

Reverse Engineering Methods for Digital Restoration Applications

Ioana Boier-Martin

IBM T. J. Watson Research Center,
Hawthorne, New York 10532
e-mail: ioana@us.ibm.com

Holly Rushmeier

Yale University,
New Haven, Connecticut 06520
e-mail: holly@acm.org

In this paper we discuss the challenges of processing and converting 3D scanned data to representations suitable for interactive manipulation in the context of virtual restoration applications. We present a constrained parametrization approach that allows us to represent 3D scanned models as parametric surfaces defined over polyhedral domains. A combination of normal- and spatial-based clustering techniques is used to generate a partition of the model into regions suitable for parametrization. Constraints can be optionally imposed to enforce a strict correspondence between input and output features. We consider two types of virtual restoration methods: (a) a paint restoration method that takes advantage of the normal-based coarse partition to identify large regions of reduced metric distortion suitable for texture mapping and (b) a shape restoration approach that relies on a refined partition used to convert the input model to a multiresolution subdivision representation suitable for intuitive interactive manipulation during digital studies of historical artifacts. [DOI: 10.1115/1.2356497]

1 Introduction

Reverse engineering is the process of taking an object apart for the purpose of analyzing its workings in detail, usually with the intention of constructing a new object of similar or extended functionality. In the case of 3D scanning, the objects are physical items for which digital replicas are created so that the original objects can be further studied, manipulated, and/or reproduced with the help of the computer. Today, 3D scanning has become a commodity service, with multiple applications, ranging from industrial design, to retail, to entertainment, to archeology and cultural heritage. Most often, the scanner software allows for the acquired data to be converted to an unstructured 3D mesh. This representation is typically difficult to manipulate beyond basic display and conversions to other representations are needed.

In this paper we focus on the demands of restoring digital objects for cultural heritage applications. However, the methods we present are relevant to many other areas.

This is a revised and extended version of our previous paper [1]. In that paper we specifically considered the problem of converting a triangle mesh to a Catmull-Clark multiresolution subdivision surface. Here we place that work in the context of cultural heritage restoration applications. This context motivates the problem of partitioning to facilitate texture map editing in addition to the partitioning needed for conversion to a subdivision surface. We reformulate our techniques to deal with the different partitions needed for intuitive texture map and shape editing.

1.1 Related Work. Our work is motivated by experiences in virtual restoration applications and we build on previous work in geometric processing.

Scanning and virtual restoration. The pipeline for processing to form a texture mapped triangle mesh from sensed data is well established [2]. Commercial products (e.g., Polyworks®, Rapid-Form®) are available to convert point clouds to meshes. These products also include some standard CAE/CAD analysis tools such as measurements on surfaces and comparisons of as-built to originally digitally designed parts. However, the needs of cultural heritage applications [3–5,2] are different from traditional CAE/

CAD in several different respects. First, the objects of interest are typically not regularly machined objects, but rather complex free-form shapes that were produced manually and have been worn or broken over time through natural processes. Second, while shape alone is sometimes of interest, the material appearance of the object as well as its shape are being studied. Third, several alteration iterations are needed to evaluate the validity of a change. Oftentimes, judgments are made on the basis of appearance and matching images and textual descriptions of appearance, rather than on more readily quantified requirements such as making parts fit together or minimizing weight. Fourth, the population using digital models in cultural heritage is much different from the population using CAE/CAD tools. While cultural heritage specialists may have extensive technical training, their primary interest is in the history of the objects being studied and the people that used them and not in the technology itself. The cost of hiring technical specialists and the difficulty of working through a mediator make it impractical to include pure technicians on most teams. Finally, the hardware and software resources available to cultural heritage projects are generally much more limited than those of engineering design projects.

Geometry processing. Considerable literature exists on building high-quality parametrizations over triangulated base domains (see [6] for a recent example and references). In contrast, very little work has been done on deriving quadrilateral base complexes for arbitrary meshes. The most recent work is that of Levy et al. [7]. Their periodic parametrizations are used to derive high-quality quadrilateral meshes. However, these meshes may be dense (the size of the quads is directly related to the size of the smallest tubular feature) and typically contain T-junctions so they are not directly usable as base meshes for subdivision surfaces. Eck et al. [8] describe an automatic method for fitting B-splines to meshes of arbitrary topology. A quad base domain is generated indirectly from a triangulated one by simplification followed by pairing of neighboring triangles. In general, approaches based on mesh simplification suffer from several shortcomings: it is not clear how to determine when to stop the simplification, geometric error typically drives the process with little or no control over the resulting topology and connectivity, the simplified mesh is a triangle mesh for which a quadrilateral decomposition has to be found, and constraints are difficult to enforce [9]. Generating quadrilateral meshes by pairing of triangles (see also 4-8 subdivision schemes [10]) pose additional problems: complete pairings may not always

Contributed by the Computer-Aided Product Development (CAPD) Committee of ASME for publication in the JOURNAL OF COMPUTING INFORMATION SCIENCE AND ENGINEERING. Manuscript received August 9, 2005; final manuscript received May 30, 2006. Guest Editors: A. Fischer, R. Levi, D. Dutta, and E. Cohen.

exist and finding ones that minimize distortion is expensive. Arbitrary meshes are typically parametrized by cutting and flattening. An extreme example is the parametrization of an entire mesh over a square [11]. Maintaining consistency across seams is difficult, especially if the model is to be modified (e.g., edited, compressed). In [12] conformal parametrizations of complex surfaces are computed without cutting. However, the resulting parametrizations are highly nonuniform and controlling them requires manual topology modification. Recent methods have targeted restricted classes of models. Mappings to either a sphere [13] or a plane [14] were used to recover quadrilateral meshes for genus 0 models with and without boundary, respectively. For the latter class, an interesting quad-dominant anisotropic remeshing method was described in [15]. Previous methods also include the user-driven approaches of Krishnamurthy and Levoy [16] to fit tensor-product B-spline patches to irregular meshes and that of Guskov et al. [17] to build quad base domains as part of their hybrid mesh representation. Techniques for converting given models to quadrilateral meshes have also been proposed in the mesh generation community. Advancing front and packing are among the most common strategies [18,19]. Typically, the resulting meshes have a large number of quads and are not suitable as parametrization domains.

Also related to our approach are mesh partitioning methods. The fuzzy clustering technique of [20] produces patches that are not homeomorphic to disks and cannot be directly used for parametrization. Other methods generate disklike patches according to various criteria (e.g., [21–24]), but no single method addresses the combination of patch shape, distortion, number of neighbors, alignment to features, and constraints.

