



HAL
open science

Intuitive and efficient roof modeling for reconstruction and synthesis

Jing Ren, Biao Zhang, Bojian Wu, Jianqiang Huang, Lubin Fan, Maks Ovsjanikov, Peter Wonka

► **To cite this version:**

Jing Ren, Biao Zhang, Bojian Wu, Jianqiang Huang, Lubin Fan, et al.. Intuitive and efficient roof modeling for reconstruction and synthesis. ACM Transactions on Graphics, 2021, 40 (6), pp.1-17. 10.1145/3478513.3480494 . hal-04479835

HAL Id: hal-04479835

<https://hal.science/hal-04479835>

Submitted on 27 Feb 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Intuitive and Efficient Roof Modeling for Reconstruction and Synthesis

JING REN*, KAUST and Alibaba Group, Saudi Arabia

BIAO ZHANG, KAUST, Saudi Arabia

BOJIAN WU, JIANQIANG HUANG, LUBIN FAN, Alibaba Group, China

MAKS OVSJANIKOV, LIX, École Polytechnique, France

PETER WONKA, KAUST, Saudi Arabia

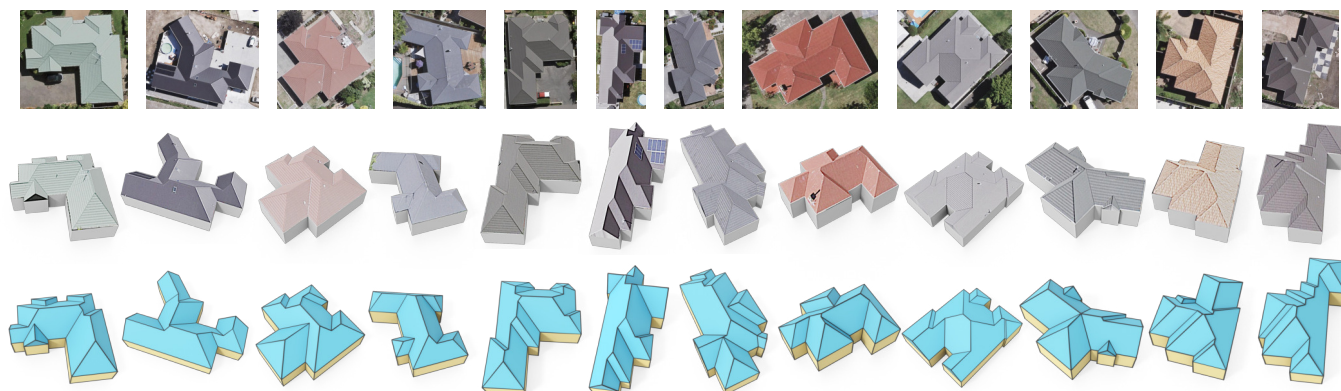


Fig. 1. Our method can be used to reconstruct buildings from aerial images. *Top*: aerial images, on which a user can use our UI to specify the roof topology. *Middle*: the reconstructed buildings obtained by our geometric optimization (textured by the input image). *Bottom*: 3D geometry of the reconstructed buildings.

We propose a novel and flexible roof modeling approach that can be used for constructing planar 3D polygon roof meshes. Our method uses a graph structure to encode roof topology and enforces the roof validity by optimizing a simple but effective planarity metric we propose. This approach is significantly more efficient than using general purpose 3D modeling tools such as 3ds Max or SketchUp, and more powerful and expressive than specialized tools such as the straight skeleton. Our optimization-based formulation is also flexible and can accommodate different styles and user preferences for roof modeling. We showcase two applications. The first application is an interactive roof editing framework that can be used for roof design or roof reconstruction from aerial images. We highlight the efficiency and generality of our approach by constructing a mesh-image paired dataset consisting of 2539 roofs. Our second application is a generative model to synthesize new roof meshes from scratch. We use our novel dataset to combine machine learning and our roof optimization techniques, by using transformers and graph convolutional networks to model roof topology, and our roof optimization methods to enforce the planarity constraint.

*Work was done during the author's research internship at Alibaba Group

Authors' addresses: Jing Ren, KAUST and Alibaba Group, Saudi Arabia, jing.ren@kaust.edu.sa; Biao Zhang, KAUST, Saudi Arabia, biao.zhang@kaust.edu.sa; Bojian Wu, Jianqiang Huang, Lubin Fan, Alibaba Group, China, ustcbjwu@gmail.com, jianqiang.hjq@alibaba-inc.com, lubinfan@gmail.com; Maks Ovsjanikov, LIX, École Polytechnique, France, maks@lix.polytechnique.fr; Peter Wonka, KAUST, Saudi Arabia, pwonka@gmail.com.

Table 1. Different solutions for roof modeling.

Property	straight skeleton	weighted straight skeleton	commercial software	Ours
<i>Easy to use for beginners?</i>	✓	✓	✗	✓
<i>Efficient for roof construction?</i>	✓	✓	✗	✓
<i>Accurately reconstruct roofs from images?</i>	✗	✗	✓	✓
<i>Light user input?</i>	✓	✓	✗	✓
<i>Allow editing operations?</i>	✗	✓	✓	✓
<i>Intuitive editing operations?</i>	✗	✗	✓	✓
<i>Insensitive to noisy user input?</i>	✗	✗	✓	✓

CCS Concepts: • **Computing methodologies** → **Shape modeling; Shape mesh modeling.**

Additional Key Words and Phrases: Roof modeling, Optimization, Interactive editing, Roof synthesis

ACM Reference Format:

Jing Ren, Biao Zhang, Bojian Wu, Jianqiang Huang, Lubin Fan, Maks Ovsjanikov, and Peter Wonka. 2021. Intuitive and Efficient Roof Modeling for Reconstruction and Synthesis. *ACM Trans. Graph.* 40, 6, Article 249 (December 2021), 17 pages. <https://doi.org/10.1145/3478513.3480494>

1 INTRODUCTION

Roof modeling is an important topic in urban reconstruction. In our work, we propose a novel and simple formulation for roof modeling which is expressive enough to handle a large range of roofs. Our formulation is also suitable for two applications: interactive reconstruction of roofs from aerial images (See Fig. 1) and the generative modeling of new roofs.

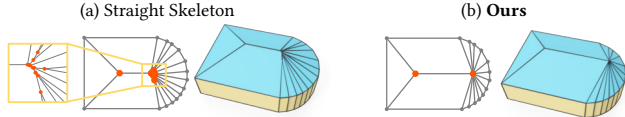


Fig. 2. For a given outline with 16 vertices, we run the straight skeleton method to obtain the 2D roof embedding and construct the corresponding building as shown in (a). The straight skeleton computes a planar but complicated 3D roof with 12 roof vertices (colored in red). As a comparison, our method constructs a more plausible roof with 2 roof vertices.

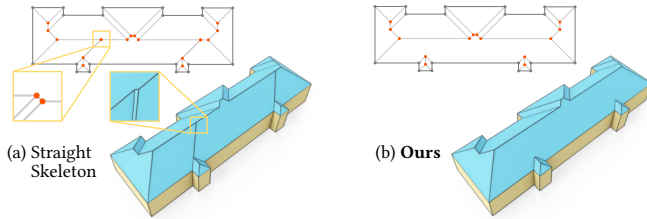


Fig. 3. The straight skeleton method cannot handle the cases where there exists a face that contains multiple outline edges. *Top*: we show the roof graph computed by the straight skeleton and our method. *Bottom*: we show the corresponding constructed roofs.

The main challenges of roof modeling are how to mathematically describe the roof structure and how to enforce the planarity of the roof faces. A popular tool for roof construction is to use the straight skeleton [Aichholzer and Aurenhammer 1996; Aichholzer et al. 1996] or one of its extensions to increase its modeling expressiveness [Biedl et al. 2015; Eppstein and Erickson 1999; Kelly and Wonka 2011]. These methods typically use a closed roof outline to describe the roof structure assuming that the interior topology can be determined by the roof outline. The roof face planarity is enforced during the computation of roof topology from the input outline. Another solution is to use general purpose commercial software such as 3ds Max and SketchUp to model roofs. Though the commercial software can construct large variations of realistic roofs, the roof planarity is either ignored (e.g., 3ds Max) or implicitly enforced by triangulation (SketchUp).

We observe that these solutions have limitations in different aspects (see Table 1). For example, using only the outline for roof structure specification is simple but not sufficient for roof modeling, since different roofs can have the same outline. Moreover, determining the roof topology from the outline can be error-prone. For example, for some outlines, the straight skeleton based methods create spurious additional vertices close-by (Fig. 2), and fail to recover the correct roof topology when there exists a face corresponding to multiple outline edges (Fig. 3). On the other hand, commercial software provides large freedom, but modeling is significantly more complicated especially for non expert users and it is not easy to enforce geometric constraints.

In our work, we propose to use a two-step procedure where we first model the roof topology and optionally the approximate geometry and then refine the geometry using optimization. Specifically, we propose to use a *roof graph* to specify roof topology which is simple and flexible enough to represent a large range of roofs including residential buildings (Fig. 1) and architecture with different styles (Fig. 4). We then propose an optimization-based method for roof

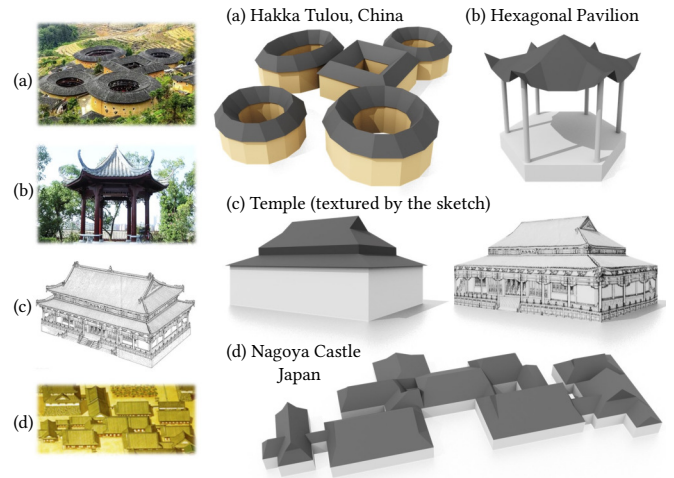


Fig. 4. Building meshes of Asian architecture created by our method. The reference images on the left are collected from internet.

modeling from an input roof graph, where we introduce a simple planarity metric. Our method is generic and can be adapted to different settings such as including user-specified regularizers. Compared to the straight skeleton based methods, our solution has stronger expressiveness with fewer assumptions placed on the underlying roofs. Meanwhile, our solution is more robust and can better reconstruct roofs from an image with higher accuracy. Compared to the general purpose 3D modeling software, our method has a more natural and systematic roof structure representation and is explicitly designed to output planar 3D *polygon* roofs. Our roof modeling framework can be easily used by novices requiring light user input.

Our method has two practical applications, interactive roof editing and roof synthesis from scratch. Specifically, our roof graph representation allows different editing operations for modifying roof topology, while our roof optimization efficiently updates the roof embedding by enforcing planarity constraints. These features allow a user to interactively model a planar roof by iteratively modifying roof topology and optimizing the roof planarity. Another useful and novel application is to automatically synthesize realistic roofs. Roof synthesis in general is a difficult problem, since it has a mixture of discrete (roof topology) and continuous (roof planarity) constraints. To tackle this issue, we train neural networks to generate a discrete roof graph then run our roof optimization method to enforce the continuous geometric constraints.

In summary, our main contributions are:

- A novel roof modeling method including a simple roof graph representation that encodes roof topology and a new planarity metric that can be used to enforce geometric constraints for generating 3D polygonal roof meshes.
- An optimization-based framework that is complementary to learning algorithms and user input for automatic roof synthesis and interactive editing.
- We created a dataset consisting of 2539 roof meshes paired with images without topological or geometric errors (see

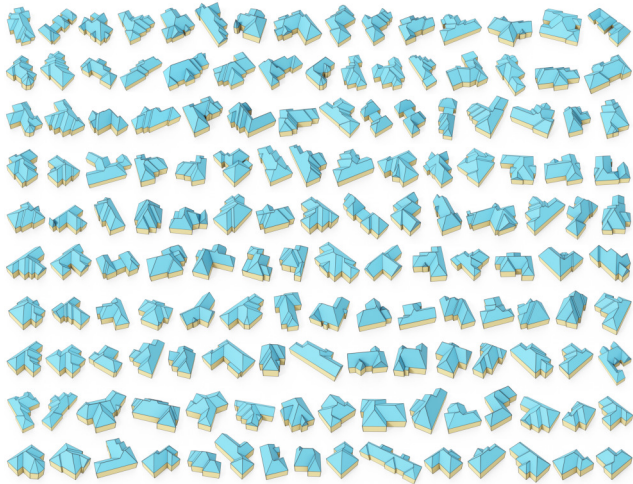


Fig. 5. We created a dataset of 3D planar roofs paired with the corresponding aerial images. Here we show 150 example roofs and highlight the topology. The paired images are shown in Fig. 17 in the supplementary materials.

Fig. 5) using our method, which can be helpful for different visual computing tasks. Code and data are available.¹

2 RELATED WORK

We review related work in three categories: roof construction algorithms, roof reconstruction, and generative modeling.

2.1 Roof Construction Algorithms

In computer graphics, previous work proposes solutions to model various aspects of architectural models, such as the room layout and arrangement [Hu et al. 2020; Merrell et al. 2010] and indoor scene synthesis [Fisher et al. 2015; Yu et al. 2011]. Roofs are one of the components previous works attempted to model using algorithmic or procedural approaches.

