COMPARATIVE ANALYSIS OF POLYGONAL MESH SIMPLIFICATION METHODS FOR REAL-TIME INTERACTIVE RENDERING

Beknazarova Saida Safibullayevna[1], Makhammadjonov Makhammadjon Alisher ogli[2]

[1] Professor, Dsc. technical science, professor of department Television and media technologies Tashkent University of Information Technologies named after Muhammad Al- Khwarizmi, Tashkent, Uzbekistan
[2] MSc, Lecturer at the department of Computer Science in Central Asian University, Tashkent, Uzbekistan
E-mail: saida.beknazarova@ gmail.com, m.maxammadjonov@centralasian.uz
ORCID: 0000-0001-7708-7616,
0009-0009-2609-9321

K E Y W O R D S

polygonal mesh simplification, level of detail, quadric error metric, vertex clustering, topology simplification, geometric error, real-time rendering

A B S T R A C T

Polygonal mesh simplification is a core geometric operation of the real-time rendering pipeline, since the number of primitives submitted to the graphics processor in a single frame must be bounded in order to sustain an interactive frame rate. Several decades of research have produced method families that differ fundamentally in the geometric quantity they minimize and in the constraints they impose on the surface, yet a comparison of these families against a single set of criteria is largely absent from the literature. This paper presents a comparative analysis of four method families on the basis of representative studies: adaptive octree vertex clustering with an extended quadratic error matrix, iterative contraction driven by quadric error metrics, alpha-shape-based topology simplification, and voxel-based simplification employing morphological operators. A unified evaluation framework of seven criteria is proposed, comprising the class of error metric, the treatment of topology, the attainable reduction, computational cost and scalability, robustness to mesh degeneracies, preservation of surface attributes, and support for continuous level-of-detail transitions. The comparison is performed on the quantitative results reported in the source studies, reduced to a common form. It is shown that relaxing the topology constraint is the single largest lever on the attainable reduction, yielding improvement factors of up to 117.7 on mechanically structured models, whereas clustering methods retain a decisive advantage in execution time. It is further established that every error metric examined is purely object-space and observer-independent, which is formulated as the direction of further research.

1. INTRODUCTION

The production of large geometric models has grown faster than the capacity of graphics hardware to render them. Remote sensing generates terrain datasets of enormous extent; computed tomography and magnetic resonance scanners yield volume data from which isosurfaces of hundreds of thousands of triangles are extracted; computer-aided design systems routinely produce assemblies whose polygonal representation runs into millions of faces [14]. In the design of gas, oil and power installations the concurrent use of geometric databases and interactive three-dimensional tools requires that the number of displayed polygons, and therefore of displayed objects, be explicitly bounded [15]. The mismatch between the complexity of available models and the throughput of the rendering pipeline is thus structural rather than transient, and it is not removed by advances in hardware alone.

The established response to this mismatch is the level-of-detail technique. For each object several polygonal representations of differing complexity are precomputed, and at each frame the representation whose cost the frame budget can absorb is selected [15, 17]. The technique separates into two distinct problems. The first is geometric: given a triangle mesh, construct an approximating mesh with substantially fewer faces whose deviation from the original is controlled. The second is a scheduling problem: given a set of objects, a set of available representations and a bounded frame time, choose an assignment that maximizes the delivered detail [6]. The present

(cc) BY *dtai.uz*

work is concerned with the first problem, which is the one belonging properly to engineering geometry, and treats the second only insofar as it constrains the properties an admissible approximation must possess.

The geometric problem has been attacked from several directions. Rossignac and Borrel merged all vertices falling within a cell of a spatial subdivision into a single representative vertex, obtaining a method that is extremely fast and indifferent to the degeneracies present in real data [17]. Schroeder, Zarge and Lorensen removed vertices lying in comparatively flat regions and retriangulated the resulting holes [18]. Hoppe formulated the progressive mesh, a representation in which a coarse base mesh together with a sequence of vertex splits encodes a continuum of resolutions [9], later extended to view-dependent refinement [10]. Garland and Heckbert introduced the quadric error metric, which stores at each vertex a symmetric matrix summarizing the squared distances to the planes of the adjacent triangles and drives an ordered sequence of vertex-pair contractions [7]; the metric was subsequently extended to surfaces carrying color and texture [8]. Cohen and colleagues bounded the approximation error explicitly by constructing offset envelopes about the original surface [2]. Turk redistributed points over the surface by mutual repulsion and retiled it [21], while Low and Tan improved clustering accuracy by removing its dependence on the position of the model in space [12].

