

SHARP GEOMETRY RECONSTRUCTION OF BUILDING FACADES USING RANGE DATA

Eric Turner and Avidah Zakhor

University of California Berkeley
Department of Electrical Engineering and Computer Science
{elturner,avz}@eecs.berkeley.edu

ABSTRACT

In this paper we describe a method for detailed geometry reconstruction of building façades in an urban environment, given a 3D point-cloud of LiDAR range data. Our approach separates planar faces and interpolates their shape with Moving Least-Squares (MLS) sampling. A method is then proposed to reconstruct occluded areas of the building whereby gaps in the building surface are modeled with axis-aligned planes fit to the gap boundary vertices. This approach reconstructs unsampled areas of building surfaces under the assumption that buildings have 3D rectilinear, axis-aligned features. We demonstrate the effectiveness of our approach on a number of building façades.

Index Terms— Laser radar, Urban areas, Least squares approximation, Clustering methods, Signal processing algorithms

1. INTRODUCTION

There has been a growing interest in generating realistic 3D models for existing buildings. This technology would be applicable to industries such as GPS navigation devices, video games, and Computer-Generated Imagery [1]. The current method of collection requires an acquisition vehicle to drive down a street while taking Light Detection And Ranging (LiDAR) scans and panoramic photographs of the surround area.

This process results in an incomplete 3D point-cloud of exterior building façades. Buildings are only partially captured because these scans are typically taken from street-level. Gaps in the LiDAR point-cloud can be caused by obstructions from the building or surrounding clutter. Figure 1 shows an example point-cloud where window ledges and columns cause occlusion. Parts of the building that are transparent to LiDAR, such as window glass, also result in holes.

In many applications, it is desirable to generate a triangulated geometry for a collected point-cloud. This geometry needs to be accurate to the original architecture so that texturing and other processing can be applied. In this paper, we propose an approach for rectilinear surface reconstruction of building façades with incomplete 3D point-clouds. Specifically, our proposed algorithm generates geometry for any




Fig. 1. An example point-cloud of a façade. Note gaps near windows, ledges, and columns.




Fig. 2. Locations on the building A, B, C are occluded from the LiDAR collector on the vehicle.

holes in the point-cloud while preserving flat planes and sharp corners, which are common in modern architecture.

For buildings of an appreciable height, the street-level LiDAR scans are taken at a very acute angle with respect to the surface of the building. Any protrusions from the building cause shadows in the LiDAR and therefore the building is not collected in its entirety, as shown in Figure 2. For many applications, a typical solution would be to gather scans from multiple angles [2]. For the urban reconstruction problem, this tactic cannot be used since the scanner is restricted to street-level.

The LiDAR scans typically occur at uniformly spaced angles, so as the scans move higher up the building the resulting points are spread further apart. The final point-cloud becomes sparser, and with lower sampling density of these areas, any




Fig. 3. Demonstration of two-plane sharp hole-filling; (a) point-cloud; (b) MLS sampling; (c) planes used to fit hole; (d) sharp hole-filling.

error that exists in the position measurements has a greater impact on the resulting geometry.

Previous attempts to model buildings have assumed that architecture takes a simple polyhedron geometry. This approach ignores all minor details of a building façade, such as windows and ledges [3]. The goal of this paper is to preserve as much detail as possible in the reconstructed model geometry. Any interpolation must preserve the sharp edges of corners that occur within these holes. This requirement diverges from the typical assumptions of most surface completion methods, which usually result in smoothed surfaces [4]. While developing 3D models with sharp features is a well-explored topic, most current approaches require a sufficient density of points near the corner regions, which would not be applicable here [5]. By contrast, our goal is to generate sharp, axis-aligned, planar features in locations where the point-cloud is sparse or not sampled at all.

2. MLS INTERPOLATION

Our approach to geometry reconstruction considers only a single building façade at a time. While this simplification prevents one façade from being used when modeling another, it allows us to assume a relatively flat point-cloud while creating a triangulated geometry.