The *straight skeleton* algorithm [Aichholzer and Aurenhammer 1996; Aichholzer et al. 1996] is a popular geometric algorithm for the generation of complex roof structures from user-specified roof outlines [Laycock and Day 2003]. The weighted straight skeleton [Biedl et al. 2015; Eder and Held 2018; Eppstein and Erickson 1999] is an extension to improve the modeling expressiveness to facilitate the modeling of roof faces at different angles specified by users. Several large-scale urban modeling projects are inspired by these techniques. Larive and Gaildrat [2006] combine 3D building descriptions, including the height of footprints and roofs, with GIS information to create building models. Kelly and Wonka [2011] propose a subsequent extension to the weighted straight skeleton for the interactive modeling of complete buildings. Besides, Buron et al. [2013] consider to bring parallelism to grammar-based roof generation in order to improve the computational efficiency. Other extensions to the straight skeleton include the works of Sugihara [2013; 2019] and Held and Palfrader [2017]. An alternative approach to modeling complex roofs is to create a combination of elementary roof primitives, e.g. as done in procedural modeling [Müller et al. 2006].

¹Demo Code: https://github.com/lloz/SGA21_roofOptimization

Complementary to our work in roof modeling, several magnificent real-world buildings feature smooth roof shapes. The design of such roofs [Liu et al. 2006; Pottmann et al. 2007, 2008] requires different and specialized tools and should be considered as a separate topic.

2.2 Roof Reconstruction

Urban reconstruction aims at automatically generating 3D models from real physical measurements, such as multi-view images or point clouds [Demir et al. 2015; Musialski et al. 2013]. We focus our review on recent methods that have roof reconstruction as a major component. One category of algorithms uses optimization. For example, Zhou and Neumann published a series of papers to reconstruct coarse building models including roofs from LiDAR point clouds [2008; 2010; 2011]. Lin et al. [2013] propose to find a combination of planar primitives that best explains the input LiDAR data for roof reconstruction. Dehbi et al. [2021] propose an active sampling strategy for RANSAC to fit plane approximations to input point clouds. Nan and Wonka [2017] and Kelly et al. [2017] employ integer programming to reconstruct coarse planar building models. Arikian et al. [2013] propose an initial automatic method to extract candidate planes and estimate coarse polygons from unorganized point clouds, and then allow users to interactively edit the model by optimization-based snapping operations. Verdie et al. [2015] and Zhu et al. [2018] focus on segmenting parts belonging to the roof surface, and then fitting a collection of piece-wise planar planes [Liu et al. 2018] or prior defined templates for roof reconstruction. Habbecke and Kobbelt [2012] propose an interactive geometric modeling system that can be used for polygonal roof editing with geometric regularities or constraints. Salinas et al. [2015] present a mesh decimation approach that generates planar abstractions of roof meshes. Bauchet and Lafarge [2020] adopt a kinetic data structure to partition the 3D space into convex polyhedra from which the underlying surface mesh of the input point cloud can be extracted. Another category of papers use deep learning [Alidoost et al. 2020; Yu et al. 2021; Zeng et al. 2018; Zhang et al. 2020] to directly output the reconstruction of 3D roof structures. However, these methods do not propose a solution for enforcing geometric constraints, while in our work, we mainly focus on enforcing geometric constraints in an optimization-based formulation during interactive roof reconstruction.

2.3 Generative Modeling

Generative models aim to generate new samples that follow a similar distribution as a collection of training samples. For example, generative adversarial networks (GANs) [Goodfellow et al. 2014] can be extended to 3D by synthesizing details on surfaces, e.g., [Kelly et al. 2018] create textures on coarse building models. Volumetric GANs [Wu et al. 2016] can create 3D models with a voxel grid representation. Normalizing flows [Dinh et al. 2015; Rezende and Mohamed 2015] (NF) have been extended to 3D by modeling the distribution of point clouds as an invertible parameterized transformation from a probability density embedded in 3D space [Kim et al. 2020; Stypulkowski et al. 2019; Yang et al. 2019]. In the following, we mainly discuss variational autoencoders (VAEs) and

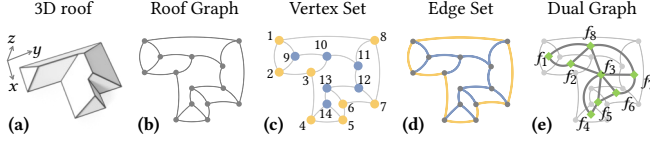


Fig. 6. 3D roof represented as a *roof graph*: (a) a simple 3D roof with 8 faces; (b) the topology of the roof represented as a graph $G = (V, E)$; we highlight the vertex set V in (c) and the edge set E in (d), where the *outline* vertices (edges) are colored in orange, and the *roof* vertices (edges) are colored in blue. (e) We can also construct a dual graph of the roof graph, where each face of the roof is represented as a node, and two nodes are connected to each other if the corresponding roof faces are adjacent.

auto-regressive models (ARs) for 3D tasks that are closely related to our roof synthesis application.

Variational autoencoders [Kingma and Welling 2014] parameterize latent variable models with deep neural networks. Compared to GANs, VAEs are easier to train but they typically produce images of lower visual quality [Razavi et al. 2019; Van Den Oord et al. 2017]. While extending to 3D geometry [Brock et al. 2016], they have the potential to synthesize 3D shapes, especially on domain-specific tasks that require limited topological variations, such as face [Ranjan et al. 2018] or human body [Tan et al. 2018] modeling. VAEs also have been used for generating structured models, such as furniture [Gao et al. 2019; Mo et al. 2019; Yang et al. 2020]. In these instances, there are separate generative models for the individual object parts and the part arrangement.

Auto-regressive models (ARs) factorize a distribution over a sequence into several conditional densities. Each conditional density models a single element in the sequence conditioned on all previous elements. Language models are auto-regressive in nature [Dai et al. 2019; Radford et al. 2019]. Some works also model images as sequences, e.g., PixelRNN [Van Oord et al. 2016] and its follow-up works PixelCNN [Van den Oord et al. 2016], PixelCNN++ [Salimans et al. 2017], and PixelSNAIL [Chen et al. 2018]. With the rise of transformers, autoregressive models have been increasingly applied to 3D data, including polygonal meshes [Nash et al. 2020], floor plans [Para et al. 2020], and scenes [Wang et al. 2020]. We argue that the discrete nature of ARs is a unique advantage for modeling 3D data that does not consist of smoothly varying models, but needs a latent representation that can handle discontinuities. Therefore, we propose to adapt ARs to the generation of roof outlines. We will demonstrate that ARs are more suitable than VAEs (the currently most popular generative model) in the results.

3 BACKGROUND & PROBLEM FORMULATION

3.1 Definitions

For a planar 3D roof, we can use a roof graph to represent its topology (as shown in Fig. 6 (b)). We first categorize the vertices V in the roof graph into two sets, the outline vertices V_O and the roof vertices V_R . The outline vertices are on the boundary of the (2D projection of the) 3D embedding of the roof graph and the remaining vertices are roof vertices. We denote by $n_O = |V_O|$, the number of outline vertices, and $n_R = |V_R|$, the number of roof vertices. We also use $n_v = n_O + n_R$ to denote the total number of roof vertices. For example, for the roof graph shown in Fig. 6, we have

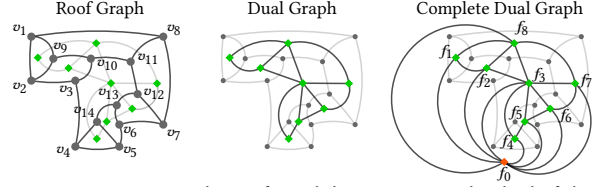


Fig. 7. We can recover the roof graph by computing the dual of the dual graph. *Left*: a given roof graph. *Middle*: its corresponding dual graph. *Right*: we add a node f_0 indicating the outside region in the dual graph. We can see that the roof graph (*left*) is the dual of the complete dual graph (*right*).

$V_O = \{v_1, \dots, v_8\}$, $V_R = \{v_9, \dots, v_{14}\}$. Similarly, we can also categorize the edges in the roof graph into two sets, the outline edges E_O (both of the two endpoints of the edge are outline vertices), and the roof edges E_R (at least one of the endpoints is a roof vertex). A region bounded by a set of edges and vertices in the embedding of the roof graph defines a *roof face*. For example, we have $f_1 = \{v_1, v_2, v_9\}$ and $f_3 = \{v_3, v_{10}, v_{11}, v_{12}, v_{13}, v_{14}, v_4\}$. Therefore, we can equivalently represent the roof graph G by (V, F) , where the edge set E can be easily extracted from the face set F . We denote $n_f = |F|$, the number of roof faces in the roof graph.

Dual graph. For each roof graph $G = (V, F)$, we can construct its *dual graph* $G^D = (V^D, E^D)$, where each face in the roof graph G is represented as a node in the dual graph G^D , i.e., $|V^D| = |F|$. Two nodes in V^D are connected to each other by an edge (stored in E^D) if and only if the corresponding two faces are adjacent (i.e., share an edge) in the roof graph G . We can store this connectivity information into an adjacency matrix $A^D \in \{0, 1\}^{n_f \times n_f}$, i.e., $A_{ij}^D = 1$ if f_i is adjacent to f_j . We can therefore equivalently represent the dual graph as $G^D = (V^D, A^D)$. In Fig. 6 (e) we show the dual graph which is placed above the original roof graph. Note that it is possible not only to construct a dual graph G^D from the primal roof graph G , but also to recover the primal roof graph from a dual graph by computing the dual of the dual graph as shown in Fig. 7 (see Sec. 4.3 for more details).

We can embed a roof graph in 3D (2D) by assigning each vertex a 3D (2D) coordinate. For a vertex set V , we denote its 2D embedding as $\bar{X} \in \mathbb{R}^{n_v \times 2}$ and its 3D embedding as $X \in \mathbb{R}^{n_v \times 3}$, where we store the vertex positions in rows. See Fig. 6 (a) for a 3D roof that stems from a 3D embedding of the roof graph in Fig. 6 (b).

3.2 Valid Roofs

An important constraint for a roof is that all faces have to be planar. We can therefore use the planarity constraint to define valid 2D and 3D embeddings of roof graphs:

DEFINITION 3.1. We call a 3D embedding of a roof graph valid if each 3D roof face is planar and the roof has non-zero height.

DEFINITION 3.2. We call a 2D embedding of a roof graph valid if there exists a valid 3D embedding such that the projection of the 3D embedding in the xy plane is exactly the same as the 2D embedding.

Therefore, we can always obtain a valid 2D embedding by projecting a valid 3D embedding to the xy plane. At the same time, we can get a valid 3D embedding by lifting up a valid 2D embedding along the z -axis (i.e., assigning each vertex a z -axis value).

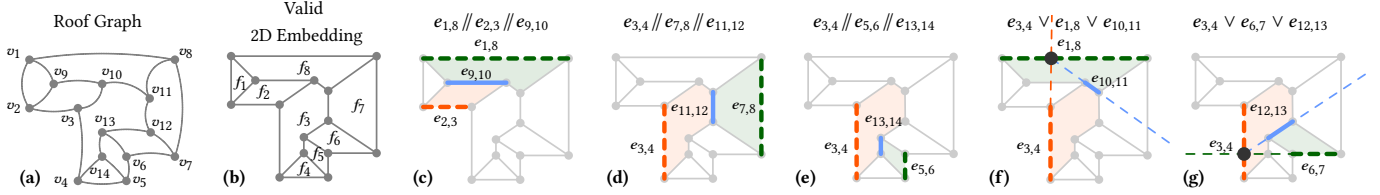


Fig. 8. **Valid 2D embedding of a roof graph.** For a pair of adjacent faces in the input roof graph (a), we consider their outline edges and their shared edge. If for every pair of adjacent faces (c-g), these three edges are either *parallel to each other* or *intersect at the same point*, according to Remark 3.1, this 2D embedding (b) is *valid*. Here the notation $e_1 \vee e_2 \vee e_3$ means these three edges e_1, e_2, e_3 intersect at the same point.

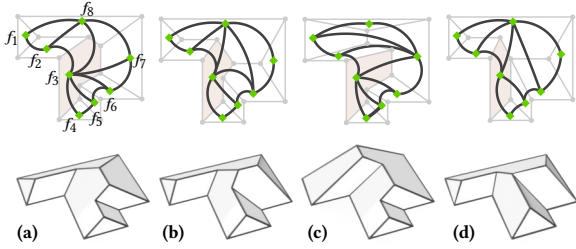


Fig. 9. **Different roof styles with the same outline.** The roof style can be encoded into the dual graph of the roof graph. *Top:* different dual graphs (with green nodes and black edges) on top of their corresponding primal roof graph (gray nodes and gray edges). *Bottom:* we show the corresponding reconstructed roofs with different styles.

Verification of the validity. To verify the validity of a given 3D embedding, we can simply check if each 3D face is planar or not. To verify the validity of a given 2D embedding, according to the definition, we need to check if there exists a set of z -values that combined with the given 2D embedding forms a valid 3D embedding. An alternative and much easier way to verify the validity of a given 2D embedding is to use basic geometry:

REMARK 3.1. *The intersecting line of two adjacent 3D planar faces with fixed outline edges, is either parallel to both outline edges, or intersects the two outline edges at the same point. The same conclusion holds when we project the 3D planar faces to xy -plane.*

See Fig. 32 and Appendix A for a simple proof. Remark 3.1 gives both a necessary and sufficient condition. Therefore, we can use it to check the validity of a 2D embedding. See Fig. 8 for an example of a *valid* 2D embedding.

3.3 Background: Straight Skeleton Methods

The straight skeleton [Aichholzer and Aurenhammer 1996; Aichholzer et al. 1996] or its extension the weighted straight skeleton [Biedl et al. 2015; Eppstein and Erickson 1999; Kelly and Wonka 2011] are popular tools for roof construction. The straight skeleton based methods take a 2D roof outline as input, and output a valid 2D/3D roof embedding by solving for the roof *topology* and roof *embedding* at the same time. Specifically, the straight skeleton methods formulate the roof construction problem as determining how roof planes with given slope from a given roof outline intersect with each other. See the inset figure for an example, where blue planes stemming from the outline intersect and form the roof structure colored in red. This can be

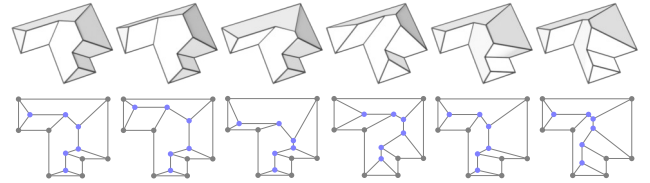
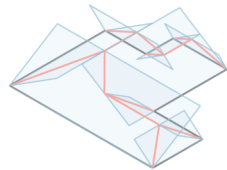


Fig. 10. **Multiple valid embeddings for the same roof graph.** *Top:* a set of 3D planar roofs with the same outline and topology. *Bottom:* xy projection of the corresponding roofs.

equivalently formulated as shrinking the input outline edges with a constant rate and determining how the resulting interior polygon changes. Once there is a change, an interior roof vertex or roof edge can be detected accordingly. See [Felkel and Obdrzalek 1998] for a more detailed description.

