



Tetrahedral meshing via maximal Poisson-disk sampling



Jianwei Guo^a, Dong-Ming Yan^{a,*}, Li Chen^b, Xiaopeng Zhang^a,
Oliver Deussen^{c,d}, Peter Wonka^e

^a National Laboratory of Pattern Recognition, Institute of Automation, Chinese Academy of Sciences, Beijing 100190, China

^b School of Software, Tsinghua University, Beijing 100084, China

^c University of Konstanz, Konstanz 78457, Germany

^d Shenzhen Key Lab of Visual Computing and Visual Analytics/SIAT, Shenzhen 518055, China

^e Visual Computing Center, King Abdullah University of Science and Technology, Thuwal 23955-6900, Saudi Arabia

ARTICLE INFO

Article history:

Available online 15 February 2016

Keywords:

Tetrahedral mesh generation
Maximal Poisson-disk sampling
Sliver removal
Mesh optimization

ABSTRACT

In this paper, we propose a simple yet effective method to generate 3D-conforming tetrahedral meshes from closed 2-manifold surfaces. Our approach is inspired by recent work on maximal Poisson-disk sampling (MPS), which can generate well-distributed point sets in arbitrary domains. We first perform MPS on the boundary of the input domain, we then sample the interior of the domain, and we finally extract the tetrahedral mesh from the samples by using 3D Delaunay or regular triangulation for uniform or adaptive sampling, respectively. We also propose an efficient optimization strategy to protect the domain boundaries and to remove slivers to improve the meshing quality. We present various experimental results to illustrate the efficiency and the robustness of our proposed approach. We demonstrate that the performance and quality (e.g., minimal dihedral angle) of our approach are superior to current state-of-the-art optimization-based approaches.

© 2016 Elsevier B.V. All rights reserved.

1. Introduction

Mesh generation aims to approximate a given closed domain with simple discrete elements (e.g., triangle/quad in 2D and tetrahedron/pyramid/prism/hexahedron in 3D). It has numerous applications ranging from engineering to scientific research, e.g., simulation of mechanical parts or architecture, medical and biological data analysis, geographical science, computational fluid dynamics, and animation in computer graphics (Klingner et al., 2006; Ando et al., 2013). Here we focus on 3D tetrahedral mesh generation.

Although many robust commercial software (e.g., Ansys) and open source packages exist for mesh generation (e.g., TetGen (Si, 2015), CGALmesh (Jamin et al., 2014), DelPSC (Cheng et al., 2007), Gmsh (Geuzaine and Remacle, 2009), Qhull (Barber et al., 1996), etc.), research on tetrahedral meshing remains active because no available approach is able to fulfill all conflicting quality requirements of various applications.

There are many criteria to evaluate the quality of a tetrahedral mesh, e.g., the approximation quality, the gradation, the dihedral angle, and the radius ratio of a single tetrahedron. These criteria are typically in conflict with each other and are difficult to satisfy simultaneously. Among all the quality criteria, the dihedral angle is the most important for simulation

* Corresponding author.

E-mail addresses: jianwei.guo@nlpr.ia.ac.cn (J. Guo), yandongming@gmail.com (D.-M. Yan), chenlee@tsinghua.edu.cn (L. Chen), xpzhang@ia.ac.cn (X. Zhang), Oliver.Deussen@uni-konstanz.de (O. Deussen), pwonka@gmail.com (P. Wonka).

<http://dx.doi.org/10.1016/j.cagd.2016.02.004>

0167-8396/© 2016 Elsevier B.V. All rights reserved.

because it is directly related to the condition number of the stiffness matrix. A single poorly shaped tetrahedron with near zero volume (called a *sliver*) can destroy the whole simulation (Shewchuk, 2002). Nevertheless, recent approaches still face the difficulty of completely removing slivers, even with careful treatment of the energy functions and with post-processing, such as simulated annealing or sliver exudation (Alliez et al., 2005; Tournois et al., 2009; Chen et al., 2014).

The main contribution of this paper is to propose a very simple algorithm that improves 3D tetrahedral mesh generation compared with existing algorithms that optimize complex objective functions. Our work is inspired by recent work in maximal Poisson-disk sampling (MPS) (Ebeida and Mitchell, 2012; Yan and Wonka, 2013), which generates uniformly and randomly distributed point sets while exhibiting blue-noise properties. These extracted meshes from sampled point sets have very good geometric properties, e.g., angle bounds, edge-length bounds, etc. Our approach is a combination of randomized sampling and several effective optimization strategies that improve the lower/upper bounds of dihedral angles. Tetrahedra near the boundary surface are treated carefully such that the quality of the resulting tetrahedral mesh is provably protected. Our method works better than other greedy approaches for generating samples (e.g., Delaunay insertion). It is more efficient than optimization-based algorithms while at the same time it can improve the quality of the smallest dihedral angles.

2. Related work

Delaunay mesh generation. There are three main categories of tetrahedral meshing algorithms: (i) advancing-front-based approaches (Schöberl, 1997; Choi et al., 2003; Ito et al., 2004), (ii) octree-based methods (Labelle and Shewchuk, 2007; Bronson et al., 2014; Yu et al., 2014), and (iii) Delaunay-based meshing algorithms (Cheng et al., 2012). A complete review of these algorithms is beyond the scope of this paper. We refer the reader to a survey paper (Owen, 1998) and textbooks (Bern and Plassmann, 2000; Frey and George, 2001; Edelsbrunner, 2001) (and references therein) for a more comprehensive introduction to tetrahedral meshing algorithms. In the following, we focus on Delaunay-based meshing algorithms, which are most closely related to our work.

Delaunay-based tetrahedral meshing algorithms have been shown to be the most successful in recent years (Cheng et al., 2012). From an algorithmic perspective, Delaunay meshing techniques can be further classified into two major types: Delaunay insertion/refinement-based methods (Chew, 1997; Jamin et al., 2014; Si, 2015) and optimization-based (or variational) approaches (Alliez et al., 2005; Tournois et al., 2009; Dardenne et al., 2009; Vanderzee et al., 2010; Yan et al., 2010, 2013; Chen et al., 2014).

Delaunay insertion/refinement-based approaches usually start with an initial mesh and gradually insert new vertices at the center of a circumscribed sphere of the existing tetrahedra with the worst quality. The algorithm stops once all the user-specified criteria are satisfied, e.g., smallest dihedral angle, local edge length, approximation quality of the domain boundary, and so on. The framework of TetGen (Si, 2015) uses a series of such approaches, including incremental Delaunay algorithms for inserting vertices, constrained Delaunay algorithms for inserting constraints and a Delaunay refinement algorithm for quality mesh generation. In summary, this type of meshing technique can usually provide theoretical guarantees of the angle bound, element size and approximation error. On the other hand, it is difficult to explicitly control the desired number of vertices. The Delaunay insertion method usually inserts more vertices than necessary to reach the desired quality criteria.

The optimization-based approaches start with an initial set of randomly sampled vertices, where the number of vertices is usually specified by the user. Then, different types of energy functions are proposed to characterize the desired meshing quality. A well-known optimization-based approach is *Centroidal Voronoi Tessellation* (CVT) (Du et al., 1999) mesh generation (Du and Wang, 2003; Yan et al., 2013). CVT-based algorithms optimize the compactness of the dual Voronoi cells instead of the shape of the tetrahedra, which does not suppress any poorly shaped elements. Alliez et al. (2005) generalize *Optimal Delaunay Triangulation* (ODT) (Chen, 2004) for 2D mesh smoothing to 3D tetrahedral meshing. They propose to minimize the global mesh-dependent energy to update both vertex positions and connectivity. The work of Tournois et al. (2009), called natural ODT (NODT), further extends the ODT energy to the domain boundaries. This extension ensures consistency of the energy function on the boundary and inside the domain. The number of badly shaped elements is reduced near the domain boundary. Recent work of Chen (2004) proposes an extended formulation for ODT that is able to generate graded meshes efficiently using quasi-Newton solvers. This type of algorithm generates higher-quality meshes in practice, but they are more time consuming and post-processing for mesh improvement is necessary to eliminate badly shaped elements (Cheng et al., 2000; Klingner and Shewchuk, 2007). In addition, Jiao et al. (2011) present a variational smoothing approach for optimizing surface and volume meshes simultaneously. They define energy functions based on conformal and isometric mappings between actual elements with ideal reference elements and develop a simple algorithm to minimize the energies.

From a boundary approximation perspective, Delaunay meshing can be classified into constrained meshing or conformal meshing. Our approach falls into the latter category, i.e., instead of interpolating the boundary, we approximate the boundary by resampling, which allows for higher tetrahedron quality at the boundary, especially when the quality of the input surface is very poor. Here, we present a simple method based on randomized sampling that is able to generate high-quality meshes that are comparable with state-of-the-art techniques, while being more efficient than variational approaches.

