3D Building Roof Reconstruction from Point Clouds via Generative Models

Hai Huang, Claus Brenner, Monika Sester
Institute of Cartography and Geoinformatics, Leibniz University Hannover
Appelstr. 9a, D-30167 Hannover, Germany
{Hai.Huang, Claus.Brenner, Monika.Sester}@ikg.uni-hannover.de

ABSTRACT

This paper presents a generative statistical approach to 3D building roof reconstruction from airborne laser scanning point clouds. In previous works bottom-up methods, e.g., points clustering, plane detection, and contour extraction, are widely used. Since the laser scanning data of urban scenes often contain extra structures and artefacts due to tree clutter, reflection from windows, water features, etc., bottom-up reconstructions may result in a number of incomplete or irregular roof parts.

We propose a new top-down statistical method for roof reconstruction, in which the bottom-up efforts mentioned above are no more required. Based on a predefined primitive library we conduct a generative modeling to construct the target roof that fit the data. Allowing overlapping, primitives are assembled and, if necessary, merged to present the entire roof. The selection of roof primitives, as well as the sampling of their parameters, is driven by the Reversible Jump Markov Chain Monte Carlo technique. Experiments are performed on both low-resolution (1m) and high-resolution (0.18m) data-sets. For high-resolution data we also show the possibility to reconstruct smaller roof features, such as chimneys and dormers. The results show robustness despite the clutter and flaws in the data points and plausibility in reconstruction.

Categories and Subject Descriptors

I.4.8 [**Computing Methodologies**]: Image Processing and Computer Vision—*Scene Analysis*[Object Recognition]

General Terms

Algorithms, Theory

Keywords

3D reconstruction, LIDAR, Point cloud, Building, Statistical modeling, RJMCMC, Model selection

ACM SIGSPATIAL GIS '11, November 1-4, 2011. Chicago, IL, USA

1. INTRODUCTION

Many approaches for the reconstruction of 3D city models from measurement data have been reported in the past decades. The introduction of laser scanning makes the acquisition of 3D data easier and more accurate. Overviews are given by Brenner [2], Schnabel et al. [9] and Vosselman [10].

Current works include Rottensteiner et al. [7], in which a roof plane delineation from LIDAR data is presented. Statistical tests and robust estimation are employed for stable edge detection against the clutter. Using manually generated geometric constraints, topological correction is ensured without additional 2D data. Sampath and Shan [8] segment and reconstruct more complicated buildings from airborne LIDAR point clouds based on polyhedral models. First, non-planar points are detected by means of the eigenanalysis making the roof planar segmentation more robust. The latter is implemented through a clustering with extended fuzzy k-means. An adjacency matrix is derived after the segmentation. For reconstruction, the roof vertices, ridges, and edges are determined by intersecting the corresponding planes, which include roof segments and possibly vertical walls or roof boundaries as the imposed constraints.

Lafarge et al. [5] present building reconstruction from a Digital Surface Model (DSM) combining generic and parametric methods. Buildings are considered as an assemblage of 3D parametric blocks. 2D-supports are firstly extracted manually or automatically (Ortner et al. [6]). 3D blocks are then assembled based on 2D-support and optimized within a Bayesian framework.

Airborne laser scanning data of urban scenes often has the following artefacts: (1) clutter by vegetation, e.g., trees, (2) reflection from windows and waterlogged depressions on the roof, and (3) HVAC (Heating, Ventilating and Air Conditioning) equipment, and generally lower point cloud density. The segmentation of relatively small roof structures and an accurate determination of roof edges are always hard. Results of the bottom-up reconstruction may thus be limited to a number of incomplete and irregular roof facets or building parts. A regularization during the extraction or afterwards is always needed and in many cases it is not easy to be conducted. For the issue of regularized plane detection a probability-driven edge sweeping method is proposed by Huang and Brenner [4]. Although it works robustly in spite of clutter and data flaws, it encounters difficulties when processing complex roofs.

In this paper we present a top-down statistical reconstruction of building roofs via generative models. Similar to [5],

a library of roof primitives is predefined and a building roof is considered as a variant of one primitive or a combination of a set of primitives. In this work, however, overlapping of primitives is allowed in the combination for a more flexible reconstruction. Unlike most of the related researches, the bottom-up effort, e.g., points clustering, plane detection ([8]), or 2D building contour extraction ([5]), are not used in the proposed work. Model selection mechanisms are integrated into the statistical sampling perceptively determining appropriate candidate models to compensate the absence of the initial information provided by bottom-up analysis.

The paper is organized as follows. In Section 2 the definition of roof primitives and the combination rules are described. The overall strategy of the statistical reconstruction of roofs is given in Section 3. Hypothesis models are generated by means of Markov Chain Monte Carlo (MCMC) sampling. An estimation of model size (ground area) and reversible jumps between different primitives are integrated into the search. As shown in Section 4, the optimal reconstruction is achieved by Maximum A Posteriori (MAP) estimation integrating prior information of the parameters into the goal function. After presenting results for data with varied resolutions in Section 5, the paper ends up with conclusions.

2. ROOF PRIMITIVES

Generative modeling works based on parametrized primitives. In this work we give flexible definitions of primitives and their interactions to gain more stable reconstruction results. The basic idea is to allow overlapping in the combination of multiple primitives. Basic rules are given to ensure plausibility.

2.1 Primitives vs. facets

In comparison with facet-based reconstruction, modeling with assembled roof primitives has the following advantages:

- There are no more irregular and incomplete roof facets or building blocks (forced to fit the bottom-up features like plane segment or 2D-support)
- A large scope of buildings can be represented by a limited number of primitives and their combinations.
- Complex roofs can possibly be interpreted more easily.

