Modeling Residential Urban Areas from Dense Aerial LiDAR Point Clouds

Qian-Yi Zhou and Ulrich Neumann

University of Southern California

Abstract. We present an automatic system to reconstruct 3D urban models for residential areas from aerial LiDAR scans. The key difference between downtown area modeling and residential area modeling is that the latter usually contains rich vegetation. Thus, we propose a robust classification algorithm that effectively classifies LiDAR points into trees, buildings, and ground. The classification algorithm adopts an energy minimization scheme based on the 2.5D characteristic of building structures: buildings are composed of opaque skyward roof surfaces and vertical walls, making the interior of building structures invisible to laser scans; in contrast, trees do not possess such characteristic and thus point samples can exist underneath tree crowns. Once the point cloud is successfully classified, our system reconstructs buildings and trees respectively, resulting in a hybrid model representing the 3D urban reality of residential areas.

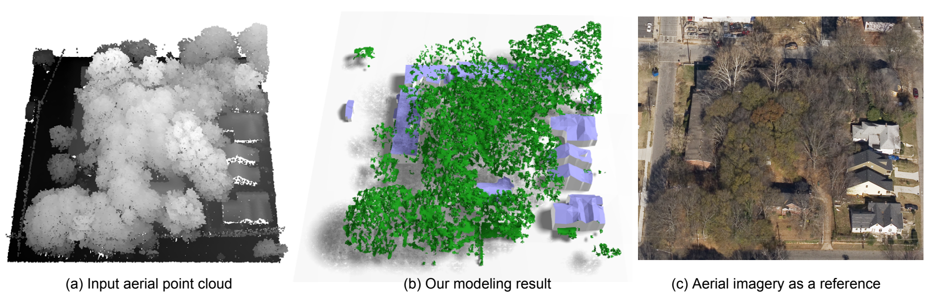

Fig. 1. Given (a) a dense aerial LiDAR scan of a residential area (point intensities represent heights), we reconstruct (b) 3D geometry for buildings and trees respectively. (c) Aerial imagery is shown as a reference.

1 Introduction

Urban modeling from aerial LiDAR scans has been an important topic in both computer graphics and computer vision. As researchers mainly focus on downtown areas containing various building structures such as skyscrapers, modern

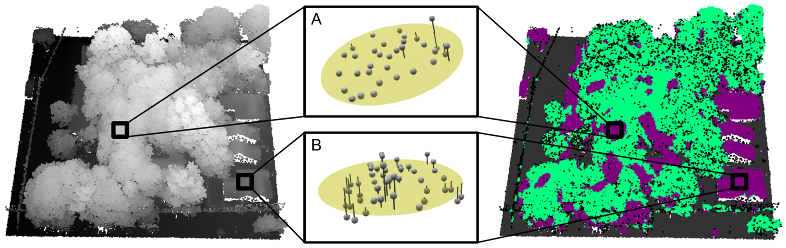

Fig. 2. Local geometry features become unreliable when dealing with residential areas with rich vegetation. In closeups of (A) a tree crown region and (B) a rooftop region, points are rendered as spheres while a locally fitted plane is rendered in yellow. Right: classification results of [18], trees in green, buildings in purple, and ground in dark grey.

office buildings, stadiums and convention centers; *building reconstruction* is believed to be the core of urban modeling, which has attracted much attention such as [4,5,8,10,15,18,19,20,21]. In these efforts, trees are usually considered as an interference to the urban modeling problem, and thus are detected and removed from the input by classification in pre-processing. Existing classification algorithms apply heuristics or machine learning approaches on point features including height, intensity, and local geometry information.

However, two new challenges emerge when the urban modeling problem extends to residential areas. First, as shown in Figure 1(a), vegetation is a major component of urban reality in residential areas. An urban modeling method for residential areas should detect and reconstruct both buildings and trees, *e.g.*, as we did in Figure 1(b). The second challenge lies in the classification method: dense LiDAR scans capture the detailed geometry of tree crowns, which may have similar height and local geometry features as rooftops of residential buildings. Figure 2 shows such an example where part of the tree crown shows similar or even better *planarity* than part of the rooftop (see closeups illustrating local points as spheres together with the optimal plane fitted to them). Classification algorithms based on local geometry features may fail and produce significant modeling errors. *E.g.*, Figure 2 right.

