial LiDAR data. Section 3 provides a detailed description of the

* Corresponding author.
E-mail address: yunting.su2015@gmail.com (Y.-T. Su).

http://dx.doi.org/10.1016/j.isprsjprs.2016.01.001
0924-2716/

proposed octree-based split and merge segmentation algorithm. Section 4 shows the results of the proposed algorithm on several LiDAR datasets, comparing with two other published techniques. Section 5 presents the conclusions of this work.

2. Background

The earliest studies on the segmentation of range data used sensors that acquired 2.5D range images, also referred to as depth maps, that lie on a uniform spatial grid – each regularly spaced point (x, y) on a rectangular grid has a range measurement. Henderson and Bhanu (1982) developed a planar region growing algorithm for range images using a spatial proximity graph. Hebert and Ponce (1982) proposed a method to segment depth maps into primitives such as planes and cylinders by mapping point surface normals to the Gaussian sphere. An edge-region segmentation ring operator was developed by Inokuchi et al. (1982). Besl and Jain (1985) provided an excellent literature review on the studies involving range image analysis. Unlike range images, point clouds from LiDAR sensors are 3D data that do not lie on a uniform spatial grid. While many concepts can be adopted from the earlier work with range images, there has been substantial progress recently in the segmentation of LiDAR point cloud data.

As mentioned in Section 1, LiDAR sensors are categorized as either airborne or terrestrial. Many techniques have been developed for segmentation of airborne LiDAR data. Arefi and Hahn (2005), Zhao et al. (2011), Li et al. (2013), Yan et al. (2015), Poullis and You (2009), and McLaughlin (2006) are a few representative techniques. Due to the top-down perspective of airborne LiDAR platforms, airborne LiDAR point clouds predominantly contain planar surfaces, especially for urban scenes. Therefore, airborne LiDAR data are often considered to be 2.5D data – for example, a substantial fraction of the points in an airborne LiDAR point cloud of an urban scene would lie on roof surfaces, but few points would be acquired of building walls (Zhou and Neumann, 2010). Terrestrial LiDAR can be subdivided into mobile platforms or static systems, which is the focus of this work. However, segmentation of terrestrial LiDAR data, especially scenes of piping systems, has received relatively less attention. The remainder of this section is devoted to discussing the relevant works in this area.

Rabbani (2006) introduced a smoothness constraint based segmentation technique that is one of the most widely cited works on segmentation of terrestrial LiDAR data. Rabbani's technique is a bottom-up approach with two main steps: normal vector estimation and region growing. In the first step, the surface normal for each point is estimated by fitting a plane to its neighbors, found through the k-nearest neighbors method. The residual of the plane fitting to the neighbors of a point is used by Rabbani (2006) to approximate the local surface curvature. A small residual indicates that the neighbors lie on a planar surface, while a large residual indicates a more curved surface. However, a large residual may also be due to noise. Following computation of the surface normal of every point, the second step of region growing is initiated with a seed point that has the smallest residual from the first step. The neighboring points of this seed point with residual below a set threshold are added to the list of available points for consideration, and a surface smoothness constraint determines whether these available points are added to the current region. The surface smoothness constraint is implemented by considering the angle between the surface normal of a seed point and the surface normals of its neighbors. If this angle is below a certain threshold, typically set at 15° (Rabbani, 2006), this point is added to the region and updated to be the current seed point. The process continues iteratively until the list of available points is exhausted, and then a new region is initiated using the point with the smallest residual from the remaining points.

Rabbani's segmentation technique has two limitations, as we observed through experimentation using our Matlab implementation of Rabbani's algorithm. First, regions linked together by a smooth connector are segmented as a single region. For example, a vertical pipe connected to a horizontal pipe via a long radius elbow joint would exhibit smoothly varying surface normals from one end to the other, and would be segmented as a single region instead of three separate regions, as typically would be desired. Rabbani (2006) also recognized this concern, but reasoned that this under-segmentation is more preferable to over-segmentation. The second limitation is the computational complexity of the algorithm, which requires the k-nearest neighbors (KNN) for every point in the dataset to be computed. The linear search solution for KNN has a running time of $O(Nd)$, where N is the number of points and d is the dimensionality of the data. Elseberg et al. (2012) provide an excellent comparison of different nearest-neighbor search strategies. For typical LiDAR point clouds that contain hundreds of thousands to tens of millions of points, computing the KNN of every point is computationally prohibitive. Space partitioning methods such as k–d trees have been applied to KNN search (Friedman et al., 1977), reducing the search complexity to $O(\log N)$, but involve an offline phase to construct the k–d tree.

Schnabel et al. (2007) proposed an efficient random sample consensus (RANSAC) algorithm for large scale point cloud shape detection, using a hierarchically structured sampling strategy for generating different types of primitive shapes which can significantly reduce computational runtime. The major deficiency of RANSAC is its computational demand (Schnabel et al., 2007), and it must be heavily optimized for practical processing.

Wang and Tseng (2010) introduced an incremental segmentation technique using an octree-structured voxel space. Their octree based split and merge segmentation algorithm first divides the input point cloud into octree subspaces (i.e. voxels) until each voxel only contains coplanar points during the splitting process. Coplanarity is measured by computing the residuals of plane fitting in a voxel, similar to the plane fitting procedure of Rabbani (2006). If the variance of the residuals exceeds a user defined threshold, indicating that the points do not form a coplanar surface, the node is further subdivided into eight child voxels. Following the splitting procedure, Wang and Tseng (2010) perform a merging procedure, checking whether adjacent planes have similar surface normal orientations and are sufficiently proximate to be merged into a single plane. Since this technique only focuses on coplanarity during the split and merge steps, it is more suited for segmentation of airborne LiDAR data than for industrial scenes. Wang and Tseng (2011) extended their 2010 technique, proposing a four step procedure. First, Wang and Tseng (2011) use an octree decomposition to divide the point cloud into voxels with dimensions similar to the LiDAR point spacing, resulting in relativity small voxels. Then, voxels are connected together based on spatial proximity into groups of points via connected component labeling (CCL). The third step is split and merge, based on their 2010 work on airborne LiDAR (Wang and Tseng, 2010) which introduces the co-planarity criterion. Their co-planarity criterion is used to segments groups of points which lie on the same plane. The fourth step of Wang and Tseng (2011) then combines planes with similar orientations together using an angular threshold, enabling this 4-step procedure to be applied to terrestrial LiDAR point clouds. One major difference between the proposed approach and Wang and Tseng (2011) is the criterion for the merge procedure. While Wang and Tseng (2011) use the coplanar criterion, we base our approach on the orientation criterion. Though cylinders can be represented as being composed of a ring of planar strips which are then merged together into a cylinder in the fourth step (Wang and Tseng, 2011), we do not believe this is an ideal representation. Piping systems typically consist of pipes of varying sizes, so a single angular