

Partial Procedural Geometric Model Fitting for Point Clouds

Zongliang Zhang^{1*} Jonathan Li^{1,2} Yulan Guo³ Yangbin Lin¹ Ming Cheng¹ Cheng Wang¹

¹Fujian Key Laboratory of Sensing and Computing for Smart Cities, Xiamen Key Laboratory of Geospatial Sensing and Computing, Innovation Center of Sensing and Computing for Smart Cities, School of Information Science and Engineering, Xiamen University, Xiamen, Fujian 361005, China

²Mobile Sensing and Geodata Science Lab, Department of Geography and Environmental Management, University of Waterloo, Waterloo, ON N2L 3G1, Canada

³College of Electronic Science and Engineering, National University of Defense Technology, Changsha, Hunan 410073, China




Figure 1: Model evolution of fitting a procedural geometric model \mathcal{M}_8 to a point cloud Q_{11} (green) (Section 6.4). From left to right: fitted models at different iterations (time in seconds): 0 (0), 683 (979.4786), 1782 (2458.3916), 5566 (5968.5444), 16666 (11487.1016) and 43299 (23074.8750), respectively.

Abstract

Geometric model fitting is a fundamental task in computer graphics and computer vision. However, most geometric model fitting methods are unable to fit an arbitrary geometric model (e.g. a surface with holes) to incomplete data, due to that the similarity metrics used in these methods are unable to measure the rigid partial similarity between arbitrary models. This paper hence proposes a novel rigid geometric similarity metric, which is able to measure both the full similarity and the partial similarity between arbitrary geometric models. The proposed metric enables us to perform partial procedural geometric model fitting (PPGMF).

The task of PPGMF is to search a procedural geometric *model space* for the model rigidly similar to a *query* of non-complete point set. Models in the procedural *model space* are generated according to a set of parametric modeling rules. A typical *query* is a point cloud. PPGMF is very useful as it can be used to fit arbitrary geometric models to non-complete (incomplete, over-complete or hybrid-complete) point cloud data. For example, most laser scanning data is non-complete due to occlusion. Our PPGMF method uses Markov chain Monte Carlo technique to optimize the proposed similarity metric over the model space. To accelerate the optimization process, the method also employs a novel coarse-to-fine model dividing strategy to reject dissimilar models in advance. Our method has been demonstrated on a variety of geometric models and non-complete data. Experimental results show that the PPGMF method based on the proposed metric is able to fit non-complete data, while the method based on other metrics is unable. It is also shown that our method can be accelerated by several times via early rejection.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric algorithms, languages, and systems;

*e-mail: zhangzongliang@stu.xmu.edu.cn

†e-mail: junli@xmu.edu.cn (Corresponding author)

Keywords: inverse procedural modeling, partial shape fitting, rigid geometric similarity, incomplete point cloud reconstruction, mean measure

1 Introduction

A geometric model is a continuous point set (e.g. a surface). A geometric model space is a set of geometric models. A procedural model space is defined by a set of parametric modeling rules [Dang et al. 2015]. Retrieving desired model from procedural space is called inverse procedural modeling (IPM). As an important but challenging problem in computer graphics and computer vision, IPM has been actively studied in recent years [Musialski et al. 2013]. PPGMF is a special case of IPM, as it searches a procedural geometric model space for the model which is rigidly similar to a query of non-complete point set. PPGMF can be used in a number of applications including pattern recognition, shape matching, geometric modeling and point cloud reconstruction.

The task of basic geometric model fitting (BGMF) is to rigidly fit basic geometric models to the input geometric data (i.e. a point set). The most famous BGMF method is RANSAC [Fischler and Bolles 1981], which can be used to fit several basic geometric models such as planes and spheres. Following RANSAC, a lot of BGMF methods have been proposed [Isack and Boykov 2011]. However, BGMF methods cannot be used to fit arbitrary models. To fit arbitrary models, the rigid geometric similarity between arbitrary models has to be calculated. A rigid geometric similarity metric should ensure that a model is most similar to itself than any other models. To the best of our knowledge, symmetric Hausdorff distance (SHD) is the only rigid geometric similarity metric. However, it is time consuming to calculate SHD, making it difficult to perform arbitrary model fitting. This paper hence proposes a novel efficient rigid geometric similarity metric to perform arbitrary model fitting.

Given the similarity metric, the remaining task of arbitrary model fitting is to optimize over a given arbitrary model space. An arbitrary model space is usually defined by a procedural modeling

approach. That is, a set of procedural modeling rules are used to generate models [Smelik et al. 2014]. Retrieving desired models from procedural model space is called IPM. Many IPM methods have been proposed using different retrieval criteria such as indicator satisfying [Vanegas et al. 2012b] and image resembling [Teboul et al. 2013] [Lake et al. 2015]. In this paper, we focus on geometric criterion based IPM (GIPM) methods, which aim at fitting procedural geometric models to the input geometric data. As one geometric criterion, voxel difference (VD) has been investigated in GIPM methods [Talton et al. 2011] [Ritchie et al. 2015]. However, VD is an approximate geometric similarity criterion. It is obvious that part information is lost by voxelization, as the voxelization resolution cannot be as small as 0. Theoretically, our metric does not rely on a resolution. In other words, VD-based GIPM method cannot be used for rigid model fitting, which requires the calculation of rigid geometric similarity. To the best of our knowledge, our PPGMF method is the first GIPM method which can be used for arbitrary rigid model fitting.

PPGMF aims at rigidly fitting procedural geometric models to a query of non-complete (incomplete, over-complete or hybrid-complete) geometric object. “Hybrid-complete” means hybrid incomplete and over-complete. Typical queries are point clouds. There are two key processes in a PPGMF method. The first key process is to calculate the rigid geometric similarity between a model and the query. We have found that a SHD or VD based method is unable to fit non-complete point clouds. We hence propose a novel partial rigid geometric similarity metric to fit non-complete data. The second key process is to optimize over the procedural model space. Optimizing over the procedural space is challenging due to the hierarchical and recursive nature of modeling rules. Markov chain Monte Carlo (MCMC) technique is used to perform optimization. Although the similarity calculation based on our metric is faster than SHD, it is still too slow for practical applications. We hence propose a novel coarse-to-fine model dividing strategy to reject dissimilar models in advance to accelerate this optimization.

PPGMF is of important significance, and is very useful as it can be used to fit arbitrary geometric models to non-complete data. For example, most laser scanning data is non-complete (cluttered) due to occlusion [Guo et al. 2014]. We have tested our metric and PPGMF method on a variety of geometric models and non-complete data. Experimental results show that the PPGMF method based on the proposed metric is able to fit non-complete data, while the SHD or VD based method fails. Experimental results also show that our method can be accelerated by several times using early rejection.

In summary, our contributions are: (1) A novel rigid geometric similarity metric is proposed to measure the similarity between two geometric models. (2) An effective method is proposed to rigidly fit arbitrary geometric models to non-complete data. (3) A coarse-to-fine geometric model dividing strategy is proposed to reject dissimilar models in advance for the acceleration of PPGMF.

The rest of this paper is organized as follows. Sections 2 and 3 present related work and the overview of our method, respectively. Section 4 introduces our rigid geometric similarity metric. Section 5 presents the MCMC-based optimization approach and our coarse-to-fine model dividing strategy. Sections 6 and 7 present experimental results and conclusion, respectively.

