

2.5D Building Modeling by Discovering Global Regularities

Qian-Yi Zhou

University of Southern California

qianyizh@usc.edu

Ulrich Neumann

University of Southern California

uneumann@graphics.usc.edu

Abstract

We introduce global regularities in the 2.5D building modeling problem, to reflect the orientation and placement similarities between planar elements in building structures. Given a 2.5D point cloud scan, we present an automatic approach that simultaneously detects locally fitted plane primitives and global regularities. While global regularities are extracted by analyzing the plane primitives, they adjust the planes in return and effectively correct local fitting errors. We explore a broad variety of global regularities between 2.5D planar elements including both planer roof patches and planar facade patches. By aligning planar elements to global regularities, our method significantly improves the model quality in terms of both geometry and human judgement.

1. Introduction

Building modeling is a critical problem of 3D urban reconstruction which has attracted broad interests in computer vision community. The 2.5D characteristic of this problem has been exploited by a variety of research work either implicitly [7, 9] or explicitly [5, 14]. These 2.5D building modeling approaches benefit from the fact that building structures are usually composed of complicated roof surfaces and vertical facades; and reconstruct these two types of features individually while preserving the consistency between them. Aerial LiDAR point clouds are commonly adopted as a reliable source of geometric information for 2.5D building modeling.

A popular strategy in attacking the 2.5D building modeling problem is to introduce primitives (*e.g.*, planes, spheres, cones) to represent building shapes. In particular, planes receive the most attention since they are the commonest structures in man-made objects, especially in buildings. Planar roof patches are locally fitted from raw points, and are later combined with vertical facades aligning with roof boundaries, to construct a compact mesh model while maintaining low geometric fitting error rate. The main difficulty of this strategy is that local plane fitting can become unstable

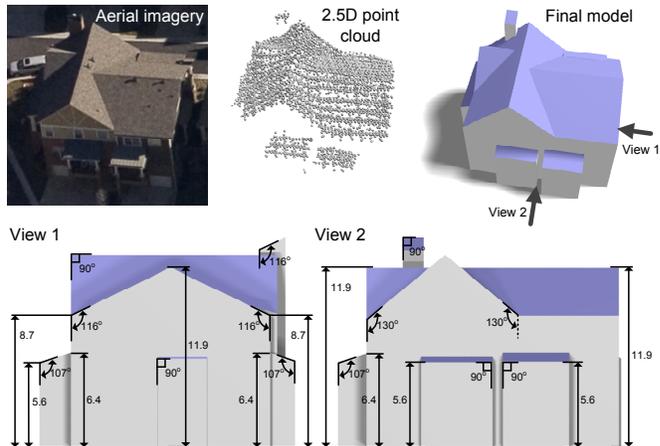


Figure 1. Our method automatically discovers global regularities from a noisy 2.5D point cloud, and use them to create a convincing 2.5D building model. Two orthogonal projections illustrate a subset of the global regularities in this model (lengths in meters).

when dealing with noisy or incomplete point clouds. Artifacts are inevitably created from unreliable plane primitives. To alleviate this problem, existing methods typically introduce strong urban priors to prune the fitted planes, such as roof topology [11] and Manhattan-world grammars [7, 9]. While prior knowledge successfully increases the robustness of these methods, it tends to be overstrict and thus limits their applicability when dealing with moderately complex building structures such as the one shown in Figure 1.

We propose *global regularities*, a moderate yet informative structure to organize planar roof patches and roof boundary segments. As illustrated in Figure 1 bottom, we explore both orientation and placement regularities between planar roof patch pairs (*e.g.*, slope angle equality), between roof boundary segment pairs (*e.g.*, segment height equality), and between a planar roof patch and its boundary segments (*e.g.*, orthogonality between their orientations). These global regularity patterns reveal the inter-element similarities and relations that intrinsically exist in building models because of human design and construction. With these patterns, the complexity of the building modeling problem can be significantly reduced for complicated build-

ing models such as the one in Figure 1.

We present an automatic algorithm to detect global regularities and utilize them to calibrate plane primitives. Unlike the strong priors introduced by previous methods, global regularities offer a more flexible representation of the global knowledge in 2.5D building models, and thus enable our algorithm to handle more complicated building shapes.

Another significant advantage of global regularities is that they characterize the intrinsic structures of building models, to which human vision is sensitive. For instance, Figure 2 right shows two models created from plane primitives. Without comparing model geometry with input points, human vision immediately finds the top-right model more convincing since it conforms to more global regularities.

Contributions:

1. We propose global regularities for 2.5D building models, involving orientation and placement similarities among 2.5D elements. These elements consist of both planar roof patches and roof boundary segments. We define roof-roof regularities, roof-boundary regularities, boundary-boundary regularities and integrate them into a unified framework.
2. We present an automatic algorithm to discover and enforce global regularities through a series of alignment steps, resulting in 2.5D building models with high quality in terms of both geometry and human judgement.

2. Related Work

We review the related work from two aspects: LiDAR-based building modeling, and shape from symmetry.

2.1. LiDAR-based Building Modeling

Many research efforts have attempted to address the challenging problem of urban reconstruction from aerial LiDAR point clouds. A common modeling pipeline proposed in recent research work [5, 7, 9, 11] includes three major steps: (1) *classification* detects and removes vegetation and noise points; (2) *segmentation* splits building point clusters from ground; and (3) *building modeling* generates mesh models from building point clusters.

