Automation in Construction 88 (2018) 44–58

Contents lists available at ScienceDirect

Automation in Construction

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

6D DBSCAN-based segmentation of building point clouds for planar object classification

T. Czerniawski[a,*], B. Sankaran[a], M. Nahangi[b], C. Haas[b], F. Leite[a]

[a] University of Texas at Austin, Center for Construction Industry Studies, Ernest Cockell Jr. Hall 5.202, 301 E. Dean Keeton Stop C1752, Austin, TX 78712, United States
[b] University of Waterloo, Waterloo, Canada

A R T I C L E I N F O

Keywords:
Point cloud
Segmentation
DBSCAN
Classification
Machine learning
Plane detection
Building

A B S T R A C T

Due to constraints in manufacturing and construction, buildings and many of the manmade objects within them are often rectangular and composed of planar parts. Detection and analysis of planes is, therefore, central to processing point clouds captured in these spaces. This paper presents a study of the semantic information stored in the planar objects of noisy building point clouds. The dataset considered is the Scene Meshes Dataset with aNNotations (SceneNN), a collection of over 100 indoor scenes captured by consumer-grade depth cameras. All planar objects within the dataset are detected using a new point cloud segmentation method that applies Density Based Spatial Clustering of Applications with Noise (DBSCAN) in a six dimensional clustering space. With all planes isolated, an extensive list of features describing the planes is extracted and studied using feature selection. Then dimensionality reduction and unsupervised learning are used to explore the discriminative ability of the final feature set as well as emergent class groupings. Finally, we train a bagged decision tree classifier that achieves 71.2% accuracy in predicting the object class from which individual planes originate.

1. Introduction

Laser scanning has profoundly affected project surveying in the architectural, engineering and construction industries. These 3D imaging sensors capture existing structural and terrestrial conditions accurately, objectively, and with greater continuity than any manual metrology methods. Current applications of laser scanning by construction firms include schedule and progress tracking [27], creating as-built documents [1,16], path planning, clearance evaluation, and quality assurance [3,18,26]. Despite these benefits, widespread adoption by industry is slow because manually extracting information from raw 3D images and running analysis is painstaking, requires many person-hours, and specialized personnel training.

Plane detection from 3D point clouds aims to simplify and explain point cloud data by discovering shape primitives. It turns a large amount of raw point data into higher-level representations. Planarity and rectangularity are particularly fundamental to the built environment. Steadman [24] discussed the many practical reasons why most buildings are predominantly rectangular (Fig. 1). Floors are flat so pieces of furniture can stand easily on them. Walls and columns are vertical so they are structurally stable. Architectural plans layout rooms in rectangular patterns to avoid spatial interstices. Rectangularity pervades even the smaller components of buildings, including bricks, doors, windows, floorboards, and furniture. The cutting stock problem is famous within the furniture industry because much of the stock material comes in planar sheets, and so the final products are typically compositions of planar parts [21].

Algorithms designed to process planes in as-built data of buildings will then consequently cover the vast majority of the data. This paper presents a study of the semantic information stored in the planar objects of noisy building point clouds. More capably extracting information from noisy data, specifically, will improve the utility of low-cost depth cameras. Semantic information is important because it enables targeted point cloud pruning and targeted component isolation. For example, if a user is interested in detecting a pipe spool within an industrial laser scan [6], planar objects only represent clutter within the input. Hence, if planes could easily be isolated and semantically labeled, then the user could shrink their search space by deleting all planes except for perhaps the tables, on which the pipe spool might be located. Alternatively, the user might wish to isolate planes that represent the major architectural features of a space such as walls, floor, and ceiling to generate an as-built architectural model. Semantic labels would identify the planes in a laser scan that correspond with these components and aid in this process.

For this study, point clouds are sourced from SceneNN, a Scene Meshes Dataset with aNNotations [13]. SceneNN is a collection

* Corresponding author.
E-mail address: tczernia@utexas.edu (T. Czerniawski).

https://doi.org/10.1016/j.autcon.2017.12.029
Received 1 May 2017; Accepted 19 December 2017

Nomenclature

ρ	point cloud density
ρ_d	target density and density of downsampled point cloud
ρ_o	density of original point cloud
$D_{d(xyz)}$	point spacing in xyz space space for downsampled point cloud
$D_{o(xyz)}$	point spacing in xyz space space for original point cloud
Eps	Eps neighborhood (DBSCAN parameter)
N_d	sample size dictating number of points in downsampled point cloud
N_o	number of points in original point cloud
MaxC	Maximum Curvature accepted into DBSCAN clusters
minPts	minimum number of points in Eps neighborhood (DBSCAN parameter)
NN	nearest neighbor(s)
S	Normal Vector Scaling Factor