3.4 Observations & Challenges

The straight skeleton based methods can construct planar roofs from given roof outline efficiently. However, they still have some limitations. For example, the roof topology is determined at the same time as the roof embedding, with an implicit assumption that a single roof topology corresponds to the input roof outline, which is not the case in practice. For example, we can observe that:

- **Roofs with the same outline can have different styles.** In Fig. 9 we show four different roof constructions for the *same* outline. We can observe that these four roofs have different style and structure.
- **Roofs with the same outline and topology can have different embeddings.** In Fig. 10 we show a set of 3D planar roofs with exactly the same outline and the same adjacency between faces. All of the shown roofs are valid but the 3D roof embeddings are different, i.e., the roof vertices have different locations.

These observations suggest that it is *not* enough to only use the roof outline for roof construction, as the straight skeleton based methods do. We also need to specify the *roof topology* and *geometry* in some way. As a result, to solve the roof construction problem, we need to tackle the following challenges:

- How to specify the *roof topology* or *roof style*, i.e., the geometry of the interior vertices of a roof?
- How to formulate or enforce *roof planarity*?
- How to *avoid undesirable but valid* roof embeddings?

In the following, we will discuss how the proposed primal-dual roof graph representation can help to tackle these challenges and solve for desirable planar roofs.

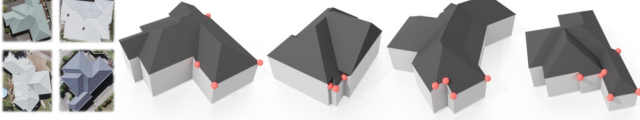


Fig. 11. Roofs with outline vertices in different height. We highlight the outline vertices with non-zero height in red.

4 METHODOLOGY: ROOF OPTIMIZATION

In this section, we introduce an optimization-based method to construct a 3D planar roof where the roof structure/style is encoded into a primal or dual roof graph. We first discuss how to formulate the roof planarity, where we introduce a planarity metric to measure the validity of an arbitrary 3D embedding of a roof graph (Sec. 4.1). We then discuss how to reconstruct a planar roof from its primal roof graph (Sec. 4.2) or its dual graph (Sec. 4.3) respectively. For simplicity of method description, we make the same assumptions as the straight skeleton method: (1) the outline vertices of a roof are in the same height, and (2) each roof face stems from one of the outline edges. We then discuss in Sec. 4.4 how to relax these assumptions to deal with roofs with outline vertices in different height (e.g., Fig. 11) and roofs with faces having multiple outline edges (e.g., Fig. 3), which are not supported by straight skeleton based methods.

4.1 Roof Planarity Formulation

Assume we have a 3D embedding X for the roof graph $G = (V, F)$, i.e., X_i is a 3D position for the vertex v_i in the roof graph, how can we evaluate the validity of the embedding X ?

Planarity metric on a point set. We first propose to use the following metric to measure the planarity of a set of 3D points Z : $g(Z) = \sigma_1(\text{Cov}(Z))$, where $\sigma_1(M)$ is the *smallest eigenvalue* of the square matrix M , $\text{Cov}(Z)$ gives the *covariance* matrix of the set of points Z . We know that $g(Z) \geq 0$ since the covariance matrix is positive semi-definite. If the 3D points in Z are coplanar, the rank of the Z is at most 2. Then, the smallest eigenvalue of the covariance matrix is 0 and we have $g(Z) = 0$. Therefore, for an arbitrary set of 3D points Z , the smaller $g(Z)$ is, the more coplanar the points Z are. We therefore use $g(Z)$ to measure the planarity of a set of points Z . Note that $g(Z)$ is *differentiable* with respect to Z .

Polygonal roof planarity. With the planarity metric in hand, we can easily measure the validity (planarity) of a roof embedding X . Specifically, we denote as X_{f_i} the corresponding 3D embedding of the face f_i in F . As discussed above, $g(X_{f_i})$ measures the planarity of the embedding of face f_i . We can sum over the planarity metric on each face f_i as a measure of the roof planarity:

$$E_{\text{planarity}}(X) = \sum_{i=1}^{n_f} \sigma_1(\text{Cov}(X_{f_i})) \quad (1)$$

Further, we can construct a valid roof by solving for an embedding X that has zero planarity error as defined above, which can be formulated as an optimization problem. In the following, we will discuss in details how to achieve a roof construction from a primal or a dual roof graph, respectively.

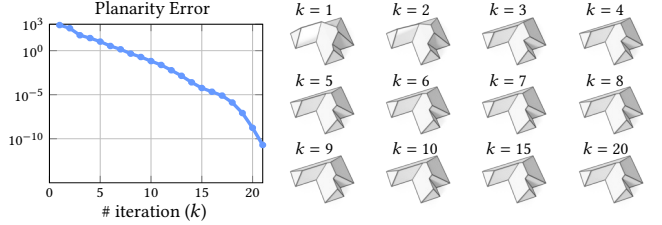


Fig. 12. Planarity error over iterations. We visualize the 3D embedding updated over iterations by minimizing the planarity measure.

4.2 Roof Construction from Primal Graph

We assume that we are given a primal roof graph $G = (V, F)$. For example, a user can draw a roof graph similar to Fig. 6 (b). In this case, a 2D embedding \bar{X}^{user} is also provided by the user. Recall that we use \bar{X} to denote a 2D embedding and use X to denote a 3D embedding of the vertex set V . Note that this 2D embedding \bar{X}^{user} is *unlikely* to be valid due to the noise in the user annotations, but it provides a strong prior of the expected positions of the roof vertices from the user. Then our goal is to solve for a *valid* 3D embedding X from the user input.

Preprocessing. We first lightly regularize the 2D positions of the outline vertices from user input, i.e. \bar{X}_O^{user} , to promote the accuracy of the parallel edges. Specifically, the outline edges labeled or drawn by users can be inaccurate: some outline edges that should be parallel are only approximately parallel. Therefore, for a pair of edges that has a smaller angle than the threshold θ , we modify the outline vertex positions a bit to make them parallel to each other numerically, and this leads to new outline vertex positions \bar{X}_O .

Problem formulation. To find a valid 3D embedding, without loss of generality X , we can fix the outline vertices V_O to $X_O = [\bar{X}_O, \mathbf{0}]$ (i.e., with 0 z-axis value). Our goal is to find a 3D embedding $X_{\mathcal{R}}$ for the roof vertices $V_{\mathcal{R}}$. We propose to solve the following problem:

$$\min_{X_{\mathcal{R}}} E_{\text{planarity}}(X) + \lambda \|\bar{X}_{\mathcal{R}} - \bar{X}_{\mathcal{R}}^{\text{user}}\|_F^2 \quad \text{s.t. } x_{z^*} = h \quad (2)$$

where x^* is a randomly selected roof vertex in $X_{\mathcal{R}}$, h is a pre-defined roof height parameter. We can optimize the above problem with initialization $X_{\mathcal{R}} = [\bar{X}_{\mathcal{R}}^{\text{user}}, \mathbf{h}]$, i.e., set the z-axis value of all the roof vertices to h . Our objective function promotes the planarity of the embedding X and at the same time promotes its corresponding 2D embedding \bar{X} to be close to the user input. See Fig. 12 for an example where we show the intermediate roofs over iterations. We also include a hard constraint that enforces one of the roof vertices to have z-value (height) as h . This design choice has two advantages: (1) it helps to avoid degenerate global minimizers. Without this hard constraint, we can see that any arbitrary 2D embedding of the roof graph with zero z-axis values leads to a 3D embedding of the graph where the planarity of each face is satisfied. To avoid this type of degenerate solutions, we can force that at least one roof vertex has non-zero height. (2) this hard constraint also provides the user a way to control the overall height of the constructed 3D roof.

4.3 Roof Construction from Dual Graph

Another scenario is that we are given a dual roof graph $G^{\mathcal{D}} = (V^{\mathcal{D}}, A^{\mathcal{D}})$. Recall that each roof face f_i is represented as a node

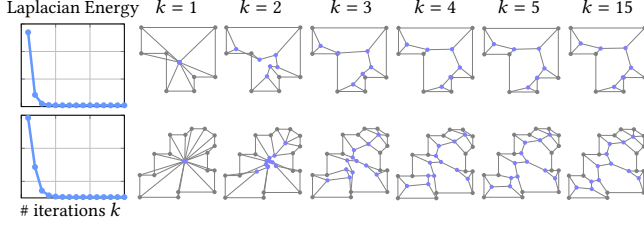


Fig. 13. 2D Spectral Embedding. Here we show two examples of embedding a roof graph into 2D with fixed outline by minimizing the Laplacian energy. We initialize all the roof vertices at the center of the roof outline.

in $V^{\mathcal{D}}$, and $A^{\mathcal{D}}$ stores the face adjacency. In practice, such a dual graph can be described by the roof outline V_O and $A^{\mathcal{D}}$. Specifically, the 2D roof outline is given as a list of consecutive 2D points. Once stored in a matrix we have $\bar{X}_O \in \mathbf{R}^{n_O \times 2}$. I.e., we have n_O outline vertices V_O embedded in 2D with vertex positions \bar{X}_O . Then we can obtain n_O outline edges, $E_O = \{e_{1,2}, \dots, e_{i,i+1}, \dots, e_{n_O,1}\}$, where the edge $e_{i,i+1}$ connects the outline vertices v_i and v_{i+1} . We assume each roof face stems from one of the outline edges, and we denote f_i as the face that is associated with the outline edge $e_{i,i+1}$. Then the face adjacency is specified in the matrix $A^{\mathcal{D}} \in \{0, 1\}^{n_O \times n_O}$.

The roof outline \bar{X}_O can either be drawn by a user or generated by a transformer. Similarly the face adjacency $A^{\mathcal{D}}$ can either be specified by a user or predicted by a trained network.

Recovering primal from dual. Since the roof planarity is defined on the primal roof graph representation, we need to first recover the primal graph from its dual graph. Specifically, we first add an outside node to the dual graph, and connect all the node in $G^{\mathcal{D}}$ to the outside node to obtain a complete dual graph. Then the primal graph can be recovered by computing the dual of the complete dual graph (see Fig. 7).

Problem formulation. With the recovered primal roof graph, we can solve for a roof embedding X by optimizing the roof planarity as before. However, there is a big difference from the previous case, where we do not have the user-specified roof interior structure $\bar{X}_{\mathcal{R}}^{\text{user}}$ for initialization and for guiding the roof optimization to a preferred structure. We therefore propose a new energy:

$$\min_{X_{\mathcal{R}}} E_{\text{planarity}}(X) + \gamma E_{\text{aesthetic}}(X) \quad \text{s.t. } X_z^* = h \quad (3)$$

where $E_{\text{aesthetic}}(X)$ encodes some additional aesthetic constraints which can help to solve for a *planar* roof with *preferred* properties. For example, we can categorize the *roof edges* into three categories according to Remark 3.1: the roof edges parallel to the corresponding outline edges are colored green (1), the roof edges that connect to an outline vertex are colored yellow (1), and the other roof edges are colored red (1), as illustrated in (b) of the inset figure. In practice, we would like to have (1) the green roof edge has equal distance to the corresponding outline edges, i.e., in the medial axis; (2) the yellow roof edge is an angle bisector that equally splits the angle formed by the two corresponding outline

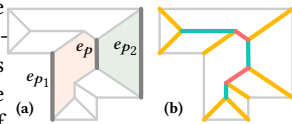


Fig. 14. Constructing the roof of the hexagonal pavilion shown in Fig. 4, where the outline vertices have different height. Starting from the initial embedding (left), we optimize the roof planarity with (right) and without (middle) the variance energy defined on vertex height, i.e., with $\eta = 1$ and $\eta = 0$ respectively. The red dashed and curved lines highlights the fact that the roof embedding is more symmetric with the variance energy. We also report the planarity error ("err") of the three embeddings.

edges. We therefore have the following aesthetic constraints:

$$E_{\text{aes.}} = \sum_{p \in \{1\}} \|\langle \vec{e}_p, \vec{e}_{p_1} \rangle - \langle \vec{e}_p, \vec{e}_{p_2} \rangle\|_F^2 + \sum_{q \in \{1\}} \|\text{dist}(\vec{e}_q, \vec{e}_{q_1}) - \text{dist}(\vec{e}_q, \vec{e}_{q_2})\|_F^2$$

where \vec{e} is a unit vector on edge e ; for a roof edge e , we can find the outline edge e_{p_1}, e_{p_2} in its neighboring faces as illustrated in (a) of the inset figure; $\text{dist}(\vec{a}, \vec{b})$ gives the distance between two parallel unit vectors \vec{a} and \vec{b} .

Initialization: 2D embedding by spectral drawing. Without user inputs, we propose to use spectral graph drawing in 2D as initial embedding for planarity optimization as opposed to random initialization, which can help to avoid self-intersections. Specifically, we first find a 2D embedding $\bar{X}_{\mathcal{R}}$ for the roof vertices by minimizing the Dirichlet energy using the graph Laplacian [Ren et al. 2018], then we initialize the 3D embedding by $X_{\mathcal{R}} = [\bar{X}_{\mathcal{R}}, \mathbf{h}]$. We can then update the 3D embedding by minimizing the roof planarity as discussed above.

