Intrinisic Mesh Simplification

RANDY SHOEMAKER, College of William & Mary, USA
SAM SARTOR, College of William & Mary, USA
PIETER PEERS, College of William & Mary, USA

This paper presents a novel simplification method for removing vertices from an intrinsic triangulation corresponding to extrinsic vertices lying on near-developable (i.e., with limited Gaussian curvature) and general surfaces. We greedily process all intrinsic vertices with an absolute Gaussian curvature below a user selected threshold. For each vertex, we repeatedly perform local intrinsic edge flips until the vertex reaches the desired valance (three for internal vertices or two for boundary vertices) such that removal of the vertex and incident edges can be locally performed in the intrinsic triangulation. Each removed vertex's intrinsic location is tracked via (intrinsic) barycentric coordinates that are updated to reflect changes in the intrinsic triangulation. We demonstrate the robustness and effectiveness of our method on the *Thingi10k* dataset and analyze the effect of the curvature threshold on the solutions of PDEs.

CCS Concepts: • **Computing methodologies** → **Shape analysis**.

Additional Key Words and Phrases: Intrinsic Triangulation, Simplification

1 INTRODUCTION

Geometric data encountered in the wild are often *"messy"* from a geometry processing perspective, necessitating the need for robustified processing methods [Hu et al. 2018; Qi et al. 2022; Sawhney and Crane 2020; Sellén et al. 2019; Zhou et al. 2016]. Intrinsic triangulation frameworks [Fisher et al. 2007; Gillespie et al. 2021a; Sharp et al. 2019] have been proposed as an alternative strategy for robust geometry processing in the wild. Intrinsic geometry processing approaches geometry processing from an intrinsic view where all operations are expressed as combinations of atomic operations (e.g., edge flipping, face splitting, inserting edges, etc.) defined on distances between points over the 2D manifold. While existing intrinsic triangulation frameworks differ in data-structure and efficiency of certain atomic operations, they all share that each "extrinsic" vertex has an immutable counterpart in the intrinsic triangulation, and consequently, vertex removal of initial extrinsic vertices is not supported. Many in-the-wild triangle meshes are finely triangulated in order to faithfully approximate curved surfaces in $\mathbb{R}^3$ by piecewise planar surfaces. However, a significant portion of in-the-wild triangle meshes are designed with CAD tools or are the result of 3D scans of real-world surfaces formed by combining developable patches. Intrinsically, such developable parts are isometric to a plane. For example consider a cap-less cylinder which can be fully modeled using just two intrinsic triangles (see inset). However, to faithfully capture the curvature, thousands of extrinsic triangles are needed in $\mathbb{R}^3$. Not only does blindly embedding a heavily tessellated

Authors' addresses: Randy Shoemaker, rwshoemaker@wm.edu, College of William & Mary, Williamsburg, VA, USA; Sam Sartor, slsartor@wm.edu, College of William & Mary, Williamsburg, VA, USA; Pieter Peers, ppeers@siggraph.org, College of William & Mary, Williamsburg, VA, USA.
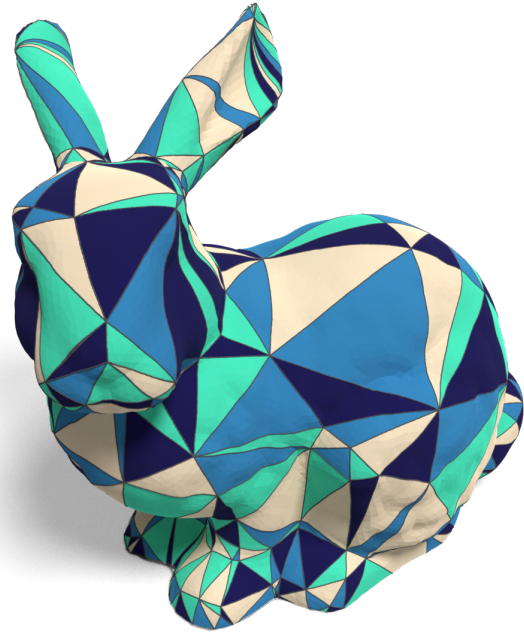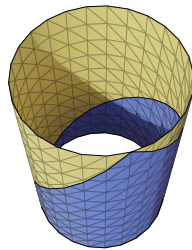
Fig. 1. Our Intrinsic Simplification algorithm simplifies the metric of a surface by approximating the intrinsic geometry with developable patches.

developable surface incur significant overhead, it also impacts the efficiency of many downstream processing algorithms such as optimal Delaunay triangulations [Chen and Xu 2004], adaptive intrinsic mesh refinement [Sharp et al. 2019], computing geodesics [Sharp and Crane 2020], geodesic distances, and other tasks.

In this paper we present a topology-preserving method for simplifying a triangle mesh directly on the intrinsic manifold. A key insight is that vertices with zero Gaussian curvature can be removed without impacting the accuracy of the metric defined on the embedding. Moreover, a locally developable approximation can be obtained by allowing vertices with small Gaussian curvature to be removed as in Figure 1. However, classic vertex merging and edge collapse require updating of edge lengths in the intrinsic setting which can be non-trivial when removing a vertex with non-zero curvature. Instead, we introduce a novel intrinsic simplification method based on edge flipping, a stable atomic intrinsic operation. For each vertex we would like to remove, we perform edge flips until the vertex has valance 3 (or valance 2 for vertices on the boundary). We can then remove the vertex if the resulting triangulation remains valid. If the triangulation becomes invalid, we undo the edge flips in reverse order and reschedule the vertex for later evaluation. We process vertices in a greedy *"intrinsically-flattest-first"* order until no more vertices can be removed. In addition, we keep track of the intrinsic

location of deleted vertices by their intrinsic barycentric coordinates. Compared to extrinsic simplification which approximates the surface with fewer planar triangles, intrinsic simplification can be seen as approximating the surface with developable patches. Since the metric of an intrinsic triangulation is fully described by its edge lengths, intrinsic simplification can be seen as simplification of the metric rather than the extrinsic mesh.

We demonstrate and validate the robustness of our method on the *Thingi10k* dataset [Zhou and Jacobson 2016], and evaluate the impact of relaxing the Gaussian curvature threshold on PDE solutions.

2 RELATED WORK

Mesh Simplification. There exists a vast body of work on *extrinsic* mesh simplification and an exhaustive enumeration is beyond the scope of this paper. Instead we will focus on seminal papers in this area and contrast them against our intrinsic mesh simplification method; we refer the reader to [Khan et al. 2020] (Sec. 4.1) for an in-depth systematic review. Extrinsic mesh simplification methods aim to reduce the number of vertices in the mesh such that some quality metric is best preserved. The most commonly preserved quality is the visual appearance of a mesh [Cohen-Steiner et al. 2004; Garland and Heckbert 1997; Popović and Hoppe 1997; Rossignac and Borrel 1993; Schroeder et al. 1992]. Vertex decimation [Schroeder et al. 1992] iteratively deletes vertices according to an extrinsic criterion and the resulting hole is carefully re-triangulated. Our method of vertex deletion is similar to vertex decimation, except that we employ intrinsic edge flips until the ring of a vertex is a triangle which can be removed without retriangulation. In their seminal work, Garland and Heckbert [1997] greedily merge vertices to minimize a quadric error metric (QEM) via edge contraction. We also follow a greedy approach, but instead of QEM, we use Gaussian curvature to drive the simplification.

