

Proofs for Appendix A

Proof of Theorem 11:

PROOF. \Rightarrow : By contradiction. Assume that there exist two boundary edges of P , e_0 and e_1 , that project onto each other. Without loss of generality, we assume that e_0 is encountered prior to e_1 in a counter-clockwise boundary walk (see Def. 3). Consider the subset of quadrilaterals in \bar{P} that can be completely separated from \bar{P} by a straight path from the first ending vertex of e_0 to the second ending vertex of e_1 . We denote this subset as Q . Since P is regular, Q is a subset of a parallelogram. Since e_0 and e_1 are opposite to each other, the walking direction of Q at e_1 is the walking direction at e_0 plus two, implying that there exists at least a pair of consecutive convex corners in between. This pair of consecutive convex corners in Q is equivalent to a pair of consecutive concave corners in P , which contradicts the assumption of no consecutive concave corners.

\Leftarrow : Contrapositively, if there is a pair of consecutive concave corners, v_1 and v_k , we show that there exist two boundary edges of P that project one to another. Consider the two edges (v_0, v_1) and (v_k, v_{k+1}) before and after these two concave corners on the boundary of P , respectively. There exists a parallelogram Q , $Q \subset \bar{P}$, with four corners as v_0, v_1, v_k , and v_{k+1} in clockwise order. Consequently, (v_0, v_1) projects onto (v_k, v_{k+1}) . \square

To prove Theorem 14, we first need the following lemma:

LEMMA 18. *There exists at least one simple cave adjacent to a set of consecutive concave corners.*

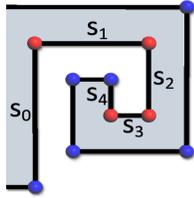


Fig. 2: The five sides adjacent to four consecutive concave corners.

PROOF. We denote sides adjacent to a set of consecutive concave corners as s_i with $0 \leq i \leq C$, where C is the number of consecutive concave corners. s_0 is the side from a convex corner to the first concave corner, and s_C is the last side from the last concave corner to a convex corner. An example is shown in Figure 2. Assume that a simple cave does not exist at both the beginning and the end of the consecutive concave corners; otherwise, the proof is done. This implies that s_0 is longer than s_2 and s_C is longer than s_{C-2} , since the first and the last convex corner cannot emanate a straight path to hit P 's boundary. This clearly cannot hold for $C = 2$. For $C > 2$, recall that we assume that \bar{P} is regular; thus, s_{i+2} cannot be longer than s_i for $1 \leq i$; otherwise, s_{i+2} will penetrate s_{i-1} . \square

Proof of Theorem 14:

PROOF. \Rightarrow : The walking direction along the boundary edges before the first end and after the second end of the closure backtracks by at least 2, implying that there are two consecutive concave corners in between.

\Leftarrow : According to Lemma 18, there is at least a simple cave adjacent to a set of consecutive concave corners. A simple cave is also a cave by definition. \square

Proof of Theorem 16:

PROOF. \Rightarrow : A simple cave is also a cave by definition.

\Leftarrow : By Theorem 14 \Rightarrow we know that there will be consecutive concave corners adjacent to a cave. By Lemma 18 we know that there exists at least one simple cave adjacent to these consecutive concave corners. \square

Proof of Theorem 17:

PROOF. Each simple cave-filling operation subtracts a rectangular patch from \bar{P} and adds it to P . Consequently, any new edge introduced to P is opposite to some existing edge in P and hence \hat{P} remains unchanged. \square

Proofs of Theorem 7 and 8

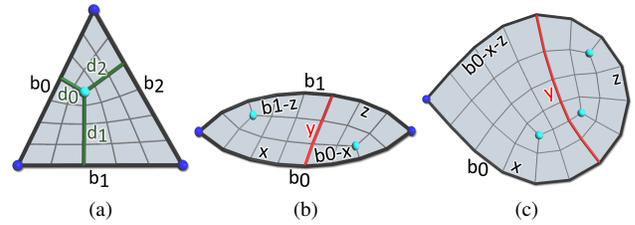


Fig. 3: (a) Notations of a quadrangulated triangle. (b) Notations of a quadrangulated 2-sided patch subdivided into two triangles. (c) Notations of a quadrangulated 1-sided patch subdivided into a triangle and a 2-sided patch.