Surazhsky et al. [25] use centroidal Voronoi tessellations to generate dense isotropic triangulations. The centroid updates are performed in 2D and require computing local parametrizations of model regions. This is a remeshing approach that does not produce a parametrization domain for the input model.

1.2 Goals and Contribution. The consideration of the needs of cultural heritage leads us to the following goals:

- Object representations should make minimal demands on computational resources during editing.
- Editing operations should be as direct as possible on the target representation.

The first goal requires that we have representations that do not require loading the full model to make edits. The second goal requires that we have representations that make sense to a viewer. In the case of the material or texture, we want textures that are partitioned along features that are meaningful to the user and are not distorted. This allows the user to update the texture map directly in a standard image editor (such as PhotoShop® or PaintShopPro®).

To achieve these goals we propose a segmentation approach that allows the user to focus on minimal representations of the object necessary to achieve a particular task. We present a novel unifying framework for manipulating both geometric and attribute data to allow the shape and attributes of the object to be edited directly and efficiently, without loading the full 3D model into memory. When shape edits are necessary, we rely on the ease-of-use of semi-regular parametric representations to perform the editing in an intuitive way. We reformulate our clustering-based techniques developed in [1] to allow us to handle different types of data partitioning required by virtual restoration and we illustrate the application of these techniques with multiple examples in this domain.

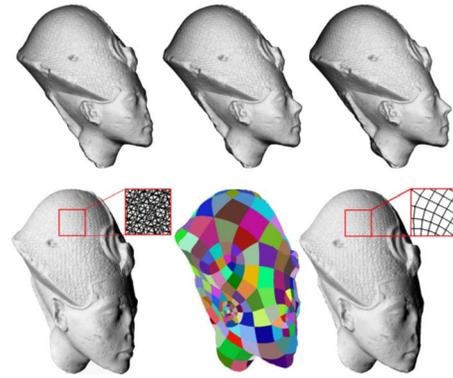


Fig. 1 Top: Akhenaton statue with broken nose (left) and two virtual reconstructions (middle, right). Bottom: underlying parametrization computed using our method and used for interactive shape editing.

2 Digital Model Creation

In this section we first describe some specific cultural heritage applications and their workflow. Next we review the main requirements of virtual restoration applications as considered in the present work.

2.1 Cultural Heritage Applications. Photography has long been an essential tool for documenting and communicating information. Nevertheless, photographs provide limited insight into the construction of objects, as they do not allow any interaction. Artists use photographs as starting points for rendering objects in their original states, but such depictions cannot be tested for physical feasibility. The power of 3D digital objects is that both the shape and material of the objects can be altered in physically feasible ways, and then critically evaluated to examine the validity of the changes.

An example of changes applied to a 3D scanned object is shown in Fig. 1. Figure 1 shows a scanned model of a sculpture of the head of the pharaoh Akhenaton and two proposed restorations of the sculpture’s nose. Figure 2 shows a small Greco-Roman queen statue and a proposed restoration of the surface colors. Other examples include a study of the breaking and repair of Michelangelo’s Florence Pietaà [26], the restoration of the original shape and finish of the Nara Great Buddha [5], and the virtual restoration of the Parthenon friezes [27].

The past two decades have seen tremendous efforts in obtaining the sensed data and building the initial 3D models. However, the second part of the process, editing and evaluating the models, has received less attention. Ironically, it is the second part that poses greater challenges as it requires multiple repetitions.

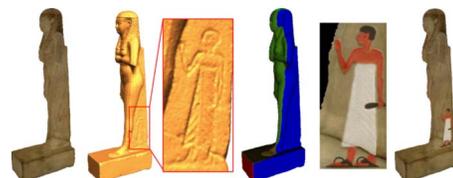


Fig. 2 Left to right: input model with textures; geometry shaded without textures; human character detail on the queen’s robe is enlarged for illustration purposes; segmentation using our method produces large height-field patches, well aligned with the main features of the model; previously painted 2D restoration of the character; the painted image is applied to the model by a simple cut-paste-and-blend operation in 2D texture space without the need for 3D intervention

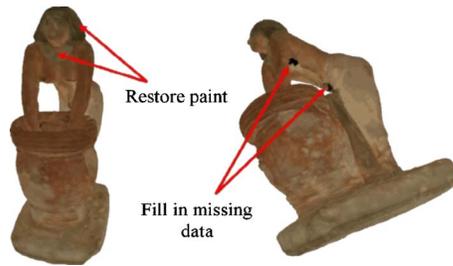


Fig. 3 Virtual paint restoration is needed in different contexts: to remove artifacts introduced by scanning (right) or to study different hypotheses on how the original might have looked like (left)

2.2 The 3D Scanning Pipeline. A wide variety of technologies are available for measuring 3D shape and appearance [28]. Two major types of scanners are frequently used for shape measurement—triangulation and time-of-flight scanners. The principle of triangulation scanners is to view a spot of light from a known source position with a sensor in a known position. Knowing the baseline distance between emitter and sensor and the angles of emission and viewing with respect to this baseline allow the calculation of the location of the spot. The principle of time-of-flight scanners is to estimate distance to a spot along a known ray direction by measuring the time between the emission of a pulse of light and the sensing of its return after reflection from the target to be measured. Using either technology, the output of a 3D scanner is typically a range image—a 2D array of measured 3D points for one surface patch on the object.

A pipeline of operations is necessary to convert collections of range images into a single object representation [2]. The pipeline begins by registering the range images into a global coordinate system. The points in overlapping range images are then adjusted along lines of sight to the scanner to account for measurement errors. The adjusted points are then joined into a single triangular mesh. Since scanners frequently oversample surfaces to capture small details, a simplification step usually follows the mesh generation step to reduce the number of vertices in the mesh.

Surface appearance properties, such as color and the magnitude of diffuse and specular reflectance, are generally captured using digital photography. The digital images are aligned to the geometric model either by calibration or by finding correspondences between geometric features that can be computed from both the digital images and the geometric model. Once aligned, the digital images need to be processed to estimate the intrinsic surface properties [29,30].

Once the digital images are processed to estimate surface properties, they need to be projected onto the geometric surface and combined into a single texture. To map a single image onto a 3D surface, the surface needs to be parametrized. In general, surfaces need to be partitioned into pieces and then projected or flattened onto a plane. Most commonly, early methods either used each triangle as a separate partition or relied on greedy approaches to produce groups of triangles forming nearly flat areas (e.g., [29]).

2.3 Virtual Restoration Requirements. The 3D scanning pipeline produces a triangle mesh model that is difficult to modify for virtual restoration. We can make restoration considerably easier for the user by partitioning the model into large patches for texture mapping and into quadrilaterals for conversion to a Catmull-Clark subdivision surface.