Figure 1 demonstrates two example roofs represented by the two schemes mentioned above. The target roofs can be seen as a group of facets, which are labeled with numbers. A derived Region Adjacency Graph (RAG) shows the organization of the facets (right) and could be used to guide the further reconstruction. Some roof parts, e.g., facet 10 in the first model and facet 3 in the second, however, are hard to be interpreted from the graph. In primitive-based modeling, on the other hand, the building roof is considered as an assembly of regular basic parts. The interpretation (left) is much simpler and clear. It is because not only the belonging facets have already been embedded in the primitives, but also the irregular facets caused by the plane intersection no longer need to be studied.

Another advantage of the primitive-based modeling is that the predefined constraints of member facets in the primitive ensure regularized reconstructions. Allowing overlapping, furthermore, makes every primitive to maintain complete during the reconstruction and assembly (cf. Figure 1, colorful blocks) instead of being cropped to fit the ground plan or neighbors. In the final roof model, the redundant parts are always hidden inside the assembly and they could be easily removed afterwards, e.g., by intersecting with CAD tools.

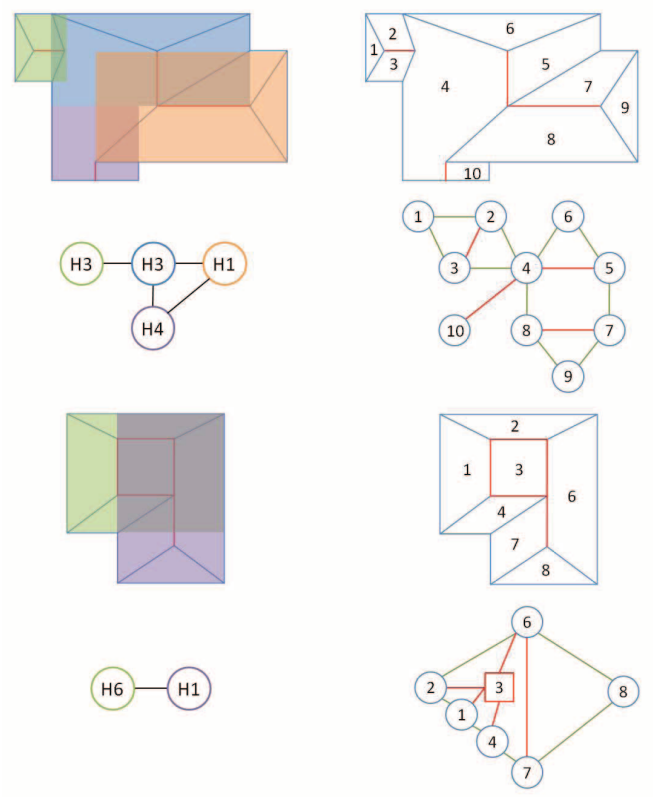

Figure 1: Roof configurations (top) can be simpler interpreted by the primitive-based modeling (bottom left) than the facet-based one (bottom right).

Beside the horizontal intersection we allow the vertical one as well. By this means some combined roofs, e.g., platform roofs (multi-level flat roofs, cf. Figure 12, buildings 1 and 8) and complex buildings (cf. Figure 12, buildings 3 and 5) are possible to be reconstructed without adding any more particular models in the library.

2.2 Library of primitives

The parameters of roof primitives are defined as:

$$\theta \in \Theta; \Theta = \{\mathcal{P}, \mathcal{C}, \mathcal{S}\}, \tag{1}$$

where Θ the parameter space: position parameters $\mathcal{P} = \{x, y, azimuth\}$, contour parameters $\mathcal{C} = \{length, width\}$ as all primitives are defined to have a rectangle footprint. $\mathcal{P}$ and $\mathcal{C}$ have fixed members. $\mathcal{S}$ contains shape parameters, e.g., ridge/eave height, depth of hip, deviation of ridge from the center, and is varied for different primitives.

As shown in Figure 2, we provide 3 groups including 11 types of primitives:

- Group F: generalized flat roofs including flat roof (F1) and shed roof (F2).

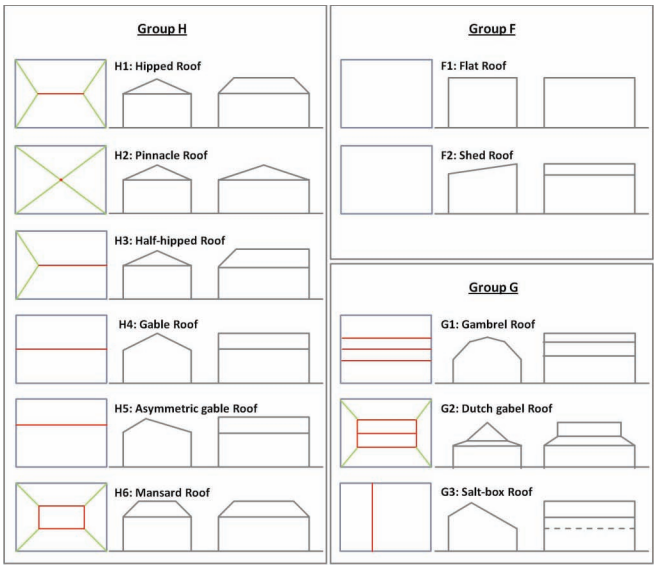

Figure 2: Library of roof primitives with top- (left), front- (middle), and side-view (right)

All possible shape parameters: $\mathcal{S}_{F,max} = \{z_1, z_2\}$ with z_1 the roof/eave height for flat roof and z_2 the additional different eave height for shed roof.

– Group H: generalized hipped roofs including all variants of hipped roofs. Gable roof (H4) and mansard roof (H6) are considered as special cases of hipped roofs.

$\mathcal{S}_{H,max} = \{z_1, z_2, hip_{l1}, hip_{l2}, hip_{d1}, hip_{d2}\}$ with z_1 the eave height, z_2 the ridge height, hip_{l1} and hip_{l2} the longitudinal hips, and hip_{d1} and hip_{d2} the lateral hips.

– Group G: Gambrel roof (G1) and others. The common point of this group is every roof has three different height levels. E.g., although the salt-box roof (G3) is very similar to the asymmetric gable roof (H5), it cannot be represented as an H-type because the parameter number of "heights" is limited to two in Group H. In this library we do not define elliptic roofs, which might be approximated by gambrel roofs.