A second line of work relaxed the requirement that the topology of the surface be preserved. Topology preservation is indispensable where the tunnels and cavities of a model carry meaning, as in the analysis of molecular structure or in finite-element modelling, but in interactive visualization it is a limiting factor: a plate perforated by a hundred small holes cannot be reduced below three hundred triangles by any topology-preserving method, although six faces represent its shape adequately once the holes fall below the size of a pixel [5, 14]. El-Sana and Varshney addressed this by extending the concept of alpha-shapes [4] to polygonal meshes under the $L\infty$ metric, filling those regions that a cube of a prescribed size cannot enter, and thereby reducing genus, removing protuberances and repairing cracks within one framework [5]. Schroeder attacked the same limitation from the side of decimation by admitting vertex-splitting operations [19], and Popović and Hoppe generalized the vertex split so as to represent changes of topology as well as of geometry [16]. Nooruddin and Turk moved the problem into the volume domain altogether, converting the mesh to voxels, applying morphological opening and closing, and recovering a polygonal surface by isosurface extraction [14].

Despite the volume of this literature, the individual studies are difficult to place against one another. Each reports its results on different test models, on different hardware, and, more importantly, under error tolerances that are defined differently and are therefore not directly commensurable: as a percentage of the bounding-box diagonal in one case, as a distance in voxel units in another, as the value of a quadratic form in a third. A comparison referred to a single explicit set of criteria is largely absent. The relevance of the present work follows from this gap, and its novelty consists in the construction of such a framework and in the reduction of the reported quantitative results to a common form.

The aim of the study is to perform a comparative analysis of the principal families of polygonal mesh simplification methods and to establish which properties of the geometric error criterion determine their applicability in real-time rendering. The tasks subordinate to this aim are: to formalize the simplification problem and the error measures used to control it; to construct a system of evaluation criteria; to apply that system to four representative studies; and to identify the limitation common to the examined approaches that defines the direction of further research.

2. THEORETICAL BACKGROUND

2.1. Statement of the simplification problem

Let a triangle mesh be given as the pair $M = (V, T)$, where $V = \{v_1, \ldots, v_n\}$ is a set of vertices in three-dimensional Euclidean space and T is a set of triangles, each triangle being an ordered triple of indices into V. The simplification problem consists in constructing a mesh $M' = (V', T')$ with $|T'|$ substantially smaller than $|T|$ such that the deviation of M' from M, measured by a chosen error functional E, is minimal. In the budgeted

(cc) BY

formulation, which is the one relevant to real-time rendering, the number of faces is fixed in advance and the error is minimized subject to that constraint:

$$E(M, M') \rightarrow min, \ |T'| \leq n_{max} \quad (1)$$

The dual, tolerance-driven formulation fixes an admissible error ε and minimizes the number of faces subject to $E \leq \varepsilon$. The two formulations are not equivalent in practice: the first guarantees a rendering cost and lets the quality float, the second guarantees a quality and lets the cost float. Interactive systems require the first, because the frame budget is a hard constraint, whereas engineering analysis generally requires the second.

The most general geometric measure of deviation between two surfaces is the two-sided Hausdorff distance, which for the meshes M and M′ regarded as point sets is defined as

$$d_H(M, M') = \max\{h(M, M'), h(M', M)\}, \ h(A, B) = \max_{a \in A} \min_{b \in B} \|a - b\| \quad (2)$$

The Hausdorff distance is a strict bound and is therefore the natural reference measure, but computing it exactly at every candidate operation is prohibitively expensive. All practical algorithms consequently substitute a local, incrementally maintainable surrogate, and the choice of that surrogate is precisely what distinguishes the method families from one another.

2.2. Local error measures

In the approach of Garland and Heckbert the error at a candidate position is expressed as a quadratic form. Writing a vertex in homogeneous coordinates as $v = [x\ y\ z\ 1]^T$, the error associated with placing a vertex at v is

$$\Delta(v) = v^T Q v \quad (3)$$

where the symmetric matrix Q of order four accumulates the contributions of the planes incident on the vertex. If a plane is written as $p = [a\ b\ c\ d]^T$ with $a^2 + b^2 + c^2 = 1$, so that $p^T v$ is the signed distance from v to that plane, then

$$Q_v = \sum_{p \in planes(v)} K_p, \ K_p = pp^T \quad (4)$$

The matrices of two vertices are summed when the vertices are contracted, so that the accumulated error can be maintained incrementally at constant cost per operation, and the optimal position of the replacement vertex is obtained by solving a linear system of order four [7]. This construction is the reason for the wide adoption of the metric: it converts a global geometric quantity into a local algebraic one without abandoning a defensible geometric interpretation.

