

Robust Low-Poly Meshing for General 3D Models

ZHEN CHEN, The University of Texas at Austin & LightSpeed Studios, USA

ZHERONG PAN, LightSpeed Studios, USA

KUI WU, LightSpeed Studios, USA

ETIENNE VOUGA, The University of Texas at Austin, USA

XIFENG GAO, LightSpeed Studios, USA

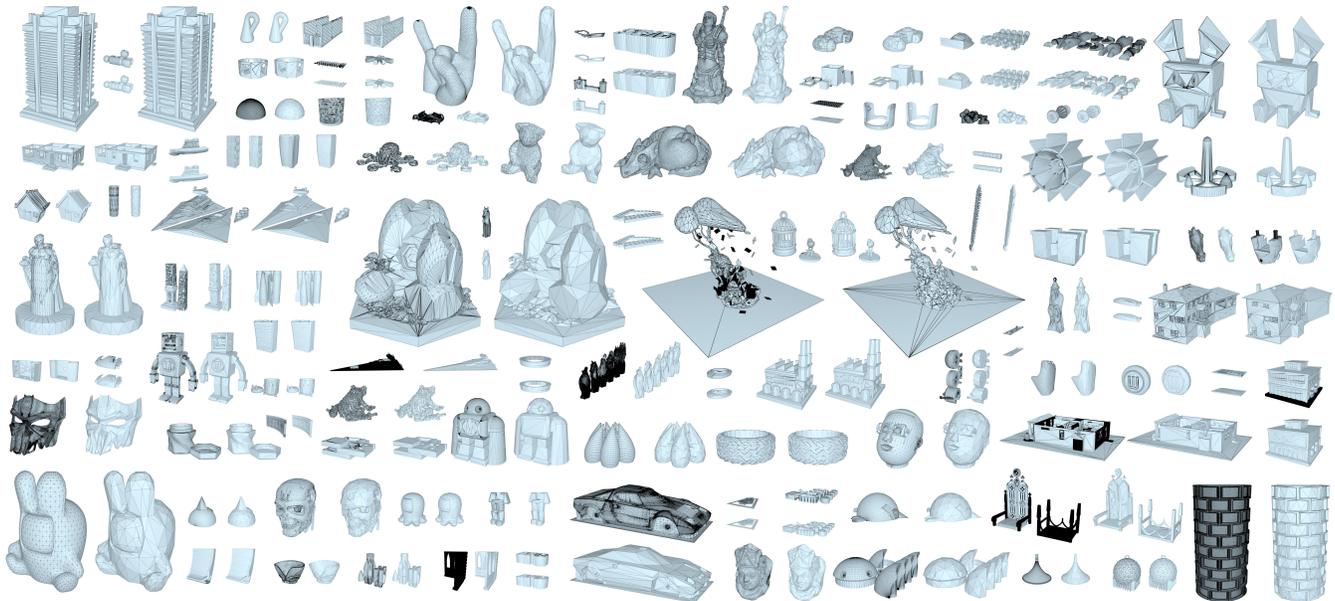


Fig. 1. A gallery of *wild* high-poly input meshes and their corresponding low-poly outputs generated by our method, where the low-polys are manifold, watertight, and self-intersection free, and have a small visual difference from their high-poly counterparts.

We propose a robust re-meshing approach that can automatically generate visual-preserving low-poly meshes for any high-poly models found in the wild. Our method can be seamlessly integrated into current mesh-based 3D asset production pipelines. Given an input high-poly, our method proceeds in two stages: 1) Robustly extracting an offset surface mesh that is feature-preserving, and guaranteed to be watertight, manifold, and self-intersection free; 2) Progressively simplifying and flowing the offset mesh to bring it close to the input. The simplicity and the visual-preservation of the generated low-poly is controlled by a user-required target screen size of the input: decreasing the screen size reduces the element count of the low-poly

Authors' addresses: Zhen Chen, The University of Texas at Austin & LightSpeed Studios, Austin, Texas, USA, 78712, zchen96@cs.utexas.edu; Zherong Pan, LightSpeed Studios, Bellevue, Washington, USA, 98004, zrpan@global.tencent.com; Kui Wu, LightSpeed Studios, Los Angeles, California, USA, 90066, kwu@global.tencent.com; Etienne Vouga, The University of Texas at Austin, Austin, Texas, USA, 78712, evouga@cs.utexas.edu; Xifeng Gao, LightSpeed Studios, Bellevue, Washington, USA, 98004, xifgao@global.tencent.com.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2023 Copyright held by the owner/author(s).

0730-0301/2023/1-ART1

<https://doi.org/10.1145/3592396>

but enlarges its visual difference from the input. We have evaluated our method on a subset of the Thingi10K dataset that contains models created by practitioners in different domains, with varying topological and geometric complexities. Compared to state-of-the-art approaches and widely used software, our method demonstrates its superiority in terms of the element count, visual preservation, geometry, and topology guarantees of the generated low-polys.

CCS Concepts: • **Computing methodologies** → **Mesh geometry models**.

Additional Key Words and Phrases: iso-surface extraction, remeshing, geometry processing

ACM Reference Format:

Zhen Chen, Zherong Pan, Kui Wu, Etienne Vouga, and Xifeng Gao. 2023. Robust Low-Poly Meshing for General 3D Models. *ACM Trans. Graph.* 1, 1, Article 1 (January 2023), 20 pages. <https://doi.org/10.1145/3592396>

1 INTRODUCTION

Mesh is a ubiquitously employed representation of 3D models for digital games. While a mesh with a large number of polygons (high-poly) is required to express visually appealing details, rendering its low-poly approximation at distant views is a typical solution to achieve real-time gaming experience, especially on low-end devices. High-polys, no matter whether they are manually created

through modeling software or automatically converted from CSG and implicit functions, often have complex topology and geometries, such as numerous components, high genus, non-manifoldness, self-intersections, degenerate elements, gaps, inconsistent orientations, etc. These complexities can pose significant challenges to the design of automatic low-poly mesh generation algorithms.

Over past decades, two typical ways are developed to obtain low-poly textured models: automatic *mesh reduction* that preserves original textures [Hoppe 1999; Liu et al. 2017]; manual mesh modeling followed by UV-generation and texture baking that creates new textures. Mesh reduction usually removes triangles through iterative application of local operations [Garland and Heckbert 1997] or element clustering [Cohen-Steiner et al. 2004], which relies on existing mesh vertices and connectivity. As a result, this method is only suitable for generating the medium-levels of LOD, while introducing serious artifacts when generating low-polys for meshes with excessive topology complexity as illustrated in Fig. 12 and Fig. 13. In the second pipeline, while UV-generation and texture baking can be done via semi-automatic tools (e.g. Maya and Marmoset), manual meshing is the most labor(cost)-intensive step. Therefore, we address this most urgent and challenging problem, aka, low-poly meshing.

Many automatic *re-meshing* approaches exist to represent the original mesh with a proxy and simplify the proxy mesh via a row of different techniques, e.g., polygonal mesh construction by plane fitting and mixed-integer optimization [Nan and Wonka 2017], cage mesh generation through voxel dilation and mesh simplification [Calderon and Boubekeur 2017], shape abstraction by feature simplification [Mehra et al. 2009], extremely low-poly meshing using visual-hull boolean operations [Gao et al. 2022], mesh simplification through differentiable rendering [Hasselgren et al. 2021], enclosing mesh generation through alpha-wrapping with an offset [Portaneri et al. 2022], and learning based approaches [Chen et al. 2020, 2022b]. However, these re-meshing approaches either rely on heavy user interactions, need careful parameter tweaking, or work for a limited type of model. Commercial software also has low-poly mesh functions, but generates unsatisfactory results in many cases. For example, in Fig. 11, and Table 3, we show the generated results using different low-poly construction modules in Simplifyon [AB 2022]. It generates meshes with either triangle intersections, or non-satisfactory visual appearances. It remains to be a challenging task to automatically and robustly generate low-poly meshes for general 3D models used in the industry.

In practice, artists still manually craft low-poly meshes to ensure that they have a small number of triangles and preserve the visual appearance of the original mesh as much as possible. This often involves multiple iterations of manual adjustments, which incurs intensive labor work and prolonged project periods and remains to be a bottleneck for the current fast-changing game industry. Robust and automatic approaches that can generate satisfactory low-polys are in high demand.

From an input mesh with arbitrary topology and geometry properties, our goal is to generate its low-poly counterpart that is visually indistinguishable from faraway views. The visual appearance of a 3D shape can be evaluated by its silhouette and surface normal, while the simplicity of a low-poly mesh is typically measured by the

number of triangles. We propose a new approach to generate low-polys that is both simple and visual-preserving, with the additional guarantees of being manifold, watertight, and self-intersection free. These additional properties are essential for artists to conveniently perform UV-generation and texture baking on the bare low-poly mesh.

Our method combines the idea of mesh-reduction and re-meshing. During the first stage, we re-mesh the high-poly to a “clean” proxy mesh and remove all the topological complexities therein. We then aggressively reduce the element count of the proxy mesh, while geometrically deforming the proxy to maintain visual preservation, leading to our low-poly output.

Our method specifically requires two inputs: a high-poly mesh and a parameter n_p , which represents the screen-space size of the high-poly mesh when rasterized onto a 2D screen. In practical terms, let l_p denote the 2D screen’s pixel length, and l represent the diagonal length of the high-poly mesh’s bounding box. The parameter n_p can be calculated as l/l_p , signifying the maximum number of pixels that the high-poly mesh’s diagonal could occupy across all potential rendering views. During the first stage of our method, we build an unsigned distance field for the input and introduce a novel offset surface extraction method to extract a d -isosurface with $d = l/n_p$. Our algorithm not only guarantees the offset mesh is manifold, watertight, and self-intersection-free, but also recovers the normal approximated sharp features. During our second stage, we alternate among three steps, i.e. mesh simplification, mesh flow process, and feature alignment, to reduce the element count of the extracted mesh, while bringing the mesh close to the input and maintaining the aforementioned guarantees. All three steps contain only local operations, such as edge collapse and vertex optimization. Therefore, any local operation that violates a guarantee can be easily rolled back.

By construction, our algorithm is robust and automatic. The effectiveness of our approach is demonstrated by comparing it with state-of-the-art methods and popularly used software on a subset of the Thingi10K dataset [Zhou and Jacobson 2016] containing 3D models with varying complexities. All the data shown in the paper and the executable program can be found here¹.

2 RELATED WORKS

We first review low-poly mesh generation methods and then summarize the methods for iso-surface extraction.

2.1 Low-poly Meshing

Obtaining a low-poly mesh has been a research focus in computer graphics for several decades. Early works use various mesh reduction techniques that directly operate on the original inputs through iterative local element removal. Examples involve geometric error-guided techniques [Garland and Heckbert 1997; Hoppe 1996; Lescot et al. 2020], structure-preserving-constrained technique [Saliinas et al. 2015], volume-preserving technique [Lindstrom and Turk 1998], and image-driven technique [Lindstrom and Turk 2000], to name just a few. Clustering-based approaches [Cohen-Steiner et al. 2004; Li and Nan 2021] provide another direction for reducing the

¹<https://robust-low-poly-meshing.github.io/>

element count. An inclusive survey is given in Khan et al. [2022]. While these approaches are well recognized in game production pipelines, they are better suited for reducing the mesh size of the original models to a medium level, e.g. reducing the number of faces by 20% - 80%. Unfortunately, for 3D graphics applications running on lower-end devices, often a much coarser low-poly mesh is desired. Such extremely low-poly meshes require topologic and geometric simplifications that are far beyond the capabilities of these mesh reduction techniques. Unlike mesh reduction techniques, a parallel effort aims at re-meshing, i.e., completely reconstructing a new mesh mimicking the original one. These methods vary drastically in their techniques and we classify them by their main feature into *voxelization base re-meshing*, *primitive fitting*, *visual-driven*, and *learning-based*.

Voxelization Based Re-meshing. Both Mehra et al. [2009] and Calderon and Boubekeur [2017] rely on a voxelization of the raw inputs to obtain a clean voxel surface. While Mehra et al. [2009] requires feature-guided re-triangulation, deformation, and curve-network cleaning to generate shape abstractions for architectural objects, Calderon and Boubekeur [2017] assumes the input meshes come with clear separation of the inside and outside space and heavily depends on user interactions to generate the final low-polys. Recently, Wu et al. [2022] combine voxelization-based remeshing with patch-based simplification to generate occluders for building models to pre-cull unseen meshes before online rendering.

Primitive Fitting. Various primitives can be composed to fit an object. For example, methods in [Bauchet and Lafarge 2020; Chauve et al. 2010; Fang and Lafarge 2020; Fang et al. 2018; Kelly et al. 2017; Nan and Wonka 2017] first compute a set of planes to approximate patch features detected in point clouds or 3D shapes, and then select a faithful subset of the intersecting planes to obtain the desired meshes. However, the key challenges of this type of method are: 1) properly computing a suitable set of candidate planes is already a hard problem by itself; 2) the complexity of the resulting mesh is highly unpredictable, requiring many trial-and-error to find a possible good set of parameters. Works using other primitives [Huang et al. 2014; Mehra et al. 2009; Wei et al. 2022; Yang and Chen 2021], such as boxes, convex shapes, curves, etc., are also promising directions, but none of them has been specifically dedicated for generating low-polys.

Visual-driven Approaches. Differentiable rendering [Laine et al. 2020] rises as a hot topic that enables the continuous optimization of scene elements through the guidance of rendered image losses. However, most of them [Hasselgren et al. 2021; Luan et al. 2021; Nicolet et al. 2021] require an initial mesh that is typically a uniformly discretized sphere. The key obstacle to generating low-polys via differentiable rendering is that the mesh reduction cannot be modeled as a differentiable optimization process. Although the analysis-by-synthesis type of optimizations [Luan et al. 2021] could be employed, the Laplacian regularization term used by most differentiable rendering techniques can guide the mesh far from the groundtruth, especially in a low-poly setting. A visual hull-based approach [Gao et al. 2022] has been recently proposed to generate extremely low-polys for building models, however, it not only

creates sharp creases for organic shapes, but also makes it hard to control the target element number.

Learning-based Methods. A conventional 3D mesh reconstruction pipeline is composed of three steps: plane detection, intersection, and selection, while learning-based methods enable alternative pipelines. As an example, by converting it to a BSP-net, Chen et al. [2020] demonstrates that low-polys can be extracted from images. However, this method shares the common shortcomings of learning approaches: a large dataset is required for the network training, and the learned model works only for meshes of a similar type. It further requires the voxelizations of the dataset to have well-defined in/out segmentation. Furthermore, the generated meshes inherit the issues of polyfit-like [Nan and Wonka 2017] approaches: it creates sharp creases that are not present in the high-poly, and parameter tuning is difficult. By embedding a neural net of marching tetrahedral into the differentiable rendering framework, Munkberg et al. [2022] can optimize the meshes and materials simultaneously. As demonstrated in their work, by controlling the rendered image resolution, it can generate 3D models in a LOD manner. But these extended features brought by learning approaches are beyond the scope of our investigation.

2.2 Iso-surfacing Algorithms

The marching cubes (MC) algorithm was proposed concurrently by Lorensen and Cline [1987] and Wyvill et al. [1986] for reconstructing iso-surfaces from discrete signed distance fields. Several follow-up works were proposed to solve the tessellation ambiguities in each cube [Chernyaev 1995; Dürst 1988; Matveyev 1994; Nielson 2003, 2004; Nielson and Hamann 1991]. One of the best methods is MC33 [Chernyaev 1995], which enumerates all possible topologic cases based on trilinear interpolation in the cube. The follow-ups resolve non-manifold edges in MC33 [Custodio et al. 2013; Lopes and Brodlié 2003]. MC33 was correctly implemented by Vega et al. [2019] after resolving the defective issues of the previous implementations [Custodio et al. 2013; Lewiner et al. 2003]. However, none of these methods is able to recover sharp features.

