

Structure Recovery via Hybrid Variational Surface Approximation

Jianhua Wu Leif Kobbelt

Computer Graphics Group, RWTH Aachen University, Germany

Abstract

Aiming at robust surface structure recovery, we extend the powerful optimization technique of variational shape approximation by allowing for several different primitives to represent the geometric proxy of a surface region. While the original paper only considered planes, we also include spheres, cylinders, and more complex rolling-ball blend patches. The motivation for this choice is the fact that most technical CAD objects consist of patches from these four categories. The robust segmentation and global optimization properties which have been observed for the variational shape approximation carry over to our hybrid extension. Hence, we can use our algorithm to segment a given mesh model into characteristic patches and provide a corresponding geometric proxy for each patch. The expected result that we recover surface structures more robustly and thus obtain better approximations with a smaller number of primitives, is validated and demonstrated on a number of examples.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Curve, surface, solid and object representations

1. Introduction

Recovering the inherent structure and finding a faithful approximation of a given surface's geometry have always been challenging tasks in computer graphics research. Due to the steadily increasing availability and complexity of geometric models, better understanding of the underlying surface structures and characteristic geometric features becomes more and more critical nowadays. Accurately recovered surface structures will provide best blueprints for most geometry processing tasks which require optimal surface approximations, like multiresolution modeling, optimal domains for subdivision modeling, remeshing, high quality LODs, progressive compression, hybrid shape representations, etc. Moreover, traditional CAD applications including reverse engineering and rapid prototyping can benefit from such representations as well.

In this paper we present an automatic and robust surface structure recovery algorithm (cf. Fig. 1) which extends the powerful optimization technique of variational surface approximation [CSAD04]. In addition to planes, we allow for higher order surface elements like spheres, cylinders and more complex rolling-ball blend patches to embrace a new hybrid variational framework. By this, surface structures can




Figure 1: The face model (16K faces, left) approximated by 5 spherical (middle) and 5 planar (right) proxies. Our hybrid scheme can more faithfully separate the distinct face parts like the forehead, the eyes, the cheeks and the nose.

be recovered very elegantly and more robustly than with typical reverse engineering processes [VMC97], especially for mechanical CAD models which inherently consist of clear geometry structures. Moreover, as a *side product* of the improved structure recovery, our hybrid approximation method typically provides a higher approximation quality compared to the standard variational approximation scheme in particular when a very coarse segmentation is utilized.




Figure 2: The rocker arm model (60K faces) approximated from left to right, by 150, 75, 30 planar and 30 hybrid proxies. Top row shows their respective partitioning (gray color for planes, red for spheres, green for cylinders and blue for blend patches), middle row: approximation results via vertex projection onto the corresponding proxies and bottom row: remeshed outputs with 578, 498, 508 and 496 triangles respectively and absolute maximum Hausdorff errors 0.18, 0.23, 0.26 and 0.19. Note the drastical decrease of the approximative quality when fewer planar proxies are used (center right), comparing to more favorable results of our hybrid scheme (far right) where most geometric structures can be captured with only 30 proxies.

1.1. Related Work

Many approaches have been proposed trying to find the *optimal* structure to approximate a given surface's geometry, with certain constraints either on the approximation tolerance or the target complexity. Unfortunately, this optimization problem has been proven to be NP-hard [AS94], which is also the reason why most approaches employ heuristics based on local geometry characteristics.

There are quite a lot well-studied mesh simplification and decimation methods, some of which address the above approximation problem by incrementally performing atomic decimation operators like vertex removal or (half)edge collapse according to a certain priority ordering [Hop96, GH97, Gar99, GGK02, LRC*03]. The resulting coarse mesh already reveals in some sense the structure of the original dense mesh and can function, e.g., as the base mesh of a multiresolution surface representation. An alternative approach are the "dual" methods to mesh simplification, i.e. the face clustering approximations [KT96, She01, GWH01, SSGH01, SWG*03, MK05], where progressive region merging operators are applied in the same greedy way to create a set of face clusters for geometric segmentation.

