

ROOF RECONSTRUCTION FROM POINT CLOUDS USING IMPORTANCE SAMPLING

William Nguatem, Martin Drauschke, Helmut Mayer

Institute of Applied Computer Science
Bundeswehr University Munich
Werner-Heisenberg-Weg 39, 85577 Neubiberg, Germany
william.nguatem,martin.drauschke,helmut.mayer@unibw.de

Commission III/4

KEY WORDS: Building Reconstruction, Point Cloud Segmentation, MCMC, Model Selection, Non-linear Regression Fitting

ABSTRACT:

We propose a novel fully automatic technique for roof fitting in 3D point clouds based on sequential importance sampling (SIS). Our approach makes no assumption of the nature (sparse, dense) or origin (LIDAR, image matching) of the point clouds and further distinguishes, automatically, between different basic roof types based on model selection. The algorithm comprises an inherent data parallelism, the lack of which has been a major drawback of most Monte Carlo schemes. A further speedup is achieved by applying a coarse to fine search within all probable roof configurations in the sample space of roofs. The robustness and effectiveness of our roof reconstruction algorithm is illustrated for point clouds of varying nature.

1 INTRODUCTION

The generation of accurate 3D city models from point clouds is a very active research topic in computer graphics, photogrammetry and computer vision, (Wahl et al., 2008; Huang et al., 2013; Lafarge et al., 2013). The main focus of these related contributions has been to achieve (a) a significant amount of data reduction, (b) the automation of object recognition and reconstruction from unstructured 3D point clouds, and (c) the integration of semantics into building and 3D city models. However, most of the efforts in this domain are based on the analysis of terrestrial and airborne LIDAR data sets. The two major reasons for this are (1) the maturity of the LIDAR technology and (2) the resulting high accuracy and very dense 3D data.

However, due to the general availability of digital cameras, there is an immense interest in the generation of 3D point clouds from multiple images (Snively et al., 2006; Frahm et al., 2010; Bartelsen et al., 2012). Though this data acquisition pipeline is being improved due to the developments of new algorithms for dense matching such as (Strecha et al., 2008; Hirschmueller, 2008; Kuhn et al., 2013), the acquired 3D points are still inferior to data acquired from LIDAR as far as accuracy is concerned. However, the 3D data are good enough to reliably recognize objects and their parts. Furthermore, these algorithms also benefit from the inherent flexibility using a standard consumer camera. We see a trend in an increasing quality of the 3D point clouds from matching and thus a necessity for devising methods for roof reconstruction from point clouds, which are robust with respect to the different data acquisition process.

The goal of our work is to find, classify and fit basic roof types above a given quadrilateral building footprint. Thus, we aim at a level-of-detail (LOD) 2 model, consisting of the basic 3D shape such as a gable or a mansard roof. This will act as an intermediate state in a hierarchical modeling of buildings. A wide variety of applications in robotics, automotive navigation, tourism or city planning can benefit from such an automated building roof reconstruction scheme. We illustrate our algorithm with experimental results for data from matching images. We employ terrestrial images combined with images obtained from unmanned aircraft system collected in Southern Germany. Additionally, LIDAR data sets and synthetic data sets generated from sampling a

CAD model with additive noise are used in our experiments.

The paper is structured as follows. Section 2 discusses existing approaches for reconstructing roofs in 3D point clouds. Beginning with an overview, we present our algorithm in Section 3. This is followed by the discussion of experimental results in Section 4. In Section 5, we conclude and give directions for future work.

2 RELATED WORK

There is a significant amount of work on the reconstruction of building roofs from point clouds. Most often, this is done based on domain knowledge, e.g., using shape priors of the objects of interest (Huang et al., 2013). Sampath and Shan (2010) take curvature as a feature within a fuzzy k-Means framework to cluster points defining roof segments in aerial LIDAR data. A similar region growing approach (Rottensteiner, 2010) detects planar surfaces by combining information from multiple images and point clouds. These surfaces can further be classified into various roof types. A Markov Chain Monte Carlo (MCMC) approach to roof reconstruction is presented in (Huang et al., 2013). Here, samples of possible roof constellations are proposed from within a predefined library of planar shapes. Proposals may or may not be accepted within a Reversible Jump MCMC (RJMCMC) setting depending on their computed likelihoods. However, the generative nature of the algorithm makes it very sensitive to noise and the RJMCMC sampler means a bad scalability of the algorithm. In (Taillandier and Deriche, 2004), a bayesian approach is used to fit polyhedral models from aerial images. The resulting planar patches are used to construct graphs from which buildings are reconstructed.