No assumption is made about the surface density of points along the building face or the amplitude of noise at each point position. Most triangulation techniques, such as the Ball-Pivoting Algorithm, require points to be uniformly spaced and have little error associated with their position information [6]. So rather than use the original points to form the triangulation, we choose to uniformly interpolate the input point-cloud before triangulation occurs. In our application, the sampling

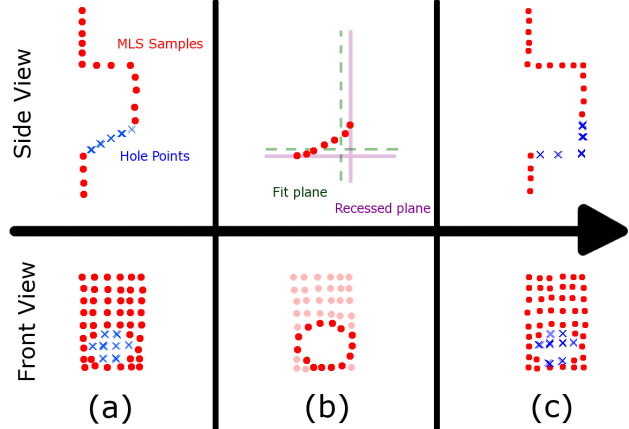


Fig. 4. Diagram of how a window ledge is sharpened; (a) MLS sampling points around hole; (b) the edge points used to fit axis-aligned planes; (c) sampling after snapping hole positions to axis-aligned planes.

distance for this interpolation is typically on the order of 1-10 cm, and is sampled in a grid parallel to the dominant plane of the façade. This dominant plane is found using Principle Components Analysis (PCA), which is reasonable given that the face is roughly planar.

We use Moving Least-Squares (MLS) to interpolate since it allows for a mitigation of noise in areas of high sampling, interpolation of surface in areas corresponding to gaps in the point-cloud, and overall uniformity of the resulting geometry [7].

During MLS sampling, each point in the original point-cloud is weighted with a decaying exponential with respect to lateral distance from the 2D sample location. If this distance is large compared to the decay of the exponential, then the interpolation asymptotically approaches a nearest-neighbor interpolation. Since this property yields undesirable discontinuities in the geometry, sample points within gaps in the point-cloud require post-processing to ensure that they maintain sharp features of the original architecture. We propose the following sharp hole-filling algorithm for this purpose.


3. SHARP HOLE FILLING

To determine whether an interpolation sample occurs in a gap area of the point-cloud, we use its distance from the nearest LiDAR point. This distance threshold is dependent on the sampling density of the LiDAR. For all examples shown in this paper, this threshold is set at 20 cm, while the grid sample size is 5 cm.

As discussed above, these gaps in the point-cloud are either caused by the transparency of windows or occlusions from protrusions. If the artificial geometry to be inserted needs to mimic realistic architecture, then it is most probable that the missing mesh is composed of a vertical and horizon-



(a)



(b)


Fig. 5. Window ledge before and after sharp reconstruction for façade from Figure 1; (a) MLS sampling; (b) with sharp hole-filling.

tal plane. An example result is shown in Figure 3. Note that the planes shown are recessed into the building. Since the geometry is missed due to being in the shadow of another part of the building, the actual architecture must be indented with respect to the building façade. Otherwise, it would have been captured by the LiDAR and there would be no gap.


Note that the planes fit are axis-aligned, rather than oriented based on local geometry. This simplification is reasonable, since the façade is assumed to be well-modeled by a flat surface. Since the overall façade is flat, then any missing components are either parallel or orthogonal to the full façade for most modern architecture.

The process to find and reshape these holes in the surface is shown in Figure 4. For each MLS interpolated sample, we check its distance to the closest point from input point-cloud. If this distance is greater than a specified threshold, we mark the sample as “hole”, as shown in Figure 4 (a).

Since MLS samples reside on grid, we use Union-Find to group connected hole positions, so that the entirety of each hole is considered at a time. [8]. For each connected “hole” region, we determine the sample points that compose the exterior border of the region. We then use K-Means clustering to fit horizontal and vertical plane to hole edges [9]. These planes are shown as dotted lines in the side-view of Figure 4 (b). Then we move the final plane locations by one standard deviation of edge sample heights into the building face, to simulate a recessed cavity in the building. Vertical planes are moved back, while horizontal planes representing tops of ledges are moved down, and horizontal planes representing



(a)



(b)