Zhou and colleagues observed that the quadratic form above accounts only for the distance from the newly created vertex to the planes of the original surface, and not for the distance from the removed vertices to the planes of the new surface. They therefore defined the error of clustering a vertex set as the sum of two terms,

$$\varepsilon(Vs_i) = \varepsilon_{i,1} + \varepsilon_{i,2} \quad (5)$$

in which the first term is the sum of squared distances from the representative vertex to the triangles of the related triangle set and the second is the sum of squared distances from the clustered vertices to the triangles incident on the representative vertex. Since the position of the representative vertex and the second term depend on one another, the pair is resolved by fixed-point iteration, which the authors report as converging to an acceptable accuracy within three steps [22]. The measure is two-sided in the same sense as the Hausdorff distance of equation (2) and is therefore geometrically stricter than the one-sided form of equation (3), at the cost of an iterative evaluation.

2.3. The selection problem as a constraint on the geometry

Once a family of approximations has been constructed, the rendering system must decide at each frame which member of the family to draw. Associating with every triple consisting of an object, a level of detail and a rendering algorithm a benefit B and a cost C, the selection problem takes the form

$$\max, \ \sum_{s \in S} B(O, L, R) \rightarrow \sum_{s \in S} C(O, L, R) \leq T_{target} \quad (6)$$

This is a continuous multiple-choice knapsack problem and is therefore NP-complete; a

(cc) BY

greedy approximation that orders the candidates by the ratio of benefit to cost and selects until the budget is exhausted is used in practice [6, 15]. Two consequences bear directly on the geometric problem. First, the family of approximations must be reasonably dense in the resolution axis, since a coarse family forces the scheduler into poor choices; four to six levels are typical [15]. Second, and less obviously, an abrupt substitution of one level for another is perceived as a discontinuity, and the accepted remedy is to introduce hysteresis by using two distinct switching distances, one for refinement and one for coarsening [15]. The geometric method must therefore support either continuous interpolation between adjacent levels or a nested relationship among them; a method that produces unrelated meshes at different resolutions is difficult to employ regardless of the accuracy of any single level.

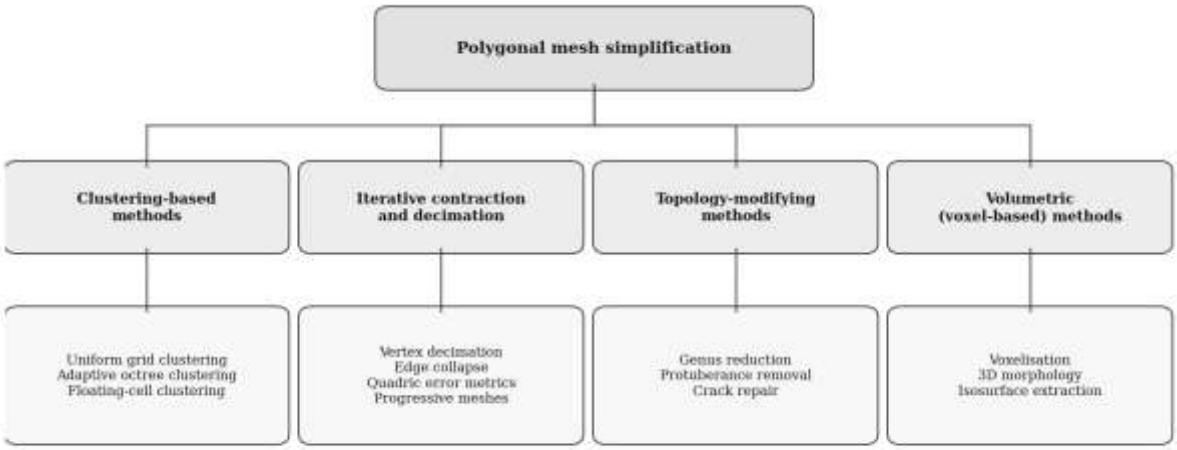


Figure 1. Classification of polygonal mesh simplification methods by the class of the underlying geometric operator.

3. MATERIALS AND METHODOLOGY

The material of the study consists of four primary sources, each representative of one of the families shown in Figure 1, supplemented by the body of work cited in Section 1. The primary sources were selected on three conditions: that the method be described in sufficient detail for its geometric operator to be identified unambiguously; that quantitative results be reported on named test models with stated face counts and execution times; and that the four sources together cover the whole classification rather than several variants of a single family. Their principal characteristics are summarized in Table 1.