Maximal Poisson-disk sampling. Poisson-disk sampling techniques have been extensively studied in past decades (Lagae and Dutré, 2008), but most previous work cannot achieve maximality of the sampling. In 2006, Jones first proposed an unbiased algorithm using the Voronoi diagram for gap detection/filling. Later, Ebeida et al. (2011a, 2012) improved the efficiency of MPS by using a uniform grid for acceleration. Meanwhile, they showed that MPS point sets are able to generate high-quality triangulations (Ebeida et al., 2011b) and Voronoi meshes (Ebeida and Mitchell, 2012) for the purpose of simulation. More recently, Yan and Wonka (2013) and Guo et al. (2015) extended MPS to mesh surfaces and adaptive sampling. They demonstrated that MPS can generate high-quality surface remeshing as well. In this paper, we further study the use of MPS for tetrahedral mesh generation. To the best of our knowledge, this is the first application of MPS to volumetric meshing.

3. Preliminaries and overview

In this section, we first introduce the concept of maximal Poisson-disk sampling and then briefly describe the MPS framework for tetrahedral mesh generation.

3.1. Maximal Poisson-disk sampling

Suppose $\mathbf{X} = \{(\mathbf{x}_i, r_i)\}$ is a Poisson-disk distribution in a compact domain Ω , where \mathbf{x}_i is the position of the i th point and r_i is the associated radius of \mathbf{x}_i . A *maximal Poisson-disk sampling* distribution should satisfy the following three properties: 1) the *minimal distance property* requires that any two disks are at least a minimum distance apart, i.e., $\forall \mathbf{x}_i, \mathbf{x}_j \in P, \|\mathbf{x}_i - \mathbf{x}_j\| \geq \max(r_i, r_j)$; 2) the *bias-free property* requires that each point in the domain has a probability that is proportional to the sizing at this point to receive a sampling point; and 3) the *maximal sampling property* requires that the union of the disks covers the entire sampling domain, i.e., $\bigcup (\mathbf{x}_i, r_i) \supseteq \Omega$. Furthermore, if the radius, r_i , for all disks is a constant, then the sampling is classic uniform sampling (Lagae and Dutré, 2008). Otherwise, it is adaptive sampling, where each r_i is defined with respect to an underlying sizing function, $\mu(\mathbf{x})$, defined over the sampling domain.

3.2. MPS framework

Two recent and efficient MPS frameworks (Ebeida et al., 2012; Guo et al., 2015) are useful for tetrahedral meshing. The 2D example in Fig. 2 illustrates how to apply an MPS framework for mesh generation. There are four main steps in our sampling/meshing pipeline as follows:

1. Detect sharp corners and generate MPS on the boundary of the input domain. To avoid introducing vertices arbitrarily close to the domain boundary, we apply edge-length optimization to ensure that the boundary disks are deeply intersected (Cheng et al., 2007).
2. Keep the boundary samples fixed and then sample the interior of the domain to achieve maximality.
3. Triangulate the sampled points using Delaunay or regular triangulation for uniform or adaptive sampling, respectively, and extract a mesh from the triangulation.
4. To improve the meshing quality, we perform optimization operations to remove bad elements and perform resampling/retriangulation locally. Multiple criteria are used to perform the mesh optimization.

4. Tetrahedral meshing

Here, we seek to generate tetrahedral meshes from piecewise-smooth surfaces based on MPS. The details of each step are provided in the following pseudo-code (Algorithm 1).

Algorithm 1 Tetrahedral meshing based on MPS.

Input: 3D domain, Ω , bounded by a surface \mathcal{M} , the minimal sampling radius, r_{min}

Output: tetrahedral mesh T

- 1: Build a uniform grid, G , to voxelize Ω and design a sizing field inside the domain
 - 2: Generate samples on surface \mathcal{M} , and perform boundary optimization
 - 3: Sample the interior of Ω while conforming the boundary
 - 4: Optimize bad dihedral angles
-

The input of our algorithm is a 3D domain, Ω , bounded by a 2-manifold watertight triangular mesh, $\mathcal{M} = \{f_i\}_{i=1}^m$, where f_i is a triangle facet of the boundary \mathcal{M} . Similar to the 2D case, the core of our tetrahedral meshing algorithm consists of MPS in the input domain and a series of mesh optimization operations for improving the meshing quality. However, there are still some difficulties that make tetrahedral meshing more difficult than its 2D counterpart. First, MPS always leads to high-quality triangulations in the 2D domain. In 3D, even well-spaced vertices could create degenerate 3D elements (e.g., slivers). Besides, dealing with boundary conformity is also more difficult in 3D. In the 2D case, there always exists a 2D triangulation that conforms to a set of non-intersecting line segment constraints, but this is no longer true in 3D (Alliez et al., 2005). Fig. 3 shows the pipeline of the proposed framework in 3D.

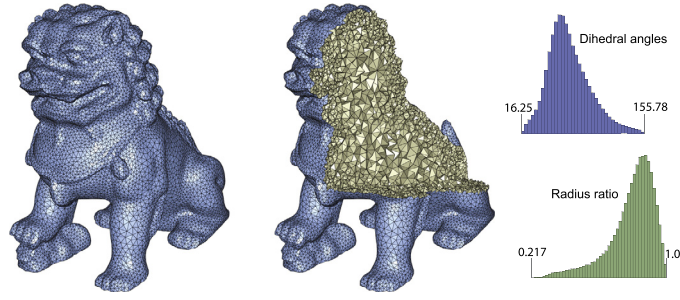


Fig. 1. An example of our tetrahedral meshing result on the Lion model (input 100K triangles). Left and middle are the meshing result (55K vertices and 300K cells) and a cutaway view, respectively. Right are the dihedral angle and radius ratio distributions.

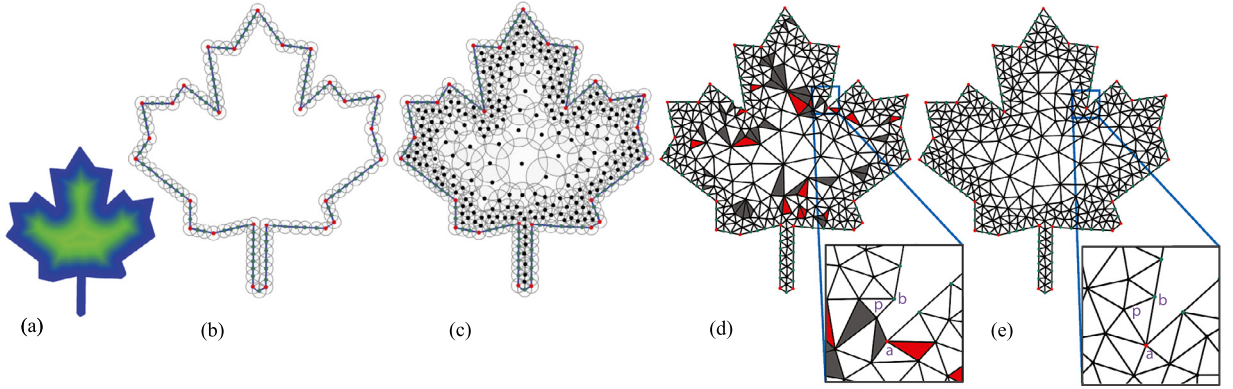


Fig. 2. A 2D illustration of the MPS framework. (a) Input boundary and the sizing field defined in this domain, where warmer colors indicate a larger radius, and cooler colors indicate a smaller radius. (b) Boundary sampling and optimization with disks. The red points indicate sharp corners, while green points are the optimized samples. (c) Interior sampling (black points). (d) Regular triangulation. Triangles with the minimal angle smaller than 30° are shown in dark gray and triangles with undesirable angles by the user input are shown in red. The zoom-in view shows the non-conformal point, p , which causes the boundary part between a and b to disappear. (e) Samples and remeshing after optimization. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

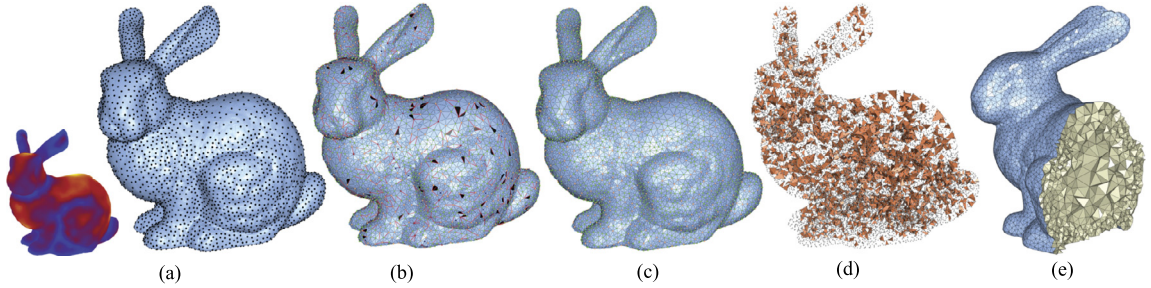


Fig. 3. Overview of our algorithm. Given the input boundary mesh, we first generate surface samples (a) and extract a surface mesh (b). We optimize the boundary remeshing to improve its quality (c). After the boundary optimization, we perform volume sampling while showing slivers (d). Then we optimize the dihedral angles to remove these slivers. (e) is a cutaway view of the final mesh.