Poullis and You (2009) presented a generic parametrization for the reconstruction of basic roof types. The likelihood for a specific roof type is assumed to underlie a Gaussian mixture model (GMM) whose parameters are inferred. However, the input data is assumed to be in perfect raster form. This is achieved by preprocessing including local energy minimization which can be highly error prone and is sensitive to noise. Moreover, the GMM assumption for a roof type implies perfect symmetry which will often not hold in real world. Recently, Henn et al. (2013) presented a model-based roof detection and classification algorithm

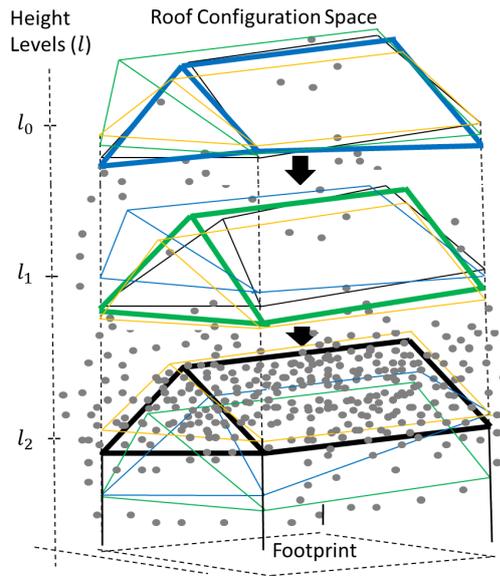


Figure 1: At each height level, the configuration with the highest likelihood is marked in bold. The black arrows describe the transition of the configurations from one level to the next and the gray dots are data points. Different colors correspond to different configurations. Please note that the space between height levels is widely exaggerated for a better comprehension.

that can cope with very sparse data sets. It assumes an underlying rectangular footprint. M-estimator random consensus (MSAC) proposed by Torr and Zisserman (2000) is used as a robust estimator to fit basic roof models. To differentiate between competing roof types, a support vector machine was used for classification. However, the approach is not suited for point clouds from image matching which are often highly irregular distributed and more complex roofs (mansard, gambrel etc.). Additionally, the assumption of a perfect rectangular footprint may be too strong a constraint. Besides the above mentioned deficits the majority of the mentioned algorithms are tailored for LIDAR data sets. Our algorithm contains an extension of the parametrization of (Poullis and You, 2009), combined with a likelihood model similar to (Henn et al., 2013) within a Monte Carlo (MC) setting similar to (Huang et al., 2013), to optimally explore the search space for roof types above a detected general quadrilateral footprint.

3 PROPOSED METHOD

3.1 Overview of the Algorithm

The input to our system is an unstructured 3D point cloud, \mathcal{D} , containing a building as viewed from outside and above. The output is a set of connected polygonal surfaces above a general quadrilateral footprint with a defined roof type, e.g., mansard, gable or pinnacle. We suppose the 3D region above the footprint defines the most plausible search space for a roof. Appropriately, we transform \mathcal{D} to a coordinate frame where the direction vector $(0, 0, 1)$ points upwards. Starting at the highest point above the footprint and descending in discrete steps, we estimate probability distributions over possible roof configurations at several height levels. Making no assumptions about the form of these probability distributions over roofs at each level, our objective is to find the configuration that “best” explains \mathcal{D} .

In Fig. 1, we show a gable roof illustration of this probabilistic roof fitting scheme using four roof configurations (coloured black, green, blue and yellow) and three height levels (l_0 , l_1 , l_2).

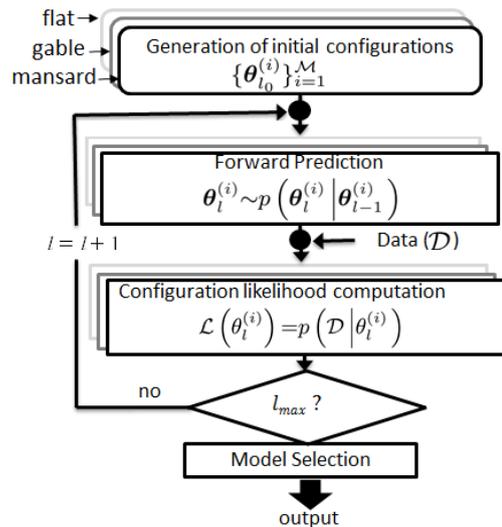


Figure 2: Roof fitting using sequential importance sampling.

First, we build initial roof configurations (level l_0) and estimate their likelihoods. Then, iteratively a transition is made to the next level and the roof configurations likelihood is updated. At each level, the configuration with the highest likelihood (marked bold in Fig. 1) explains our data better than all other configurations present. At height level l_2 , configuration θ_1 (black in Fig. 1) provides a better fit than all other configurations. The problem is defined in a Bayesian setting and a sequential Monte Carlo (SMC) approach is used to search for a solution.