To capture sharp features of the iso-surface, Kobbelt et al. [2001] first introduced an extended marching cubes method (EMC) to insert additional feature points, given that the normals of some intersecting points are provided. Dual contouring (DC) [Ju et al. 2002] adapted this idea with Hermite data (the gradient of the implicit surface function). They proposed to insert one dual feature point inside a cube and then connect the dual points to form an iso-surface. DC does not need to perform the edge-flip operations required by [Kobbelt et al. 2001], but often generates non-manifold surfaces with many self-intersections. The non-manifold issue was later addressed in [Ju and Udeshi 2006], while the self-intersection issue was addressed in [Schaefer et al. 2007]. However, none of these two methods solves both the non-manifold and self-intersection problems simultaneously. Dual Marching Cubes (DMC) [Schaefer and Warren 2004] considers that the dual grid aligns with features of the implicit function, and extracts the iso-surface from the dual grid. DMC can preserve sharp features without excessive grid subdivisions as required by DC. Unfortunately, DMC still does not guarantee the generated mesh is self-intersection-free. Manson and

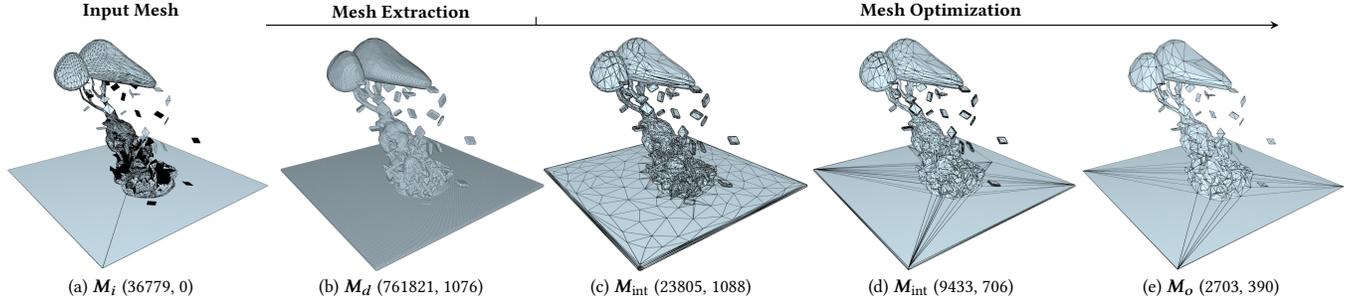


Fig. 2. The pipeline of our algorithm. The notation (\bullet, \bullet) represents (number of faces, light field distance to the input mesh), where the latter is a measure of visual similarity between two 3D shapes, introduced by Chen et al. [2003]. (a): The input high-poly surface, which is not necessarily manifold, orientable, or self-intersection free. In this example, the input surface M_i is not 2-manifold, with 114 non-2-manifold edges, has 751 components and is not orientable (the back faces are rendered in black). It has over 36k faces, among which 23345 faces are self-intersected. (b): The extracted iso-surface M_d with $d = \frac{l}{200}$, where l is the bounding box diagonal size of M_i (Section 3.1). It is an orientable, water-tight, self-intersection free mesh with 19-components and 761k faces. (c-e) Our mesh optimization step (Section 3.2), during which we apply mesh simplification, flow, and alignment. While the simplification step may result in a slight increase in LFD, the subsequent flow and alignment steps enhance visual similarity. Consequently, the overall optimization step progressively reduces the light field distance and simplifies the mesh, and the intermediate meshes denoted as M_{int} . The output M_o has only 2703 faces. Our approach resolves the existing topologic (non-manifoldness) and geometric issues (self-intersection), and approximates the high-poly with 7.3% faces.

Schaefer [2010] avoided self-intersections by subdividing each cube into multiple tetrahedra, and then applying marching tetrahedra (MT) to extract the iso-surface [Doi and Koide 1991]. This approach solves the self-intersection problems in the DMC approach, but the division of multiple tetrahedra, together with the employed octree structure, makes the algorithm either generate an overly dense mesh or require trial-and-error for suitable octree depth parameter settings. A survey about these approaches can be partially found in [de Araújo et al. 2015]. Recently, Portaneri et al. [2022] proposed an algorithm to generate watertight and orientable surfaces that strictly enclose the input. Their output is obtained by refining and carving the 3D Delaunay triangulation of the offset surface, however, still without the feature-preserving property.

There are also several learning-based approaches for iso-surface extraction. Deep marching cubes [Liao et al. 2018] and deep marching tetrahedra [Shen et al. 2021] learn differentiable MC and MT results. However, none of them can capture the sharp features of the initial surface. Neural marching cubes [Chen and Zhang 2021] and Neural dual contouring [Chen et al. 2022a] train the network to capture the sharp features without requiring extra Hermite information. However, the former generates self-intersected meshes, and the latter leads to non-manifold results. In Table 1, we summarize these methods and show their strength and weakness in terms of topologic and geometric properties: manifoldness, self-intersection-free, and sharp feature preservation.

3 METHOD

Given the input of a polygonal mesh M_i , a maximum number of screen size n_p (i.e. the number of pixels covered by the diagonal length of the input’s bounding box), and an optionally user-specified target number of triangles n_F , we seek to generate a triangle mesh M_o with the following properties:

²Although the initial paper results in non-manifold edges, this artifact was fixed by the follow-up works [Custodio et al. 2013; Lopes and Brodlie 2003]

Table 1. A Brief summary of the existing methods by their capabilities of maintaining geometry and topology properties. A more detailed survey can be found in [de Araújo et al. 2015].

Method	Manifold	Free of Self-Intersection	Preserve Features
Lorensen and Cline [1987]	✓	✓	×
Wyvill et al. [1986]	✓	✓	×
Chernyaev [1995]	✓ ²	✓	×
Doi and Koide [1991]	✓	✓	×
Kobbelt et al. [2001]	✓	×	✓
Ju et al. [2002]	×	×	✓
Ju and Udeshi [2006]	✓	×	✓
Schaefer et al. [2007]	×	✓	✓
Manson and Schaefer [2010]	✓	✓	✓
Portaneri et al. [2022]	✓	✓	×
Liao et al. [2018]	✓	×	×
Shen et al. [2021]	✓	×	×
Chen and Zhang [2021]	✓	×	✓
Chen et al. [2022a]	×	✓	✓

- Pro.I The number of triangles in M_o is either minimized or equals to n_F if provided as a parameter;
- Pro.II M_o is indistinguishable from M_i when rendered from a far-away view (a view where the diagonal length of the bounding box of M_i is less than n_p pixels);
- Pro.III M_o is both topologically and geometrically clean, i.e., water-tight, manifold, and intersection-free.

These three properties of M_o ensure rendering quality and enable any downstream geometric processing on it to have high computational efficiency, requiring no mesh repairing steps. The level of visual preservation in our second property is measured by Silhouette difference and the normal difference between M_i and M_o . A similar normal indicates M_o preserves the sharp features of M_i as much as possible.

We follow several principles to design our approach: 1) We make no assumptions on the topologic or geometric properties of the input, allowing our approach to handle any models created in the wild; 2) We adopt an interior-point optimization-like strategy to realize the topology and geometry properties of Pro.III one by one: once a property is satisfied, it will be maintained for the rest of the steps; 3) We value robustness with the highest priority, so that our approach can process any inputs created by different domains of applications. Under guaranteed robustness, we further attempt to improve the computational efficacy to the greatest extent possible.

Overview. We tackle this problem in two main stages (Fig. 2), namely, *mesh extraction* (Section 3.1), and *mesh optimization* (Section 3.2). During the mesh extraction stage, we first compute an unsigned distance field for M_i , then introduce a novel iso-surface mesh extraction approach for a positive offset distance d ($d = l/n_p$ as mentioned in Section 1), and finally remove all invisible disconnected components from the extracted iso-surface to obtain a mesh M_d . Our generated M_d optimally recovers the sharp features implied by the d -iso-surface of the distance field, and guarantees watertightness, manifoldness, and free of self-intersections. The purpose of this stage is to generate a “clean” proxy mesh M_d of the input M_i that possibly has “dirty” topologic and geometric configurations. Our second mesh optimization stage involves a while-loop of three sequential steps: simplification, flow, and alignment. The simplification step aims to reduce the number of triangles of M_d by performing one pass of quadric edge-collapse decimation for the entire mesh; the flow step aims to pull M_d close to M_i via a per-vertex distance minimization; and the alignment step aims to optimize the surface normal of M_d so that the sharp features are maintained, which is achieved through local surface patch optimization. When the while loop stops, we output the final mesh M_o . Since all three steps contain only local operations, the guarantees of M_d achieved during the first stage can be easily maintained by rolling back or skipping any operations that violate a guarantee. In the following sections, we provide technical details for each stage.

3.1 Mesh Extraction

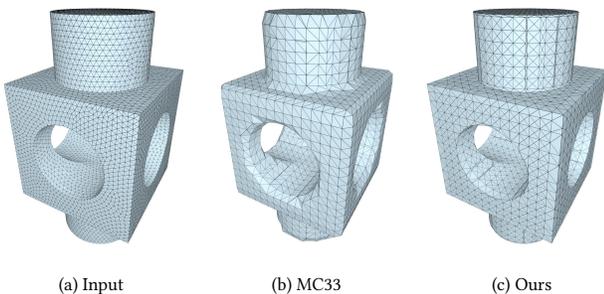


Fig. 3. The iso-surface meshes of a CAD model “block”. Compared with classic MC33 algorithm [Chernyaev 1995], our iso-surfacing approach achieves a better visual similarity by recovering sharp features.

Given M_i and d , our goal is to extract an d -iso-surface mesh M_d that is watertight, manifold, feature-preserving, and self-intersection-free. Ensuring all these properties simultaneously is a challenging task. As an example, simply applying the well-known algorithm of MC33 cannot capture sharp features of the iso-surface as shown in Fig. 3. We tackle the mesh extraction problem by re-meshing the extracted local surface patches from templates of MC33. We selectively insert additional points to refine these local surface patches. Our key technique lies in the proposed mesh refinement technique that 1) guarantees topologically watertight and manifold properties and 2) captures geometric sharp features without causing self-intersection. We first compute a proper *discretization* of an unsigned distance function defined for M_i , then analyze the connectivity changes when inserting new points to the MC33 templates to maintain the *topology guarantees* of the resulting mesh. After that, we focus on the *geometry fidelity* of the resulting mesh, i.e. feature-preserving and self-intersection-free. For the extracted mesh M_d , we finally remove invisible components.

Discretization. An unsigned distance field is defined as a function:

$$f(\mathbf{p}) := \min_{q \in M_i} \|\mathbf{p} - \mathbf{q}\|, \quad (1)$$

where $\mathbf{p} \in \mathcal{R}^3$. The implicit function of d -iso-surface is $f(\mathbf{p}) = d$. Since the explicit representation of the d -iso-surface is intractable, we follow the general pipeline of prior iso-surfacing approaches that first voxelize the ambient space around M_i , and then approximate the solution through extracting local surface patches within each voxel. Since the local patches are typically simple, the voxel size plays an important role in the to-be-extracted mesh. A coarse voxel size can miss important solutions, such as the one illustrated in Fig. 4, where no d -iso-surface could be extracted if a grid size larger than $2d$ is used, while an excessively small voxel size will result in a dense grid that is time-consuming to compute (Fig. 15). By default, we choose the edge length of a voxel to be $d/\sqrt{3}$ to avoid geometric feature losses as illustrated in the Fig. 4.



Fig. 4. The d -iso-surface (green lines) of the input mesh (red line) cannot be captured if the voxel size is larger than $2d$.

Topologic Guarantees. For each voxel, we employ existing templates to decide the iso-contours [Chernyaev 1995; Custodio et al. 2013; Lorensen and Cline 1987], and then insert an additional point for each contour. The templates of either the original MC [Lorensen and Cline 1987] or MC33 [Chernyaev 1995; Custodio et al. 2013] can be used to generate the iso-contours since they both ensure the extracted mesh is watertight and manifold. We choose MC33 in this work since it covers more linear interpolation cases and resolves the ambiguity in MC, thus extracting iso-surface meshes with generally smaller genus [Chernyaev 1995; Custodio et al. 2013]. As illustrated in Fig. 23, we insert one vertex per iso-contour surface if it is homomorphic to a disk, where the iso-contour surfaces of cases 4.1.2, 6.1.2, 7.4.2, 10.1.2 and 12.1.2, and one iso-contour of case 13.5.2 are excluded since they have two boundaries. This scheme

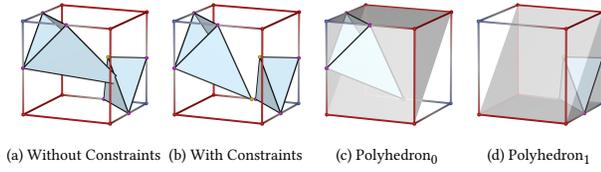


Fig. 5. Illustration of our feature-aware iso-surface extraction step for MC33 case 4.1.1. The cube vertices with iso-value smaller than d is marked as blue, while the red vertices are the opposite. The iso-points are marked as pink, and the feature points are marked as yellow. Without forcing the feature points within their belonging polyhedra (c, d), it is easy to obtain a mesh with self-intersections (a).

ensures that the additional vertices neither bring any non-manifold configurations nor create holes in the resulting mesh.

Feature Vertex Insertion. We use one vertex per local patch with a disk-topology to provide more degrees of freedom to capture sharp features. Its position can be computed by minimizing the distance to patch vertices along the normal directions:

$$\arg \min_{\mathbf{x}} \sum_i (\mathbf{n}_{\mathbf{p}_i} \cdot (\mathbf{x} - \mathbf{p}_i))^2, \quad (2)$$

where \mathbf{x} is the desired position, \mathbf{p}_i and $\mathbf{n}_{\mathbf{p}_i}$ denote an iso-contour vertex and its normal, respectively, and the summation goes through all patch vertices. However, without any constraints, \mathbf{x} may be arbitrarily positioned and cause surface intersections in the extracted mesh. This issue is deteriorated by the template cases with multiple iso-contours. Fig. 5 illustrates such as example using MC33 case 4.1.1.

We propose a simple approach to recover feature vertices without mesh intersections. Given a voxel, we subdivide it into convex polyhedra within which the feature vertices are constrained. As illustrated in Fig. 24, the subdivision is performed according to the number and the different configurations of the iso-contours with a disk-topology. For example, for cases with only a single disk-topology iso-contour, such as case 1, no subdivision is involved and the polyhedron is the voxel itself, and for those with multiple iso-contours, such as case 4.1.1, convex and planar polygons are needed to partition the voxel into non-overlapping polyhedra. If we constrain the position of the inserted vertex stays inside its belonging polyhedron, the extracted mesh is guaranteed to incur no self-intersections.

Accordingly, for each inserted vertex, we obtain its coordinate \mathbf{x} by solving a linear constrained quadratic programming:

$$\arg \min_{\mathbf{x}} \sum_i (\mathbf{n}_{\mathbf{p}_i} \cdot (\mathbf{x} - \mathbf{p}_i))^2 \quad (3)$$

$$s.t. \quad \mathbf{n}_s \cdot (\mathbf{x} - \mathbf{c}_f) < 0, \quad \forall \mathbf{n}_s \in N_s$$

where N_s is the set of face normals of the corresponding polyhedron of \mathbf{x} , and \mathbf{c}_f is the center of the corresponding polyhedron face. To compute $\mathbf{n}_{\mathbf{p}_i}$, if $\mathbf{p}_i \in \mathbf{M}_i$, we simply use the mesh normal, otherwise, we first solve for any $\mathbf{p}_i^* \in \arg \min_{\mathbf{p} \in \mathbf{M}_i} \|\mathbf{p}_i - \mathbf{p}\|$ and then let $\mathbf{n}_{\mathbf{p}_i} := (\mathbf{p}_i - \mathbf{p}_i^*) / \|\mathbf{p}_i - \mathbf{p}_i^*\|$.