$\mathcal{S}_{G,max} = \{z_1, z_2, z_3, hip_{l1}, hip_{l2}, hip_{d1}, hip_{d2}\}$.

All geometrical elements, i.e., vertices, edges, and facets, and their relationships are encapsulated in the primitives.

2.3 Primitive combination and merging

By combining primitives we propose a context-sensitive regularization with the following rules:

Rule 1: Intersection angle of adjacent primitives are conditionally regularized (cf. also below) to 0°or 90°, if they are close to each other.

Rule 2: Heights of flat roofs or ridge heights of hipped roofs are harmonized if they are close to each other.

Rule 3: Eave heights of all non-flat roofs are harmonized if they are close to each other.

The Rule 1 helps us to solve another general problem in flat roof extraction. As the roof points (especially in low density data) on the corners are very likely missing and that on the edges are not perfect either, slight deviation in azimuth (rotation in the roof plane) may lead to similar or even better evaluation and therefore cannot be detected. This is especially true for flat roofs with (1) less length to width ratio (square shape), (2) small size, and (3) low data density. This is hard to be solved without any prior information. Figure 3 (top) shows the reconstruction errors.

With the assumption that shed roofs and hipped roofs, which are composed by non-horizontal plane(s), are extracted normally with more reliably azimuth. We implement Rule 1 for flat roofs as follows:

R1.1: If a flat roof is adjacent to a shed/hipped roof, then align to the latter (cf. Figure 3, left).

R1.2: If two flat roofs are combined, the azimuths are jointly adjusted, weighted according to their areas (cf. Figure 3, right).

R1.3: A single flat roof can possibly be adjusted according to neighbor buildings in the multiple-building scene (cf. Section 3.1 Point D).

The derived parameter constraints are employed to guide the generation of new roof parts (primitives) and the post-processing of already found ones.

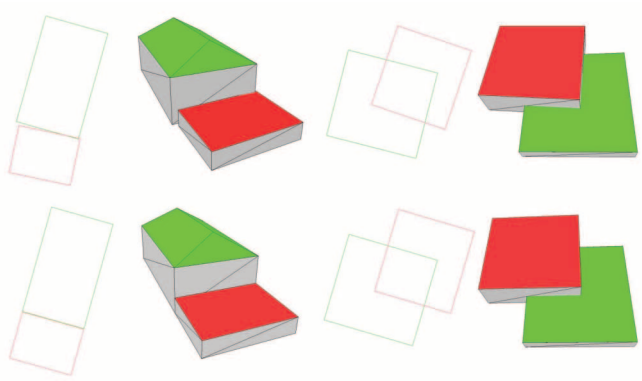

Figure 3: Azimuth deviation of flat roofs can be adjusted jointly with the adjacent roofs.

Although the primitives are regularized with the combination constraints, deviations still very often exist because of the random sampling. Most of them do not jeopardize the reconstruction results as they are hidden inside the intersected domain. However, for the primitives that share multiple planes, e.g., the L-form hipped roof, the deviation can be crucial (cf. Figure 5, top). A geometrical adjustment is needed to "merge" the primitives into a plausible model. We conduct the following rules for primitive merging:

– To choose arbitrary edge of each of the both primitives l_i and l_j with end points $(P_{0,i}, P_{1,i})$ and $(P_{0,j}, P_{1,j})$

– To compare their lengths: e.g., $|l_i| > |l_j|$

– If l_j near to l_i ($dist(P_{0,j}, l_i)$ and $dist(P_{1,j}, l_i) < \epsilon$) then it should be merged to the latter.

– By merging, as shown in Figure 4, if P_j is close to one of the end points of l_i then align to the point (left), if not then align to the line (right).

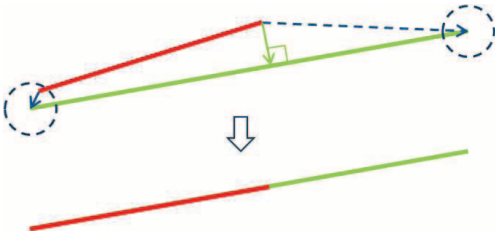

Figure 4: Primitive merging: the shorter edge (red) is merged to the adjacent longer edge (green). If their end points are close enough (inside dashed circles) then align to the end point, otherwise align to the edge.

Figure 5 (bottom) shows the result. Three pairs of edges are found and merged while the corresponding vertices are also aligned to each other.

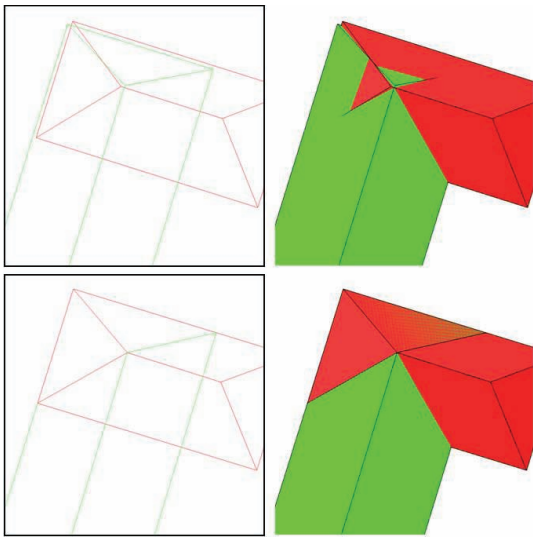

Figure 5: Primitive merging: geometrical conflicts by primitive combination (top) can be adjusted by merging corresponding edges (bottom).

3. HYPOTHESES GENERATION

The question to be answered in the top-down extraction is if the target roof can still be stably found without any initial information from bottom-up analysis. Besides the heavier search task, there are also two difficulties:

1 Local extrema: Enlarged search area means also many other objects, e.g., trees, are possibly involved and lead to many more local extrema. An improved search strategy is needed to tackle this (cf. Section 3.1).