3.2 Notation

Random points \mathbf{p}_j are considered in the subspace $\Phi \in \mathbb{R}^3$ of the 3D Euclidean space. An ordered sequence $\theta \equiv (\mathbf{p}_1 \cdots \mathbf{p}_q) \in \Phi^q$ uniquely defines a 3D surface in some standard way, e.g., in our experiments the \mathbf{p} 's are the vertices of a 3D polygonal chain. Thus, θ represents a single roof configuration. The symbol l represents a discrete height level index. \mathcal{P}_l defines a plane with normal vector $(0, 0, 1)$, passing through height level l . The symbol \sim means “distributed according to” as in $x \sim \mathcal{N}(m, \sigma^2)$ defining the Gaussian random variable x with mean m and standard deviation σ and $y \sim \mathcal{U}(a, b)$ refers to a uniform random variable y with support $[a, b]$. Draws are collected, unless otherwise stated, in an independent and identically distributed (i.i.d.) way. The same notation will be used for a random variable and its realization to reduce notational clutter.

Being in an MC setting, we use a set of discrete roof configurations to approximate a distribution over roofs. For example, at level l , $\{\theta_l^{(i)}\}_{i=1}^{\mathcal{M}}$ are a finite set of \mathcal{M} such 3D polygonal chains used as a discrete representation of $p(\theta_l | \mathcal{D})$, the distribution over polygonal chains at level l . The likelihood of a configuration, $\mathcal{L}(\theta_l^{(i)})$, denotes how the data is explained by that polygonal chain. The normalized likelihood defines the weight of a configuration, $w_l^{(i)}$. This is used as an approximation of the posterior probability, $p(\theta_l^{(i)} | \mathcal{D})$. At each height level, the configuration with the highest weight is reckon on as the maximum a-posteriori probability (MAP) configuration, denoted θ_{MAP} (bold lines in Fig. 1).

3.3 Roof Fitting using a Sequential Monte Carlo Method

Having the complete set of data, \mathcal{D} at our disposal, we define the sequence as height levels above the foot print. Starting at the highest point in \mathcal{D} above the footprint and descending in discrete

levels along the vertical direction, we employ sequential importance sampling (SIS) to consistently keep track of roof configurations at these height levels. This iteratively approximates the distribution over roof configurations. The goal is to move the starting configurations based on a forward transition distribution, $p(\theta_i|\theta_{i-1})$, which retains expected properties of the 3D polygonal chains to be fitted and on an observation model, $p(\mathcal{D}|\theta_i)$, that provides evidence about whether a measured 3D point in \mathcal{D} is in the close vicinity of the “true” 3D polygonal chain. The complete algorithm summarized in the flow chart given in Fig. 2 can be described as follows:

1. In initialization, we generate \mathcal{M} configurations for every considered roof type (mansard, gable, flat, etc.)
2. In forward prediction, these \mathcal{M} configurations per roof type are propagated using the transition probability distribution, $p(\theta_i|\theta_{i-1})$. This results in a change of configurations from one level to another
3. Within an update stage, likelihoods of the configurations, $\mathcal{L}(\theta_i^{(i)})$, are estimated and normalized, thus giving a discrete approximation of the posterior probability, $p(\theta_i|\mathcal{D})$
4. Steps 2 and 3 are iterated until a minimum level well above the ground plane (denoted by l_{max} index) is reached
5. At the minimum level, we are confronted with three problems: (1) Which of the considered roof types explains our data better than the others? (2) Amongst the levels, which posterior distribution, $p(\theta_i|\mathcal{D})$, should be chosen? (3) Considering all the \mathcal{M} configurations accounting for the chosen posterior distribution, which configuration explains our data “best”?

We divide the above algorithm into four major parts: Initialization, forward prediction, weight estimation (roof configuration likelihood estimation) and model selection. Details of these major parts of our algorithm are described in the following.

3.3.1 Initialization The generation of initial configurations means to choose an appropriate parametrization for the various basic roof types and corresponding prior distributions for them. We could extend the generic mansard roof parametrization of (Poullis and You, 2009) to include all symmetric and asymmetric variants of these basic roof types. However, this would be too computational intensive for a scene with few mansard roofed buildings. Therefore, we have devised more distinctive parametrization for different types of roofs. We use non-informative priors for the parameters defining the building heights, $\mathbf{p}_0, \mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3$ in Fig. 3 and Fig. 4. This prior compensates for uncertainty in the estimates of heights of the buildings walls. Therefore, for the points $\mathbf{p}_0, \mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3$, we make four i.i.d. draws given by

$$\mathbf{p}^{(i)} \sim \mathcal{U}(\mathbf{p}_k - 0.5h_d, \mathbf{p}_k + 0.5h_d), \quad (1)$$

where $k \in \{0, 1, 2, 3\}$ stands for the respective points \mathbf{p}_k , h_d is the height tolerance as depicted in Figs. 3 and 4 and i is the index of the i -th of \mathcal{M} configurations. Since the direction vector $(0, 0, 1)$ has been defined as the vertical direction of the scene, equation (1) means a 1-dimensional search along the z -axis. These four draws representing a prior over the height of the walls are fundamental for all roof types.