The third step, *i.e.*, building modeling, is the most complicated step of the pipeline. Zhou and Neumann [14] identify the 2.5D characteristics of this problem and propose a data-driven method named 2.5D dual contouring, which is later extended to support 2.5D building topology in [15]. 2.5D models are optimized solely targeting at small quadratic fitting errors, and thus lose beauty and simplicity from a human vision perspective (Figure 2 bottom left).

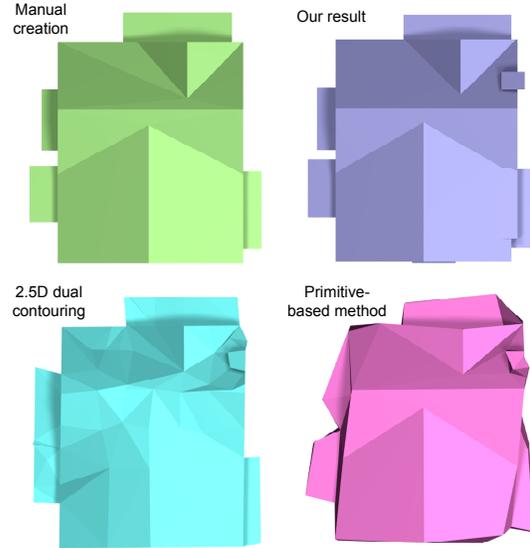


Figure 2. Modeling results generated from the same input point cloud by manual creation, our method, 2.5D dual contouring with principal direction snapping [14], and a primitive-based method [5]. Our method creates the most visually convincing result among all three automatic methods since our building model conforms to the most global regularities.

Another way to attack the building modeling problem is with primitive fitting approaches. In particular, local plane fitting is a popular strategy in extracting simple roof primitives [7, 9, 11, 13]. As discussed in Section 1, artifacts can be created due to misaligned plane primitives, and thus strong urban priors are frequently introduced to restrict the plane primitives. Typical priors include roof topology [11], Manhattan-world grammars [7, 9], and principal directions [13, 14]. Other research work aims at extending the ability of representing complicated shapes, by introducing additional primitive shapes and optimizing junctions between fitted primitives [4, 5, 12].

Nevertheless, we are the first to exploit global regularities and significantly improve the modeling quality in terms of both fitting accuracy and human vision judgement. A comparison between our method and two recent approaches is shown in Figure 2.

2.2. Shape from Symmetry

In both computer vision and computer graphics, symmetry has been identified as reliable global knowledge in 3D reconstruction. For instance, Fisher [2] introduces domain knowledge of standard shapes and relationships into reverse engineering problems. Thrun and Wegbreit [10] detect symmetries and utilize them to extend partial 3D models into occluded space. Gal *et al.* [3] adopt 1D wires to carry both local geometry information and global mutual relationships in man-made objects. Li *et al.* [6] extract re-

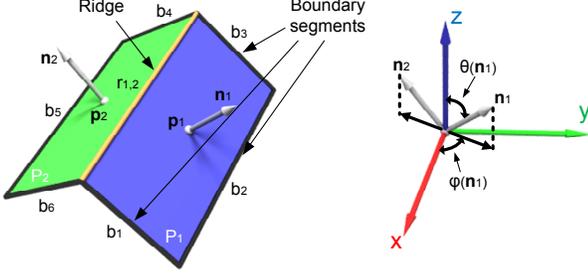


Figure 3. A typical gable-shaped building roof containing a set of 2.5D elements (e.g., plane primitive, boundary segments, and ridges)

relationship graphs among primitives to encode intra-part relations and use them to further improve the reconstruction quality.

These methods are similar in spirit to our method. While previous research work focuses on 3D man-made objects, we are the first to explore global regularities in 2.5D building models composed of roof patches and vertical walls.

3. Global Regularities

In 2.5D building models, global regularities characterize the inter-element similarities and relations arising from human design and construction. They are particularly useful in correcting plane primitives and creating more visually convincing building models. In this section, we explore various global regularity patterns that commonly exist in 2.5D building models and demonstrate them using a typical gable-shaped building roof shown in Figure 3.

Considering a 2.5D building model composed of plane primitives including planar roof patches and planar facade patches, it can be equivalently represented by a set of planar roof patches together with their boundary segments; because given the 2.5D constraints that roof surfaces are always bounded by vertical facades [14], planar facade patches and linear roof boundary segments have the same projection on the x-y plane. In particular, we denote the planar roof patch set as $\mathcal{P} = \{P_i : (\mathbf{v} - \mathbf{p}_i) \cdot \mathbf{n}_i = 0\}$ in which each plane P_i is determined by a normal-position pair $(\mathbf{n}_i, \mathbf{p}_i)$. A boundary segment set for each planar roof patch is collected by intersecting P_i with its surrounding planar facade patches, denoted as \mathcal{B}_i (e.g., in Figure 3, $\mathcal{B}_1 = \{b_1, b_2, b_3\}$). We explore global regularities among the 2.5D element set $\mathcal{P} \cup (\cup_i \mathcal{B}_i)$ from three aspects: roof-roof regularities, roof-boundary regularities, and boundary-boundary regularities.