Table 1.

Primary sources of the comparative analysis

No.	Source	Method family	Core geometric operator	Quantity controlling the error
1	Rossignac and Borrel [17]; Zhou et al. [22]	Clustering	Merging of all vertices within a cell of a spatial subdivision into one representative vertex	Cell size; two-sided quadratic error matrix
2	Schroeder et al. [18]; Garland and Heckbert [7]; Hoppe [9]	Iterative contraction and decimation	Ordered sequence of vertex removals or edge and vertex-pair collapses	Distance to the average plane; quadric error

(cc) BY

No.	Source	Method family	Core geometric operator	Quantity controlling the error
3	El-Sana and Varshney [5]	Topology-modifying	Filling of regions inaccessible to an L∞ cube of side 2α, followed by retriangulation	Tolerance α, expressed as a percentage of the bounding-box diagonal
4	Nooruddin and Turk [14]	Volumetric	Morphological opening and closing on a distance map, followed by isosurface extraction	Erosion and dilation distance in voxel units

The methodology is that of a criterion-based comparative analysis of published results. Seven criteria were formulated on the basis of the requirements derived in Section 2, namely: K1, the class of the error metric and whether it is one-sided or two-sided; K2, the treatment of topology; K3, the attainable reduction; K4, computational cost and its scaling with input size; K5, robustness to the degeneracies present in real data, that is, cracks, T-junctions, double walls, non-manifold edges and interpenetrating components; K6, preservation of surface attributes such as color and texture; and K7, support for continuous transitions between adjacent levels of detail, as required by equation (6) and the accompanying discussion.

The quantitative results reported by the source studies were reduced to a common form by expressing the outcome of each experiment as a percentage of faces removed relative to the original model, this being the only quantity defined identically across all four sources. Two limitations of this procedure must be stated explicitly. First, the execution times reported by the sources were obtained on hardware of different generations and are therefore not comparable in absolute terms; they are used only for the ratios established within a single source, where the hardware is held constant. Second, the error tolerances are defined differently in each source and are consequently not reduced to a common scale; the comparison of accuracy is therefore conducted qualitatively, on the basis of the analytical properties of the metrics established in Section 2, and not by direct numerical juxtaposition. Within these limits the comparison is valid, and the limits themselves are a finding of the study rather than an artefact of it.

4. RESULTS

4.1. Clustering-based methods

The defining property of clustering is that the cost of the method is governed by the number of vertices and by the cost of the spatial query, and not by the length of a sequence of dependent local operations. This yields the highest throughput of any family and an indifference to input degeneracies, since a cluster is well defined whether or not the incident triangles form a manifold [17]. The corresponding weaknesses follow from the same property. A uniform subdivision distributes cells without regard to the distribution of vertices, so that some cells receive many vertices and are over-simplified while others receive few and are simplified not at all; the representative vertex is commonly taken as the weighted centre of the cluster, which is not optimal in any stated error sense; and vertices belonging to distinct components of the scene may fall into a common cell and be merged, producing visible distortion [22].

The refinements introduced by Zhou and colleagues address each of these points within the clustering framework. The bounding box is subdivided adaptively by an octree, a leaf being split when it exceeds a prescribed size, when it contains more vertices than permitted, or when the merge would invert the normals of triangles in the related triangle set; the representative vertex is placed at the minimizer of the two-sided error of

(cc) BY

equation (5) rather than at the centroid; and where a cluster contains boundary vertices, planes perpendicular to the boundary edges are added to the accumulation of equation (4), which preserves boundary discontinuities. Color and texture are carried through simplification by constructing a texture map for the simplified mesh, each texel being assigned the attribute value of the nearest point on the original surface. The authors report that the resulting method runs at approximately half the speed of the uniform-grid method, requiring 44.02 seconds against 20.60 seconds on a bunny model of some 69,000 faces, while avoiding the block distortion that the uniform method exhibits at reduction levels above 95 per cent [22]. The trade-off is thus explicit and bounded: a factor of two in time purchases the removal of the characteristic clustering artefact.