4.1. Initialization

Initially, a minimal sampling radius, r_{min} , is specified by the user. In addition, a uniform grid, G , is built to voxelize the 3D domain, Ω . The cell size of G equals $\frac{r_{min}}{\sqrt{3}}$, which guarantees that each grid cell can contain at most one sample in the context of Poisson-disk sampling. The grid cells are classified into three types: boundary cells (intersecting the boundary surface), interior cells, and exterior cells. This grid is useful in two ways. First, it is used to accelerate the sampling operations during the MPS process. In this case, each grid cell stores a flag (initialized as ‘true’) to indicate whether or not this cell is valid. Here, a valid cell means that it is not fully covered by a sample. Further, we utilize the grid G to design a sizing field in the 3D domain for graded meshing, which we describe later.

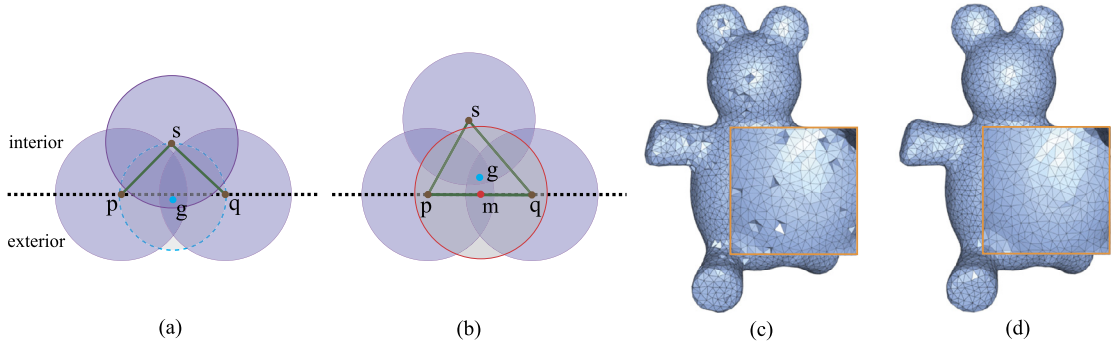


Fig. 4. (a) and (b) illustrate our boundary-protection approach. (c) and (d) show the comparison of tetrahedral meshing results without and with our boundary-protection approach, respectively.

4.2. Surface sampling

At the beginning, we perform MPS only on the boundary surface. The purpose of this step is to approximate the domain boundary by a new surface mesh with better quality and to make the final tetrahedral mesh conform with it. We apply a surface sampling strategy that includes two stages: dart throwing with grid and gap filling. Dart throwing continually chooses a random triangle facet, f_i , from \mathcal{M} and generates a random sample in this facet. To make it unbiased, the probability of selecting the triangle, f_i , is proportional to its area, A_i . Then we reject or accept this new sample based on whether or not it conflicts with any previously accepted samples. The conflict check is fast because we need to check only the sample's local neighboring boundary cells (in the uniform case, it is a $5 \times 5 \times 5$ template of cells). Once a sample is accepted, the flags of the corresponding cells fully covered by this sample are marked as “false”. After observing k (e.g., $k = 300$) consecutive rejections, we switch to the gap-filling stage. In this stage, we clip the triangle facets in \mathcal{M} by the spheres of the samples and collect the gap regions, as described in Yan and Wonka (2013). Then we rethrow darts only in these gaps to generate new samples. This process of gap computation and dart throwing repeats until there are no gaps. The result of this step is shown in Fig. 3(b).

In addition, the input boundary may include sharp edges and feature vertices, especially in mechanical models. To handle these features, we implement a feature-aware sampling step before the surface sampling. Specifically, the features are either provided by the user as 1D curved skeletons or automatically detected by our algorithm. The feature vertices (e.g., corners, cusps) are directly added to the sampling set. Then, we perform 1D MPS to get samples on the creases that are composed by a chain of sharp edges. These feature samples remain fixed in all later steps.

Boundary optimization. Once the surface sampling is done, we extract the surface mesh using the algorithm described in Yan et al. (2009). The extracted mesh might contain undesirable elements, such as irregular vertices whose degrees are less than 5 or larger than 7, and elongated triangles with long edges and small angles (less than 30°). All such elements will result in interior samples that are too close to the boundary, which further result in the tetrahedra with bad shapes. To make the surface remeshing as compact as possible, we propose an approach that interleaves between the valence, angle and edge length optimization. Take the edge length optimization as an example. We iteratively remove the vertices of the long edges whose length is larger than $\frac{\sqrt{3}}{2}(r_1 + r_2)$, where r_1 and r_2 are the sampling radii of the vertices. Then we resample the domain boundary until there are no long edges or until the maximal iteration number is reached. The valence and angle optimization are performed in the same way as in Yan and Wonka (2013), Guo et al. (2015). Fig. 3(c) shows the result of boundary optimization.

4.3. Volume sampling

Next, we generate samples in the interior of the domain. However, if we sample the volume directly after surface sampling, it will generate many slivers near the boundary regions that are difficult to remove in the later steps. Besides, it is difficult to maintain the conformation of the boundary of the final tetrahedral mesh (Fig. 4(c) shows a non-conforming case), since we use Delaunay triangulation instead of constrained Delaunay triangulation to extract the final mesh. To protect the input boundary, we relax the maximal property of MPS near the boundary regions by adding some “virtual samples”. Before introducing our approach in 3D, we first use a simple 2D illustration to make the concept of “virtual sample” clear. As shown in Fig. 4(a), the black dashed line is one edge of the input boundary, points p and q are two boundary samples and they are locally maximal (there is no gap between p and q). If a new sample is generated at point s , a triangle $\triangle spq$ will be extracted. However, since the circumcenter, g , of $\triangle spq$ is outside the domain, this triangle will be discarded, resulting in the boundary part of pq not being maintained in the final mesh. In our approach, we insert a “virtual sample”, m , at the middle point of pq . Then, a newly generated sample, s , should not conflict with m (Fig. 4(b)). In this case, the circumcenter, g , of $\triangle spq$ will be inside the domain, and the edge, pq , is kept. To generalize this approach to a 3D volume, we compute

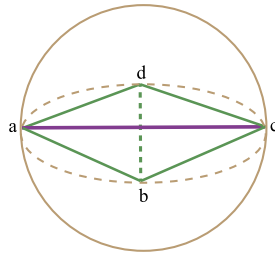


Fig. 5. An example of a sliver with 4 near-cocircular vertices.

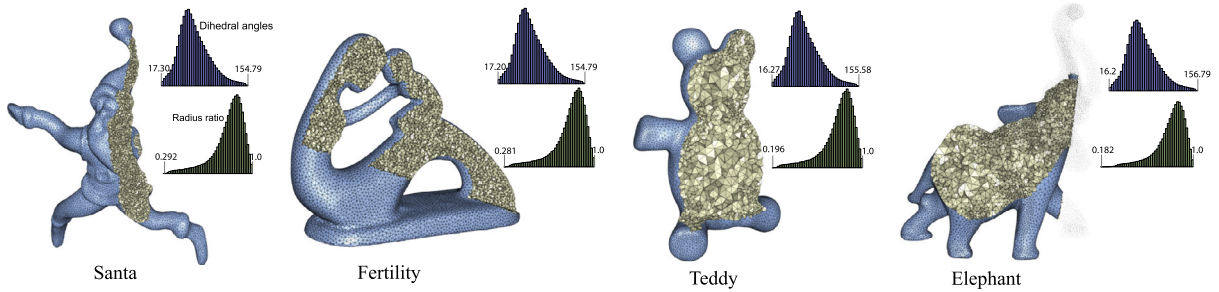


Fig. 6. Tetrahedral meshing results of our algorithm. The two images of the left (Santa and Fertility) are uniform meshes, and the two images of the right (Teddy and Elephant) are graded meshes.

the restricted Voronoi diagram (RVD) (Yan et al., 2009) of the surface samples and put a “virtual sample” at each restricted Voronoi vertex. These “virtual samples” will not be added to the sample set. They are only used to mark the grid cells as invalid, which are fully covered by themselves. Therefore, we cannot generate samples in these cells any more. Fig. 4(c) and Fig. 4(d) show a comparison of meshing results without and with our boundary protection approach.