Virtual paint restoration. In the case of virtual paint restoration, there are two main categories of applications (see Fig. 3): (a) cultural heritage studies, which require color modifications across regions of the object (e.g., to test hypotheses regarding the original appearance of the objects or to simulate restoration procedures before they are performed) and (b) digital content creation appli-



Fig. 4 Drawback of 3D paint programs that require simplification: severe artifacts appear when the texture is remapped onto the simplified model

cations, which require touch up to remove scanning artifacts.

At first it would seem that emulating physical painting on a 3D object would be the most natural and efficient method to perform paint restoration. Several commercial packages allow direct painting on digital objects. Painting on relatively simple digital objects is a straightforward task for trained digital artists. However, for the detailed models of free-form shapes obtained by scanning, direct painting is not practical. In many cases the full models cannot even be loaded in memory for interactive viewing, let alone for more elaborate interactions.

One alternative is to load a simplified version of the geometric object and paint the texture using this simplified version. Except for very limited simplification, this technique causes distortions in the texture that make it difficult to see the paint adjustments that will appear correct on the full model (see Fig. 4).

Another alternative is to use the *projection paint* mode used in products such as BodyPaint® and DeepPaint®. A simplified version of the model is positioned while a high-resolution version of the model is rendered. The user paints on this high-resolution image. The results are then reprojected back onto the model. This approach still requires the user to interact with a 3D system. It also resamples and reprojects texture that may not be of as high a resolution as the captured texture.

Our goal is to produce a texture map with large undistorted patches, and to insure that critical areas for editing are in undivided regions. Large patches make it easy to edit textures without introducing discontinuities. Distorted textures allow very large patches, but are very difficult to edit to apply specific details. For users who are not technical experts it is difficult to explain why a texture map shows a distorted view of the object, and how they should mentally transform the change in flat texture to the three-dimensional object. The result is that when they edit a distorted texture they need to iteratively be reprojected onto the object in a 3D system to observe the effect.

In Sec. 3.3 we describe a method that allows us to generate texture maps that can be edited within a standard 2D paint program, without the need for cumbersome 3D manipulation.

Virtual shape restoration. 3D shape restoration requires the ability to edit the geometry of the digital model to produce a desired shape. While it is conceivable to selectively edit the triangle mesh output by the scanning software directly [31], interactive methods for general purpose editing with constraints are yet to be found for such representations. In contrast, parametric surfaces have been extensively used for modeling and design. In particular, subdivision surfaces have become a fundamental representation for styling and have obtained validation through repeated use in entertainment as well as industrial applications [32]. Hence, we proceed to automatically derive a semi-regular representation that approximates the input mesh. The underlying base domain construction for this representation is described in Sec. 3.3.

3 Partitioning

We now describe the overall clustering-based framework that produces the two different types of partitions discussed in the previous section.

3.1 Overview of Voronoi Partitioning Methods. We focus our attention on partitioning techniques that allow for easy-to-do virtual touch-ups. Our methods rely on concepts from clustering and quantization theory. In particular, we consider *Constrained Centroidal Voronoi Tessellations* (CCVTs) [33] to create decompositions suitable for editing. Our key idea is to use CCVTs to segment a given two-manifold into regions of reduced normal variation (i.e., almost flat) and then to use these resulting regions for the purpose of texturing the model, as well as for further decomposition for subsequent remeshing to support shape editing tasks. We briefly review the basic concepts related to CCVTs. Detailed properties can be found in [33]. Their applicability to segmentation tasks was studied in detail in [1].

Given a bounded domain $\Omega \subset \mathbb{R}^n$ and a set of K sample points $\{c_{ij}\}_{i=1}^K \in \Omega$, the *Voronoi partition* or *tessellation* \mathcal{V} induced by $\{c_{ij}\}_{i=1}^K$ on Ω is defined as:

$$\mathcal{V} = \{V_i = \{x \in \Omega : |x - c_i| < |x - c_j|, \\ j = 1, \dots, K, \quad j \neq i\}, i = 1, \dots, K\}$$

The points $\{c_{ij}\}_{i=1}^K$ are referred to as *generators* of the corresponding Voronoi regions V_i .

DEFINITION 1. A *centroidal Voronoi tessellation* (CVT) is a Voronoi tessellation in which the centroids (i.e., centers of mass) of the regions serve as their generators.

CVTs are closely related to statistical clustering algorithms. For discrete data sets, CVTs can be identified with K-means clustering methods [34]. This concept can be adapted to produce CVTs over arbitrary surfaces, by restricting $\Omega = S \subset \mathbb{R}^n$ to be a compact continuous surface and letting $\{c_{ij}\}_{i=1}^K \in S$ be a set of sample points on it [33].

DEFINITION 2. A *constrained centroidal Voronoi tessellation* (CCVT) is a Voronoi tessellation for which the *constrained mass centroid* c_i^* of each region V_i serves as its generator. The constrained mass centroid of a region $V_i \in S$ is defined as:

$$c_i^* = \operatorname{argmin}_{c \in S} F_i(c), \quad \text{where } F_i(c) = \int_{V_i} \rho(x) |x - c|^2 dx$$

where $\rho(x)$ is a given density function over Ω .

3.2 Our Partitioning Framework. We consider input models represented as two-manifold triangle meshes M_I of arbitrary topology, possibly with boundaries and possibly with an associated texture image. Feature curves along edges of M_I may be tagged as constraints. Such curves may be specified through user input or as a result of an automatic detection procedure.

We seek to build an atlas of patches covering the input model and suitable for texturing and subsequently for interactive shape editing. As our final representation we target a Catmull-Clark [35] multiresolution subdivision surface that allows for intuitive interactive shape editing operations to be performed [36–38]. We define the target representation through a control mesh hierarchy with L levels $M_{H_0}, \dots, M_{H_{L-1}}$. Automatic conversion from an input triangle mesh (with feature constraints) to such a representation is a difficult problem. We propose a Voronoi-based clustering technique such that:

1. M_{H_0} is a coarse quadrilateral mesh with a predominant number of regular control points (i.e., valence 4 in the interior and valence 2 on the boundary; we define the *valence* of a control point as the number of faces adjacent to it).

2. $M_{H_{L-1}}$ is a fine mesh such that its control points (or alternatively, their projections onto the limit surface of subdivision) are located on the input mesh M_I .
3. Each mesh M_{H_i} is obtained from the coarser mesh $M_{H_{i-1}}$ by applying the Catmull-Clark [35] subdivision rules, for $i = 1, \dots, L-1$. To accurately capture the input shape, the positions of control points on each level may be perturbed from the locations computed by subdivision using multiresolution detail vectors.
4. The texture of the input model can be transferred to the control meshes M_{H_i} , $i = 1, \dots, L-1$, with reduced distortion.