Feature Extraction. The previous step recovers vertices on sharp features of the d -iso-surface. But their connections may not align well with the sharp edges, the highlighted region in Fig. 6b demonstrates this issue. Furthermore, since the sharp features exist in a small fraction of voxels, we aim at a minimal increase in the additional feature edges and vertices by inserting only those feature vertices on sharp features and using the original MC33 templates as much as possible. However, we do not know the sharp features of the d -iso-surface as prior. Therefore, we introduce a posterior approach to recover the necessary feature curves, which involves two phases: *Feature Edge Adjustment*, and *Feature Filtering*. The first phase generates an iso-surface mesh by considering all inserted feature vertices as sharp features. With this iso-surface mesh, we can use existing automatic feature identification approaches to obtain sharp features. The second phase generates the actual iso-surface mesh \mathbf{M}_d by blending the iso-contour patches containing the detected feature vertices with those original MC33 patches that do not contain any sharp features.

Feature Edge Adjustment. During the first phase, we insert a feature vertex for every disk-topologic iso-contour patch. We then perform an edge-flip operation for every mesh edge if the flipped edge connects two inserted feature vertices (see Algorithm 2 for more details). To ensure the self-intersection-free guarantee, we skip those edge-flips that may cause self-intersections. This phase can already produce an iso-surface mesh that satisfies the desired topologic and geometric properties. However, as mentioned earlier, this iso-surface mesh contains more elements than desired and those unnecessary “fake” sharp features are noisy and not visual-appealing (see the top two zoomed-in figures in Fig. 6c).

Feature Filtering. During the second phase, we first extract a feature graph from the resulting mesh of the first phase. The feature graph is composed of a set of feature curves where each curve is a sequence of mesh edges with its dihedral angle smaller than θ_0 (see [Gao et al. 2019] for details). We then mark a feature curve as valid if it is composed of more than l_0 mesh edges. The valid feature curves are considered to contain “real” sharp features to recover. After that, for each iso-contour patch, we keep those inserted feature vertices if they are on the valid feature curves, otherwise, we use their original template. This step removes a lot of noisy, “fake” feature edges. Finally, we perform the edge flip algorithm Algorithm 2 once more to extract the d -iso-surface mesh \mathbf{M}_d .

Our feature recovery algorithm performs well for the models with various features that can be represented by piecewise line segments, e.g. sharp curves in Fig. 3 and the eye and beak contours in Fig. 12.

Interior Removal. Since we use an unsigned distance function, our final extracted iso-surface \mathbf{M}_d may have interior components, which are totally invisible. Given the generated mesh \mathbf{M}_d are watertight and free of self-intersection, we can apply the in-and-out test and remove the components which is purely inside of any of the others.

3.2 Mesh Optimization

Starting from a mesh \mathbf{M}_d that is watertight, manifold, feature preserving, and self-intersection-free, we now introduce an iterative mesh optimization approach to obtain a final \mathbf{M}_o that satisfies our

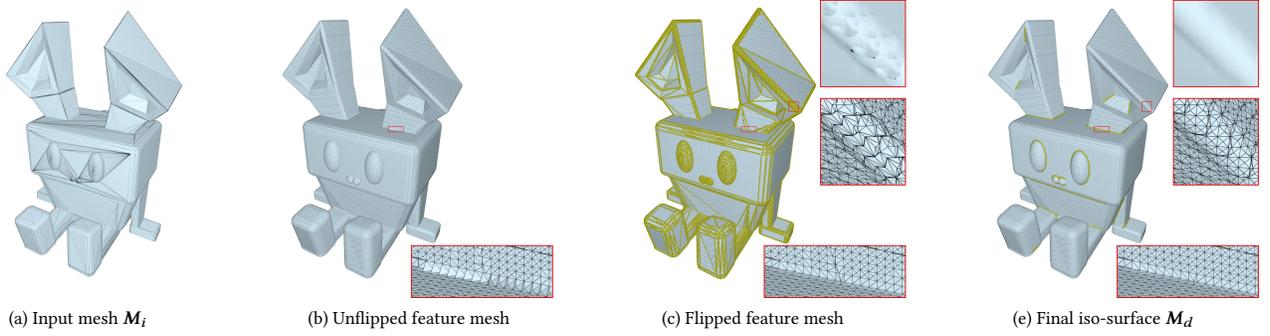


Fig. 6. Our iso-surface extraction pipeline. We mark the initial feature points in the third figure and the remaining ones after applying feature graph filter in the last figure (the yellow points). Our post process successfully resolves the sawtooth artifacts around the ear of the character (the top two zoomed-in figures in the third and last columns, with the top ones are rendered without wireframes), but still keeps the major sharp features, for example, the bottom zoomed-in figures.

Algorithm 1 Iso-surface Extraction

Input: M_i, d, θ_0, l_0
Output: M_d

- 1: $G \leftarrow \text{gridDiscretization}(M_i, d)$ ▷ generate the grids
- 2: Compute $f(\mathbf{p})$ for all grid points \mathbf{p} in G ▷ Equation 1
- 3: **for each** $\text{cube} \in G$ **do** ▷ iso-surface extraction
- 4: Lookup for the template case ▷ [Chernyaev 1995], Fig. 23
- 5: **for each** Disk-topologic patch in cube case **do**
- 6: Compute the iso-points on cube edges
- 7: Form the quadratic program ▷ Equation 3, Fig. 24
- 8: Solve for feature vertices ▷ Fig. 23
- 9: **end for**
- 10: **end for**
- 11: $M_d \leftarrow \text{edgeFlip}(M_d)$ to connect feature vertices ▷ Algorithm 2
- 12: $M_d \leftarrow \text{featureFilter}(M_d, \theta, l_0)$
- 13: $M_d \leftarrow \text{removeInterior}(M_d)$

Algorithm 2 Edge Flip

Input: M_d
Output: M_d ▷ mesh after edge-flips

- 1: $Q \leftarrow \{\}$
- 2: **BVHTree** T
- 3: $T.\text{build}(M_d)$ ▷ [Karras 2012]
- 4: **for each** edge $e \in M_d$ **do**
- 5: **if** $e.\text{oppVs}$ are feature vertices **then**
- 6: $Q.\text{push}(e)$ ▷ opposite vertices are feature vertices
- 7: **end if**
- 8: **end for**
- 9: **while** $Q \neq \{\}$ **do**
- 10: $e \leftarrow Q.\text{top}()$
- 11: **if** e was not flipped before **then**
- 12: **if** $\text{isIntersectionFree}(M_d, T, e)$ **then**
- 13: $\text{flipEdge}(M_d, e)$
- 14: $T.\text{refit}(M_d)$ ▷ update BVH [Karras 2012]
- 15: **end if**
- 16: **end if**
- 17: **end while**

three desired properties, i.e., Pro.I-III. As shown in Algorithm 3, our optimization involves a maximum of N iterations of three sequential steps: *simplification*, *flow*, and *alignment*. We stop the iterations until either the Hausdorff distance between the simplified meshes of two consecutive loops (relative change) is smaller than ϵ , the loop number reaches N , or the target face number reduces below n_F , where the first condition has the highest priority by default. Each of the three steps involves only local operations. To ensure our optimization proceeds towards the generation of M_o with the desired properties, we perform the following checks for the meshes before and after applying a local operation:

- (1) Topology consistency: the updated mesh is manifold, watertight, and has the same genus and the number of components as the mesh before applying the local operation;
- (2) Self-intersection-free: the updated mesh is free of intersections.

Mesh Simplification. This step aims to achieve the first property of M_o , i.e., M_o contains as few triangles as possible. We perform an entire pass of the standard edge-collapse operation for all edges of M_o to reduce as many faces as possible or match the target face number n_F , where the coordinate of the newly generated vertices are determined by the quadratic edge metric (QEM) [Garland and Heckbert 1997] weighted by virtual planes for each edge to avoid the degeneracy in planar regions. Importantly, the topologic and geometric validity of M_o is maintained during the simplification process by skipping those edge-collapse operations that may violate the aforementioned checks. Moreover, to ensure we get closer to M_i , we also skip the collapse operations which increase the distance between affected local triangle patches and M_i . For efficiency concerns, we only compute the one-sided distance from the local patch to the input mesh. This one-sided check may result in the acceptance

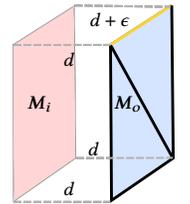


Fig. 7. Collapsing the yellow edge reduces the distance from M_o to M_i ($d + \epsilon \rightarrow d$), but leads to an undesirable visual appearance.

of unexpected collapses, as illustrated in Fig. 7. To overcome this, we further skip the operations leading to a Hausdorff distance larger than d , where the two involving meshes are the ones before and after the local operation and the Hausdorff distance is computed approximately by sampling points on the local triangle patches as in [Cignoni et al. 1998]. We show the comparison in Fig. 8. Without the guarantee of a distance decrease, we will lose some important information. Without the guarantee of a small Hausdorff distance, we may end up with larger silhouette difference and normal difference, that is, worse visual similarity.

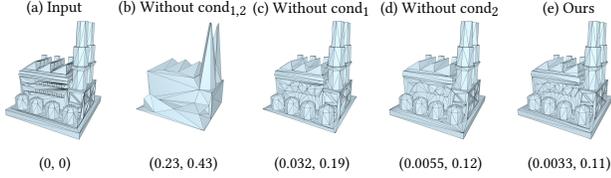


Fig. 8. The results of different simplification conditions. (\bullet, \bullet) denotes (silhouette difference, normal difference). cond_1 : skip the collapse which increases the vertex-surface distance to M_i ; cond_2 : skip the collapse which results a large Hausdorff distance. After applying the both conditions, we achieve a better visual score.

Mesh Flow. Our mesh flow step brings M_o geometrically close to M_i and reduces the silhouette visual differences between the two meshes. The detailed algorithm is provided in Algorithm 3 (Line 7-11).

Algorithm 3 Mesh Optimization Process

Input: $M_i, M_d, d, n_F, N, r, \epsilon$
Output: M_o

- 1: $M_o \leftarrow M_d$
- 2: $l \leftarrow \text{bboxSize}(M_i)$ ▷ bounding box diagonal size
- 3: **for** $i = 0$ to N **do**
- 4: $M' \leftarrow M_o$
- 5: $M_o \leftarrow \text{meshSimplification}(M_i, M_o, d, n_F)$ ▷ Algorithm 5
- 6: $\tilde{M} \leftarrow M_o$
- 7: **for each** vertex $v \in M_o$ **do** ▷ mesh flow step
- 8: $v^* \leftarrow \text{argmin}_{u \in M_i} \|u - v\|$
- 9: $d_v \leftarrow r(v^* - v)$ ▷ successive flow, $r < 1$
- 10: $v \leftarrow \text{localUpdate}(M_o, v, d_v)$ ▷ Algorithm 4
- 11: **end for**
- 12: **for each** vertex $v \in M_o$ **do** ▷ feature alignment step
- 13: $v_{\text{opt}} \leftarrow \text{featureAlignment}(\tilde{M}, M_o, v)$ ▷ Algorithm 6
- 14: $d_v \leftarrow v_{\text{opt}} - v$
- 15: $v \leftarrow \text{localUpdate}(M_o, v, d_v)$ ▷ Algorithm 4
- 16: **end for**
- 17: **if** $\text{Hausdorff}(M_o, M') < \epsilon \cdot l$ **then**
- 18: **Break** ▷ update is small enough
- 19: **end if**
- 20: **end for**

When actually applying the mesh flow process, for each vertex v in M_o , we find its Euclidean-distance-wise closest point v^* of M_i

Algorithm 4 Local Update

Input: M, v, d_v
Output: updated v
Note: v is a vertex of M .

- 1: $\alpha \leftarrow 1$
- 2: **while** $v + \alpha d_v$ leads to self-intersections **do**
- 3: $\alpha \leftarrow \alpha/2$
- 4: **end while**
- 5: **return** $v + 0.95\alpha d_v$ ▷ Using 0.95 to avoid numerical error

Algorithm 5 Mesh Simplification

Input: M_i, M_o, d, n_F
Output: simplified mesh M_o
Notes: M_i is the reference mesh, n_F is optional

- 1: Form priority queue Q ▷ [Garland and Heckbert 1997]
- 2: **BVHTree** T
- 3: **T.build**(M_o) ▷ [Karras 2012]
- 4: **while** $Q \neq \{\}$ **do**
- 5: $e \leftarrow Q.\text{top}()$
- 6: **if** e has been visited before **then**
- 7: **continue** ▷ has been collapsed
- 8: **end if**
- 9: **if** $\text{topologyConsistencyCheck}(M_o, e)$ failed **then**
- 10: **continue** ▷ [Cignoni et al. 2008]
- 11: **end if**
- 12: **if not** $\text{isIntersectionFree}(M_o, T, e)$ **then**
- 13: **continue** ▷ collapse will introduce intersections
- 14: **end if**
- 15: $M_e \leftarrow$ sub-mesh of M_o adjacent to e
- 16: $M'_e \leftarrow M_e$ after collapse
- 17: **if** $\text{dist}(M_e \rightarrow M_i) < \text{dist}(M'_e \rightarrow M_i)$ **then**
- 18: **continue** ▷ collapse increases the distance to M_i
- 19: **end if**
- 20: **if** $\text{dist}(M_e \rightarrow M'_e) > d$ or $\text{dist}(M'_e \rightarrow M_e) > d$ **then**
- 21: **continue** ▷ distance update is too large
- 22: **end if**
- 23: $\text{collapseEdge}(M_o, e)$ ▷ satisfy all desired properties
- 24: $T.\text{refit}(M_o)$ ▷ update BVH [Karras 2012]
- 25: **if** n_F is given and $M_o.\text{faceNumber} \leq n_F$ **then**
- 26: **break**
- 27: **end if**
- 28: **end while**
- 29: **return** M_o

and successively push v to v^* along the vector $d_v = v^* - v$. Instead of updating v to v^* directly, we deform v towards v^* based on a constant fractional ratio r of the vector, which allows more moving space for the entire mesh and reduces the chance of optimization stuck when M_o is still far from M_i . We also apply a simple line search for the self-intersection-free check to find the maximum step size during the local deformation (Algorithm 4).

