Computer Vision and Image Understanding 000 (2016) 1–12

Contents lists available at ScienceDirect

Computer Vision and Image Understanding

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

Error-aware construction and rendering of multi-scan panoramas from massive point clouds

Marc Comino, Carlos Andújar*, Antonio Chica, Pere Brunet

VirVIG, Computer Science Department, Universitat Politecnica de Catalunya Jordi Girona 1-3, Barcelona, Spain

A R T I C L E I N F O

Article history:
Received 6 December 2015
Revised 14 September 2016
Accepted 19 September 2016
Available online xxx

Keywords:
3D reconstruction
Range data
Massive point clouds
Error-aware reconstruction
Compression
Panoramas
Interactive inspection

A B S T R A C T

Obtaining 3D realistic models of urban scenes from accurate range data is nowadays an important research topic, with applications in a variety of fields ranging from Cultural Heritage and digital 3D archiving to monitoring of public works. Processing massive point clouds acquired from laser scanners involves a number of challenges, from data management to noise removal, model compression and interactive visualization and inspection. In this paper, we present a new methodology for the reconstruction of 3D scenes from massive point clouds coming from range lidar sensors. Our proposal includes a panorama-based compact reconstruction where colors and normals are estimated robustly through an error-aware algorithm that takes into account the variance of expected errors in depth measurements. Our representation supports efficient, GPU-based visualization with advanced lighting effects. We discuss the proposed algorithms in a practical application on urban and historical preservation, described by a massive point cloud of 3.5 billion points. We show that we can achieve compression rates higher than 97% with good visual quality during interactive inspections.

1. Introduction

Reconstruction from accurate range data is becoming a more and more important topic. Processing of massive point clouds from lidar sensors involves a number of challenges, from data management to noise removal, model compression and interactive visualization.

When dealing with point clouds created by registering various range scans, a single surface can be represented by points coming from both nearby and distant sensors, with highly different expected measurement noise. Moreover, by merging multiple scans, the size of the resulting cloud rapidly grows, making it harder to handle in terms of resources.

In this paper we present a new methodology for the reconstruction of 3D scenes from massive point clouds coming from range lidar sensors. Our proposal is based on an error-aware, robust normal estimation and a panorama-based compact reconstruction and visualization.

Robust normal estimation is a key issue in most point-based tasks such as visualization and surface reconstruction. Existing methods are able to deal with noisy data under the assumption that every point has a measurement error with similar statistical properties. However, this assumption does not hold for point clouds coming from lidar scanners, because the expected measurement noise depends on the point's distance from sensor and its reflective properties. We propose and discuss a normal estimation method that is able to deal with this problem by considering, for each point, a 1D directional probability distribution with variance proportional to its associated measurement error.

Streaming and visualizing a massive point cloud is not a trivial task since it might not even fit in memory. Our solution to this problem is to build panoramas at different points of interest and enabling navigation between them. For this, we propose an efficient method to build high-quality panoramas from arbitrary view points by combining data from multiple 3D scans.

The main contributions of our approach are:

- An error-aware normal estimation algorithm for each point of the point cloud. We have developed a method which is able to deal with noisy point clouds with non-uniform and anisotropic scanning error.
- A panorama-based 3D representation for gigantic point clouds that encodes 3D data in a compact way, with an easy and inexpensive color matching algorithm.
- An efficient algorithm for interactive navigation, supporting both realistic and illustrative rendering techniques.

As a test case, we processed a collection of 3D scans of the Mercat de Sant Antoni including the archaeological remains that were

* Corresponding author.
E-mail address: andujar@cs.upc.edu (C. Andújar).

http://dx.doi.org/10.1016/j.cviu.2016.09.011
1077-3142/

found underground, resulting in a point cloud consisting of 3.5 billion points.

The rest of the paper is organized as follows. Section 2 reviews the most relevant previous work on the subject. Section 3 provides an overview of our approach, and Sections 4, 5 and 6 present the preprocessing phases of our method, including the estimation of normal vectors, generation of panoramas and estimation of consistent panorama colors. Section 7 is devoted to visualization and interactive rendering of the panoramas. Section 8 discusses our results with the test dataset. Concluding remarks are provided in Section 9.

2. Previous work

2.1. Normal estimation

Surface reconstruction from unstructured point clouds is an area widely studied as explained by Berger et al. (2014). In this work we only focus on the problem of normal estimation since we do not need an actual surface for rendering. Moreover, we assume that we can trivially orient these normals in a consistent way since we know the position of the scanner head from which each point was captured.

Hoppe et al. (1992) presented an algorithm for surface reconstruction from clouds of points which is based on estimating the signed geometric distance to the unknown surface and then applying a variant of marching cubes in order to get the reconstructed geometry. Particularly, they estimate normals by fitting the least squares tangent plane to the local neighbourhood of each point, which can be efficiently computed by principal component analysis (PCA). This approach is robust in the presence of noise but, as noted by Mitra et al. (2004), choosing the right neighbourhood size is key to obtain a smooth result whilst preserving the local curvature.