2 Related Work

Most GIPM methods take either a particular type of geometric model or geometric data as input. BGMF methods such as [Fischler and Bolles 1981] [Isack and Boykov 2011] work on basic geometric models. [Debevec et al. 1996] [Mathias et al. 2011] rely on image information to achieve IPM while our work does not rely

on images. [Ullrich et al. 2008] assumes the number of model parameters is fixed. [Bokeloh et al. 2010] takes symmetry as an assumption. [Wan and Sharf 2012] is limited to facade point clouds and split grammar. [Vanegas et al. 2012a] takes Manhattan-World as an assumption. [Boulch et al. 2013] is limited to constrained attribute grammar. [Toshev et al. 2010] [Lafarge et al. 2010] [Huang et al. 2013] work well on airborne laser scanning data, however, it is hard to extend them to other types of data. [Stava et al. 2014] works on tree models. [Demir et al. 2015] relies on semi-automatic segmentation operations. Our method is full-automatic and makes no assumption about the type of input geometric model and geometric data. Consequently, similar to [Talton et al. 2011] [Ritchie et al. 2015], our method can be used for general-purpose GIPM.

It is worth noting that there are a lot of other types of geometric data reconstruction methods but not GIPM methods such as [Pu and Vosselman 2009] [Zheng et al. 2010] [Nan et al. 2010] [Li et al. 2011] [Lafarge and Mallet 2012] [Poullis 2013] [Lin et al. 2013] [Lin et al. 2015] [Monszpart et al. 2015] [Wang and kai Xu 2016]. A GIPM method results in procedural models, while other types of methods usually result in mesh models. Procedural model is more powerful than mesh model [Weissenberg et al. 2013], as it perceives the abstract causal structure of input data [Lake et al. 2015].

3 Method Overview

Figure 2 shows our PPGMF pipeline. The task of PPGMF is to search a procedural geometric model space for the model which is rigidly geometrically similar to a query of non-complete point set. Hence the input of a PPGMF method consists of a non-complete query and a set of parametric geometric procedural modeling rules, which defines the target model space. A query can be a continuous or discrete point set (i.e. a point cloud). In this paper, we focus on point cloud queries.

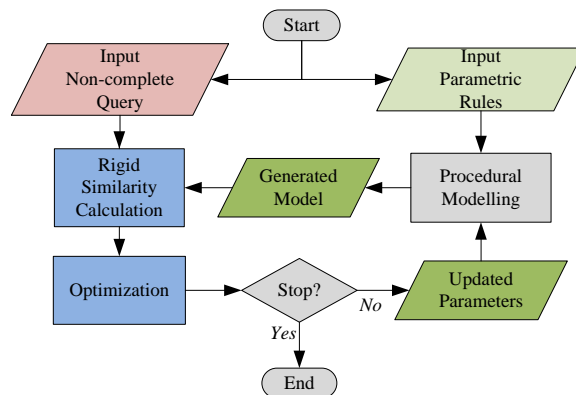


Figure 2: Our PPGMF pipeline.

An example set of procedural modeling rules is shown in Table 1, where $p(\cdot)$ denotes parameter prior. There are 3 rules A , B and D in this example. The axiom rule A manages a non-recursive parameter x_D and a recursive parameter x_B . It is straightforward to optimize for non-recursive parameters. However, it is challenging to deal with recursive parameters. When the rules are executed, a recursive parameter will spawn a family of non-recursive parameters under the same name of this recursive parameter. We have to individually identify every non-recursive parameter spawned from the same recursive parameter. To this end, a calling trace can be used. For example, we can identify a parameter x_B in a calling level 3 like this $x_B: A_B2_B1_B2$.

As shown in Fig. 2, given the query and the rules with parameter

rule $A()$	rule $B(\mathbf{x}_B)$	rule $D(\mathbf{x}_D)$
Sample $\mathbf{x}_B \sim p(\mathbf{x}_B)$	Sample $\mathbf{x}_B \sim p(\mathbf{x}_B)$...
Call $B(\mathbf{x}_B)$	Call $B(\mathbf{x}_B)$	end rule
Sample $\mathbf{x}_D \sim p(\mathbf{x}_D)$	Sample $\mathbf{x}_B \sim p(\mathbf{x}_B)$	
Call $D(\mathbf{x}_D)$	Call $B(\mathbf{x}_B)$	
...	...	
end rule	end rule	

Table 1: An example set of rules.

\mathbf{x} , a similarity calculation procedure is used to calculate the rigid geometric similarity between the query and the model generated by the procedural modeling procedure according to the rules. Based on the calculated similarity, \mathbf{x} is iteratively updated by the optimization procedure. It is worth noting that the number of parameters may vary during the optimization. Based on Bayesian inference theory, the optimization problem can be formulated as follows:

$$\max_{\mathbf{x}} p(\mathbf{x}|Q) \propto L(Q|\mathbf{x})p(\mathbf{x}) \quad (1)$$

where Q is the query, $p(\cdot)$ is the posterior of the parameters given the query, $L(\cdot)$ is the likelihood of the query given the parameters, and $p(\cdot)$ is the parameter prior.

The prior is directly drawn from the input modeling rules. Assume that $M^{\mathbf{x}}$ represents the model corresponding to \mathbf{x} , the likelihood depends on the rigid similarity between $M^{\mathbf{x}}$ and Q (Section 4), the optimization process is then performed (Section 5).

4 Rigid Geometric Similarity

A rigid geometric similarity metric should ensure that a geometric model is most rigidly similar to itself than any other models. Let \mathcal{M} be the universal geometric model space, the self rigidly similar property is formally stated as:

$$\forall M \in \mathcal{M}, \forall M' \neq M, s(M', M) < s(M, M) \quad (2)$$

where $s(\cdot, \cdot)$ denotes the similarity metric.

4.1 Full Similarity

SHD is the only metric satisfying the self rigidly similar property. The SHD between two point sets P and P' is defined as:

$$d_s(P, P') = \max \{d(P, P'), d(P', P)\} \quad (3)$$

where $d(\cdot, \cdot)$ is the one-sided Hausdorff distance (OHD):

$$d(P, P') = \max_{a \in P} \min_{b \in P'} \|a - b\| \quad (4)$$

where $\|\cdot\|$ is Euclidean norm. Note that, in general, $d(P, P') \neq d(P', P)$.

Usually, we can use SHD to exactly measure the rigid similarity between two point sets P and P' . If $d_s(P, P')$ is 0, then P and P' are the same. However, to calculate SHD, we have to compute OHD two times, i.e., one time from P to P' and another time from P' to P . In other words, only one OHD is insufficient for similarity assessment [Aspert et al. 2002]. If only one OHD calculation is required, the computational cost for similarity calculation can be reduced. Fortunately, one of the point sets involved in our similarity calculation is a geometric model. The measure of the model allows us to compute OHD only once to assess similarity. It is worth noting that different types of models have different types of measures.

For example, the measure of a curve is its length, the measure of a surface is its area.