2 Model-minimizing tendency: Using the average deviation to evaluate the reconstruction result, the finally found "best" model tends to shrink to a very small piece as smaller comparison area (corresponding to fewer data points) means less error and there is nothing to constrain the minimum model size ($A = length \times width$). In this case the size of the model should be reasonably estimated (cf. Section 3.2).

We conduct MCMC sampling with Sequential Monte Carlo (SMC) scheme for a stable and efficient search. Information criterion is employed to give instructive values for A, which guide the search of parameter-set $\mathcal{C}$. The choice of different primitives is simulated by reversible jumps integrated in MCMC.

3.1 Search strategy

We summarize the strategy as follows:

A Finding the first primitive (mostly the largest one or that near the start position) in the scene

- A1 Rough search with relaxed MCMC sampling and one simple primitive: gable roof (cf. Figure 6, a). Gable roof is chosen as initial primitive instead of flat roof because it is more sensitive for azimuth and it can actually also represent flat roof by harmonizing ridge and eave heights.
- A2 To refine the model parameters (cf. Figure 6, b) and possibly switch to more sophisticated primitives with reversible jumps in MCMC (cf. Figure 6, c and d).
- A3 After this primitive is accepted, we do not as usual delete the corresponding points from the source data as these points may be shared with other parts of this roof. This is more meaningful for low density data. Instead, we record its parameters and update the prior distributions to avoid this combination of parameters being sampled twice, i.e., the same primitive will not be proposed again in the next rounds of search.

B Finishing the whole building

- B1 Iteratively do step A finding all possible primitives near to (supposed to compose the target roof with) the found one(s) (cf. Figure 7)
- B2 Stop criterion: the maximum acceptable error is no more satisfied.
- B3 Jointly adjustment of primitives with combination rules

C Reconstructing superstructures if the point cloud is dense enough

- C1 Search in the roof area focusing the points above roofs
- C2 Simple flat and gable roof models with modified priors are used as primitives for chimneys and dormers.

D Possibly finding further buildings in the scene by repeating steps A to C with the points of the found building

removed. The parameters of the found roofs, e.g., azimuths and heights, can be (optionally) used as references for the search of their potential neighbors. The generic parameter priors can thereby be refined for this scene.

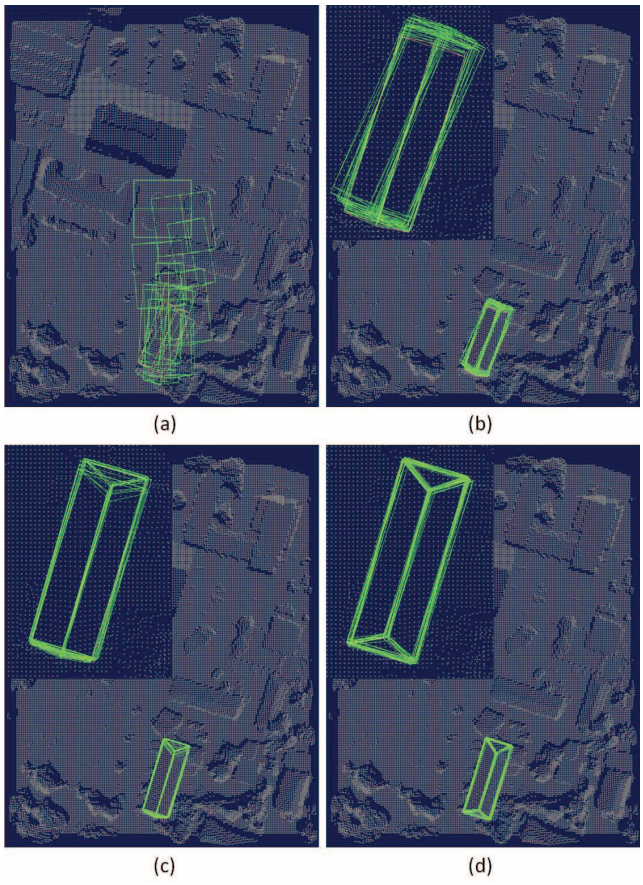

Figure 6: Search of the roof primitive: (a) rough search with simple gable roof, (b) local search refining the model, (c) jump to the half-hipped roof, and (d) jump to the hipped roof.

Figure 6 shows traces of the iterative search described in Step A. The rough search (a) started simply from the center of the scene (no initial position information) and employs generic priors for the other parameters. As the parameter sampling is constrained by the underlying primitive model, the disturbance from other non-building objects nearby is reduced. In relatively simple scenes MCMC is powerful enough to travel to the target. A refined search is then conducted locally (b) with reversible jumps (c, d) to finished the reconstruction.

3.2 Model size estimation

To overcome the model-minimizing tendency mentioned above, an instructive constraint for the model size A (corresponding to parameter-set $\mathcal{C}$) is needed. We conduct a perceptive estimation employing information criterion to balance the goodness of fit and the size of the model.

Figure 8 (top) shows the average deviation (blue) from the proposed model to the data points while the model size

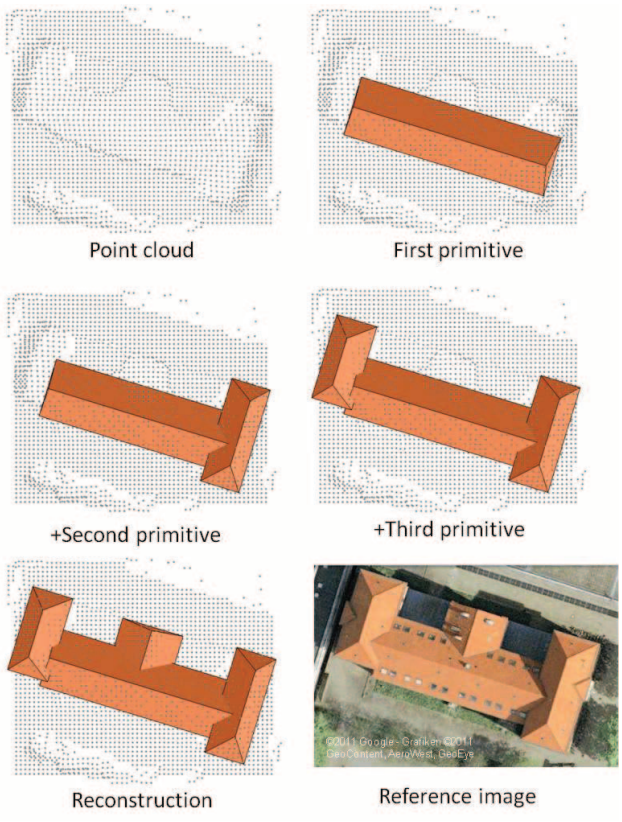

Figure 7: Reconstruction of an "E"-form building

increases. Please note that this function is neither linear nor monotonic increasing because of the influences coming from the target building and other objects in the scene.