Giraudot et al. (2013) tackle the problem of having different levels of noise within the same point set by estimating a noise-adaptive robust distance function which is used to reconstruct the underlying surface. However, they treat all the points within a region equally without taking into account that, in a point cloud formed by a mixture of registered scans, some points might be more reliable than others. Moreover, although normals could be computed from the resulting meshes, they do not directly do normal estimation.

In the same sense, the following methods do surface reconstruction and require normals associated to each point as input. Nevertheless they introduce the idea of using extra per-point cues related to their reliability. Curless and Levoy (1996) propose to associate to each point a confidence value that depends on the ranging scanning technology. In their case, they associate lower confidence values to higher scanning grazing angles. Similarly, Fuhrmann and Goesele (2014) introduce a new method for surface reconstruction from oriented sample points with an associated scale cue and some optional confidence information. The scale of a point refers to the finite surface area the point represents. Surfaces are approximated by the zero set of an implicit function which is the result of the weighted sum of a set of basis functions parametrized by each sample point.

In addition to noise, range scanners might suffer from other defects, such as outliers, and there exist methods designed to be robust against them (Campos et al., 2013; Giraudot et al., 2013; Huang et al., 2009; Lipman et al., 2007; Nurunnabi et al., 2014). Nurunnabi et al. (2014) use the Minimum Covariance Determinant in order to compute a robust estimation of the covariance matrix for the neighbourhood of a point, on which they apply PCA to obtain the normal. Campos et al. (2013) iteratively fit d-dimensional splats (d-jets) to the local neighbourhood of each point in a RANSAC-like loop in order to minimize the impact of outliers. In fact, they do surface reconstruction although normals could be extracted from their estimated splats. However, both methods consider a general noise model which is not sensor-dependent.

Another family of methods aims at *consolidating* the point cloud prior to performing the surface reconstruction process , as a way to deal with noise and outliers. The process of consolidation comprises a projection operator and resamples the original point cloud producing a surface implicitly defined as the fixed point of this operator. In this direction, Alexa et al. (2003) introduce moving least squares (MLS) surfaces for point-based methods. Using a Gaussian weighting function that gives more importance to closer samples in the cloud, points are projected onto the local tangent space defined by a locally-fitted low-degree bivariate polynomial. Guennebaud and Gross (2007) extend this method by using spheres for shape approximation, achieving higher robustness. Nevertheless, these approaches are not considering specific directional assumptions on the noise and may require input normals for each point which are usually computed by using PCA.

Lipman et al. (2007) propose a method for point cloud consolidation requiring no normals. Their parameterization-free Local Projection Operator (LOP) tries to project a set of points onto the local multivariate median of the original cloud by ensuring that the resulting ones are distributed as uniformly as possible. Huang et al. (2009) extend this approach to deal with non-uniform sampled input clouds. Then, an initial guess for normals is computed using traditional weighted PCA and improved by using a corrector loop consisting in one consistent normal orientation step and one orientation-aware PCA step. One shortcoming for this approach is that data redundancy is reduced by uniformly placing samples regardless of the original sampling density, however some other works, such as Fuhrmann and Goesele (2014), with the help of some extra cue exploit this redundancy in order to capture finer detail.

The work of Shi et al. (2012) is also related to consolidation but, instead of defining a projection operator, they integrate multiple overlapping range images by performing a normal-aware clustering of the input samples and generate a simplified point cloud composed by cluster representatives computed using the mean-shift algorithm. Nevertheless, the result is limited by the quality of the normals resulting from the choice of the algorithm used to compute them.

One common flaw to all previous approaches is the assumption that the underlying surface is smooth. Other methods specialize in preserving sharp features, such as Boulch and Marlet (2012) who randomly draw triplets of points from a given neighbourhood, compute the plane going through them and use a Hough accumulator to estimate a robust normal. Still, this technique requires high point density near sharp features to produce reliable results and highly depends on a correct choice of parameters in order to be able to smooth out the noise. Another sharp feature-preserving algorithm is given by Castillo et al. (2013) who reformulated the least squares fitting problem and use non-linear least squares solvers in order to estimate, at the same time, the denoised points and corresponding normals. For this, they introduce new weights which penalize neighbours that lie far away from the estimated planes and favour an even distribution of the output cloud. Zhang et al. (2013) estimate a confidence for the PCA normals and those under a certain threshold are considered to potentially belong to a feature. This information is used to guide a low-rank subspace clustering that segments the neighbourhood of each feature point into planar subspaces. Finally, the normal for these points is computed using PCA on the samples belonging to the subspace with minimum fitting residual. These methods apparently achieve good results when estimating normals in sharp harp