4.2. Iterative contraction and decimation

Methods of this family repeatedly apply a local operator, a vertex removal or an edge collapse, in an order determined by an error criterion, and each application invalidates and requires the recomputation of the criterion in a neighborhood. The quadric error metric of equations (3) and (4) makes this update inexpensive, which is why it has become the reference method against which others are measured [7]. Because the operator is local and the criterion is accumulated, the family produces the most favorable relation between geometric error and face count of any examined here at moderate reduction levels, and the sequence of operations is itself a multiresolution representation, since reversing it refines the mesh; this is the basis of the progressive mesh [9] and of its view-dependent refinement [10], and it satisfies criterion K7 more completely than any other family.

The limitations are equally structural. The operator requires a well-formed neighborhood, so that cracks, T-junctions and non-manifold configurations either block the operation or must be handled by special cases [14]. More significantly, a strictly topology-preserving contraction cannot create non-manifold edges, and this blocks further simplification at a face count far above the geometric optimum. El-Sana and Varshney report the case directly: an object of 612 triangles, simplified by edge collapse under a fixed tolerance, is reduced only to 120 triangles, whereas removing the protuberances that constitute its topological detail yields 12 triangles at the same tolerance [5]. The limiting quantity, in other words, is not the geometry but the connectivity.

4.3. Topology-modifying methods

The alpha-shape approach of El-Sana and Varshney identifies those regions of a model that cannot be reached by a cube of side 2α rolled over the surface and fills them with a valid triangulation, the $L\infty$ metric being chosen over the Euclidean one because the resulting axis-aligned primitive makes the union and intersection operations robust and inexpensive, at the cost of rotational invariance of the tolerance. Candidate regions are located by the heuristic that genus-reducing simplifications occur in the neighborhood of sharp edges, taken as those whose dihedral angle exceeds seventy degrees; the same machinery removes protuberances when the primitive is rolled over the interior of the object instead of the exterior, and repairs cracks when the one-sided edges are treated as the boundary of a region to be filled [5].

The reported effect of alternating genus-reducing and genus-preserving stages, rather than applying the genus-preserving stage alone at the summed tolerance, is shown in Figure 2. The improvement is strongly dependent on the structure of the model: it is negligible on a sphere, which has no topological detail to remove, and reaches a factor of 117.7 on a mechanical fixture, whose small holes and tunnels dominate its face count. This dependence is itself the principal result of the family. Topology relaxation is not a uniform improvement but a lever whose length is set by the proportion of the face count devoted to features that will fall below the size of a pixel.

(cc) BY

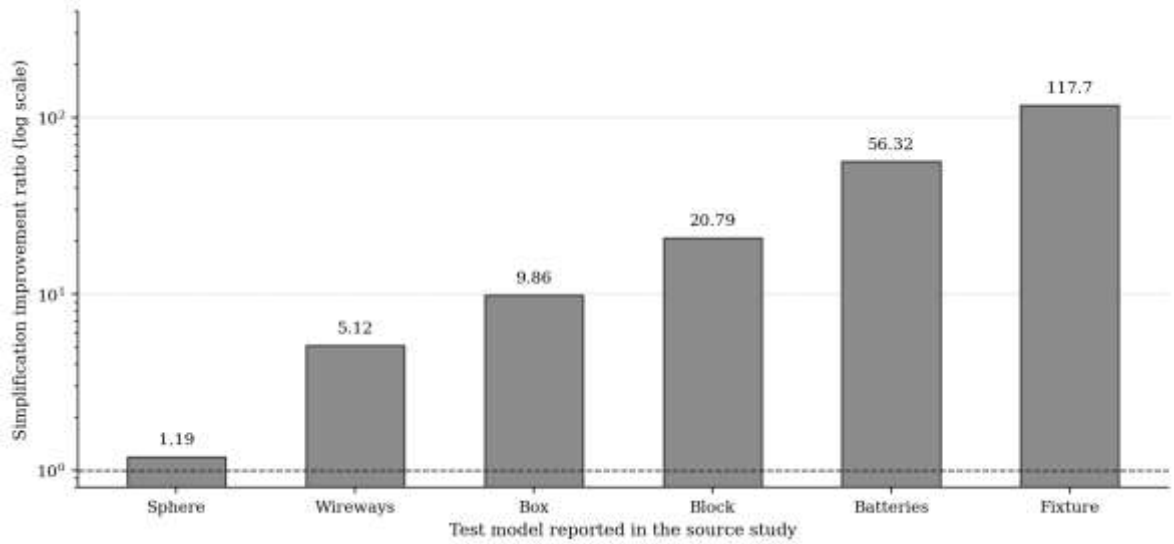


Figure 2. Improvement in the simplification ratio obtained by combining genus-reducing with genus-preserving simplification, by test model, according to the data of El-Sana and Varshney [5].

