Computers & Graphics 44 (2014) 20–32

Contents lists available at ScienceDirect

Computers & Graphics

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

Technical Section

Automatic room detection and reconstruction in cluttered indoor environments with complex room layouts ☆

Claudio Mura [a,*], Oliver Mattausch [a], Alberto Jaspe Villanueva [b], Enrico Gobbetti [b], Renato Pajarola [a]

[a] Visualization and MultiMedia Lab, University of Zürich, Switzerland
[b] Visual Computing Group, CRS4, Sardinia, Italy

ARTICLE INFO

Article history:
Received 9 May 2014
Received in revised form
6 July 2014
Accepted 8 July 2014
Available online 1 August 2014

Keywords:
Indoor scene reconstruction
LIDAR reconstruction
Point cloud processing

ABSTRACT

We present a robust approach for reconstructing the main architectural structure of complex indoor environments given a set of cluttered 3D input range scans. Our method uses an efficient occlusion-aware process to extract planar patches as candidate walls, separating them from clutter and coping with missing data, and automatically extracts the individual rooms that compose the environment by applying a diffusion process on the space partitioning induced by the candidate walls. This diffusion process, which has a natural interpretation in terms of heat propagation, makes our method robust to artifacts and other imperfections that occur in typical scanned data of interiors. For each room, our algorithm reconstructs an accurate polyhedral model by applying methods from robust statistics. We demonstrate the validity of our approach by evaluating it on both synthetic models and real-world 3D scans of indoor environments.

1. Introduction

In architecture and engineering, there is a substantial need for semantically rich 3D models of buildings. As 3D designs are most often not available, or significantly different from the “as-is” condition of a given building, technology for creating models from observations is of primary importance. 3D acquisition devices such as laser scanners are now available for fast, accurate and cost-effective acquisition of 3D data. However, efficient methods must be devised to extract higher-level models from the acquired raw point-cloud data.

Of particular interest is the problem of determining the architectural structure of indoor environments (e.g., room walls, floors and ceilings). Indoor reconstruction exhibits a number of distinctive challenges that make it significantly harder to manage than the more well-studied problem of building shape reconstruction from outdoor scans (see also Section 2). First of all, indoor reconstruction methods must be significantly more tolerant to missing data than their outdoor counterparts, since environments such as offices and apartments exhibit extremely high levels of clutter. This typically results in heavy occlusions of walls and other structures of interest (see also Fig. 1). Secondly, windows and other highly reflective surfaces are often present in such scenes. As a result, the acquired model is heavily affected by large-scale artifacts, measurement noise and missing data, due to the critical interaction properties of the reflective elements with the measurement devices (see also Fig. 1). Finally, creating structured 3D models of typical indoor environments, such as apartments and office buildings, poses the challenge of recognizing their interior structure in terms of a graph of connected rooms and corridors.

Much of the work on interior environments has focused so far on the analysis and classification of the objects in the scene [1,2], while the problem of recovering architectural components is less developed, and has concerned mostly floor plan reconstruction and wall boundary determination (see Section 2). Most current methods rely on the implicit assumption that the architectural components are well sampled. Even those approaches that include an explicit filtering stage in their pipeline are only able to tolerate small amounts of clutter and can fail in many situations that are commonly found in real world scenes. Moreover, many of the existing solutions are targeted at simply connected environments such as corridors and cannot reconstruct the shape of individual rooms within more complex environments.

In this paper, we present a robust pipeline for reconstructing a clean architectural model of an indoor environment from a set of cluttered 3D input scans that partially cover the scene of interest

☆This article was recommended for publication by Hongbo Fu.

* Corresponding author. Tel.: +41 44 6354365.
E-mail address: claudio@ifi.uzh.ch (C. Mura).

http://dx.doi.org/10.1016/j.cag.2014.07.005
0097-8493/

(typically 1 or 2 panoramic scans per room). Our method only assumes that the scanner positions are known, and the environment is composed of multiple rooms bound by vertical walls, which holds true for a vast majority of buildings, and is capable to recover a room graph, as well as an accurate polyhedral representation of each room.

The whole pipeline is depicted in Fig. 2. An occlusion-aware process extracts vertical planar patches as candidates for genuine wall segments, separating them from clutter and coping with missing data by using efficient *viewpoint-based* visibility computations on a per-scan basis. Starting from a space partitioning induced by the candidate walls, we use a robust heat diffusion process to propagate similarities between cells of the partitioning which belong to the same room. We then cluster the area into multiple rooms using an iterative binary subdivision algorithm. Unlike standard methods like *k*-means, our solution automatically finds the correct number of rooms without a termination threshold by exploiting the knowledge of the scanner positions.

This work is a significantly extended version of our *CADCG 2013* contribution [3]. Besides supplying a more thorough exposition, we provide here significant new material and a number of important novel contributions. Our main improvements are the following:

- a thorough description of the room detection process, which shows in more detail the properties of the computed diffusion embedding and provides a comprehensive analysis of the subdivision scheme;
- important methodological improvements, including a post-processing stage that corrects possible imperfections in the clustering results and a more effective robust technique based on *M-estimators* [4] for the reconstruction of the final wall planes;
- an extended evaluation, where we perform both a qualitative and quantitative analysis on a wider set of inputs, including two large real-world datasets and two new synthetic datasets that feature more complex room layouts and high variability in the shapes of the rooms.