Algorithm 6 Feature Alignment**Input:** \tilde{M}, M_o, v **Output:** updated v which preserves features**Require:** \tilde{M} and M_o has the same mesh connectivity

```

1: for each  $f \in N^1(v)$  do                                ▶ loop over adjacent faces
2:    $\mathbf{n}_f \leftarrow e_0 \times e_1$                           ▶ unnormalized face normal of  $M_o$ 
3:    $c_n \leftarrow \|\mathbf{n}_f\|$                                ▶ get the initial norm
4:    $\tilde{\mathbf{n}}_f \leftarrow \tilde{e}_0 \times \tilde{e}_1$                     ▶ unnormalized face normal of  $\tilde{M}$ 
5: end for
6: Fix  $c_n$ , get  $v_{\text{opt}}$  by minimizing Equation 5 ▶ quadratic program
7: return  $v_{\text{opt}}$ 

```

Feature Alignment. The previous mesh flow can stretch the mesh unanimately, breaking features and creating dirty inputs for subsequent mesh simplification and flow procedure (see Fig. 9 for an example). We thus introduce a feature alignment step. For each vertex v , we seek an optimized position v_{opt} by minimizing the shape difference between the local surface of v_{opt} and that of v before mesh flow:

$$E(v) := \sum_{f \in N^1(v)} \left\| \frac{\mathbf{n}_f}{\|\mathbf{n}_f\|} - \frac{\tilde{\mathbf{n}}_f}{\|\tilde{\mathbf{n}}_f\|} \right\|^2, \quad (4)$$

where we use the normal disagreement to approximate the shape difference (Line 6 in Algorithm 6), $N^1(v)$ is the faces within the 1-ring neighbor of v , and $\mathbf{n}_f, \tilde{\mathbf{n}}_f$ are the unnormalized face normal of the current mesh and the one before the flow respectively. The summation takes over all faces within the 1-ring neighborhood of vertex v . This face normal difference summation approximates the vertex normal difference. Notice that Equation 4 is a nonlinear function, which can be solved by the classical Newton's Method. In order to improve the efficiency, we instead treat the $\|\mathbf{n}_f\|$ as constant (equal to the value at the beginning of the alignment step, denoted as c_n) and solve a quadratic approximation of Equation 4:

$$E(v) := \sum_{f \in N^1(v)} \left\| \frac{\mathbf{n}_f}{c_n} - \frac{\tilde{\mathbf{n}}_f}{\|\tilde{\mathbf{n}}_f\|} \right\|^2. \quad (5)$$

Once we obtain the corresponding v_{opt} that minimizes Equation 5, we update v to be v_{opt} with the line search Algorithm 4 to prevent self-intersections. Given this local operation only slightly updates the mesh, our quadratic approximation leads to small errors, but in turn, significantly boosts performance (turning a non-convex problem to an unconstrained quadratic program).

3.3 Self-intersection Check Acceleration

Starting from an intersection-free 3D triangle mesh, our low-poly re-meshing pipeline could introduce intersections when performing the edge flips during mesh extraction, the edge collapses during mesh simplification, and the vertex optimization during the mesh flow and the feature alignment steps. Note that we say a mesh has intersections when any of its two triangles overlaps, or any of its two non-adjacent triangles touch or intersect.

Before discussing our accelerated check of self-intersections, we first introduce the necessary notations. For a vertex v , we denote

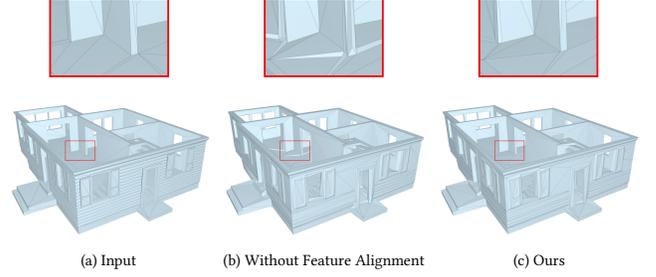


Fig. 9. Without feature alignment step, we will end up with the results with "spikes" (see zoomed-in region for details).

$N^i(v)$ as the set of all triangles that are bounded within its i -th ring neighborhood. For example, for the bottom left image in Fig. 10, $N^1(v)$ is the red region, and $N^2(v)$ is the union of red and green region³. We further denote M_e as the local neighborhood related to a certain local operation, where M_e endows different definitions. For edge flip, we define $M_e = \{f_1, f_2\}$ where f_1 and f_2 are the two neighboring triangles of e . For edge collapse, $M_e = N^1(v')$ where v' is the newly created vertex. For vertex optimization (otherwise known as smoothing), we let $M_e = N^1(v)$. Moreover, we let M_s be the sub-mesh formed by all faces that share at least one vertex with M_e but not in M_e (the green regions in Fig. 10). Finally, we let $M_1 = M_s \cup M_e$. The rest of the mesh are denoted as M_r (the blue regions in Fig. 10).

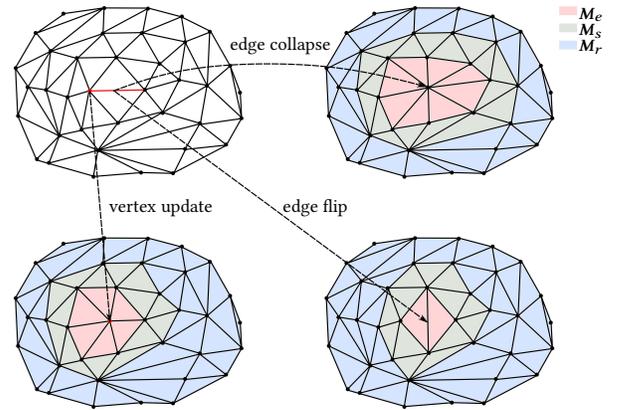


Fig. 10. Edge or vertex operation illustration. $M_1 = M_e \cup M_s$.

We note that a mesh-reduction operation does not introduce self-intersection iff the following two conditions hold:

- (1) M_e does not intersect M_r ;
- (2) M_1 is self-intersection free.

Here we skip the intersection check within M_s , M_r and between M_s and M_r , because M_s and M_r are the unchanged sub-mesh of the mesh before the local operation, which is free of self-intersections. In general, the two conditions above can be check by conventional

³These sets of triangles are called "topological neighborhoods", introduced in [Attene 2010]

Table 2. A speedup summary of self-intersection checks (used in edge flip, edge collapse, and vertex optimization steps) for some of the figures shown in the paper. The upper index of the figures indicates the corresponding row of that figure. T_n : the time cost of the neighboring triangle intersection check only using surface normal test. T'_n : the same time information with full normal cone test (surface normal + contour test). T_n^* : the same time information but applying parallel triangle pairs intersection check. T_t, T'_t, T_t^* : the total time information of the whole intersection-check process (neighboring triangle intersection check + BVH check), corresponding to T_n, T'_n, T_n^* .

Figures	T_n (s)	T'_n (s)	T'_n/T_n	T_n^* (s)	T_n^*/T_n	T_t (s)	T'_t (s)	T'_t/T_t	T_t^* (s)	T_t^*/T_t
Fig. 2	23.80	2587.09	108.71	111.59	4.69	119.31	2689.11	22.54	213.63	1.79
Fig. 6	16.15	2847.69	176.31	119.82	7.42	115.86	2955.23	25.51	224.50	1.94
Fig. 11(d) ⁰	21.98	2653.21	120.70	158.18	7.20	120.35	2769.56	23.01	266.17	2.21
Fig. 12(g) ⁰	13.51	2409.66	178.40	70.66	5.23	93.06	2497.08	26.83	156.08	1.68
Fig. 12(g) ¹	10.24	1377.05	134.51	43.98	4.30	57.25	1421.92	24.84	87.36	1.53
Fig. 13(j) ⁰	58.93	7958.25	135.05	552.99	9.38	360.12	8291.89	23.03	885.49	2.46
Fig. 13(j) ¹	5.92	819.00	138.37	39.56	6.68	29.79	846.21	28.41	65.04	2.18
Fig. 14(d)	13.92	2284.73	164.08	106.71	7.66	99.87	2382.30	23.86	200.90	2.01
Fig. 22(a) ²	15.65	2283.99	145.93	96.42	6.16	98.12	2376.82	24.22	187.28	1.91
Average	20.01	2802.30	144.67	144.43	4.69	121.52	2914.46	24.69	254.05	1.79

triangle-triangle intersection test. However, checking the first condition above is computationally inefficient, especially when M_r contains lots of triangles. Given M_e does not share any vertex or edge with M_r , this part can be handled by standard BVH-based collision detection. The detailed algorithm is given in Algorithm 7.

Algorithm 7 BVH Meshes Intersection Check

Input: M_e, M_r, T (BVH tree of M_r)

Output: whether M_e intersects with M_r

Notes: all faces M_e do not share vertices with the faces in M_r

```

1: for each face  $f \in M_e$  do
2:    $f_1 \leftarrow T.\text{closestFace}(f)$             $\triangleright$  get the closest face
3:   if  $\text{triTriIntersection}(f, f_1)$  then
4:     return true                            $\triangleright$  does intersect
5:   end if
6: end for
7: return false                              $\triangleright$  does not intersect

```

Unfortunately, for the second case, all the faces in M_e share at least one vertex with the faces in M_s . The BVH-based acceleration is no longer efficient, as the shared features always lead to failure in BVH culling. In this scenario, the naive approach involves $|M_e| \cdot |M_s|$ pairs of triangle-triangle intersection check. Although $|M_e|$ and $|M_s|$ are usually small for one local operation, the three edge flip, edge collapse, and vertex optimization operations will typically be executed for a massive number of times during the entire re-meshing pipeline. In practice, we find that this M_1 intersection-free check takes $\sim 50\%$ of the computational time of the whole intersection check process. Avoiding unnecessary triangle-triangle intersection checks, which is expensive to compute, will lead to a dramatic speedup. To this end, we note that M_1 is open, and to check whether a mesh with boundaries has self-intersection, Volino and Thalmann [1994] introduce a theory providing a sufficient condition: Let M be a continuous surface, bounded by ∂M , M is self-intersection free if there exists a vector \mathbf{n} , such that:

- (1) *Surface Normal Test:* For every point $\mathbf{p} \in M$, $\mathbf{n}_p \cdot \mathbf{n} > 0$, where \mathbf{n}_p is the surface normal at \mathbf{p} ;

- (2) *Contour Test:* The projection of the contour ∂M along the \mathbf{n} is not self-intersected.

They also provide a discrete version for triangle meshes:

- (1) *Surface Normal Test:* The angle of the normal cone formed by all triangle face normals is less than $\frac{\pi}{2}$;
- (2) *Contour Test:* The projection of the mesh boundary ∂M along the normal cone axis is not self-intersected.

For the first test, one can use the tight normal cone merging algorithm mentioned by Han et al [2021], and for the second test, Wang et al [2017] proposed a side-sign based unprojected contour test.

Surface normal test only need $|M_1|$ times normal cone expansion [Han et al. 2021]. As shown in Table 2, only applying *Surface Normal Test* results in $\sim 145\times$ speedup compared with the full normal cone test, and $4.69\times$ speed up compared with parallel triangle-triangle pair check. Moreover, this normal cone test acceleration speeds up the whole self-intersection check process by $24.69\times$ and $1.79\times$ compared with the full normal cone test and parallel triangle-triangle pair check, respectively. Surface normal test alone in practice is enough to generate a surface without self-intersection. We perform only the surface normal test during the self-intersection check, and if it fails, we apply direct triangle-triangle pair checks. Although the surface normal test alone is not sufficient to ensure free of self-intersection of M_1 , in practice, we find our final output M_o is always self-intersection free. We also perform a self-intersection check of M_o . If M_o intersects itself, we remove the surface normal test filter, and re-run the algorithm with direct triangle-triangle pair checks.

4 EXPERIMENTS

We implement our algorithm in C++, using Eigen for linear algebra routines, CGAL [Brönnimann et al. 2022] for exact triangle-triangle intersection check, libigl [Jacobson et al. 2018] for basic geometry processing routines. We use the fast winding number [Barill et al. 2018] for interior components identification. We implement the bottom-up BVH traversal algorithm mentioned in [Karras 2012] to refit the BVH for self-intersection check, and use Metro [Cignoni et al. 1998] for Hausdorff distance computation. Unless particularly mentioned, we set $n_p = 200$, $\theta_0 = 120^\circ$, $l_0 = 4$, $N = 50$, $r = \frac{1}{8}$, and

Table 3. Statistics of the results generated for the entire dataset by all comparing low-poly meshing approaches, including the number of vertices (#V), the number of triangles (#F), the number of components (#C), the ratios between the number of meshes being self-intersection-free (r_f), manifold (r_m), watertight (r_w), successfully generated (r_s) and the 100 models in the dataset, and the average (ave) and standard deviation (sd) of the four visual preservation metrics, i.e., PSNR, LFD, SD, ND, and HD. We treat a case as a failure if the algorithm terminated with an exception (marked as \star), or reaches the timeout threshold (1h, marked as \dagger).

Methods	#V	#F	#C	r_f	r_m	r_w	r_s	PSNR		LFD		SD		ND		HD	
								ave	sd	ave	sd	ave	sd	ave	sd	ave	sd
Simplygon ¹	703	1631	8	9.0%	37.0%	34.0%	100%	24.79	2.73	844.94	1806.07	0.013	0.016	0.049	0.030	0.024	0.016
Simplygon ²	763	1631	2	62%	93.0%	93.0%	100%	25.05	2.41	347.50	188.68	0.0048	0.0043	0.040	0.064	0.022	0.020
Blender	842	1803	12	11.0%	19.0%	14.0%	100%	23.68	3.51	1220.74	1637.12	0.030	0.058	0.071	0.092	0.040	0.052
QEM	707	1629	9	5.0%	12.0%	10.0%	100%	25.18	2.87	748.98	1059.45	0.012	0.021	0.041	0.049	0.031	0.023
Gao et al.	458	912	2	60.7%	91.0%	91.0%	89.0% [†]	22.55	2.63	1254.85	3978.50	0.020	0.055	0.063	0.059	0.078	0.073
KSR	727	1471	7	2.1%	2.1%	2.1%	96.0% [†]	22.96	3.31	3108.90	8138.43	0.058	0.12	0.089	0.13	0.029	0.020
PolyFit	69	55	5	0%	88.9%	0.0%	54.0% \star	17.31	1.41	6173.04	25941.63	0.29	0.15	0.51	0.16	0.29	0.156
Ours ^P	214	592	1	100%	100%	100%	100%	18.67	2.19	3696.22	2700.15	0.14	0.14	0.24	0.17	0.15	0.11
TetWild	753	1611	5	63.0%	32.0%	32.0%	100%	24.26	2.85	1932.26	7011.41	0.029	0.094	0.062	0.11	0.050	0.083
fTetWild	773	1643	4	73.7%	38.9%	38.9%	95.0% \star	24.21	2.77	2195.83	9562.01	0.037	0.13	0.069	0.14	0.059	0.12
ManifoldPlus	747	1610	3	24.0%	66.0%	64.0%	100%	25.14	2.49	559.84	1674.21	0.0060	0.0070	0.042	0.070	0.026	0.020
AlphaWrapping ¹	804	1631	1	93.0%	100%	100%	100%	23.18	2.28	667.40	369.27	0.018	0.0085	0.059	0.06	0.037	0.024
AlphaWrapping ²	743	1510	1	100%	100%	100%	100%	25.06	2.66	327.32	184.00	0.0046	0.0058	0.042	0.068	0.032	0.024
Ours ^Q	760	1631	2	58.0%	100%	100%	100%	22.90	2.23	716.04	413.22	0.020	0.0090	0.060	0.067	0.029	0.020
Ours	760	1631	2	100%	100%	100%	100%	25.21	2.49	310.30	169.58	0.0045	0.0045	0.037	0.067	0.022	0.020

$\epsilon = 10^{-4}$ by default and run our experiments on a workstation with a 32-cores Intel processor clocked at 3.7Ghz and 256Gb of memory, and we use TBB for parallelization.

Dataset. We test our algorithm on a subset of Thingi10K [Zhou and Jacobson 2016], where we randomly choose 100 models while filtering out those with the number of triangles smaller than 5000. For this dataset, the average number of faces and disconnected components are 120k and 10. The average number of non-manifold edges and self-intersected triangle pairs are 2197 and 6729, respectively.

4.1 Metrics

We evaluate the generated low-poly meshes from several aspects, including the number of contained triangles, topology (watertightness and manifoldness) and geometry (self-intersection-free) guarantees, and the visual preservation of the input.

Similarity Metrics. For visual similarity measurement, we employ the following metrics:

- (1) Hausdorff distance (HD), used to measure the geometrical distance between two 3D shapes;
- (2) Light field distance (LFD) [Chen et al. 2003], which measures the visual similarity between two 3D shapes;
- (3) Silhouette and normal differences [Gao et al. 2022], denoted as SD and ND respectively;
- (4) Peak signal-to-noise ratio (PSNR) computed by rendering the high- and low-poly meshes with 48 camera views and averaging the PSNR of the 48 pairs of images.