LEMMA 19. *The number of quads in a triangle with side lengths b_0, b_1 , and b_2 , when quadrangulated with a single v_3 , is $(2*b_1*b_2 + 2*b_2*b_0 + 2*b_0*b_1 - b_0*b_0 - b_1*b_1 - b_2*b_2)/4$.*

PROOF. We denote the (topological) lengths of the three separatrices emanated from the v_3 to the three sides with lengths b_0, b_1 , and b_2 as d_0, d_1 , and d_2 (see Figure 3a). $b_0 = d_1 + d_2$, $b_1 = d_2 + d_0$, and $b_2 = d_0 + d_1$ since the three separatrices decompose the triangle into three parallelograms. The number of quads in the triangle is $d_0 * d_1 + d_1 * d_2 + d_2 * d_0$, which equals $(2*b_1*b_2 + 2*b_2*b_0 + 2*b_0*b_1 - b_0*b_0 - b_1*b_1 - b_2*b_2)/4$. \square

LEMMA 20. *Assuming that a quadrangulated 2-sided patch of side lengths b_0, b_1 is subdivided into two triangles, one with side lengths x, y , and $b_1 - z$, and the other with side lengths $y, b_0 - x$, and z (see Figure 3b). The number of quads is maximized when $x = \lfloor b_0/2 \rfloor$, $y = (b_0 + b_1)/2$, and $z = \lfloor b_1/2 \rfloor$.*

PROOF. The formula for the number of quads in the 2-sided patch can be constructed as the sum of the two triangles (Lemma 19). It is a convex function with x, y , and z as variables. The values at which the function is maximized can then be found by taking partial derivatives of x, y , and z over the function. \square

LEMMA 21. *Assuming that a quadrangulated 1-sided patch of side length b_0 is subdivided into a triangle of lengths x, y , and $b_0 - x - z$, and a 2-sided patch of side lengths y and z (see Figure 3c), the number of quads is maximized when $x = b_0/2 - 1$, $y = b_0 - 2$, and $z = 2$.*

PROOF. The formula for the number of quads in the 1-sided patch can be constructed as the sum of the triangle (Lemma 19) and the 2-sided patch (Lemma 20). It is a convex function with $x,$

y , and z as variables. The values at which the function is maximized can then be found by taking partial derivatives of x , y , and z over the function. \square

It is then straightforward to verify Theorem. 7 by Lemma 20 and Theroem 8 by Lemma 21.

Comparing Parameterizations

In this section, we discuss how the choice of parameterization methods may affect the results available to the users. Although exhaustive enumerations are independent of the parameterization, in practice, the temporal order of topologies being enumerated is determined by geometric ranking heuristics and is thus influenced by parameterization. In Figure 4, we compare the first enumerated topologies of a bumpy circular patch parameterized by LSCM (a conformal method) and a discrete authalic method with a fixed circular domain.

Geometric Optimization

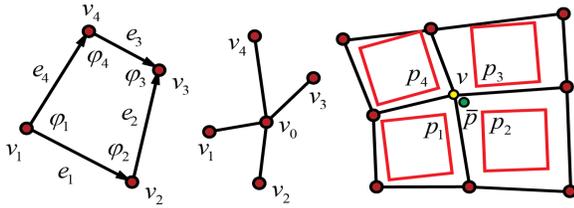


Fig. 5: Variables in the geometric optimization.

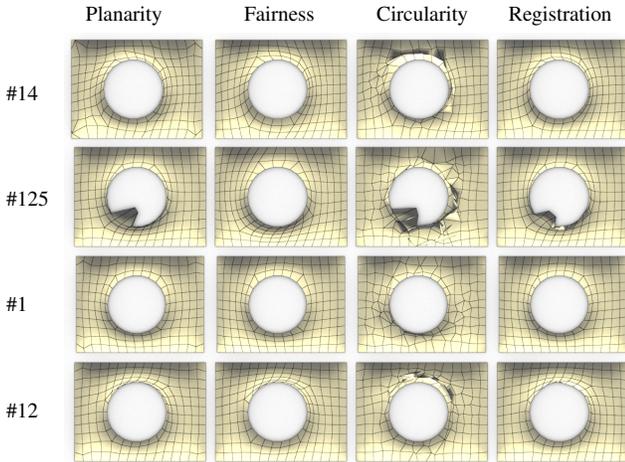


Fig. 6: Geometric optimization vs. topology. Each column corresponds to one specific optimization towards one geometric quality, namely planarity, fairness, circularity and closeness to the reference object. Additionally, in all four tests, the planarity of faces was required with a higher weight ($10\times$) and all the boundary vertices were constrained to be fully fixed. The meshes on the main diagonal are the highest ranked according to a specific optimization. It is interesting to see that a mesh, e.g., #125, can be best for one optimization, but produce errors in another one.