Our insight is that, in real world, two models M and M' are identical if and only if every point of M is in M' (i.e. $d(M, M') = 0$) and the measure of M is equal to the measure of M' . That means the measure can be used for similarity assessment. We hence propose a Mean Measure (MM) to represent the rigid similarity between a model M and a query point set Q . Formally, MM is defined as the ratio of the measure of M to the OHD from M to Q :

$$r(M, Q) = \frac{|M|}{\epsilon + d(M, Q)} \quad (5)$$

where $|\cdot|$ denotes the measure of M , ϵ is a small positive number used to derive different similarities for the models with different measures but the same OHD of 0. For example, as shown in Fig. 3, both $d(C_1, C_1)$ and $d(C_2, C_1)$ are equal to 0. If ϵ is 0, then both $r(C_1, C_1)$ and $r(C_2, C_1)$ are infinite despite C_1 is more similar to C_1 than C_2 . Theoretically, when ϵ is sufficiently small, the similarity metric defined by MM can ensure that a model is most similar to itself than any other models.




Figure 3: An illustration of rigid similarity. (a) Curve C_1 , (b) Curve C_2 , (c) Curve C_3 , and (d) the overlap of C_1 , C_2 and C_3 . The overlapping part (black) shows that C_2 is a part of C_1 or C_3 .

Theorem. MM is a rigid geometric similarity metric for real-world geometric models. Note that, a real-world model M has a positive finite measure, i.e. $0 < |M| < +\infty$. Let \mathcal{M}_W be the real-world geometric model space, this theorem is proved as follows.

Proof. $\forall M \in \mathcal{M}_W, \forall M' \neq M$, (1) If $M' \subset M$, then $|M'| < |M|$ and $d(M', M) = d(M, M)$, so $r(M', M) < r(M, M)$ because $\epsilon > 0$; (2) If $M' \not\subset M$, then $r(M', M) < +\infty$ because $|M'| < +\infty$ and $d(M', M) > 0$. Meanwhile, $r(M, M) \rightarrow +\infty$ because, $d(M, M) = 0$, $0 < |M|$ and ϵ is assumed to be sufficiently small, i.e. $\epsilon \rightarrow 0$. So $r(M', M) < r(M, M)$. So MM is a rigid similarity metric for real-world models according to Eq. (2).

Note that, the values of MM are comparable for the same query, but are incomparable for different queries. That is, it makes no sense to compare the MM values across different queries. For example, as shown in Fig. 3, it makes no sense to compare $r(C_2, C_2)$ and $r(C_2, C_1)$, although $r(C_2, C_2) = r(C_2, C_1)$. It is worth noting that, the query is unnecessary to have a geometric measure. That is, the query can be a discrete point set, i.e., a point cloud. If the query Q is discrete, then ϵ is trivial because $d(M, Q)$ is always larger than 0. In practice, we use squared mean measure (SMM), which is a variant of MM and is defined as:

$$r_s(M, Q) = \frac{|M|}{\epsilon + d^2(M, Q)} \quad (6)$$

4.2 Partial Similarity

SHD and MM are defined as full similarity metrics as they assume that the query is complete. However, if the query is non-complete, we have to calculate partial similarity, which is challenging. Partial similarity is not straightforward and is fundamentally different from full similarity. If two point sets have a common part, then these two

point sets are partially similar. As shown in Fig. 3, each pair of C_1 , C_2 and C_3 are partially similar, while they are not fully similar. We expect that the partial similarity between C_1 and C_2 is equal to the partial similarity between C_2 and C_2 . Because the common part between C_1 and C_2 is the same as the common part between C_2 and C_2 .

Consequently, we propose a Weighted Mean Measure (WMM) to represent the partial similarity between a geometric model M and a query point set Q . We divide M into N non-overlapping sub-models: $M = \bigcup_{i=1}^N M_i$, and define WMM as:

$$r_w(M, Q) = \frac{\sum_{i=1}^N w_i |M_i|}{\epsilon + d_w(M, Q)} \quad (7)$$

where w_i is the weight: $w_i = \exp(-d(M_i, Q)h)$, where h is the weighting factor, which is a non-negative number. When h is 0, WMM becomes a full similarity metric. $d_w(\cdot, \cdot)$ is the weighted mean error [Aspert et al. 2002]:

$$d_w(M, Q) = \frac{\sum_{i=1}^N w_i d(M_i, Q)}{\sum_{i=1}^N w_i} \quad (8)$$

By weighting, the sub-models of M far away from Q have less contribution to the computation of WMM than those close sub-models. In other words, the common part of M and Q makes major contribution to WMM, making WMM plausible to measure partial similarity. One merit of WMM is that it has only one argument h , as ϵ is trivial. Similar to MM, it can be easily proved that WMM is a rigid geometric similarity metric.

Let C_1 , C_2 and C_3 be model spaces containing models C_1 , C_2 and C_3 (Fig. 3), respectively. There are 3 cases of partial similarity for PPGMF. (1) Target model is a part of query. For example, target model C_2 is a part of query C_1 . This case corresponds to partially fitting C_2 to over-complete data C_1 . (2) Query is a part of target model. For example, query C_2 is a part of target model C_1 . This case corresponds to partially fitting C_1 to incomplete data C_2 . (3) Target model and query have common part. For example, target model C_1 and query C_3 have common part. This case corresponds to partially fitting C_1 to hybrid-complete data C_3 .

4.3 Similarity Calculation

To compute MM (SMM or WMM) between a model and a point set, we have to compute OHD from the model to the point set, which consists of two steps. First, the model is uniformly divided into sub-models, and the center points of the sub-models are sampled (Section 5.2). Second, the nearest point is searched in the point set for a query point. This is time-consuming if the point set contains a large number of points (e.g. a laser scanning point cloud consisting of millions of points). We employ the FLANN [Muja and Lowe 2014] algorithm to perform nearest neighbour searching. The computational complexity for computing MM depends on the number of points sampled from the model and the size of the point set.

5 Optimization

Given the rigid geometric similarity defined by MM, we empirically define the likelihood in the optimization problem (see Eq. (1)) as:

$$L(Q|\mathbf{x}) = \exp\left(\sqrt{r(M^{\mathbf{x}}, Q)}\right) \quad (9)$$

Eq. (1) defines a derivative-free optimization problem, for which traditional mathematical optimization methods are not applicable. We use the Metropolis-Hastings (MH) algorithm [Metropolis et al. 1953] [Hastings 1970] to solve Eq. (1). MH algorithm is a general and popular MCMC optimization algorithm [Talton et al. 2011].

5.1 Metropolis-Hastings Algorithm

Let x^i be the value of variable x in iteration i , the MH algorithm works as follows. First, x is randomly initialized as x^0 . To determine x^{i+1} in each iteration, \tilde{x} is sampled from a proposal density function $p(x|x^i)$. The probability of accepting \tilde{x} as x^{i+1} is defined as:

$$\alpha(x^i \rightarrow \tilde{x}) = \min\left\{1, \frac{p(\tilde{x}|Q) p(x^i|\tilde{x})}{p(x^i|Q) p(\tilde{x}|x^i)}\right\} \quad (10)$$

That is, the probability for $x^{i+1} = \tilde{x}$ is α , and the probability for $x^{i+1} = x^i$ is $(1 - \alpha)$.