Since what we need is just an approximate value (actually a lower limit) to guide the search, we simply follow the basic idea of Akaike Information Criterion (AIC, [1])

$$AIC = 2k - 2ln(L) \tag{2}$$

to build our goal function. In Equation 2, k indicates the number of parameters, which implies the complexity of the model, and L the maximum likelihood while the employed parameters have been optimized. In our case, on the other hand, what we want to prevent is not the model being too complicated but the size of the model being too small. We use the assessment of model for the L (cf. Section 4) and $-K^{\alpha}$ to represent the influence of the model size. K is the number of data points involved to the proposed model implying the size of the model (linear proportion in raster data). We employ the actual number of points instead of the model size because the former is more sensitive as the likelihood is also calculated with these points. α is the influence factor, which is empirically determined $\alpha = 0.1$. The total information entropy of the proposed model ($\mathcal{M}$) is then be updated as:

$$H_{\mathcal{M}} = -K^{\alpha} - 2ln\left(L(\mathcal{D}|\mathcal{M})\right). \tag{3}$$

By these means better fit is rewarded while size decrease gaining trivial improvement is discouraged. The conducted

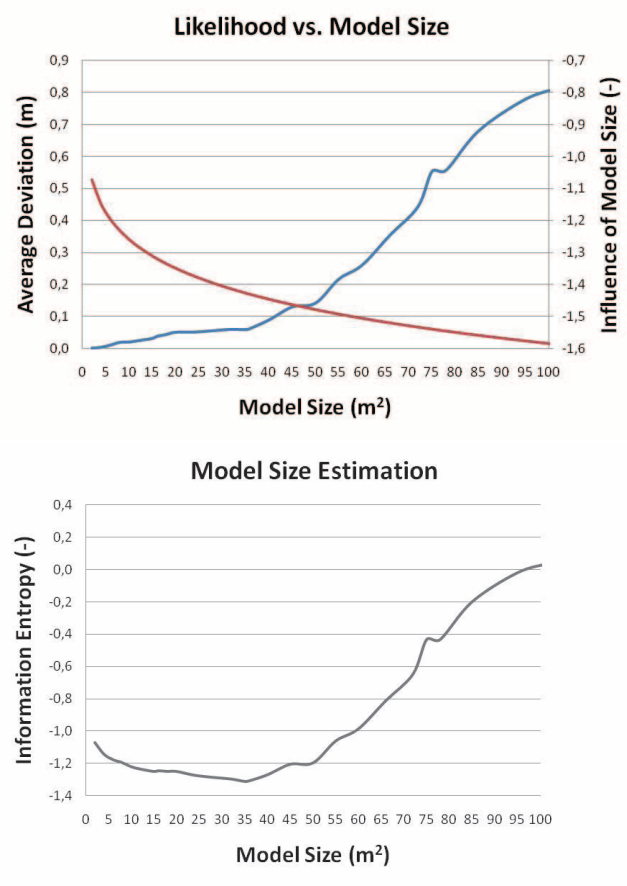

Figure 8: Plots of average deviation (blue), the influence of model size $-K^{\alpha}$ (red) and Akaike Information Criterion (black) over model size. The minimum AIC indicates the optimal model size balancing goodness of fit and complexity of the model.

information criterion, as shown in Figure 8 (bottom), is employed to guide the search of parameter-set $\mathcal{C}$. Please note, unlike AIC, this is not a trade-off between reconstruction accuracy and model complexity because less error does not mean better reconstruction either in this case. Besides, there is no general threshold for acceptable error. The optimization of the entropy function finds the maximum possible size of the model by perceiving its influence on the tendency of error change.

3.3 Reversible jumps

Reversible Jump Markov Chain Monte Carlo (RJMCMC) is an extension of the MCMC algorithm to handle solution spaces of variable dimensions. As introduced by Green in [3], the Markov Chain sampler is allowed to switch between subspaces with variable dimensions in the search. The dimension jumps, i.e., modeling with varying numbers of parameters, is employed in this work to simulate the choice of different roof primitives.

RJMCMC has a mixed transition kernel, which defines all possible movements. Studying the primitives, we narrow down the possible moves as shown in Figure 9. The given "jump routine" makes sure that each jump step only change a limited number and more sensible parameters. A transition matrix $\mathcal{T}$ indicating the jump probabilities between primitives is derived from the routine. Please note most elements in $\mathcal{T}$ are thus zero.

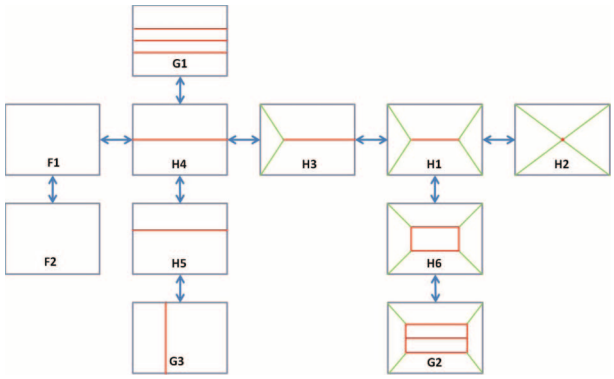

Figure 9: Possible jumps between primitives are limited by a "jump routine".