Recently, Lescoat et al. [2020] proposed spectral mesh simplification, a greedy extrinsic mesh simplification strategy, that aims to preserve the intrinsic geometry (i.e., minimize the change in the first k eigenvectors of the Laplace-Beltrami operator). Spectral mesh simplification is able to reduce the number of extrinsic triangles while minimizing errors when computing spectral distances. However, spectral mesh simplification is relatively computationally expensive and limited to reducing extrinsic surfaces. In contrast, intrinsic mesh simplification is computationally light weight and more efficient in reducing the number of elements in developable patches while preserving the intrinsic geometry.

In concurrent work, Liu et al. [2023] propose a similar method for intrinsic mesh simplification via intrinsic error metrics. Their method tracks an approximation of the accumulated error during simplification, informing the order of vertex deletion. Our method uses the (dynamic) magnitude of the Gaussian curvature of a vertex to decide deletion order, which is easily updated after deletion. Our method removes vertices by reducing them to a desired valence whereby they can be safely removed. Once the valence of a vertex has been reduced, our method implicitly intrinsically flattens the neighborhood of a vertex via a simple update whereas Liu et al. [2023] flattens the neighborhood of a vertex prior to valence reduction, requiring iterative optimization.

Intrinsic Triangulations. Intrinsic triangulation frameworks [Fisher et al. 2007; Gillespie et al. 2021a; Sharp et al. 2019] provide tools and atomic operations to perform geometry processing algorithms that only rely on intrinsic information directly on the intrinsic mesh (e.g., geodesic distance [Sharp and Crane 2020], computing distortion minimizing homeomorphisms [Takayama 2022], or algorithms that rely on the Laplace-Beltrami operator [Botsch et al. 2010]). All existing intrinsic triangulation frameworks support edge flipping, and the Signpost [Sharp et al. 2019] and Integer Coordinate frameworks [Gillespie et al. 2021a] support additional atomic operations such as adding vertices, repositioning (added) vertices, and computing common subdivisions. However, none of the existing intrinsic triangulation frameworks currently support the removal of an extrinsic vertex's intrinsic counterpart. Our method only relies on edge flipping to remove vertices, opening the door to possible adaptation to other current and future intrinsic triangulation frameworks.

Global Parameterization. Global mesh parameterization is similar to intrinsic mesh simplification in that both produce a base domain with a mapping to the original mesh. Various in depth surveys discuss the variety of parameterization methods [Floater and Hormann 2005; Hormann et al. 2007; Sheffer et al. 2006]. In our case the base domain is an intrinsic triangulation and the mapping is specified by the barycentric coordinates of removed extrinsic vertices with respect to the intrinsic triangulation. Global parameterization methods utilizing simplicial and quadrilateral complexes as the base domain that apply iterative decimation schemes are most relevant to this work, for example [Bommes et al. 2013; Khodakovsky et al. 2003; Lee et al. 1998]. Such methods typically require some kind of embedding of mesh vertices as the algorithm progresses (usually in two dimensions) and are therefore unable take advantage of the additional degrees of freedom offered by intrinsic triangulations [Sharp et al. 2019] which are unencumbered by the requirement to maintain an embedding. Furthermore global parameterization methods, such as those producing integer grid maps (IGM) are computationally intensive whereas our edge flipping procedure is fast. Ebke et al. [2016] offer a framework for computing an IGM on large meshes via a decimation objective based on the change in Gaussian curvature between the coarse and fine mesh. While our vertex ordering criteria are similar, their method is not designed to support intrinsic triangulations since both the parameterization domain and simplified meshes must be embedded in $\mathbb{R}^2$ and $\mathbb{R}^3$ respectively.

3 BACKGROUND

We briefly review the data structure required to support intrinsic triangulations. For a more detailed discussion of intrinsic triangulations we refer the reader to [Sharp et al. 2019].

Starting from a mesh $M = \{V, E, F\}$ encoded as a Δ-complex, an intrinsic triangulation requires the lengths ℓ_{ij} of the ij-th edge in E between vertices i and $j \in V$. Consistent with previous approaches, we represent the intrinsic triangulation as a Δ-complex since it supports phenomena such as self-edges (edges that connect a vertex to itself) and degree one vertices. The lengths ℓ_{ij} describe the shape of the triangles. It can be easily seen that the edge lengths fully define the intrinsic geometry. Other relevant intrinsic information

can be directly computed from the edge lengths:

$$\theta^i_{jk} = \arccos\left(\frac{\ell^2_{ij} + \ell^2_{ik} - \ell^2_{jk}}{2\ell_{ij}\ell_{ik}}\right), \tag{1}$$

$$A_{ijk} = \sqrt{s(s - \ell_{ij})(s - \ell_{jk})(s - \ell_{ki})}, \tag{2}$$

$$s = (\ell_{ij} + \ell_{jk} + \ell_{ki})/2 \tag{3}$$

where θ^i_{jk} is the interior angle at $i \in \mathsf{V}$ in the triangle $ijk \in \mathsf{F}$ and A_{ijk} is the area of the triangle $ijk \in \mathsf{F}$.

4 INTRINSIC SIMPLIFICATION BY EDGE FLIPPING

Our goal is to remove intrinsic vertices that are part of a (near) developable patch in the corresponding extrinsic mesh. Removing such vertices will not alter the intrinsic geometry since a developable patch is isometric to a planar neighborhood. A surface is developable around a vertex i if it has zero Gaussian curvature: $\kappa_i = 2\pi - \alpha_i$, where $\alpha_i = \sum_{ijk} \theta^i_{jk}$ is the cone angle. Ideally, only vertices with zero Gaussian curvature should be removed such that the intrinsic geometry is not changed. However, developable surfaces are often highly tessellated for accurate approximation in $\mathbb{R}^3$. Depending on the exact triangulation and/or numerical round-off errors, an exact zero Gaussian curvature might not be reached. Therefore, in practice we try to remove all vertices with an absolute value Gaussian curvature less than some predetermined threshold κ_{max}. Setting κ_{max} to a larger threshold, allows us to obtain an intrinsic approximation where portions of the triangulation are replaced with developable patches.