To embed the roof graph in 2D by spectral embedding, we first construct the adjacency matrix $A_V \in \mathbf{R}^{n \times n}$ between the *vertices* in V , i.e., $A_V(p, q)$ equals to 1 if (v_p, v_q) is an edge in some face f_i , and equals to 0 otherwise. We can then construct the graph Laplacian $L_V = \text{diag}(\mathbf{1}_n^T A_V) - A_V$. Then we can embed the roof graph with fixed outline by minimizing the Laplacian energy:

$$\min_{\bar{X}_{\mathcal{R}}} \left\| \begin{pmatrix} \bar{X}_O \\ \bar{X}_{\mathcal{R}} \end{pmatrix}^T L_V \begin{pmatrix} \bar{X}_O \\ \bar{X}_{\mathcal{R}} \end{pmatrix} \right\|_F^2 \quad (4)$$

Note that, the spectral energy is considered at the complete roof graph while we only solve for the roof vertices $\bar{X}_{\mathcal{R}}$ with fixed outline \bar{X}_O . In this case, we can obtain a planar 2D embedding without self-intersections (see Fig. 13 for some examples).

4.4 Relaxing the Assumptions on Roof Graphs

Here we discuss how to use our optimization-based formulation to handle roofs with outline vertices at different height and roofs with faces containing multiple outline edges.

Roof with outline vertices at different heights. We can simply extend our method to handle a roof with outline vertices at different heights by setting the outline vertices as free variables for the optimization. Our method can handle common cases such as two roof outline edges of different heights emanating from the same vertex or sloped roof outline edges. See Fig. 14 for an example, where we label the outline vertices in three categories colored in green (●), yellow (●), and gray (●) respectively. To construct a realistic pavilion, the

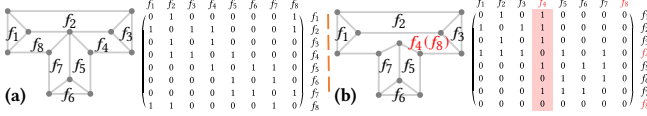


Fig. 15. **Roof face containing multiple outline edges.** We show two examples (a-b) of using dual graph (*right*) to encode the roof topology (*left*). In (b): we merge the face f_4 and f_8 in (a) into a single face, and we highlight the changes in the face adjacency matrix on the right.

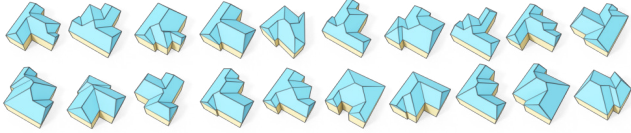


Fig. 16. **Our synthesized buildings from scratch.** Our method can automatically generate realistic roof outlines and correctly predict face adjacency. We then run our roof optimization method to construct roofs from the learned components.

green and yellow outline vertices are expected to be higher than the gray outline vertices. Therefore, we propose to solve the following problem:

$$\min_{x_{xyz}, x_z} E_{\text{planarity}}(X) + \lambda \|\bar{X} - \bar{X}^{\text{user}}\|_F^2 + \eta \text{Var}(x_z^{\bullet}) + \eta \text{Var}(x_z^{\circ})$$

where x_{xyz}^{\bullet} means that the xyz -axis values of the red vertices are variables for optimization, and x_z^{\bullet} means that only the z -axis value of the green/yellow vertices are variables while their xy -axis values are fixed. There are two modifications made to Eq. (2): (1) we set the z -axis value of the green and yellow outline vertices as variables for the optimization besides the positions of the roof vertices. (2) we add extra energy terms to regularize the *variance* of the height of the green/yellow vertices, $\text{Var}(x_z^{\bullet})$ and $\text{Var}(x_z^{\circ})$. The additional regularizers can help to construct a more symmetric and realistic pavilion (see Fig. 14 with $\eta = 0$ and $\eta = 1$). In summary, our optimization-based formulation is flexible to address user preferences by adding variables to the optimization and including different types of regularizers for different types of roofs. Fig. 11 shows more examples of roofs with outline vertices with different heights.

Roof face containing multiple outline edges. Our roof optimization from the primal graph can handle the case where a face contains multiple outline edges directly. Here we only discuss how to handle this with the dual graph as input. The main issue is how to represent such a roof in a dual graph. This can be easily handled by modifying the face adjacency matrix $A^{\mathcal{D}}$. Specifically, each outline edge corresponds to a roof face; and for the set of outline edges that correspond to the same face, the corresponding rows in $A^{\mathcal{D}}$ are merged to the first outline edge, while the rest outline edges are ignored by setting all the entries to 0. See Fig. 15 for an example of how $A^{\mathcal{D}}$ is constructed. In this way, we can use a dual graph to represent the roof topology where faces can have multiple outline edges. Note that the straight skeleton methods do not support this feature (e.g., see Fig. 3, Fig. 29, and Fig. 30).

Alternative solutions. In Appendix B we discuss three alternative planarity metrics that can be used for planar roof modeling as well.

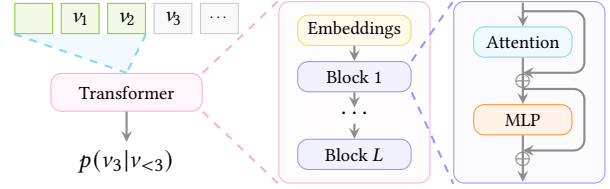


Fig. 17. Our auto-regressive transformer with input of a flattened vertex sequence, and output of the probability distribution of the next token.

$$\begin{aligned} \text{MLP} \left[\max \begin{pmatrix} \bar{f}_i^{(l-1)}, \bar{f}_j^{(l-1)} \\ n_O \end{pmatrix} \right] \oplus \begin{pmatrix} \bar{f}_{i,j}^{(l-1)} \\ n_O \end{pmatrix} \oplus g^{(l-1)} &= \begin{pmatrix} \bar{f}_{i,j}^{(l)} \\ n_O \end{pmatrix} \\ \text{MLP} \left[\begin{pmatrix} \bar{f}_{i,j}^{(l)} \\ n_O \end{pmatrix} \right] \oplus \begin{pmatrix} \bar{f}_i^{(l-1)} \\ n_O \end{pmatrix} \oplus g^{(l-1)} &= \begin{pmatrix} \bar{f}_i^{(l)} \\ n_O \end{pmatrix} \\ \text{MLP} \left[\begin{pmatrix} \bar{f}_i^{(l)} \\ n_O \end{pmatrix} \right] \oplus g^{(l-1)} &= g^{(l)} \end{aligned}$$

Fig. 18. **Face adjacency prediction building block.** D is the feature dimensions. \oplus is the concatenation operator. From top to bottom, we show the adjacency model, the edge model, and the global model

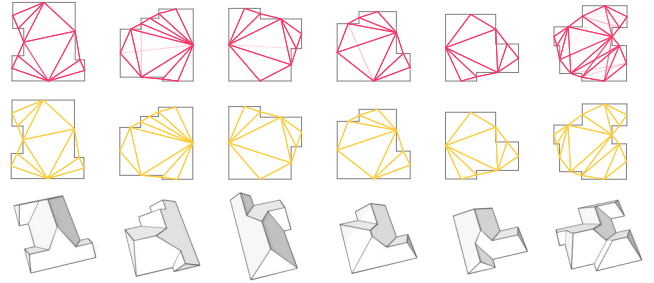


Fig. 19. *Top*: the predicted adjacency with probability using our transformer. *Middle*: post-processed adjacency that forms a valid dual graph. *Bottom*: the corresponding constructed 3D roofs using our method.

In Appendix C we introduce another solution to optimize for a valid 2D embedding directly based on Remark 3.1.

5 GENERATIVE MODELS FOR ROOF SYNTHESIS

One direct application of our method is roof synthesis from scratch (see Fig. 16). We believe that synthesizing a valid roof directly can be hard since the model needs to take care of the discrete constraints (roof topology) and the continuous constraints (roof embedding) at the same time. We propose to tackle the roof synthesis problem by combining generative models for roof topology generation (dealing with discrete constraints only) with roof optimization (dealing with continuous constraints only). Specifically, we propose a *transformer* for roof outline generation and a *graph neural network* for face adjacency prediction. The architectures for both networks can be found in the supplementary materials.

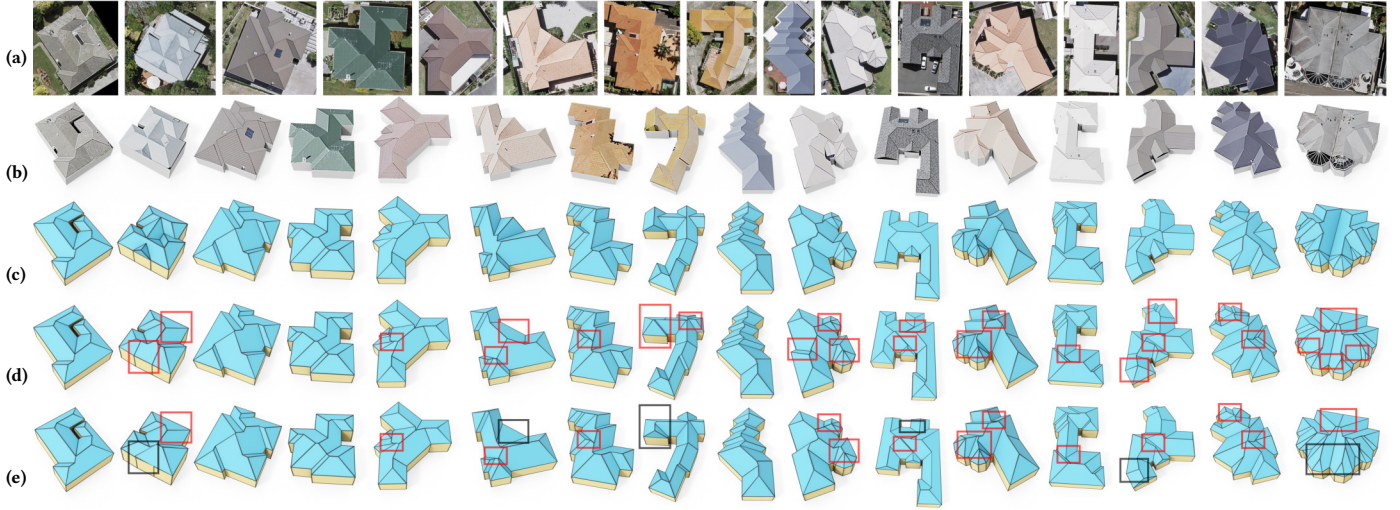


Fig. 20. **Roof reconstruction from aerial images.** (a) 16 input images of roofs with different complexity. (b) our reconstructed roofs with texture. (c) the topology of our reconstructed roofs. (d) the results of the straight skeleton where the visually erroneous regions are colored in red. (e) we use the weighted straight skeleton method to refine the results in (d) aiming to make them more consistent with the input image. The weighted straight skeleton successfully resolved 5 inconsistencies. However it also introduced new inconsistencies (colored in black) and failed to resolve several inconsistencies (colored in red).

Outline generation. Our goal is to model a distribution of the 2D roof outline $X_{\mathcal{O}} \in \mathbf{R}^{n_{\mathcal{O}} \times 2}$, where we assume the outline vertices are in counter-clockwise order and the first vertex is the one closest to the lower left corner. We flatten the coordinate matrix to $N^{seq} = \{v_1, v_2, \dots, v_{2n_{\mathcal{O}}}\}$. The vertex values are first normalized to range $[0, 1]$ and then quantized to b -bits, i.e., v_i belongs to the set $\{1, 2, \dots, 2^b\}$ for any i . We also append the sequence N^{seq} with a stopping token s . Consequently, the sequence has the length of $2n_{\mathcal{O}} + 1$ and each entry of the sequence has $2^b + 1$ kinds of tokens. We train a transformer [Vaswani et al. 2017] to convert the input tokens N^{seq} to roof outline embeddings. The probability of N^{seq} can be factorized into a chain of conditional probabilities: $p(N^{seq}; \varphi) = \prod_{i=1}^{2n_{\mathcal{O}}} p(v_i | v_{<i}; \varphi)$, where φ are the parameters of the model. The model is an auto-regressive network implemented with a transformer. The network outputs a probability p at time step i based on $v_{<i} = \{v_1, v_2, \dots, v_{i-1}\}$. See Fig. 17 for the structure of our transformer. We train this model by minimizing the negative log-likelihood over all training sequences.

Face adjacency prediction. With the generated roof outline, we propose a GCN [Kipf and Welling 2016] to predict the face adjacency, i.e., $p_{i,j}$, the probability of having the face stemming from the edge $e_{i,i+1}$ being adjacent to the face stemming from the edge $e_{j,j+1}$, for all $1 \leq i, j \leq n_{\mathcal{O}}$. The network is built by L basic building blocks. The l -th block updates 3 types of representations: (1) an edge model updates the feature representation $\tilde{f}_i^{(l)}$ for the edge $e_{i,i+1}$; (2) an adjacency model updates the feature representation $\tilde{f}_{i,j}^{(l)}$ for the adjacency $(e_{i,i+1}, e_{j,j+1})$; (3) a global model updates the global feature representation $g^{(l)}$. See Fig. 18 for an illustration of these three building blocks. Specifically, the input roof outline is transformed by L blocks and we obtain the final adjacency representation $\tilde{f}_{i,j}^{(L)}$. We project the representation through a fully-connected layer which

outputs the adjacency probability: $p_{i,j} = \text{Sigmoid}(\text{FC}(\tilde{f}_{i,j}^{(L)})) \in [0, 1]$. The loss function of our GCN is the binary cross entropy between the predicted probability $p_{i,j}$ and the ground-truth adjacency A_F . See Fig. 19 for some qualitative examples, where we visualize the probability $p_{i,j}$ via opacity (top row). We can extract the dual graph with the highest predicted probability (second row) and construct planar roofs by using our roof optimization method (bottom row). See the supplementary materials for more details and discussions about our generative models.

6 RESULTS

In this section, we show results of our roof optimization method from the primal graph and dual graph. We demonstrate the advantages of our method over the straight skeleton based methods and commercial software.