In our selected applications, specifically architecture, planar quad meshes are of great interest due to their manufacturing advantages. We have a configurable optimization framework that can

topology	planarity rank/value	fairness rank/value	circularity rank/value	registration rank/value
Mesh # 14	1 / 0.0004	8 / 0.0110	90 / 0.0713	40 / 0.0078
Mesh #125	26 / 0.0009	1 / 0.0094	16 / 0.0166	25 / 0.0071
Mesh #1	3 / 0.0007	6 / 0.0109	1 / 0.0012	2 / 0.0043
Mesh #12	9 / 0.0008	28 / 0.0122	32 / 0.0303	1 / 0.0038

Table II. : Optimization ranking and objective function values for four selected British museum models. Various geometric optimization objectives were applied on a set of 130 topologies and these were ranked accordingly. The ranking values and the values of each particular objective function are shown.

combine multiple objective functions (described next) by specifying weights in a user interface. Our geometric optimization is an extension of Yang et al. [Liu et al. 2006]. Vertices on the boundary can be constrained to remain fixed (c_{fix}). Another option is to let the boundary vertices glide along the boundary (c_{glide}). The optimization is solved using Levenberg-Marquardt or Gauss-Newton for the examples in the paper.

Planarity: For a quad with vertices v_1, v_2, v_3 , and v_4 , planarity is computed as $f_{planarity} = \Delta_1 + \Delta_2 + \Delta_3 + \Delta_4$, where $\Delta_1 = Det(e_1, e_2, e_3)$, $\Delta_2 = Det(e_1, e_2, e_4)$, $\Delta_3 = Det(e_1, e_3, e_4)$, $\Delta_4 = Det(e_2, e_3, e_4)$. The e_i are unit vectors along the edges of a quad (Figure 5 left).

Circularity: Circularity is computed as

$$f_{circularity} = \max(\text{abs}(\varphi_1 + \varphi_3 - \pi), \text{abs}(\varphi_2 + \varphi_4 - \pi)). \quad (2)$$

Fairness: For a valence 4 vertex (Figure 5 middle), fairness is defined as

$$f_{fairness} = (v_1 + v_3 - 2v_0)^2 + (v_2 + v_4 - 2v_0)^2. \quad (3)$$

Registration Closeness: $f_{registration}$ describes the similarity of the quads to a prescribed *master panel*, e.g., a square with edge length one. First, the master panel is registered to each quad by rigid transformations. Each valence, n , vertex will correspond to n master panel vertices p_i . The registration closeness is then computed as $\|v - \bar{p}\|$, where \bar{p} is the weighted average of p_i (See Figure 5 right).

Smoothness: The smoothness f_{smooth} of a vertex v is $\|\Delta(v)\|$, the magnitude of the cotangent Laplacian, $\Delta(v)$ [Botsch et al. 2010].

Closeness: For a vertex, v , its distance to a reference mesh, or closeness, is

$$f_{closeness} = \|(v - w) \cdot n_w\|, \quad (4)$$

where w is the nearest point to v on the reference mesh and n_w is the unit normal vector at point w .

In Figure 6 we compare how different topologies perform under changing optimization criteria. The errors for the models are given in Table II. Again, we can observe that no single topology is well suited to all geometric optimizations and that exploring topologies can greatly contribute to geometric optimization problems. More examples are shown in Figure 7.

Additional Results and Figures

An example of shape-space exploration with a fixed non-planar boundary is shown in Figure 8. We rank the shapes according to one geometric objective and color code the faces according to another geometric objective. We also visualize the best shape according to six geometric criteria. While one topology turns out to be the best