Surface remeshing techniques are another family of surface approximation approaches that can adapt to the underlying surface structures. They commonly re-sample the geometry carefully with well-shaped surface elements (often triangles or quads) in either isotropic or anisotropic fashion [KVL99, GVSS00, BK01, AMD02, SAG03, ACSD*03, MK04]. Although delicate information like sharp features or curvature lines can be used to guide the remeshing process, the approximation quality often can not be guaranteed especially for very coarse scale approximations. In addition, a global minimization of some specific error metric can hardly be achieved in this kind of schemes.

Recently, the powerful optimization technique of variational shape approximation was proposed in [CSAD04]. The approximation task is casted as a discrete, variational geometric partitioning problem which is driven by minimizing the global approximation error. As we will extend this scheme, a more detailed description will be given in Section 2. The variational framework is simple yet effective and elicits a provable good trade-off between conciseness and geometric distortion. However, its structure recovery power and the approximation capability are somewhat limited by

the restriction to planar elements, yielding unfavorable results as the number of proxies gets very small.

We extend the standard variational approximation method by allowing for more different primitives (e.g. spheres and cylinders) to represent the local geometry of a surface region. In this sense, our work is also closely related to reverse engineering [VMC97, PR98, BKV*02, PLH*04], which converts raw 3D data points into concise and explicit geometry representations. While for reverse engineering the segmentation phase and the surface fitting phase are usually separated, we merge both of them in the variational framework, i.e., segmentation and fitting are *alternatingly improved* driven by Lloyd iteration [Llo82], hence leading to an automatic process and better recovery of surface structures.

2. Variational Surface Approximation

In the following we give a more detailed description of the variational surface approximation scheme [CSAD04] in order to set up the notations for our discussion. Trying to find an optimal piecewise-linear approximation of the input geometry, two major phases, *partitioning* and *fitting*, are repeated alternatively based on Lloyd’s clustering idea to minimize the total approximation error. Variational geometric partitioning is first applied to the input mesh \mathcal{M} . Each region \mathcal{R}_i in the partitioning \mathcal{R} will contain a set of triangles $\{T_i\}$ with barycenters $\{\mathbf{g}_i\}$ (note we can use \mathbf{g}_i to fit surfaces based on triangles rather than vertices). Then each \mathcal{R}_i is approximated by a planar *shape proxy* $P_i = (\mathbf{x}_i, \mathbf{n}_i)$, where \mathbf{x}_i and \mathbf{n}_i represent the “average” point and normal. After the fitting phase, the proxy set $\mathcal{P} = \{P_i\}$ approximates the whole geometry. As an alternative to the standard \mathcal{L}^2 error metric, the normal-based $\mathcal{L}^{2,1}$ metric is also used to evaluate the partition quality as well as to guide the fitting. For more details, the reader is referred to the original paper [CSAD04].

The second ingredient of the variational surface approximation framework is geometric re-partitioning where each triangle is assigned to the best fitting proxy. The two steps, fitting and re-partitioning, are iterated until convergence. This global optimization procedure is easy to implement and usually leads to better results compared to those obtained by greedy approaches. The *strengths* of variational shape approximation are twofold. Locally, the iterative technique is very sensitive to anisotropic bending of the surface which leads to an almost perfect alignment of the surface partition to the principle curvature directions. Globally, the technique reliably detects and merges flat regions which can be approximated by a single planar proxy.

The goal of our extension to this technique is to add more shape detection *power* such that higher order shape primitives like spherical or cylindrical regions can be properly recovered as well. We achieve this by integrating other classes of proxy geometries while leaving the rest of the algorithm almost unchanged. The resulting *hybrid* variational shape approximation techniques proves to converge to high-quality

partitionings and faithful shape approximations even when the number of proxies is very small (cf. Fig. 2).

3. Hybrid Variational Surface Approximation

Aiming at accurate surface structure recovery, we derive our hybrid scheme from the standard variational surface approximation. In addition to the standard planar surface elements, we also employ other higher order primitives like spheres, cylinders and more complex rolling-ball blend patches, in the variational framework (cf. Fig. 2 and 3). As hybrid surface elements lead to more geometric fidelity, we can recover more surface structures while using much fewer primitives (e.g., fewer regions in the partitioning). This results in more flexibility and even better blueprints for subsequent geometry processing tasks.