Among all these metrics, a smaller HD, LFD, SD, or ND indicates a better visual similarity, while for PSNR the higher the better.

4.2 Comparisons

To demonstrate the effectiveness of our approach, we compare against ten competing methods, including two modules of the state-of-the-art commercial solution—Simplygon [AB 2022], denoted as Simplygon¹ and Simplygon², four academic approaches, and four

baselines by combining mesh repairing and simplification. Since only Simplygon² cannot exactly control the element count of the generated mesh, we compare all methods by matching the element counts of their results to those generated by Simplygon² with 200 as its parameter value.

Comparison with Commercial Software. Simplygon [AB 2022] can automatically generate simplified meshes and is popularly used by game studios. We compare our approach with both its mesh reduction (Simplygon¹) and re-meshing (Simplygon²) modules. As shown in the Simplygon¹ and Simplygon² rows of Table 3, Simplygon can robustly process all meshes in the dataset, while Simplygon² generates better results than Simplygon¹ from basically all aspects but still introduces self-intersections and non-manifoldness for some models. In comparison, our approach not only guarantees the outputs are topologically clean and free of surface intersections, but also preserves the visual appearance much better, e.g., with 63.2% and 10.7% higher LFD, on average over the tested dataset than Simplygon^{1,2} respectively. Fig. 11 illustrates their visual comparisons on two models.

Comparison with Academic Approaches. We compare our algorithm with three state-of-the-art low-poly mesh generation methods, i.e., PolyFit [Nan and Wonka 2017], KSR [Bauchet and Lafarge 2020] and Gao et al. [2022], and two typically used mesh simplification approaches, i.e., QEM module in MeshLab [Cignoni et al. 2008] and the Blender decimation modifier. For PolyFit and KSR, we use the uniform sampling filter in MeshLab [Cignoni et al. 2008] to simple 1M points on the input mesh. We use the built-in PolyFit API in CGAL with default parameters for final mesh generation, For the KSR method, in accordance with the authors' suggestion, we utilize the plane-extraction approach proposed by Yu and Lafarge [2022] and subsequently employ KSR for surface reconstruction. For all of these, we use the executable program provided in the authors'

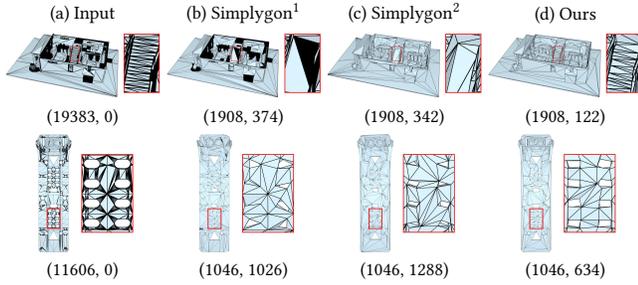


Fig. 11. Comparison with Simplygon. (\bullet, \bullet) denotes (face number, light field distance). Notice that, some input meshes may have inconsistent face orientations, such as the mesh shown in the top row where back faces are rendered in black. From the zoomed-in regions in the top row, only our method keeps the features of the stairs. From the bottom zoomed-in row, our approach has the best match of the input.

website⁴ with default parameters. Note that PolyFit often generates meshes with much fewer triangles than the target value. In this case, we further simplify our algorithm to match the triangle numbers of their outputs, which are denoted as Ours^P. In contrast, KSR generates more triangles than the target value. In this case, we apply a post QEM step to simplify its output to the target triangle number. For QEM [Cignoni et al. 2008], we first try to match the target triangle number with the topology preservation option turned on. We then turn it off if the simplification cannot reduce the element count to the desired value. As shown in Table 3, PolyFit fails to generate results for 46 out of 100 models due to the failure of planar feature detection, which is a challenge by itself; KSR and Gao et al. [2022]’s approach fail to provide any results for 4 and 13 out of 100 models respectively, within the computing time limit of 1h; QEM [Cignoni et al. 2008] and Blender generate considerably worse results in terms of topology and geometry guarantees. As shown in Table 3, our approach not only has geometrical and topological guarantees, but also achieves the best visual similarity scores, with an LFD 95.0%, 90.0%, 58.6%, 74.5%, and 75.3% smaller than those generated by PolyFit, KSR, QEM, Blender, and Gao et al. [2022], respectively. This behaves similarly to the other metrics. We also demonstrate some visual results in Fig. 12.

Comparison with Alternative Pipeline. One alternative approach for low-poly meshing is to first repair the input surface to get a high-quality surface mesh through the various mesh repair methods [Diazi and Attene 2021; Hu et al. 2020, 2018; Huang et al. 2020; Portaneri et al. 2022], then apply a mesh simplification step (for example QEM [Cignoni et al. 2008]) to reduce the element count to a specific number. We also show the comparison between our approach and four variants of this two-step process, i.e., TetWild [Hu et al. 2018] + QEM, fTetWild [Hu et al. 2020] + QEM, ManifoldPlus [Huang et al. 2020] + QEM, AlphaWrapping [Diazi and Attene 2021] + QEM (AlphaWrapping¹). For QEM, we first turn on the topology and normal preservation options, and set the target face number as the one from Simplygon². If the QEM fails to simplify the

mesh under these conditions, we turn off topology and normal options and simplify the mesh again. It is worth noting that removing the interior of other mesh repairing results did not affect the results since most of these approaches, such as TetWild [Hu et al. 2018], and AlphaWrapping [Portaneri et al. 2022], have either implicitly or directly removed the interior. However, ManifoldPlus [Huang et al. 2020] produced inconsistent face orientations, requiring more complex interior removal approaches that we are not aware of. As demonstrated in Table 3 and Fig. 13, the main drawback of this idea is that, although mesh-repairing approaches can fix the mesh to some extent, the follow-up simplification step will break the desired properties especially when the desired element count is small. For example, although ManifoldPlus [Huang et al. 2020] and AlphaWrapping [Portaneri et al. 2022] generate manifold and watertight mesh, respectively, the following simplification step breaks these guarantees. Notice that some mesh repairing methods also introduce issues in the meshes. For example, there are lots of self-intersections in ManifoldPlus’s output of the first example of Fig. 13. AlphaWrapping [Portaneri et al. 2022] always generates a self-intersection-free surface, but it does not capture the sharp features in the input mesh, which leads to undesired visual appearances after simplification. TetWild [Hu et al. 2018] and fTetWild [Hu et al. 2020] can generate meshes with non-manifold configurations, which could be further repaired to be manifold at the cost of surface intersections [Attene et al. 2009].

Additionally, we experiment with replacing parts of our algorithm with alternative methods. For instance, combining our mesh extraction with QEM (Ours^Q) leads to significantly inferior results compared to our original approach. Furthermore, when combining our mesh optimization with other mesh repair methods, such as AlphaWrapping (denoted as AlphaWrapping²), the results exhibit comparable LFD values but worse HD outcomes.

To sum up, comparing to these baseline variants, our method ensures the generated mesh is topologically clean, geometrically self-intersection-free, and visually appearance preserving.

Timings. We take about 7 minutes on average to finish the re-meshing of the entire tested dataset while the others take less than 2 minutes, except for Gao et al. [2022] (over 10 minutes) and the KSR method (over 15 minutes). In Fig. 17, we further analyze the time costs of different stages of our approach: the edge flip step in the iso-surface extraction step (t_e); the other iso-surface steps (t_i); interior removal step (t_r); mesh simplification step (t_s); and the others, like flow, alignment, and I/O (t_o). We find that the most time-consuming part is the edge-flip (in iso-surface extraction) and mesh simplification (in mesh optimization) with self-intersection checks involved. It turns out that these two parts take over 80% of our entire process (See Fig. 17), of which self-intersection check takes over 70% of the time. For this reason, the more faces the extracted iso-surface, the more edge flip and collapse operations will be conducted, ultimately leading to a higher cost. The right image of Fig. 17 also reveals this positive correlation.

4.3 Parameters

Screen Size n_p . We conduct a performance analysis in terms of user-specified screen size n_p and the corresponding iso-values d

⁴GoCopp and KSR in <https://team.inria.fr/titane/software/>

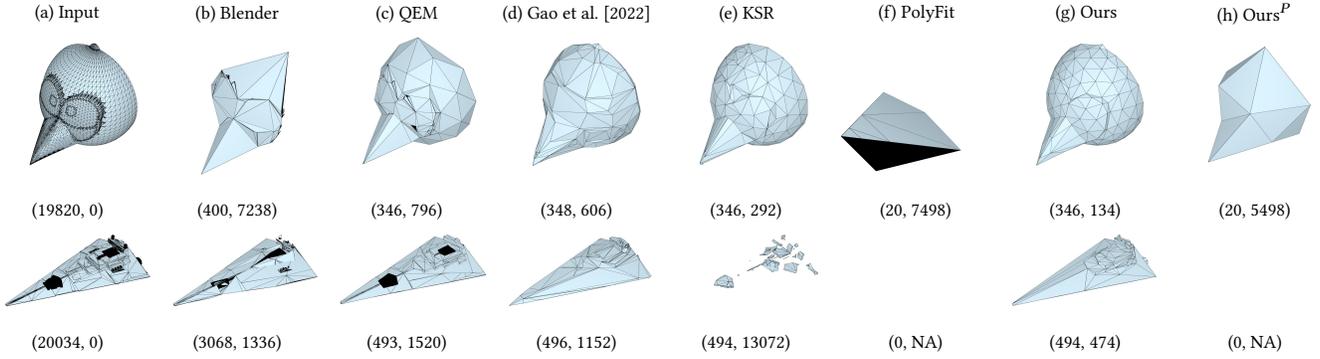


Fig. 12. Comparison with academia and open-source solutions, where (\bullet, \bullet) denotes (face number, light field distance). The inverted faces are rendered as black. Note that, even after re-orientation using MeshLab [Cignoni et al. 2008], inverted faces appears in the results of PolyFit. Besides, PolyFit also fails in the second example.

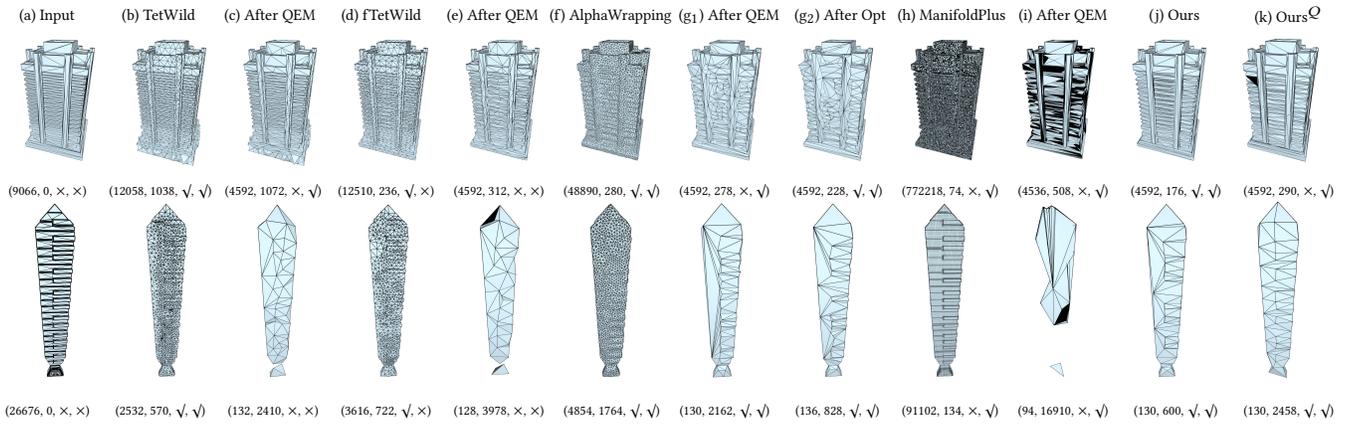


Fig. 13. Comparison with variants of the pipeline of first mesh repairing and then mesh simplification. For (g_2) , we apply the proposed mesh optimization on the AlphaWrapping output; for (k) , we combine our mesh extraction with QEM simplification. All the back faces are rendered black. $(\bullet, \bullet, \bullet, \bullet)$ indicates (face number, light field distance, self-intersection-free flag, manifoldness flag).

($d = l/n_p$ as mentioned in Section 1). In Table 4, we report the average face number, timing and the visual metrics for 3 different choices of n_p . We notice that increasing n_p improves the visual similarity between our output and the input high-poly mesh, but at the same time, it will cost more time and end up with a larger number of faces. Fig. 14 also provides an illustration.

Voxel Size. Given a screen size n_p , different voxel size will lead to different results. In Fig. 15, we compared the extracted iso-surface results using different voxel sizes for a fixed iso-value $d = l/n_p$, where l is the diagonal length of the bounding box and $n_p = 200$. As we argued in Section 3.1, too large voxels may lead to the missing parts of the extracted iso-surface (second left image), while too small voxels will slow down the extraction (rightmost image). To achieve a trade-off between efficiency and performance, we set the diagonal length of voxels to be equal to our offset distance.

Flow Step Fractional Ratio r . As we state before, during the geometric flowing process, we multiply the flow direction by a fraction r to allow more moving space for the entire mesh and to achieve a

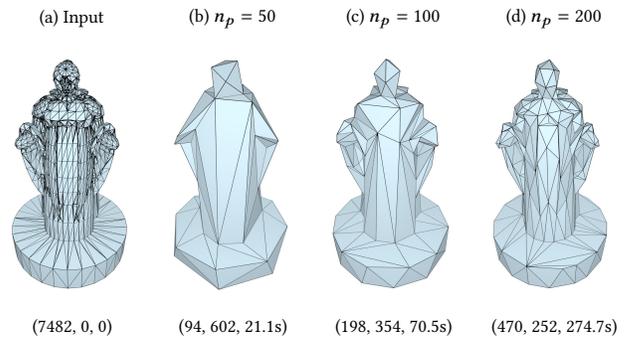


Fig. 14. The re-meshed results w.r.t. different user-specified distance tolerance, where l is the diagonal length of the bounding box of input mesh. $(\bullet, \bullet, \bullet)$ denotes (face number, light field distance, time cost). The smaller the tolerance is, the better re-meshed result we will get, but at the same, the computational cost grows.

better-optimized result. In practice, we find that a smaller step size will lead to a better visual similarity (a smaller *light field distance*)

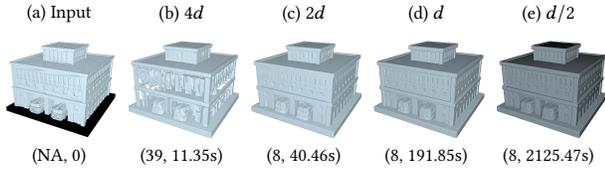


Fig. 15. The different extracted iso-surface for a fixed offset distance using different voxel sizes. (\bullet, \bullet) denotes $(\#genus, \text{time cost})$. The “NA” means that the input mesh is non-manifold. The black bottom in the first figure is due to inverted face orientation. As observed, a larger voxel size (e.g., 4d) produces a high-genus surface. Reducing the voxel size captures finer details but increases the computational cost.

between the output mesh M_o and the input mesh M_i , but at the cost of a larger number of triangles. From the test shown in Fig. 16, we empirically choose $\frac{1}{8}$ as the default value to achieve a good balance between a low element count and a high visual similarity to the input.