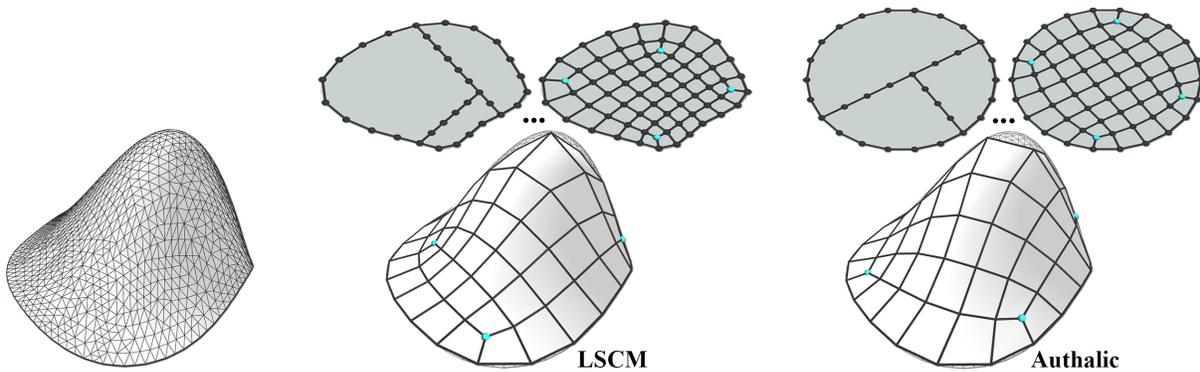


Fig. 4: Left: the input mesh, which is a bumpy surface with an ellipse-shaped boundary. Middle and right: the first enumerated topologies when parameterized by LSCM (middle) and the discrete authalic method (right). The 2D visualizations of the topologies after the first two subdivisions and when fully quadrangulated are shown on top. We observe that the parameter domain of LSCM has less distortion (mainly because it is open boundary) and the resulting first enumerated topology better approximates the input mesh.

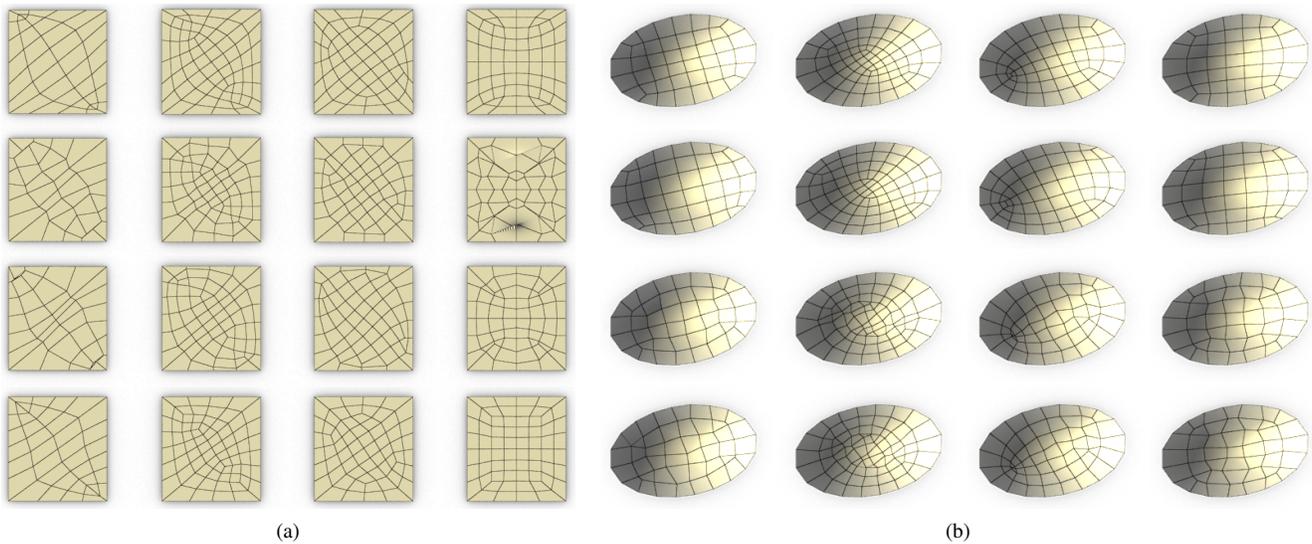


Fig. 7: Two examples of geometric optimization vs. topology. Each row corresponds to one specific optimization towards one geometric quality, namely, from top to bottom, planarity, fairness, circularity and closeness to the reference object. Boundary vertices are fixed. The meshes on the main diagonals are the highest ranked according to the specific optimization

in four categories, the two remaining categories are optimized by two other topologies. The important insight from these examples is that different geometric objectives favor different topologies.

In Figure 11, we explore three alternative low-resolution requadrangulations to approximate an input dome model. Note that the three distinctive quadrangulations all share the same boundary configuration. However, the densities increase with the number of additional v_3 - v_5 pairs, and approximation errors are reduced. This is one example where redundant v_3 - v_5 pairs could be desirable.