Figure 3: The fan disk model (13K faces) approximated with 24 hybrid proxies. Left two are partitioning outputs and right the remeshing result with 168 triangles.

To compose our new variational approximation scheme, in the following, we will first present the definitions of the new hybrid surface elements and then explain how they can be embedded into the standard framework. Error metrics will be modified correspondingly and finally diverse surface fitting techniques are discussed.

New Surface Elements In order to faithfully approximate a local surface region \mathcal{R}_i , we use the following elements as basic shape proxies:

- Planes $P_i = (\mathbf{x}_i, \mathbf{n}_i)$.
- Spheres $S_i = (\mathbf{c}_i, r_i)$, \mathbf{c}_i for sphere center and r_i for radius.
- Cylinders $C_i = (\mathbf{x}_i, \mathbf{d}_i, r_i)$, \mathbf{x}_i a point on the axis, \mathbf{d}_i its direction and r_i radius.
- Rolling-ball blend patch $B_i = (\mathbf{c}_i(t), r_i)$, $\mathbf{c}_i(t)$ the center trajectory and r_i the ball radius. Specifically $\mathbf{c}_i(t)$ is a uniform quadratic B-Spline curve plus two rays extending on the two ends in its tangent directions .

The motivation for this choice is mainly due to the fact that most technical CAD objects consist of patches from these four categories. Accordingly, the full proxy set now will be $\mathcal{P} = \{P_i\} \cup \{S_i\} \cup \{C_i\} \cup \{B_i\}$. While obtaining higher approximation power, we still note that, except the blending patch element, the other two new surface primitives have similar memory consumptions as planes and hence will not bring more overhead to the approximation setting. Even though the cylinder is just a special case of the rolling-ball

blend patch, we keep both of them for compactness reasons, since cylinders need less data than blend patches.

New Elements Embedding We can embed the above proxy geometries into the variational framework with modified error metrics (explained below). Otherwise the geometric partitioning phase is kept untouched.

The major changes appear in the proxy fitting phase. For each partitioned region \mathcal{R}_i , we compute not only the best fitting plane P_i , but at the same time the best fitting sphere S_i , cylinder C_i and blend patch B_i . From those fitted basic elements, we choose the one as the *optimal* shape proxy SP_i that minimizes the following fitting energy:

$$E(\mathcal{R}_i, SP_i) = \sum_{T_i \in \mathcal{R}_i} \mathcal{L}^*(T_i, SP_i).$$

where \mathcal{L}^* refers to the \mathcal{L}^2 or $\mathcal{L}^{2,1}$ metric.

To start the Lloyd iteration, we set all initial elements to be planes and perform a flood-based partitioning. The initial seeds can be randomly selected or by farthest point initialization as in [CSAD04]. Then in the first proxy fitting phase, better new types of surface elements can possibly be found and the iteration can continue.

Since the fitting algorithms for the new surface elements are more complex and thus slower, we use *progressive partitioning* in our implementation to speed up the whole process, i.e. we only use planes for fitting in the first iterations until the partitioning is close to stable (e.g., the change of approximation errors becomes very small for two consecutive iterations, this is also the terminating condition for the whole procedure). Then we also allow spheres and cylinders to fit surface regions and continue the iterations. Finally we permit the rolling-ball blend patches in the iterations to reach the final stable geometric partitioning.

Modified Metrics For the \mathcal{L}^2 distance metric, the extension to the new proxy shapes is straightforward: let $T_i = (v_1, v_2, v_3)$ be a triangle of area $|T_i|$ in the region \mathcal{R}_i associated with a shape proxy $SP_i = P_i \mid S_i \mid C_i \mid B_i$ and d_1, d_2, d_3 be the distances of vertices v_1, v_2, v_3 to the proxy SP_i . Then the \mathcal{L}^2 metric can be locally approximated by the exact solution for linear proxies [CSAD04],

$$\mathcal{L}^2(T_i, SP_i) = \frac{1}{6}(d_1^2 + d_2^2 + d_3^2 + d_1d_2 + d_1d_3 + d_2d_3)|T_i| \quad (1)$$

The only non-trivial question is how to compute the distance $d(v, B_i)$ from a vertex to a rolling-ball blend proxy. Here we first calculate the distance $d(v, \mathbf{c}_i(t))$ to the center trajectory, and then subtract from it the radius r_i , thus $d(v, B_i) = |d(v, \mathbf{c}_i(t)) - r_i|$.