To sample the interior of the domain, we utilize the uniform grid, G , as the implicit data structure for efficient 3D sampling, as inspired by Ebeida et al. (2012). The sampling process also has two steps. In the first step, we repeatedly choose a random valid grid cell that is not outside of the domain. Then, a random sample is generated in this cell. This new sample is accepted if it is inside the domain and does not conflict with previous samples (including the boundary samples). In the second step, the background grid is refined, and each valid cell is subdivided into eight smaller cells, which are referred to as fragments. We collect the fragments that are not fully covered into a flat array and perform the dart throwing on it. This process repeats until the maximal sampling is achieved. The result of this step is shown in Fig. 3(d).

4.4. Mesh optimization

Finally, given the sampled point set, we can extract the tetrahedral mesh by using 3D Delaunay triangulation, while discarding the Delaunay tetrahedra whose circumscribed sphere centers are outside the domain. However, the quality of this tetrahedral mesh is usually not sufficient to fulfill all quality requirements. We improve the mesh quality through optimizing the sample locations, both on the boundary surface and the domain volume.

Dihedral angle optimization. Although the meshing results based on MPS in 2D and on surfaces exhibit nice properties like the triangle's angle bound $[30^\circ, 120^\circ]$, such properties are not evident in our tetrahedral mesh for dihedral angles. Slivers with very small dihedral angles may exist (see Fig. 3(d)). Further, traditional sliver perturbation (Tournois et al., 2009) or the sliver exudation method (Cheng et al., 2000) is not suitable for our case, because it breaks the Poisson-disk sampling criteria or has little effect. To remove slivers, we use a similar optimization framework as above. For each sliver, the two vertices of the longest edge are first removed (one sliver and its longest edge, ac , are displayed in Fig. 5). Then, we collect the gap fragments caused by removing these vertices and refill them using the method described in Section 4.3. The optimization step terminates when there are no slivers.

4.5. Graded mesh generation

To generate a graded 3D mesh (corresponding to adaptive sampling), we design a sizing field in the whole domain that indicates the local sampling radius. A sizing function, $\mu(\mathbf{x})$, is first defined over the boundary surface, which is also used for adaptive boundary sampling. Then, it is further extrapolated to the grid cells in the interior of the domain by using a fast marching construction method (Alliez et al., 2005). As a result, the elements will be denser on the boundary where high resolution is generally desired, especially in the high-curvature regions. In the interior of the domain, the elements will be coarser and vary in size.

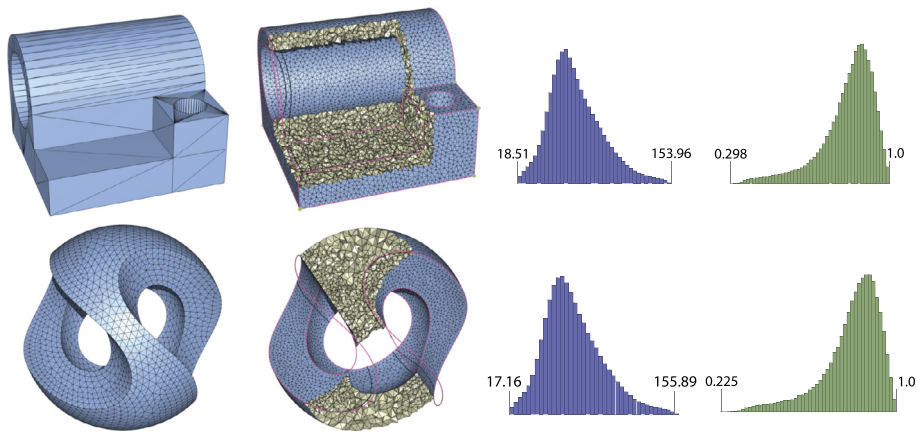


Fig. 7. Tetrahedral meshing results on the Joint (uniform) and Sculpt (adaptive) models with sharp features preserved. Left: input boundary surfaces. Middle: cutaway view of the meshing results, while showing the features (green points are the feature vertices and pink lines are the sharp creases). Right: dihedral angle distributions (blue) and radius ratio distributions (green). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

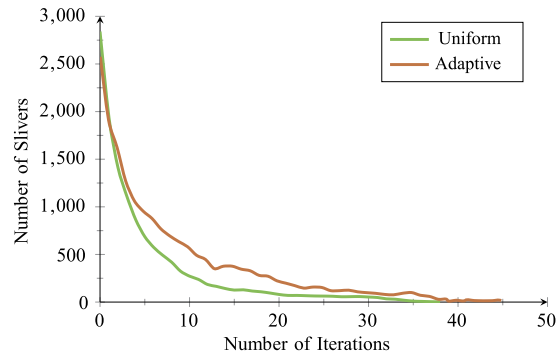


Fig. 8. Number of slivers in relation to the number of iterations. The statistics are obtained on the Santa (uniform with 136K tetrahedra) and Elephant (adaptive with 133K tetrahedra) models.

Table 1

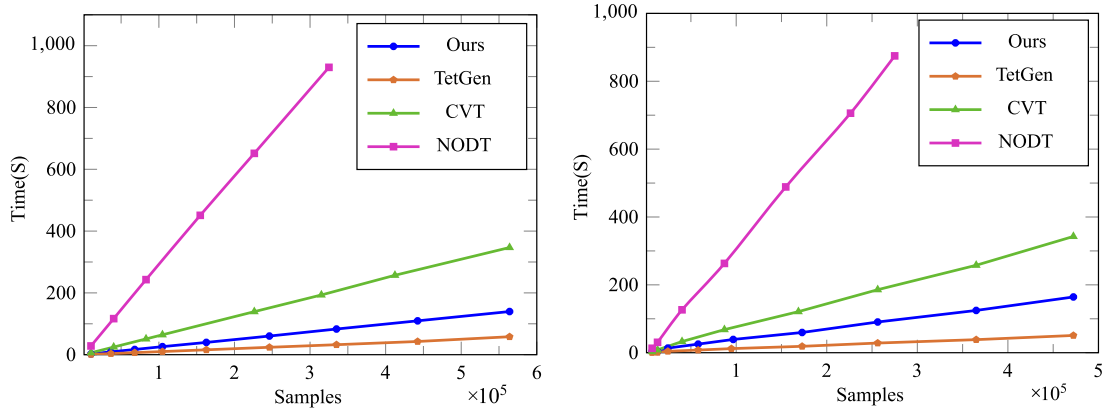
Comparison of uniform tetrahedral meshing qualities. $|X|$ is the number of sampled points, $|\tau|$ is the number of generated tetrahedra.

Model	Method	$ X $	$ \tau $	θ_{min}	θ_{max}	κ_{min}	$\bar{\kappa}$	d_H	$\theta_{<10^\circ}(\%)$	$\theta_{<18^\circ}(\%)$
Sphere	TetGen	4.8K	25K	11.80	159.91	0.197	0.795	0.31	0	0.46
	CVT	5.0K	26K	12.24	161.67	0.236	0.910	0.32	0	0.1
	RODT	5.0K	26K	6.16	169.47	0.133	0.915	0.32	7.5×10^{-3}	0.02
	NODT	4.9K	26K	17.60	152.36	0.320	0.882	0.28	0	0.002
	Ours	4.6K	25K	19.02	149.91	0.346	0.821	0.32	0	0.0
Fertility	TetGen	42K	241K	6.84	165.73	0.104	0.782	0.31	3.7×10^{-4}	0.35
	CVT	41K	227K	7.43	162.37	0.187	0.894	0.26	0.06	0.16
	NODT	40K	216K	15.82	156.76	0.217	0.879	0.32	0	0.01
	Ours	41K	230K	17.20	154.97	0.281	0.808	0.34	0	0.04
Santa	TetGen	37K	192K	7.47	165.98	0.086	0.781	0.27	1.2×10^{-4}	0.29
	CVT	38K	203K	8.45	169.61	0.193	0.911	0.25	0.04	0.19
	NODT	39K	210K	15.05	158.41	0.208	0.863	0.24	0	0.01
	Ours	39K	209K	17.30	154.79	0.292	0.810	0.33	0	0.03
Joint	TetGen	20K	103K	8.79	161.20	0.181	0.789	0.23	3.3×10^{-4}	0.38
	CVT	21K	106K	10.01	165.24	0.191	0.886	0.20	0	0.12
	NODT	21K	105K	16.13	156.42	0.205	0.868	0.22	0	0.008
	Ours	21K	107K	18.51	153.96	0.298	0.815	0.27	0	0.0

Table 2

Comparison of graded tetrahedral meshing qualities.