Center-Based Segmentation. Our segmentation method makes use of CCVTs in two distinct stages. An initial decomposition is computed for the purpose of texture mapping. The requirements in this case are that the resulting regions be relatively flat height fields to ensure reduced distortion during mapping. Constraints must be considered to ensure natural partitioning according to features and to simplify the paint restoration processes. A second decomposition proceeds to further break down the clusters found in the first stage into regions suitable for remeshing. In this case, the requirements for the regions are much stricter, as the boundaries of the regions will be directly used to extract a base mesh for the multiresolution representation. In the first case, the CCVT is computed based on normal information. In the second case, a combination of positional and normal information is used to generate a CCVT with well-shaped regions for remeshing. For texturing and texture editing we look for partitions that satisfy the following requirement:

(R1) For each region there exists a direction \vec{h} in space such that the corresponding geometry defines a height field along \vec{h} and can be approximated within some tolerance by a plane perpendicular to \vec{h} .

Subsequently, to parameterize the model for shape editing we further process the initial partition to also satisfy the following requirements:

(R2) Each region is homeomorphic to a disk.

(R3) The closed piecewise linear curve defining the boundary of each region can be approximated within some tolerance by a convex polygon.

(R1) is necessary because of the way in which textures are synthesized (see Sec. 2.2). Also, a mapping-dependent distortion is likely to be introduced when the mesh geometry is parameterized onto the face. To reduce the amount of distortion and to avoid resampling problems, we enforce (R1). (R2) is needed if mesh faces are to correspond to regions and (R3) ensures that they are well shaped.

Algorithm 1 provides a general-purpose method for segmentation according to generic properties associated with each triangular face of the input mesh (e.g., normal or location information). The algorithm computes centroidal decompositions based on these properties using a MacQueen-type approach [39]:

Algorithm 1 (center-based clustering of mesh faces):

Given an input mesh M with F faces, per-face property data items $\{d_{ij}\}_{j=1}^F$, and an initial set of K generator items $\{c_{ij}\}_{i=1}^K$: Repeat For each face f of M do

1. Find c_{i^*} among $\{c_{ij}\}_{i=1}^K$ closest to d_f
2. $j = \text{Label}(f)$, $\text{Label}(f) = i^*$
3. $c_{i^*} \leftarrow (|C_{i^*}|c_{i^*} + d_f) / (|C_{i^*}| + 1)$, $c_j \leftarrow (|C_j|c_j - d_f) / (|C_j| - 1)$,

until (*convergence*)

($|C|$ denotes the number of faces of cluster C ; the label of each face is the index of the cluster to which it belongs).

Depending on the application (e.g., paint or shape restoration), we apply Algorithm 1 to obtain partitions suitable for the restoration operations.

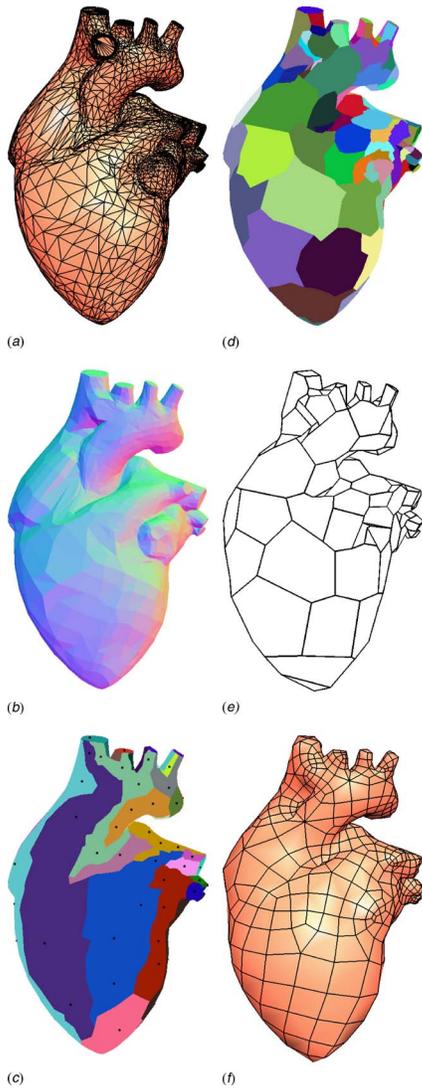


Fig. 5 Segmentation steps: (a) Input. (b) Variation of normals across the model. (c) Normal-based clustering and initial placement of generators (black dots). (d) Final partition after spatial-based clustering (using the generators from (c)) and cluster refinement. (e) Coarse polyhedral approximation extracted from the partition in (d). (f) Coarse level of a Catmull-Clark representation after resampling and multiresolution analysis.

3.3 Partitioning for Paint Restoration. We propose forming a texture that simplifies the work flow for restoring the surface texture and minimizes the amount of interaction the user needs to

have in a 3D viewer. By appropriately partitioning the surface, we can form a texture map that can be updated using any standard 2D paint program.

An initial classification of input mesh faces into approximately flat regions/clusters suitable for texturing is obtained through center-based clustering of normals. Face normals are computed and Algorithm 1 is applied with $d_f = \text{normal}(f)$, for all input faces f . The initial generators are chosen to be a small fixed subset of the set of all possible unit normal vectors (those pointing from the origin to the vertices, mid-edges, and face centers of a cube centered at the origin). To reduce the influence of geometry discretization, the normals are smoothed prior to classification. The result is a partition \mathcal{P}_N of the input mesh into regions of reduced normal variation.

This process is illustrated in Fig. 5(b). The variation of normals over the heart model is shown using a linear mapping of normal vector components to RGB color. The resulting mesh decomposition is shown in Fig. 5(c).

In the absence of constraints, the resulting partition tends to conform to salient features of the model. However, the boundaries of the regions may not conform to semantic components of the model (e.g., follow a hairline). To enforce separation along such boundaries we allow the user to outline features of interest and we use clustering with constraints (see Sec. 3.4) to segment the model.

Figures 2 and 6 illustrate the results of virtual paint restoration on two digital models corresponding to artifacts from the Egyptian National Museum collection (see Sec. 4). For models that only need texture restoration, no further partitioning is necessary.