Although not strictly necessary, we start by performing an intrinsic Delaunay retriangulation to ensure a well behaved mesh. Next, we sort all vertices by the magnitude of their Gaussian curvature in a queue P by smallest Gaussian curvature first, and process the vertices in P until no vertices with a Gaussian curvature less than κ_{max} can be removed. For each vertex, we perform edge flips on all incident edges until the desired valence (three for interior vertices or two for vertices on boundaries) is reached. We avoid removing degree one vertices or those incident to self-edges and avoid flipping edges that would modify mesh topology. We record each edge flip in a FIFO queue Q for additional post-processing detailed below. For vertices with zero Gaussian curvature, we are guaranteed to reach the desired valence [Gillespie et al. 2021a; Sharp and Crane 2020]. However this is not the case for vertices with non-zero curvature. If the Gaussian curvature is extremely negative the valence can not be achieved. In practice models typically have few if any vertices with large negative curvature and we found that we can usually achieve the desired valance. When the desired valence is reached, simplification can be easily achieved by removing the vertex i and all incident edges ij, ik and il, such that the resulting triangle jkl forms a developable approximation (Figure 2). However, depending on the configuration of the 1-ring, removing the vertex can lead to an invalid triangulation. We therefore only remove vertices if jkl strictly adheres to the triangle inequality. When we are unable to reach the desired valence or if the triangle inequality condition is not met, we undo the edge flips recorded in Q in reverse order to restore the triangulation. It is possible that after removing more
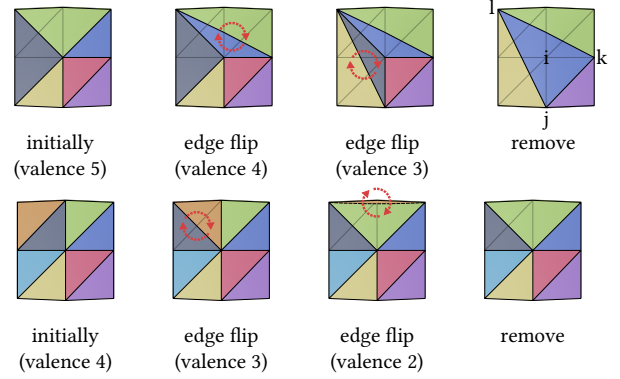

Fig. 2. Illustration of intrinsic vertex removal by edge flipping for an interior (top) and a boundary (bottom) vertex.

vertices, the triangulation is more favorable for removing the vertex. Therefore, we re-queue the vertex for re-processing after all outstanding vertices with a Gaussian curvature less than κ_{max} have been processed.

If a vertex i can be successfully removed (i.e., it has the desired valence after edge flipping, and it satisfies the triangle inequality condition), then we perform the following steps:

- We update the Gaussian curvature of the vertices of jkl in the processing queue P (either by updating the order if j, k, or l was in the queue, or by adding the vertex if not yet queued).
- We perform a post-removal edge flipping procedure. During edge flipping to achieve the desired valence, it is sometimes possible to create degenerate triangles. Instead of trying to figure out a safe flipping order, we instead 'repair' the triangulation after vertex removal. For each edge recorded in Q, we check (in reverse order) if the resulting edge is Delaunay. If not, then we flip the edge, and add the four edges of the resulting triangles to Q. We repeat this process until Q is empty and the resulting triangulation meets the intrinsic Delaunay property again.

Discussion. A benefit of our simplification algorithm is that it preserves the Euler characteristic $\chi = V - E + F$ of the mesh. For a closed mesh (without boundary), each vertex removal step results in a deletion of 1 vertex, 3 edges, and 3 faces, while adding 1 new face ($\Delta\chi = -1 + 3 - 3 + 1 = 0$). In case there is a boundary, we remove 1 vertex, 2 edges and 1 face ($\Delta\chi = -1 + 2 - 1 = 0$) and boundary loops are unaffected. By virtue of the Gauss-Bonnet theorem, we know that since the Euler characteristic is preserved, the sum of the Gaussian curvature is unchanged. Thus deleting a vertex i implies that its Gaussian curvature is redistributed to its neighbors. This also justifies why after each vertex removal, we update the Gaussian curvature sorted processing queue P.

The above edge flipping strategy reduces the 1-ring polygon around each candidate vertex for removal to the trivial re-triangulation case; i.e., such that no re-triangulation is needed. We can apply this
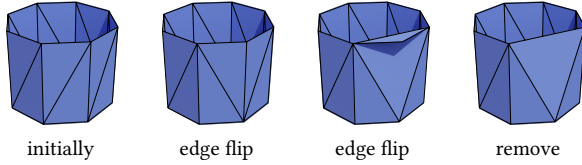
Fig. 3. Applying edge flipping for vertex removal in the extrinsic setting can significantly alter the metric.

strategy only on the intrinsic triangulation, not in an extrinsic setting where it can significantly alter the metric. The metric can be altered when an edge flip is applied to the extrinsic mesh (Figure 3) whereas applying the same operation to an intrinsic edge preserves the metric. Only the vertex removal step alters the surface metric, occurring only if the vertex has nonzero Gaussian curvature.

While vertex removal typically requires retriangulation in the extrinsic setting, it is unclear in the intrinsic domain how to update the edge lengths when the Gaussian curvature is not zero. First, the 1-ring polygon needs to be flattened into to a developable surface, altering the global intrinsic geometry. Second, it is unclear what the updated edge lengths should be such that the candidate vertex is projected onto the developable approximation. By reducing the valence to three we do not need to update any edge lengths in the intrinsic triangulations yielding a computationally light strategy that avoids both retriangulation and computing new edge lengths via optimization. While our approach avoids updating edge lengths during vertex removal it requires us to carefully consider defining the barycentric coordinates of the removed vertex with respect to the remaining triangle should they be required by downstream applications.

Intrinsic-Extrinsic Correspondence. Similar to the Signpost data structure[Sharp et al. 2019], we keep track of the correspondence between intrinsic and extrinsic triangulations. Special care needs to be taken for tracking the correspondence of deleted vertices. This is achieved by maintaining *intrinsic* barycentric coordinates for each extrinsic vertex deleted from the intrinsic triangulation. There are three scenarios to consider:

(1) *Defining the barycentric coordinate* (c_i^j, c_i^k, c_i^l) *of a deleted vertex* i *inside a triangle* jkl. When the threshold κ_{max} is not zero, the vertex is 'projected' onto a developable approximation. Since this projection is an approximation, a number of viable options exist. In our implementation, we opt for a conformal projection, which are known to have low distortion [Gillespie et al. 2021b; Springborn et al. 2008]. We fix the edge lengths of the bounding triangle jk, kl, lj and seek a uniform scale of the edges ij, ik, and il incident to i that ensures the projection of i is intrinsically flat (as in [Springborn et al. 2008] with fixed boundary). Since i is degree three we can compute the uniform scale by ensuring that the corner angles of the projections of ikl, ilj, and ijk agree with the corner angles of jkl. For example, for corner j of jkl we require that the scale s satisfies:

$$\theta_{kl}^j = \theta_{ki}^j + \theta_{il}^j$$

where $\theta_{ki}^j = \arccos\left(\frac{l_{jk}^2 + s^2(l_{ij}^2 - l_{ki}^2)}{2l_{jk}s\ell_{ij}}\right)$ (similarly for θ_{il}^j). In fact if the formula holds for any corner of jkl, say j then the projection of edge ji must lie in the (unfolded) plane of ijk and similarly for i and the projections of edge ki and kl, ensuring that the new Gaussian curvature of i is zero. Solving for s in the corner angle equation yields a quadratic in s^2. We then choose the solution for s^2 (and hence s) that ensures the Gaussian curvature at the projection of i is zero and set the barycentric coordinate (c_i^j, c_i^k, c_i^l) accordingly. If s would cause the projections of any of ikj, ijl, and ilk to violate the triangle inequality and i (or any of its dependent vertices) is required (for downstream applications) we do not remove i and instead restore the intrinsic mesh by undoing the edge flips recorded in Q in reverse order and reschedule it for deletion.