4.4. Volumetric methods

Nooruddin and Turk convert the mesh to a solid voxel representation, simplify its topology by morphological operators acting on a Euclidean distance map computed by the algorithm of Danielsson [3], extract an isosurface by marching cubes [11], smooth the voxelization artefacts by the low-pass filter of Taubin [20], and reduce the resulting face count by quadric-driven contraction [7, 14]. Two voxelization procedures are given. Parity counting classifies a voxel by the parity of the number of intersections between a ray through its center and the surface, and is extended to models with holes by voting among thirteen projection directions; ray stabbing retains only the first and last intersection along each ray and is required for models composed of interpenetrating parts or double walls, for which parity counting classifies the overlapped interior as exterior.

The strength of the approach is that the volumetric representation admits no topological ambiguity, so that the extracted isosurface is everywhere manifold by construction and the degeneracies of the input are removed as a side effect of the conversion; a bunny model with several large holes in its base is returned watertight, and a hand-built figure of fifteen interpenetrating components is returned as a single manifold surface whose parts cannot separate under further simplification [14]. Morphological opening and closing also give the user a direct and metrically meaningful control over the size of the features to be removed, which filtering approaches do not. The cost is the conversion itself: for a turbine blade of 1,729,892 faces the authors report that voxelization consumes 78 per cent of a total pipeline time of 24.35 minutes, and where the volume is large relative to the input mesh the morphological stage becomes the bottleneck instead. The family is therefore a preprocessing tool rather than an interactive one.

4.5. Cross-family comparison

Table 2 sets the four families against the seven criteria of Section 3. Table 3 collects the quantitative results reported by the sources, reduced to the common form described in the methodology.

(cc) BY

Table 2.

Comparative evaluation of the method families against the criteria K1–K7

Criterion	Clustering	Iterative contraction	Topology-modifying	Volumetric
K1. Error metric	Cell size; two-sided quadratic form in the adaptive variant	Quadric error, one-sided; distance to the average plane	Accessibility to an L∞ primitive of prescribed size	Distance map in voxel units; two-sided by construction
K2. Topology	Not preserved, and not controlled	Preserved; the principal limit on reduction	Modified under explicit control	Modified under explicit control
K3. Attainable reduction	High, but with distortion beyond about 95 %	Moderate; blocked by connectivity	Very high on models rich in topological detail	Very high; limited only by the voxel resolution
K4. Computational cost	Lowest; governed by the vertex count	Moderate; a sequence of dependent local updates	Low for the topological stage taken alone	Highest; dominated by voxelization
K5. Robustness to degeneracies	High	Low; requires a well-formed neighborhood	Handles T-junctions, T-edges and cracks	Highest; degeneracies removed by the conversion
K6. Attributes	Preserved by construction of a texture map	Preserved by extending the quadric to color	Not addressed	Color carried on surface voxels
K7. Continuous level of detail	By direct interpolation between levels	Complete; the operation sequence is the representation	Not addressed directly	Inherited from the final contraction stage

Table 3.

Quantitative results reported by the source studies, reduced to a common form

Source	Test model	Faces, original	Faces, simplified	Removed, %	Reported time
Zhou et al. [22]	Head	1 355	597	55.9	0.52 s
Zhou et al. [22]	Terrain	8 192	2 057	74.9	4.56 s
Zhou et al. [22]	Bunny	69 473	2 772	96.0	44.02 s
El-Sana, Varshney [5]	Box	1 612	14	99.1	2.3 s
El-Sana, Varshney [5]	Fixture	74 396	62	99.9	171 s

(cc) BY

Source	Test model	Faces, original	Faces, simplified	Removed, %	Reported time
El-Sana, Varshney [5]	Batteries	158 728	2 216	98.6	5.4 min
Nooruddin, Turk [14]	Al Capone	13 476	200	98.5	20.46 min
Nooruddin, Turk [14]	Motor	140 113	200	99.9	8.88 min
Nooruddin, Turk [14]	Turbine blade	1 729 892	200	99.99	24.35 min

Three regularities are visible in Table 3. The first is that the topology-modifying and volumetric families reach residual face counts one to two orders of magnitude below those of the clustering family on models of comparable size, and that this gap is not attributable to a difference in geometric accuracy but to the removal of the connectivity constraint. The second is that execution time grows with the input size far more steeply for the volumetric family than for the others, the turbine blade of 1.73 million faces requiring some twenty-four minutes against the seconds required by clustering on models two orders of magnitude smaller. The third, and the one that qualifies the other two, is that the residual face counts of the volumetric rows are not the outcome of an error criterion at all but a fixed target of two hundred faces chosen by the authors in order to exercise every stage of the pipeline; they measure what the method can be driven to, not what it selects.