To address these two challenges, we present a robust classification method to classify input points into trees, buildings, and ground. Building models and trees are created from these points using a state-of-the-art building reconstruction algorithm [19] and a novel leaf-based tree modeling approach, respectively. The heart of our classification method is a simple, intuitive, but extremely effective measurement. In particular, we observe that residential buildings usually show a strong 2.5D characteristic, *i.e.*, they are composed of skywards roofs and vertical walls; both are opaque and thus prevent the laser beams from penetrating the building structure. Therefore, there is no point sample inside the building structure. The rooftops (or ground) become the lowest visible surface at a certain x-y position, as illustrated in Figure 3 left. In contrast, trees, composed of

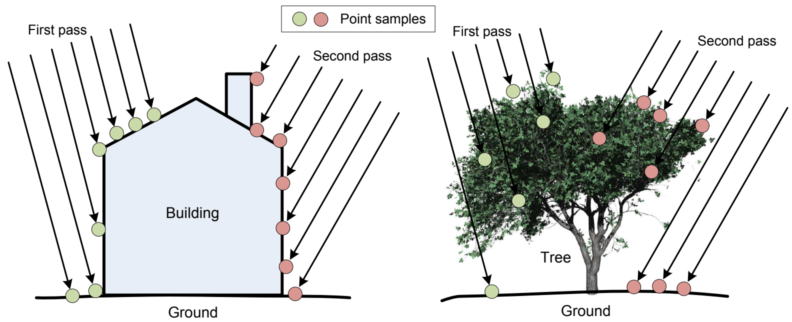

Fig. 3. While building structures have a 2.5D characteristic, trees do not possess such property. Dense laser scans may capture surface points under the tree crown (right).

branches and leaves, do not have this 2.5D structure. With multiple passes of scanning from different angles, the point cloud captures not only the top surface of the tree crown, but also surfaces inside and underneath the crown, as shown in Figure 3 right.

Contributions: To the best of our knowledge, we are the first to address the urban modeling problem for residential areas with rich vegetation from aerial LiDAR scans. We specifically list our novelties as follows:

1. We observe the key difference between building structures and trees from the perspective of the 2.5D characteristic. Based on this observation, we propose an effective algorithm to classify trees, building roofs, and ground.
2. We propose a complete system for urban reconstruction of residential areas. A hybrid model containing both 2.5D building models and leaf-based tree models is generated in an automatic and robust manner.

2 Related Work

Urban modeling from aerial LiDAR is an important topic that has drawn much attention in both computer graphics and computer vision communities. Recent research work [8,10,15,18] introduces an automatic urban modeling pipeline involving three key steps: *classification* detects and removes trees from the input point cloud; *segmentation* splits individual building roof patches out of the ground; and *building reconstruction* focuses on creating compact and accurate mesh models to represent the geometry of building structures.

Since downtown areas are usually the main target of reconstruction, modern urban modeling methods emphasize on building structures. For instance, Verma *et al.*[15] explore the roof topology graph connecting planar roof patches. Lafarge *et al.*[4] find the optimal configuration of 3D building primitives using a RJMCMC sampler. Matei *et al.*[8] and Poullis and You [10] create building models adapted to Manhattan-World grammars via different approaches. Zebedin *et al.*[17] generate both planar roof patches and surfaces of revolution. Toshev *et* *al.*[14] propose parse trees as a semantic representation of building structures. Lafarge and Mallet [5] combine primitives and a general mesh representation to achieve hybrid reconstruction. Zhou and Neumann develop both data-driven modeling approaches [19,20] and primitive-based method that supports global regularities [21].

In urban modeling systems, trees are often recognized as outliers and thus are classified and removed in the first step. Most of the classification algorithms rely on point-wise features including height [5,7,11,14] and its variation [2,7,11], intensity [7,11], and local geometry information such as *planarity* [5,15,18], *scatter* [5,14,18], and other local geometry features. Heuristics or machine learning algorithms are introduced as classifiers based on the defined feature set. To further identify individual building roof patches, segmentation is either introduced in a post-classification step, or combined with classification in the form of energy minimization such as [5].

Computer graphics and remote sensing communities have made great efforts in modeling trees from ground LiDAR and imagery, such as [3,6,9,12,13,16]. A general tree model is broadly adopted in these literatures, composed of skeletal branches and leaves attached to them. Inspired by these efforts, we propose leaf-based tree modeling from aerial LiDAR scans.

The **2.5D characteristic of building models** is first formally observed and defined in [19], as "building structures being composed of detailed roofs and vertical walls connecting roof layers". Many research efforts exploit this characteristic to help building reconstruction either implicitly [8,10,15] or explicitly [5,19,20,21]. Nevertheless, we are the first to introduce the 2.5D characteristic of building structures into the classification problem. We propose a simple, efficient and effective classification algorithm that gains great accuracy in residential areas with rich vegetation.

