Advanced Engineering Informatics 38 (2018) 605–619

Contents lists available at ScienceDirect

Advanced Engineering Informatics

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

Full length article

Point cloud occlusion recovery with shallow feedforward neural networks

Luigi Barazzetti

Dept. of Architecture, Built Environment and Construction Engineering (ABC), Politecnico di Milano, Piazza Leonardo da Vinci 32, Milan, Italy

ARTICLE INFO

Keywords:
Approximation
Artificial intelligence
Machine learning
Neural network
Occlusion
Point cloud real occlusions

ABSTRACT

3D scenes reconstructed from point clouds, acquired by either laser scanning or photogrammetry, are subject to data voids generated by occluding objects. Modeling from incomplete data is usually a manual process in which human interpretation plays an essential role. This paper presents a machine learning algorithm based on neural networks capable of recovering point cloud occlusions for surfaces that can be approximated with injective functions. Starting from the point clouds acquired around the occlusion, a set of single-layer feedforward networks with a variable number of neurons is trained and validated with a subset of the original cloud, which is preliminarily decimated using local curvature to reduce CPU cost. The averaged result of the best neural networks is evaluated on a spatial domain that contains the 2D projection of the void, obtaining a complete 3D point cloud for the occluded volume. Criteria for choosing the number of neurons and the activation function for hidden and output layers are illustrated and discussed. Results are presented for both simulated and real occlusions, describing the pros and cons of the proposed method.

1. Introduction

In the latest two decades, laser scanning technology has become very popular in several disciplines such as civil engineering, archeology, architecture, cultural heritage documentation, mechanics, forensics, and geology, among others [16,35,29]. Nowadays, dense point clouds are used to generate accurate digital reconstructions in different formats such as CAD drawings (e.g., plans, cross sections, elevations), digital elevation models (DEMs), orthophotos, 3D models based on mesh or NURBS surfaces, or building information modeling (BIM) [3,4]. Recent work [1] has demonstrated that small “secondary” building components can be reconstructed from laser scans. On the other hand, users of laser scanning technology are aware that acquiring a large number of dense laser scans is not sufficient to generate complete and detailed deliverables. For instance, complex scenes with different object textures and materials provide variable quality of the measured return pulses. Specular reflective objects (e.g., mirrors) cannot be captured through point clouds, and transparent materials could result in refracted pulses.

Point clouds can also be generated using photogrammetric techniques [26,28]. Nowadays, the high level of automation of commercial software enables the creation of dense point clouds from large image datasets [21,36,32]. Although photogrammetric point clouds are similar to laser scanning point clouds, images require additional data processing that could involve a long processing time. Overall, complex and detailed surveys often require the combined use of photogrammetry and laser scanning [18,17,13].

This paper aims at considering one of the issues during point cloud generation (i.e., occlusions and the consequent lack of data). Laser scanning pulses cannot penetrate opaque objects, and 3D coordinates are measured only for the first reflective surface along the emission direction. Acquiring a single point cloud that reveals the whole scene is not possible in most practical applications. The acquisition of several scans from various station points (registered in a global reference system) allows users to fill in parts of the occlusions [33], notwithstanding the fact that complete scene acquisition (i.e., the whole object) is rarely achievable. Also, multiple scans of the same scene increase the time required for data acquisition and provide redundant data.

Photogrammetry does not solve the problems related to occlusions [38]. Although the acquisition of additional images can partially overcome the problems, real scenes can reveal occlusions only after image acquisition and processing, especially for large projects. The acquisition of specific images limited to the occlusions is not always feasible, especially if a new image acquisition campaign has to be carried out for a limited occluded part.

Modeling with incomplete point clouds is a common situation. Manual measurements carried out using the additional algorithms able to fill holes in the reconstruction are used in modeling projects aimed at creating reliable deliverables. Manual work and human interpretation are still mandatory to check the quality of the final result because automatic algorithms are prone to producing wrong reconstructions, as illustrated in the following sections.

E-mail address: luigi.barazzetti@polimi.it.

https://doi.org/10.1016/j.aei.2018.09.007
Received 13 March 2018; Received in revised form 20 September 2018; Accepted 26 September 2018
1474-0346/

2. Overview of the occlusion recovery problem

The technical literature reports several automatic and semiautomatic solutions for object occlusion detection and filling. Such methods were mainly developed to deal with mesh surfaces. They are based on various input data, such as neighborhood points, topology, change of local curvature, and distribution of hole texture [19].

Direct occlusion filling in point clouds (i.e., before mesh generation) is a challenging field of research due to the lack of a continuous surface, which means the absence of connectivity information in unorganized point datasets. As described by Chi and Bisheng [10], previous research work in point cloud occlusion recovery has been carried out by several authors, such as Becker et al. [6], Cai et al. [8], Centin and Signoroni [9], Doria and Radke [12], Friedman and Stamos [15], and Lozes et al. [27], among others. The prevailing idea was to recover (i.e., to fill) the occluded areas with a synthetic point cloud (i.e., a point cloud that completes the original laser scanning dataset). In this case, recovering means to provide a dense set of 3D points (according to the laser scanner used) for areas that appear shadowed in a point cloud due to occluding elements. Most methods are based on specific geometric constraints to provide realistic results. For instance, the approach proposed by Tagliasacchi et al. [34] assumes that the volume of an occluded part varies smoothly and that each point cloud sample is visible from outside the shape. The introduction of a specific energy term based on volumetric smoothness allows such an approach to deal with large holes in objects with deep concavities.