Model	Method	$ X $	$ \tau $	θ_{min}	θ_{max}	κ_{min}	$\bar{\kappa}$	d_H	$\theta_{<10^\circ}(\%)$	$\theta_{<18^\circ}(\%)$
Botijo	TetGen	39K	199K	7.15	167.15	0.089	0.773	0.32	9.2×10^{-4}	0.38
	CVT	38K	171K	4.94	171.10	0.085	0.768	0.32	0.12	0.98
	NODT	37K	185K	14.63	156.25	0.197	0.874	0.31	0	0.02
	Ours	38K	195K	16.31	155.70	0.201	0.796	0.33	0	0.11
Elephant	TetGen	55K	291K	8.72	162.23	0.148	0.772	0.28	1.2×10^{-4}	0.26
	CVT	56K	289K	4.17	169.89	0.073	0.742	0.33	0.02	0.89
	NODT	56K	294K	13.36	157.38	0.192	0.881	0.30	0	0.01
	Ours	57K	301K	16.20	156.79	0.182	0.796	0.31	0	0.13
Sculpt	TetGen	24K	128K	10.98	157.66	0.149	0.783	0.18	0	0.34
	CVT	23K	127K	5.31	168.54	0.085	0.761	0.28	0.21	0.96
	NODT	24K	128K	16.13	157.11	0.178	0.869	0.24	0	0.01
	Ours	23K	127K	17.16	155.89	0.225	0.802	0.19	0	0.05

**Fig. 9.** Performance comparison of our method with TetGen, CVT and NODT. Left: uniform tetrahedral mesh; right: graded tetrahedral mesh. The running time of our algorithm also contains building the uniform grid.

In our approach, we adopt the local feature size (lfs) as the boundary sizing function, $\mu(\mathbf{x})$, but other functions can also be used with the same framework. We also set a maximal sampling radius, $r_{max} = \lambda r_{min}$ (in default $\lambda = 8$), to avoid very big radii. The adaptive sampling process is similar to the uniform case, except for some differences in the details. In the dart-throwing stage, we evaluate the radius, r_i , for each newly generated sample, s_i . Then, we perform a conflict check using a subset of $\lceil \frac{2r_i}{r_{min}/\sqrt{3}} \rceil^3$ grid cells. In the mesh extraction stage, we adopt regular triangulation instead of Delaunay triangulation, where each sample, s_i , is associated with a weight, $w_i = r_i^2$. Instead of using the traditional Euclidean distance, the distance function used by the regular triangulation is the power distance, represented by $d(\mathbf{s}_i, w_i, \mathbf{s}_j, w_j) = \|\mathbf{s}_i - \mathbf{s}_j\|^2 - w_i - w_j$.

5. Experimental results

Our algorithm is implemented in C++ using the CGAL library (CGAL, 2015) for 3D Delaunay/regular triangulation. The results shown in this paper are obtained on an Intel i7 3.07 GHz desktop with 12 GB memory.

Tetrahedral meshing. To verify the robustness and the validity of the proposed approach, we ran our algorithm on various input domains with arbitrary topologies. Intuitively, our algorithm worked well because Poisson-disk samplers are well distributed, leading to well-shaped tetrahedra. Figs. 1 and 6 illustrate the tetrahedral meshing results, as well as the histograms of the dihedral angle and radius ratio distributions. The radius ratio, κ , is defined as $\kappa = \frac{3r_{in}}{r_{cir}}$, where r_{in} and r_{cir} are the inscribed and circumscribed sphere radius of a tetrahedron, respectively. Furthermore, our algorithm can preserve the features of the input domain. Two examples of feature preserving are illustrated in Fig. 7.

In the mesh optimization stage, we set the default value of the desired minimal dihedral angle to $16^\circ \sim 18^\circ$. The optimization usually converges within 30 ~ 50 iterations. Fig. 8 shows the convergence behavior of the mesh optimization process. Though we do not have a formal theoretical proof of the convergence, our algorithm works well in practice with various input domains.

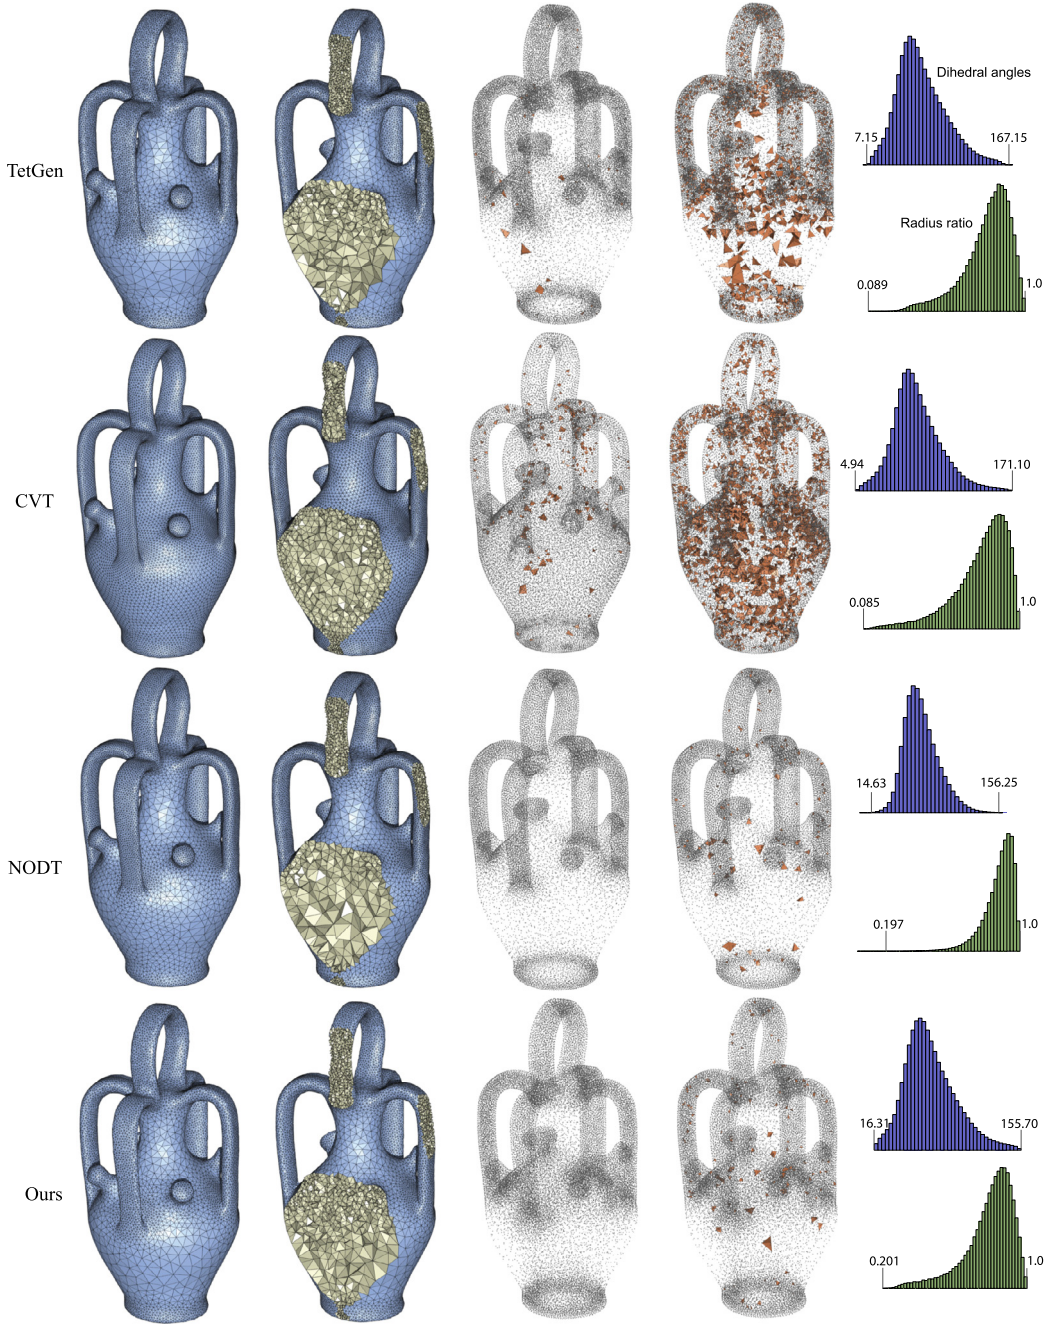


Fig. 10. Comparison between different algorithms in generating a graded tetrahedral mesh. For each row, from left to right we show the meshing result, its cutaway view, slivers with dihedral angles smaller than 10° and 18° , and distributions of dihedral angles (blue) and radius ratios (green). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Quality comparison. Next, we compare our approach with current state-of-the-art techniques, including the interleaved Delaunay refinement and optimization method (NODT) (Tournois et al., 2009), Centroidal Voronoi Tessellation (CVT) (Yan et al., 2013), revisited optimal Delaunay triangulation (RODT) (Chen et al., 2014) and a Delaunay-insertion-based approach – TetGen (Si, 2015). Note that for TetGen, we provided as input the boundary of our optimized mesh for fair comparison, while NODT is implemented in CGAL with sliver perturbation.