Let $\mathcal{M}_i$ and $\mathcal{M}_j$ be models in the family of primitives, $\{\mathcal{M}_n; n = 1, ..., N\}$. The move from i to j will be accepted according to the probability:

$$\mathcal{A}_{i\to j} = \min\left\{1, \frac{p(\mathcal{D}|\Theta_{\mathcal{M}_j})p(\Theta_{\mathcal{M}_j})}{p(\mathcal{D}|\Theta_{\mathcal{M}_i})p(\Theta_{\mathcal{M}_i})} \cdot \frac{\mathcal{J}_{i\to j}}{\mathcal{J}_{j\to i}}\right\}. \quad (4)$$

$\mathcal{J}_{i\to j}$ is the Jacobian of the transform from i to j. For simplicity, we use a fixed transition matrix $\mathcal{T}$ instead of the Jacobian matrix $\mathcal{J}$. Employing model selection mechanism for the transition kernel, the acceptance probability is expressed as:

$$\mathcal{A}_{i\to j} = \min\left\{1, \frac{H^{-1}_{\mathcal{M}_j}}{H^{-1}_{\mathcal{M}_i}} \times \mathcal{T}_{ij}\right\}. \quad (5)$$

The information entropy of the model $H_{\mathcal{M}}$ is calculated as:

$$H_{\mathcal{M}} = k^{\beta} - 2ln\left(L(\mathcal{D}|\mathcal{M})\right), \quad (6)$$

where β is used to tune the tolerance for the model complexity. In this work we prefer better reconstruction than simple model, so that β is given relatively small value, 1/12, to reduce the sensitivity to the parameter number.

As long as the jump kernels keep the balance condition, i.e., being reversible – possible to return to the previous state, the RJMCMC sampler can search in a great wider variety of hypothesis models. This makes RJMCMC itself has powerful model selection effect, but also usually time-consuming. We conduct the sampling process in this work as follows:

- Jumps are proposed till the maximum likelihood for the current primitive has been reached instead of switching primitives by random probability in every MCMC move (exception is the step A1).

- Model selection technique is integrated in the transition kernel for jumps.

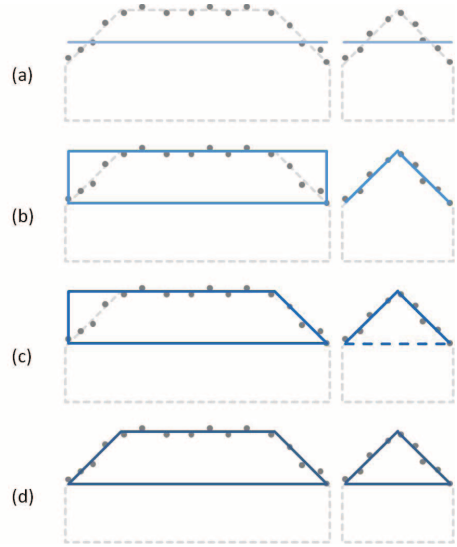

Figure 10: An example jump trace: (a) flat roof, (b) gable roof, (c) half-hipped roof to (d) hipped roof with increasing goodness of fit to the data points (gray) as well as the number of parameters. The gray dashed line shows the underlying building model.

By these means the computing time is reduced by "cooling down" the search entropy while the model selection ability of RJMCMC is compromised at the same time. The loss is, however, compensated by the explicit model selection mechanism.

4. LIKELIHOOD AND OPTIMIZATION

We use the average absolute deviation in z-direction (z-error) from the proposed model ($\mathcal{M}$) to the data points ($\mathcal{D}$) to assess new candidates.

$$\Delta_z = \frac{\sum_{f\in\mathcal{F}}(\sum_{i\in\Omega_f}|z_\mathcal{M} - z_\mathcal{D}|_i)}{K} \tag{7}$$

with f an individual facet from the facet-group ($\mathcal{F}$) of the primitive, i the data points in the domain of f: Ω_f, and K again the number of the involved data points.

Let X be the observations, the likelihood function for model with parameters Θ can be expressed as:

$$\Theta \mapsto L(\mathcal{D}) = L(X|\Theta) \propto \exp(-\Delta_z) \tag{8}$$

For optimization we employ the posterior of the proposed model integrating priors of the roof parameters: $p(\Theta)$, which are supposed to provide the following information:

- generic value ranges of parameters – implausible candidates, e.g., roofs with the height near ground, can be sorted out.
- recorded parameters of the already found primitives – the joint distribution is simplified by labeling particular areas. Possible further (often smaller) primitives are thereby easier to be found.

The Maximum A Posteriori (MAP) estimate can be expressed as:

$$\widehat{\Theta}_{MAP} = \underset{\Theta}{argmax}\left\{\frac{L(\mathcal{D}|\Theta)p(\Theta)}{P(\mathcal{D})}\right\} = \underset{\Theta}{argmax}\left\{L(\mathcal{D}|\Theta)p(\Theta)\right\} \tag{9}$$

with $p(\Theta) = p(\mathcal{M}_{found}|\Theta_0)\cdot p(\Theta_0)$ and $P(\mathcal{D})$ the marginal probability. $P(\mathcal{D})$ can be seen as a constant in the optimization, as it does not depend on Θ.

In practice, z-errors for different roofs in the same scene may not be close to each other. Some qualified primitives have relatively large deviation to the data points because of (1) the "occlusion" of other intersected primitive(s) and (2) the clutter of trees. This part of deviation cannot be seen as reconstruction error. On the contrary, it implies the regularity of the primitive is maintained in spite of the clutter. E.g., in the example scene shown in Figure 12 the error range of individual accepted primitives is [0.05, 1.10] meters, i.e., a fixed threshold for acceptance is no more feasible. As shown in Figure 11, however, the z-error for individual primitive converges well in the MCMC sampling. We therefore use the following stop criteria:

- The z-error becomes stable or
- A predefined maximum number of iterations is reached.