In addition to equation (1), we use domain knowledge and derive informative priors over the ridges, peaks, hips etc. to give a complete prior for a certain roof type.

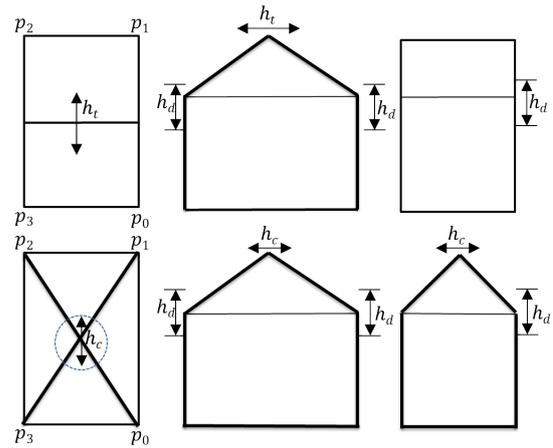


Figure 3: Top row: gable roof model including variance in building heights (h_d) and ridge position (h_t). Bottom row: Pinnacle roof model showing hip intersection point tolerance circle with diameter h_c .

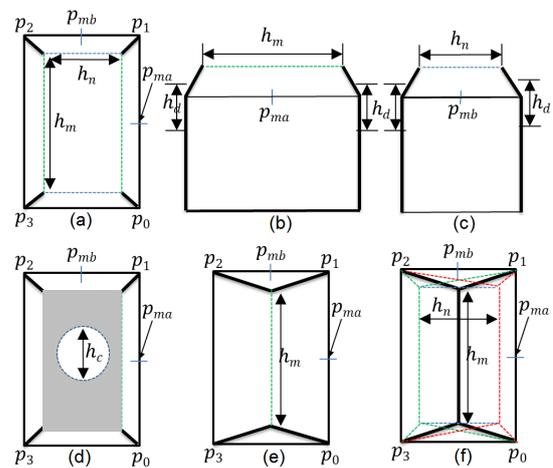


Figure 4: Top row (a), (b) and (c): Illustration of mansard roof parametrization. Bottom row: In (d) the gray region illustrates the actual search space for the mansard roof excluding pinnacle roof candidates. (e) and (f) illustrate the symmetric and asymmetric (half)-hipped roof parametrization respectively.

Gable Roof Prior Fig. 3 illustrates the parametrization of the gable roof. A prior over the ridge combined with equation (1) defines our informative prior model for the gable roof type. A realization of the ridge is defined by the line segment between points \mathbf{p}_{ra} and \mathbf{p}_{rb} projected on plane \mathcal{P}_i , where

$$\mathbf{p}_{ra}^{(i)} \sim \mathcal{N}(\mathbf{p}_{ma}, h_t) \quad \text{and} \quad \mathbf{p}_{rb}^{(i)} \sim \mathcal{N}(\mathbf{p}_{mb}, h_t) \quad (2)$$

with $\mathbf{p}_{ma} = 0.5 \cdot (\mathbf{p}_0 + \mathbf{p}_1)$ and $\mathbf{p}_{mb} = 0.5 \cdot (\mathbf{p}_2 + \mathbf{p}_3)$, h_t defines the tolerance in the ridge and i is the index of the i -th configuration. These are 1-dimensional draws along the direction of the walls. Since Equations (1) and (2) are i.i.d. samples, this prior model for the gable roof captures the complete family of gable roofs (symmetric and asymmetric). It is constructed twice such that in each case, the two sloping roof planes are oriented according to wall pairs. This ensures the right orientation of the roof planes to the pair of opposite walls.

Pinnacle Roof Prior Similarly, for the pinnacle roof shown in Fig. 3, a combination of Equation (1) and a prior over the peak location, \mathbf{p}_{peak} , defines the informative prior model. A realiza-

tion of the peak is defined by

$$\mathbf{p}_{peak}^{(i)} \sim \mathcal{N}(\mathbf{p}_{mid}, h_c) \quad (3)$$

with $\mathbf{p}_{mid} = 0.25 \cdot (\mathbf{p}_0 + \mathbf{p}_1 + \mathbf{p}_2 + \mathbf{p}_3)$ the centroid of the four points defining the footprint, h_c is the tolerance of the intersection point of all four hips and i is the configuration index. \mathbf{p}_{peak} is projected on \mathcal{P}_l to get the true peak.

