ISPRS Journal of Photogrammetry and Remote Sensing 129 (2017) 118–130

Contents lists available at ScienceDirect

ISPRS Journal of Photogrammetry and Remote Sensing

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

A voting-based statistical cylinder detection framework applied to fallen tree mapping in terrestrial laser scanning point clouds

Przemyslaw Polewski [a,b,*], Wei Yao [a], Marco Heurich [c], Peter Krzystek [a], Uwe Stilla [b]

[a] Dept. of Geoinformatics, Munich University of Applied Sciences, 80333 Munich, Germany
[b] Photogrammetry and Remote Sensing, Technische Universitaet Muenchen, 80333 Munich, Germany
[c] Dept. for Conservation and Research, Bavarian Forest National Park, 94481 Grafenau, Germany

ARTICLE INFO

Article history:
Received 13 February 2017
Received in revised form 27 April 2017
Accepted 27 April 2017

Keywords:
TLS
Statistical shape detection
Fallen trees
LIDAR
Cylinder

ABSTRACT

This paper introduces a statistical framework for detecting cylindrical shapes in dense point clouds. We target the application of mapping fallen trees in datasets obtained through terrestrial laser scanning. This is a challenging task due to the presence of ground vegetation, standing trees, DTM artifacts, as well as the fragmentation of dead trees into non-collinear segments. Our method shares the concept of voting in parameter space with the generalized Hough transform, however two of its significant drawbacks are improved upon. First, the need to generate samples on the shape's surface is eliminated. Instead, pairs of nearby input points lying on the surface cast a vote for the cylinder's parameters based on the intrinsic geometric properties of cylindrical shapes. Second, no discretization of the parameter space is required: the voting is carried out in continuous space by means of constructing a kernel density estimator and obtaining its local maxima, using automatic, data-driven kernel bandwidth selection. Furthermore, we show how the detected cylindrical primitives can be efficiently merged to obtain object-level (entire tree) semantic information using graph-cut segmentation and a tailored dynamic algorithm for eliminating cylinder redundancy. Experiments were performed on 3 plots from the Bavarian Forest National Park, with ground truth obtained through visual inspection of the point clouds. It was found that relative to sample consensus (SAC) cylinder fitting, the proposed voting framework can improve the detection completeness by up to 10 percentage points while maintaining the correctness rate.

1. Introduction

With the advent of commercial terrestrial laser scanning (TLS) systems and the steady decrease in their price and size, the remote sensing community gained an invaluable tool for mapping 3D scenes with unprecedented precision and resolution. Over the last 15 years, TLS systems have found numerous applications in tasks as diverse as cultural heritage conservation (Abmayr et al., 2005), damaged building monitoring (Pesci et al., 2013), precision forestry (Kelbe et al., 2015) and many others (Lichti et al., 2008). Particularly in vegetation mapping applications, TLS has changed the way forest inventories are carried out (Liang et al., 2016). The dense point clouds acquired from a terrestrial perspective are a natural complement to aerial laser scanning (ALS) technology since the latter usually cannot capture understory information in the presence of a thick canopy cover (Amiri et al., 2016). Among various forest structure parameters which could be successfully recovered using TLS data to date are: tree position, stem volume, total biomass, leaf area index, gap fraction, etc. (Kankare et al., 2013; Zheng et al., 2013). However, it appears that thus far little attention has been devoted to the automatic detection and mapping of fallen trees. On the other hand, it should be noted that the distribution of downed trees and coarse woody debris in forests is highly interesting from an ecological standpoint, due to the crucial role of dead wood in several forest processes. Perhaps the most prominent is its role in supporting wildlife biodiversity, because up to one third of plant and animal species in forests depend on dead wood for survival (Stokland et al., 2012).

With this study, we aim to bridge this gap by providing an automatic approach capable of delineating individual fallen trees directly from dense point clouds. We target challenging scenarios with ground vegetation, irregular terrain, presence of standing

* Corresponding author at: Dept. of Geoinformatics, Munich University of Applied Sciences, 80333 Munich, Germany.

E-mail addresses: polewski@hm.edu (P. Polewski), yao@hm.edu (W. Yao), marco.heurich@npv-bw.bayern.de (M. Heurich), krzystek@hm.edu (P. Krzystek), stilla@tum.de (U. Stilla).