(2) *Updating the barycentric coordinates of a vertex* v *that depends on a deleted vertex* i. When a vertex is deleted, it is possible that a previously deleted vertex's barycentric coordinates depends on it. Because our vertex deletion is only performed after the vertex has the desired valance, it follows that the dependent vertex must lie in one of the faces that will be merged. Hence, both the barycentric coordinates of the deleted and dependent vertex will depend on the same corner vertices after deletion. Thus, we can easily substitute the barycentric coordinates of the deleted vertex. For example, if $v \in ijk$, and i is deleted, then the updated coordinates are: $(c_v^j + c_v^i c_i^j, c_v^k + c_v^i c_i^k, c_v^i c_i^l)$.

(3) *Updating the barycentric coordinate if the vertex is contained in a triangle whose edge is flipped.* In this case, we unfold both triangles isometrically into a local plane and recompute the barycentric coordinates with respect to the new face containing the vertex.

5 RESULTS AND EVALUATION

We have implemented our intrinsic simplification method in C++ using a half-edge data structure annotated with edge lengths. We have validated our method on a subset of the *Thingi10k* [Zhou and Jacobson 2016] dataset containing valid manifold and oriented triangle meshes. Processing all $7k$ triangle meshes takes approximately 25 hours to run the whole algorithm for the 7 threshold values for κ_{max}: $10^{-9}, 10^{-6}, 10^{-4}, 10^{-2}, 10^{-1}, 1.0, \pi$ (in radians) on a Intel i5-8265U (1.60GHz) CPU with 16GB of memory (using a single core).

Figure 4 shows a selection of 3 meshes from the test set. For each mesh we show the original mesh, and 3 thresholds $\kappa_{max} = \{10^{-4}, 10^{-2}, 1.0\}$. For each visualization the simplified mesh is projected on the original extrinsic mesh, and thus shows an approximation of the intrinsic geometry of the simplified mesh. Typically, a common subdivision is utilized to provide a visualization of an intrinsic triangulation [Gillespie et al. 2021a; Sharp et al. 2019] by making use of the fact that both triangulations have the same intrinsic geometry. When vertices with non-zero Gaussian curvature are deleted the correspondence between the two triangulations breaks down and these methods fail. To visualize a simplified mesh we first perform a refinement of the original extrinsic mesh. We then
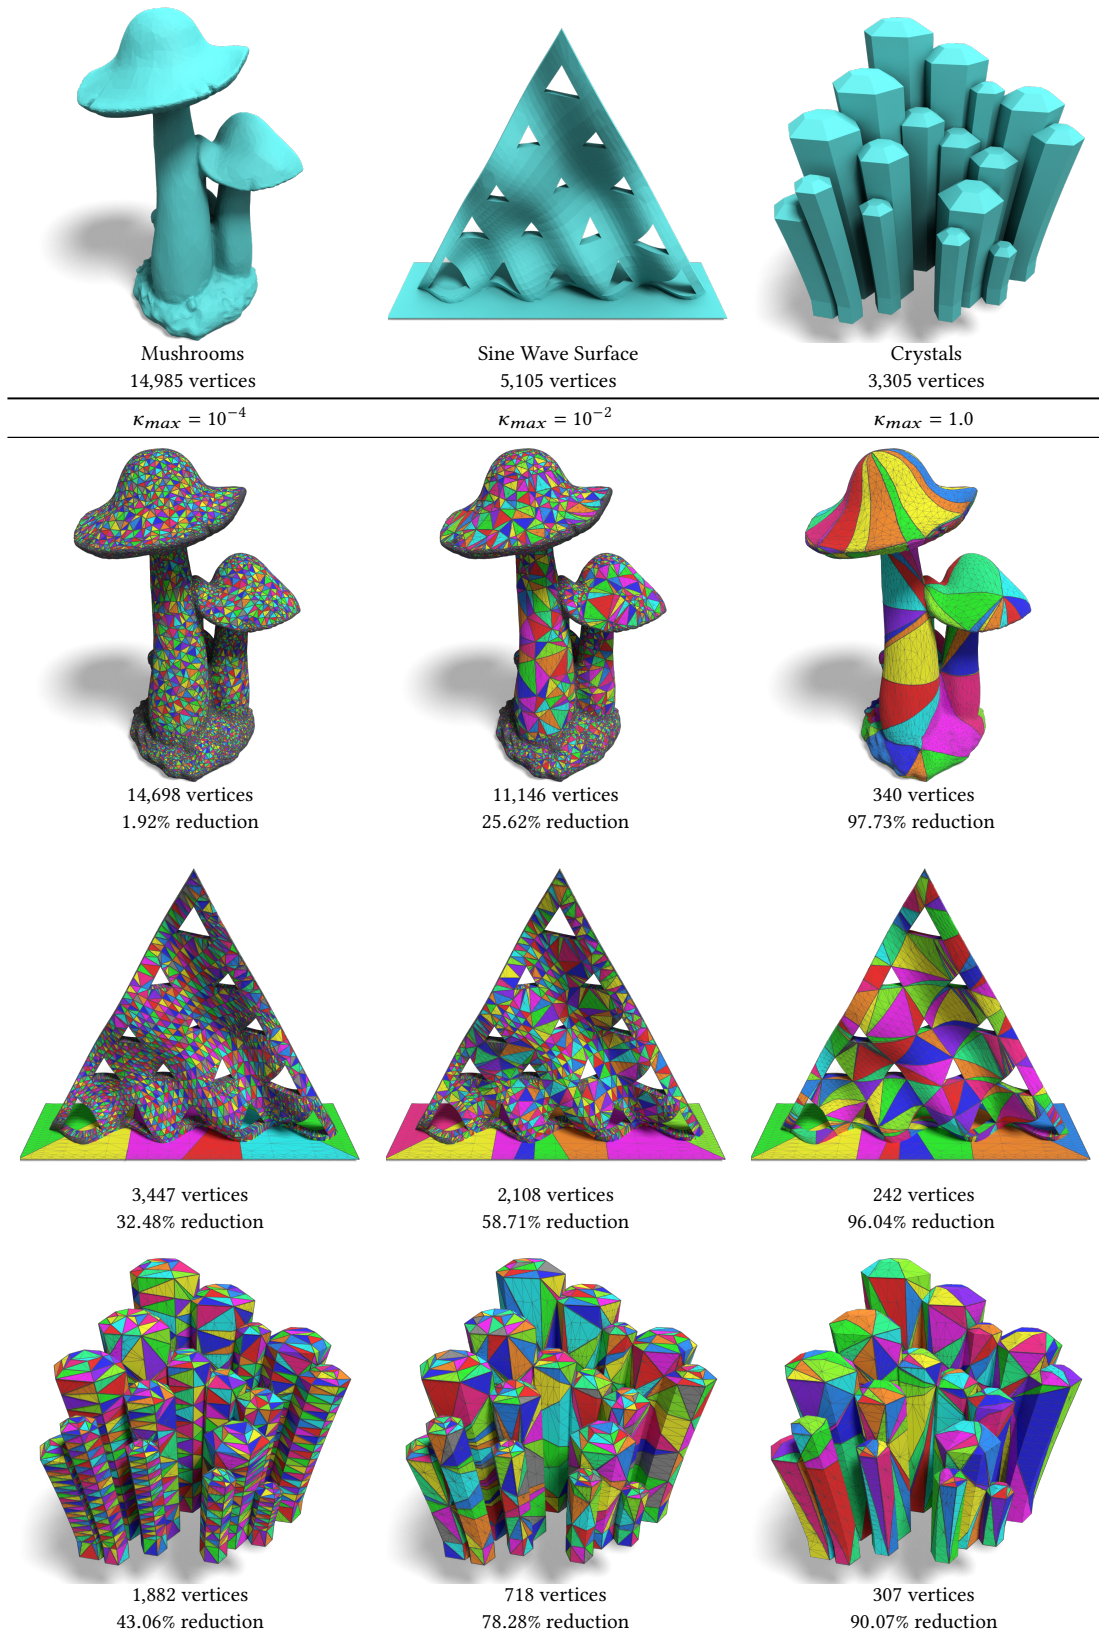
Mushrooms	Sine Wave Surface	Crystals
14,985 vertices	5,105 vertices	3,305 vertices
$\kappa_{max} = 10^{-4}$	$\kappa_{max} = 10^{-2}$	$\kappa_{max} = 1.0$