Tables 1 and 2 list the mesh quality comparison results for uniform and adaptive meshing, respectively. Here θ_{min} and θ_{max} are the minimal and maximal dihedral angles of a tetrahedron, κ_{min} and $\bar{\kappa}$ are the minimal and average radius ratios. d_H is the Hausdorff distance (divided by the diagonal length of the input mesh bounding box) between the input boundary and remeshing result. $\theta_{<10^\circ}$ and $\theta_{<18^\circ}$ are the percentages of dihedral angles smaller than 10° and 18° , respectively. The

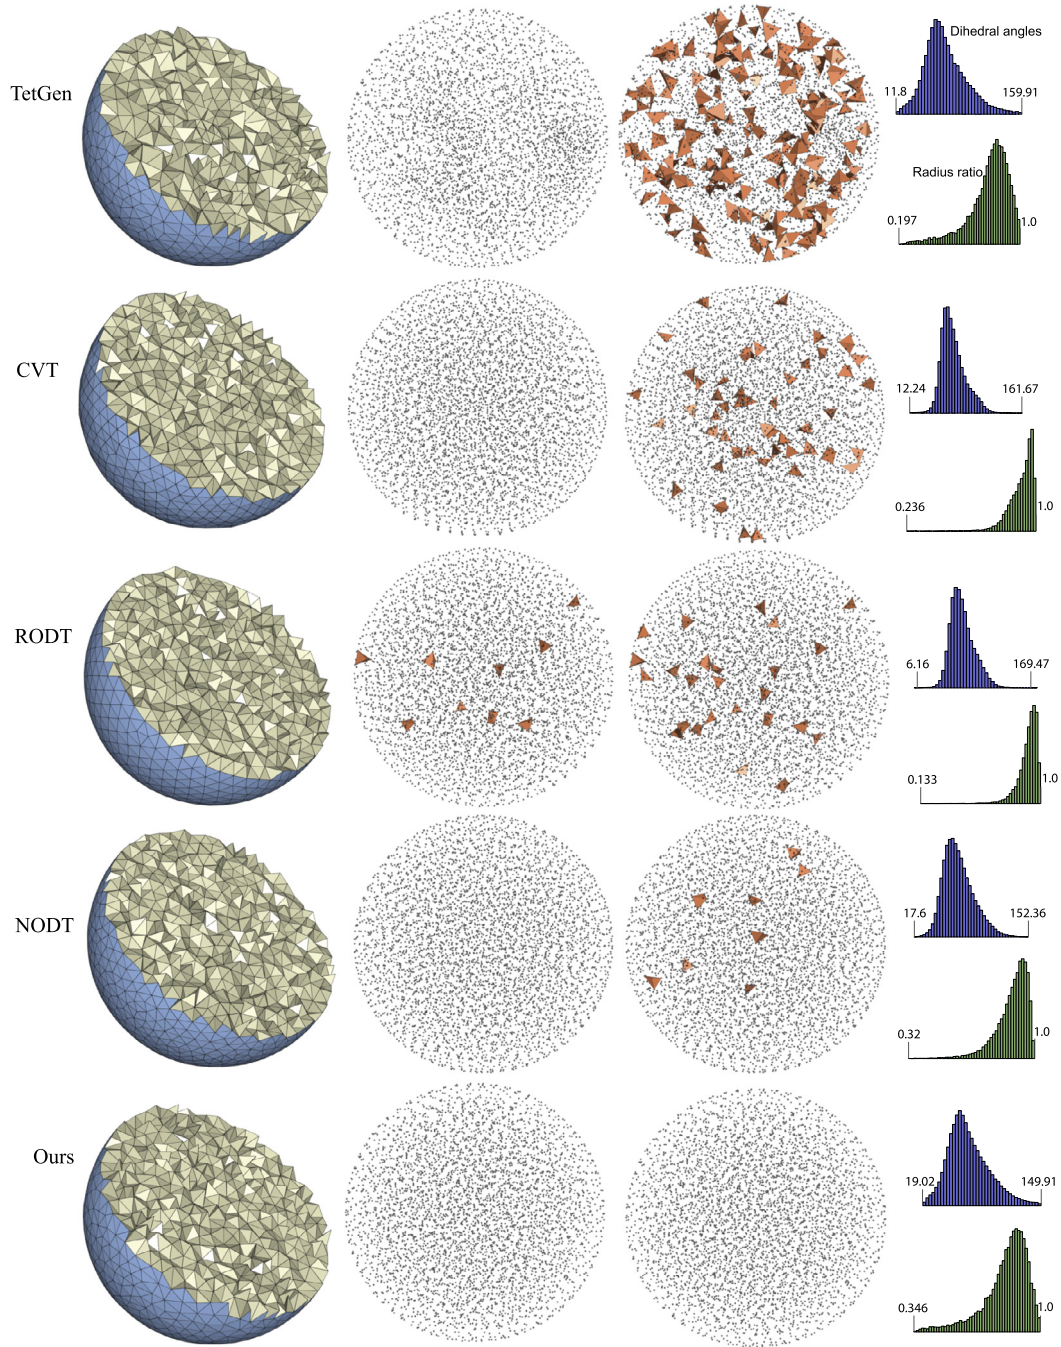


Fig. 11. Comparison of our uniform tetrahedral mesh with those produced by other algorithms. For each row, from left to right we show the cutaway view of meshing result, slivers with dihedral angles smaller than 10° and 18° , and distributions of dihedral angles (blue) and radius ratios (green). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

best results are highlighted using bold font. Both of these tables suggest that our method achieves best min/max dihedral angles, θ , and the minimal radius ratios, κ_{min} . Our $\bar{\kappa}$ is not as good as NODT, CVT, RODT, while our $\theta_{<18^\circ}$ is comparable with NODT method. The reason is that we generate the point set by employing the randomized sampling, while the other three methods tend to generate more regular structures. The histogram distributions of dihedral angles and radius ratios shown in Figs. 10 and 11 also confirm this observation. However, randomized meshes are preferred in many simulation applications (Ebeida et al., 2011b), such as fracture simulations (Bolander and Saito, 1998) and fluid simulation (de Goes et al., 2015). Additional comparison results are provided in the supplemental material.

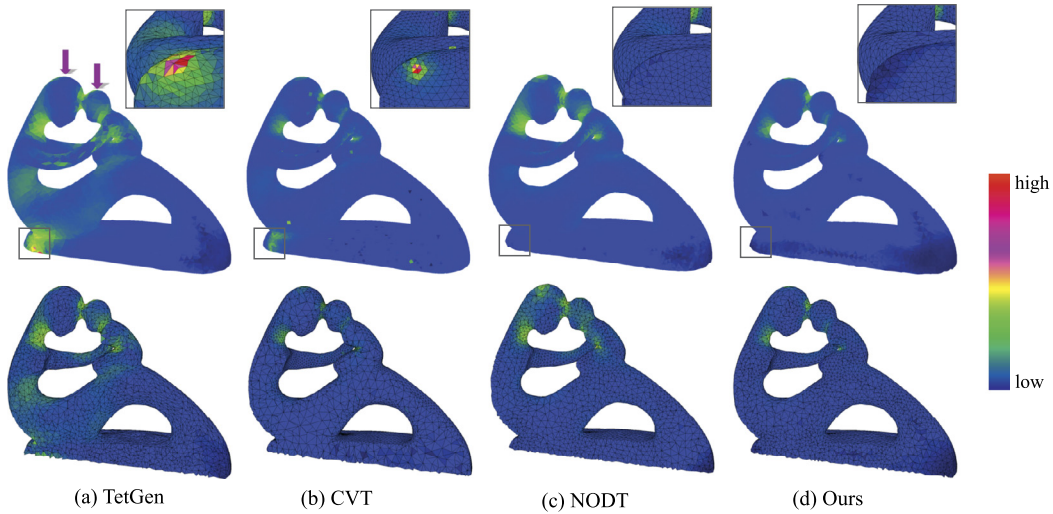


Fig. 12. Structural analysis with the same force settings (arrows in (a)). The top row shows the stress map with color coding (zoom-in figures show a slightly different view), and the bottom row shows the interior structure.

Performance. We now evaluate the performance of our algorithm. The processing time mainly contains three components: sampling, Delaunay/regular triangulation, and mesh optimization. With the uniform grid as the implicit flat-array data structure, the sampling stage could achieve good performance. In the Delaunay/regular triangulation, we build a full triangulation once after the surface sampling. Then, in the later mesh extraction and optimization steps, we update dynamically for point removal and insertion. Note that in adaptive sampling, the resolution of the sub-grid used for conflict checking is much larger than that used for uniform sampling. Therefore, it reduces the performance of the graded mesh generation.

Fig. 9 presents a comparison of the time our algorithm consumes with that of previous methods. All the timing curves are computed on the Botijo model, with an increasing number of sample points. In all these algorithms, the dihedral angle bound is set to 15 degrees. As illustrated in this figure, our algorithm is reasonably fast. Though TetGen is the most efficient method, its meshing results contain many slivers. In summary, our algorithm is capable of generating high-quality tetrahedral meshes, while maintaining good efficiency.