To compute the $\mathcal{L}^{2,1}$ error metric, the normal information \mathbf{N}_i of a shape proxy SP_i near a triangle T_i with normal \mathbf{n}_T and barycenter \mathbf{g}_i is required, i.e., $\mathbf{N}_i = \mathbf{n}_i$ for planes P_i , $\mathbf{N}_i = \mathbf{g}_i - \mathbf{c}_i$ for spheres S_i , and $\mathbf{N}_i = \mathbf{g}_i - \mathbf{c}_\perp$ for cylinders C_i and blend patches B_i where \mathbf{c}_\perp is the projection of \mathbf{g}_i onto the

axis of the cylinder or onto the center trajectory of the blend patch. Note that \mathbf{N}_i has to be normalized and its direction might have to be flipped ($\mathbf{N}_i = -\mathbf{N}_i$) when for most triangles T_i in the region \mathcal{R}_i , $(\mathbf{n}_T \cdot \mathbf{N}_i) < 0$ holds. Finally the $\mathcal{L}^{2,1}$ error can be computed as follows:

$$\mathcal{L}^{2,1}(T_i, SP_i) = \|\mathbf{n}_T - \mathbf{N}_i\|^2 |T_i|. \quad (2)$$

Surface Element Fitting In practice we find that, though fitting a plane according to the $\mathcal{L}^{2,1}$ metric is fairly simple as pointed out in [CSAD04], it is not clear how this can be extended to our new types of shape proxies, since for spheres and cylinders, the normals do not change under scaling which implies that the radius of the proxy cannot be determined. Hence we will always use the \mathcal{L}^2 metric to guide the fitting step for both metrics and fortunately it works well for all of our test scenarios. The specific fitting techniques will be presented in the following sections.




Figure 4: A surface patch (orange) approximated with a sphere (left) and a cylinder (right).

3.1. Sphere Fitting

We use the robust least-square method of [Pra87] for sphere fitting (cf. Fig. 4). The sphere is represented implicitly as

$$f(x, y, z) = A(x^2 + y^2 + z^2) + Bx + Cy + Dz + E = 0.$$

For each region \mathcal{R}_i , we have to find a sphere that best approximates the barycenters $\{\mathbf{g}_i\}$ which are *weighted* by the corresponding triangle area $|T_i|$ to compensate for irregular sampling. The solution can be found by minimizing the quadratic form

$$F(A, B, C, D, E) = \sum_i [f(\mathbf{g}_i)]^2 |T_i|,$$

under the constraint

$$Q(A, B, C, D, E) = B^2 + C^2 + D^2 - 4AE = 1$$

which enforces a unit gradient of F near the sphere's surface.

For this, only a 5×5 general eigensystem has to be solved by using Lagrange multipliers. Usually the point set $\{\mathbf{g}_i\}$ will not degenerate to a line or a single point due to the teleporting mechanism in the variational partitioning, and hence we always get a proper solution. Once the parameters are computed, the approximating sphere $S_i = (\mathbf{c}_i, r_i)$ is specified by the coefficients A, \dots, E .

3.2. Cylinder Fitting

To geometrically fit a cylinder C_i to the region \mathcal{R}_i , we first use the curvature tensor field Γ [CSM03] to determine the direction \mathbf{d}_i of the cylinder axis. A robust statistical analysis is applied to the minimum curvature directions $\{\gamma_{\min}\}$ measured at all vertices in \mathcal{R}_i to get the average direction $\bar{\gamma}$ and their deviation $\Delta\gamma$. If the deviation $\Delta\gamma$ is larger than some threshold, the region is considered isotropic thus there is no need to further approximate it with an anisotropic cylinder. Otherwise, we let $\mathbf{d}_i = \bar{\gamma}$ and continue the fitting procedure.