Mansard Roof Similarly, Fig. 4 shows the parametrization of the mansard roof.

$$h_{mlength}^{(i)} \sim \mathcal{U}(0, h_n), \quad h_{mwidth}^{(i)} \sim \mathcal{U}(0, h_m) \quad (4)$$

where h_n and h_m are tolerance values defining the lengths of the ridges and again, i is the index representing the i -th configuration. The true ridges are a projection of the line segments $h_{mwidth}^{(i)}$ and $h_{mlength}^{(i)}$ on \mathcal{P}_l . If the outcome of Equation (4) lies within a circle with diameter h_c (cf. Fig. 4d), i.e., a highly probable pinnacle roof candidate, it is rejected. The final valid search space for mansard roof is marked gray in Fig. 4d.

Other Roof Types The most basic roof types are the flat and shed roofs. Their parametrization is obtained by varying the four points defining the building height, i.e., equation (1). Figs. 4e and 4f shows our template for the general family of hipped roofs, symmetric, asymmetric and (half)-hipped. Obviously, from this template the parametrization is obtained by a combination of gable and mansard roofs. To increase the effective number of samples and avoid “inverse” and thus invalid roof types, a constraint is added to all parameterizations that guarantees the ridge being higher than the building walls.

3.3.2 Forward Prediction In Fig. 2, we describe the movement of roof configurations from one height level to another as sampling from a forward transition distribution, $p(\theta_l | \theta_{l-1})$ given by the following simple first-order model

$$\theta_l^{(i)} \sim \theta_{l-1}^{(i)} \quad (5)$$

Having present the initialization and the forward transition probability, we now proceed to present how the configuration likelihoods of the 3D polygonal chains are computed.

3.3.3 Roof Configuration Likelihood Computation We use MSAC (Torr and Zisserman, 2000) to determine if a point in \mathcal{D} is within the close vicinity of a 3D polygonal chain defined by a configuration, θ . MSAC underlies a likelihood function proportional to a truncated Gaussian bounded by the RANSAC threshold, T^2 . Henn et al. (2013) used a similar likelihood model on aerial LIDAR data. Yet we optimally explore the configuration space in a Bayesian setting. Thus at level l , for a certain configuration, $\theta_l^{(i)}$, assuming measurements being conditionally independent, the likelihood function is given by

$$\mathcal{L}(\theta_l^{(i)}) \propto \exp\left(-\sum_j \rho(e_j^2)\right), \quad (6)$$

with

$$\rho(e_j^2) = \begin{cases} \frac{e_j^2}{T^2} & e_j^2 < T^2 \\ 1 & e_j^2 \geq T^2 \end{cases} \quad (7)$$

and e_j^2 the shortest Euclidean distance from point \mathbf{p}_j to the surface of the 3D polygonal chain defined by the configuration $\theta_l^{(i)}$. The summation index j runs over all the points in \mathcal{D} .

3.4 Model Selection

Equation (6) defines the likelihood for a single configuration of a specific roof type. We define the configuration weight w , as the

normalized likelihood, given by

$$w_l^{(i)} = \frac{\mathcal{L}(\theta_l^{(i)})}{\sum_{i=1}^{\mathcal{M}} \mathcal{L}(\theta_l^{(i)})} \quad (8)$$

Therefore, the posterior probability at level l for a roof type is approximated by

$$p(\theta_l | \mathcal{D}) \approx \hat{p}(\theta_l | \mathcal{D}) = \sum_{i=1}^{\mathcal{M}} w_l^{(i)} \delta(\theta_l - \theta_l^{(i)}) \quad (9)$$

where $\delta()$ defines the Dirac function. So far, three major questions are unanswered: (1) the roof type to choose (2) the height level to select and (3) within the chosen height level and roof type, the configuration that fits \mathcal{D} “best”. For the first two questions, we compare weights of the MAP configurations at each level for all the competing roof types and select the roof type whose θ_{MAP} has the highest weight. The height level within which the selected, θ_{MAP} resides defines the selected height level and is used to answer the third question. For this, we adopt the \mathcal{M} -opened modelling perspective (Smith and Bernardo, 2000) i.e., we believe the “best” configuration may not be within the \mathcal{M} configurations of the selected height level for the chosen roof type, however, through Bayesian model averaging, we compute the minimum mean square error estimator (MMSE) configuration as the inferred “best” configuration. For the chosen roof type within the selected height level, this is given by

$$\hat{\theta}_{MMSE} = \sum_{i=1}^{\mathcal{M}} \theta^{(i)} \hat{p}(\theta | \mathcal{D}). \quad (10)$$

The MMSE estimator in this setting is analogous to taking a weighted least square fit with the weights proportional to the estimated likelihoods as defined in section 3.3.3.