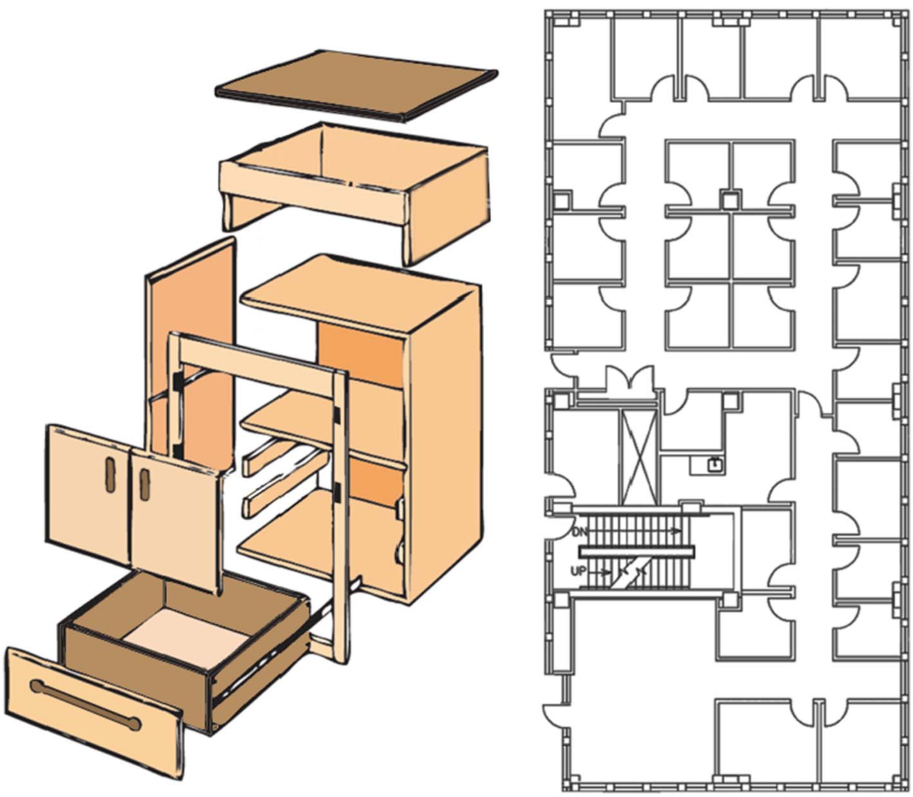

Fig. 1. Rectangularity and planarity in the built environment, furniture assembly diagram and architectural plans.

of > 100 indoor scenes, reconstructed from RGB-D videos captured by consumer-grade depth cameras. First, all planar objects are detected and extracted from each scene using a new point cloud segmentation method that applies Density Based Spatial Clustering of Applications with Noise (DBSCAN) in a six dimensional clustering space. Then an extensive list of features is formed and used to describe each planar object, and the information stored within these features is explored. Finally, we train a bagged decision tree classifier to categorize planar objects into twenty-four different classes and discuss the results.

The contributions of this paper are twofold: (1) first a new method for segmenting point clouds and detecting planes is presented and (2) the semantic information stored in planes is investigated alongside the feasibility of targeted point cloud plane pruning and targeted plane isolation in noisy building point clouds.

2. Background

2.1. Plane detection

The increasing availability of 3D imaging devices such as laser scanners, structured light cameras, and stereo cameras has enabled development of analysis techniques dependent on as-built information. The process of synthesizing, summarizing, or otherwise decluttering is an important part of extracting information from the raw geometrical data provided by these devices. Much like data compression in signal processing, this process seeks to eliminate redundancy by finding commonalities in data. If, for example, a collection of discrete points exhibit features that indicate they all fall on the same planar patch in a point cloud, it is desirable to agglomerate them into a single planar object. Plane detection methods generally, fall into one of three categories: primitive fitting, parameter space clustering, or region growing.

2.1.1. Primitive fitting

Fitting geometric primitives is a specific case of fitting mathematical models to data. Primitive fitting methods iteratively search for regions that demonstrate consistency with prior shape descriptions or models. The most prominent of these methods is Random Sample Consensus (RANSAC) [11], which is widely used for primitive shape detection. Basic RANSAC is comprised of two repeating steps: (1) minimal set selection and (2) minimal set evaluation. The minimal set for plane removal is three points or a single point along with its normal vector, as this provides a complete description for a plane. RANSAC randomly samples minimal sets from the scan data, fits a plane using their description, and counts the number of points in the scan that are consistent with the fitted plane. After a given number of trials, a plane is considered recognized at the locations defined by the minimal sets that achieved a score higher than a predefined threshold. Although basic RANSAC is conceptually simple, a direct application to plane