14,698 vertices	11,146 vertices	340 vertices
1.92% reduction	25.62% reduction	97.73% reduction
3,447 vertices	2,108 vertices	242 vertices
32.48% reduction	58.71% reduction	96.04% reduction
1,882 vertices	718 vertices	307 vertices
43.06% reduction	78.28% reduction	90.07% reduction

Fig. 4. Intrinsic mesh simplification results on three meshes for different thresholds κ_{max}. Note that the simplified intrinsic mesh is projected on the original mesh, showing only an approximation of the intrinsic triangulation's geometry.

Table 1. Mesh simplification statistics over a subset of over 7,000 manifold and oriented triangle meshes in the Thingi10K dataset comparing the percentage of vertices with a Gaussian curvature less then κ_{max} versus the percentage removed after simplification.

κ_{max}	Removable		Successfully Removed		Mean Time (s)		
	mean	std. dev.	mean	std. dev.	remove	track	total
10^{-9}	5.57%	13.84	99.56%	4.73	0.15s	0.61s	0.76s
10^{-6}	7.48%	16.20	99.37%	5.15	0.16s	0.87s	1.03s
10^{-4}	12.67%	20.23	95.58%	8.14	0.20s	0.84s	1.03s
10^{-2}	33.96%	33.53	91.41%	10.12	0.33s	1.24s	1.57s
10^{-1}	58.44%	37.81	89.18%	12.62	0.41s	1.79s	2.20s
1.00	88.75%	22.34	94.87%	7.87	0.45s	2.39s	2.84s
π	99.81%	2.21	94.57%	9.17	0.46s	2.70s	3.15s

track the elements of the refined mesh during simplification. After simplification we color the elements of the refined mesh according to the intrinsic face they map to post-simplification. By tracking the elements of a heavily refined mesh we are able to produce a visualization that approximates the intrinsic triangulation. A more robust method of visualizing a simplified intrinsic mesh, perhaps by producing a (pseudo) common subdivision is an interesting avenue for future research. For each example in Figure 4 we list κ_{max} and the number of vertices in the simplified mesh. Note how our method simplifies most in regions that are (approximately) developable. For example, in the *Sine Wave Surface* example, we can see that significant simplification takes place on the base of the model and similarly for the sharp edges of the *Crystals* model. Note how some of the simplified intrinsic triangles form a curved surface in $\mathbb{R}^3$. For the lowest threshold $\kappa_{max} = 10^{-4}$ there is virtually no simplification in the doubly-curved regions of *Sine Wave Surface* and *Mushrooms* (along the cap). In general there is little simplification in *Mushrooms* for $\kappa_{max} = 10^{-4}$, even in the stalk region, due to noisy vertex placement in the model. However, as the threshold increases, our algorithm is able to achieve a large reduction in the noisy model. As expected, for each example, a higher threshold removes much move vertices; particularly at the highest threshold $\kappa_{max} = 10^{1.0}$ where our algorithm achieves a reduction of 97.73%, 96.04%, and 90.07% for *Mushrooms*, *Sine Wave Surface*, and *Crystals* respectively. Notice that the doubly curved regions of *Sine Wave Surface* and *Mushrooms* (on the mushroom caps) are approximated by larger developable patches as κmax increases. Figure 5 shows 2 meshes simplified to less than 2% of their original vertex count, approximating each surface with a small number of developable patches.

To gain further insight in the efficacy of our intrinsic simplification method, Table 1 reports the mean percentage and standard deviation of ideally removable vertices (i.e., those with a Gaussian curvature less than κ_{max}) and the mean percentage (and standard deviation) of actually removed vertices for seven different thresholds $\kappa_{max} = \{10^{-9}, 10^{-6}, 10^{-4}, 10^{-2}, 10^{-1}, 1.0, \pi\}$. As expected, with an increasing threshold, more vertices can be removed, and in general vertex removal succeeds for almost all candidate vertices. In addition, we also report the mean running time (excluding data IO) when tracking all removed vertices via intrinsic barycentric coordinates.

Typically 80% of the compute time is spent tracking barycentric coordinates.

To demonstrate the impact of removing non-zero Gaussian curvature vertices, we visualize the solution to a Poisson equation with a spike centered on a chosen vertex on 2 selected meshes for 3 κ_{max} thresholds (Figure 6); we report the number of vertices and MSE error over the Poisson solution for each of the vertices in the original mesh, where the solution at a removed vertex is interpolated based on the intrinsic barycentric coordinates. We deliberately did *not* apply any refinement and instead directly solve the equation on the simplified mesh to better illustrate the impact of the simplification. In the first example, we deliberately placed the source term on a vertex with low Gaussian curvature. As expected, as the threshold κ_{max} increases, so does the error. In the second example, we placed the source term on a vertex with high Gaussian curvature, which results in a smaller increase in error because most intrinsic triangles around the source term will not be simplified. In a practice, for maximal accuracy, we would first insert vertices via an intrinsic optimal Delaunay algorithm [Sharp et al. 2019] or adaptive intrinsic mesh refinement [Sharp et al. 2019]. While both methods will increase the number of intrinsic vertices, neither would be encumbered by sub-optimally placed intrinsic vertices from the original mesh. Indeed, in the concurrent work of Liu et al. [2023] they perform refinement prior to simplification to achieve a better vertex distribution post-simplification.