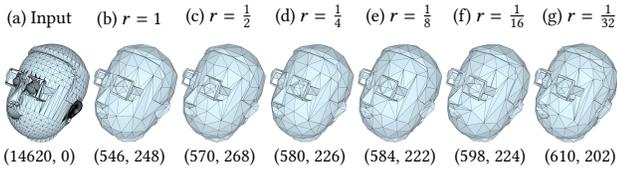


Fig. 16. The different results using different flow step fractional ratio r . (\bullet, \bullet) denotes $(\#faces, \text{LFD})$. As we can see, decreasing r will lead to a smaller LFD, but in turns, it will produce an output with larger number of faces.

Feature Curve Length l_0 . In Fig. 18, we show different iso-surface results based on different choices of feature-line length. Increasing this threshold does gradually solve the “saw-tooth” issue of the initially extracted iso-surface (Fig. 18b). At the same time, it will blur some sharp features. In practice, we find $l_0 = 4$ is a choice of ideal trade-off. A better understanding of choosing these parameters needs further exploration, we leave this as future work.

5 ADDITIONAL APPLICATIONS

5.1 Iso-surface Extraction Comparison

The mesh extraction step of our algorithm can be independently useful, where many competing algorithms have been proposed in the past as shown in Table 1. We show the advantage of our mesh extraction algorithm by comparing our approach with: 1) MC33 [Chernyaev 1995], 2) EMC [Kobbelt et al. 2001], 3) DC [Ju et al. 2002], and 4) Manson and Schaefer [2010]’s approach. The first three serve as the baselines, and the last one meets all the desired properties listed in Table 1. In order to apply these algorithms to any input mesh M_i , we convert the input mesh M_i to be an implicit function by Equation 1, and the corresponding Hermite data (Section 3) for DC. We modified the EMC algorithm provided in Mario Botsch [2015], adapt the Vega et al. [2019]’s implementation of MC33, use the embedded DC function in libigl [Jacobson

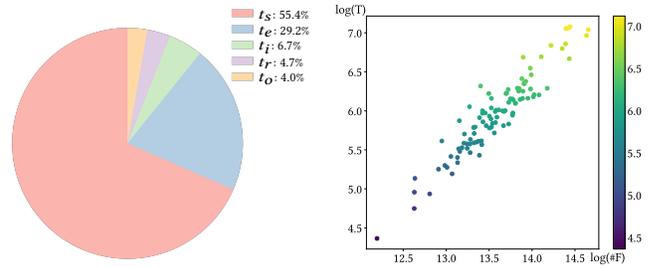


Fig. 17. The time statistics for our method. *left*: the pie chat of time costed the different stages of our algorithm. All the symbols are defined in Section 4.2. *right*: the log-log (base e) plot of our time cost T (in seconds) and the face number $(\#F)$ of extracted iso-surface. We find that the most time consuming parts (over 80%) are the edge-flip step (t_e) in iso-surface generation step and the mesh simplification (t_s) during the mesh optimization, where massive self-intersection checks are applied to ensure the desired intersection-free properties. This explains the strong positive correlation between the time consumption and the face number in extracted iso-surface. Indeed, the more faces you have, the more edge flip and collapse operations will be applied.

et al. 2018], and choose the Manson and Schaefer [2010]’s own implementation⁵ to generate the corresponding results. In Fig. 22, we show the extracted iso-surface in terms of the different iso-values: $l/50$, $l/100$, $l/200$ and $l/400$, where l is the diagonal length of the bounding box of M_i . Among all of these examples, we use the same grid resolution as ours, except for Manson and Schaefer [2010]’s approach, where the default octree settings are used. One thing to point out is that although all approaches generate reasonable results, prior works suffer from several drawbacks: MC33 can generate a closed and self-intersection-free manifold surface, but cannot capture the sharp creases especially when the grid resolution is low (see the zoom-in of Fig. 22); EMC and DC recover the sharp features, but they either may lead to self-intersections or have no guarantees of the manifoldness and self-intersection-free properties (see Section 2.2 for more detailed discussion); Manson and Schaefer [2010]’s approach may generate iso-surface with an undesired high genus (circled regions in Fig. 22).

5.2 Cage Generation

Our re-meshing scheme can be easily adapted to generate cages for the input mesh, without any requirements for it to be manifold, watertight, or self-intersection-free. The cage mesh has to fully enclose but not penetrate a 3D model. We can easily achieve this by adding a penetration check during our mesh optimization step. More specifically, every time we update a vertex position in the flow and alignment steps, we simply modify Algorithm 4 by adding one more intersection check between the current mesh and the input. We reject any edge collapse that leads to the intersection with the input. These additional checks can be handled efficiently by classical BVH-based collision detection [Karras 2012]. In Fig. 19, we show the cage generated by our algorithm. Unlike the automatic

⁵http://josiahmanson.com/research/iso_simplicial/

Table 4. The statistics for low-poly meshes generated with different screen sizes (n_p). Increasing the screen size results in a better re-meshing result, but leads to a larger face number and a slower solving speed.

n_p	#F	Time(s)		PSNR		LFD		SD		ND		HD	
		ave	sd	ave	sd	ave	sd	ave	sd	ave	sd	ave	sd
50	139	30.40	8.82	20.67	2.26	1298.14	956.31	0.030	0.015	0.093	0.070	0.047	0.027
100	408	99.92	39.13	22.45	2.33	656.88	320.81	0.014	0.0086	0.062	0.067	0.033	0.022
200	1193	440.12	234.35	24.40	2.48	367.14	190.88	0.0065	0.0058	0.042	0.066	0.022	0.020

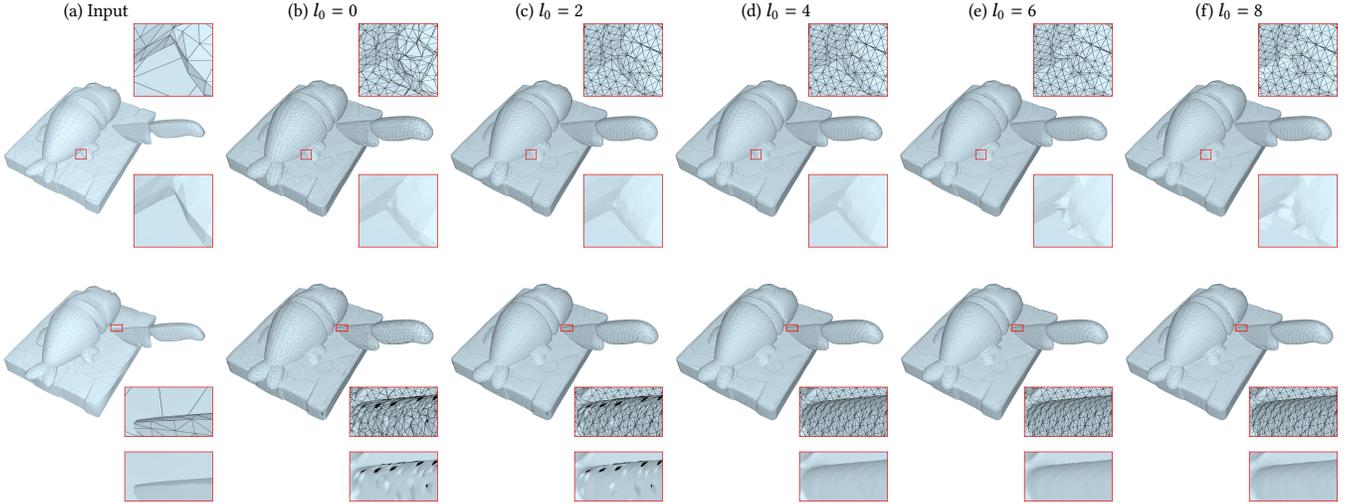


Fig. 18. Different feature-line length threshold l_0 leads to different results, where we also show the zoomed-in regions without the wireframes for a better visualization. The red-framed top and bottom rows are the same models with different rendering. As we can see, larger threshold does smooth out the geometry (the black regions disappears), but at the same time, some of sharp features are blurred. $l_0 = 4$ achieves a trade-off between these considerations, thus we choose this as our default parameter.

caging algorithm Sacht et al. [2015] requiring the input to be watertight, self-intersection-free, and manifold, our algorithm makes no assumptions of the input mesh. For example, the input model in Fig. 19 has 32 non-manifold edges and 264 intersecting triangle pairs.

To compare with [Sacht et al. 2015] more thoroughly, we run both approaches to generate cages for a dataset [Gao et al. 2019] containing 93 meshes with clean topology and geometry that is required by Sacht et al. [2015]’s approach. We use the author-provided code to generate a cage with their E_{varap} energy (see Section 3.2 of Sacht et al. [2015] for details). We run both methods by setting the number of triangles of the final cage to be 2000 and the computing time limit of 1h. As shown in Table 5, Sacht et al. [2015]’s solution returns run time error for 20 models and fails to produce any results within the time limit for 5 models. At the same time, our approach successfully generates a tighter cage (smaller Hausdorff distance) for the entire dataset. Notice that, some cages generated by Sacht et al. [2015] have really bad artifacts, for example, the “spikes” shown in Fig. 20. We also reported the updated statistics after manually removing these models in the last two rows of Table 5.

Table 5. The Hausdorff distance statistics. $HD_{c \rightarrow i}$ is the Hausdorff distance from generated cage to the input mesh, $HD_{i \rightarrow c}$ is the distance from the opposite direction, and $HD = \max(HD_{c \rightarrow i}, HD_{i \rightarrow c})$. r_s is the successful ratio. Sacht et al. [2015] failed to produce the results for 20 out of 93 models due to the run time error, for 5 out of 93 models since exceeding the time threshold.

Methods	r_s	$HD_{c \rightarrow i}$		$HD_{i \rightarrow c}$		HD	
		ave	sd	ave	sd	ave	sd
Ours	100%	0.21	0.58	0.20	0.53	0.22	0.59
Sacht et al.	73.1%	1.11	7.17	38.03	306.89	38.03	306.89
Ours*	100%	0.14	0.36	0.13	0.32	0.15	0.37
Sacht et al.*	74.7%	0.15	0.33	0.19	0.45	0.19	0.45

6 CONCLUSION, LIMITATIONS, AND FUTURE WORKS

In this paper, we propose a robust approach to generate low-poly representations of any input mesh. Our approach can be decomposed into two independently useful stages: 1) the iso-surface extraction stage (re-meshing), where we extract a water-tight, feature-preserving, and self-intersection-free iso-surface of the input mesh with any user-specific iso-value; 2) a mesh optimization stage, where we alternatively re-mesh and flow the extracted the surface to meet the desired properties: low-resolution and visually close to the input mesh. Although we currently cannot guarantee the deviation bound,

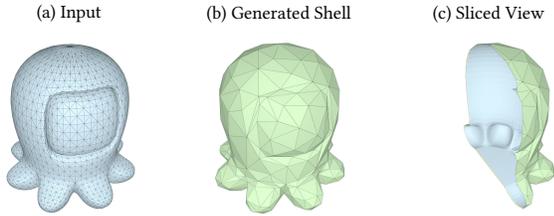


Fig. 19. The generated shell for a cartoon octopus.

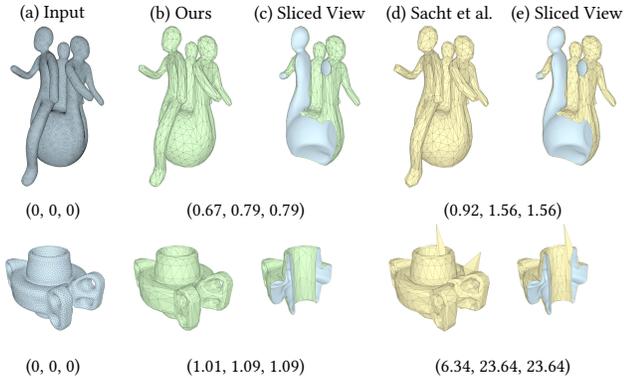


Fig. 20. Comparison with Sacht et al. [2015]. $(\bullet, \bullet, \bullet)$ denotes the Hausdorff distance from cage to input, from input to cage, and between input and cage, respectively. Even if the input mesh is water-tight, Sacht et al. [2015] may end up with bad cage shape (bottom row).

our algorithm effectively adheres to it, with a Hausdorff distance (HD) of $4.4d$ for the dataset, where d represents the offset distance.

Scalability. Fig. 21 illustrates the relationship between screen size and the time and memory requirements for the tree model shown in Fig. 2. While our approach successfully produces results for larger screen sizes, it does not demonstrate optimal scalability in terms of memory and time efficiency. The primary reason for memory consumption is the dense grid generation, which accounts for over 70% of memory usage. We believe that using a sparse grid implementation will alleviate this issue, and we plan to explore this as a future engineering improvement. Regarding efficiency, as shown in Table 4, our approach spends the majority of its time (over 80%) on simplification and edge flip steps during iso-surface extraction. These steps involve numerous intersection checks, and a parallel implementation could significantly accelerate the process. Additionally, our current iso-surface extraction relies on a CPU-based algorithm, and we aim to develop a GPU-based version in future work.

Manifoldness, Self-intersection-freeness, Watertightness. Our approach consistently produces intersection-free, manifold, and watertight outputs. Ensuring intersection-free and manifold properties facilitates easier UV unwrapping and minimizes visible appearance artifacts during texture baking. However, when dealing with an open input mesh, watertightness might not be essential, where a

post mesh segmentation process could be employed to remove the redundant faces to carve out the open region.

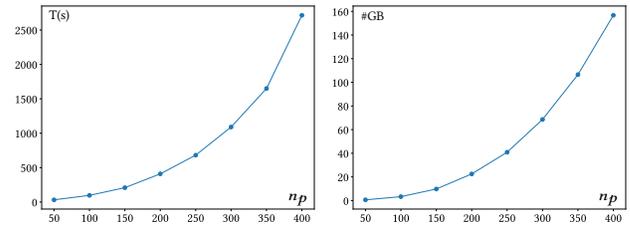


Fig. 21. Time and memory consumption in terms of screen size for the tree model in Fig. 2.

ACKNOWLEDGMENTS

We would like to thank our colleagues from LightSpeed Studios, Fengquan Wang and Dong Li, for their valuable support on this project. This project was partially funded by NSF grant HCC-2212048 and Adobe Systems.