4 EXPERIMENTS

4.1 Implementation Details

The major computational bottleneck is Equation (6), since the likelihood has to be estimated \mathcal{M} times per roof type for every level. To speed up this computation, we use a coarse to fine search. First \mathcal{D} is downsampled using a voxel grid with leaf size $0.4m \times 0.4m \times 0.4m$ defining the input point density. A search is conducted with a large enough inter-height level value to capture the height level where the roof most probably lies. Around the selected level, we apply a finer inter-level value for the final search. We also use a fast version of the polygon hit test from (Schneider and Eberly, 2002) for the various polygon segments of a 3D polygonal chain defining a configuration. Furthermore, since the heights of walls of a building are correlated, a pair-wise sampling of the building height priors was preferred to four i.i.d. draws of Equation (1). In our formulation of SIS, a strong data parallelism is implied due to the 3D space partitioning in discrete height levels. Using alternative sampling schemes such as Gibbs, Metropolis-Hastings or sequential importance resampling samplers, this proper decoupling of samples from one level to another is impossible.

4.2 Parameters and Data Sets

For all performed experiments, we assumed the input data is on a metric scale and the following parameters were used:

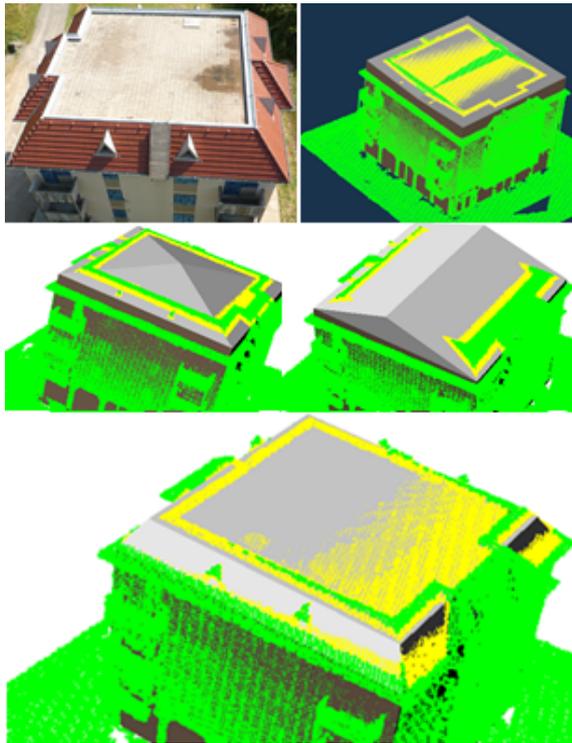


Figure 5: MMSE fit for flat, pinnacle, gable and mansard roof types on a mansard roofed building data set obtained from image matching. The yellow points are the inliers to the “best” 3D polygonal chain. As expected, the mansard roof model explains the data best by providing the best fit.

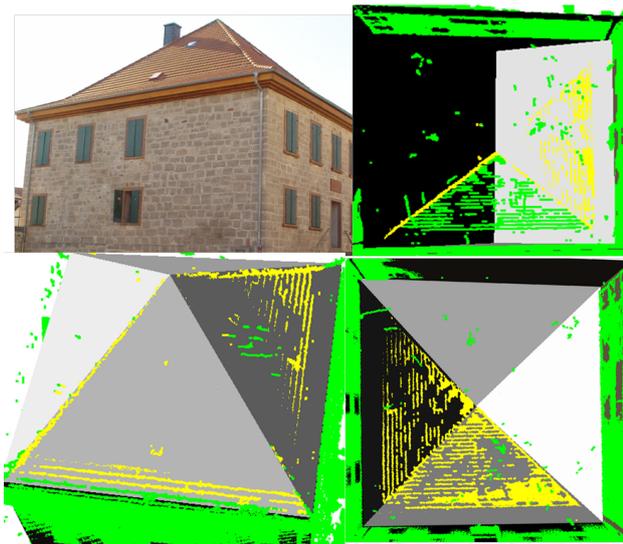


Figure 6: Top row: Image and MMSE fit using a gable roof type. The model is viewed from inside the building. Bottom row: MMSE fit using a pinnacle roof model. The model is viewed from outside (left) and inside (right) the building.