We project the barycenters $\{\mathbf{g}_i\}$ onto the plane P passing through the origin \mathbf{O} with normal \mathbf{d}_i (cf. Fig. 5). We create a local frame Ψ on P and the projections $\{\mathbf{g}_i^\perp\}$ are fitted with a 2D circle (\mathbf{c}_i^*, r_i^*) to finalize the geometric scale the cylinder (cf. Fig. 4). This can be computed in a similar way as the sphere fitting by minimizing

$$F(A, B, C, D) = \sum_i [f(\mathbf{g}_i^\perp)]^2 |T_i|,$$

under the constraint

$$Q(A, B, C, D) = B^2 + C^2 - 4AD = 1,$$

where

$$f(x, y) = A(x^2 + y^2) + Bx + Cy + D = 0$$

is the circle representation.




Figure 5: Cylinder fitting. Barycenters are projected to a plane (left) and a 2D circle is fitted (right).

3.3. Rolling-Ball Blend Patch Fitting

Compared to the above methods, fitting a rolling-ball blend patch is more complicated and hence we need to derive a heuristic. We present here a three step fitting algorithm which is summarized in Fig. 6.

(1) Fit a curve $\mathbf{h}(t)$ to the barycenters $\{\mathbf{g}_i\}$ of the region \mathcal{R}_i . $\mathbf{h}(t)$ is defined as a uniform quadratic B-Spline curve with double end-knots:

$$\mathbf{h}(t) = \sum_{k=0}^n N_{k,2}(t) \mathbf{p}_k, \quad 0 \leq t \leq 1.$$

The number $(n+1)$ of control points \mathbf{p}_k is restricted to 15 in most cases the suppress oscillations as shown in Fig. 7. The initial parameterization $\{u_i\}$ for $\{\mathbf{g}_i\}$ is computed by projecting \mathbf{g}_i to a least-square line L^* approximating the point




Figure 6: Rolling-ball blend patch fitting. A surface patch (top left) is first fitted with a B-Spline curve (bottom left, blue) then a set of circles is found (top right, centers in black) and finally the center trajectory is fitted with another B-Spline curve (bottom right).

set $\{\mathbf{g}_i\}$ [Jol86]. Then a fitting routine is used to find the positions of \mathbf{p}_k by minimizing the following least-squares functional:

$$f(\{\mathbf{p}_k\}) = \sum_i \|\mathbf{h}(u_i) - \mathbf{g}_i\|^2.$$

Since the initial parameter values may not be distributed uniformly, and because the tangent directions at the end points of the fitted curve tend to point towards the barycenters with parameters 0 or 1, we add linear constraints with large weights ω at these two ends ($\omega * (\mathbf{p}_0 - 2\mathbf{p}_1 + \mathbf{p}_2) = 0$ and $\omega * (\mathbf{p}_{n-2} - 2\mathbf{p}_{n-1} + \mathbf{p}_n) = 0$) into the above system to compensate for the instability (cf. Fig. 8).

(2) Fit a set of circles $\{(c_i, r_i)\}$. First we compute corrected parameter values t_i for each barycenter \mathbf{g}_i , such that: $[\mathbf{g}_i - \mathbf{h}(t_i)] \perp \mathbf{h}'(t_i)$. Then we sort $\{\mathbf{g}_i\}$ according to $\{t_i\}$ and for each subsequence $\mathbf{g}_i \dots \mathbf{g}_{i+m}$ (m is usually 30), we project them to a plane Q_i defined by the point $\mathbf{h}((t_i + t_{i+m})/2)$ and the normal $\mathbf{h}'((t_i + t_{i+m})/2)$. A 2D circle is fitted similarly to Section 3.2 and mapping it back to 3D we get a circle $b_i = (c_i, r_i)$ within the same 3D plane Q_i .

(3) Fit the trajectory $\mathbf{c}(t)$ to the circle centers. We first have to check the validity of these circles by computing the average radius \bar{r} and its deviation Δr . If Δr is larger than some threshold, there would be no consistent blending patch for the region \mathcal{R}_i and we stop the fitting procedure. Otherwise, we simply fit another B-Spline curve, the center trajectory $\mathbf{c}(t)$, to the centers $\{c_i\}$ and set the final radius of the blend patch to $r_i = \bar{r}$.