Approaches for hole filling are divided into (i) volume- and (ii) surface-based solutions. Volume-based approaches subdivide the space using an octree grid. Processing is carried out for those grids classified as belonging to the object, such as in the work of Franchini et al. [14], who developed a solution based on preliminary uniform sampling to work with large datasets. Surface-based methods are instead based on local properties (neighbor information) of the surface, such as curvature, topology, convergence, and texture. For instance, Doria and Radke [12] proposed an approach based on the depth gradient, whereas Wu and Chen [37] used both boundary extension and convergence. More recently, Nguyen et al. [30] developed a rapid strategy that determines the hole boundary and fills the hole through tangent planes for each boundary point. Then, the process is repeated from the hole boundary toward the center of the void. Altantsetseg et al. [2] added locally uniform points by generating contour lines inside the boundary edges of the hole. For an extensive review and comparison of mesh-based filling techniques, refer to Pérez et al. [31].

The commercial market offers several 3D modeling tools able to work with point clouds and their occlusions. Most commercial software allows the user to create a mesh from a set of point clouds, which is then inspected to check for the presence of holes. Then, filling algorithms can automatically close holes by adding new triangles to the existing mesh. Hole detection and filling algorithms are only some of the tools available in such software, which also have functions for point cloud preprocessing (denoising, filtering, subsampling, sparse point removal, etc.), alignment, registration, classification, meshing, texture generation, draping, dimensioning, segmentation, profile extraction, and so on. Some solutions for mesh processing and editing are also available as open source systems (e.g., MeshLab, OpenFlipper).

An example of occlusion detection and removal is shown in Fig. 1. The aim is to clarify the results achievable with such software in the case of occlusion filling. The vault was captured with multiple laser scans. However, complete data acquisition of the vault surface was not possible because a metal stiffener partially occluded the object.

The point cloud was imported in different software after the part with the metal stiffener was removed. Different commercial software programs were tested for the generation of a mesh surface using the available interpolation techniques and the automatic filling algorithms. Results are shown in Fig. 2 and reveal clear errors in the reconstruction achieved. In fact, automation based only on geometric aspects could be the source of gross errors. Although most software programs for point cloud editing have automated algorithms for hole identification and filling, they were mainly developed to deal with relatively small holes when compared to the whole object. In other words, the achieved result is undoubtedly wrong because the new mesh does not take into consideration the constructive logic of the vault.

It is therefore beneficial to compare the results achieved using commercial software with the traditional manual work carried out by an expert (human) operator. Manual (interactive) modeling was therefore carried out by a human operator, who had to:

- inspect the geometry of the vault;
- identify the rib and understand that is a continuous constructive element occluded by the stiffener;
- trace lines and surfaces while taking into consideration such constructive aspects, that is, the logic of construction of the vault (how the vault has been built).

Fig. 3 shows the results obtained through manual measurements and NURBS lines and surfaces using Rhinoceros. Here, the expert (human) operator was able to understand the geometry of the object and generate a surface without a stiffener.

The approach presented in this contribution is an alternative (automatic) solution that tries to replicate the work of a human operator. The implemented solution fills occluded areas through the generation of a synthetic point cloud for the occluded part. It is distinguished from previous approaches due to the use of machine learning techniques able to learn the geometry of the considered object from unstructured point clouds with inhomogeneous density. Because the method was developed through an automatic learning procedure, it can deal with various 3D scenes (i.e., it can be applied to various objects without prior knowledge of the considered geometry). In particular, the proposed solution tries to replicate the work of a human operator who is manually tracing (modeling) from point clouds.

Starting from a set of laser scans in the field, the input for digital reconstruction is a set of registered point clouds that captured the metal reinforcement installed on the vault. If the goal is the reconstruction of the entire vault, an expert (human) operator has to identify the occluding objects, delete the corresponding laser points, and complete the reconstruction using only the points measured on the vault.

The proposed approach starts from a simple consideration: The human operator was able to provide a consistent reconstruction of the object because particular attention was paid to the constructive logic of the vault. After a visual inspection, the solution was found using the geometrical features of the surface and occlusion, as well as the perception of a static occlusion on a continuous surface. The solution obtained using the proposed approach is shown in Fig. 4. As can be seen, the method involved understanding the geometry of the object thanks to machine learning procedures. In other words, can we develop and use an algorithm that can replicate human interpretation for this particular case study? Can such a method be extended to deal with other categories of objects in an automated way? The next sections describe and discuss the proposed algorithm.

3. Point cloud approximation with shallow neural networks

Let us suppose that the whole surface of the vault in Fig. 1 (without reinforcement) can be approximated by a function $f: D \rightarrow C \subseteq \mathbb{R}$, in which D is the domain and C the codomain. The traditional way to define f is to analyze the shape of the vault and implement a (static) computer program that generates the required mapping. This is not a trivial task, especially in the case of irregular vaults like those of historic buildings, which have geometric anomalies and do not fit basic geometric primitives.

We may define the domain as a 2D space $D = \{(x_{1D}, x_{2D}) \in \mathbb{R}^2\}$, notwithstanding the fact that the method can be extended to