Pattern Recognition 48 (2015) 2043–2053

Contents lists available at ScienceDirect

Pattern Recognition

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

Real-time detection of planar regions in unorganized point clouds

Frederico A. Limberger, Manuel M. Oliveira *

Universidade Federal do Rio Grande do Sul, Instituto de Informática - PPGC - CP 15064, 91501-970 - Porto Alegre - RS - Brazil

ARTICLE INFO

Article history:
Received 28 April 2014
Received in revised form
15 November 2014
Accepted 23 December 2014
Available online 14 January 2015

Keywords:
Hough transform
Real-time plane detection
Unorganized point clouds

ABSTRACT

Automatic detection of planar regions in point clouds is an important step for many graphics, image processing, and computer vision applications. While laser scanners and digital photography have allowed us to capture increasingly larger datasets, previous techniques are computationally expensive, being unable to achieve real-time performance for datasets containing tens of thousands of points, even when detection is performed in a non-deterministic way. We present a deterministic technique for plane detection in unorganized point clouds whose cost is $O(n \log n)$ in the number of input samples. It is based on an efficient Hough-transform voting scheme and works by clustering approximately co-planar points and by casting votes for these clusters on a spherical accumulator using a trivariate Gaussian kernel. A comparison with competing techniques shows that our approach is considerably faster and scales significantly better than previous ones, being the first practical solution for deterministic plane detection in large unorganized point clouds.

1. Introduction

Automatic plane detection in point clouds is a key component in many graphics, image processing, and computer vision applications. These include, among others, model reconstruction for reverse engineering [1–5], camera calibration [6], object recognition [7,8], augmented reality [9,10], and segmentation [11,12]. The recent popularization of laser scanners has led to an increasingly growth in the sizes of the available datasets, and point clouds containing tens of millions of samples are now commonplace. Software applications like *SynthExport* [13] and *Photosynth* [14] also allow us to extract point clouds from large collections of digital images. However, existing techniques for detecting planar regions in point clouds are computationally expensive and do not scale well with the size of the datasets. For performance improvement, they often exploit non-deterministic strategies, such as working on a randomly-selected sub-set of the original samples. While this can reduce execution time, these techniques are still unable to achieve real-time performance even on datasets containing just tens of thousands of points. More importantly, their results depend on the selected sample sub-sets and, therefore, there is no guarantee that all relevant planes will be detected, or that such results will be consistent across multiple executions.

We present an efficient technique to perform deterministic plane detection in unorganized point clouds whose cost is $O(n \log n)$ in the number of input samples. Our approach scales well with the size of the datasets, is robust to the presence of noise, and handles point clouds with different characteristics in terms of dimensions and sampling distributions. While the actual running times depend on specific features of the dataset (*e.g.*, the number of planar regions), our technique is several orders of magnitude faster than previous ones. For instance, it processes an entire point cloud with 20-million samples (Bremen dataset) in just 2.1 s on a typical PC. In contrast, efficient versions of RANSAC can take from 12 min to more than 2 h to process the same dataset, while the Randomized Hough transform takes 42.8 s to process only 10% of the samples.

Our technique is based on a robust and fast algorithm to segment point clouds into approximately planar patches, even in the presence of noise or irregularly distributed samples. For this, we use a subdivision procedure to refine an octree and cluster groups of approximately coplanar samples. We use the identified clusters to obtain an efficient Hough-transform voting scheme by casting votes for each of these clusters (instead of for individual samples) on a spherical accumulator. For voting, we use a Gaussian kernel centered at the cluster's best fitting plane, which takes into account the cluster's variances. In this sense, our approach extends the kernel-based voting scheme proposed by Fernandes and Oliveira [15] using a trivariate Gaussian distribution defined over spherical coordinates (θ, ϕ, ρ). While, at first, plane detection in unorganized point clouds might seem as an immediate extension of line detection in images, the lack of explicit neighborhood information among samples imposes significant challenges, requiring new clustering and accumulation-management strategies.

* Corresponding author. Tel.: +55 51 3308 6821; fax: +55 51 3308 7308.
E-mail addresses: falimberger@inf.ufrgs.br (F.A. Limberger), manuel.menezes@gmail.com, oliveira@inf.ufrgs.br (M.M. Oliveira).
URLS: http://www.inf.ufrgs.br/~falimberger (F.A. Limberger), http://www.inf.ufrgs.br/~oliveira (M.M. Oliveira).

http://dx.doi.org/10.1016/j.patcog.2014.12.020

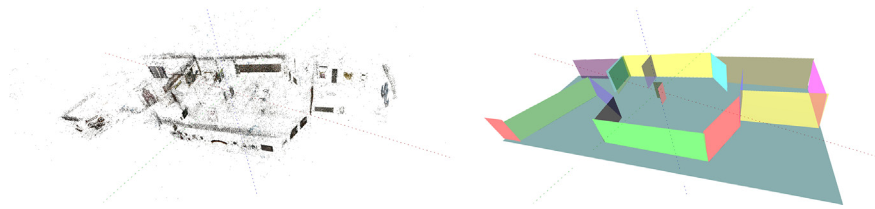

Fig. 1. Example of plane detection using our technique. Left: Museum dataset: point cloud consisting of 179,744 samples obtained from a set of photographs using SynthExport and Photosynth. Right: planes automatically detected by our technique in just 0.025 s on a 3.4 GHz PC. They were manually resized to better represent the original model.