As the center trajectory $\mathbf{c}(t)$ is only a finite curve with two endpoints, it is easy to imagine that during the partitioning phase, the rolling-ball blend patch could not grow at these ends. Hence we extrapolate the blend patch by extending two rays at the two ends of $\mathbf{c}(t)$ in tangent direction. By this an infinitely extended cylinder is added to make region growing possible. This strategy also recalls the importance of the linearity constraint shown in Fig. 8.




Figure 7: A surface patch fitted by B-Spline curves (blue) with 10 (left) and 40 (right) control points (red). More control points leads to jaggy curves and thus to an incorrect parameterization for rolling-balls fitting.




Figure 8: Portion of a cylinder patch fitted by a B-Spline curve (blue, 15 control points) without (left) and with (right) linearity constraint at the end points. Note the improved quality of final fitted center trajectory (orange).

4. Results

By adopting more types of surface primitives, our hybrid variational surface approximation scheme can accurately recover more surface structures and geometric details. As a side product, it also can offer higher approximation quality than the original variational method if the number of proxies is fixed, especially for technical CAD models. To compare our results with other previous work, we first introduce two comparison criteria:

Remeshing Quality The segmentation which emerges from the hybrid variational approximation can be exploited to control a decimation algorithm such that it generates coarse meshes (as remeshing output) which reflect the global structure of the input mesh very well. We prefer this method over the original one in [CSAD04] because it proves to be robust even in situations when the number of proxies is extremely low.

Our decimation algorithm for *surface remeshing* is a very simple variant of the standard QEM technique [GH97]. All we have to do is to multiply the priority values by a factor F which depends on the status of the corresponding vertex, i.e., $F = 1$ for vertices in the interior of a region, $F = 100$ for vertices on region boundaries, and $F = 1000$ for anchor vertices which belong to the boundary of three or more regions. Although this decimation technique turns out to produce fairly good meshes, in general we use this remeshing procedure mainly to generate coarse triangle meshes which allow us to compare our results to the results obtained by [CSAD04].

Vertex to Proxy Projection We not only project anchor vertices onto surface proxies as [CSAD04], but instead we map *all* vertices of the input mesh onto the geometric proxies that they are associated with. Projections of those vertices shared by several proxies are averaged accordingly. This operator is used to compare the *pure* geometric information captured by the surface proxies which are no longer planar in our case.

Figure 2 shows a complete comparison using above two quality criteria. We always use the $\mathcal{L}^{2,1}$ metric if not otherwise specified. The original variational approximation could not provide plausible results when the number of proxies is chosen too low. In contrast, hybrid approximation can recover more surface structures and geometric details than the standard variational approximation even when only one fifth of the proxies are used (30 hybrid proxies vs. 150 planar ones). Based on the well recovered surface structure, our hybrid scheme can also produce remeshed approximations with higher quality compared to the original method.

The global optimization properties for the variational shape approximation naturally carry over to our hybrid extension. To demonstrate this, we compare our scheme to QEM mesh simplification [GH97] in Fig. 9. It is obvious that both in terms of visual quality as well as approximation error, our scheme performs better than the greedy approach.




Figure 9: The rocker arm model remeshed based on 30 hybrid proxies (left, 496 faces) and decimated by QEM [GH97] (right, 500 faces) with respective Hausdorff errors 0.19 and 0.23. The bottom row shows the color coded error plots. Note that our hybrid approximation result preserves more surface features and geometric structures.

More complicated examples are shown in Fig. 10 and Fig. 11 for technical CAD models where the color coding scheme is same as in Fig. 2. Notice how well the original surface's geometric structures are recovered and expressed by only 50 or 29 hybrid proxies respectively. Although we are targeting at CAD models as our primary application, we




Figure 10: A CAD model (110K faces, top); its hybrid variational partitioning with 50 proxies (middle right); approximation via vertex projection (middle left) and remeshed to 500 triangles (bottom). Almost all geometric information is preserved with only 50 hybrid proxies.




Figure 11: From left to right, top to bottom are respectively the original input model (58K faces), its hybrid variational partitioning with 29 proxies, the approximation via vertex projection and the remeshed model with 400 triangles.