3.4 Partitioning for Virtual Shape Restoration. For models that require shape restoration, we continue our Voronoi-based decomposition approach, starting from the partition \mathcal{P}_N previously obtained. For smooth, relatively simple shapes, \mathcal{P}_N has nice regions that are suitable for remeshing. However, this cannot be guaranteed. Since only a small number of regions are generated at this step, we can quickly check if all of them can be approximated by convex polygons. If this is not the case, we use the resulting decomposition for further refinement, as described next. In the interest of clarity, we defer the description of the shape check to Sec. 4.

Cluster refinement. In contrast to the normal-based partitioning scheme just described, spatial CCVT-type decompositions of the mesh are guaranteed to produce almost circular regions centered around their generators. The main issue to be addressed is choosing the number of generators and their initial locations.

A common strategy to select initial generators is to place random samples over the mesh. Two things must be specified: the number of samples and their distribution. For our problem, it is difficult to estimate in advance how many generators to use: too many lead to base meshes with many faces, whereas too few fail to properly represent the input. Given a number of samples, their

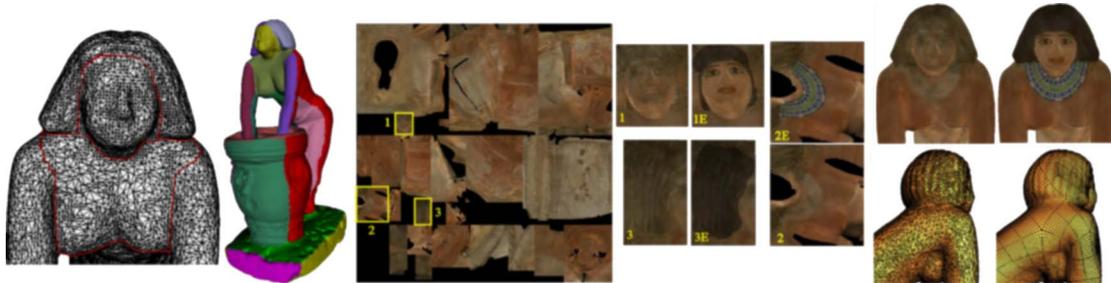


Fig. 6 Left to right: constrained segmentation followed by image-based editing (edited regions are shown numbered before and after editing; editing results are marked “E”); before and after comparison for paint restoration; underlying mesh before and after quad parametrization (thick lines indicate patch boundaries)

placement can be controlled using Monte Carlo methods or an error diffusion strategy [40], both of which can be computationally expensive for arbitrary models.

As an alternative, we use \mathcal{P}_N to estimate both the number of generators needed and their initial locations. Since regions of \mathcal{P}_N correspond to relatively flat model parts, the idea is to produce a refined decomposition in which clusters are centered on such flat spots. Ideally, we want to place samples along the medial axis of each region. In practice, we use the following heuristic method to avoid medial axis computation:

Algorithm 2 (region sampling):

Given a mesh region R with F_R faces:

1. Identify the set of boundary faces B
2. Uniformly distribute a set S_R of N random samples on $R \setminus B$
3. Select sample $s_1 \in S_R$ with the largest minimum distance to B
4. for $i=2, \dots, N$ do
5. Select $i \in S_R$ farthest from s_1, \dots, s_{i-1}
6. Retain N_0 of the N selected samples

After generators have been placed on the mesh according to Algorithm 2, a new partition \mathcal{P}_S is computed with Algorithm 1, this time based on spatial information. We classify mesh faces according to their proximity (measured using Euclidean distance) to generators. We use $d_f = \text{barycenter}(f)$, for all input faces f .

Cluster cleanup. By construction, clusters obtained after the refinement stage are isotropic, almost circular in shape. If the generators are sufficiently dense, (R2) is guaranteed to be satisfied by all clusters [41]. However, since we use sparse generator sets, some clusters may not conform to (R2). The purpose of the cleanup phase is to enforce (R1), (R2), and (R3) on all clusters. For this, we treat the geometry of each region of \mathcal{P}_S individually. The height-field condition (R1) is checked first and, if violated, the region is split using normal-based clustering. If any of the resulting subregions does not satisfy (R2) or (R3), the subregion is further split using spatial clustering. The normal-based split guarantees that the resulting regions are height fields. Subsequent spatial decompositions ensure that each height field is decomposed into regions with disk-topology and of approximately circular shape. Figure 5(d) shows the final decomposition of the heart model.

Coarse mesh extraction. Having found a partition of the input model that satisfies all of our requirements, the algorithm proceeds to generate a coarse mesh corresponding to it. Mesh vertices are placed at the points where three or more regions meet. Boundaries between regions define the edges of the coarse mesh.

Resampling and multiresolution analysis. Having found a coarse polygon mesh M_B that approximates the input model, we build a Catmull-Clark multiresolution representation by parametrizing the model over M_B . The faces of M_B correspond to regions that satisfy requirements (R1)–(R3). By virtue of the properties of CCVTs and Gershgorin's conjecture [33], the faces of M_B are predominantly hexagonal. Since all faces are convex polygons by construction, we can always find a partial quasi-conformal decomposition [42] as follows: split each face with $2k$ edges into $k-1$ quads; split each face with $2k+1$ edges into $k-1$ quads and 1 triangle. We apply the decomposition recursively: first we find the quad with the lowest shape number (see Sec. 4) that shares an edge with the face being quadrangulated, then we repeat the process for the remaining portion of the face. The decomposition is also perfect [42], as no Steiner points are introduced. This leads to a mesh consisting of predominantly quadrilateral faces (since the starting mesh was predominantly hexagonal) and a small number of triangular faces. One step of subdivision leads to a quadrilateral base mesh with a predominant number of regular vertices. Let M_{H_0} denote the quad mesh after one step of subdivision. M_{H_0} becomes a parametrization domain for the input model. To produce a hierarchy with subdivision connectivity, we subdivide M_{H_0}

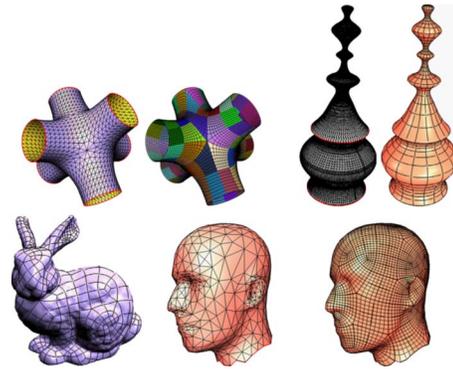


Fig. 7 Quadrilateral parametrization examples (constraints: indicated by shading)

to the desired number of levels L and compute data on the finest level by resampling followed by multiresolution analysis similar to [43,1].

Constrained parametrization. Our segmentation framework supports constraints in the form of closed curves along edges of the input model. Edges along feature curves are tagged and are used as clip boundaries during clustering. They also receive special handling during simplification and resampling. Figure 7 shows examples of segmentation with interior and boundary constraints.