http://dx.doi.org/10.1016/j.isprsjprs.2017.04.023

trees and overlapping fallen stems. There are several factors which complicate the detection task: occlusion from standing trees and shrubs, incompleteness of captured objects due to scanning positions, as well as varying point density associated with distance to the scanner. In this setting, the most reliable cue for recognizing the fallen trunks is their cylindrical shape. However, the lying stems are only locally cylindrical, because they can be fragmented into non-collinear pieces. Also, many stems may overlap on the ground. Due to all these factors, a method is required which can robustly detect multiple cylinders having arbitrary spatial relationships, in the presence of occlusions and outlier points. Two well-known classes of shape detection approaches meeting these criteria are sample consensus (SAC) and generalized Hough transform type procedures. The primary contribution of this paper is to introduce a robust cylinder detection framework based on voting in parameter space, designed to overcome two problems associated with Hough methods: (i) the high computational complexity due to exhaustive generation of samples from different parametrizations of the target shape and (ii) the sensitivity to the accumulation bin width. To solve (i), our framework does not generate shape samples and instead relies on votes cast by pairs of neighboring points from the input point cloud to determine the parameter space based on relationships between adjacent surface normals intrinsic to cylindrical surfaces. Regarding (ii), we propose to not discretize the parameter space and instead to construct a kernel density estimator (KDE) and locate its maxima using continuous optimization methods. We take advantage of the established and theoretically well founded multivariate KDE bandwidth estimation methods to alleviate the difficulty of manual bandwidth selection. This methodology is an extension of our prior work presented in Polewski et al. (2016), where only the cylinders' orientation was estimated through voting. Together with the idea to decompose the search of the 5D cylindrical parameter space into two sequential steps of separately estimating the orientation and position (Rabbani and van den Heuvel, 2005), the proposed modifications ensure that the voting-based cylinder reconstruction procedure is both robust and computationally feasible. The second major contribution of this work is the entire processing pipeline for obtaining object-level semantic information about individual fallen trees from TLS point clouds. We utilize graph-cut segmentation to merge the detected cylindrical primitives into semantically connected objects. To the best of our knowledge, this is the first attempt to perform such detection and reconstruction from TLS data in a fully automatic manner, although there have been attempts to *manually* delineate the stems (Grigillo et al., 2015). We examine the performance of our method on 3 diverse and challenging test plots from the Bavarian Forest National Park. A comparison with several parametrizations of SAC is carried out, and the influence of the method's substeps on the final segmentation result is investigated.

2. Related work

From the application side, we are not aware of any published prior work regarding automatic fallen tree mapping from TLS point clouds, although several authors describe methods for dead tree detection in ALS data (Polewski et al., 2015; Nyström et al., 2014; Lindberg et al., 2013). Note that the methodology for dense and sparse point clouds is fundamentally different, since the former relies on recognition of cylindrical shapes, which requires a sufficiently high point density currently unavailable in standard ALS data. The closest related application is the segmentation of standing tree stems in TLS point clouds. Lindberg et al. (2012) project candidate stem points onto a 2D plane and apply the Hough transform to find circles corresponding to potential stems. Schilling et al. (2011) employ a similar strategy. In Wang et al. (2016), the authors first perform RANSAC based circle fitting of projected stem points, followed by RANSAC cylinder fitting in 3D and subsequent growing of the initial result. Olofsson et al. (2014) voxelize the point cloud and analyze the occupancy of different height layers to determine stem positions, followed by a RANSAC procedure for radius determination. A common feature of these studies is to utilize only points lying over 1 m over the ground, which eliminates most of the ground vegetation. In Liang et al. (2012), the stem points are found using classification based on features of the local covariance matrix. We utilize a similar pre-processing step to filter out ground points, shrubs and foliage. The authors then perform cylinder fitting with an M-estimator based on Tukey's biweight influence function. Overall, it appears that the detection and modeling of standing stems is a somewhat easier task compared to fallen tree mapping due to the availability of good prior knowledge of the cylinder's orientation (world Z axis), the ability to remove the influence of ground vegetation through DTM filtering, and the fact that neighboring tree stems are usually well separated by empty space. From a methodological perspective, our approach shares some conceptual similarity with the work of Rabbani and van den Heuvel (2005). We also decompose the task of determining all of the cylinder's parameters into simpler subproblems of first determining only the orientation, followed by estimation of the position and radius. An important difference is that in their approach, both of these steps are conducted via Hough transform, which implies generating a high number of point samples. Our method does not generate new samples, but rather considers pairs of nearby input points, reducing the computational effort. Also, the choice of the Hough accumulation space quantization is a deciding factor in the final performance. In our case, the voting is based on a KDE with an automatically determined bandwidth. Note that the idea to apply kernel density estimation for modeling the parameter space in shape detection has been explored before, e.g. Dahyot (2008) defines a statistical Hough transform framework, where the parameters are estimated as modes of a variable-bandwidth KDE. Tran et al. (2015) propose to calculate local cylinder axis orientations by finding vectors perpendicular to subsets of adjacent surface normals, which is similar to our method. They follow an iterative scheme where an initial point set is augmented with points based on the distance to the currently estimated cylinder and surface normal deviation. In contrast, in our method the locally determined orientations are collected over all points of the object candidate and vote for the aggregate axis. This eliminates the dependency on initialization, inherent in iterative approaches. Moreover, in Tran et al. (2015), the estimation of radius and center parameters is done by directly fitting a circle to *all* points considered inliers, whereas our approach estimates these quantities by voting.

3. Method overview

The proposed method accepts a dense 3D point cloud as input. It is assumed that the dead trees are lying on the ground, and their crowns/branches have already decomposed or fallen off, leaving only the stem. The point cloud should be detailed enough to capture at least part of the cylindrical surfaces representing the stems, but a full coverage of each stem from all sides is not required. Our processing pipeline consists of three major steps (Fig. 1). The goal of the first stage is to determine the subset of probable stem point candidates from the entire input point cloud, which is usually dominated by non-relevant points belonging to the ground, the shrub layer, or foliage. This is achieved through DTM filtering, a supervised classification at point level, as well as a connected component segmentation. Section 4 describes these procedures in detail. In the next part of the method, outlined in Section 5, the