REFERENCES

- Donya Labs AB. 2022. Simplygon 9. <https://www.simplygon.com/Home/Index#section-solutions>
- Marco Attene. 2010. A lightweight approach to repairing digitized polygon meshes. *The Visual Computer* 26 (2010), 1393–1406.
- Marco Attene, Daniela Giorgi, Massimo Ferri, and Bianca Falcidieno. 2009. On converting sets of tetrahedra to combinatorial and PL manifolds. *Computer Aided Geometric Design* 26, 8 (2009), 850–864.
- Gavin Barill, Neil G. Dickson, Ryan Schmidt, David I. W. Levin, and Alec Jacobson. 2018. Fast Winding Numbers for Soups and Clouds. *ACM Trans. Graph.* 37, 4, Article 43 (jul 2018), 12 pages.
- Jean-Philippe Bauchet and Florent Lafarge. 2020. Kinetic Shape Reconstruction. *ACM Trans. Graph.* 39, 5, Article 156 (jun 2020), 14 pages.
- Hervé Brönnimann, Andreas Fabri, Geert-Jan Giezeman, Susan Hert, Michael Hoffmann, Lutz Kettner, Sylvain Pion, and Stefan Schirra. 2022. 2D and 3D Linear Geometry Kernel. In *CGAL User and Reference Manual* (5.5.1 ed.). CGAL Editorial Board.
- Stéphane Calderon and Tamy Boubekeur. 2017. Bounding Proxies for Shape Approximation. *ACM Trans. Graph.* 36, 4, Article 57 (jul 2017), 13 pages.
- Anne-Laure Chauve, Patrick Labatut, and Jean-Philippe Pons. 2010. Robust piecewise-planar 3D reconstruction and completion from large-scale unstructured point data. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*. 1261–1268.
- Ding-Yun Chen, Xiao-Pei Tian, Yu-Te Shen, and Ming Ouhyoung. 2003. On Visual Similarity Based 3D Model Retrieval. *Computer Graphics Forum* 22, 3 (2003), 223–232.
- Zhiqin Chen, Andrea Tagliasacchi, Thomas Funkhouser, and Hao Zhang. 2022a. Neural Dual Contouring. *ACM Trans. Graph.* 41, 4, Article 104 (jul 2022), 13 pages.
- Zhiqin Chen, Andrea Tagliasacchi, and Hao Zhang. 2020. BSP-Net: Generating Compact Meshes via Binary Space Partitioning. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, Los Alamitos, CA, USA, 42–51.
- Zhiqin Chen, Kangxue Yin, and Sanja Fidler. 2022b. AUV-Net: Learning Aligned UV Maps for Texture Transfer and Synthesis. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, Los Alamitos, CA, USA, 1455–1464.
- Zhiqin Chen and Hao Zhang. 2021. Neural Marching Cubes. *ACM Trans. Graph.* 40, 6, Article 251 (dec 2021), 15 pages.
- Evgeni V Chernyaev. 1995. Marching cubes 33: construction of topologically correct isosurfaces. *Tech. Rep. CERN CN/95-17* (1995).
- Paolo Cignoni, Marco Callieri, Massimiliano Corsini, Matteo Dellepiane, Fabio Ganovelli, and Guido Ranzuglia. 2008. MeshLab: an Open-Source Mesh Processing Tool. In *Eurographics Italian Chapter Conference*, Vittorio Scarano, Rosario De Chiara, and Ugo Erra (Eds.). The Eurographics Association.
- P. Cignoni, C. Rocchini, and R. Scopigno. 1998. Metro: Measuring Error on Simplified Surfaces. *Computer Graphics Forum* 17, 2 (1998), 167–174.
- David Cohen-Steiner, Pierre Alliez, and Mathieu Desbrun. 2004. Variational Shape Approximation. *ACM Trans. Graph.* 23, 3 (aug 2004), 905–914.

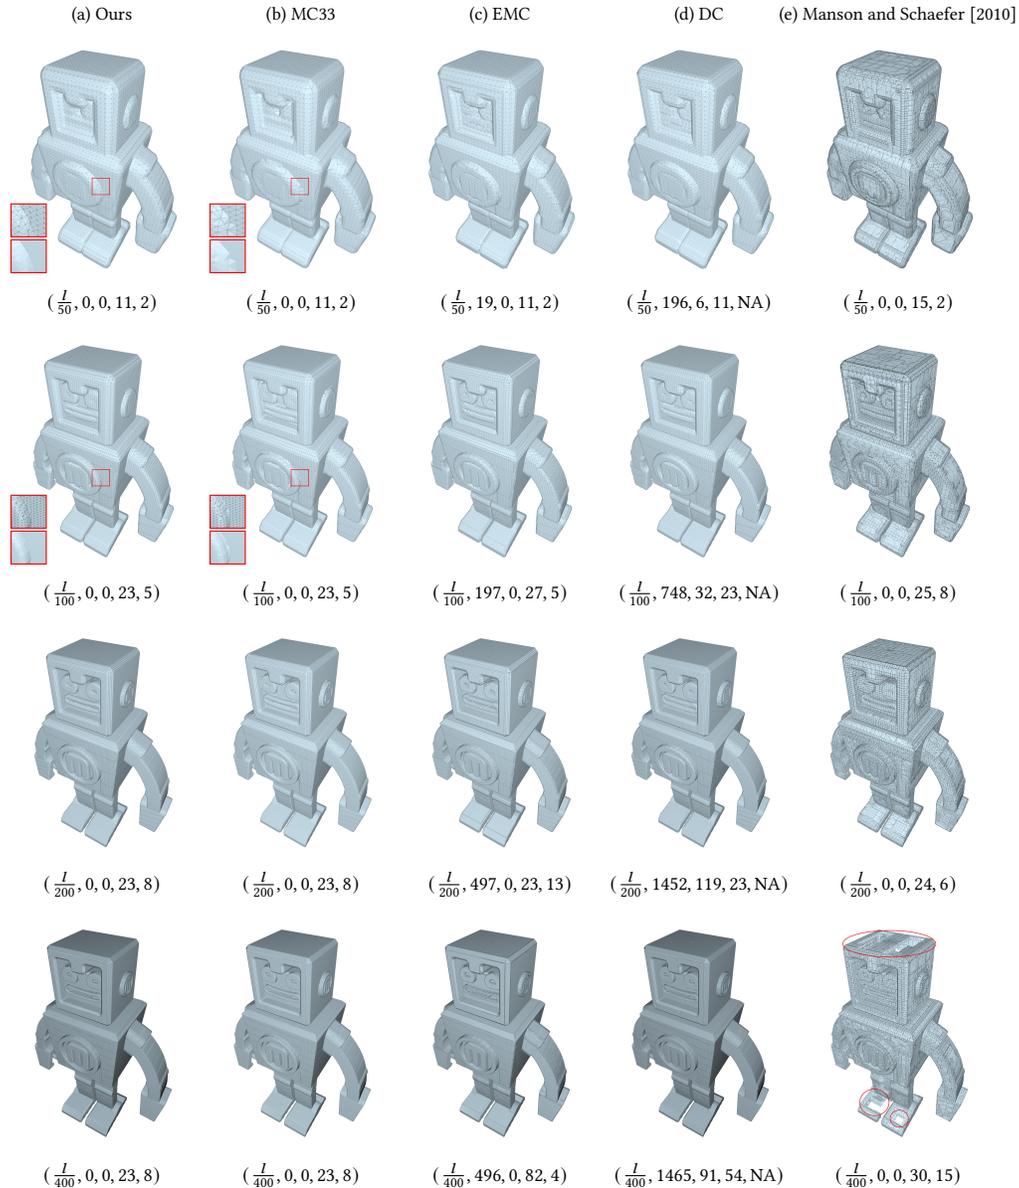


Fig. 22. The comparison of different iso-surface extraction method. $(\bullet, \bullet, \bullet, \bullet, \bullet)$ are (iso-value, #self-intersected faces, #non-manifold edges, #comps, #genus), where l is the diagonal length of the bounding box of the input mesh. “NA” means the genus is not well defined given the mesh is not manifold

Lis Custodio, Tiago Etienne, Sinesio Pesco, and Claudio Silva. 2013. Practical considerations on Marching Cubes 33 topological correctness. *Computers and Graphics* 37, 7 (2013), 840–850.

B. R. de Araújo, Daniel S. Lopes, Pauline Jepp, Joaquim A. Jorge, and Brian Wyvill. 2015. A Survey on Implicit Surface Polygonization. *ACM Comput. Surv.* 47, 4, Article 60 (may 2015), 39 pages.

Lorenzo Diazzi and Marco Attene. 2021. Convex Polyhedral Meshing for Robust Solid Modeling. *ACM Trans. Graph.* 40, 6, Article 259 (dec 2021), 16 pages.

Akio Doi and Akio Koide. 1991. An efficient method of triangulating equi-valued surfaces by using tetrahedral cells. *IEICE Transactions on Information and Systems* (1991).

Matin J. Dürst. 1988. Letters: additional reference to marching cubes. *Computer Graphics* (1988).

Hao Fang and Florent Lafarge. 2020. Connect-and-Slice: An Hybrid Approach for Reconstructing 3D Objects. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 13487–13495.

Hao Fang, Florent Lafarge, and Mathieu Desbrun. 2018. Planar Shape Detection at Structural Scales. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2965–2973.

Xifeng Gao, Hanxiao Shen, and Daniele Panozzo. 2019. Feature Preserving Octree-Based Hexahedral Meshing. *Computer Graphics Forum* 38 (08 2019), 135–149.

Xifeng Gao, Kui Wu, and Zherong Pan. 2022. Low-Poly Mesh Generation for Building Models. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Proceedings (Vancouver, BC, Canada) (SIGGRAPH22 Conference Proceeding)*. Association for Computing Machinery, New York, NY, USA, Article 3, 9 pages.

Michael Garland and Paul S. Heckbert. 1997. Surface Simplification Using Quadric Error Metrics. In *Proceedings of the 24th Annual Conference on Computer Graphics*

- and Interactive Techniques (SIGGRAPH '97). ACM Press/Addison-Wesley Publishing Co., USA, 209–216.
- Dong-Hoon Han, Chang-Jin Lee, Sangbin Lee, and Hyeong-Seok Ko. 2021. Tight Normal Cone Merging for Efficient Collision Detection of Thin Deformable Objects. In *Eurographics 2021 - Short Papers*, Holger Theisel and Michael Wimmer (Eds.). The Eurographics Association.
- Jon Hasselgren, Jacob Munkberg, Jaakko Lehtinen, Miika Aittala, and Samuli Laine. 2021. Appearance-Driven Automatic 3D Model Simplification. In *Eurographics Symposium on Rendering - DL-only Track*, Adrien Bousseau and Morgan McGuire (Eds.). The Eurographics Association, Prague, Czech Republic, 19 pages.
- Hugues Hoppe. 1996. Progressive Meshes. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '96)*. Association for Computing Machinery, New York, NY, USA, 99–108.
- Hugues Hoppe. 1999. New Quadric Metric for Simplifying Meshes with Appearance Attributes. In *Proceedings of the Conference on Visualization '99: Celebrating Ten Years* (San Francisco, California, USA) (VIS '99). IEEE Computer Society Press, Washington, DC, USA, 59–66.
- Yixin Hu, Teseo Schneider, Bolun Wang, Denis Zorin, and Daniele Panozzo. 2020. Fast Tetrahedral Meshing in the Wild. *ACM Trans. Graph.* 39, 4, Article 117 (July 2020), 18 pages.
- Yixin Hu, Qingnan Zhou, Xifeng Gao, Alec Jacobson, Denis Zorin, and Daniele Panozzo. 2018. Tetrahedral Meshing in the Wild. *ACM Trans. Graph.* 37, 4, Article 60 (July 2018), 14 pages.
- Jin Huang, Tengfei Jiang, Zeyun Shi, Yiyang Tong, Hujun Bao, and Mathieu Desbrun. 2014. ℓ_1 -Based Construction of Polycube Maps from Complex Shapes. *ACM Trans. Graph.* 33, 3, Article 25 (jun 2014), 11 pages.
- Jingwei Huang, Yichao Zhou, and Leonidas Guibas. 2020. ManifoldPlus: A Robust and Scalable Watertight Manifold Surface Generation Method for Triangle Soups. (2020). arXiv:2005.11621
- Alec Jacobson, Daniele Panozzo, et al. 2018. libigl: A simple C++ geometry processing library. <https://libigl.github.io/>.
- Tao Ju, Frank Losasso, Scott Schaefer, and Joe Warren. 2002. Dual Contouring of Hermite Data. *ACM Trans. Graph.* 21, 3 (jul 2002), 339–346.
- Tao Ju and Tushar Udesi. 2006. Intersection-free contouring on an octree grid. *Pacific Graphics Poster* (01 2006).
- Tero Karras. 2012. Maximizing Parallelism in the Construction of BVHs, Octrees, and k-d Trees. In *Proceedings of the Fourth ACM SIGGRAPH/Eurographics Conference on High-Performance Graphics* (Paris, France) (EGGH-HPG'12). Eurographics Association, Goslar, DEU, 33–37.
- Tom Kelly, John Femiiani, Peter Wonka, and Niloy J. Mitra. 2017. BigSUR: Large-Scale Structured Urban Reconstruction. *ACM Trans. Graph.* 36, 6, Article 204 (nov 2017), 16 pages.
- Dawar Khan, Alexander Plopski, Yuichiro Fujimoto, Masayuki Kanbara, Gul Jabeen, Yongjie Jessica Zhang, Xiaopeng Zhang, and Hirokazu Kato. 2022. Surface Remeshing: A Systematic Literature Review of Methods and Research Directions. *IEEE Transactions on Visualization and Computer Graphics* 28, 3 (2022), 1680–1713.
- Leif P. Kobbelt, Mario Botsch, Ulrich Schwanecke, and Hans-Peter Seidel. 2001. Feature Sensitive Surface Extraction from Volume Data. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '01)*. Association for Computing Machinery, New York, NY, USA, 57–66.
- Samuli Laine, Janne Hellsten, Tero Karras, Yeongho Seol, Jaakko Lehtinen, and Timo Aila. 2020. Modular Primitives for High-Performance Differentiable Rendering. *ACM Trans. Graph.* 39, 6, Article 194 (nov 2020), 14 pages.
- Thibault Lescoat, Hsueh-Ti Derek Liu, Jean-Marc Thiery, Alec Jacobson, Tamy Boubekeur, and Maks Ovsjanikov. 2020. Spectral Mesh Simplification. *Computer Graphics Forum* 39, 2 (2020), 315–324.
- Thomas Lewiner, Hélio Lopes, Antônio Wilson Vieira, and Geovan Tavares. 2003. Efficient Implementation of Marching Cubes' Cases with Topological Guarantees. *Journal of Graphics Tools* 8, 2 (2003), 1–15.
- Minglei Li and Liangliang Nan. 2021. Feature-preserving 3D mesh simplification for urban buildings. *ISPRS Journal of Photogrammetry and Remote Sensing* 173 (2021), 135–150.
- Yiyi Liao, Simon Donné, and Andreas Geiger. 2018. Deep Marching Cubes: Learning Explicit Surface Representations. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2916–2925.
- Peter Lindstrom and Greg Turk. 1998. Fast and Memory Efficient Polygonal Simplification. In *Proceedings of the Conference on Visualization '98* (Research Triangle Park, North Carolina, USA) (VIS '98). IEEE Computer Society Press, Washington, DC, USA, 279–286.
- Peter Lindstrom and Greg Turk. 2000. Image-Driven Simplification. *ACM Trans. Graph.* 19, 3 (jul 2000), 204–241.
- Songrun Liu, Zachary Ferguson, Alec Jacobson, and Yotam Gingold. 2017. Seamless: Seam Erasure and Seam-Aware Decoupling of Shape from Mesh Resolution. *ACM Trans. Graph.* 36, 6, Article 216 (nov 2017), 15 pages.
- Adriano Lopes and Ken Brodlie. 2003. Improving the robustness and accuracy of the marching cubes algorithm for isosurfacing. *IEEE Transactions on Visualization and Computer Graphics* 9, 1 (2003), 16–29.
- William E. Lorensen and Harvey E. Cline. 1987. Marching Cubes: A High Resolution 3D Surface Construction Algorithm. In *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '87)*. Association for Computing Machinery, New York, NY, USA, 163–169.
- Fujun Luan, Shuang Zhao, Kavita Bala, and Zhao Dong. 2021. Unified Shape and SVBRDF Recovery using Differentiable Monte Carlo Rendering. *Computer Graphics Forum* 40, 4 (2021), 101–113.
- Josiah Manson and Scott Schaefer. 2010. Isosurfaces Over Simplicial Partitions of Multiresolution Grids. *Computer Graphics Forum* 29, 2 (2010), 377–385.
- Jan Möbius Mario Botsch. 2015. IsoEx. <https://www.graphics.rwth-aachen.de/IsoEx/index.html>
- Sergey V. Matveyev. 1994. Approximation of Isosurface in the Marching Cube: Ambiguity Problem. In *Proceedings of the Conference on Visualization '94* (Washington, D.C.) (VIS '94). IEEE Computer Society Press, Washington, DC, USA, 288–292.
- Ravish Mehra, Qingnan Zhou, Jeremy Long, Alla Sheffer, Amy Gooch, and Niloy J. Mitra. 2009. Abstraction of Man-Made Shapes. *ACM Trans. Graph.* 28, 5 (dec 2009), 1–10.
- Jacob Munkberg, Wenzheng Chen, Jon Hasselgren, Alex Evans, Tianchang Shen, Thomas Muller, Jun Gao, and Sanja Fidler. 2022. Extracting Triangular 3D Models, Materials, and Lighting From Images. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, Los Alamitos, CA, USA, 8270–8280.
- Liangliang Nan and Peter Wonka. 2017. PolyFit: Polygonal Surface Reconstruction from Point Clouds. In *2017 IEEE International Conference on Computer Vision (ICCV)*. 2372–2380.
- Baptiste Nicolet, Alec Jacobson, and Wenzel Jakob. 2021. Large Steps in Inverse Rendering of Geometry. *ACM Trans. Graph.* 40, 6, Article 248 (dec 2021), 13 pages.
- Gregory M. Nielson. 2003. On marching cubes. *IEEE Transactions on Visualization and Computer Graphics* 9, 3 (2003), 283–297.
- Gregory M. Nielson. 2004. Dual Marching Cubes. In *Proceedings of the Conference on Visualization '04* (VIS '04). IEEE Computer Society, USA, 489–496.
- Gregory M. Nielson and Bernd Hamann. 1991. The asymptotic decider: resolving the ambiguity in marching cubes. In *Proceeding Visualization '91*. 83–91.
- Cédric Portaneri, Mael Rouxel-Labbé, Michael Hemmer, David Cohen-Steiner, and Pierre Alliez. 2022. Alpha Wrapping with an Offset. *ACM Transactions on Graphics* 41, 4 (June 2022), 1–22. <https://hal.inria.fr/hal-03688637>
- Leonardo Sacht, Etienne Vouga, and Alec Jacobson. 2015. Nested Cages. *ACM Trans. Graph.* 34, 6, Article 170 (nov 2015), 14 pages.
- D. Salinas, F. Lafarge, and P. Alliez. 2015. Structure-Aware Mesh Decimation. *Computer Graphics Forum* 34, 6 (2015), 211–227.
- Scott Schaefer, Tao Ju, and Joe Warren. 2007. Manifold Dual Contouring. *IEEE Transactions on Visualization and Computer Graphics* 13, 03 (may 2007), 610–619.
- Scott Schaefer and Joe Warren. 2004. Dual marching cubes: Primal contouring of dual grids. *Proceedings - Pacific Conference on Computer Graphics and Applications*, 70–76.
- Tianchang Shen, Jun Gao, Kangxue Yin, Ming-Yu Liu, and Sanja Fidler. 2021. Deep Marching Tetrahedra: a Hybrid Representation for High-Resolution 3D Shape Synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- David Vega, Javier Abache, and David Coll. 2019. A Fast and Memory Saving Marching Cubes 33 Implementation with the Correct Interior Test. *Journal of Computer Graphics Techniques (JCGT)* 8, 3 (8 August 2019), 1–18.
- Pascal Volino and Nadia Magnenat Thalmann. 1994. Efficient self-collision detection on smoothly discretized surface animations using geometrical shape regularity. *Computer Graphics Forum* (1994).
- Tongtong Wang, Zhihua Liu, Min Tang, Ruofeng Tong, and Dinesh Manocha. 2017. Efficient and Reliable Self-Collision Culling Using Unprojected Normal Cones. *Computer Graphics Forum* 36, 8 (2017), 487–498.
- Xinyue Wei, Minghua Liu, Zhan Ling, and Hao Su. 2022. Approximate Convex Decomposition for 3D Meshes with Collision-Aware Concavity and Tree Search. *ACM Trans. Graph.* 41, 4, Article 42 (jul 2022), 18 pages.
- Kui Wu, Xu He, Zherong Pan, and Xifeng Gao. 2022. Occluder Generation for Buildings in Digital Games. *Computer Graphics Forum* 41, 7 (2022), 205–214.
- Geoff Wyvill, Craig McPheeters, and Brian Wyvill. 1986. Data Structure for Soft Objects. *The Visual Computer - VC* 2 (08 1986), 227–234.
- Kaizhi Yang and Xuejin Chen. 2021. Unsupervised Learning for Cuboid Shape Abstraction via Joint Segmentation from Point Clouds. *ACM Trans. Graph.* 40, 4, Article 152 (jul 2021), 11 pages.
- Mulin Yu and Florent Lafarge. 2022. Finding Good Configurations of Planar Primitives in Unorganized Point Clouds. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, Los Alamitos, CA, USA, 6357–6366.
- Qingnan Zhou and Alec Jacobson. 2016. Thingi10K: A Dataset of 10,000 3D-Printing Models. (2016). arXiv:1605.04797