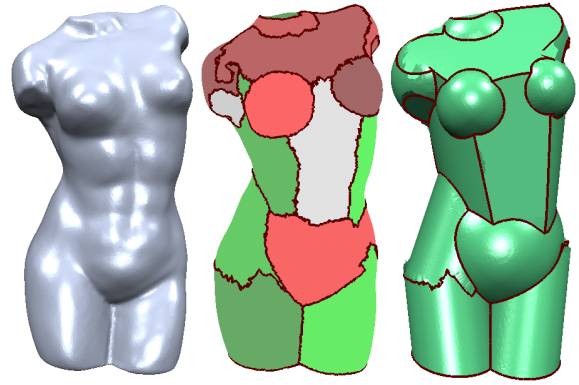


Figure 12: A torso model (80K faces, left), its hybrid partitioning with only 22 proxies (middle) and the recovered structures with approximation via vertex projection (right).




Figure 13: The hybrid partitioning of the bunny model (70K faces) with 31 proxies (left) and the recovered structures with approximation via vertex projection (right).

also tested our algorithm with more organic mesh models as shown in Fig. 12 and Fig. 13. It is interesting to see that equally good structure recovery capability can be observed.

Due to the increased number of surface primitive types and their more involved surface fitting algorithms with progressive partitioning, our hybrid method needs about 3–5 times more running time than the original scheme. However all datasets shown in the paper have been processed in less than 3 minutes. This seems worthwhile because the higher computational costs are traded for better surface structure recovery and higher approximation quality.

5. Conclusions

Dedicated to faithfully and reliably recovering surface structures, we have extended the powerful optimization technique of variational shape approximation by allowing different types of surface primitives to represent the local geometry of clustered regions. In addition to planes, we also include spheres, cylinders, and more complex rolling-ball

blend patches to form a new hybrid variational framework. In addition to the robust segmentation and global optimization properties which have been observed for the standard variational approximation scheme already, our hybrid extension can better extract the geometric configurations and hence produce better approximations with fewer primitives. This mostly automatic and high quality structure recovery property has been verified by the partitioning and approximation results for various mechanical and organic objects.

Future work can be extended in many directions. As in some rare cases, the variational framework still needs some user interaction (e.g. manual teleporting) to converge to the optimal result, we want to exploit more sophisticated heuristics to arrive at a fully automatic scheme. Finding the blend regions is sometimes difficult and our current blend patch fitting procedure is relatively slow. Last but not least, once having the variational partitioning and their respective hybrid proxies, we want more applications to benefit from it, like compact representations, remeshing, compression, optimal multiresolution models, and so on.