4 Implementation and Results

Implementation details Our method has several parameters, which can be set with default values. We briefly describe our choices. To check (R1), we compute the average normal of a region \vec{h} and we test for the deviation of face normals from it. We also compute a plane perpendicular to \vec{h} that passes through the center of the bounding box of the region. The region passes the test if the maximum distance from region vertices to this plane is within some tolerance (5% of the bounding box diagonal). In practice, we observed this latter test to be rarely needed, due to the way regions are created (normal-based clustering first).

We use a simple check to decide whether a region obtained by clustering satisfies (R3). We apply this check following normal-based clustering, to regions that are relatively flat, so it makes sense to measure the convexity of their boundaries. A coarse polygon is defined for each cluster with vertices at the confluence of three or more clusters. A polygon is considered convex if each of its interior angles is at most π . If a polygon is found to be convex, we then consider how well it approximates the boundary of the region. For this, we measure the maximum distance between points along the boundary curves and the corresponding coarse lines approximating them. If the distance is below a tolerance threshold, the approximation is acceptable. Since clusters are groups of faces, boundaries between them may appear jagged (a face equidistant from two generators is arbitrarily categorized in one of the corresponding clusters). We define the approximation tolerance as a constant times the average edge length in the cluster and we ignore such jaggedness in shape testing.

We use a fixed number of Laplacian smoothing iterations on the normal field prior to clustering to attenuate noise. The default number in our implementation is 10, however this value should be increased for very noisy data. For *random sampling* according to Algorithm 2, we followed the following strategy: uniformly distribute [44] a large number of samples over the region (we use $F_R/2$ samples) and of these select $N = F_R/4$ using a farthest-first approach; N_0 of the N samples are kept, where N_0 is the index of the sample for which the largest jump in the max-min distance is observed [20].

Paint restoration. Figure 6 shows the result of segmenting an



Fig. 8 Left: naive segmentation of the model in Fig. 2 without constraints results in many regions. Right: the result of our segmentation method.

Egyptian artifact representing a woman making beer. The original statuette is 26.7 cm tall and was scanned at 1 mm resolution. The data were processed to remove redundant points in overlapping scans and to integrate the points into a single triangle mesh. The model has 112,478 faces and was segmented into 30 regions (second from the left) in 3.6 min. User-imposed constraints (curves; left) ensure that the face, chest, and hair of the model are properly separated to allow subsequent editing (i.e., without constraints the semantics of these parts is lost and they are arbitrarily split between regions, making it difficult to edit them). The middle images illustrate the corresponding texture map (note clearly identifiable head and chest textures). A comparison of the model before and after paint restoration is also shown (top right).

Figure 8 shows comparatively the texture maps obtained with a naive segmentation without constraints and with our technique. Features of interest such as face, hair, and necklace appear split between regions, making image-based editing difficult. In contrast, our segmentation produces a higher quality partition, with a much smaller number of regions and with region boundaries that respect user-imposed constraints, as well as salient geometric features.

Figure 2 shows a segmentation and detail painting of an Egyptian statue from the Greco-Roman period. Note that, in the absence of constraints, semantic features such as the face or the leg are not preserved.

Shape restoration. In Fig. 1 we illustrate a high-resolution scan of Akhenaton's head for which two different restorations of the nose have been created. We use a free-form variational editing tool (similar to [38]) to reconstruct the nose interactively. The user can smoothly deform the shape while preserving the Catmull-Clark connectivity. Many possibilities can be generated and evaluated quickly.

Additional parametrization examples are illustrated in Figs. 7 and 6 (bottom right) for meshes with and without sharp features and boundaries.

Performance statistics. We evaluate the quality of our output meshes using four different measures:

Hausdorff distance: provides a numerical estimate of the maximum distance between two meshes as the largest of two directed max-min distances from each mesh to the other. We use Metro [45] to report the results as a percentage of the mesh bounding box diagonal.

Volume ratio: quantifies the change in volume after remeshing (closed meshes only). It is defined as the ratio between the volume of the finest-level Catmull-Clark control mesh and that of the original.

Face shape distribution: characterizes the deviation of remeshed faces from an ideal shape (square in the isotropic case). Element shape quality is paramount in virtually all mesh-based computations, from finite element computations [46] to texture mapping and parametrization [22]. We plot a histogram of face shape measures for the faces of the base mesh. We quantify the deviation of a 3D quadrilateral mesh face $(p_i, p_j, p_{j+1}, p_{j+2})$ from a square by considering the affine mapping of the triangle (p_i, p_j, p_{j+2}) to the right triangle with unit-length legs (O, q_1, q_2) .

Table 1 Quality statistics: # F=number of input faces; # CF =number of coarse faces extracted with our method (before quadrangulation); HD=Hausdorff distance between input mesh and M_{H_3} (fourth subdivision level); $R_V = V_{M_{H_3}} / V_i$ =ratio of output mesh volume (M_{H_3}) to the input mesh volume; \mathcal{H}_S =histogram of face shape numbers for the quad base mesh M_{H_0} ; \mathcal{H}_V =histogram of vertex valences for M_{H_0}

| | heart | bunny | woman | pharaoh |
|-----------------|--------|--------|---------|---------|
| Size: | | | | |
| # F | 7,412 | 71,040 | 112,478 | 315,462 |
| # CF | 173 | 155 | 219 | 71 |
| Quality: | | | | |
| HD | 0.01 | 0.02 | 0.008 | 0.008 |
| R_V | 0.9994 | 0.9981 | 0.9991 | 0.9990 |
| \mathcal{H}_S | | | | |
| \mathcal{H}_V | | | | |

As pointed out in [47,22], the singular values of the Jacobian of this map characterize the local distortion between the right triangle and its 3D counterpart. We use the ratio of the singular values, i.e., the condition number of the Jacobian, as our shape measure. Simple calculations give $\mathcal{K}(J) = (|p_j - p_i|^2 + |p_{j+2} - p_i|^2) / (2A)$, where A is the area of (p_i, p_j, p_{j+2}) . We define the *shape number* S of a quadrilateral face as the average of the four condition numbers at its vertices, normalized so that the shape number of a square equals 1.

Valence distribution: characterizes the regularity of the base mesh.

Table 1 shows results for remeshed models with four levels of subdivision. The histograms confirm the quality of the base meshes, with a majority of well-shaped, almost square faces (shape numbers close to 1), a high percentage of regular vertices, and no highly irregular ones (valence 10 or higher). Face shape variation is also illustrated in Fig. 9.