6.1 Roof Reconstruction from Aerial Images

6.1.1 Comparison to Straight Skeleton based methods. In Fig. 20 we compare to the straight skeleton based methods on roof reconstruction, which take user-specified roof outlines as input [Aichholzer and Aurenhammer 1996; Eppstein and Erickson 1999]. We test on 16 aerial images containing roofs with different structure and complexity. Then the primal roof graph of the input image is specified by a user. We run our roof optimization method to reconstruct the roofs, and report the runtime including user labeling and optimization in Table 2 and show our reconstructed roofs in Fig. 20 (b-c). Note that, for the most complicated roof with more than 50 of roof vertices, it only takes less than 5 minutes to label and reconstruct a realistic roof from the aerial image.

We then compare to the straight skeleton and the weighted straight skeleton using the same roof outline as ours. In row (d) of Fig. 20, we show the results of the straight skeleton. Though globally,

Table 2. **Comparison to straight skeleton based methods.** We report the complexity of the roofs shown in Fig. 20, including the number of vertices (n_v) and faces n_f . For each image shown in Fig. 20 (with underlying roof having n_v vertices and n_f faces), We report the number of visual inconsistencies (#err) between the constructed roof and the image, the number of vertices (\bar{n}_v) and faces (\bar{n}_f) on the reconstructed roof, and the construction time (t) of different methods.

No.	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
n_v	22	24	25	25	27	27	32	33	33	33	34	35	36	39	39	51
n_f	12	17	14	14	15	16	17	20	18	20	18	21	19	22	21	38
Ours	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
#err																
ss	0	2	0	0	1	2	1	2	0	3	2	2	1	3	2	4
wss	0	2	0	0	1	2	1	1	0	2	2	2	1	2	2	2
Ours	22	24	25	25	27	27	32	33	33	33	34	35	36	39	39	51
\bar{n}_v																
ss	22	22	26	26	30	28	34	40	34	38	38	40	38	40	44	50
wss	22	22	26	26	30	28	34	40	34	38	38	40	38	40	44	50
Ours	12	17	14	14	15	16	17	20	18	20	18	21	19	22	21	38
\bar{n}_f																
ss	12	12	14	14	16	15	18	21	18	20	20	21	20	21	23	26
wss	12	12	14	14	16	15	18	21	18	20	20	21	20	21	23	26
Ours	89.4	115	153	97.9	114	92.9	151	145	155	167	148	158	179	199	178	284
t																
(s)																
ss	16	20	24	18	20	16	28	29	27	33	32	35	31	33	38	38
wss	-	300	-	-	60	180	180	300	-	120	180	60	60	480	180	360

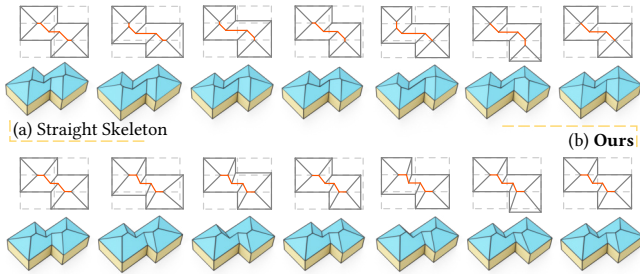


Fig. 21. *Top*: the straight skeleton algorithm is sensitive to the input and can lead to different roof structure for very similar roof outlines. *Bottom*: as a comparison, our model allows to fix the roof structure with different outlines. We overlay the roof graph with a fixed-scale grid colored in gray to better visualize the difference between the outlines.

the obtained results appear reasonable, note that the reconstructed roofs using straight skeleton contain a lot of structural inconsistencies w.r.t. the input images and unrealistic errors (highlighted in red). We then use the weighted straight skeleton to fine-tune the results in order to fix the errors by tuning the edge weights. We use the GUI provided by [Kelly and Wonka 2011] for the weighted straight skeleton where the user is allowed to change the weight for each outline edge. We asked a well-trained user to tune the edge weights until the structure of the reconstructed roof is as consistent as possible with the one shown in the image. We can see that the weighted straight skeleton can fix some inconsistencies and errors in the roofs constructed using the straight skeleton (for those successful edits, we changed the highlighting color from red to black). However, there are still many structural errors that cannot be fixed by changing the weights (highlighted in red).

In Table 2 we report the topology of the reconstructed roofs from different methods, including the number of vertices (\bar{n}_v) and faces (\bar{n}_f). We can see that the roofs reconstructed by the standard or

Table 3. **Comparison to commercial software.** We compare to the roofs constructed in 3ds Max (3D) and SketchUP (SU) and report if the constructed roof is planar (valid), the number of topological errors ("err"), the ratio of the polygon faces in the roof ("poly%"), the number of vertices (\bar{n}_v) and faces (\bar{n}_f) of the constructed roof, and the modeling time in minutes (t).

No.	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
Ours	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
valid																
3D	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
SU	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Ours	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
#err																
3D	9	12	10	11	12	12	16	15	15	11	16	15	16	18	17	20
SU	20	16	21	19	16	27	33	21	25	38	28	30	33	33	29	42
Ours	75	65	71	79	67	69	82	75	83	60	72	57	79	73	71	42
poly (%)																
3D	75	62	71	79	75	67	87	75	83	52	76	60	80	72	74	56
SU	0	24	2.6	3.3	20	2.3	2.0	22	7.0	10	11	9.8	3.8	15	14	0
Ours	22	24	25	25	27	27	32	33	33	33	34	35	36	39	39	51
\bar{n}_v																
3D	22	31	37	26	28	28	44	34	34	35	44	35	36	40	39	53
SU	22	29	26	25	28	30	36	45	33	52	39	40	37	46	42	58
Ours	12	17	14	14	15	16	17	20	18	20	18	21	19	22	21	38
\bar{n}_f																
3D	12	21	14	14	16	18	18	20	18	21	21	25	20	25	23	36
SU	32	33	35	33	31	43	50	41	43	58	46	51	52	55	50	82
Ours	1.5	1.9	2.6	1.6	1.9	1.6	2.5	2.4	2.6	2.8	2.5	2.6	3.0	3.3	3.0	4.7
t																
(min)																
3D	6	7	7	6	6	12	12	11	7	14	6	9	10	8	10	23
SU	12	16	15	12	22	20	21	22	21	25	19	32	22	14	25	36

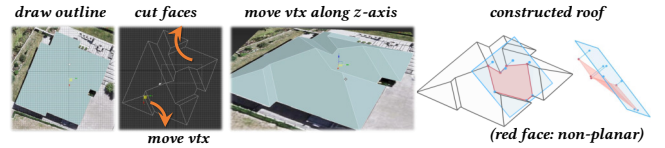


Fig. 22. Workflow of using 3ds Max for roof modeling (No.3 roof in Fig. 20).

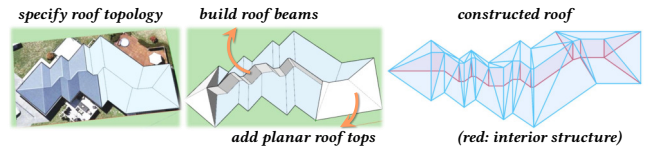


Fig. 23. Workflow of using SketchUp for roof modeling (No.9 roof in Fig. 20)

the weighted straight skeleton method always have the same number of vertices and faces. It suggests that although the weighted straight skeleton can fix some visual inconsistencies from the standard method, it cannot change the roof topology. The structural errors of the straight skeleton based methods in Fig. 20 show that these methods have much less expressiveness power in roof topology representation than our method. As suggested in Fig. 21, the straight skeleton method is very sensitive to the input condition and can lead to different topology for similar building outlines.

6.1.2 *Comparison to Commercial Software.* We also compare to two commercial software frameworks for roof reconstruction. We asked two experts, one with 5-year experience of modeling in 3ds Max, and the other with 3-year experience of modeling in SketchUp, to reconstruct the roofs shown in Fig. 20. The only instruction we gave to the experts was to model a polygonal roof that is as consistent as possible with the given image and as simple as possible. The two experts that worked independently followed a similar logic: they

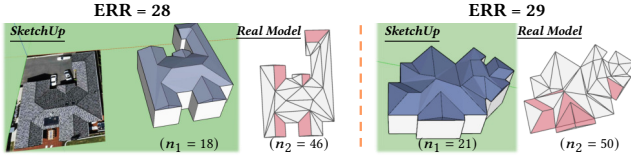


Fig. 24. **Evaluating roofs constructed in SketchUp.** The expert working in SketchUp hid some edges to make the constructed roofs visually consistent with the input image (*left*). However, the real models shown on the *right* are more complicated. We therefore measure the *unwanted complexity*, i.e., the difference between the number of faces in the real model (n_2) and the number of faces in the visually expected model (n_1), to evaluate SketchUp. We highlight the faces in the real models that are not triangles in red.

first specified the roof topology on top of the imported image by drawing the outline and then by adding/cutting faces; then they constructed a 3D roof based on the roof topology. Specifically, the 3ds Max expert chose to move vertices mainly along the z-axis and checked in different views until satisfied with the reconstructed roofs, as shown in Fig. 22. The SketchUp expert chose to first build the roof beams, i.e., vertical planes such that the rooftop planes can be placed on top of it. See Fig. 23 for an example. We show some quantitative comparison in Table 3. For the roofs that are constructed in SketchUp, we ignore the constructed interior structures and only evaluate the rooftops.

In general, commercial software provides more modeling tools and can model a larger variety of polygonal meshes. There are still some limitations for roof modeling. First, commercial software needs domain knowledge to be used efficiently, while our user input is light and friendly for novice users. For example, the 3ds Max expert used different types of operations including creating polygons by adding edges, moving vertex positions in a 2D plane, cutting faces, extruding faces, translating grouped vertices and so on. During the modeling process, the expert had to frequently change the views or even switch to the four-view editing mode to operate and check the constructed 3D roof. As a comparison, our method allows the user to specify the roof topology in 2D with one simple operation, i.e. clicking on the image to construct a primal or dual roof graph. More importantly, it is hard to explicitly enforce the planarity of the 3D roof faces using 3ds Max or SketchUp. For example, 3ds Max allows to create a polygon from a set of non-coplanar 3D points. Therefore, it relies on the user to adjust the vertex positions to make the polygonal roof faces planar, which is extremely hard to achieve even visually. See Fig. 22 for an example where we highlight a non-planar polygon (colored red) w.r.t. a reference plane (colored blue). Also as reported in Table 3, all the roofs constructed in 3ds Max are not planar numerically. On the other hand, SketchUp does not allow to construct a non-planar polygon. Therefore, to form a rooftop from a set of non-coplanar vertices, a user need to manually triangulate the rooftop. We show the constructed roofs by 3ds Max and SketchUp in Fig. 33 and Fig. 34 in Appendix. In Table 3 we report the ratio of polygonal faces in the constructed roof. We can see that this ratio of the roofs constructed using SketchUp is significantly lower than our method or using 3ds Max due to the implicit planarity constraint in SketchUp.

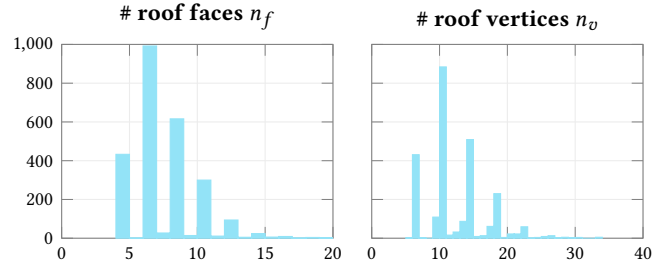


Fig. 25. Variations of our constructed dataset: we report the number of faces and vertices on each of the 2539 constructed buildings via a histogram. The number of faces ranges from 4 to 20, and the number of vertices on the roofs ranges from 5 to 34.

Another issue of roof modeling using the commercial software is that it heavily depends on the user’s preference of how to specify roof topology. See the inset figure for an example, the roof constructed using SketchUp is visually consistent with the input image. However, it would be more plausible to have the roof faces colored in red being connected to the main roof body colored in blue. In practice, it is typically preferable to avoid self-intersections as highlighted by the gray part. As a comparison, our proposal of using a roof graph to describe roof topology is general and can result in roofs with simpler topology. For example, as shown in Fig. 3, the roofs constructed by our method have a smaller number of vertices and faces. Please see the supplementary video for more examples.

We quantitatively measure the topological errors (“#err”) of the constructed roofs and report them in Table 3. Specifically, we measure how many *non-planar* faces in each roof constructed in 3ds Max. In SketchUp, the expert hid some edges of the constructed roofs as shown in Fig. 24. We can then use the “unwanted face complexity” as a measure to evaluate the topological errors in SketchUp. In summary, compared to the commercial software for roof modeling, our method is more efficient and simpler for novice users. The roofs constructed using our method have a simpler topology while the roof planarity and the visual consistency to the input image are enforced.

6.2 Image-Building Paired Dataset

We created an image-building paired dataset using our roof optimization method, where a complete building is constructed by adding facade planes along the roof outline and a base plane at the bottom. Specifically, we created a dataset consisting of 2539 buildings paired with the input aerial image and face labels (including roof faces, facade faces, and base face). The annotations are cleaned automatically by merging close-by vertices, removing duplicate or redundant edges/vertices, and our method is then used for roof reconstruction. See Fig. 5 for some example buildings and Fig. 25 for a summary of the roof complexity including the number of roof vertices and roof faces in this dataset. We believe this dataset can be helpful for different visual computing tasks. Specifically, this is a mesh-image paired dataset, which can be used for learning-based roof mesh



Fig. 26. Generated roof outlines with our auto-regressive model. We use our model to generate a sequence of 2D vertices and connect the tail vertex to the head by an orange line.

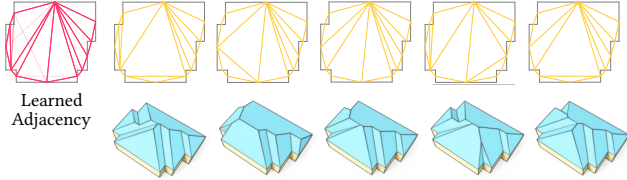


Fig. 27. We can extract multiple valid dual graphs from the learned adjacency (top). We show the corresponding constructed roof on the bottom.

detection. We also assign each polygon face in the building mesh a label from the set of roof face, body face, or base face. We can sample from the polygon mesh to obtain a point cloud as well (see inset figure). This dataset also contains roofs with a larger range of complexity. For example, in the following we discuss how to use this dataset for roof synthesis from scratch.