We now define the proposal function for the modeling parameter \mathbf{x} . Similar to [Vanegas et al. 2012b], each parameter $x \in \mathbf{x}$ is required to be within a range of $[x_{min}, x_{max}]$. For a continuous parameter, we randomly select one of the following two proposal functions, i.e., local move function and global move function in each iteration. The local move function is a Gaussian function, that is, $\tilde{x} \sim \mathcal{N}(x^i, \sigma_x^2)$, where $\sigma_x = \sigma(x_{max} - x_{min})$, and σ is the standard deviation ratio. The global move function is a uniform function, that is, $\tilde{x} \sim [x_{min}, x_{max}]$. We use β to denote the probability for selecting local move function, and $1 - \beta$ to denote the probability for selecting global move function. For a discrete parameter, we always perform global move. Since both local move and global move functions are symmetric, the probability of accepting \tilde{x} is simplified as:

$$\alpha(x^i \rightarrow \tilde{x}) = \min\left\{1, \frac{p(\tilde{x}|Q)}{p(x^i|Q)}\right\} \quad (11)$$

5.2 Early Rejection

The acceptance probability indicates that the proposed model with a larger similarity is more likely to be accepted than the model with a smaller similarity. More time will be consumed to obtain more accurate similarity since more points have to be sampled from the model. However, we observe that, it is sufficient to determine the dissimilarity by sampling only one point from the model. As shown in Fig. 4, Curve C_4 consists of one horizontal line segment, and Curve C_5 consists of two vertical line segments, these two curves are dissimilar. The similarities computed by sampling one point (Fig. 4c) and four points (Fig. 4a) are the same and equal to the true similarity. However, if two points are sampled (Fig. 4b), the computed similarity will be incorrect as it shows that C_4 and C_5 are similar. It can be inferred that a small similarity between two objects means that these two objects are dissimilar. However, a large similarity between two objects does not mean that these two objects are really similar. In other words, a proposed model should be accepted carefully but rejected boldly.

Consequently, to reduce computational time, we propose a coarse-to-fine model dividing strategy for similarity calculation to reject dissimilar models in advance. We take a square surface for example (as shown in Fig. 5), and the conclusions can be easily adapted to other types of geometric models. Assuming that the length of the square surface is γ , given a predefined minimal dividing resolution δ , the top dividing level is:

$$\eta_{top} = \log_2(\gamma/\delta + 1) \quad (12)$$




Figure 4: *Overlap between Curves C_4 (green) and C_5 (blue). Black dots represent the points sampled from C_4 . (a), (b) and (c) show that 4, 2 and 1 point(s) are sampled, respectively.*

At each level η , we uniformly divide the surface into $2^{2\eta}$ sub-surfaces, and sample only one point (center point) from each sub-surface to calculate OHD. The similarity is then calculated to decide whether to accept or reject the proposed surface. If it is accepted, then the surface is divided into more sub-surfaces and more points are sampled at a higher level to obtain more accurate similarity. Otherwise, a new surface is proposed.




Figure 5: *A square surface to illustrate the coarse-to-fine model dividing. From left to right, the dividing level is 0, 1 and 2, respectively. The black dots represents the points sampled in the current level, and the white dots represents the points sampled in previous levels.*

5.3 Pseudo Code

The pseudo code of our MH-PPGMF method is presented in Algorithm 1, where $p_\eta(\cdot)$ denotes the posterior computed at dividing level η . The minimal model dividing resolution δ should be set as small as possible to obtain accurate similarity. To achieve better performance, parallel tempering with the same configuration as in [Talton et al. 2011] is also used. That is, 10 Markov chains are run with different temperatures and the chains are randomly swapped.

6 Results

We implemented our method in C++ and conducted our experiments on a machine running Ubuntu 14.04 with Intel Core i5-3470 3.20GHz CPU and 12GB RAM. In all experiments, we set $\epsilon = 10^{-8}$, $\beta = 0.8$, $\sigma = 0.05$. δ should be at least 2 times smaller than query resolution.

6.1 Metric Comparison

We compare several metrics with our WMM metric by fitting 4 models (Fig. 6) to 4 queries (Fig. 7). Model $M_1^x \in \mathcal{M}_1$ is a ring-like surface between an outer square and an inner square. The outer and inner squares have the same center. The length of the outer and inner squares are 4 and $2x$, respectively. Models $M_2^x \in \mathcal{M}_2$, $M_3^x \in \mathcal{M}_3$ and $M_4^x \in \mathcal{M}_4$ are 0.75, 0.5 and 0.25 part of M_1^x , respectively. As shown in Fig. 7, for $i = 1$ to 4, the ground-truth model of Q_i is $M_i^{x=1}$. In this paper, we refer to the target model of a query as the model which is partially similar to the ground-truth model. Therefore, for each query in Fig. 7, there is an target model

Algorithm 1 MH-PPGMF with early rejection

input: a set of modeling rules with parameter \mathbf{x} , query Q , posterior function $p(\mathbf{x}|Q)$, computational budget I , standard deviation ratio σ , local move probability β , and minimal model dividing resolution δ .

output: a maximum a posteriori estimate of \mathbf{x} : \mathbf{x}^* .

Randomly initialize \mathbf{x}^0 , $\mathbf{x}^* \leftarrow \mathbf{x}^0$

for $i = 0$ to I **do**

Randomly select a parameter $x \in \mathbf{x}^i$

Sample $t \sim [0, 1]$

if $t < \beta$ **and** x is continuous **then**

Sample $\tilde{x} \sim \mathcal{N}(x^i, \sigma_x^2)$

else Sample $\tilde{x} \sim [x_{min}, x_{max}]$

Compute η_{top} of $M^{\tilde{x}}$ according to δ

for $\eta = 0$ to η_{top} **do**

$\alpha \leftarrow \min \left\{ 1, \frac{p_\eta(\tilde{x}|Q)}{p_\eta(x^i|Q)} \right\}$

Sample $t \sim [0, 1]$

if $t < \alpha$ **then** $\mathbf{x}^{i+1} \leftarrow \tilde{x}$

else $\mathbf{x}^{i+1} \leftarrow \mathbf{x}^i$, **break**

if $p(\mathbf{x}^{i+1}|Q) > p(\mathbf{x}^*|Q)$ **then** $\mathbf{x}^* \leftarrow \mathbf{x}^{i+1}$

existing in each model space (as shown in Fig. 6). That is, for $i = 1$ to 4 and $j = 1$ to 4, the target model of Q_i in \mathcal{M}_j is $M_j^{x=1}$.




Figure 7: *Queries. From left to right: Queries Q_1 , Q_2 , Q_3 and Q_4 . For $i = 1$ to 4, Q_i is a point cloud uniformly sampled from Model $M_i^{x=1}$ with 0.02 resolution. Q_1 , Q_2 , Q_3 and Q_4 consist of 12288, 9216, 6144 and 3072 points, respectively.*

The metrics used for comparison include negative SHD (-SHD), negative VD (-VD), negative OHD from query to model (-OHDQM), and inlier ratio (IR). VD has been used in [Talton et al. 2011] [Ritchie et al. 2015], while OHDQM has been used in [Ullrich et al. 2008]. As the foundation of many BGMF methods such as [Fischler and Bolles 1981] and [Isack and Boykov 2011], IR is defined as:

$$s_{IR}(M, Q) = \frac{z(Q \cap M)}{z(Q)} \quad (13)$$

where $z(\cdot)$ denotes the size of discrete point set. The comparison results of fitting the models (Fig. 6) to the queries (Fig. 7) are shown in Fig. 8. To compute SHD, OHDQM and WMM, we uniformly sample points from the models with a resolution of 0.01, which is half of the query resolution. In these 16 experiments, the weighting factor h for WMM calculation is 2.5, and the resolution for VD calculation is 0.04.