Regarding time complexity, Algorithm 1 runs linearly through the input mesh faces and, for each face, through the set of generators. A simple test checking if any faces have changed clusters during an iteration is used as a *stopping criterion*. This is combined with a limit on the number of iterations to deal with cases when a face on the border between two clusters flips back and forth between them. We used a limit of 50 iterations for all our tests. The number of iterations to convergence for our models was

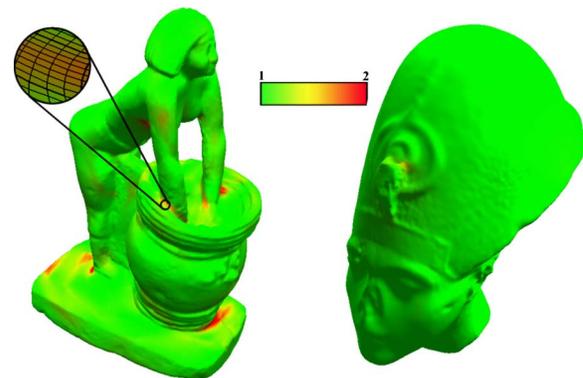


Fig. 9 Visualization of the face shape numbers across two models

less than 10. Algorithm 2 is also linear in the number of faces in each region. Multiresolution resampling is the bottleneck of the computation. To speed it up, we use a uniform grid approach as in [45], which gives reasonable results. In our tests, running times for the extraction of quadrilateral parametrization domains ranged from 6 to 40 s. Typical resampling times ran between a few seconds and 25 min for four levels of subdivision. The measurements were performed on a Pentium 4 3 GHz PC with 2 GB of RAM.

5 Conclusions and Future Work

In this paper we outlined the challenges of virtual restoration applications and we introduced novel methods for reverse engineering 3D objects for the purpose of interactive editing.

This work opens interesting possibilities for further exploration. Building anisotropic domains according to given direction fields is a natural next step. Another challenge is to use our constrained parametrization approach to generate maps between different models of similar shape with correspondence between features.

Acknowledgment

We thank Jingyi Jin and Richard Giantisco for their help. Portions of this work were performed as part of a collaboration with Dassault Systèmes. The Egyptian models are courtesy of the Egyptian Center for Documentation of Cultural and Natural Heritage and are available through the “Eternal Egypt” website (www.eternegypt.org).