6.3 Application 1: Roof Synthesis from Scratch

As discussed in Sec. 5, our roof modeling formulation can simplify the roof synthesis problem. Specifically, we propose learning-based techniques for roof topology synthesis, and then use our roof optimization method to enforce geometric planarity constraints. We use the constructed dataset in Sec. 6.2 to train models on roof graph generation, including a transformer for roof outline generation and a graph neural network for face adjacency prediction. See Fig. 26 for some example roof outlines generated by our transformer. Our graph neural network can predict the probability of adjacency for face pairs, from which we can extract either a single dual graph with highest probability (see Fig. 19) or *multiple* valid dual graphs for roof construction as shown in Fig. 27. See the supplementary materials and videos for more discussions and results.

We compare to a Variational Auto-Encoder (VAE) based generative model [Kingma and Welling 2014] for roof graph generation. We used the trained VAE to synthesize 360 roof graphs, and only 119 of them are fully connected graphs while the remaining graphs have up to 19 disconnected components. We then only focus on fully connected cases for potentially valid roof graphs. Fig. 28 shows some example roof graphs synthesized by the VAE-based model. Even most of the fully connected roof graphs do not have a reasonable topology. Among the few synthesized roofs that do have a reasonable topology (e.g., the first and the last one) the geometry is not valid and violates aesthetic constraints. We therefore conclude that the task of constraint geometry generation is very difficult for a VAE. This shows that our strategy of separating the continuous constraints from the discrete constraints can simplify the problem and make it easier for training a generative model to learn roof topology.

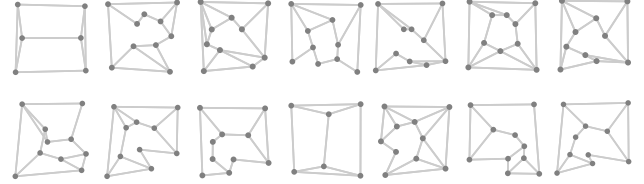


Fig. 28. **Synthesized roof graph via VAE.** Synthesizing a valid roof directly can be hard since the model needs to take care of the discrete constraints and the continuous constraints at the same time.

6.4 Application 2: Interactive Roof Editing & Optimization

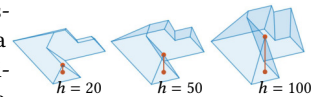
One of the biggest advantages of our roof construction method is its flexibility. Specifically, the optimization based planarity formulation makes it possible to incorporate different regularizers. Moreover, the primal-dual roof graph representation can support different editing operations. Therefore, our method can be used for interactive roof editing and optimization. Specifically, a user can (1) modify/edit a (valid) roof graph (either the primal or dual graph) (2) starting from the modified roof graph, run our optimization method to obtain a *valid* roof graph. The user can then go back to step (1) and edit again. In this way, one can edit the roof graph until satisfied.

Our primal-dual roof graph representation can naturally support different types of operations including moving a vertex or an edge, snapping an edge, merging two faces, splitting a face, forcing two faces to be adjacent, and so on. See Fig. 29 for an example. We can see that our editing operations are expressive and our optimization-based formulation is well suited for interactive editing since after applying different operations the updated roof embedding does not change too much from the previous embedding while staying valid.

As a comparison, the weighted straight skeleton, which allows a user to change edge weights for roof editing, is not trivial or efficient enough for interactive editing, since the change of the roof structure is not continuous or easily predictable w.r.t. the change of edge weights. See Fig. 30 for an example, where we compare our interactive editing power to the weighted straight skeleton in correcting topological errors. Our method can easily fix all the errors while the weighted straight skeleton can only fix one out of four topological errors. See Appendix D for more discussions.

6.5 Additional Justification of Our Formulation

Adjustable roof height. As discussed in Eq. (3), we optimize for a valid 3D roof embedding by minimizing the planarity energy w.r.t. a hard constraint such that a randomly selected roof vertex should have fixed height (z -axis value) h . This can help to avoid degenerate solutions where all the roof vertices have zero height, which leads to valid roofs with zero planarity error. In the inset figure we show that this parameter h is not critical and we can set it to an arbitrary value to obtain valid 3D roofs. Additionally, this design choice can also benefit the interactive roof editing where a user can tune the roof height h during the construction. In our experiments, we usually set $h = \sqrt{S}/2$, where S is the roof area.



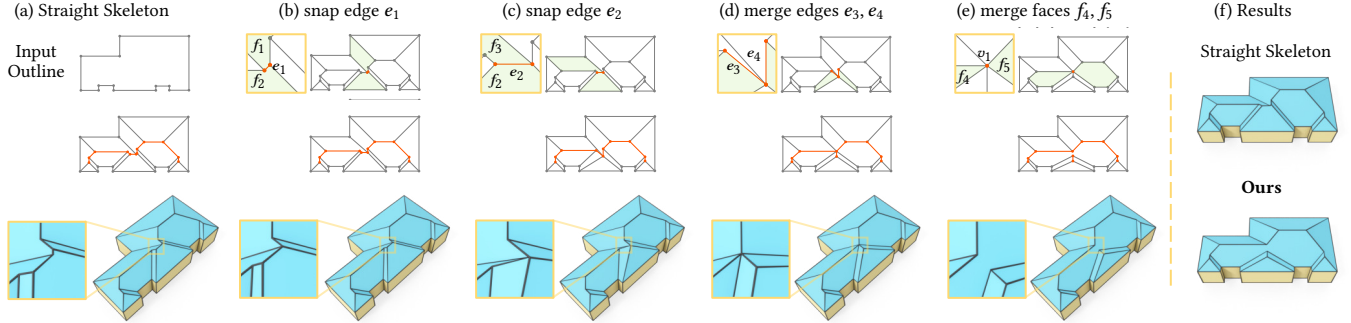


Fig. 29. The roof constructed from the straight skeleton can be undesirable or unrealistic as shown in (a). Moreover, the straight skeleton formulation only supports a limited set of edits. As a comparison, our optimization-based formulation supports different types of edits (b-e). In (f), we show that after a set of edits to the result from the straight skeleton, we can obtain a more realistic roof. For (b-e), *Top*: we visualize the edits made to the roof graph and we show a zoom-in version on the left to highlight the region of interest. *Middle*: we show the valid 2D embedding after editing using our method and we color the roof edges in red to highlight the roof structure. *Bottom*: we show the corresponding constructed planar roofs with a zoomed-in version of the regions with changes.

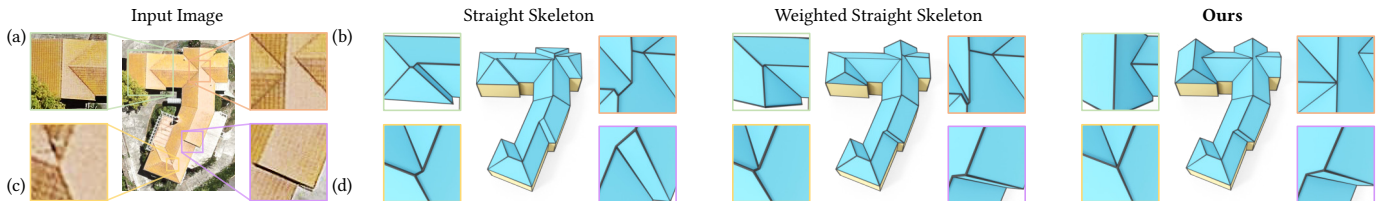


Fig. 30. The roofs constructed using straight skeleton can be different from the input image. For example, the highlighted region with a zoom-in view of (a) and (b) contains errors; region (b) and (c) contains extremely short edges which are unlikely to exist in reality; region (d) has inconsistent vertex positions. We can consider using the weighted straight skeleton to fix these issues by changing the weight for each outline edge. However, changing the edge weights is not trivial and the user could only solve the vertex inconsistency in (d), while the other structural errors in (a,b,c) are mediated, but not fully resolved. As a comparison, our operations for interactive editing are explicitly defined on the roof graph and are much easier to apply. Our method successfully fixes all the errors and obtains a consistent roof.

Usefulness of spectral embedding. Taking the 2D spectral embedding as initialization can help to avoid self-intersections for roof constructions from a dual graph. In the inset figure, we show the roofs optimized from zero embedding (left) and our spectral embedding (right). Both roofs are valid with planarity error ("err") smaller than 10^{-9} . However, the roof shown on the right is more plausible with no self-intersections. In Fig. 31, we show more results of our roof optimization algorithm from a dual graph. On top we show the 2D spectral embedding as the initialization, in the middle we show the optimized valid 2D embedding, and at the bottom we show the corresponding reconstructed roofs (buildings).

6.6 Implementation & Runtime

We implemented the roof optimization methods (including spectral embedding and planarity optimizations) in MATLAB and used the build-in function "fmincon" with Quasi-Newton solver for optimization. We designed a web-based GUI with two modes for collecting user inputs of primal-dual roof graph specification for the application of roof reconstruction from images. The roof synthesis application is implemented using Pytorch [Paszke et al. 2019].

In Table 4 and Table 5, we report the roof complexity and computation time of our method for each roof in Fig. 20 and Fig. 31

Table 4. **Runtime of roof construction from primal graph.** For each of the examples in Fig. 20, we report the time for annotating the vertices t_o and faces t_f in the image for roof reconstruction by the user. t_o reports the runtime in seconds of our roof optimization algorithm. $t_{ours} = t_o + t_f + t_o$ shows the total amount of time for the roof reconstruction of our method.

No.	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
n_v	22	24	25	25	27	27	32	33	33	33	34	35	36	39	39	51
n_f	12	17	14	14	15	16	17	20	18	20	18	21	19	22	21	38
t_o	27.6	36.6	39.3	29.4	31.4	25.5	43.9	40.8	43.5	49.8	49.9	52.5	50.9	55.5	58.3	66.6
t_f	61.4	77.7	112	67.9	82.2	66.5	106	102	109	116	97.3	104	126	143	118	214
t_o	0.42	1.20	0.73	0.55	0.61	0.89	0.94	2.23	1.47	1.05	0.72	1.58	1.24	1.34	2.52	3.59
t_{ours}	89.4	115	153	97.9	114	92.9	151	145	155	167	148	158	179	199	178	284

respectively. We can see that our algorithms, both 2D spectral embedding and planarity optimization, are efficient and robust w.r.t. various roof outlines with different complexity.

7 CONCLUSION, LIMITATION & FUTURE WORK

We proposed an optimization-based roof construction method that first designs a primal or dual roof graph as input and then optimizes the geometry to output a planar 3D polygonal roof. Our formulation is flexible and can be adapted easily to different settings such as incorporating user-specified regularizers. Our method has two practical applications, interactive roof editing and roof synthesis from scratch. Our method of roof reconstruction is more expressive than the straight-skeleton based methods, and is much easier

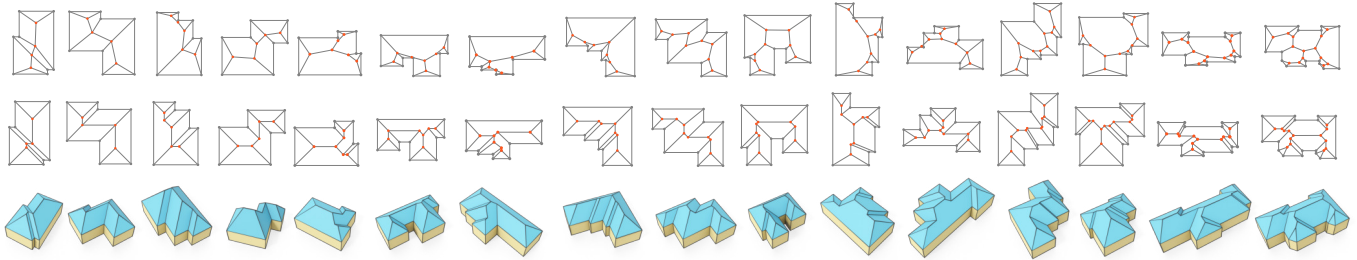


Fig. 31. **Roof construction from dual graph.** *Top:* the initial 2D spectral embedding. *Middle:* the valid 2D embedding of each roof obtained by minimizing the planarity. *Bottom:* the constructed planar roof meshes.

Table 5. **Runtime of roof construction from dual graph.** For the 16 example roofs shown in Fig. 31, we report the complexity of each roof, including the number of vertices (n_v) and faces (n_f) in the roof. We also report the runtime in seconds for the 2D embedding via the spectral method (t_1) and 3D embedding via optimization w.r.t. the planarity and aesthetic constraints (t_2).

No.	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
n_v	12	12	13	13	14	17	17	17	18	18	18	21	22	22	26	30
n_f	8	8	8	8	8	10	10	10	10	10	10	12	12	12	14	16
t_1	0.04	0.03	0.03	0.04	0.04	0.05	0.05	0.04	0.05	0.05	0.06	0.06	0.08	0.06	0.08	0.09
t_2	0.20	0.17	0.26	0.24	0.29	0.41	0.39	0.37	0.43	0.44	0.45	0.60	0.80	0.69	1.90	1.15

for novice users to use than commercial software. We also use our method to construct a image-building paired dataset with 2539 roof meshes, that can be helpful for different visual computing tasks. Although our method can be used to model roofs with different styles, including buildings with inner courtyards or vertical facades inside the roof, and buildings with outline edges in different height as shown in Fig. 4, it cannot directly handle curved roofs including stadiums and skyscrapers. One possible solution is to approximate curved roofs with planar sub-faces, which can be reconstructed via our method as shown in Fig. 4 (a). Another limitation of our work is that we do not model roof textures. While it would be interesting to research a GAN for the generation of roof textures in our framework, image synthesis is mainly orthogonal to the core topics of our paper. Further, the interactive roof modeling is still fairly slow so that it is hard to scale up the dataset construction by another order of magnitude. For practical reasons, it might be very important to investigate efficient combinations of automatic and interactive reconstruction that would be time-saving without increasing the error rate compared to a human only baseline. Finally, we did not touch on automatic reconstruction in this paper. In future work, we would like to investigate transformers using images for cross attention following to the work of Dosovitskiy et al. [2021], and investigate how to predict roof graphs from images directly by using the image-mesh dataset we constructed. It would also be interesting to study practical constraints for roof fabricability using our optimization-based formulation, such as incorporating slope requirements of roof faces during the construction, which can be addressed by either hard or soft constraints.