Structural analysis. To demonstrate that our meshing method is useful in real-world engineering applications, we perform structural analysis on the generated tetrahedral meshes as shown in Fig. 12. The stress maps are computed by using the OOFEM library (Patzák and Rypl, 2012), an object-oriented finite element solver. First, an isotropic linear elastic material model is assigned to each mesh. After the user's exterior forces (defined on the boundary facets of some specific tetrahedra) are included, a non-linear static analysis is performed. The output data provide a quantitative description of the stress over all tetrahedra, which is visualized using a continuous stress map. As demonstrated in this figure, the stress distribution over our meshing structure is more uniform than over the others, indicating that our mesh can resist more forces.

Comparison to sliver removal approaches. There are several approaches that are able to improve the meshing quality by suppressing slivers as a post-processing step. Two well-known representative approaches are sliver exudation (Cheng et al., 2000) and sliver perturbation (Tournois et al., 2009). The former tunes a weighted Delaunay triangulation that is free of slivers (minimal dihedral angles smaller than 5° in their method) based on a weight reassignment computation. This method eliminates slivers without adding new points or changing the positions of current points, while updating their connectivity. Sliver perturbation performs a gradient ascent search over the radius of the circumscribed sphere of a sliver as well as a gradient descent over the sliver's volume.

To justify the validity of our mesh optimization strategy, we compare our results with sliver exudation (Cheng et al., 2000) and sliver perturbation (Tournois et al., 2009). For a fair comparison, we first generate an initial tetrahedral mesh using our method without any optimization. Then, we compare with sliver exudation and sliver perturbation for improvement on this mesh. Fig. 13 shows the comparison results using the Bimba model, where a graded mesh is generated. We observe that sliver exudation has a limited effect, while sliver perturbation cannot always remove the slivers for a given threshold. Our strategy is able to generate meshes with larger minimal dihedral angles than these two state-of-the-art approaches.

Limitations. The main limitation of our proposed framework is that we do not have a theoretical proof of the lower bound of the dihedral angle that we can achieve. Our algorithm cannot guarantee the generation of a pleasant result if the required minimal dihedral angle bound is too large (e.g., 18°). In addition, the proposed method has similar difficulties as previous methods (ODT, CVT, Delaunay refinement, etc.) in handling thin-sheet regions. We cannot generate a valid mesh if the sampling density does not meet the local feature size of the domain. Fig. 14 shows such an example. Finally, although

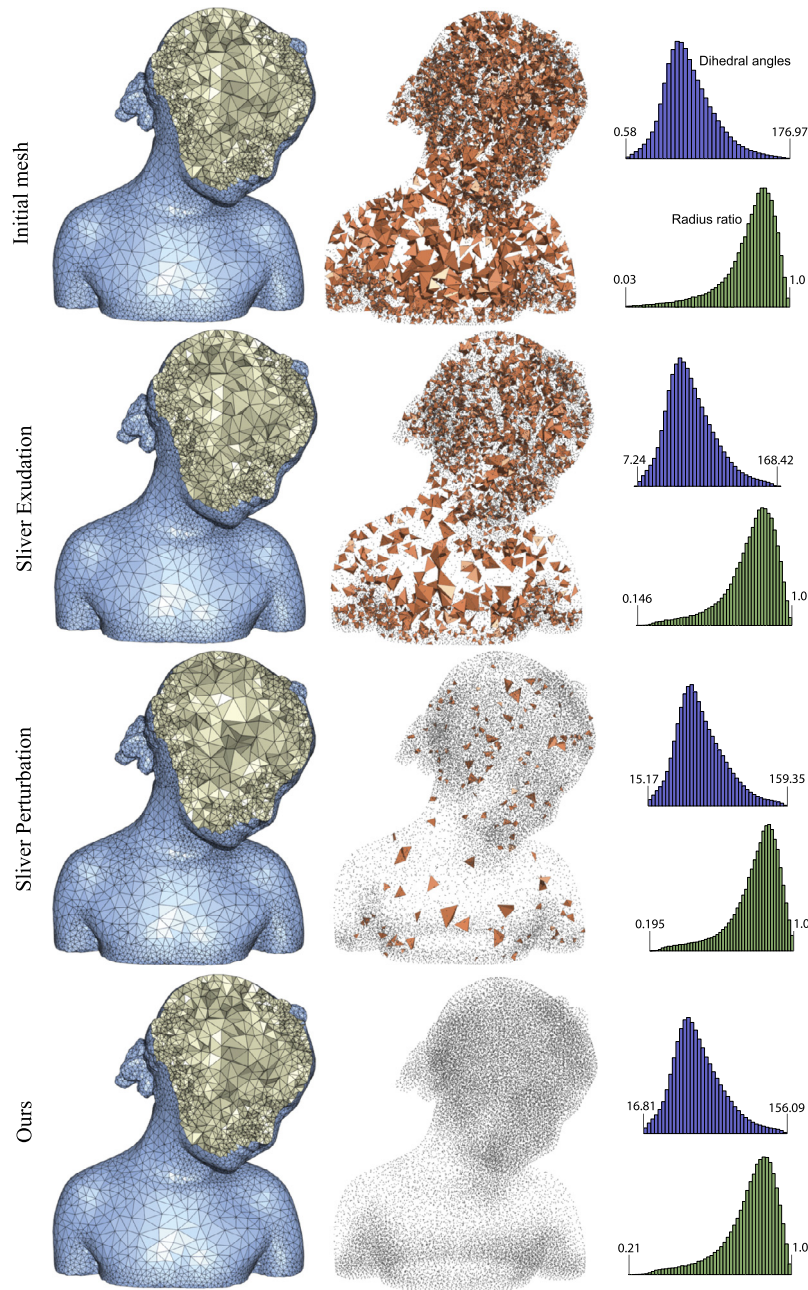


Fig. 13. Comparison with sliver exudation (Cheng et al., 2000) and sliver perturbation (Tournois et al., 2009) on the Bimba model. Left: cutaway view of the initial mesh and the optimized mesh; middle: sliver tetrahedra for a dihedral angle of 17 degrees; right: dihedral angle distributions (blue), and radius ratio distributions (green). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

sampling has been demonstrated to be useful in a variety of graphics applications including remeshing and modeling, exploring the relationship between blue-noise properties and mesh generation is still an open problem.

6. Conclusion

In this paper, we introduced a simple yet efficient method for high-quality tetrahedral mesh generation that is superior to the current state-of-the-art approaches. Our approach is based on the maximal Poisson-disk sampling framework, while taking boundary preservation and mesh grading into account. Several optimization steps are performed to greatly improve the meshing quality. We demonstrated the validity of our algorithm by meshing a variety of complex domains even with

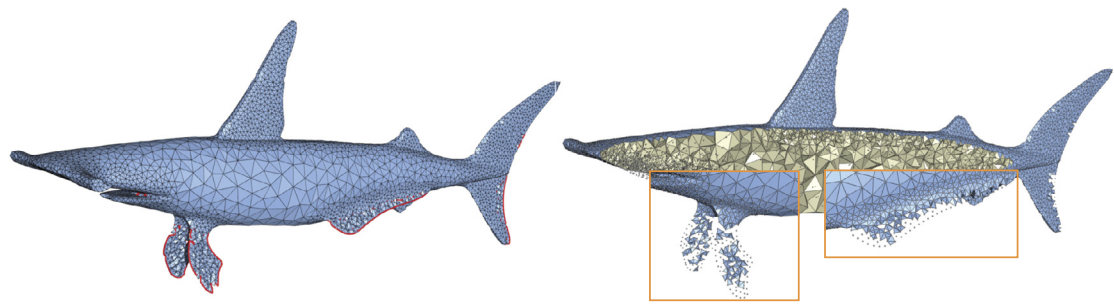


Fig. 14. Our algorithm cannot guarantee a pleasant result when the minimal radius is larger than the local feature size of the domain. Left: surface remeshing after the boundary sampling stage. The non-manifold edges are shown in red. Right: tetrahedral meshing result with two zoom-in views. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

sharp features, as well as comparing it with the state-of-the-art approaches. In the future, we plan to extend our approach to anisotropic mesh generation. We are also interested in meshing implicit surfaces such as isosurfaces.

Acknowledgements

This work was partially funded by the National Natural Science Foundation of China (Nos. 61372168, 61331018, 61272225, 61572274 and 61271431), the China National 863 Program (No. 2015AA016402), and Foreign 1000 Talent Plan (WQ201344000169).

Appendix A. Supplementary material

Supplementary material related to this article can be found online at <http://dx.doi.org/10.1016/j.cagd.2016.02.004>.