References

- [1] Boier-Martin, I., Rushmeier, H., and Jin, J., 2004, “Parameterization of Triangle Meshes Over Quadrilateral Domains,” in *Proc. SGP 04*, July 8–10, Nice, France.
- [2] Bernardini, F., and Rushmeier, H., 2002, “The 3D Model Acquisition Pipeline,” *Comput. Graph. Forum*, **21**(2), pp. 149–172.
- [3] Godin, G., Beraldin, J. A., Taylor, J., Courmoyer, L., Rioux, M., El-Hakim, S., Baribeau, R., Blais, F., Boulanger, P., Domey, J., and Picard, M., 2002, “Active Optical 3d Imaging for Heritage Applications,” *IEEE Comput. Graphics Appl.*, **22**(5), pp. 24–36.
- [4] Levoy, M., Pulli, K., Curless, B., Rusinkiewicz, S., Koller, D., Pereira, L., Ginzton, M., Anderson, S., Davis, J., Ginsberg, J., Shade, J., and Fulk, D., 2000 “The Digital Michelangelo Project: 3D Scanning of Large Statues,” in *SIGGRAPH '00*, July 23–28, New Orleans, LA, pp. 131–144.
- [5] Ikeuchi, K., Nakazawa, A., Hasegawa, K., and Ohishi, T., 2003, “The Great Buddha Project: Modeling Cultural Heritage for VR Systems Through Observation,” in *IEEE ISMAR03*, October 7–10, Tokyo, Japan.
- [6] Khodakovskiy, A., Litke, N., and Schröder, P., 2003, “Globally Smooth Parameterizations With Low Distortion,” *ACM Trans. Graphics*, **22**(3), pp. 350–357.
- [7] Ray, N., Li, W. C., Levy, B., Sheffer, A., and Alliez, P., 2005, “Periodic Global Parameterization,” Technical Report TR, submitted to ACM TOG.
- [8] Eck, M., and Hoppe, H., 1996, “Automatic Reconstruction of B-Spline Surfaces of Arbitrary Topological Type,” in *Proc. SIGGRAPH 96*, August 4–9, New Orleans, LA, pp. 325–334.
- [9] Kho, Y., and Garland, M., 2003, “User-guided Simplification,” in *Proc. ACM I3D Symp.*, April 27–30, Monterey, CA, pp. 123–126.
- [10] Velho, L., 2000, “Quadrilateral Meshing Using 4–8 Clustering,” in *Proc. CILANCE'00*, December 6–8, pp. 61–64.
- [11] Gu, X., Gortler, S., and Hoppe, H., 2002, “Geometry Images,” *ACM TOG SIGGRAPH*, **21**(3), pp. 355–361.
- [12] Gu, X., and Yau, S.-T., 2003 “Global Conformal Surface Parameterization,” in *Proc. SGP03*, June 23–25, Aachen, Germany, pp. 127–137.
- [13] Praun, E., and Hoppe, H., 2003, “Spherical Parameterization and Remeshing,” *ACM Trans. Graphics*, **22**(3), pp. 340–348.
- [14] Hormann, K., and Greiner, G., 2000, “Quadrilateral Remeshing,” in *Proc. Vision, Modeling, and Viz 2000*, November 22–24, Saarbrücken, Germany, pp. 153–162.
- [15] Alliez, P., Cohen-Steiner, D., Devillers, O., Lévy, B., and Desbrun, M., 2003, “Anisotropic Polygonal Remeshing,” *ACM Trans. Graphics*, **22**(3), pp. 485–493.
- [16] Krishnamurthy, V. and Levoy, M., 1996, “Fitting Smooth Surfaces to Dense Polygon Meshes,” in *Proc. SIGGRAPH 96*, August 4–9, New Orleans, LA, pp. 313–324.
- [17] Guskov, I., Khodakovskiy, A., Schröder, P., and Sweldens, W., 2002, “Hybrid Meshes: Multiresolution Using Regular and Irregular Refinement,” in *Proc. 18th Annual Symp. Comp. Geom.*, June 5–7, Barcelona, Spain, pp. 264–272.
- [18] Owen, S., Staten, M., Canann, S., and Saigal, S., 1998, “A Survey of Unstructured Mesh Generation Technology,” in *Proc. 7th Int. Meshing Roundtable*, October 26–28, Dearborn, MI, pp. 239–267.
- [19] Bern, M. W., and Eppstein, D., 2002, “Quadrilateral Meshing by Circle Packing,” *Int. J. Comput. Geom. Appl.*, **10**(4), pp. 347–360.
- [20] Katz, S., and Tal, A., 2003, “Hierarchical Mesh Decomposition Using Fuzzy Clustering and Cuts,” *ACM Trans. Graphics*, **22**(3), pp. 954–961.
- [21] Garland, M., Willmott, A., and Heckbert, P., 2001, “Hierarchical Face Clustering on Polygonal Surfaces,” in *Proc. of ACM Symp. on Interactive 3D Graphics*, March, Raleigh, NC.
- [22] Sander, P. V., Snyder, J., Gortler, S. J., and Hoppe, H., 2001, “Texture Mapping Progressive Meshes,” in *Proceedings of SIGGRAPH 01*, August 12–17, Los Angeles, CA, pp. 409–416.
- [23] Lévy, B., Petitjean, S., Ray, N., and Maillot, J., 2002, “Least Squares Conformal Maps for Automatic Texture Atlas Generation,” *ACM TOG SIGGRAPH*, July 23–26, San Antonio, TX, pp. 362–371.
- [24] Cohen-Steiner, D., Alliez, P., and Desbrun, M., 2004, “Variational Shape Approximation,” *ACM Trans. Graphics*, **23**(3), pp. 905–914.
- [25] Surazhsky, V., Alliez, P., and Gotsman, C., 2003, “Isotropic Remeshing of Surfaces: A Local Parameterization Approach,” in *Proc. 12th Intl. Meshing Roundtable*, September 14–17, Santa Fe, NM.
- [26] Wasserman, J., 2003, *Michelangelo's Florence Pieta*, Princeton U. P., Princeton, NJ.
- [27] Stumpfel and Ch. Tchou, J., Yun, N., Martinez, P., Hawkins, T., Jones, A., Emerson, B., and Debev, P., 2003, “Digital Reunification of the Parthenon and its Sculptures,” in *4th International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage*, November 5–7, Brighton, United Kingdom, pp. 41–50.
- [28] Blais, F., 2004, “Review of 20 years of range sensor development,” *J. Electron. Imaging*, **13**(1), pp. 231–240.
- [29] Bernardini, F., Martin, I. M., and Rushmeier, H., 2001, “High-Quality Texture Reconstruction From Multiple Scans,” *IEEE Trans. Vis. Comput. Graph.*, **7**(4), pp. 318–332.
- [30] Hendrik, Lensch, P. A., Kautz, J., Goesele, M., Heidrich, W., and Seidel, H.-P., 2003, “Image-Based Reconstruction of Spatial Appearance and Geometric Detail,” *ACM Trans. Graphics*, **22**(2), pp. 234–257.
- [31] Sorkine, O., 2005, Eurographics'05 State-of-the-Art Report.
- [32] Zorin, D., Schröder, P., DeRose, T., Kobbelt, L., Levin, A., and Sweldens, W., 2000, SIGGRAPH'00 Course Notes.
- [33] Du, Q., Gunzburger, M. D., and Ju, L., 2003, “Constrained Centroidal Voronoi Tessellations for Surfaces,” *SIAM J. Sci. Comput. (USA)*, **24**(5), 1488–1506.
- [34] Du, Q., Faber, V., and Gunzburger, M., 1999, “Centroidal Voronoi Tessellations: Applications and Algorithms,” *SIAM Rev.*, **41**, pp. 637–676.
- [35] Catmull, E., and Clark, J., 1978, “Recursively Generated B-spline Surfaces on Arbitrary Topological Meshes,” *CAD*, **10**(6), pp. 350–355.
- [36] Biermann, H., Martin, I., Bernardini, F., and Zorin, D., 2002, “Cut-and-paste Editing of Multiresolution Surfaces,” in *Proc. SIGGRAPH 02*, July 23–26, San Antonio, TX, pp. 312–321.
- [37] Biermann, H., Martin, I., Zorin, D., and Bernardini, F., 2001, “Sharp Features on Multiresolution Subdivision Surfaces,” in *Proceedings of Pacific Graphics 01*, October 16–18, Tokyo, Japan, IEEE, New York.
- [38] Boier-Martin, I., Ronfard, R., and Bernardini, F., 2004, “Detail-Preserving Variational Surface Design With Multiresolution Constraints,” in *Proc. SMI 04*, June 7–9, Genova, Italy, pp. 119–128.
- [39] McQueen, J., 1967, “Some Methods for Classification and Analysis of Multivariate Observations,” *5th Berkeley Symp. Math. Stat. Prob.*, edited by L. Le Cam and J. Neyman, Univ. of California, June 21–July 18, 1965 and December 21, 1965–January 7, 1966, Univ. of California Press, Vol. 1, pp. 281–297.
- [40] Alliez, P., de Verdière, E. C., Devillers, O., and Isenburg, M., 2003, “Isotropic Surface Remeshing,” in *Proc. SMI03*, May 12–16, Seoul, Korea, pp. 49–58.
- [41] Amenta, N., Bern, M., and Kamvysselis, M., 1998, “A New Voronoi-Based Surface Reconstruction Algorithm,” *Proc. SIGGRAPH 98*, July 19–24, Orlando, FL, pp. 415–421.
- [42] Müller-Hannemann, M., and Weihe, K., 1997, “Minimum Strictly Convex Quadrangulations of Convex Polygons,” in *Symp. Comp. Geom.*, June 4–6, Nice, France, pp. 193–202.
- [43] Zorin, D., Schröder, P., and Sweldens, W., 1997, “Interactive Multiresolution Mesh Editing,” *Proceedings of SIGGRAPH 97*, August, Los Angeles, CA, pp. 259–268.
- [44] Osada, R., Funkhouser, T., Chazelle, B., and Dobkin, D., 2002, “Shape Distributions,” *ACM Trans. Graphics*, **21**(4), pp. 807–832.
- [45] Cignoni, P., Rocchini, C., and Scopigno, R., 1998, “Metro: Measuring Error on Simplified Surfaces,” *Comput. Graph. Forum*, **17**(2), pp. 167–174.
- [46] Ciarlet, P. G., 1991, “Basic Error Estimates for Elliptic Problems,” *Handbook of Numerical Analysis*, North-Holland, Amsterdam.
- [47] Hormann, K. and Greiner, G., 2000, “MIPS: An Efficient Global Parameterization Method,” in *Curve and Surface Design: Saint-Malo 1999 (Innovations in Applied Mathematics)*, edited by P.-J. Laurent, P. Sablonniere, and L. L. Schumaker, Vanderbilt University Press, pp. 153–162.