Fig. 1 shows an example of planar regions detected using our technique. The point cloud shown on the left consists of 179,744 samples obtained from a set of photographs taken inside a museum. The samples were extracted using SynthExport [13] and Photosynth [14]. The image on the right shows the planes detected by our technique in just 0.025 s on a 3.4 GHz PC and illustrates the effectiveness of our approach.

The *contributions* of this paper include:

- An $O(n \log n)$ deterministic Hough-transform-based technique for detecting planar regions in unorganized point clouds (Section 3). Our solution is robust to noise and to sampling distributions. It is a few orders of magnitude faster and scales significantly better than existing approaches. A software implementation of our technique handles datasets with up to 10^5 points in real time on a typical PC.
- A fast Hough-transform voting strategy for planar-region detection (Section 3.3). Our solution uses a robust segmentation strategy to identify clusters of approximately coplanar samples. Votes are cast for clusters as opposed to for individual samples, greatly accelerating the detection process.

2. Related work

The most popular techniques to detect planes in point clouds are the *Hough transform*, *RANSAC*, and *region growing*. This section discusses these algorithms and their various optimizations intended to accelerate plane detection in point clouds.

2.1. Hough transform

The Hough transform (HT) [16,17] is a feature-detection technique. For any given input sample, it casts a vote for each instance of the feature one wants to detect that could possibly contain that sample. The votes are accumulated over all samples and the detected features correspond to the ones with most votes. The time and space complexity of the algorithm both depend on the discretization used for the accumulator, whose dimensionality varies with the number of parameters used to describe the features to be detected. For instance, plane detection requires a 3-D accumulator to represent the three parameters that characterize a plane.

The Hough transform was introduced by Paul Hough [16] for the detection of lines in images. Today, the universally used version of the HT is the *generalized Hough transform* (GHT) proposed by Duda and Hart [17], which replaced the slope-intercept with an angle-radius parameterization based on the normal equation of the line (1):

$$\rho = x\ \cos(\theta) + y\ \sin(\theta). \quad (1)$$

Here, x and y are the coordinates of a sample pixel, ρ is the distance from a line (passing through the pixel) to the origin of image's coordinate system, and θ is the angle between the normal of the line and the x-axis. This parameterization naturally extends to 3-D, supporting plane detection in the (θ, ϕ, ρ) Hough Space:

$$\rho = x\ \cos(\theta)\ \sin(\phi) + y\ \sin(\phi)\ \sin(\theta) + z\ \cos(\phi). \quad (2)$$

In (2), x, y and z are the Cartesian coordinates of the samples, $\theta \in [0^\circ, 360^\circ)$ and $\phi \in [0^\circ, 180^\circ]$ are the polar coordinates of the plane's normal vector, and $\rho \in \mathbb{R}_{\geq 0}$ is the distance from the plane to the origin of the coordinate system.

The *Standard Hough transform* (SHT) for plane detection uses (2) and iterates over each sample in the point cloud casting votes in the accumulator for all possible planes passing through that sample. More specifically, for given x, y and z coordinates, it iterates over all combinations of θ and ϕ, computing the value of the parameter ρ (2) and casting a vote at the corresponding accumulator cell (or bin). To make the computation feasible, one needs to discretize the θ and ϕ parameter values (defining angular steps). Thus, the computational cost of the SHT is $O(|P| N_\theta N_\phi)$, where $|P|$ is the number of points in the point cloud P, and N_θ and N_ϕ are the number of bins in the discretization of the θ and ϕ angles, respectively.

Given the high computational cost of the SHT, many techniques have been proposed to accelerate its voting procedure. *Common to most of these techniques is the focus on reducing the execution time by using a subset of the points in P, as opposed to designing new algorithms that effectively reduce the asymptotic cost of the voting process*. Next, we briefly review these strategies.

The *Probabilistic Hough transform* (PHT) [18] randomly selects m points ($m < |P|$) and uses them, instead of the entire point cloud, for voting. Since m is a percentage of $|P|$, the asymptotic cost is still $O(|P| N_\theta N_\phi)$. The PHT needs to find an optimal value for m to achieve good results. Small values tend to cause some relevant planes not to be detected, while large values do not result in significant reduction in execution time. As opposed to the SHT, the PHT is not deterministic.

Finding the optimal value for m is not a simple task, as it depends on many characteristics of the point cloud. To overcome this difficulty, the *Adaptive Probabilistic Hough transform* (APHT) [19] monitors the accumulator during the voting procedure. As stable structures emerge, they are stored in a list of potential maximum cells and only this list needs to be monitored. Since the process is adaptive, there is no need for an initial m value. The algorithm ends when the list of potential peaks becomes stable (*i.e.*, when the list does not change for a few iterations). The APHT is sensitive to noise, as the choice of the points is probabilistic and may lead to the detection of planes not present in the dataset. Its asymptotic cost is the same as SHT's.

The *Progressive Probabilistic Hough transform* (PPHT) [20] tries to avoid the influence of random noise by only detecting structures whose number of votes exceeds a threshold defined as a percentage of the total number of votes. Once a structure has been detected, the votes from all samples that support it are removed from the accumulator. Like the previous techniques, PPHT is non-deterministic and its asymptotic cost is the same as SHT's.

The Randomized Hough transform (RHT) [21] reduces the SHT's voting-processing time by exploiting the fact that a plane can be