References

- Alliez, P., Cohen-Steiner, D., Yvinec, M., Desbrun, M., 2005. Variational tetrahedral meshing. In: *Proceedings of ACM SIGGRAPH 2005*. ACM Trans. Graph. 24 (3), 617–625.
- Ando, R., Thürey, N., Wojtan, C., 2013. Highly adaptive liquid simulations on tetrahedral meshes. *ACM Trans. Graph.* 32 (4), 103.
- Barber, C.B., Dobkin, D., Huhdanpaa, H., 1996. The quickhull algorithm for convex hulls. *ACM Trans. Math. Softw.* 22 (4), 469–483.
- Bern, M., Plassmann, P., 2000. *Mesh Generation*. Elsevier Science.
- Bolander, J., Saito, S., 1998. Fracture analyses using spring networks with random geometry. *Eng. Fract. Mech.* 61 (5), 569–591.
- Bronson, J., Levine, J.A., Whitaker, R., 2014. Lattice cleaving: a multimaterial tetrahedral meshing algorithm with guarantees. *IEEE Trans. Vis. Comput. Graph.* 20 (2), 223–237.
- CGAL, 2015. *CGAL, Computational Geometry Algorithms Library*. <http://www.cgal.org>.
- Chen, L., 2004. Mesh smoothing schemes based on optimal Delaunay triangulations. In: *IMR*, pp. 109–120.
- Chen, Z., Wang, W., Lévy, B., Liu, L., Sun, F., 2014. Revisiting optimal Delaunay triangulation for 3D graded mesh generation. *SIAM J. Sci. Comput.* 36 (3), A930–A954.
- Cheng, S.-W., Dey, T.K., Edelsbrunner, H., Facello, M.A., Teng, S.-H., 2000. Sliver exudation. *J. ACM* 47 (5), 883–904.
- Cheng, S.-W., Dey, T.K., Levine, J.A., 2007. A practical Delaunay meshing algorithm for a large class of domains. In: *16th Intl. Meshing Roundtable*, pp. 477–494.
- Cheng, S.-W., Dey, T.K., Shewchuk, J.R., 2012. *Delaunay Mesh Generation*. CRC Press.
- Chew, L.P., 1997. Guaranteed-quality Delaunay meshing in 3D. In: *Proceedings of the Thirteenth Annual Symposium on Computational Geometry*, pp. 391–393.
- Choi, W.-Y., Kwak, D.-Y., Son, I.-H., Im, Y.-T., 2003. Tetrahedral mesh generation based on advancing front technique and optimization scheme. *Int. J. Numer. Methods Eng.* 58 (12), 1857–1872. <http://dx.doi.org/10.1002/nme.840>.
- Dardenne, J., Valette, S., Siauve, N., Burais, N., Prost, R., 2009. Variational tetrahedral mesh generation from discrete volume data. In: *Proc. CGI 2009*. Vis. Comput. 25 (5), 401–410.
- de Goes, F., Wallez, C., Huang, J., Pavlov, D., Desbrun, M., 2015. Power particles: an incompressible fluid solver based on power diagrams. In: *Proc. SIGGRAPH*. ACM Trans. Graph. 34.
- Du, Q., Wang, D., 2003. Tetrahedral mesh generation and optimization based on centroidal Voronoi tessellations. *Int. J. Numer. Methods Eng.* 56 (9), 1355–1373.
- Du, Q., Faber, V., Gunzburger, M., 1999. Centroidal Voronoi tessellations: applications and algorithms. *SIAM Rev.* 41, 637–676.
- Ebeida, M., Mitchell, S., 2012. Uniform random Voronoi meshes. In: *Quadros, W. (Ed.), Proceedings of the 20th International Meshing Roundtable*. Springer, Berlin, Heidelberg, pp. 273–290.
- Ebeida, M.S., Davidson, A.A., Patney, A., Knupp, P.M., Mitchell, S.A., Owens, J.D., 2011a. Efficient maximal Poisson-disk sampling. In: *Proc. SIGGRAPH*. ACM Trans. Graph. 30 (4), 49.
- Ebeida, M.S., Mitchell, S.A., Davidson, A.A., Patney, A., Knupp, P.M., Owens, J.D., 2011b. Efficient and good Delaunay meshes from random points. *Comput. Aided Des.* 43 (11), 1506–1515.
- Ebeida, M.S., Mitchell, S.A., Patney, A., Davidson, A.A., Owens, J.D., 2012. A simple algorithm for maximal Poisson-disk sampling in high dimensions. In: *Proc. EUROGRAPHICS*. Comput. Graph. Forum 31 (2), 785–794.
- Edelsbrunner, H., 2001. *Geometry and Topology for Mesh Generation*. Cambridge University Press.
- Frey, P.J., George, P.-L., 2001. *Mesh Generation: Application to Finite Elements*. Hermes Science.

- Geuzaine, C., Remacle, J.-F., 2009. Gmsh: a three-dimensional finite element mesh generator with built-in pre- and post-processing facilities. 79 (11), 1309–1331.
- Guo, J., Yan, D.-M., Jia, X., Zhang, X., 2015. Efficient maximal Poisson-disk sampling and remeshing on surfaces. *Comput. Graph.* 46 (6–8), 72–79.
- Ito, Y., Shih, A.M., Soni, B.K., 2004. Reliable isotropic tetrahedral mesh generation based on an advancing front method. In: IMR, pp. 95–106.
- Jamin, C., Alliez, P., Yvinec, M., Boissonnat, J.-D., 2014. CGALmesh: a generic framework for Delaunay mesh generation. *ACM Trans. Math. Softw.*
- Jiao, X., Wang, D., Zha, H., 2011. Simple and effective variational optimization of surface and volume triangulations. *Eng. Comput.* 27 (1), 81–94.
- Jones, T.R., 2006. Efficient generation of Poisson-disk sampling patterns. *J. Graph. Tools* 11 (2), 27–36.
- Klingner, B.M., Shewchuk, J.R., 2007. Aggressive tetrahedral mesh improvement. In: 16th Intl. Meshing Roundtable, pp. 3–23.
- Klingner, B.M., Feldman, B.E., Chentanez, N., O'Brien, J.F., 2006. Fluid animation with dynamic meshes. In: *Proc. SIGGRAPH*. *ACM Trans. Graph.* 25 (3), 820–825.
- Labelle, F., Shewchuk, J.R., 2007. Isosurface stuffing: fast tetrahedral meshes with good dihedral angles. In: *Proc. SIGGRAPH*. *ACM Trans. Graph.* 26 (3).
- Lagae, A., Dutré, P., 2008. A comparison of methods for generating Poisson disk distributions. *Comput. Graph. Forum* 27 (1), 114–129.
- Owen, S.J., 1998. A survey of unstructured mesh generation technology. In: *International Meshing Roundtable*, pp. 239–267.
- Patzák, B., Ryppl, D., 2012. Object-oriented, parallel finite element framework with dynamic load balancing. *Adv. Eng. Softw.* 47 (1), 35–50.
- Schöberl, J., 1997. NETGEN An advancing front 2D/3D-mesh generator based on abstract rules. *Comput. Vis. Sci.* 1, 41–52.
- Shewchuk, J.R., 2002. What is a good linear element? Interpolation, conditioning, and quality measures. In: 11th Intl. Meshing Roundtable, pp. 115–126.
- Si, H., 2015. Tetgen, a Delaunay-based quality tetrahedral mesh generator. *ACM Trans. Math. Softw.* 41 (2), 11.
- Tournais, J., Wormser, C., Alliez, P., Desbrun, M., 2009. Interleaving Delaunay refinement and optimization for practical isotropic tetrahedron mesh generation. In: *Proc. SIGGRAPH*. *ACM Trans. Graph.* 28 (3), 75.
- Vanderzee, E., Hirani, A.N., Guoy, D., Ramos, E., 2010. Well-centered triangulation. *SIAM J. Sci. Comput.* 31 (6), 4497–4523.
- Yan, D.-M., Wonka, P., 2013. Gap processing for adaptive maximal Poisson-disk sampling. *ACM Trans. Graph.* 32 (5), 148.
- Yan, D.-M., Lévy, B., Liu, Y., Sun, F., Wang, W., 2009. Isotropic remeshing with fast and exact computation of restricted Voronoi diagram. *Comput. Graph. Forum* 28 (5), 1445–1454.
- Yan, D., Wang, W., Lévy, B., Liu, Y., 2010. Efficient computation of 3D clipped Voronoi diagram. In: *Advances in Geometric Modeling and Processing – GMP*, pp. 269–282.
- Yan, D.-M., Wang, W., Lévy, B., Liu, Y., 2013. Efficient computation of clipped Voronoi diagram for mesh generation. *Comput. Aided Des.* 45 (4), 843–852.
- Yu, Z., Wang, J., Gao, Z., Xu, M., Hoshijima, M., 2014. New software developments for quality mesh generation and optimization from biomedical imaging data. *Comput. Methods Programs Biomed.* 113 (1).