As the target models of the queries are models with $x = 1$, it is expected that the models with $x = 1$ have the largest similarities. As shown in Fig. 8, our WMM is the only metric to achieve this goal for all experiments. SHD is successful for full fitting (Figs. 8a, 8f, 8k and 8p), but failed for partial fitting except Fig. 8g. IR is failed to distinguish models with $x < 1$ for all experiments except Fig. 8k. The total computational time of these 16 experiments is shown in Fig. 9a, it can be seen that WMM is faster than SHD.




Figure 6: Model spaces. From left to right: Model spaces \mathcal{M}_1 , \mathcal{M}_2 , \mathcal{M}_3 and \mathcal{M}_4 . Each of these 4 spaces has only one parameter $x \in [0, 2]$.

It is worth noting that SHD, OHDQM and WMM always prefer to sample points from model with smaller resolution to obtain more accurate similarity. However, VD produces worse results with smaller resolution for a discrete query. Fine voxelization of a discrete query produces more empty voxels. Therefore, an empty model (e.g. $M_1^{x=2}$) is preferred, of that one example is shown in Fig. 9b. This indicates that voxelization is not suitable for the fine fitting of point clouds.

It is interesting to find that Fig. 8c resembles Fig. 8d. Actually, before normalizing, the original similarities are different. As shown in Table 2, the WMM similarities are comparable across different model spaces for the same query. This table along with Fig. 8 demonstrates that, a model is most similar to itself than any other models using WMM. Finally, we take the experiment of fitting M_1^x to Q_2 as an example to evaluate the effect of weighting factor h . As shown in Fig. 9c, WMM is very stable with respect to different values of h .

Model	$M_1^{x=1}$	$M_2^{x=1}$	$M_3^{x=1}$	$M_4^{x=1}$
Query				
Q_1	1056.4	792.288	528.195	264.094
Q_2	239.651	792.288	528.195	264.094
Q_3	96.4585	113.127	528.195	264.094
Q_4	32.6646	34.4811	44.3055	264.094

Table 2: WMM similarities between the target models and queries. The diagonal elements represent the similarities between the ground-truth models and queries. It is shown that, for the same query, the similarity between the ground-truth model and the query is the largest among all similarities.

6.2 Fitting Noisy Data

We evaluate our method against uniform and Gaussian noise by fitting a sphere model \mathcal{M}_5 to 4 queries (Fig. 10). We use the method in [Marsaglia 1972] to sample points from the sphere surface and replace γ (see Eq. (12)) by $2\pi R$, where R is the radius of the sphere. The sphere model \mathcal{M}_5 has 4 parameters, including 3 location parameters and 1 radius parameter. Some sample models of \mathcal{M}_5 are shown in the top row of Fig. 11. It is worth noting that all the model parameters involved in this paper are uniformly distributed. As a result, the optimization objective is reduced from posterior to likelihood and then rigid similarity. Besides, we stop model fitting casually in this paper.

The resolution of the noise-free query Q_5 (Fig. 10a) is 0.2. We set $\delta = 0.04$ and $h = 10$ in these sphere fitting experiments, the results are presented in Figs. 11 and 12. As shown in Figs. 12a and 12b, the target similarities (log likelihood) still remain the largest similarities after sufficient evolution time. This indicates that our




Figure 10: Queries. From left to right: Queries Q_5 , Q_6 , Q_7 and Q_8 . Q_5 is a noise-free point cloud consisting of 549 points sampled from a unit sphere surface. Q_6 and Q_7 are generated by adding low-level and high-level uniform noise to Q_5 , respectively. Q_6 consists of 1098 points, while Q_7 consists of 2985 points. The uniform noise is distributed within a cube with a length of 2. The cube and the unit sphere have the same center. Q_8 is generated by adding Gaussian noise with a standard deviation of 0.2 to Q_5 .




Figure 11: Fitted sphere models (blue) along with the queries (green, Fig. 10). Top row: randomly initialized models. Bottom row: final fitted models (Fig. 12). From left to right: fitting the sphere to Q_5 , Q_6 , Q_7 and Q_8 , respectively.

method is robust to uniform noise. Similarly, Fig. 12c shows that our method is also robust to Gaussian noise. Our method can only be slightly affected by Gaussian noise.

However, uniform noise influences the efficiency of our method. Particularly, tens of minutes have been consumed to generate the results shown in Figs. 12a and 12c, however, several hours have been consumed to obtain a desirable model shown in Fig. 12b. This is because high-level uniform noise introduces many local maxima to the objective function, making it difficult to find the global max-