ACKNOWLEDGMENTS

The authors thank the anonymous reviewers for their valuable comments. Parts of this work were supported by the KAUST OSR Award No. CRG-2017-3426, the ERC Starting Grant No. 758800 (EXPROTEA), the ANR AI Chair AIGRETTE, and Alibaba Innovative Research (AIR) Program. We would like to thank *Guangfan Pan* and *Jiacheng Ren* for helping modeling roofs in 3ds Max and SketchUp, *Jialin Zhu* for helping designing the web-based roof annotation UIs, *Jianhua Guo* and *Tom Kelly* for helping with the comparison to the weighted straight skeleton. We thank *Muxingzi Li* for helping editing the supplementary video. We also thank *Chuyi Qiu*, *Tianyu He*, and *Ran Yi* for their valuable suggestions and comments.

REFERENCES

- Oswin Aichholzer and Franz Aurenhammer. 1996. Straight skeletons for general polygonal figures in the plane. In *International Computing and Combinatorics Conference*. Springer, 117–126.
- Oswin Aichholzer, Franz Aurenhammer, David Alberts, and Bernd Gärtner. 1996. A novel type of skeleton for polygons. (1996), 752–761.
- F Alidoost, H Arefi, and M Hahn. 2020. Y-Shape convolutional neural network for 3D roof elements extraction to reconstruct building models from a single aerial image. *ISPRS Annals of Photogrammetry, Remote Sensing & Spatial Information Sciences* 5, 2 (2020).
- Murat Arikian, Michael Schwärzler, Simon Flöry, Michael Wimmer, and Stefan Maierhofer. 2013. O-snap: Optimization-based snapping for modeling architecture. *ACM Transactions on Graphics (TOG)* 32, 1 (2013), 1–15.
- Jean-Philippe Bauchet and Florent Lafarge. 2020. Kinetic shape reconstruction. *ACM Transactions on Graphics (TOG)* 39, 5 (2020), 1–14.
- Therese Biedl, Martin Held, Stefan Huber, Dominik Kaaser, and Peter Palfrader. 2015. Weighted straight skeletons in the plane. *Computational Geometry* 48, 2 (2015), 120–133.
- Andrew Brock, Theodore Lim, James Millar Ritchie, and Nicholas J Weston. 2016. Generative and Discriminative Voxel Modeling with Convolutional Neural Networks. In *Neural Information Processing Conference: 3D Deep Learning*.
- Cyprien Buron, Jean-Eudes Marvie, and Pascal Gautron. 2013. GPU Roof Grammars. In *Eurographics (Short Papers)*. 85–88.
- Xi Chen, Nikhil Mishra, Mostafa Rohaninejad, and Pieter Abbeel. 2018. Pixelsnail: An improved autoregressive generative model. In *International Conference on Machine Learning*. PMLR, 864–872.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime G. Carbonell, Quoc Viet Le, and Ruslan Salakhutdinov. 2019. Transformer-XL: Attentive Language Models beyond a Fixed-Length Context. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, Anna Korhonen, David R. Traum, and Lluís Màrquez (Eds.). Association for Computational Linguistics, 2978–2988.
- Youness Dehbi, André Henn, Gerhard Gröger, Viktor Stroh, and Lutz Plümer. 2021. Robust and fast reconstruction of complex roofs with active sampling from 3D point clouds. *Transactions in GIS* 25, 1 (2021), 112–133.
- Ilke Demir, Daniel G Aliaga, and Bedrich Benes. 2015. Procedural editing of 3d building point clouds. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. 2147–2155.
- Laurent Dinh, David Krueger, and Yoshua Bengio. 2015. NICE: Non-linear Independent Components Estimation. In *International Conference on Learning Representations (ICLR)*, Yoshua Bengio and Yann LeCun (Eds.).

- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations (ICLR)*.
- Günther Eder and Martin Held. 2018. Computing positively weighted straight skeletons of simple polygons based on a bisector arrangement. *Inform. Process. Lett.* 132 (2018), 28–32.
- David Eppstein and Jeff Erickson. 1999. Raising roofs, crashing cycles, and playing pool: Applications of a data structure for finding pairwise interactions. *Discrete & Computational Geometry* 22, 4 (1999), 569–592.
- Petr Felkel and Stepan Obdrzalek. 1998. Straight skeleton implementation. In *Proceedings of Spring Conference on Computer Graphics*. Citeseer.
- Matthew Fisher, Manolis Savva, Yangyan Li, Pat Hanrahan, and Matthias Nießner. 2015. Activity-centric Scene Synthesis for Functional 3D Scene Modeling. *ACM Transactions on Graphics (TOG)* 34, 6 (2015).
- Lin Gao, Jie Yang, Tong Wu, Yu-Jie Yuan, Hongbo Fu, Yu-Kun Lai, and Hao(Richard) Zhang. 2019. SDM-NET: Deep Generative Network for Structured Deformable Mesh. *ACM Transactions on Graphics (TOG)* 38, 6 (2019), 243:1–243:15.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. *Advances in Neural Information Processing Systems* 27 (2014), 2672–2680.
- Martin Habbecke and Leif Kobbelt. 2012. Linear analysis of nonlinear constraints for interactive geometric modeling. In *Computer Graphics Forum*, Vol. 31. Wiley Online Library, 641–650.
- Martin Held and Peter Palfrader. 2017. Straight skeletons with additive and multiplicative weights and their application to the algorithmic generation of roofs and terrains. *Computer-Aided Design* 92 (2017), 33–41.
- Ruizhen Hu, Zeyu Huang, Yuhang Tang, Oliver van Kaick, Hao Zhang, and Hui Huang. 2020. Graph2Plan: Learning Floorplan Generation from Layout Graphs. *arXiv preprint arXiv:2004.13204* (2020).
- Caigui Jiang, Chengcheng Tang, Amir Vaxman, Peter Wonka, and Helmut Pottmann. 2015. Polyhedral Patterns. *ACM Transactions On Graphics (TOG)* 34, 6, Article 172 (Oct. 2015), 12 pages. <https://doi.org/10.1145/2816795.2818077>
- Tom Kelly, John Femiani, Peter Wonka, and Niloy J Mitra. 2017. BigSUR: large-scale structured urban reconstruction. *ACM Transactions On Graphics (TOG)* 36, 6 (2017).
- Tom Kelly, Paul Guerrero, Anthony Steed, Peter Wonka, and Niloy J Mitra. 2018. FrankenGAN: Guided detail synthesis for building mass models using style-Synchronized Gans. *ACM Transactions On Graphics (TOG)* 37, 6 (2018), 1–14.
- Tom Kelly and Peter Wonka. 2011. Interactive architectural modeling with procedural extrusions. *ACM Transactions on Graphics (TOG)* 30, 2 (2011), 1–15.
- Hyeongju Kim, Hyeonseung Lee, Woo Hyun Kang, Joun Yeop Lee, and Nam Soo Kim. 2020. SoftFlow: Probabilistic Framework for Normalizing Flow on Manifolds. *Advances in Neural Information Processing Systems* 33 (2020).
- Diederik P. Kingma and Max Welling. 2014. Auto-Encoding Variational Bayes. In *International Conference on Learning Representations (ICLR)*, Yoshua Bengio and Yann LeCun (Eds.).
- Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907* (2016).
- Mathieu Larive and Veronique Gaildrat. 2006. Wall grammar for building generation. In *Proceedings of the 4th international conference on Computer graphics and interactive techniques in Australasia and Southeast Asia*. 429–437.
- Robert G Laycock and AM Day. 2003. Automatically generating roof models from building footprints. (2003).
- Hui Lin, Jizhou Gao, Yu Zhou, Guiliang Lu, Mao Ye, Chenxi Zhang, Ligang Liu, and Ruigang Yang. 2013. Semantic decomposition and reconstruction of residential scenes from LiDAR data. *ACM Transactions on Graphics (TOG)* 32, 4 (2013), 1–10.
- Chen Liu, Jimei Yang, Duygu Ceylan, Ersin Yumer, and Yasutaka Furukawa. 2018. Planenet: Piece-wise planar reconstruction from a single rgb image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2579–2588.
- Yang Liu, Helmut Pottmann, Johannes Wallner, Yong-Liang Yang, and Wenping Wang. 2006. Geometric modeling with conical meshes and developable surfaces. In *ACM Transactions On Graphics (TOG)*. 681–689.
- Paul Merrell, Eric Schkufza, and Vladlen Koltun. 2010. Computer-generated residential building layouts. In *ACM Transactions On Graphics (TOG)*. 1–12.
- Kaichun Mo, Paul Guerrero, Li Yi, Hao Su, Peter Wonka, Niloy Mitra, and Leonidas Guibas. 2019. StructureNet: Hierarchical Graph Networks for 3D Shape Generation. *ACM Transactions on Graphics (TOG)* 38, 6 (2019), Article 242.
- Pascal Müller, Peter Wonka, Simon Haegler, Andreas Ulmer, and Luc Van Gool. 2006. Procedural modeling of buildings. In *ACM Transactions On Graphics (TOG)*. 614–623.
- Przemyslaw Musialski, Peter Wonka, Daniel G Aliaga, Michael Wimmer, Luc Van Gool, and Werner Purgathofer. 2013. A survey of urban reconstruction. In *Computer Graphics Forum*, Vol. 32. Wiley Online Library, 146–177.
- Liangliang Nan and Peter Wonka. 2017. Polyfit: Polygonal surface reconstruction from point clouds. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. 2353–2361.
- Charlie Nash, Yaroslav Ganin, S. M. Ali Eslami, and Peter W. Battaglia. 2020. PolyGen: An Autoregressive Generative Model of 3D Meshes. In *Proceedings of the 37th International Conference on Machine Learning (ICML) (Proceedings of Machine Learning Research)*, Vol. 119. PMLR, 7220–7229.
- Wamiq Reyaz Para, Paul Guerrero, Tom Kelly, Leonidas J. Guibas, and Peter Wonka. 2020. Generative Layout Modeling using Constraint Graphs. *CoRR abs/2011.13417* (2020).
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Eds.). Curran Associates, Inc., 8024–8035.
- Helmut Pottmann, Yang Liu, Johannes Wallner, Alexander Bobenko, and Wenping Wang. 2007. Geometry of multi-layer freeform structures for architecture. In *ACM Transactions On Graphics (TOG)*. 65–es.
- Helmut Pottmann, Alexander Schiftner, Pengbo Bo, Heinz Schmiedhofer, Wenping Wang, Niccolò Baldassini, and Johannes Wallner. 2008. Freeform surfaces from single curved panels. *ACM Transactions on Graphics (TOG)* 27, 3 (2008), 1–10.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. (2019).
- Anurag Ranjan, Timo Bolkart, Soubhik Sanyal, and Michael J Black. 2018. Generating 3D faces using convolutional mesh autoencoders. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 704–720.
- Ali Razavi, Aaron van den Oord, and Oriol Vinyals. 2019. Generating diverse high-fidelity images with vq-vae-2. In *Advances in Neural Information Processing Systems*. 14866–14876.
- Jing Ren, Jens Schneider, Maks Ovsjanikov, and Peter Wonka. 2018. Joint Graph Layouts for Visualizing Collections of Segmented Meshes. *IEEE Transactions on Visualization and Computer Graphics* 24, 9 (2018), 2546–2558.
- Danilo Rezende and Shakir Mohamed. 2015. Variational Inference with Normalizing Flows. In *International Conference on Machine Learning*. 1530–1538.
- Tim Salimans, Andrej Karpathy, Xi Chen, and Diederik P. Kingma. 2017. PixelCNN++: Improving the PixelCNN with Discretized Logistic Mixture Likelihood and Other Modifications. In *International Conference on Learning Representations (ICLR)*. Open-Review.net.
- David Salinas, Florent Lafarge, and Pierre Alliez. 2015. Structure-aware mesh decimation. In *Computer Graphics Forum*, Vol. 34. Wiley Online Library, 211–227.
- Michał Stypulkowski, Maciej Zamorski, Maciej Zięba, and Jan Chorowski. 2019. Conditional invertible flow for point cloud generation. *arXiv preprint arXiv:1910.07344* (2019).
- Kenichi Sugihara. 2013. Straight skeleton for automatic generation of 3-D building models with general shaped roofs. (2013).
- Kenichi Sugihara. 2019. Straight Skeleton Computation Optimized for Roof Model Generation. In *WSCG*, Vol. 27. 101–109.
- Qingyang Tan, Lin Gao, Yu-Kun Lai, and Shihong Xia. 2018. Variational autoencoders for deforming 3d mesh models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 5841–5850.
- Aaron Van den Oord, Nal Kalchbrenner, Lasse Espeholt, Oriol Vinyals, Alex Graves, et al. 2016. Conditional image generation with pixelcnn decoders. *Advances in Neural Information Processing Systems* 29 (2016), 4790–4798.
- Aaron Van Den Oord, Oriol Vinyals, et al. 2017. Neural discrete representation learning. In *Advances in Neural Information Processing Systems*. 6306–6315.
- Aaron Van Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. 2016. Pixel Recurrent Neural Networks. In *International Conference on Machine Learning*. 1747–1756.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*. 5998–6008.
- Yannick Verdie, Florent Lafarge, and Pierre Alliez. 2015. LOD generation for urban scenes. *ACM Transactions On Graphics (TOG)* 34, ARTICLE (2015), 30.
- Xinpeng Wang, Chandan Yeshwanth, and Matthias Nießner. 2020. SceneFormer: Indoor Scene Generation with Transformers. *arXiv:cs.CV/2012.09793*
- Jiajun Wu, Chengkai Zhang, Tianfan Xue, William T Freeman, and Joshua B Tenenbaum. 2016. Learning a probabilistic latent space of object shapes via 3D generative-adversarial modeling. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*. 82–90.
- Guandao Yang, Xun Huang, Zekun Hao, Ming-Yu Liu, Serge Belongie, and Bharath Hariharan. 2019. Pointflow: 3d point cloud generation with continuous normalizing flows. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. 4541–4550.
- Jie Yang, Kaichun Mo, Yu-Kun Lai, Leonidas J. Guibas, and Lin Gao. 2020. DSM-Net: Disentangled Structured Mesh Net for Controllable Generation of Fine Geometry. *arXiv:cs.GR/2008.05440*