A key advantage of simplifying the intrinsic representation compared to first simplifying an extrinsic mesh and then constructing an intrinsic representation is that we can produce simplified triangles that ride the underlying surface, and thus better maintain the surface metric. To better understand the differences between different simplification methods, Figure 7 compares intrinsic mesh simplification with two extrinsic simplifications methods: QEM [Garland and Heckbert 1997] and Spectral Mesh Simplification [Lescoat et al. 2020]. For each method, we perform an equal-vertex-count comparison (simplified from 14,290 to 1,715 vertices for the *Bunny*, and from 23,356 to 4,440 vertices for the *Frog*) to our intrinsically simplified mesh, and highlight the differences between all methods by visualizing the solution to a Poisson equation computed directly on the "raw" mesh with a source term placed at a low and high Gaussian curvature vertex. PDE solutions are solved on the intrinsic mesh and the barycentric coordinates are used to linearly interpolate the
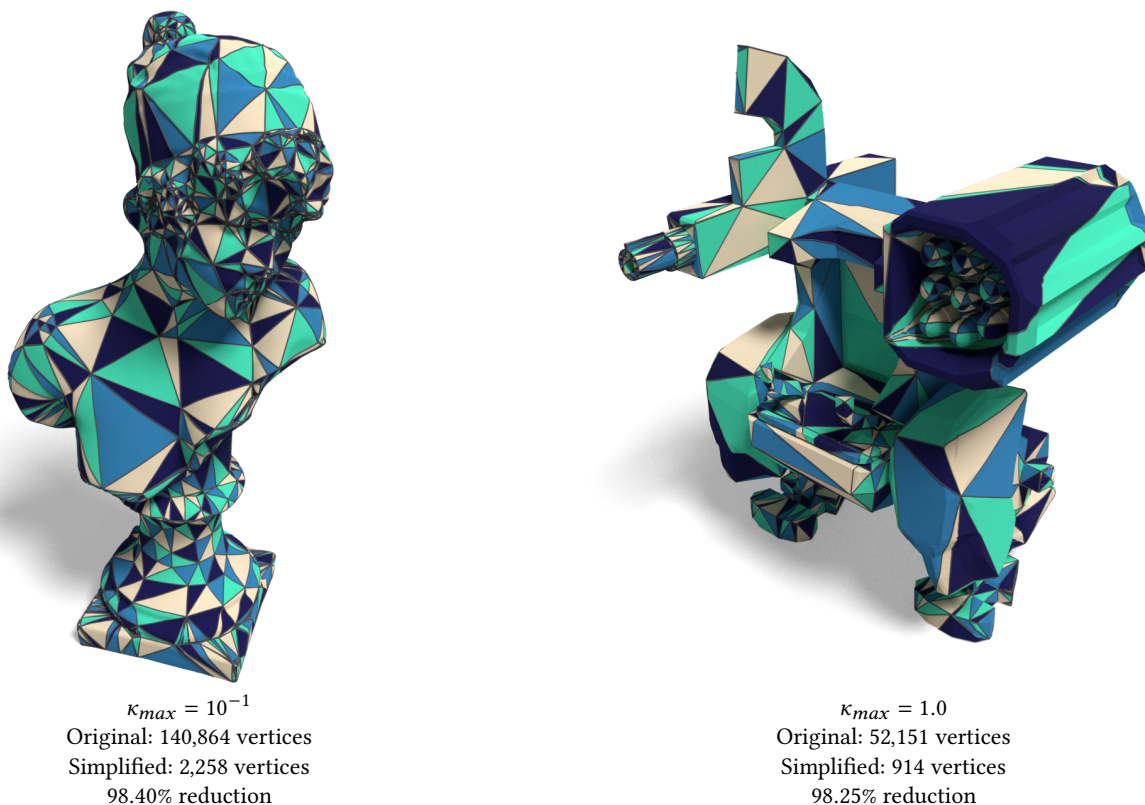
$\kappa_{max} = 10^{-1}$
Original: 140,864 vertices
Simplified: 2,258 vertices
98.40% reduction

$\kappa_{max} = 1.0$
Original: 52,151 vertices
Simplified: 914 vertices
98.25% reduction

Fig. 5. Our algorithm applied to various meshes from *Thingi10k* with large κ_{max}, achieving reduction in excess of 98%.

solution at deleted vertices. For QEM we use the implementation provided by Lescoat et al. [2020]. Note that Spectral Mesh Simplification optimizes for the intrinsic qualities of the mesh when removing vertices, and as such on average the solution to the Poisson equation is more accurate, albeit at a much higher computation cost (35 minutes versus 0.62 seconds for our method on the *Bunny*; for reference QEM took 7.6 second seconds versus 0.85 seconds for our method on the *Frog*). However, as noted before, intrinsic mesh simplification is intended as a preprocessing step, and in practice we would apply an adaptive refinement or produce an optimal Delaunay triangulation before solving the Poisson equation.

6 LIMITATIONS AND FUTURE WORK

Our method is not without limitation. First, as mentioned in Section 4, a potential limitation is imposed by reducing the valence of a vertex prior to removal. The issue arises since a vertex i with Gaussian curvature κ_i has a lower bound on its valence: $\lceil 2 - \frac{\kappa}{\pi} \rceil$. The lower bound exists since the corner angles $\theta_1, ..., \theta_n$ around i must each satisfy $\theta_j < \pi$ since each θ_j belongs to a triangle. Therefore the total angle around i must satisfy $\alpha_i = \sum_{j=1}^{n} < n\pi$; and since $\alpha_i = 2\pi - \kappa_i$, the valence n must satisfy $n > \lceil 2 - \frac{\kappa_i}{\pi} \rceil$. The lower bound means that our method is unable to remove vertices with $\kappa \leq -\pi$. However, in practice, a typical mesh only has a few, if any, vertices with such extreme negative curvature. In fact, on average,

models in our validation set (of roughly 7,000 models) fewer than 0.2% of vertices violate $\kappa > -\pi$. In practice the lower bound on the valence of a vertex does not affect the practical effectiveness of our method for the majority of real-world meshes.

Second, maintaining the correspondence between the intrinsic and extrinsic triangulations with intrinsic barycentric coordinates forces our algorithm to update barycentric coordinates of vertices in faces adjacent to a flipped edge and in one of the three faces incident to a vertex that is about to be removed. These repeated updates can cause numerical error to compound as the algorithm progresses. Furthermore, for large models that are highly decimated, e.g. highly tessellated cylinder, a large number of removed vertices must have their intrinsic coordinates updated during each edge flip and vertex removal. Exploring alternative methods of maintaining the correspondence between the two triangulations could provide an interesting avenue for future research and open the door for faster and more robust intrinsic mesh simplification.