5. DISCUSSION

The comparison establishes that the four families are not competitors on a single axis but occupy distinct positions in a trade-off among speed, attainable reduction, robustness and control. Clustering purchases speed and robustness by abandoning control over the error, and the adaptive octree variant recovers a substantial part of that control for a factor of two in time [22]. Iterative contraction purchases accuracy and a complete multiresolution representation by requiring a well-formed input and by accepting a floor on the attainable reduction imposed by connectivity [7, 9]. Topology modification removes that floor, but its benefit is a function of the model, ranging over two orders of magnitude across the test set of Figure 2 [5]. Volumetric processing purchases unconditional robustness and a metrically meaningful control over feature size at a cost that makes it a preprocessing stage rather than an interactive method [14].

It follows that the families are complementary rather than exclusive, and the source studies confirm this in their own construction. The volumetric pipeline terminates in a quadric-driven contraction stage, using the volumetric operators to remove the topological detail that would otherwise block that contraction [14]; the topology-modifying method is applied in alternation with a genus-preserving stage rather than in place of it [5]; and the adaptive clustering method imports the quadratic form of the contraction family to place its representative vertices [22]. A practical pipeline is therefore a composition: robust conversion to remove degeneracies, controlled topology reduction to remove features below the display threshold, and contraction under a geometric criterion to reach the budget.

The finding of greatest consequence, however, concerns what the examined criteria have in common rather than how they differ. Every error measure surveyed in Section 2 — the Hausdorff distance of equation (2), the quadric form of equation (3), the two-sided clustering error of equation (5), and the accessibility and distance-map criteria of the topology-modifying and volumetric families — is defined entirely in object space. Each is a function of the two surfaces alone and of no other argument. None takes as an argument the position of the observer, the angular

(cc) BY

size that the deviation subtends at the eye, or the sensitivity of the visual system at the retinal eccentricity where that deviation appears. The selection stage of equation (6) is the only point in the pipeline where the observer enters, and it enters weakly, through a benefit function that in practice depends on little beyond the projected size of the object [6, 15].

This separation was a reasonable engineering decision for the display conditions under which these methods were developed, in which the observer sits at an approximately fixed distance from a monitor subtending a small solid angle. It becomes questionable as the display conditions change. A geometric deviation of fixed magnitude is not equally perceptible when it lies on the silhouette of an object and when it lies in the interior of a smoothly shaded surface; nor when it lies at the center of the visual field and when it lies in the periphery; nor when the surface is viewed monocularly and when it is viewed under binocular disparity. A criterion that assigns the same cost to these cases must, at any given face budget, spend faces where they are not seen and withhold them where they are. Neither the direction nor the magnitude of the resulting misallocation can be estimated from the object-space error, since the two quantities are functions of different arguments.

The limitations of the present analysis should be stated. It is a comparison of published results and not of re-implementations, and it therefore inherits the incommensurability of the source tolerances discussed in Section 3; the qualitative rows of Table 2 rest on the analytical properties of the metrics rather than on measurement. The execution times of Table 3 are separated by more than a decade of hardware development and are meaningful only within a single row group. A controlled comparison, in which the four families are implemented against a common test set and a common tolerance scale, remains to be performed and is a necessary complement to the analysis presented here.

6. CONCLUSION

A comparative analysis of the principal families of polygonal mesh simplification methods has been carried out against a system of seven criteria covering the class of the error metric, the treatment of topology, the attainable reduction, computational cost, robustness to degeneracies, attribute preservation and support for continuous level-of-detail transitions.

The following results were obtained. First, the families occupy distinct and complementary positions in a trade-off among speed, reduction, robustness and control, and no family dominates the others on all seven criteria; the adaptive clustering variant exchanges a factor of two in execution time for the removal of the characteristic clustering artefact, while the volumetric family exchanges a cost that grows steeply with input size for unconditional robustness. Second, the constraint of topology preservation, and not geometric accuracy, is the binding limit on the reduction attainable by contraction-based methods, and relaxing it yields improvements ranging from negligible on topologically simple shapes to a factor exceeding one hundred on mechanically structured ones; the magnitude of the gain is a property of the model rather than of the algorithm. Third, an effective pipeline is a composition of the families rather than a choice among them, comprising a robust conversion stage, a controlled topology-reduction stage and a contraction stage driven by a geometric criterion. Fourth, and most consequentially, every error criterion examined is defined in object space alone and takes no argument describing the observer, with the result that the allocation of a fixed face budget cannot be related to the visibility of the deviations it produces.