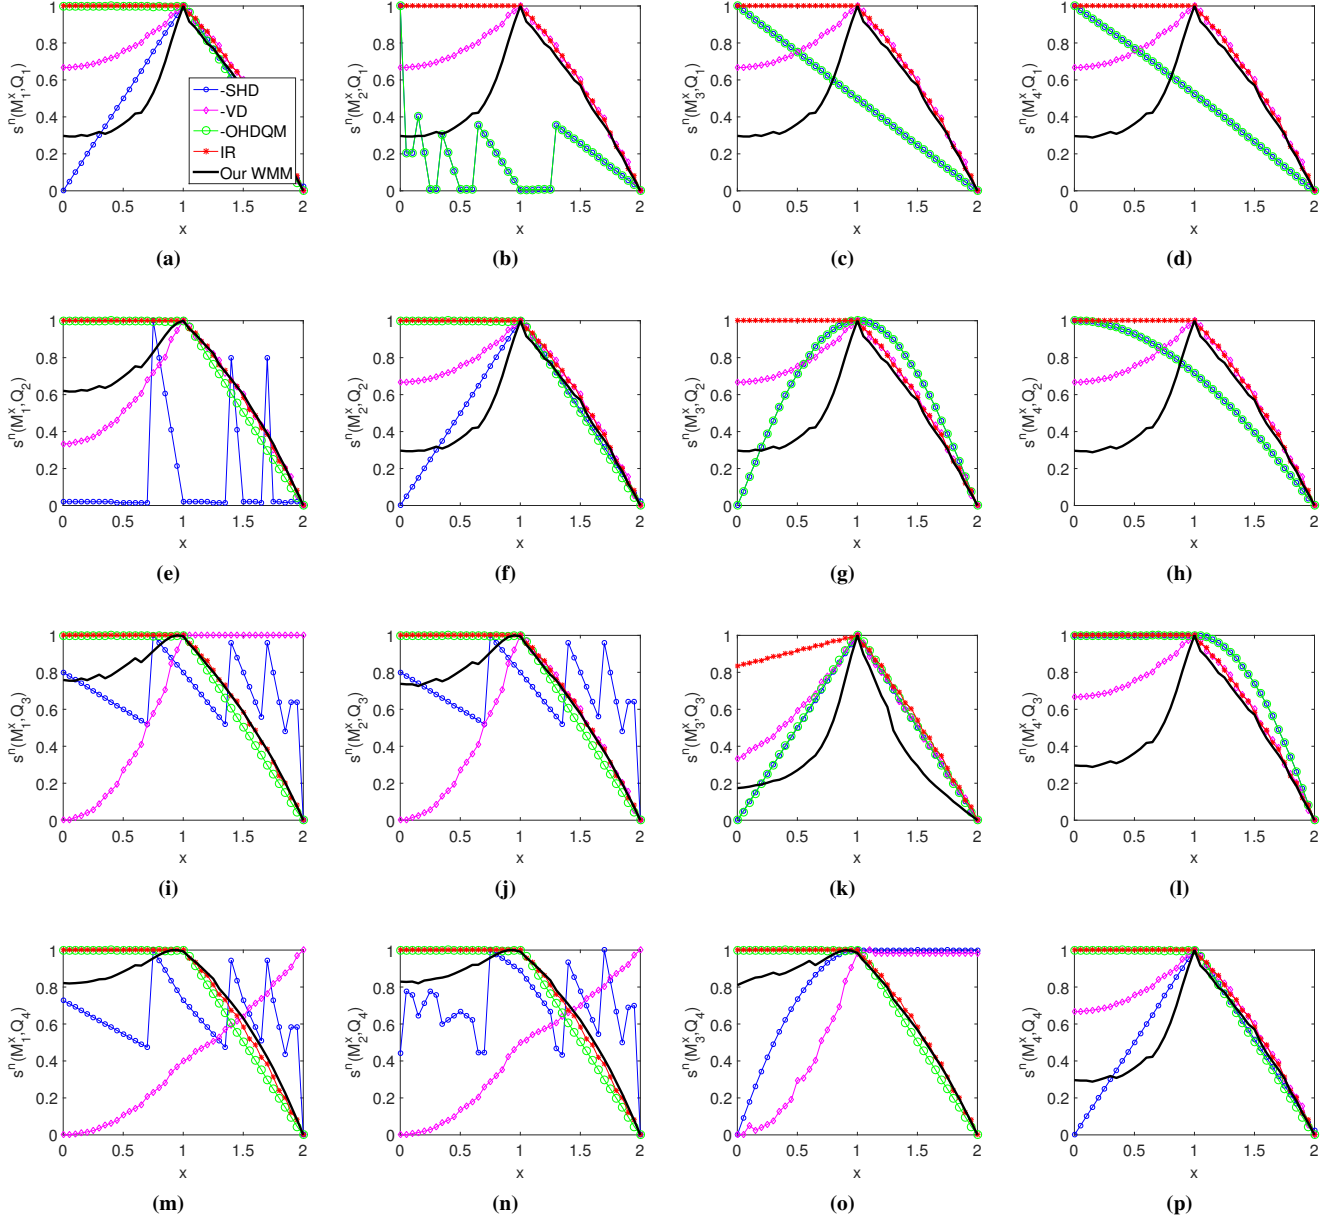


Figure 8: Metric comparison results. From left to right: the results of fitting Models M_1^x , M_2^x , M_3^x and M_4^x to the queries. From top to bottom: the results of fitting the models to Queries Q_1 , Q_2 , Q_3 and Q_4 . The vertical axis $s^n(\cdot, \cdot)$ denotes the normalized similarity. We uniformly normalize the similarities into a range of $[0, 1]$. The legend for these figures are presented in (a). The figures in diagonal, above diagonal and below diagonal represent the results of full fitting on complete data, partial fitting on over-complete data and partial fitting on incomplete data, respectively.

imum. The IPS indicator in Fig. 12 shows the influence of our early rejection strategy. It can be observed that, the optimization process is accelerated by about 3 times using early rejection. These experiments additionally demonstrate that our method is able to fit non-planar models.

6.3 Fitting Models with length-varying parameters

Fitting a model with varying number of parameters is more difficult than fitting a model with a fixed number of parameters. In this paper, we investigate two models \mathcal{M}_6 and \mathcal{M}_7 (Figs. 13 and 14) with

varying numbers of parameters, which are based on the CGA shape grammar [Müller et al. 2006]. \mathcal{M}_6 and \mathcal{M}_7 are models of buildings. The models in \mathcal{M}_6 consist of 4 facades, while the models in \mathcal{M}_7 consist of 1 facade. Let n be the number of floors, \mathcal{M}_6 has $7 + 2n$ parameters, i.e., 1 parameter for rotation, 3 parameters for location, 3 parameters for mass size (height, length and width), $2n$ parameters for window size. n depends on the height of building. Different from \mathcal{M}_6 , \mathcal{M}_7 does not have the width parameter. Windows in the same floor have the same size, but may have different sizes on different floors. Some sample models of \mathcal{M}_6 and \mathcal{M}_7 are shown in Fig. 14.




Figure 9: (a) Total time for similarity calculation on experiments shown in Fig. 8. (b) -VD similarities of fitting M_1^x to Q_1 with resolutions 0.2, 0.08, 0.02, 0.01 and 0.005. (c) WMM similarities of fitting M_1^x to Q_2 with weighting factor $h=0.5, 1, 2, 4$ and 8.

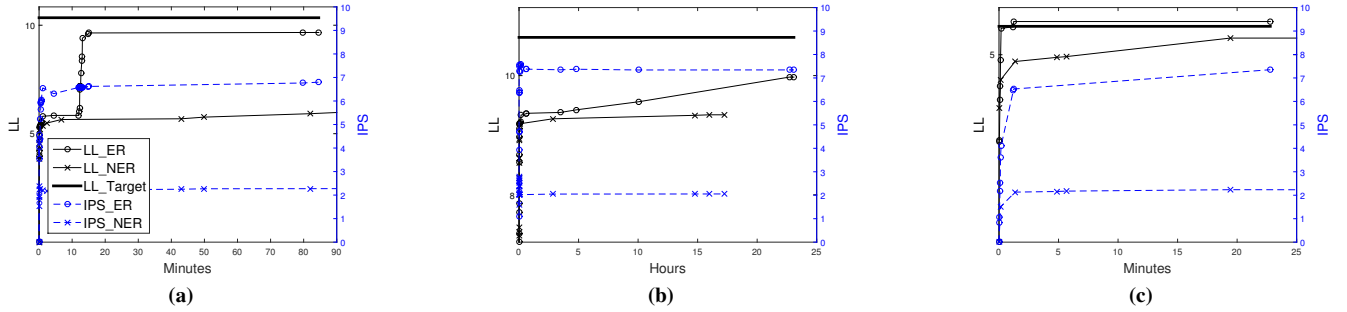


Figure 12: Sphere fitting results. (a), (b) and (c) are the results of fitting the sphere model M_5 to Q_6, Q_7 and Q_8 (Fig. 10), respectively. LL, IPS and ER denote log likelihood, iterations per second, and early rejection, respectively. LL_ER, LL_NER, IPS_ER and IPS_NER denote the evolutions of LL with ER, LL without ER, IPS with ER, and IPS without ER, respectively. LL_Target denotes the target log likelihood.




Figure 13: Queries and target models. (a) Query Q_9 , (b) Query Q_{10} . (c) and (d) are models in Model spaces M_6 and M_7 , respectively. Both Q_9 and Q_{10} are part of the point cloud which is uniformly sampled from (c) with resolution 0.2. (d) is a part of (c). Consequently, (c) and (d) are target models of Q_9 (and Q_{10}) in M_6 and M_7 , respectively. Q_9 consists of 4204 points, while Q_{10} consists of 2452 points.




Figure 14: Some sample models of M_6 or M_7 along with Q_9 or Q_{10} .