3.1. Roof-Roof Regularities

We focus on two classes of commonly encountered regularities between roof plane pair (P_i, P_j) as follows.

3.1.1 Orientation regularities

In 3D models, the orientation regularities are usually defined as the orthogonality or parallelism between plane normals (e.g., [6]). This definition, however, cannot be directly applied to 2.5D building models for two reasons: first, roof plane normals rarely show orthogonality or parallelism; second, roof inclination and direction are of more interest in building modeling. In 2.5D models, orientation regularities are not determined by the angle between plane normals, but by the projections of normals on either the x-y plane or the z-axis. For instance, although \mathbf{n}_1 and \mathbf{n}_2 in Figure 3 do not exhibit orthogonality or parallelism, they show strong orientation regularities since their projections on the x-y plane are opposite. Therefore, we choose to write plane normals in spherical coordinates $(\theta(\mathbf{n}), \varphi(\mathbf{n}))$:

$$\theta(\mathbf{n}) = \arccos(n_z), \quad (1)$$

$$\varphi(\mathbf{n}) = \arctan(n_y, n_x), \quad (2)$$

where $\theta(\mathbf{n}) \in [0, \pi/2)$ and $\varphi(\mathbf{n}) \in [0, 2\pi)$ (Figure 3 right).

Intuitively, $\theta(\mathbf{n})$ determines the inclination of the planar roof patch, and $\varphi(\mathbf{n})$ indicates the direction of the slope. We are particularly interested in roof patches having the same inclination and roof patches exhibiting regularized slope directions (either orthogonal or parallel). Four typical roof-roof orientation regularities are defined accordingly:

- **θ -equality** when $\theta(\mathbf{n}_i) = \theta(\mathbf{n}_j)$,
- **φ -equality** when $\varphi(\mathbf{n}_i) = \varphi(\mathbf{n}_j)$,
- **φ -opposite** when $\varphi(\mathbf{n}_i) = \varphi(\mathbf{n}_j) \pm \pi$,
- **φ -orthogonality** when $\varphi(\mathbf{n}_i) = \varphi(\mathbf{n}_j) \pm \frac{\pi}{2}, \frac{3\pi}{2}$.

For example, plane pair (P_1, P_2) in Figure 3 exhibits the same inclination and the opposite slope direction. Using the above formulation, these characteristics are denoted as θ -equality and φ -opposite respectively.

3.1.2 Placement regularities

Placement of roof planes (i.e., roof positions) by themselves do not contain much regularity information. However, the placement of intersections between roof plane pairs may carry meaningful structural information about the building. In particular, we define *ridges* to reveal the regularities of roof plane placements.

Ridge definition: for a neighboring plane pair (P_i, P_j) satisfying both θ -equality and φ -opposite, the intersection of P_i and P_j is defined as a *ridge*, denoted as $r_{i,j}$.

The direction of ridge $r_{i,j}$ is uniquely determined as

$$\mathbf{d}(r_{i,j}) = (\sin(\varphi(\mathbf{n}_i)), -\cos(\varphi(\mathbf{n}_i)), 0)^T. \quad (3)$$

Since $\mathbf{d}(r_{i,j})$ is parallel to the x-y plane, the placement of $r_{i,j}$ can be parameterized by a pair of real numbers $(h(r_{i,j}), p(r_{i,j}))$, denoting the height of $r_{i,j}$ and the distance from origin to $r_{i,j}$'s projection on the x-y plane respectively. They can be calculated by solving an equation system with plane equations regarding $(\mathbf{n}_i, \mathbf{p}_i)$ and $(\mathbf{n}_j, \mathbf{p}_j)$. We define two types of placement regularities for ridge pair $(r_{i,j}, r_{k,l})$:

- **Ridge-height-equality** when $h(r_{i,j}) = h(r_{k,l})$,
- **Ridge-position-equality** when $\mathbf{d}(r_{i,j}) \parallel \mathbf{d}(r_{k,l})$ and $p(r_{i,j}) = p(r_{k,l})$.

3.2. Roof-Boundary Regularities

We observe that the majority of boundary segments are aligned either orthogonally (e.g., b_2 in Figure 3) or parallel (e.g., b_1, b_3) to the normals of their owner planes (e.g., P_1), when projected on the x-y plane. Therefore, we focus on roof-boundary regularities between plane P_i and its boundary segments \mathcal{B}_i . We denote the direction of P_i on the x-y plane by a 2D vector $\mathbf{o}_i = (\cos(\varphi(\mathbf{n}_i)), \sin(\varphi(\mathbf{n}_i)))^T$, and the direction of a boundary segment b_j 's x-y projection as $\mathbf{o}(b_j) = \mathbb{P}(\mathbf{d}(b_j))$, $b_j \in \mathcal{B}_i$, where \mathbb{P} is the projection operator and $\mathbf{d}(b_j)$ is the b_j 's direction in 3D space. We define:

- **o-parallelism** when $\mathbf{o}_i \parallel \mathbf{o}(b_j)$,
- **o-orthogonality** when $\mathbf{o}_i \perp \mathbf{o}(b_j)$,

3.3. Boundary-Boundary Regularities

As the orientation regularities among boundary segments can be implied from roof-boundary regularities and roof-roof orientation regularities, we focus on placement regularities between boundary segment pairs. In particular, we present two notable regularity patterns: first, boundary segments that are parallel to the x-y plane may have similar heights (e.g., (b_2, b_5) in Figure 3); second, when projected onto the x-y plane, boundary segments with the same direction may align to the same line (e.g., (b_1, b_6) and (b_3, b_4)). We define boundary-boundary regularities as follow, where $h(b_i)$ is the height of boundary segment b_i , and $p(b_i)$ is the distance from origin to b_i 's projection on the x-y plane.