Dataset (roof type)	# footprints	correct	wrong
Laser (gable)	10	10	0
Matching (gable)	5	5	0
Synthetic (gable)	105	101	4
Laser (hipped)	5	5	0
Matching (hipped)	1	1	0
Synthetic (hipped)	15	14	1
Laser (pinnacle)	0	0	0
Matching (pinnacle)	1	1	0
Synthetic (pinnacle)	7	7	0
Laser (mansard)	0	0	0
Matching (mansard)	1	1	0
Synthetic (mansard)	3	3	0

Table 1: Results of our experiments conducted on different data sets. The number of correct versus wrong fits given the footprints.

Parameter	T	\mathcal{M}	h_c	h_t	h_d
Value	0.3m	1000	0.5m	0.5m	2m

For the mansard roof type, h_n and h_m were both chosen to be 0.5 meters less than length and width of the footprint. Furthermore, we discretized the search space above the footprint in 10 and 5 equal steps for the coarse and the fine search respectively.

Example outputs from data set obtained from image matching is shown in Figs. 5 and 6. For Fig.5, a combination of images taken from the ground and a UAV were used for matching and the mansard roof type best explains the data set. Meanwhile Fig. 6, shows a highly unequal point distribution in the data set with very few points on the roof as a result of the matching of images taken from the ground alone. We were still able to get the pinnacle roof type as the best fitted model. Since synthetic data sets are more or less perfect, for all experiments, all points were perturbed using

$$\mathbf{p}_j \sim \mathbf{p}_j + u_j \quad (11)$$

where $u_j \sim \mathcal{N}(0, 1m)$ for evaluation purpose.

4.3 Results

We tested our algorithm on a data sets containing 152 quadrilateral footprints of 3D points of varying density and origin. The results is summarized in the Table 1. All wrongly fitted roofs resulted in a flat roof type. This was due to the very low inclination angle of the roof planes. However, by tuning the parameters one can certainly get rid of these inaccuracies in the fits.

4.4 Clutter

The data sets on Figs. 5 and 6 contains dormers and a chimney respectively. However, these are fairly small roof superstructures and thus their effects on a roof fitting algorithm may be considered negligible compared to the highly unequal point distribution in the data set. Meanwhile, we demonstrate the robustness of our algorithm on a data set with substantial clutter through roof superstructures of reasonable sizes. In Fig. 7, results from a LIDAR data set containing a building with T-shaped footprint. We decompose the footprint into two separate non-overlapping quadrilaterals and apply our roof fitting algorithm 2 on each decomposed footprint. In both cases, gable roof type provided the best fit. The robustness against substantial clutter can be substantiated in twofold. First, our distinctive designed priors accounts for good discission making amongst competing models.

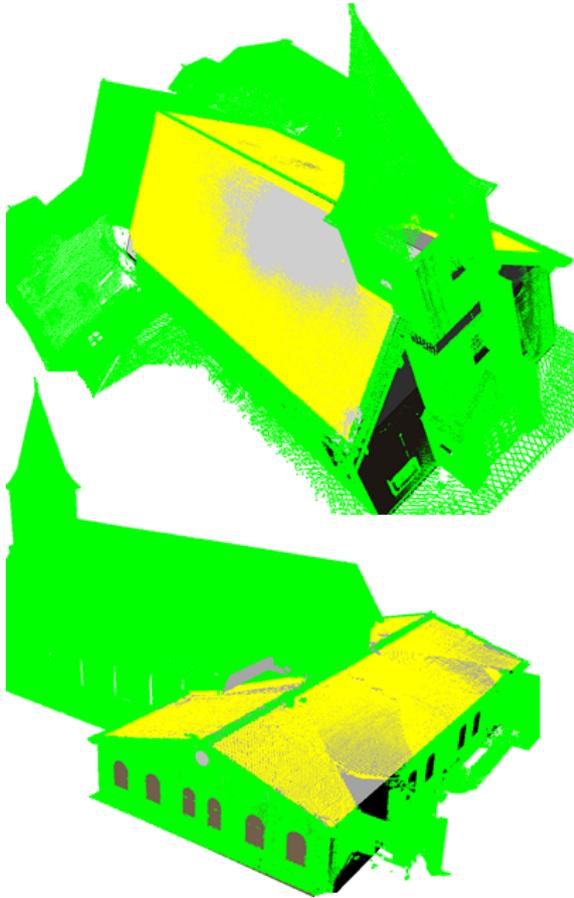


Figure 7: Data from LIDAR with a T-shaped building footprint decomposed into two separate non-overlapping quadrilaterals. The yellow points are the inliers of the MMSE fitted roof type over each footprint. Best fit in spite of substantial clutter

Secondly, the choice of the likelihood and the proper exploration of the search space within a Bayesian setting provides the optimal fit in MMSE sense.