The results of fitting M_6 to Q_9 (Fig. 13a) and Q_{10} (Fig. 13b), fitting M_7 to Q_{10} are shown in Figs. 15 and 16. In these 3 experiments, we set $h = 2.5, \delta = 0.1$. As shown in Fig. 15, after 5000 iterations, the mass parameters are correctly estimated, while the window parameters are incorrectly estimated. Finally, all the

parameters are correctly estimated. These 3 experiments show that our method is able to fit incomplete data with holes. The experiment of fitting M_7 to Q_{10} additionally demonstrates that our method is able to fit hybrid-complete data.




Figure 15: Fitted models. From left to right: results of fitting \mathcal{M}_6 to Q_9 , \mathcal{M}_6 to Q_{10} , and \mathcal{M}_7 to Q_{10} . From top to bottom: the randomly initialized models, fitted models after 5000 iterations, and final fitted models.

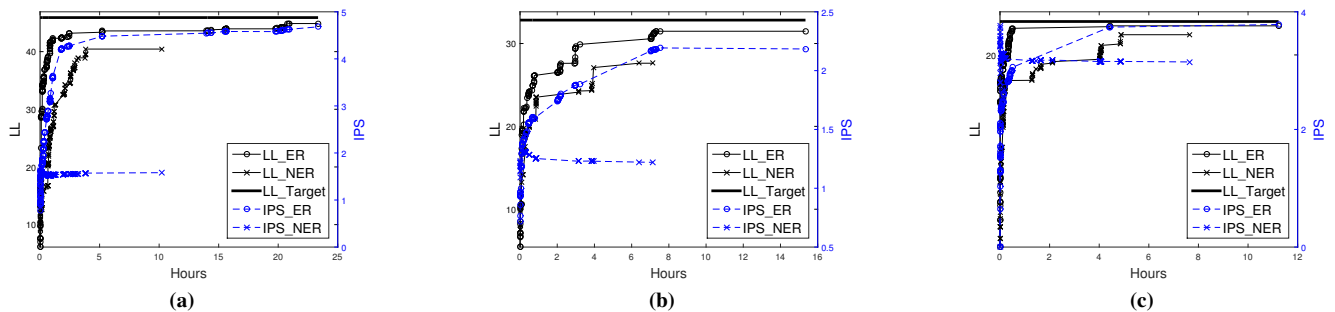


Figure 16: Fitting results. (a), (b) and (c) are results of fitting \mathcal{M}_6 to Q_9 , fitting \mathcal{M}_6 to Q_{10} , and fitting \mathcal{M}_7 to Q_{10} , respectively.

6.4 Fitting Laser Scanning Data

We also conducted experiments for fitting real-world laser scanning point clouds, which were collected by mobile laser scanners [Guan et al. 2014] [Yu et al. 2016]. The results of fitting a facade model \mathcal{M}_8 to a facade point cloud Q_{11} (Fig. 17b) are shown in Figs. 17 and 1. \mathcal{M}_8 has 15 parameters. Instead of WMM, we use SMM (Eq. (6)) to perform this experiment and set $\delta = 0.08$. In contrast to WMM, no argument h is needed to calculate SMM. We tested WMM and found that WMM is not very effective for Q_{11} due to that the holes in Q_{11} are corrupted. That is, there are undesired points within the holes. These corrupted holes are incorrectly recognized as missing data by WMM. In other words, although WMM is able to distinguish the difference between uncorrupted holes and missing data (as shown in Section 6.3), WMM is unable to distinguish the difference between corrupted holes and missing data. Fortunately, as shown in the results, SMM is able to deal with corrupted holes.

7 Conclusion

We have proposed a novel rigid geometric similarity metric to measure the similarity between geometric models. Based on the proposed metric, we presented the first method to rigidly fit arbitrary geometric model to non-complete data. We formulate the fitting problem as a Bayesian inference problem and employ MCMC technique to perform the inference. We also proposed a novel technique to accelerate the inference process. Our method has been demonstrated on various geometric models and non-complete data. Experimental results show that our metric is effective for fitting non-complete data, while other metrics are ineffective. It is also shown that our method is robust to noise. In summary, our method is able to fit non-complete data without holes (Section 6.2), non-complete data with uncorrupted holes (Section 6.3) or over-complete data with corrupted holes (Section 6.4).

We believe our work bridges the gap between inverse procedural modeling and geometric model fitting. However, several issues still remain open. For example, due to the curse of dimensionality, the fitting problem becomes intractable if the geometric model has a




Figure 17: Fitted model. (a) An original point cloud consisting of 385793 points. (b) A query Q_{11} consisting of 23266 points. Q_{11} is generated by downsampling the original point cloud (a) with a resolution of 0.2. (c) Final fitted model (after 84540 iterations and 36146.0820 seconds) for fitting \mathcal{M}_8 to Q_{11} . (d) Overlap between the query Q_{11} and the final model shown in (c).

large number of parameters. New techniques such as deep learning [Nishida et al. 2016] [Ritchie et al. 2016] are expected to address this problem. Besides, it is also challenging for our method to fit incomplete data with corrupted holes, because a corrupted hole may incorrectly be recognized as missing data. This problem may be addressed using some preprocessing process, e.g., filtering out the undesired points within each hole.

8 Acknowledgements

This work was supported by the National Natural Science Foundation of China under Grants 41471379, 61602499 and 61471371, and by Fujian Collaborative Innovation Center for Big Data Applications in Governments. The authors would like to thank the reviewers for their comments.