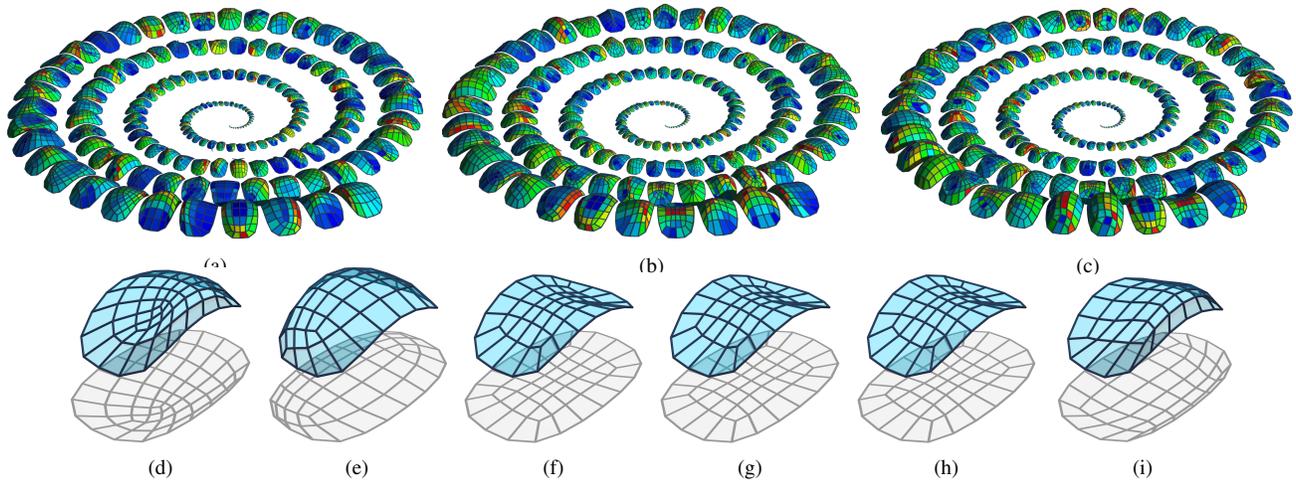


Fig. 8: Top: The explored meshes are ranked and sorted in a spiral decreasing fashion according to one selected geometric criterion. The color coding visualizes the measure of a second selected criterion, i.e., (a) planarity vs. mean curvature, (b) fairness vs. kink angle, (c) circularity vs. kink angle. Bottom: Summary of the best candidates throughout all topologies with respect to (d) planarity, (e) circularity, (f) kink angle, (g) Gaussian curvature, (h) mean curvature, (i) fairness.

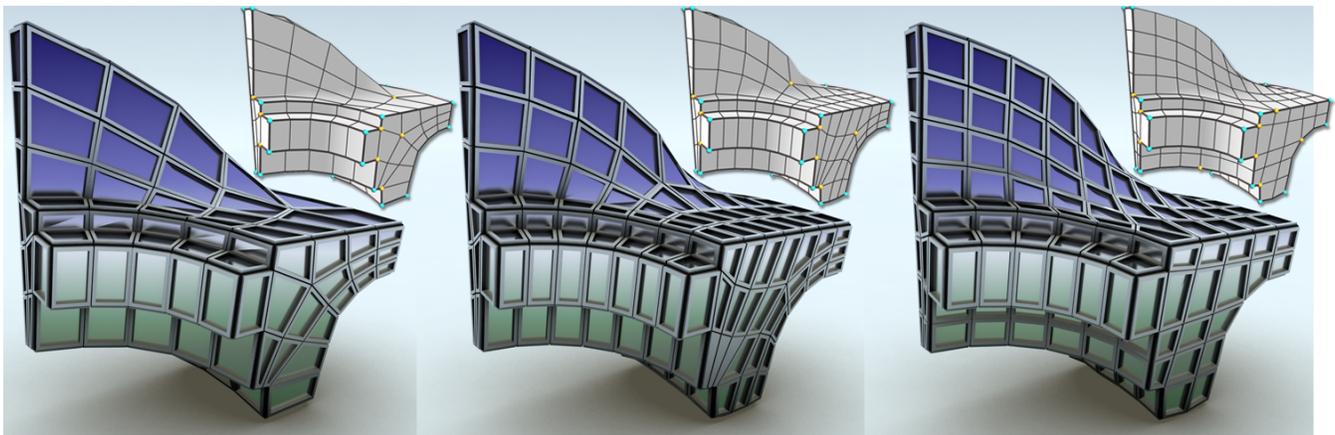


Fig. 9: The Fandisk requadrangulations rendered as glass-and-steel architecture.