- **Segment-height-equality** when $h(b_i) = h(b_j)$, and both $\mathbf{d}(b_i)$ and $\mathbf{d}(b_j)$ are parallel to the x-y plane,
- **Segment-position-equality** when $\mathbf{o}(b_i) \parallel \mathbf{o}(b_j)$ and $p(b_i) = p(b_j)$.

4. Modeling with Global Regularities

Given a noisy 2.5D point cloud as input, we present an automatic method to simultaneously detect locally fitted plane primitives and global regularities. In general, we adopt a discover-then-align strategy: once initial plane

primitives are identified, our algorithm discovers global regularities from them, and then immediately refines these initial primitives by aligning them to the global regularities. This optimization strategy is applied individually to each type of global regularities defined in Section 3. It effectively corrects the geometric errors raised by local fitting approaches, and thus significantly improves the model quality.

An overview of our approach is shown in Figure 4. Our system contains three main modules to create a 2.5D building model (bottom left) from a noisy aerial scan (top left):

1. **Planar roof patch extraction:** As shown in Figure 4 top, with plane primitives detected via local fitting, two discover-and-align steps are sequentially executed to detect the roof-roof regularities and refine the planar roof patch, namely, *orientation alignment* and *placement alignment*. Both planar roof patches and the roof-roof regularities are iteratively generated in a coarse-to-fine manner.
2. **Boundary segment production:** We immediately enforce the roof-boundary regularities by creating a rectangular bounding box for each planar roof patch, and identify boundary segments from bounding box edges, shown as the black lines in Figure 4 bottom. These boundary segments are further refined by discovering and enforcing boundary-boundary regularities.
3. **Model generation:** Vertical facades are automatically generated from boundary segments to connect roof patches and the ground. Rectangular roof patches are pruned by neighboring elements. A 2.5D building model is produced by combining both planar roof patches and vertical facades as shown in Figure 4 bottom left.

4.1. Planar Roof Patch Extraction

Given a set of points equipped with normals¹, we utilize a popular plane detection algorithm for aerial LiDAR scans [5, 11, 13] to find plane primitives: a region-growing procedure is applied to find spatially connected point clusters with similar normals; then plane primitives are locally fitted to individual point clusters. We denote the detected plane primitive set as $\mathcal{P} = \{P_i\}$ and apply orientation alignment and placement alignment sequentially.

4.1.1 Orientation alignment

By expressing plane normals in spherical coordinates, the orientation regularities can be categorized into two classes: θ -equality finds planes with similar slope angles, while φ -equality, φ -opposite and φ -orthogonality show regularized

¹Normals can be effectively estimated via covariance analysis [8].