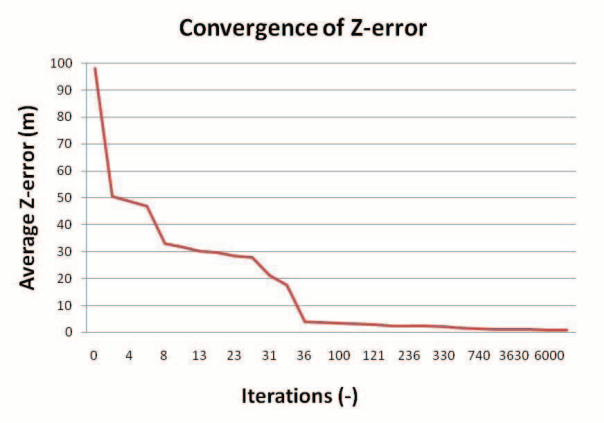

Figure 11: The convergence of z-error during the search of an individual primitive with MCMC sampler. After about 1000 iterations the reconstruction result became stable.

5. EXPERIMENTS AND RESULTS

Experiments are performed on both low-resolution (1m) and high-resolution (0.18m) data-sets.

Figure 12 shows a roof extraction from a laser scanning data (a) with the point density of 1.0 meter. The roofs are extracted in spite of the clutter from adjacent trees (red dashed circles) and reconstructed in the form of Virtual Reality Modeling Language (VRML) models (b). 2D projections of primitives on the ground are given in (c) presenting the way of the primitive combination.

Comparing the result with the reference image (Figure 12, d), all the buildings completely represented in the scene have been reconstructed with correct position and plausible

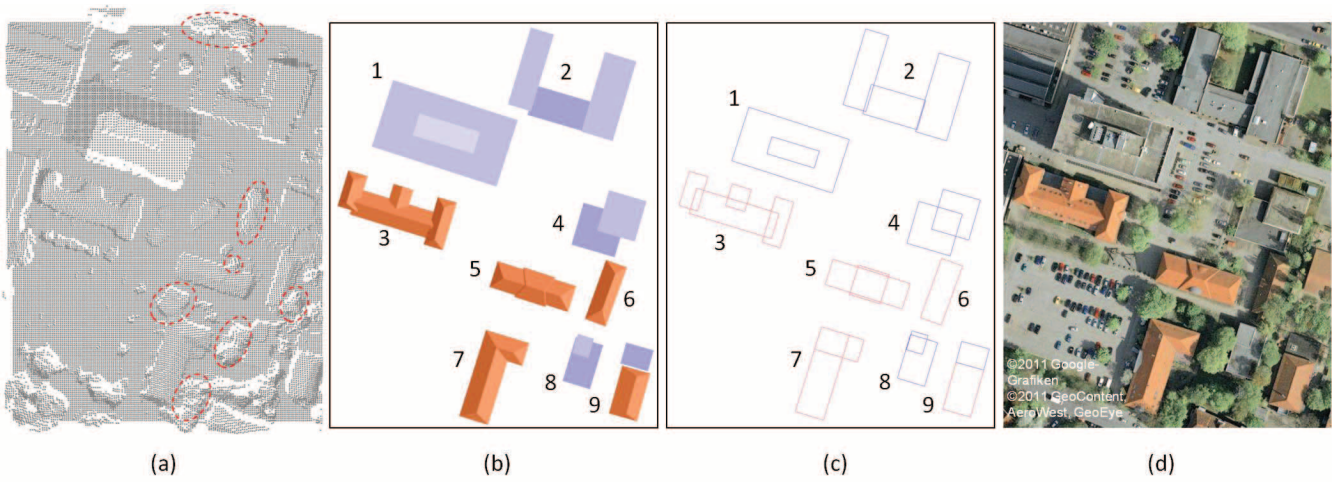

Figure 12: Experiment result: (a) laser scanning point cloud with several roofs cluttered by adjacent trees (red dashed circles), (b) the reconstruction results in the form of VRML models, (c) 2D projections of assembled roof primitives, (d) an aerial image as reference

shape. As there is no 3D ground truth data or other building models from, e.g., manual photogrammetric measurement, available, it is hard to give a quantitative evaluation for the 3D models. Quality measures, however, can be given as described in Section 4. We use the mean of z-errors for individual facets of the building (instead of the primitives) to approximate the reconstruction error, which has an average value of 0.11 meter for the scene. Please note this error still includes clutter of trees and the real reconstruction error should be less than it. The top-down reconstruction forces a plausible model by ignoring occlusions and data flaws, i.e., more accurate roof models may have larger deviation to the non-roof points. The runtime for this scene (32000 square meters, 9 buildings) is about 15 minutes in total and 1.7 minutes in average using a laptop with a 2×1.3 GHz processor. The most expensive part is the "global" search finding the first primitive of each building (cf. Section 3.1, A1), which may take more than 1 minute. Flat roofs and hipped roofs have no significant difference in runtime. The longest time for individual building was 2.5 minutes taken for building 2 as flat roof combination needs extra adjustment.

Figure 13 shows the building models derived from that of the roofs with the roof contours being extruded from the eave height to the ground.

Some narrow and small flat parts of buildings have not been extracted, e.g., that of the east wing of building 2, the fire escape shaft (top-left) of building 1, and the small structures in the north of building 4, as the point cloud is too sparse to represent them meaningfully. It is the same reason why a number of small structures on the roofs have not been reconstructed. Even relatively bigger dormers (e.g., that on the building 5) have only less than 10 data points on it. It is hard to tell if they belong to an individual structure or just data flaw or clutter.

High resolution data, relatively speaking, provides denser points representing small objects on the roof. As shown in Figure 14, superstructures, e.g., chimney and dormers, have been reconstructed from a point cloud with the density of about 0.18 meter. The average z-error is 0.05 meter.