Finally, our method of visualizing the simplified intrinsic triangulation does not accurately represent the new intrinsic geometry. A method that provides the benefits of a common subdivision can enable a better visualization as well as provide a more robust way to port data between the two triangulations, such as enabling a robust change of basis when solving PDEs. Producing a 'pseudo'-common subdivision is an interesting avenue for future research. Alternatively one could produce a visualization by computing an

embedding of the simplified intrinsic mesh, which could aide in visualizing the deformation of the intrinsic geometry of the original mesh.

7 CONCLUSION

In this paper we presented a method for intrinsic mesh simplification. Our method leverages the benefits of intrinsic edge flipping and the large space of intrinsic triangulations to significantly simplify vertex deletion to two canonical cases (i.e., valence three for an internal vertex, and valence two for a boundary vertex). We use Gaussian curvature as the deletion criterion, which, together with our removal strategy, effectively projects vertices onto a locally developable approximation. We demonstrated the robustness and effectiveness of our method on the Thingi10k dataset and demonstrated the impact of relaxing the deletion threshold on the solutions of PDEs.

8 ACKNOWLEDGEMENTS

REFERENCES

David Bommes, Marcel Campen, Hans-Christian Ebke, Pierre Alliez, and Leif Kobbelt. 2013. Integer-Grid Maps for Reliable Quad Meshing. *ACM Trans. Graph.* 32, 4, Article 98 (jul 2013), 12 pages. https://doi.org/10.1145/2461912.2462014

M. Botsch, L. Kobbelt, M. Pauly, P. Alliez, and B. Levy. 2010. *Polygon Mesh Processing.*

Long Chen and Jinchao Xu. 2004. Optimal Delaunay triangulations. *Journal of Computational Mathematics* 22, 2 (1 March 2004), 299–308.

David Cohen-Steiner, Pierre Alliez, and Mathieu Desbrun. 2004. Variational Shape Approximation. *ACM Trans. Graph.* 23, 3 (aug 2004), 905–914.

Hans-Christian Ebke, Patrick Schmidt, Marcel Campen, and Leif Kobbelt. 2016. Interactively Controlled Quad Remeshing of High Resolution 3D Models. *ACM Trans. Graph.* 35, 6, Article 218 (dec 2016), 13 pages. https://doi.org/10.1145/2980179.2982413

Matthew Fisher, Boris Springborn, Peter Schröder, and Alexander I Bobenko. 2007. An algorithm for the construction of intrinsic Delaunay triangulations with applications to digital geometry processing. *Computing* 81, 2 (2007), 199–213.

Michael S Floater and Kai Hormann. 2005. Surface parameterization: a tutorial and survey. *Advances in multiresolution for geometric modelling* (2005), 157–186.

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).* 209–216.

Mark Gillespie, Nicholas Sharp, and Keenan Crane. 2021a. Integer coordinates for intrinsic geometry processing. *ACM Trans. Graph.* 40, 6 (2021).

Mark Gillespie, Boris Springborn, and Keenan Crane. 2021b. Discrete Conformal Equivalence of Polyhedral Surfaces. *ACM Trans. Graph.* 40, 4 (2021).

Kai Hormann, Bruno Lévy, and Alla Sheffer. 2007. Mesh Parameterization: Theory and Practice. In *ACM SIGGRAPH 2007 Courses* (San Diego, California) *(SIGGRAPH '07).* Association for Computing Machinery, New York, NY, USA, 1–es. https://doi.org/10.1145/1281500.1281510

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 (jul 2018).

Dawar Khan, Alexander Plopski, Yuichiro Fujimoto, Masayuki Kanbara, Gul Jabeen, Yongjie Jessica Zhang, Xiaopeng Zhang, and Hirokazu Kato. 2020. Surface remeshing: A systematic literature review of methods and research directions. *IEEE TVCG* 28, 3 (2020), 1680–1713.

Andrei Khodakovsky, Nathan Litke, and Peter Schröder. 2003. Globally Smooth Parameterizations with Low Distortion. *ACM Trans. Graph.* 22, 3 (jul 2003), 350–357. https://doi.org/10.1145/882262.882275

Aaron WF Lee, Wim Sweldens, Peter Schröder, Lawrence Cowsar, and David Dobkin. 1998. MAPS: Multiresolution adaptive parameterization of surfaces. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques.* 95–104.

Thibault Lescoat, Hsueh-Ti Derek Liu, Jean-Marc Thiery, Alec Jacobson, Tamy Boubekeur, and Maks Ovsjanikov. 2020. Spectral mesh simplification. In *Comp. Graph. Forum*, Vol. 39. 315–324.

Derek Liu, Mark Gillespie, Benjamin Chislett, Nicholas Sharp, Alec Jacobson, and Keenan Crane. 2023. Surface Simplification using Intrinsic Error Metrics. *ACM Trans. Graph.* XX, X (2023).

Jovan Popović and Hugues Hoppe. 1997. Progressive Simplicial Complexes. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '97).* 217–224.

Yang Qi, Dario Seyb, Benedikt Bitterli, and Wojciech Jarosz. 2022. A bidirectional formulation for Walk on Spheres. *Comp. Graph. Forum* 41, 4 (2022), 51–62.

Jarek Rossignac and Paul Borrel. 1993. Multi-resolution 3D approximations for rendering complex scenes. In *Modeling in computer graphics.* Springer, 455–465.

Rohan Sawhney and Keenan Crane. 2020. Monte Carlo Geometry Processing: A Grid-Free Approach to PDE-Based Methods on Volumetric Domains. *ACM Trans. Graph.* 39, 4, Article 123 (aug 2020).

William J Schroeder, Jonathan A Zarge, and William E Lorensen. 1992. Decimation of triangle meshes. In *Proceedings of the 19th annual conference on Computer graphics and interactive techniques.* 65–70.

Silvia Sellén, Herng Yi Cheng, Yuming Ma, Mitchell Dembowski, and Alec Jacobson. 2019. Solid Geometry Processing on Deconstructed Domains. *Comp. Graph. Forum* 38, 1 (2019), 564–579.

Nicholas Sharp and Keenan Crane. 2020. You Can Find Geodesic Paths in Triangle Meshes by Just Flipping Edges. *ACM Trans. Graph.* 39, 6, Article 249 (nov 2020).

Nicholas Sharp, Yousuf Soliman, and Keenan Crane. 2019. Navigating intrinsic triangulations. *ACM Trans. Graph.* 38, 4 (2019).

Alla Sheffer, Emil Praun, and Kenneth Rose. 2006. Mesh Parameterization Methods and Their Applications. *Found. Trends. Comput. Graph. Vis.* 2, 2 (jan 2006), 105–171. https://doi.org/10.1561/0600000011

Boris Springborn, Peter Schröder, and Ulrich Pinkall. 2008. Conformal Equivalence of Triangle Meshes. *ACM Trans. Graph.* 27, 3 (aug 2008), 1–11. https://doi.org/10.1145/1360612.1360676

Kenshi Takayama. 2022. Compatible Intrinsic Triangulations. *ACM Trans. Graph.* 41, 4, Article 57 (jul 2022).