3 Point Cloud Classification

Given an aerial LiDAR point cloud of a residential area as input, the objective of classification is to classify points into three categories: trees, buildings, and ground. As mentioned in Section 1 and illustrated in Figure 3, the 2.5D characteristic is the key difference between trees and buildings (or ground). In order to formulate this concept, we discretize the point cloud by embedding it into a uniform 2D grid G. In each grid cell c, the point set $P(c)$ is segmented into multiple *layer fragments* $L(c)$, using local distance-based region growing. Ideally, a layer fragment $l_{building} \in L(c)$ lying on a 2.5D object (rooftop or ground) must have the lowest height among all layer fragments in $L(c)$, because the rooftop (or ground) is always the lowest visible surface to laser beams at a certain x-y position, as analyzed in Section 1. On the other hand, a tree layer fragment l_{tree} can exhibit any height. However, as there is usually a ground or rooftop surface underneath tree samples, l_{tree} is not expected to be the lowest layer fragment in $L(c)$. Therefore, we check all the layer fragments in each cell, assign only the lowest layer fragment as non-trees (rooftop or ground), and classify the rest layer fragments as trees. From an energy minimization perspective, this 2.5D characteristic criterion can be quantized with a data energy term $E_d(x_l)$ for each $l \in L(c)$ as:

$$E_d(x_l) = \begin{cases} \alpha & \text{if } x_l = \textit{building} \text{ or } \textit{ground}\text{, and } l \text{ is not the lowest in } L(c) \\ \beta & \text{if } x_l = \textit{tree}\text{, and } l \text{ is the lowest layer fragment in } L(c) \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

where x_l is the label of layer fragment l.

To further discriminate building and ground in the energy minimization framework, we introduce *elevation* of layer fragment $e(l)$ defined as the height difference between l and the ground elevation at c's center. Another data energy term $E_g(x_l)$ is defined accordingly:

$$E_g(x_l) = \begin{cases} \gamma \cdot max(1 - \frac{e(l)}{\sigma}, 0) & \text{if } x_l = \textit{building} \\ \gamma \cdot min(\frac{e(l)}{\sigma}, 1) & \text{if } x_l = \textit{ground} \\ 0 & \text{if } x_l = \textit{tree} \end{cases} \quad (2)$$

where σ is the normalization factor. Empirically, $\sigma = 6m$, as suggested in [5].

With a smooth energy $E_s(x_{l_1}, x_{l_2})$ defined over all neighboring layer fragment pairs (*i.e.*, layer fragments belonging to neighboring cells and satisfying certain distance criteria), we build a Markov Random Field which leads to an energy minimization problem over the labeling x of the entire layer fragment set $\mathcal{L}$:

$$E(x) = \sum_{l \in \mathcal{L}} (E_d(x_l) + E_g(x_l)) + \lambda \sum_{(l_1, l_2) \in \mathcal{N}} E_s(x_{l_1}, x_{l_2}) \quad (3)$$

where $\mathcal{N}$ is the set of neighboring layer fragment pairs, and smooth energy $E_s(x_{l_1}, x_{l_2})$ is defined as characteristic function $\mathbf{1}_{x_{l_1} \neq x_{l_2}}$.

With the energy minimization problem being solved using the well-known graph-cut method [1], point labels are determined as the label of the corresponding layer fragment. To further construct roof patches from building points, a region growing algorithm is applied based on certain distance criteria. While large building patches are adopted as rooftops, small patches are considered as outliers and removed henceforth.

4 Modeling of Urban Elements

Based on the successful classification of input points, we introduce different modeling approaches for trees, buildings, and ground respectively.

4.1 Tree Modeling

Modern tree modeling approaches adopt a general tree structure composed of skeletal branches and leaves attached to them. Tree reconstruction usually begins with a branch generation algorithm followed by a leaf modeling approach. However, unlike ground-based laser scans and imagery, aerial LiDAR data captures very few samples on branches, making branch generation a difficult task. Therefore, we choose to directly model tree leaves by fitting surface shapes around tree points having sufficient neighbors.