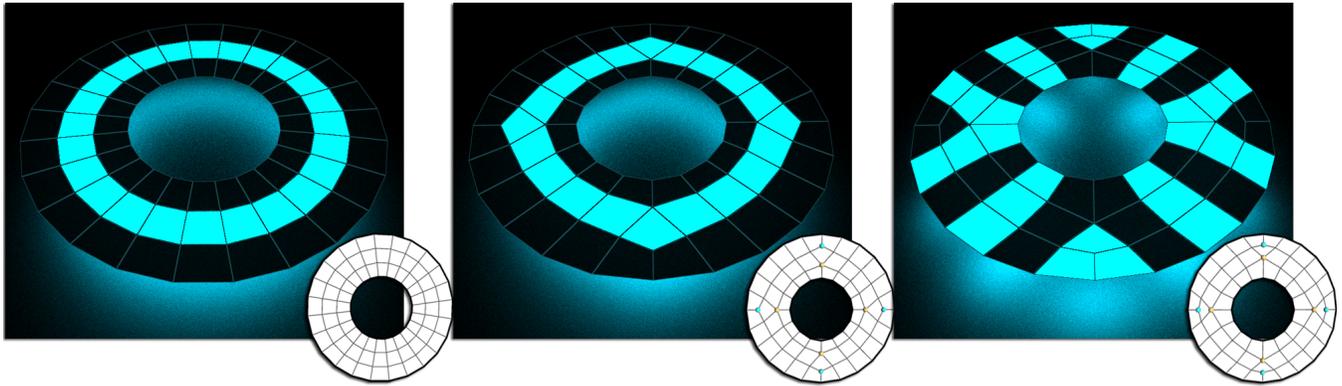


Fig. 10: We explore quadrangulations of a doughnut-shaped patch with a Tron-inspired pattern. Left: The patch is quadrangulated as a regular band. Middle: the patch is quadrangulated with four v_3 - v_5 pairs and geometrically optimized by Laplacian smoothing. Right: the same topology as the middle, but geometrically optimized by techniques described in Section 8. Each quadrangulation is rendered with a suitable pattern.

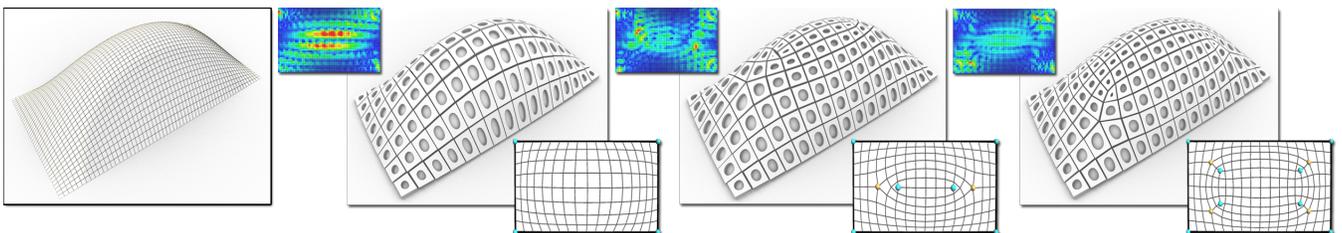


Fig. 11: Exploring alternative low-resolution requadrangulations to approximate a dome model. Left: the input dome model in high resolution. Second left to right: three ways to requadrangulate the same patch (a 16-by-8 parallelogram). The first trivially quadrangulates the parallelogram using a regular grid. However, the density is insufficient and the approximation error is high. The second and the third alternatives use two and four additional v_3 - v_5 pairs to increase the densities and reduce the approximation errors. Approximation errors, measured by closest distance to each face of the input mesh, are visualized in the accompanying figures in the upper left corners. Topologies of each alternative are illustrated in the accompanying figures in the bottom right corners.