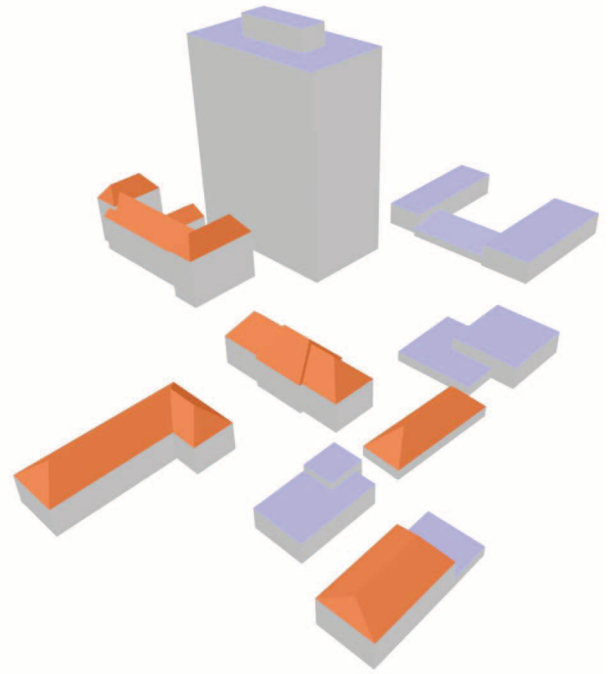

Figure 13: Building models derived from the extracted roofs

6. CONCLUSIONS AND OUTLOOK

In this paper we have proposed a generative statistical approach to the extraction and reconstruction of building roofs from laser scanning point clouds. The main contributions of this work can be summarized as follows:

– Generative modeling of building roofs with assembly of primitives allowing overlapping;

– Estimation of model size by perceptively balancing its

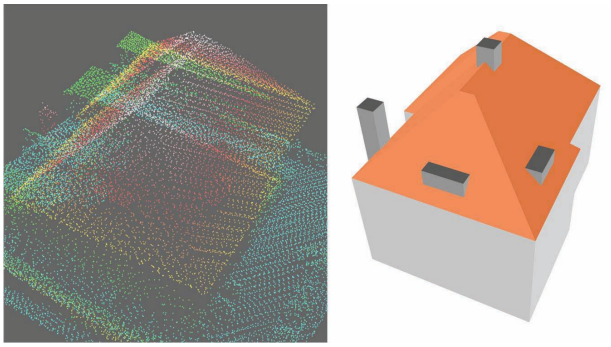

Figure 14: Roof reconstruction with superstructures (right) from high density data points (left)

influence on the reconstruction error to guide the parameter sampling;

– Scheduled reversible jumps switching between different primitives driven by model selection mechanism.

By all these means the reconstruction becomes robust against data flaws and clutter objects and a plausible result is guaranteed.

In this work we wanted to explore and demonstrate the potential of a pure top-down approach. We could show that in scenes with few buildings, i.e., lower complexity of the joint distribution landscape, the proposed algorithm can find the targets with only generic prior information.

Please note, however, that in complex scenes a complete replacement of the bottom-up process is hard. It turned out that in larger scenes, the complexity of the distribution as well as the number of disturbances is so high, that the process did not find appropriate building candidates. In practical application the real challenge is to balance the bottom-up and top-down partitions to achieve robustness as well as efficiency. This work gives more probability in the latter direction. With the proposed scheme the emphasis on the bottom-up effort can be significantly reduced and more space is given for the coordination of the both. We plan to develop a coarse segmentation scheme for building candidates in order to partition the space and apply our top-down approach to each individual partition. The segmentation can be either based on given GIS-information (e.g. road network) or by a coarse segmentation of the point cloud.

In this work our primitive library contains only planar roofs with at most 3 differing height levels. New entries, e.g., flat roofs in the forms of triangle and ellipse, domes, cones, and other curved shapes, are needed to present more sophisticated buildings like churches, exhibition centers and stadia.

7. ACKNOWLEDGMENTS

We thank the anonymous reviewers for their helpful comments. The research described in this paper is funded by the German Science Foundation (DFG).

8. REFERENCES

[1] H. Akaike. Information Theory and an Extension of the Maximum Likelihood Principle. In *Second International Symposium on Information Theory*, pages 267–281, Budapest, Hungary, 1973. Akademiai Kiado.

[2] C. Brenner. Building reconstruction from images and laser scanning. *International Journal of Applied Earth Observation and Geoinformation, Theme Issue on Data Quality in Earth Observation Techniques*, 6(3-4):187–198, 2005.

[3] P. Green. Reversible Jump Markov Chain Monte Carlo Computation and Bayesian Model Determination. *Biometrika*, 82:711–732, 1995.

[4] H. Huang and C. Brenner. Rule-based roof plane detection and segmentation from laser point clouds. In *Joint Urban Remote Sensing Event (JURSE) 2011*, pages 293–296, April 2011.

[5] F. Lafarge, X. Descombes, J. Zerubia, and M. Pierrot-Deseilligny. Structural approach for building reconstruction from a single dsm. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32:135–147, January 2010.

[6] M. Ortner, X. Descombes, and J. Zerubia. Building outline extraction from digital elevation models using marked point processes. *International Journal of Computer Vision*, 72:107–132, April 2007.

[7] F. Rottensteiner, J. Trinder, S. Clode, and K. Kubik. Automated delineation of roof planes in lidar data. In *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, volume 36(3/W19), pages 221–226, 2008.

[8] A. Sampath and J. Shan. Segmentation and reconstruction of polyhedral building roofs from aerial lidar point clouds. *IEEE Transactions on Geoscience and Remote Sensing*, 48(3):1554–1567, March 2010.

[9] R. Schnabel, R. Wessel, R. Wahl, and R. Klein. Shape recognition in 3d point-clouds. In *The 16th International Conference in Central Europe on Computer Graphics, Visualization and Computer Vision*, 2008.

[10] G. Vosselman. Advanced point cloud processing. In D. Fritsch, editor, *Photogrammetric Week '09*, pages 137–146. Heidelberg, Germany, 2009.