Qingnan Zhou, Eitan Grinspun, Denis Zorin, and Alec Jacobson. 2016. Mesh Arrangements for Solid Geometry. *ACM Trans. Graph.* 35, 4, Article 39 (jul 2016).

Qingnan Zhou and Alec Jacobson. 2016. Thingi10K: A Dataset of 10,000 3D-Printing Models. *arXiv preprint arXiv:1605.04797* (2016).
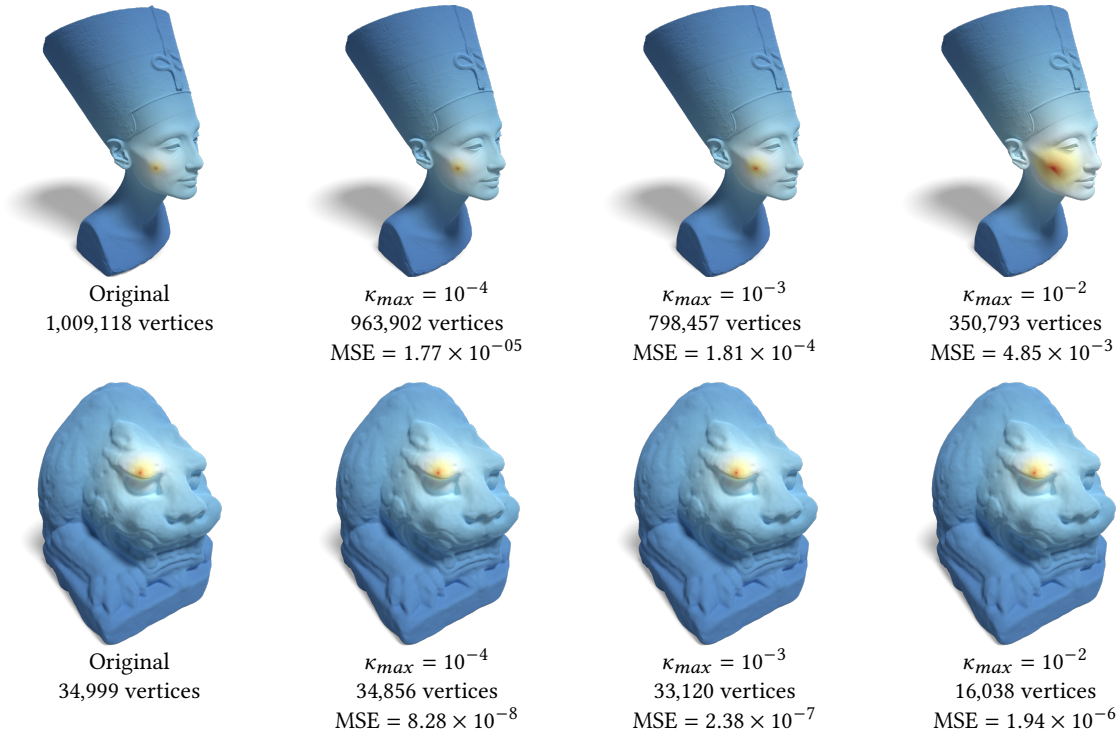
Fig. 6. Visualizations of the solution of a Poisson equation with a spike placed on a near-developable vertex (top) or on a vertex with high curvature (bottom). The equation is solved directly on the "raw" simplified mesh to better show the impact of simplification.
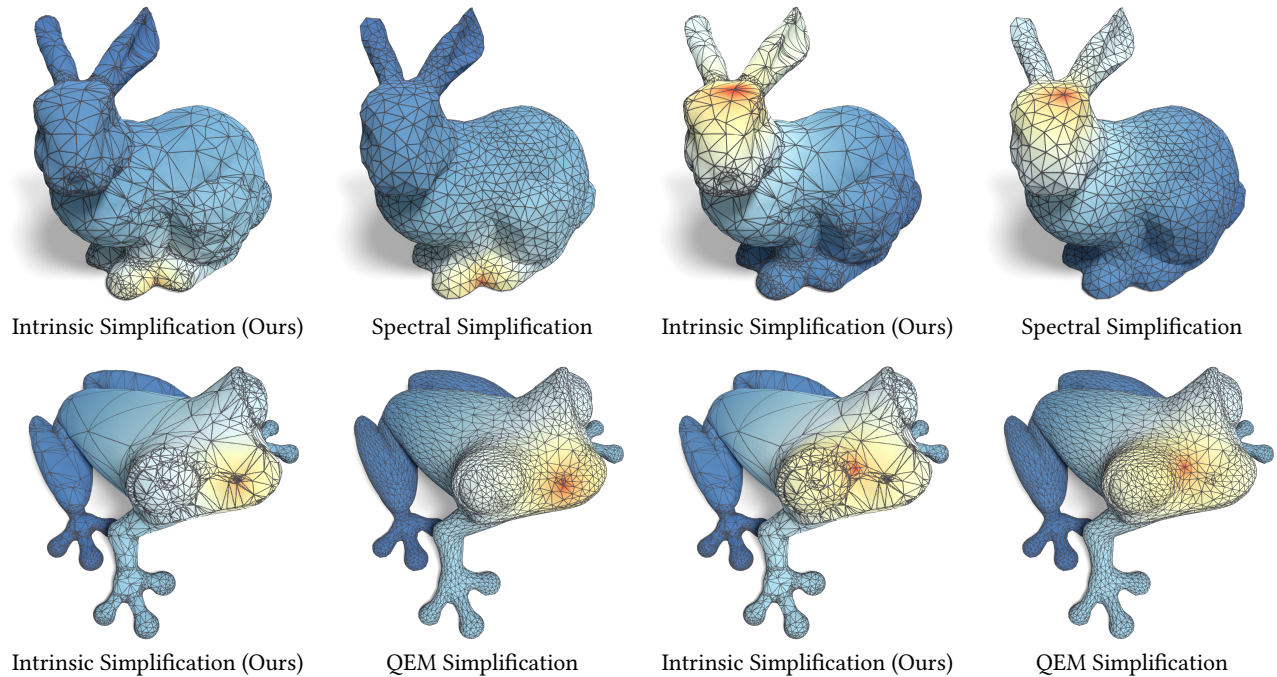

Fig. 7. Solutions of a Poisson equations computed on simplified meshes obtained with Spectral Mesh Simplification (SMS) and QEM. The *Bunny* was reduced to 1,715 vertices (88.00% reduction) via our intrinsic simplification algorithm and QEM. The *Frog* was reduced to 4,440 vertices (80.99% reduction) via our intrinsic simplification and SMS. All solutions are computed on the "raw" simplified mesh to better show the impact of simplification.