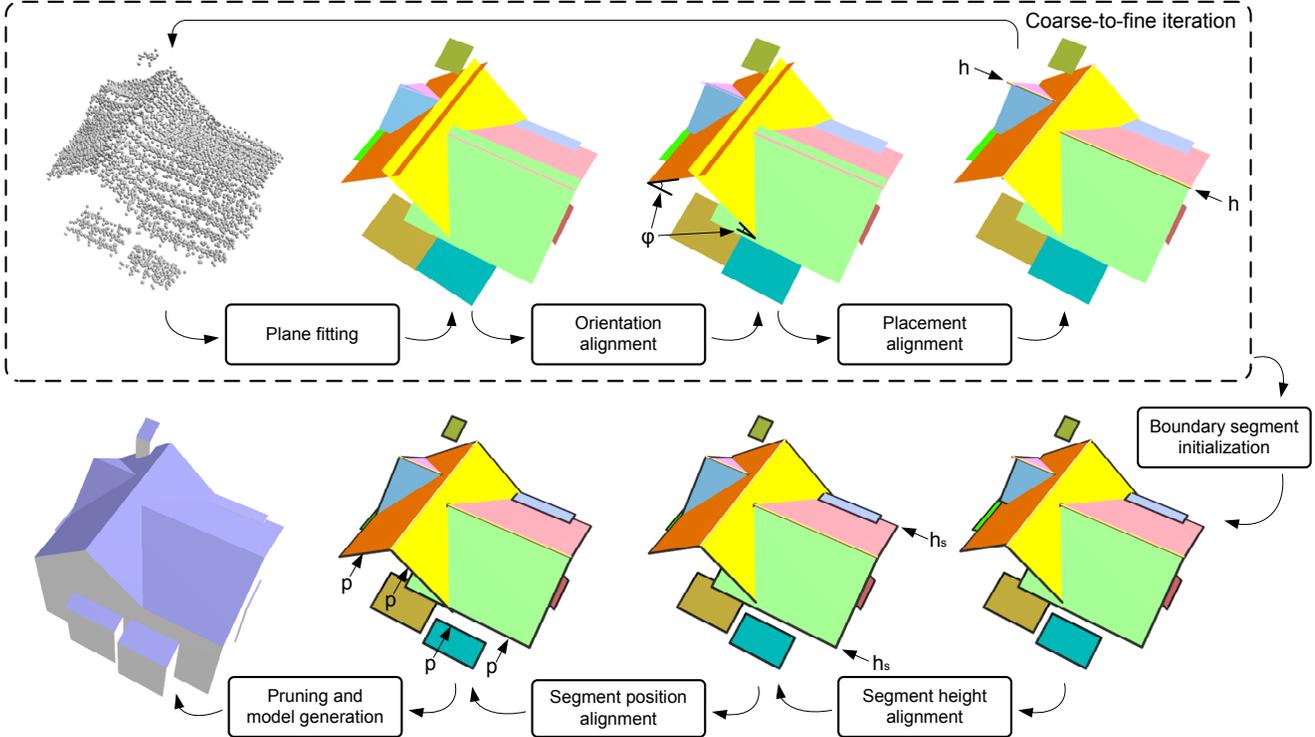


Figure 4. Pipeline of our approach: a 2.5D point cloud (top left) is transformed to a building model (bottom left) through a series of steps. Global regularities are discovered and enforced in each alignment step.

roof patch directions. These orientation regularities can be discovered by detecting clusters of $\Theta = \{\theta(\mathbf{n}_i)\}$ and clusters of $\Phi = \{\varphi(\mathbf{n}_i) \bmod (\pi/2)\}$ respectively. Each angle cluster implies a set of corresponding orientation regularities while the center of each cluster predicts the best alignment. In particular, we adopt complete-linkage clustering algorithm [1] to identify clusters in Θ and Φ . Cluster center sets \mathcal{C}_Θ and \mathcal{C}_Φ are taken as constraints in the subsequent alignment stage, in which $\theta(\mathbf{n}_i)$ and $\varphi(\mathbf{n}_i)$ are snapped to the corresponding cluster centers in \mathcal{C}_Θ and \mathcal{C}_Φ .²

4.1.2 Placement alignment

To effectively deal with placement regularities, we first detect ridges from neighboring plane pairs. Similar to orientation alignment, we decouple the placement alignment into two independent sub-problems: aligning ridge heights towards ridge-height-equality and aligning ridge positions towards ridge-position-equality. These placement regularities can be discovered by finding clusters of ridge height set $\mathcal{H} = \{h(r_{i,j})\}$; and clusters of ridge position set $\mathcal{S}(\mathbf{d}) = \{p(r_{i,j})|\mathbf{d}(r_{i,j}) \parallel \mathbf{d}\}$, regarding each ridge direction \mathbf{d} . Cluster center sets are denoted as $\mathcal{C}_\mathcal{H}$ and $\mathcal{C}_\mathcal{S}$ respectively, and used as regularity constraints henceforth. In the

²In singular cases where $\theta(\mathbf{n}_i) \approx 0$, $\varphi(\mathbf{n}_i)$ becomes unstable. Thus, we snap $\theta(\mathbf{n}_i)$ to 0 and assign the most popular $\varphi \in \mathcal{C}_\Phi$ to $\varphi(\mathbf{n}_i)$

alignment stage, ridge height $h(r_{i,j})$ and position $p(r_{i,j})$ are both aligned to their cluster centers, resulting in modifications on plane position \mathbf{p}_i and \mathbf{p}_j . In order to avoid conflicts between ridge height alignment and ridge position alignment, the former only affects the z values of position vectors, while the latter makes modifications to the x and y coordinates. Therefore, the only conflict source lies in planes that have multiple ridges, where the ridges compete in modifying the plane’s position. In this case, we allow only the longest ridge to modify the position, and ignore the effects from others.

4.1.3 Coarse-to-fine iteration

The planar roof patch extraction executes in a coarse-to-fine manner, as demonstrated in Figure 5. In particular, we fit planes to the input points, make orientation alignment and placement alignment to the plane primitives, discard points already associated with existing plane primitives, and then iterate through these steps with three modifications until no more plane primitives can be found by planing fitting:

1. We loosen the plane-fitting criterion to accept smaller plane patches. Specifically, normal variance allowance α is increased and the minimum number of points required to validate a plane primitive K is reduced.³

³Empirically, α remains $\pi/12$ while K is reduced by half in each iter-

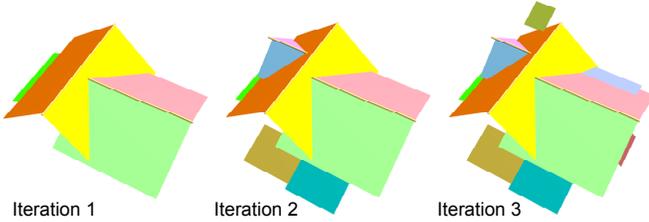


Figure 5. Coarse-to-fine planar roof patch extraction

2. In plane-fitting, instead of randomly picking region-growing seed, we start from points with normals that are close to normal $\mathbf{n}(\theta, \varphi), \forall(\theta, \varphi) \in \mathcal{C}_\Theta \times \mathcal{C}_\Phi$. These points have great potential to grow into planes conforming to existing orientation regularity constraints.
3. In alignment steps, the new plane primitives first attempt to snap to existing constraint sets $\{\mathcal{C}_\Theta, \mathcal{C}_\Phi, \mathcal{C}_H, \mathcal{C}_S\}$. Clustering is performed only for primitives that cannot align to the existing constraints given a distance threshold. New cluster centers are combined with existing centers to form regularity constraints for the next iteration.