In particular, for each tree point p with sufficient neighbors, Principal Component Analysis is applied to its neighboring point set $N(p)$ to fit an ellipsoid. Eigenvectors $\mathbf{v}_0, \mathbf{v}_1, \mathbf{v}_2$ and eigenvalues $\lambda_0, \lambda_1, \lambda_2$ of the covariance matrix represent the axes directions and lengths of the ellipsoid respectively. We employ the inscribed octahedron of the ellipsoid to represent the local leaf shape around p. Specifically, an octahedron is created with six vertices located at $\{\mathbf{v}_p \pm s\lambda_0\mathbf{v}_0, \mathbf{v}_p \pm s\lambda_1\mathbf{v}_1, \mathbf{v}_p \pm s\lambda_2\mathbf{v}_2\}$, where $\mathbf{v}_p$ is the location of p and s is a user-given size parameter.

A uniform sampling over the tree point set P_{tree} can be applied to further reduce the scale of the reconstructed models.

4.2 Building Modeling

We adopt 2.5D dual contouring method [19] to create building models from rooftop patches through three steps: (1) sampling 2.5D Hermite data over a uniform 2D grid, (2) estimating a hyper-point in each grid cell, and (3) generating polygons.

The only challenge in applying 2.5D dual contouring to residential area data lies in rooftop holes caused by occlusion. To solve this problem, we add a hole-filling step right after 2.5D Hermite data is sampled from input points. In particular, we scan the entire 2D grid to detect rooftop holes, and solve a Laplace's equation $\nabla^2 z = 0$ to fill these holes, where z represents the heights of *surface Hermite samples* at grid corners. Existing surface Hermite samples serve as the boundary condition of the Laplace's equation.

4.3 Ground Modeling

Ground models can be easily created by rasterizing ground points into a DSM (digital surface model). Holes are filled via linear interpolation.

5 Experimental Results

Figure 4 shows our urban reconstruction results for a 520m-by-460m residential area in the city of Atlanta. The input contains 5.5M aerial LiDAR points with 22.9/m^2 resolution. Our algorithm reconstructs 56K triangles for building models, and 53K octahedrons as tree leaves, in less than two minutes on a consumer-level laptop. As illustrated in the closeups of Figure 4, our classification algorithm successfully classifies points into trees, ground, and individual building patches (second column). A hybrid urban model is generated by combining 2.5D polygonal building models and leaf-based tree models (third column). Aerial imagery is given in the last column as a reference.

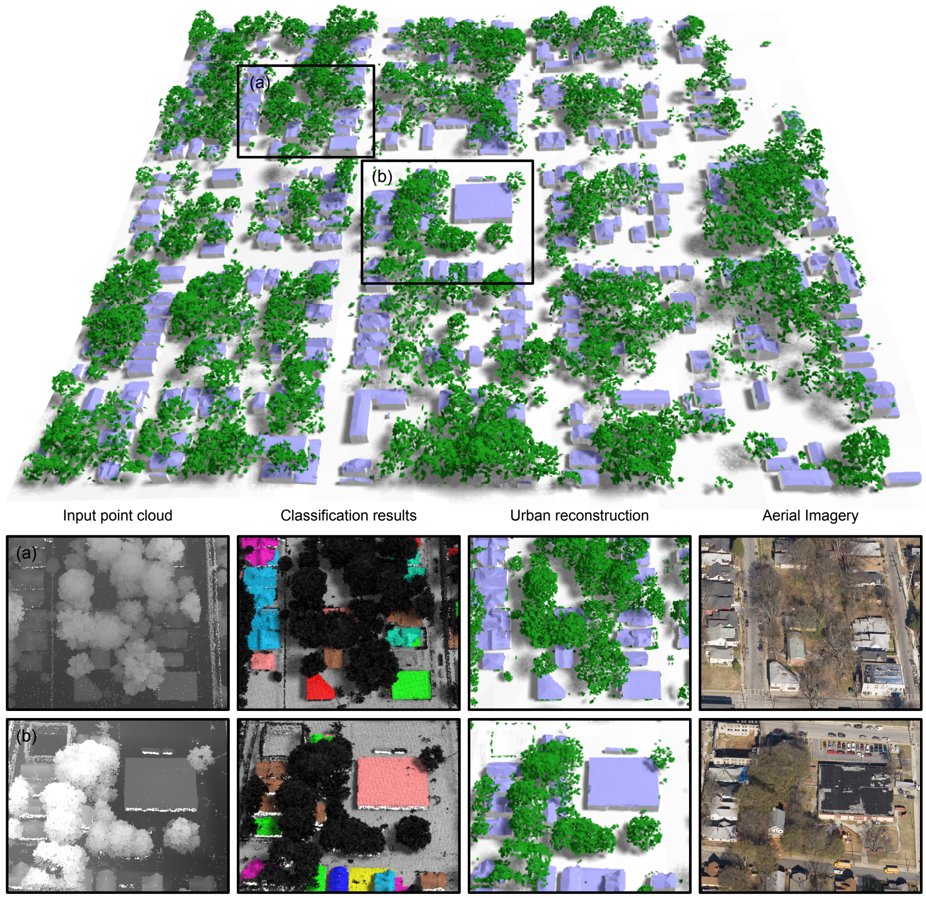

Fig. 4. Urban models reconstructed from 5.5M aerial LiDAR points for a residential area in the city of Atlanta