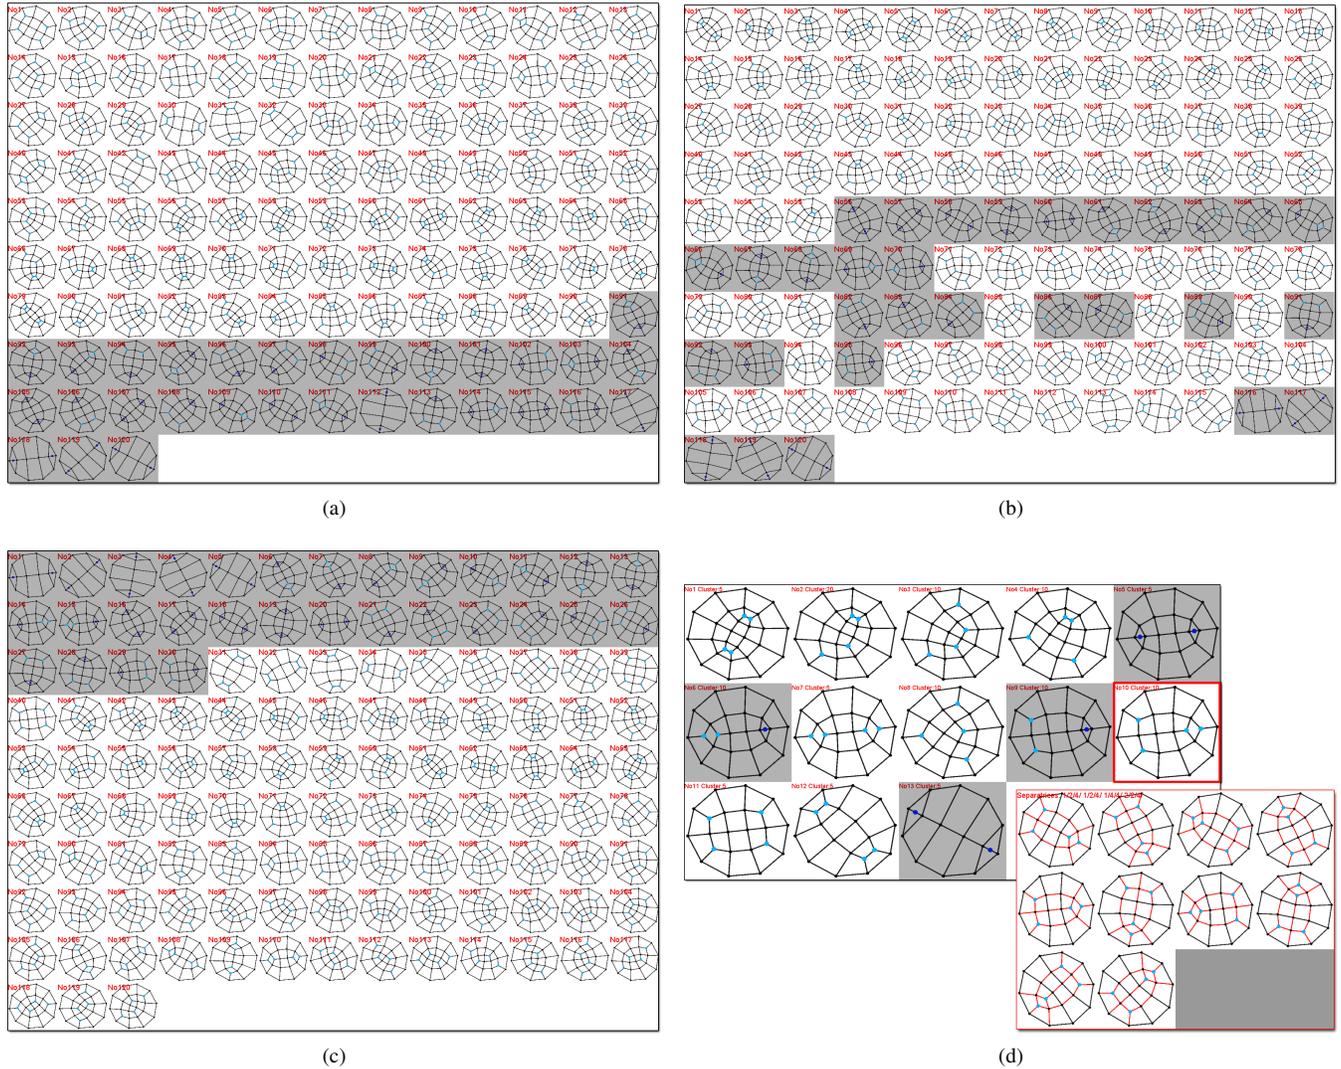


Fig. 12: Navigating topological variations of a $TVD = 4$ patch with one side of length 10 and no corners (a circle). There are infinite topologies for such a patch even under the constraint of zero $v3-v5$ pairs. Thus, we only enumerate topologies up to a reasonable density (number of quads). (a) to (c): Topological variations sorted by (a) geometric qualities, (b) densities, and (c) topology profiles. (d): Rotationally equivalent variations are clustered and each cluster is visualized by one representative topology. All rotationally equivalent variations of the selected cluster are shown in the sub-window. They all share the same topology profile: $1 - 2 - 4, 1 - 2 - 4, 1 - 4 - 4, 2 - 2 - 4$, calculated based on the length of separatrices shown as red poly-lines.