These modifications are driven by two observations: first, large plane primitives are more reliable in producing global regularity constraints, and thus we begin with large primitives and require small primitives to be snapped to large primitives if possible (modification 3); second, the regularity constraints detected from large primitives can greatly improve the robustness of plane fitting for small patches (modification 2).

4.2. Boundary segment production

Given a set of plane primitives as shown in Figure 4 top right, we create initial boundary segments for each planar roof patch with help of roof-boundary regularities, and perform segment height alignment and segment position alignment based on boundary-boundary regularities.

4.2.1 Boundary segment initialization

As discussed in Section 3.2, most boundary segments in 2.5D building models conform to the roof-boundary regularities, *i.e.*, when projected on the x-y plane, boundary segments are either orthogonal or parallel to the normals of their owner planes. On the other hand, boundary segments represent the borders of roof patches, and thus bound the patch content, *i.e.*, points associated with the roof plane. Considering a planar roof patch P_i and the set of points associated with it denoted as \mathcal{V}_i ; we can compute a rectangular bounding box R_i of point projections $\mathbb{P}(\mathcal{V}_i)$ on the x-y plane, with R_i 's orientation following

ation, from 200 to 25.

$\mathbf{o}_i = (\cos(\varphi(\mathbf{n}_i)), \sin(\varphi(\mathbf{n}_i)))^T$. The segment set \mathcal{B}_i is initialized by back projecting R_i 's edges onto P_i .

Given that the plane normals are already aligned to a small set of orientation constraints \mathcal{C}_Φ , the x-y directions of boundary segments fall into a few 2D directions, which can be regarded as building-scale *principal directions* [13, 14]. In the special case where $|\mathcal{C}_\Phi| = 1$, the boundary segments are in two orthogonal x-y directions, forming *rectilinear contours* for building models [7].

4.2.2 Segment height alignment

We now apply segment height alignment to horizontal boundary segments. Segment-height-equality is discovered and enforced by a clustering method similar to previous alignment steps with two additional rules:

1. Heights of boundary segments from the same plane patch cannot belong to the same cluster, since snapping them together risks making the roof patch degenerate.
2. For each plane pair (P_i, P_j) that creates a ridge $r_{i,j}$, the boundary segments opposite to $r_{i,j}$ (*e.g.*, b_2 and b_5 in Figure 3) are tested with a relaxed criteria, because there is high probability of a reflection-symmetry.

We mark a boundary segment as “fixed” if its height is snapped to a cluster with more than one element. The position of each fixed segment is determined accordingly by substituting the modified height value into the plane equation. These positions act as placement constraints \mathcal{C}_p .

4.2.3 Segment position alignment

Segment position alignment is applied to the remaining boundary segments including both non-horizontal segments and horizontal segments that are not marked as “fixed” in the previous step. Positions of these segments first attempt to snap to elements in \mathcal{C}_p , and only join the position clustering procedure when the snapping attempt fails.

4.3. Model Generation

With planar roof patches and their boundary segments generated and aligned to the global regularities, we can easily reconstruct a 2.5D building model by combining roof patches and vertical facades, as shown in Figure 4 bottom left. The facades are produced from boundary segments connecting roof patches to the ground, while rectangular roof patches are pruned by neighboring elements including both roof patches and facades.

5. Experimental Results

We first compare our method with two existing approaches (*i.e.*, 2.5D dual contouring [14] and a primitive-based method [5]). Figure 2 shows a qualitative comparison

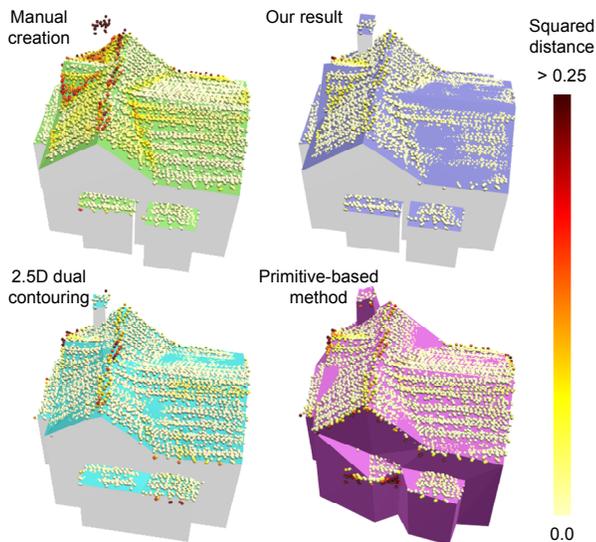


Figure 6. A comparison of geometric fitting errors (*i.e.*, squared distances from input points to the model surfaces) between models created by four different approaches.