References

- ASPERT, N., SANTA CRUZ, D., AND EBRAHIMI, T. 2002. MESH: Measuring errors between surfaces using the Hausdorff distance. *IEEE International Conference in Multimedia and Expo*, 705–708.
- BOKELOH, M., WAND, M., AND SEIDEL, H. 2010. A connection between partial symmetry and inverse procedural modeling. *ACM Transactions on Graphics* 29, 4, 104.
- BOULCH, A., HOULLIER, S., MARLET, R., AND TOURNAIRE, O. 2013. Semantizing complex 3D scenes using constrained attribute grammars. *Computer Graphics Forum* 32, 5, 33–42.
- DANG, M., LIENHARD, S., CEYLAN, D., NEUBERT, B., WONKA, P., AND PAULY, M. 2015. Interactive design of probability density functions for shape grammars. *ACM Transactions on Graphics* 34, 6, 206.
- DEBEVEC, P., TAYLOR, C. J., AND MALIK, J. 1996. Modeling and rendering architecture from photographs: a hybrid geometry- and image-based approach. *International Conference on Computer Graphics and Interactive Techniques*.
- DEMIR, I., ALIAGA, D. G., AND BENES, B. 2015. Procedural editing of 3D building point clouds. *ICCV*.
- FISCHLER, M. A., AND BOLLES, R. C. 1981. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM* 24, 6, 381–395.
- GUAN, H., LI, J., YU, Y., WANG, C., CHAPMAN, M. A., AND YANG, B. 2014. Using mobile laser scanning data for automated extraction of road markings. *ISPRS Journal of Photogrammetry and Remote Sensing* 87, 93–107.
- GUO, Y., BENNAMOUN, M., SOHEL, F., LU, M., AND WAN, J. 2014. 3D object recognition in cluttered scenes with local surface features: a survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 36, 11, 2270–2287.
- HASTINGS, W. K. 1970. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika* 57, 1, 97–109.
- HUANG, H., BRENNER, C., AND SESTER, M. 2013. A generative statistical approach to automatic 3D building roof reconstruction from laser scanning data. *ISPRS Journal of Photogrammetry and Remote Sensing* 79, 29–43.
- ISACK, H., AND BOYKOV, Y. 2011. Energy-based geometric multi-model fitting. *International Journal of Computer Vision* 97, 2, 123–147.
- LAFARGE, F., AND MALLET, C. 2012. Creating large-scale city models from 3D-point clouds: A robust approach with hybrid representation. *International Journal of Computer Vision* 99, 1, 69–85.
- LAFARGE, F., DESCOMBES, X., ZERUBIA, J., AND PIERROT-DESEILLIGNY, M. 2010. Structural approach for building reconstruction from a single DSM. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 32, 1, 135–147.
- LAKE, B. M., SALAKHUTDINOV, R., AND TENENBAUM, J. B. 2015. Human-level concept learning through probabilistic program induction. *Science* 350, 6266, 1332–1338.
- LI, Y., WU, X., CHRYSATHOU, Y., SHARF, A., COHENOR, D., AND MITRA, N. J. 2011. GlobFit: consistently fitting primitives by discovering global relations. *ACM Transactions on Graphics* 30, 4, 52.
- LIN, H., GAO, J., ZHOU, Y., LU, G., YE, M., ZHANG, C., LIU, L., AND YANG, R. 2013. Semantic decomposition and reconstruction of residential scenes from LiDAR data. *ACM Transactions on Graphics* 32, 4, 66.
- LIN, Y., WANG, C., CHENG, J., CHEN, B., JIA, F., CHEN, Z., AND LI, J. 2015. Line segment extraction for large scale unorganized point clouds. *ISPRS Journal of Photogrammetry and Remote Sensing* 102, 172–183.
- MARSAGLIA, G. 1972. Choosing a point from the surface of a sphere. *Annals of Mathematical Statistics* 43, 2, 645–646.
- MATHIAS, M., MARTINOVIĆ, A., WEISSENBERG, J., AND VAN GOOL, L. 2011. Procedural 3D building reconstruction using shape grammars and detectors. *International Conference on 3D Imaging, Modeling, Processing, Visualization and Transmission*, 304–311.

- METROPOLIS, N., ROSENBLUTH, A. W., ROSENBLUTH, M. N., TELLER, A. H., AND TELLER, E. 1953. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics* 21, 6, 1087–1092.
- MONSZPART, A., MELLADO, N., BROSTOW, G. J., AND MITRA, N. J. 2015. RApTer: rebuilding man-made scenes with regular arrangements of planes. *ACM Transactions on Graphics* 34, 4, 103.
- MUJA, M., AND LOWE, D. G. 2014. Scalable nearest neighbor algorithms for high dimensional data. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 36, 11, 2227–2240.
- MÜLLER, P., WONKA, P., HAEGLER, S., ULMER, A., AND VAN GOOL, L. 2006. Procedural modeling of buildings. *ACM Transactions on Graphics* 25, 3, 614–623.
- MUSIALSKI, P., WONKA, P., ALIAGA, D. G., WIMMER, M., GOOL, L., AND PURGATHOFER, W. 2013. A survey of urban reconstruction. *Computer Graphics Forum* 32, 6, 146–177.
- NAN, L., SHARF, A., ZHANG, H., COHENOR, D., AND CHEN, B. 2010. SmartBoxes for interactive urban reconstruction. *ACM Transactions on Graphics* 29, 4, 93.
- NISHIDA, G., GARCIA-DORADO, I., ALIAGA, D. G., BENES, B., AND BOUSSEAU, A. 2016. Interactive sketching of urban procedural models. *ACM Transactions on Graphics*.
- POULLIS, C. 2013. A framework for automatic modeling from point cloud data. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35, 11, 2563–2575.
- PU, S., AND VOSSelman, G. 2009. Knowledge based reconstruction of building models from terrestrial laser scanning data. *ISPRS Journal of Photogrammetry and Remote Sensing* 64, 6, 575–584.
- RITCHIE, D., MILDENHALL, B., GOODMAN, N. D., AND HANRAHAN, P. 2015. Controlling procedural modeling programs with stochastically-ordered sequential Monte Carlo. *ACM Transactions on Graphics* 34, 4, 105.
- RITCHIE, D., THOMAS, A., HANRAHAN, P., AND GOODMAN, N. D. 2016. Neurally-guided procedural models: Learning to guide procedural models with deep neural networks. *arXiv preprint arXiv:1603.06143*.
- SMELIK, R. M., TUTENEL, T., BIDARRA, R., AND BENES, B. 2014. A survey on procedural modelling for virtual worlds. *Computer Graphics Forum* 33, 6, 31–50.
- STAVA, O., PIRK, S., KRATT, J., CHEN, B., MĚCH, R., DEUSSEN, O., AND BENES, B. 2014. Inverse procedural modelling of trees. *Computer Graphics Forum* 33, 6, 118–131.
- TALTON, J. O., LOU, Y., LESSER, S., DUKE, J., MCH, R., AND KOLTUN, V. 2011. Metropolis procedural modeling. *ACM Transactions on Graphics* 30, 2, 11.
- TEBOUL, O., KOKKINOS, I., SIMON, L., KOUTSOURAKIS, P., AND PARAGIOS, N. 2013. Parsing facades with shape grammars and reinforcement learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35, 7, 1744–1756.
- TOSHEV, A., MORDOHAI, P., AND TASKAR, B. 2010. Detecting and parsing architecture at city scale from range data. *CVPR*, 398–405.
- ULLRICH, T., SETTGAST, V., AND FELLNER, D. W. 2008. Semantic fitting and reconstruction. *Journal on Computing and Cultural Heritage* 1, 2.
- VANEGAS, C. A., ALIAGA, D. G., AND BENES, B. 2012. Automatic extraction of Manhattan-world building masses from 3D laser range scans. *IEEE Transactions on Visualization and Computer Graphics* 18, 10, 1627–1637.
- VANEGAS, C. A., GARCIA-DORADO, I., ALIAGA, D. G., BENES, B., AND WADDELL, P. 2012. Inverse design of urban procedural models. *ACM Transactions on Graphics* 31, 6, 168.
- WAN, G., AND SHARF, A. 2012. Grammar-based 3D facade segmentation and reconstruction. *Computers & Graphics* 36, 4, 216–223.
- WANG, J., AND KAI XU, K. 2016. Shape detection from raw lidar data with subspace modeling. *IEEE Transactions on Visualization and Computer Graphics*.
- WEISSENBERG, J., RIEMENSCHNEIDER, H., PRASAD, M., AND GOOL, L. 2013. Is there a procedural logic to architecture? *CVPR*, 185–192.
- YU, Y., LI, J., WEN, C., GUAN, H., LUO, H., AND WANG, C. 2016. Bag-of-visual-phrases and hierarchical deep models for traffic sign detection and recognition in mobile laser scanning data. *ISPRS Journal of Photogrammetry and Remote Sensing* 113, 106–123.
- ZHENG, Q., SHARF, A., WAN, G., LI, Y., MITRA, N. J., COHENOR, D., AND CHEN, B. 2010. Non-local scan consolidation for 3D urban scenes. *ACM Transactions on Graphics* 29, 4, 94.