6 Conclusion

In this paper, we address the complicated problem of reconstructing urban models for residential areas with rich vegetation. We observe the key difference between buildings and trees in terms of the 2.5D characteristic: while buildings are composed of opaque skyward rooftops and vertical walls, trees allow point samples underneath the crown. This feature enables a powerful classification algorithm based on an energy minimization scheme. By combing classification, building modeling and tree modeling together, our system automatically reconstructs a hybrid model composed of buildings and trees from the aerial LiDAR scan of a residential area. Our experiments demonstrate the effectiveness and efficiency of our system.

References

1. Boykov, Y., Veksler, O., Zabih, R.: Fast approximate energy minimization via graph cuts. IEEE PAMI (2001) 5
2. Chen, G., Zakhor, A.: 2d tree detection in large urban landscapes using aerial lidar data. IEEE ICIP (2009) 4
3. Côté, J.F., Widlowski, J.L., Fournier, R.A., Verstraete, M.M.: The structural and radiative consistency of three-dimensional tree reconstructions from terrestrial lidar. Remote Sensing of Environment (2009) 4
4. Lafarge, F., Descombes, X., Zerubia, J., Pierrot-Deseilligny, M.: Building reconstruction from a single dem. In: CVPR (2008) 2, 3
5. Lafarge, F., Mallet, C.: Building large urban environments from unstructured point data. In: ICCV (2011) 2, 4, 5
6. Livny, Y., Pirk, S., Cheng, Z., Yan, F., Deussen, O., Cohen-Or, D., Chen, B.: Texture-lobes for tree modelling. In: ACM SIGGRAPH (2011) 4
7. Lodha, S.K., Fitzpatrick, D.M., Helmbold, D.P.: Aerial lidar data classification using adaboost. In: 3DIM (2007) 4
8. Matei, B., Sawhney, H., Samarasekera, S., Kim, J., Kumar, R.: Building segmentation for densely built urban regions using aerial lidar data. In: CVPR (2008) 2, 3, 4
9. Neubert, B., Franken, T., Deussen, O.: Approximate image-based tree-modeling using particle flows. In: ACM SIGGRAPH (2007) 4
10. Poullis, C., You, S.: Automatic reconstruction of cities from remote sensor data. In: CVPR (2009) 2, 3, 4
11. Secord, J., Zakhor, A.: Tree detection in urban regions using aerial lidar and image data. IEEE Geoscience and Remote Sensing Letters (2007) 4
12. Tan, P., Fang, T., Xiao, J., Zhao, P., Quan, L.: Single image tree modeling. In: ACM SIGGRAPH Asia (2008) 4
13. Tan, P., Zeng, G., Wang, J., Kang, S.B., Quan, L.: Image-based tree modeling. In: ACM SIGGRAPH (2007) 4
14. Toshev, A., Mordohai, P., Taskar, B.: Detecting and parsing architecture at city scale from range data. In: CVPR (2010) 3, 4
15. Verma, V., Kumar, R., Hsu, S.: 3d building detection and modeling from aerial lidar data. In: CVPR (2006) 2, 3, 4
16. Xu, H., Gossett, N., Chen, B.: Knowledge and heuristic-based modeling of laser-scanned trees. ACM Trans. Graph. (2007) 4
17. Zebedin, L., Bauer, J., Karner, K., Bischof, H.: Fusion of feature- and area-based information for urban buildings modeling from aerial imagery. In: ECCV (2008) 3
18. Zhou, Q.Y., Neumann, U.: A streaming framework for seamless building reconstruction from large-scale aerial lidar data. In: CVPR (2009) 2, 3, 4
19. Zhou, Q.Y., Neumann, U.: 2.5d dual contouring: A robust approach to creating building models from aerial lidar point clouds. In: ECCV (2010) 2, 4, 6
20. Zhou, Q.Y., Neumann, U.: 2.5d building modeling with topology control. In: CVPR (2011) 2, 4
21. Zhou, Q.Y., Neumann, U.: 2.5d building modeling by discovering global regularities. In: CVPR (2012) 2, 4