Metric	Our method	2.5D DC [14]	Primitive-based method [5]	Manual creation
Triangle #	130	214	89	78
Ave. dis ²	0.012	0.016	0.018	0.058
Outlier ratio	0.9%	10.0%	11.1%	17.3%

Table 1. Quantitative results for the comparison in Figure 6

between these methods. Our model has the most similar visual appearance to manual creation, because it conforms to the most global regularities that characterize the intrinsic structure of building models. In contrast, 2.5D dual contouring only considers boundary direction similarities by introducing the *principal direction snapping* algorithm, while [5] detects primitives but does not deal with the relations between them. In addition, we quantitatively compare the three methods using the metrics shown in Figure 6 and Table 1. While our result shows comparable triangle number and average squared distance compared with previous approaches, we significantly reduce the outlier ratio (*i.e.*, the ratio of points with squared distances greater than 0.25m^2), because our approach can robustly fit small plane primitives through iterations with global regularities detected and updated progressively.

We further test our method on several LiDAR scans of buildings in the city of Atlanta, as illustrated in Figure 7. The input contains aerial LiDAR point cloud with 17 samples/m^2 resolution (first column). Planar roof patches and their boundary segments are detected and aligned with global regularities learnt from them (second column). 2.5D building models are reconstructed from these primitives by pruning the roof patches (see examples in the last three rows) and creating vertical facades from boundary seg-

ments (third column). Given the aerial imagery as reference (fourth column), our results are more “realistic” than the 2.5D dual contouring results (last column). Table 2 shows the statistics of the experiments in Figure 7. Computation time is measured on a laptop with Intel i-7 CPU 1.60GHz and 6GB memory. More experimental results can be found in the supplementary video.

6. Conclusion

In this paper, we define global regularities in 2.5D building models to characterize the intrinsic structure of building models. We present an automatic algorithm to discover and enforce global regularities through a series of alignment steps. Our system creates 2.5D building models with high quality in terms of both geometry and visual judgement.

7. Acknowledgement

We thank Airborne 1 Corp. for the point cloud data, Na Chen and anonymous reviewers for their valuable comments, and Florent Lafarge for providing the results of [5].

References

- [1] B. S. Everitt, S. Landau, and M. Leese. *Cluster Analysis*. Wiley, 4th edition, 2009. 5
- [2] R. B. Fisher. Applying knowledge to reverse engineering problems. *CAD*, 2004. 2
- [3] R. Gal, O. Sorkine, N. J. Mitra, and D. Cohen-Or. iwires: An analyze-and-edit approach to shape manipulation. In *ACM SIGGRAPH*, 2009. 2
- [4] F. Lafarge, X. Descombes, J. Zerubia, and M. Pierrot-Deseilligny. Building reconstruction from a single dem. In *CVPR*, 2008. 2
- [5] F. Lafarge and C. Mallet. Building large urban environments from unstructured point data. In *ICCV*, 2011. 1, 2, 4, 6, 7
- [6] Y. Li, X. Wu, Y. Chrysanthou, A. Sharf, D. Cohen-Or, and N. J. Mitra. Globfit: Consistently fitting primitives by discovering global relations. In *ACM SIGGRAPH*, 2011. 2, 3
- [7] B. Matei, H. Sawhney, S. Samarasekera, J. Kim, and R. Kumar. Building segmentation for densely built urban regions using aerial lidar data. In *CVPR*, 2008. 1, 2, 6
- [8] M. Pauly. Point primitives for interactive modeling and processing of 3d geometry. *PhD thesis, ETH Zurich*, 2003. 4
- [9] C. Poullis and S. You. Automatic reconstruction of cities from remote sensor data. In *CVPR*, 2009. 1, 2
- [10] S. Thrun and B. Wegbreit. Shape from symmetry. In *ICCV*, 2005. 2
- [11] V. Verma, R. Kumar, and S. Hsu. 3d building detection and modeling from aerial lidar data. In *CVPR*, 2006. 1, 2, 4
- [12] L. Zebedin, J. Bauer, K. Karner, and H. Bischof. Fusion of feature- and area-based information for urban buildings modeling from aerial imagery. In *ECCV*, 2008. 2
- [13] Q.-Y. Zhou and U. Neumann. Fast and extensible building modeling from airborne lidar data. In *ACM GIS*, 2008. 2, 4, 6