The last of these results defines the direction of further research. It is required to construct an error criterion that admits the parameters of observation as arguments alongside the geometry, and to establish the conditions under which such a criterion yields a different and better allocation of a fixed face budget than the object-space criteria examined here. The construction of such a criterion, and its experimental verification under stereoscopic viewing conditions, constitute the subject of the authors' subsequent work.

REFERENCES

1. Barequet G., Kumar S. Repairing CAD models. Proceedings of IEEE Visualization '97. Phoenix, 1997, pp. 363–370.

(cc) BY

2. Cohen J., Varshney A., Manocha D., Turk G., Weber H., Agarwal P., Brooks F., Wright W. Simplification envelopes. Proceedings of SIGGRAPH '96. New Orleans, 1996, pp. 119–128.
3. Danielsson P.E. Euclidean distance mapping. Computer Graphics and Image Processing, 1980, vol. 14, no. 3, pp. 227–248.
4. Edelsbrunner H., Mücke E.P. Three-dimensional alpha shapes. ACM Transactions on Graphics, 1994, vol. 13, no. 1, pp. 43–72.
5. El-Sana J., Varshney A. Topology simplification for polygonal virtual environments. IEEE Transactions on Visualization and Computer Graphics, 1998, vol. 4, no. 2, pp. 133–144.
6. Funkhouser T.A., Séquin C.H. Adaptive display algorithm for interactive frame rates during visualization of complex virtual environments. Proceedings of SIGGRAPH '93. Anaheim, 1993, pp. 247–254.
7. Garland M., Heckbert P.S. Surface simplification using quadric error metrics. Proceedings of SIGGRAPH '97. Los Angeles, 1997, pp. 209–216.
8. Garland M., Heckbert P.S. Simplifying surfaces with color and texture using quadric error metrics. Proceedings of IEEE Visualization '98. Research Triangle Park, 1998, pp. 263–269.
9. Hoppe H. Progressive meshes. Proceedings of SIGGRAPH '96. New Orleans, 1996, pp. 99–108.
10. Hoppe H. View-dependent refinement of progressive meshes. Proceedings of SIGGRAPH '97. Los Angeles, 1997, pp. 189–197.
11. Lorensen W.E., Cline H.E. Marching cubes: a high resolution 3D surface construction algorithm. Computer Graphics, 1987, vol. 21, no. 4, pp. 163–169.
12. Low K.-L., Tan T.-S. Model simplification using vertex clustering. Proceedings of the Symposium on Interactive 3D Graphics. Providence, 1997, pp. 75–81.
13. Luebke D., Erikson C. View-dependent simplification of arbitrary polygonal environments. Proceedings of SIGGRAPH '97. Los Angeles, 1997, pp. 198–208.
14. Nooruddin F.S., Turk G. Simplification and repair of polygonal models using volumetric techniques. IEEE Transactions on Visualization and Computer Graphics, 2003, vol. 9, no. 2, pp. 191–205.
15. Pajon J.L., Guilloteau P. Geometry simplification for interactive visualization of complex engineering data. Transactions on Information and Communications Technologies, 1995, vol. 12, pp. 51–57.
16. Popović J., Hoppe H. Progressive simplicial complexes. Proceedings of SIGGRAPH '97. Los Angeles, 1997, pp. 217–224.
17. Rossignac J., Borrel P. Multi-resolution 3D approximations for rendering complex scenes. Modeling in Computer Graphics. Berlin, Springer-Verlag, 1993, pp. 455–465.
18. Schroeder W.J., Zarge J.A., Lorensen W.E. Decimation of triangle meshes. Computer Graphics, 1992, vol. 26, no. 2, pp. 65–70.
19. Schroeder W.J. A topology modifying progressive decimation algorithm. Proceedings of IEEE Visualization '97. Phoenix, 1997, pp. 205–212.
20. Taubin G. A signal processing approach to fair surface design. Proceedings of SIGGRAPH '95. Los Angeles, 1995, pp. 351–358.
21. Turk G. Re-tiling polygonal surfaces. Computer Graphics, 1992, vol. 26, no. 2, pp. 55–64.
22. Zhou K., Zhang M., Shi J., Pan Z. A new simplification algorithm for colored or textured polygonal models. The International Journal of Virtual Reality, 2000, vol. 4, no. 4, pp. 1–20.

(cc) BY