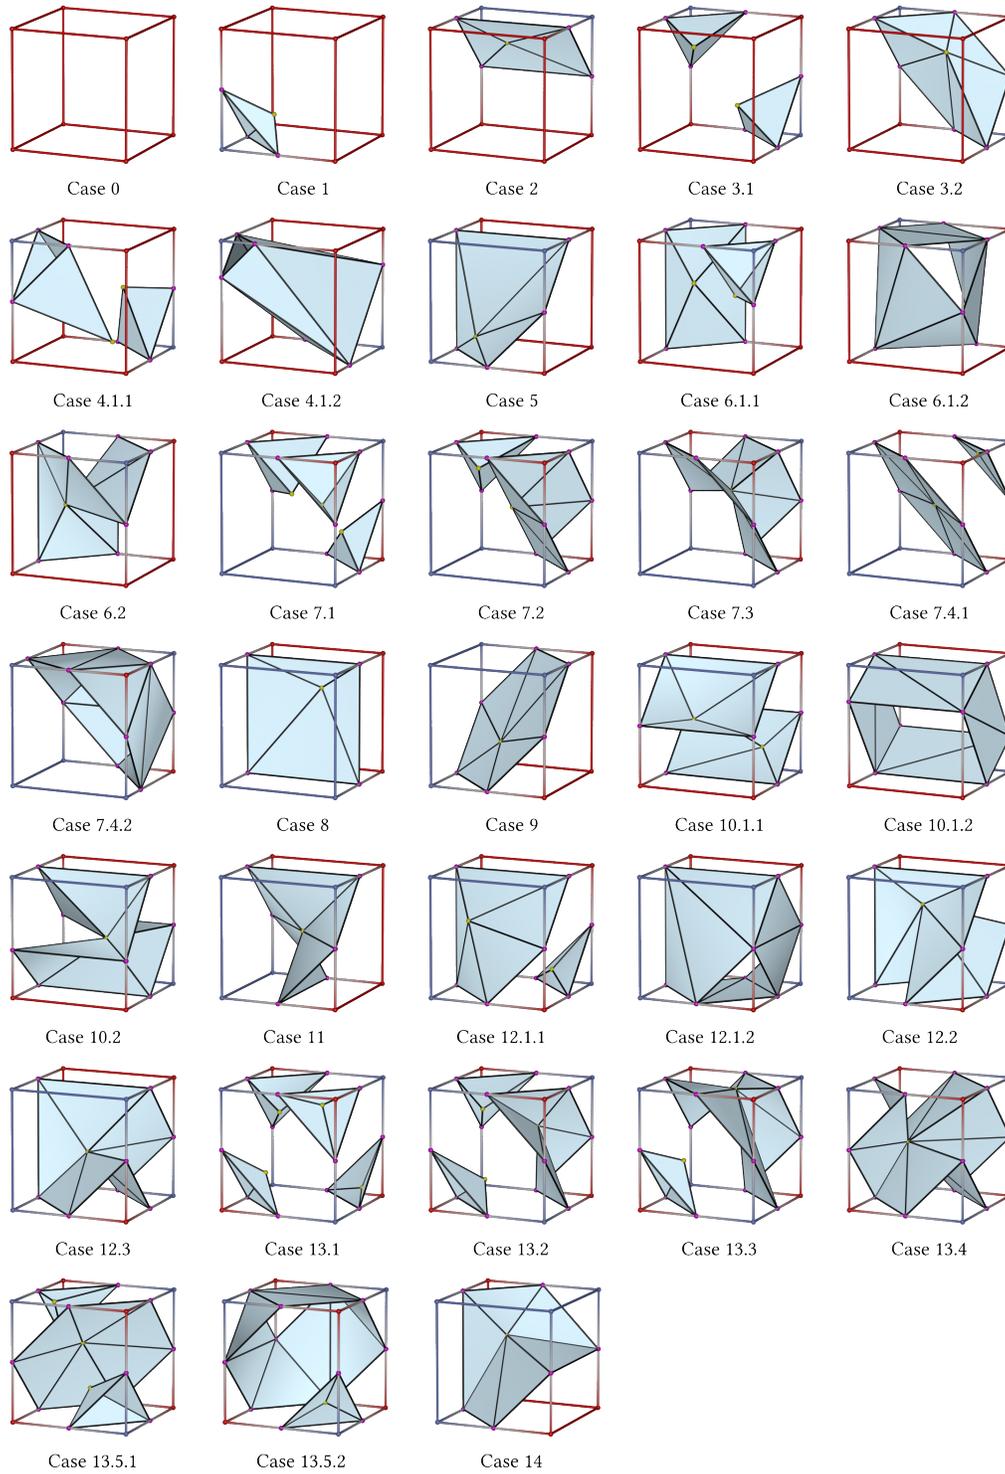


Fig. 23. We show our look up table. The cube vertices are colored by their signs, with red for positive and blue for negative. The surface-cube intersections on edges are the pink points, while the yellow points are the inserted feature points. Notice that case 12.2 and 12.3 are symmetric, as well as case 11 and 14.

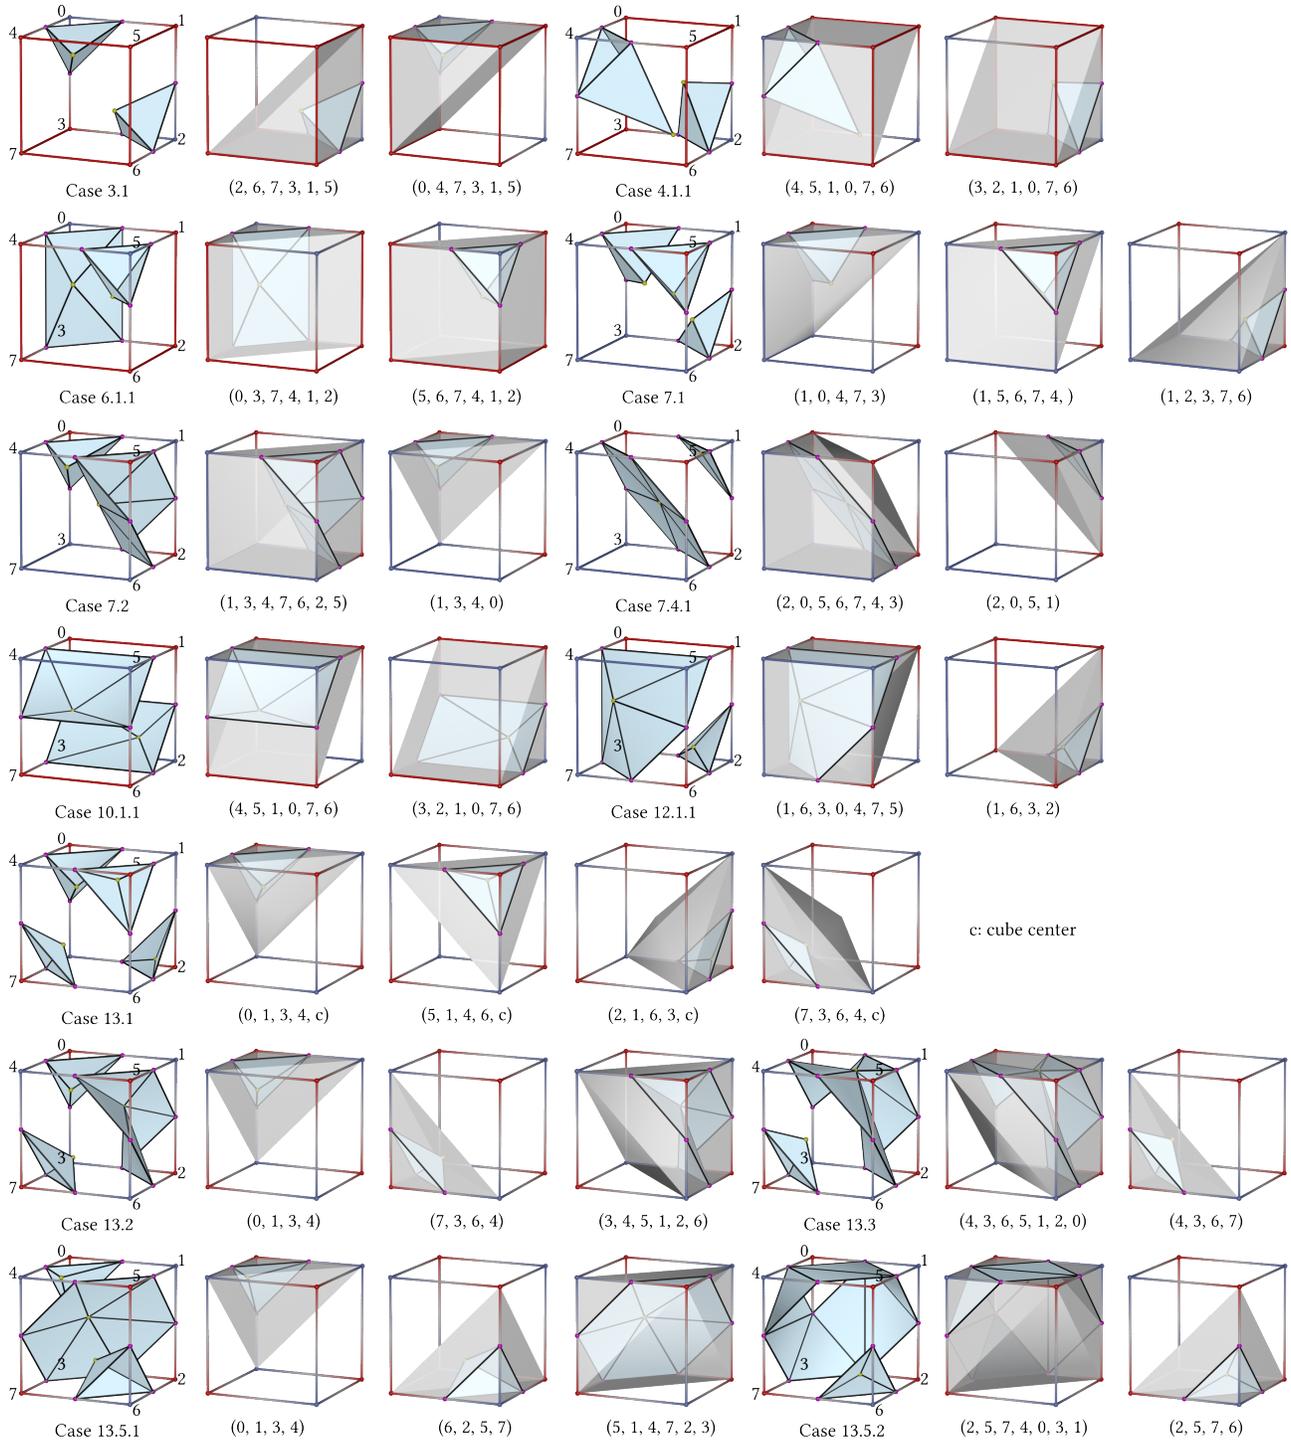


Fig. 24. Cube division policies for MC33 cases with more than one components. The numbers inside the parentheses are the cube vertices which form the constraint polyhedra (rendered in gray) for the corresponding components (rendered in blue)