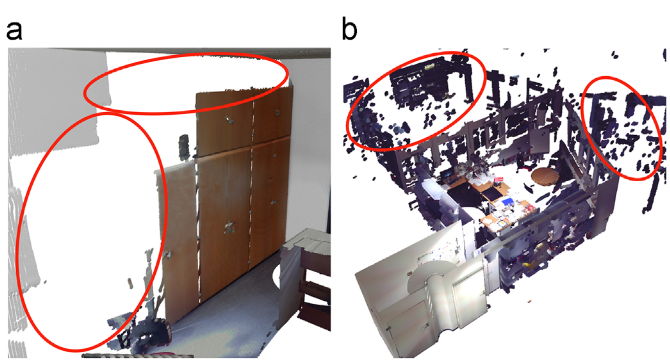

Fig. 1. Heavy occlusions (a) and large-scale artifacts (b) often occur in scanned 3D models of interior rooms.

The overall approach is the first indoor reconstruction pipeline capable of coping with heavy occlusions and missing data, while automatically recognizing different rooms as *separate* components. Such a room labeling is useful in many real-world applications, such as room asset planning and management or the definition of thermal zones for energy simulation. As demonstrated in Section 7, the method is applicable to large real-world environments with an extremely high level of clutter and is robust to scanning noise and large artifacts originating from reflecting surfaces.

2. Related work

Many researchers have studied the problem of reconstructing building structures from 3D laser range scan data. In this section, we briefly discuss only the approaches that most closely relate to ours.

Classical methods have often focused on creating visually realistic models [5,6], rather than structured 3D building models. Even though some of these 3D reconstruction algorithms extract planar patches from data [7], this has the goal of finding simplified representations of the models, rather than identifying walls, ceilings, and floors. In this context, clutter is dealt with specialized hole-filling techniques [6,8,9], which can only manage small-scale occlusions.

More recently, focus has shifted to the creation of more structured 3D models, with the purpose of simplifying the process of converting point cloud data into a building information model (the *scan-to-BIM* problem). In this sense, an important step towards the production of semantically rich models is the detection of the rooms in the input environment. Our previous work [3], which is extended and complemented by this paper, presents an occlusion-aware method that employs a diffusion process to automatically extract multiple rooms from interiors models. Recently, Turner and Zakhor [10] solved the same problem by first over-segmenting the 2D floor plan into portions of rooms and then merging adjacent segments to obtain the final partitioning. However, their approach does not perform an effective handling of occlusions. Ochmann et al. [11] also proposed a method for segmenting laser-scanned indoor models into different rooms, but their approach simply classifies the input points and does

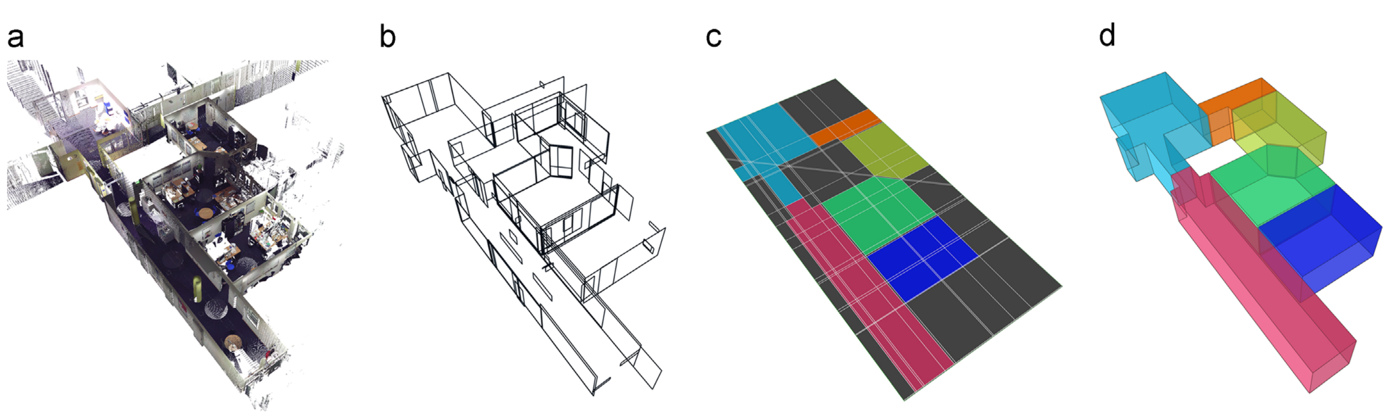

Fig. 2. The main phases of our algorithm: from the input model (a) we robustly extract candidate walls (b). These are used to construct a cell complex in the 2D floor plane. From this we obtain a partitioning into individual rooms (c) and finally the individual room polyhedra (d). Note that in (a) the ceiling has been removed for the sake of visual clarity.