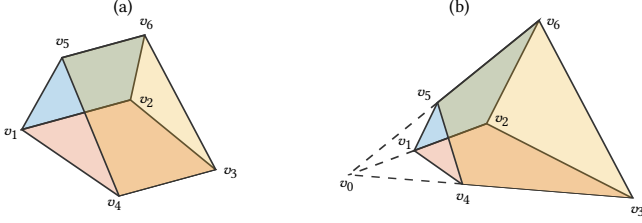


Fig. 32. Illustration of Remark 3.1

- Dawen Yu, Shunping Ji, Jin Liu, and Shiqing Wei. 2021. Automatic 3D building reconstruction from multi-view aerial images with deep learning. *ISPRS Journal of Photogrammetry and Remote Sensing* 171 (2021), 155–170.
- Lap Fai Yu, Sai Kit Yeung, Chi Keung Tang, Demetri Terzopoulos, Tony F Chan, and Stanley J Osher. 2011. Make it home: automatic optimization of furniture arrangement. *ACM Transactions on Graphics (TOG)* 30, 4 (2011).
- Huayi Zeng, Jiaye Wu, and Yasutaka Furukawa. 2018. Neural procedural reconstruction for residential buildings. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 737–753.
- Fuyang Zhang, Nelson Nauata, and Yasutaka Furukawa. 2020. Conv-mpn: Convolutional message passing neural network for structured outdoor architecture reconstruction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2798–2807.
- Qian-Yi Zhou and Ulrich Neumann. 2008. Fast and extensible building modeling from airborne LiDAR data. In *Proceedings of the 16th ACM SIGSPATIAL international conference on Advances in geographic information systems*. 1–8.
- Qian-Yi Zhou and Ulrich Neumann. 2010. 2.5 d dual contouring: A robust approach to creating building models from aerial lidar point clouds. In *Proceedings of the European Conference on Computer Vision (ECCV)*. Springer, 115–128.
- Qian-Yi Zhou and Ulrich Neumann. 2011. 2.5 D building modeling with topology control. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2489–2496.
- Lingjie Zhu, Shuhan Shen, Xiang Gao, and Zhanyi Hu. 2018. Large scale urban scene modeling from MVS meshes. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 614–629.

A PROOF OF REMARK 3.1

Recall Remark 3.1: The intersecting line of two adjacent 3D planar faces with fixed outline edges, is either parallel to both outline edges, or intersecting at the same point with the two outline edges. To prove this, we only need to discuss two settings: (1) the outline edges of the two adjacent 3D faces are parallel to each other (see case (a) in Fig. 32); (2) the outline edges of the two adjacent faces intersect with each other (see case (b) in Fig. 32).

For both cases in Fig. 32, we have two 3D planar faces $f_{1,2,6,5}$ and $f_{3,4,5,6}$, where $f_{1,2,6,5}$ has an outline edge (v_1, v_2) and $f_{3,4,5,6}$ has an outline edge (v_3, v_4) . We know that (v_1, v_2) and (v_3, v_4) are two outline edges that belong to the roof outline. Therefore, these two edges are co-planar and belong to the plane $f_{1,2,3,4}$. In case (a), we have the outline edge (v_1, v_2) being parallel to (v_3, v_4) . In case (b), we have the outline edge (v_1, v_2) intersect with (v_3, v_4) at point v_0 . We are supposed to show that, in case (a), the intersecting line $(v_5, v_6) = f_{1,2,6,5} \cap f_{3,4,5,6}$ is parallel to (v_1, v_2) and (v_3, v_4) ; and show that in case (b), the intersecting line (v_5, v_6) intersects with (v_1, v_2) and (v_3, v_4) at point v_0 . We will give the simple proof as follows.

We prove that $(v_5, v_6) \parallel (v_1, v_2)$ in case (a) by contradiction. Assume (v_5, v_6) is not parallel to (v_1, v_2) , i.e., (v_5, v_6) intersect with (v_1, v_2) at some point x . We then have $x \in (v_1, v_2)$ and $x \in (v_5, v_6) \in f_{3,4,5,6}$. Therefore, $x = (v_1, v_2) \cap f_{3,4,5,6}$, this is contradict to the fact that $(v_1, v_2) \parallel f_{3,4,5,6}$ since $(v_1, v_2) \parallel (v_3, v_4) \in f_{3,4,5,6}$. Therefore, our assumption does not hold. We then have $(v_5, v_6) \parallel (v_1, v_2) \parallel (v_3, v_4)$.

We then prove that in case (b) we have $v_0 \in (v_5, v_6)$. We already know that $v_0 = (v_1, v_2) \cap (v_3, v_4)$. Therefore, $v_0 \in (v_1, v_2) \in f_{1,2,6,5}$

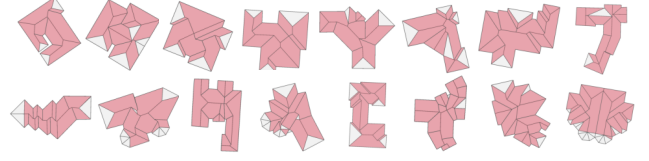


Fig. 33. Roofs constructed in 3ds Max. We highlight the roof faces that are not planar in red.

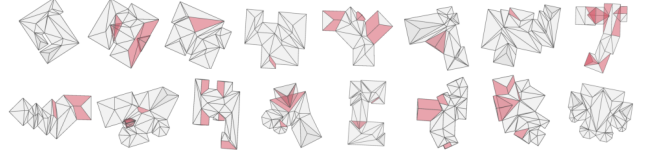


Fig. 34. Roofs constructed in SketchUp. We highlight the roof faces that are polygonal in red.

and $v_0 \in (v_3, v_4) \in f_{3,4,5,6}$. We then have $v_0 \in f_{1,2,6,5} \cap f_{3,4,5,6} = (v_5, v_6)$. This shows that the intersecting point v_0 belongs to the edge (v_5, v_6) . I.e., the intersecting edge intersects with the two outline edges at the same point. Note that the case (b) has a special case that two outline edges (v_1, v_2) and (v_1, v_3) intersect with each other at the same endpoint v_1 .

The sufficient condition can be similarly proved by contradiction.

B ALTERNATIVE PLANARITY METRICS

In our method, we propose to use the smallest eigenvalue of the covariance matrix of the vertices in each face as planarity metric (see Eq. (3)). There are other planarity metrics as well that could be considered: (1) one simple alternative is to measure the determinant of the covariance instead of the smallest eigenvalue; (2) to measure the planarity of a set of 3D points, we can first sample 3 points to form a plane, and then measure the distance from the points to the plane; (3) another commonly used planarity metric on quad meshes [Jiang et al. 2015] is to measure the distance between two diagonal lines in a quad. We can generalize this metric to general polygon meshes as well. Specifically:

$$E_{\text{det}} = \sum_{i=1}^{n_f} \det(\text{Cov}(X_{f_i})), \quad E_{\text{proj}} = \sum_{i=1}^{n_f} \sum_{x \in f_i} \text{dist}(x, P_{f_i}), \quad (5a,b)$$

$$E_{\text{diag}} = \sum_{i: f_i = (x_{i_1}, \dots, x_{i_p})} \sum_{j=1}^{p-3} \text{dist}(l_{x_{i_j}, x_{i_{j+2}}}, l_{x_{i_{j+1}}, x_{i_{j+3}}}) \quad (5c)$$

where in Eq. (5b), P_{f_i} is a plane formed by three sampled points on face f_i , and $\text{dist}(x, P)$ measures the projection distance from the point x to the plane P ; $\text{dist}(l_1, l_2)$ in Eq. (5c) measures the distance between two 3D lines l_1 and l_2 , and in our case, they are two diagonal lines connecting the vertices on the face.

In Fig. 35, we compare our planarity metric (Eq. (1)) to the three alternatives discussed above: we start with the same initial embedding (spectral initialization), use the same optimization solver (Quasi-Newton), and terminate w.r.t. the same criteria. We report the planarity error for each metric at different iterations in linear-scale (on the left) and log-scale (on the right). We also visualize the

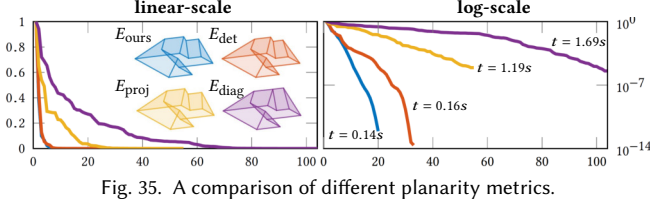


Fig. 35. A comparison of different planarity metrics.

optimized roof on the left and report the runtime on the right. We can see that, all these four planarity metrics are valid and lead to planar 3D roofs at the end. Our choice of using the smallest eigenvalue of the covariance matrix has a much simpler form and can converge faster measured by both the running time and the number of iterations.

C OPTIMIZE FOR A VALID 2D EMBEDDING

We can construct a 3D roof by optimizing for a valid 2D embedding according to Remark 3.1, then lifting up the valid 2D embedding to obtain a valid 3D roof. Specifically, we can obtain a valid 2D embedding \bar{X}_R by:

$$\min_{\bar{X}_R} \sum_{e \in \mathcal{E}_{\text{roof}}} \text{vad}(e) \quad (6)$$

where $\mathcal{E}_{\text{roof}}$ is the set of roof edges, and $\text{vad}(e)$ is a validity measure defined on the edge $e = (x_1, x_2)$ based on Remark 3.1. Specifically, assume the adjacent faces of the edge e have the outline edges e_1 and e_2 respectively. If e_1 is parallel to e_2 , then the edge e is valid if e is parallel to e_1 as well. In this case, we have $\text{vad}(e) = 1 - \langle \vec{e}, \vec{e}_1 \rangle^2$. If e_1 is not parallel to e_2 , and the two edges intersect at the point x , then this point x should be on the edge e as well, i.e., $x_1 - x$ is parallel to $x_2 - x$. In this case, $\text{vad}(e) = 1 - \left\langle \frac{x_1 - x}{\|x_1 - x\|}, \frac{x_2 - x}{\|x_2 - x\|} \right\rangle^2$ (here x_1, x_2, x are corresponding 2D positions).

D INTERACTIVE ROOF EDITING & OPTIMIZATION

Our roof modeling method can be user for interactive roof reconstruction: (1) modify/edit a valid roof graph G or its dual G^D (2) starting from the modified roof graph, run our optimization method to obtain a *valid* roof graph. We then go back to step (1) and edit again. In this way, we can edit the roof graph until we get a desirable one or the final valid roof graph is consistent with the input image. In the following, we first discuss some commonly used editing operations that are supported in step (1). We then discuss how to efficiently solve step (2) by only optimizing the position of the vertices in the *smallest affected region*.

Editing Operations. Our primal-dual roof graph formulation allows us to design editing operations to modify the roof graph or its dual directly:

- **Move a vertex.** We can modify the position of the selected vertex by an input 3D translation vector.
- **Move an edge.** We can modify the positions of the endpoint vertices of the selected edge by an input 3D translation vector.
- **Snap an edge.** We can snap an edge by merging its two endpoints into a single vertex.
- **Merge two faces.** For two faces that are adjacent to each other, we can merge them into a single face by removing the shared edge and reordering the vertices in the two faces.

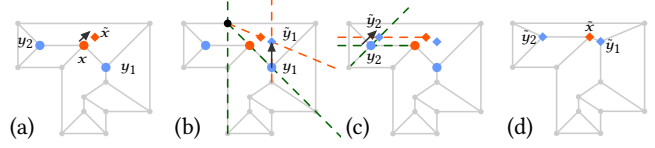


Fig. 36. **The smallest affected region.** (a) For a valid embedding of a roof graph, we change the position of the red vertex from x to \bar{x} . (b) We then need to modify the position of the right blue vertex from y_1 to \bar{y}_1 such that the line (\bar{x}, \bar{y}_1) intersecting with the two outline edges at the same point. (c) Similarly, we need to modify the position of the left blue vertex such that the intersecting line is parallel to the corresponding outline edges. (d) We then get a valid 2D embedding again after the change of x . Therefore, the smallest affected region of x is $\{y_1, y_2\}$.

- **Split a face.** We can also split a face into two by adding an extra edge to connect two non-adjacent vertices in the selected face.
- **Force two faces to be adjacent.** For two non-adjacent faces that are connected by an edge, we can force the two faces be adjacent.

Smallest affected region. After applying some editing operations, we get the modified roof graph G^{mod} with updated roof embedding X^{mod} , which is no longer valid. We need to run our roof optimization to enforce planarity constraint. Instead of rerunning our algorithm to update *all* vertex positions, we only need to update a *small set* of vertices (called "smallest affected region") to make the embedding valid again. Specifically, if we edit the position of some vertices in a valid embedding, we can detect the smallest group of vertices in the roof graph that need to be updated to satisfy the planarity constraint again by updating this group of vertices only. We therefore call this group of vertices the smallest affected region $P(x)$ of the modified vertex x . Fig. 36 illustrates how to detect the smallest affected region by investigating the validity of each roof vertex in the 2D embedding using Remark 3.1.

Then, we only need to minimize the planarity energy in this restricted region $P(x)$:

$$\min_{x_R \in P(x)} \sum_{i=1}^{n_f} \sigma_1(\text{Cov}(X_{f_i})) \quad (7)$$

The adoption of the smallest affected region has two main advantages for interactive editing: 1) less runtime of updating the positions of a smaller set of vertices instead of the complete vertex set. 2) more coherent embedding after the modification, which is more friendly for the users.