5 CONCLUSION

We presented a sampling based method for automatically fitting and classifying building roof models from 3D point clouds. Assuming a general quadrilateral footprint, the method starts with the construction of parametric models of different roof types. These are represented by an ordered sequence of 3D polygonal surfaces. A likelihood function is chosen that considers the closeness of the 3D Data to the 3D polygonal chains. These likelihoods are used within a Bayesian setting to build probability distributions for various roof types in several possible discrete height levels above the footprint. Using model averaging, a model selection scheme is chosen to get the “best” model in MMSE sense. We achieve a computational speedup by applying a coarse to fine search. The method yields geometrically accurate parametric building roof models and is robust concerning the source and nature of the input data. For the generic case of modeling 3D buildings with a quadrilateral footprint, our approach robustly produces reliable results from which further refinements can be made, progressively increasing the level of detail towards detailed model (LOD3). A parametric representation is chosen which directly encodes to state-of-the art visualization and representation environments such as CityGML. Practical problems arise particularly due to extreme occlusions and wrong priors. To

overcome this problem, it would be desirable to extend the likelihood function.

References

- Bartelsen, J., Mayer, H., Hirschmüller, H., Kuhn, A. and Michelini, M., 2012. Orientation and Dense Reconstruction from Unordered Wide Baseline Image Sets. PFG (4), pp. 421–432.
- Frahm, J.-M., Fite-Georgel, P., Gallup, D., Johnson, T., Raguram, R., Wu, C., Jen, Y.-H., Dunn, E., Clipp, B., Lazebnik, S. and Pollefeys, M., 2010. Building Rome on a Cloudless Day. In: ECCV, LNCS 6314, pp. 368–381.
- Henn, A., Gröger, G., Stroh, V. and Plümer, L., 2013. Model Driven Reconstruction of Roofs from Sparse LIDAR Point Clouds. ISPRS Journal of Photogrammetry and Remote Sensing 76, pp. 17 – 29.
- Hirschmueller, H., 2008. Stereo Processing by Semiglobal Matching and Mutual Information. PAMI 30(2), pp. 328–341.
- Huang, H., Brenner, C. and Sester, M., 2013. A Generative Statistical Approach to Automatic 3D Building Roof Reconstruction from Laser Scanning Data. ISPRS Journal for Photogrammetry and Remote Sensing 79, pp. 29–43.
- Kuhn, A., Hirschmüller, H. and Mayer, H., 2013. Multi-Resolution Range Data Fusion for Multi-View Stereo Reconstruction. 35th GCPR, Saarbrücken, Germany.
- Lafarge, F., Keriven, R., Bredif, M. and Vu, H.-H., 2013. A Hybrid Multi-View Stereo Algorithm for Modeling Urban Scenes. PAMI 35(1), pp. 5–17.
- Poullis, C. and You, S., 2009. Photorealistic Large-Scale Urban City Model Reconstruction. IEEE Transactions on Visualization and Computer Graphics 15(4), pp. 654–669.
- Rottensteiner, F., 2010. Roof Plane Segmentation by Combining Multiple Images and Point Clouds. In: The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, Vol. XXXVIII, pp. 245–250.
- Sampath, A. and Shan, J., 2010. Segmentation and Reconstruction of Polyhedral Building Roofs From Aerial Lidar Point Clouds. IEEE Transactions on Geoscience and Remote Sensing 48(3), pp. 1554–1567.
- Schneider, P. and Eberly, D., 2002. Geometric Tools for Computer Graphics. Morgan Kaufmann.
- Smith, A. and Bernardo, J., 2000. Bayesian Theory. Wiley Series in Probability and Statistics, Wiley.
- Snively, N., Seitz, S. M. and Szeliski, R., 2006. Photo Tourism: Exploring Photo Collections in 3D. ACM Transactions on Graphics 25(3), pp. 835–846.
- Strecha, C., von Hansen, W., Van Gool, L., Fua, P. and Thiennessen, U., 2008. On Benchmarking Camera Calibration and Multi-View Stereo for High Resolution Imagery. In: CVPR.
- Taillandier, F. and Deriche, R., 2004. Automatic Buildings Reconstruction from Aerial Images : a Generic Bayesian Framework. In Proceedings of the XXth ISPRS Congress, Istanbul, Turkey.
- Torr, P. H. S. and Zisserman, A., 2000. MLESAC: A New Robust Estimator with Application to Estimating Image Geometry. CVIU 78(1), pp. 138–156.
- Wahl, R., Schnabel, R. and Klein, R., 2008. From Detailed Digital Surface Models to City Models using Constrained Simplification. PFG 2008(3), pp. 207–215.