Fig. 6. Before and after sharp reconstruction on windows, viewed from below; (a) MLS sampling; (b) with sharp hole-filling.

bottoms of ledges are moved up. This movement is shown with solid lines in Figure 4 (b). We then recalculate the height of each MLS sample inside the connected hole via bilinear interpolation of the original point-cloud. The motivation for this step is to account for holes of large diameter. Since MLS devolves into nearest-neighbor interpolation, skipping this step would result in dramatic discontinuities in the sample heights. Finally, we snap these sample locations to closest fit plane of the hole, as shown in Figure 4 (c). This process finds and snaps each hole location to a sharp edge. As shown in Figures 3, 5, and 6, the process fits one or two planes as needed.

One limitation of our approach is the necessity of a full border of samples around connected hole areas. If a hole occurs on the edge of the the building façade, this technique results in sub-par performance. The most-likely cause of a hole on the edge of the façade is the sparsity of points at high elevation above street level. In this situation, it is common for a point-cloud not to represent the entirety of a building, but to be truncated at some height. Since it is desirable to be able to obtain a triangulated geometry that encompasses full buildings, façades must be extrapolated.

Our proposed approach to this problem is to set sample values to the mean depth of the point-cloud located below the sample point. This method results in flat strips of extrapolated façade. If the building curves outward or has other non-planar characteristics, these strips conform to the shape of the building, as shown in Figure 7.




Fig. 7. The curvature of the building is preserved when extrapolated; (a) building point-cloud; (b) MLS triangulation.

4. RESULTS

Once a geometry has been processed, a texture is applied using panoramic photographs collected from the same location as the LiDAR. Since these images are projected onto the geometry, any irregularities in the geometry result in dramatic disturbances in texturing. Examples of this texturing process are shown in Figure 8.

One current limitation is the possibility for highly acute triangles in the final geometry. Since vertices are snapped and repositioned, any triangles that are connected to them can be deformed and become very thin. Such geometry can cause texturing artifacts. Additionally, repositioning vertices by dramatic distances may cause singleton triangles to reorient out of the surface, which can also be a visually unappealing disturbance.

5. CONCLUSION AND FUTURE WORK

As shown, our method takes a scanned point-cloud that can contain gaps as well as high variance in point-cloud density, and generates a complete triangulated geometry whose vertices are uniformly spaced. The detail preserved in this geometry mimics the original point-cloud to arbitrary precision, and existing gaps in the original point-cloud are represented with sharp features created using assumptions about typical patterns in modern architecture. Remaining work exists in consolidating many building façades into a water-tight model of the building. One avenue would be to use existing techniques to discover all faces of a building, and then pipe that result into the process described above [3].

6. REFERENCES

[1] A. Zakhor and C. Frueh, “Automatic 3d modeling of cities with multimodal air and ground sensors,” *Mul-*




Fig. 8. Results of texturing geometry after sharp hole-filling; (a) sampling of flat façade; (b) with texture; (c) sampling of uneven façade; (d) with texture.

timodal Surveillance, Sensors, Algorithms and Systems, pp. 339–362, 2007.

- [2] P. Tang et al, “Automatic reconstruction of as-built building information from laser-scanned point clouds: A review of related techniques,” *Elsevier*, pp. 829–843, June 2010.
- [3] J. Chen and B. Chen, “Architectural modeling from sparsely scanned range data,” *Springer*, pp. 223–236, November 2007.
- [4] N. Kawai A. Zakhor T. Sato and N. Yokoya, “Surface completion of shape and texture based on energy minimization,” *ICIP*, September 2011.
- [5] A. Mhatre and P. Kumar, “Projective clustering and its application to surface reconstruction,” *SCG 2006*, June 2006.
- [6] F. Bernardini et al, “The ball-pivoting algorithm for surface reconstruction,” December 1999.
- [7] A. Nealen, “An as-short-as-possible introduction to the least squares, weighted least squares and moving least squares methods for scattered data approximation and interpolation,” May 2004.
- [8] J. Doyle and R. L. Rivest, “Linear expected time of a simple union-find algorithm,” vol. 5, pp. 146–148, November 1976.
- [9] S. P. Lloyd, “Least squares quantization in pcm,” *IEEE Transactions On Information Theory*, vol. IT-28, pp. 129–137, March 1982.