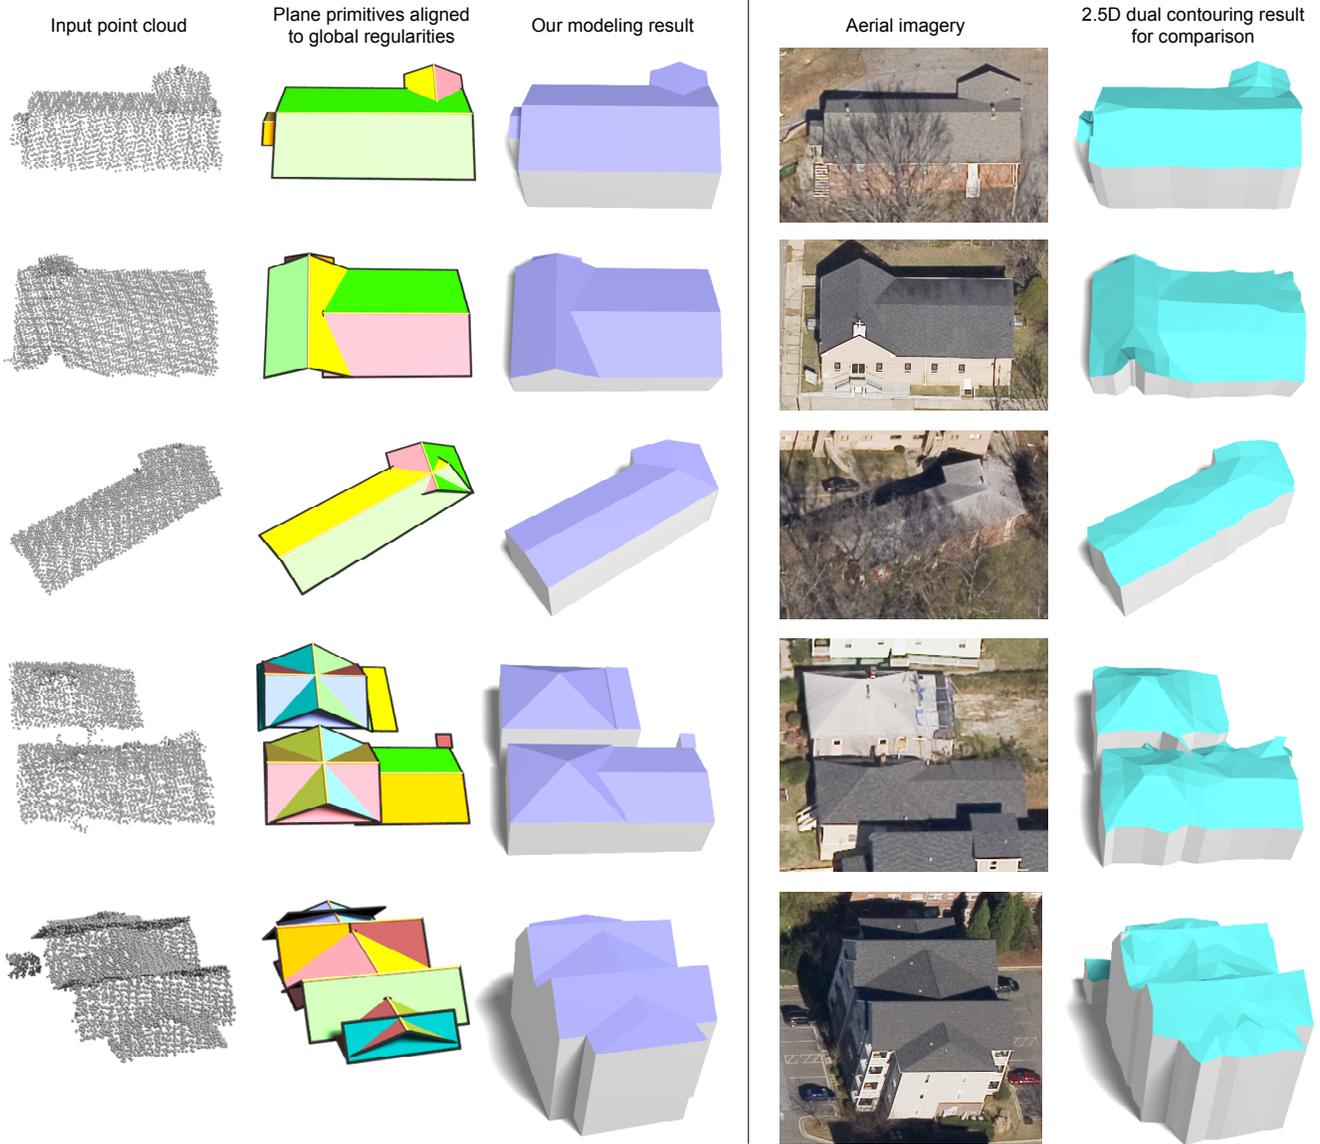


Figure 7. Experiments on several building scans. By discovering global regularities and enforcing them on the planar roof patches and their boundary segments (second column), we create visually convincing 2.5D building modelings (third column). Aerial imagery and 2.5D dual contouring results are included as reference.

Model	Point #	Plane #	$ C_{\Theta} $	$ C_{\Phi} $	Ridge #	$ C_{\mathcal{H}} $	$ C_{\mathcal{S}} $	Hor. segment #	Seg. h -cluster #	Non-hor. segment #	Seg. p -cluster #	Triangle #	time (s)
First row	3349	6	1	1	3	3	2	6	1	12	5	48	7.1
Second row	5603	6	2	1	3	3	2	6	2	12	6	48	6.8
Third row	4549	4	1	2	2	1	2	4	1	8	4	28	3.6
Fourth row	6143	13	3	1	5	3	3	18	7	24	13	110	16.6
Fifth row	6920	14	3	1	7	3	3	16	2	28	10	112	39.6

Table 2. Statistics of the experiments in Figure 7. Column 3 - 12 show statistics of the intermediate results. Since the size of constraint sets is considerably smaller than the number of primitives (in bold), the solution space (and thus the complexity) of the modeling problem is reduced significantly.

[14] Q.-Y. Zhou and U. Neumann. 2.5d dual contouring: A robust approach to creating building models from aerial lidar point clouds. In *ECCV*, 2010. 1, 2, 3, 6, 7

[15] Q.-Y. Zhou and U. Neumann. 2.5d building modeling with topology control. In *CVPR*, 2011. 2