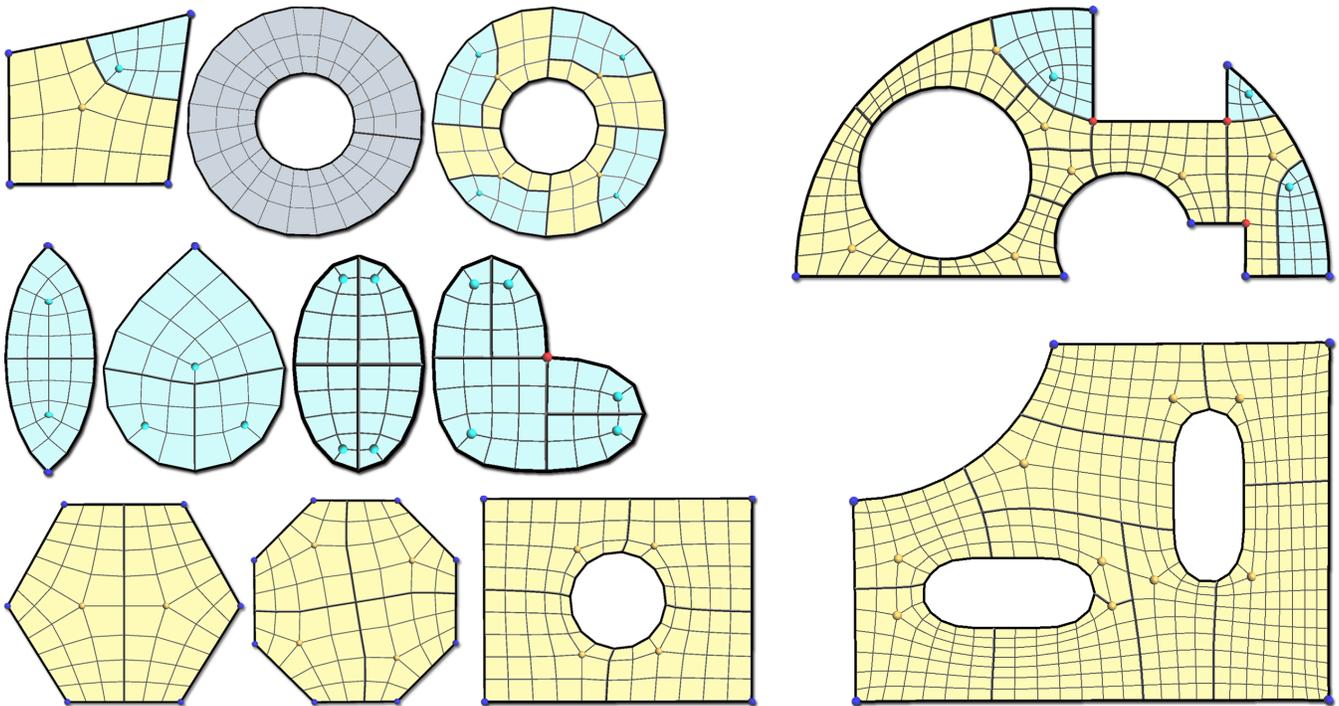


Fig. 13: The decomposition into simple patches for quadrangulations of patches with various boundary configurations. Parallelograms, triangles, and pentagons are shown as gray, cyan and yellow connected quads. The top left row shows various general 4-sided polygons ($TVD = 0$), which are quadrangulated with multiple v_3 - v_5 pairs. The middle left row shows various patches with $TVD > 2$, which are quadrangulated with multiple inner v_3 s. The bottom left row shows various patches with $TVD < -2$, which are quadrangulated with multiple inner v_5 s. On the right, we show quadrangulations of two examples that were also used in [Blacker and Stephenson 1991]. Comparing to their results, our enumeration finds many results with fewer irregular vertices. In fact, in the bottom right example, the number of irregular vertices has been minimized.