References

- [ACSD*03] ALLIEZ P., COHEN-STEINER D., DEVILLERS O., LÉVY B., DESBRUN M.: Anisotropic polygonal remeshing. *ACM Transactions on Graphics. Special issue for SIGGRAPH conference 22*, 3 (2003), 485–493.
- [AMD02] ALLIEZ P., MEYER M., DESBRUN M.: Interactive Geometry Remeshing. *ACM Transactions on Graphics. Special issue for SIGGRAPH conference 21*, 3 (2002), 347–354.
- [AS94] AGARWAL P. K., SURI S.: Surface approximation and geometric partitions. In *Proceedings of 5th ACM-SIAM Symposium on Discrete Algorithms* (1994), pp. 24–33.
- [BK01] BOTSCH M., KOBBELT L.: Resampling feature and blend regions in polygonal meshes for surface anti-aliasing. In *Eurographics 2001 Proceedings* (2001), pp. 402–410.
- [BKV*02] BENKO P., KOS G., VARADY T., ANDOR L., MARTIN R. R.: Constrained fitting in reverse engineering. *Computer Aided Geometric Design 19* (2002), 173–205.
- [CSAD04] COHEN-STEINER D., ALLIEZ P., DESBRUN M.: Variational shape approximation. *ACM Transactions on Graphics. SIGGRAPH 2004 23*, 3 (2004), 905–914.
- [CSM03] COHEN-STEINER D., MORVAN J.-M.: Restricted delaunay triangulations and normal cycle. In *19th Annual ACM Symposium on Computational Geometry* (2003), pp. 312–321.
- [Gar99] GARLAND M.: Multiresolution modeling: Survey & future opportunities. In *EUROGRAPHICS 99, State of the Art Report (STAR)* (1999), Eurographics Association, Aire-la-Ville (CH), pp. 111–131.
- [GGK02] GOTSMAN C., GUMHOLD S., KOBBELT L.: Simplification and compression of 3d-meshes. In *Tutorials on Multiresolution in Geometric Modeling* (2002), Springer.
- [GH97] GARLAND M., HECKBERT P. S.: Surface simplification using quadric error metrics. In *Proceedings of ACM SIGGRAPH 1997* (1997), Computer Graphics Proceedings, Annual Conference Series, ACM Press / ACM SIGGRAPH, pp. 209–216.
- [GVSS00] GUSKOV I., VIDIMCE K., SWELDENS W., SCHRÖDER P.: Normal meshes. In *Proceedings of ACM SIGGRAPH 2000* (2000), Computer Graphics Proceedings, Annual Conference Series, pp. 95–102.
- [GWH01] GARLAND M., WILLMOTT A., HECKBERT P.: Hierarchical face clustering on polygonal surfaces. In *ACM Symposium on Interactive 3D Graphics* (2001), pp. 49–58.
- [Hop96] HOPPE H.: Progressive meshes. In *Proceedings of ACM SIGGRAPH 1996* (1996), Computer Graphics Proceedings, Annual Conference Series, pp. 99–108.
- [Jol86] JOLLIFFE I.: *Principle Component Analysis*. Springer-Verlag, 1986.
- [KT96] KALVIN A., TAYLOR R.: Superfaces: Polygonal mesh simplification with bounded error. *IEEE Computer Graphics and Application 16*, 3 (1996), 64–77.
- [KVL99] KOBBELT L., VORSATZ J., LABSIK U., SEIDEL H.-P.: A shrink wrapping approach to remeshing polygonal surfaces. In *Eurographics 1999 Proceedings* (1999), pp. 119–130.
- [Llo82] LLOYD S.: Least square quantization in PCM. *IEEE Trans. Inform. Theory 28* (1982), 129–137.
- [LRC*03] LUEBKE D., REDDY M., COHEN J. D., VARSHNEY A., WATSON B., HUEBNER R.: *Level of Detail for 3D Graphics*. Morgan Kaufmann, 2003.
- [MK04] MARINOV M., KOBBELT L.: Direct anisotropic quad-dominant remeshing. In *Proceedings of Pacific Graphics 2004* (2004), pp. 207–216.
- [MK05] MARINOV M., KOBBELT L.: Automatic generation of structure preserving multiresolution models. In *Proceedings of Eurographics 2005* (2005). to appear.
- [PLH*04] POTTMANN H., LEOPOLDSEDER S., HOFER M., STEINER T., WANG W.: Industrial geometry: Recent advances and applications in cad. *Computer-Aided Design* (2004).
- [PR98] POTTMANN H., RANDRUP T.: Rotational and helical surface approximation for reverse engineering. *Computing 60* (1998), 307–322.
- [Pra87] PRATT V.: Direct least-squares fitting of algebraic surfaces. *Computer Graphics (SIGGRAPH 1987) 21*, 4 (1987), 145–152.
- [SAG03] SURAZHISKY V., ALLIEZ P., GOTSMAN C.: Isotropic remeshing of surfaces: a local parameterization approach. In *Proceedings of 12th International Meshing Roundtable* (2003).
- [She01] SHEFFER A.: Model simplification for meshing using face clustering. *Computer Aided Design 33* (2001), 925–934.
- [SSGH01] SANDER P., SNYDER J., GORTLER S., HOPPE H.: Texture mapping progressive meshes. In *Proceedings of ACM SIGGRAPH 2001* (2001), Computer Graphics Proceedings, Annual Conference Series, pp. 409–416.
- [SWG*03] SANDER P. V., WOOD Z. J., GORTLER S. J., SNYDER J., HOPPE H.: Multi-chart geometry images. In *Proceedings of Eurographics Symposium on Geometry Processing* (2003), pp. 146–155.
- [VMC97] VARADY T., MARTIN R., COX J.: Reverse engineering of geometric models - an introduction. *Computer Aided Design 29*, 4 (1997), 255–268.