

Anisotropic Quadrangulation

Denis Kovacs
New York University
kovacs@cs.nyu.edu

Ashish Myles
New York University
amyles@cs.nyu.edu

Denis Zorin
New York University
dzorin@cs.nyu.edu

ABSTRACT

Quadrangulation methods aim to approximate surfaces by semi-regular meshes with as few extraordinary vertices as possible. A number of techniques use the harmonic parameterization to keep quads close to squares, or fit parametrization gradients to align quads to features. Both types of techniques create near-isotropic quads; feature-aligned quadrangulation algorithms reduce the remeshing error by aligning isotropic quads with principal curvature directions. A complementary approach is to allow for anisotropic elements, which are well-known to have significantly better approximation quality.

In this work we present a simple and efficient technique to add curvature-dependent anisotropy to harmonic and feature-aligned parameterization and improve the approximation error of the quadrangulations. We use a metric derived from the shape operator which results in a more uniform error distribution, decreasing the error near features.

Categories and Subject Descriptors

I.3.5 [Computational Geometry and Object Modeling]: Geometric algorithms, languages, and systems

Keywords

parameterization, quadrangulation, remeshing, conformal parameterization.

1. INTRODUCTION

Most common techniques for generating meshes from range scans and volumetric data produce irregular meshes with complex connectivity. A surface can be stored in a much more compact form, simplifying and speeding up rendering and processing if it is converted to a predominantly regular mesh, with only a small number of irregular vertices and faces. It is desirable to minimize the number of vertices in the semiregular mesh, while keeping it close to the original mesh.

Recent quadrangulation algorithms use a *global parameterization* of a mesh; the new mesh is obtained using a regular sampling pattern in the plane. Quite often, the parameterization is optimized

to be as isometric possible. However, isometric parameterizations may be far from optimal for surface remeshing, if the goal is to obtain a surface as close as possible to the original for a given number of faces. For example, a cylinder can be mapped isometrically to the plane, resulting in a uniform sampling pattern on the surface. It can, however, also be meshed with single long quads stretched along the axial direction, with the same approximation error. We call quadrangulations that adapt the quad aspect ratio to the surface shape *anisotropic*. We present a *simple* and *robust* method for computing anisotropic quadrangulations with quad aspect ratios




Figure 1: Quadrangulations of a lion head model. Top: the original model; middle: isotropic feature-aligned quadrangulation (25% reduced) bottom: anisotropic feature-aligned quadrangulation

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SPM 2010 Haifa, Israel

Copyright 200X ACM X-XXXXX-XX-X/XX/XX ...\$10.00.

adapted to local curvature, obtaining a good surface approximation with fewer quads.

Our method utilizes a curvature-based surface metric and computes the parameterization using this metric, rather than the Euclidean metric. Our approach is compatible with most parameterization methods that only rely on intrinsic quantities and vector fields on the surface.

Defining a metric for meshes is conceptually simple: we assign a new length to each edge. However, each edge length has to satisfy local triangle inequality constraints. It is a surprisingly difficult task to ensure that no inequality is violated, and while it may still be possible to compute a parameterization, the results may not have the desired anisotropic behavior (Section 5). We solve this problem using the idea of a high-dimensional embedding [20, 4]: the Euclidean metric in the higher-dimensional space defines the new edge lengths for the mesh. The embedded vertex coordinates consist of the original positional and normal coordinates, making the new edge length computation straightforward.

2. RELATED WORK

The literature on parameterization, remeshing and quadrangulation is vast; [20], [4] and [9] are the most closely related to our work. Our key observation is that the high-dimensional embedding proposed in [4] to obtain anisotropic quadrangulations with the quad aspect ratio determined by the ratio of principal curvatures can be applied in the context of a particular class of parameterization techniques, and yields robust results while preserving fine surface features.

Many recent quadrangulation methods (in contrast to the work based on the construction of base complexes by simplification [12, 18, 17, 8]) have similar structure: a global parameterization is obtained by solving equations for gradients of parametric functions, and a new mesh is generated by following parametric lines. The two main categories of methods of this type are harmonic and feature-aligned.

Harmonic and conformal methods (for brevity we will refer to both as harmonic) are robust, efficient and typically produce good results even for complex meshes for a suitable choice of singularities and boundary conditions. Some quadrangulation methods use harmonic maps directly [10, 29]. These methods can be viewed as minimizing nonconformality of the map, while allowing significant area scaling; nonlinear methods such as [25, 27] are needed to guarantee a one-to-one parameterization. Extreme area distortion is reduced by adding singularities (or “cones”) to the parameterization, with several methods for automatic placement of singularities proposed in [10, 1, 27]. These techniques allow explicit user control over the number of irregular points on the mesh. The downside of harmonic techniques, especially in the context of remeshing, is that non-intrinsic shape information is not used directly.

The shape information can be taken into account in two distinct ways to minimize the approximation error. Locally, a smooth shape can be characterized by its shape operator. Figure 2 show two ways of taking the shape operator into account (with principal curvature directions scaled by inverse principal curvatures shown in red).

A “perfect” quad of a given area approximating a surface is *aligned*, i.e., has edges parallel to principal curvature directions




Figure 2: Quad alignment and anisotropy

and *anisotropic* i.e., has aspect ratio inversely proportional to the ratio of principal curvatures. This corresponds to two classes of feature-aware parameterization techniques.

Feature-alignment methods [22, 15, 2] adapt the parameterization to the shape by aligning new mesh elements with a feature field, typically derived from the principal curvature direction field, either by smoothing, or interpolation of salient features. The singularities of the parameterization are determined by the singularities of the field, so the feature field cannot match the actual curvature field too closely: substantial smoothing is needed to keep the number of singularities small. The shape of the quads generated by these techniques tends to be uniform, rather than anisotropic: one can view these techniques as minimizing non-isometry, while aligning with the feature field. [2] permits a degree of anisotropy, penalizing changes in length less than changes in the direction, but without relating these to curvature.

In geometric modeling, *anisotropic parameterization* was introduced as signal-specialized parameterization [24, 28]. This work uses a metric derived from the Hessian of the signal to adapt the parameterization to a signal defined on the surface; in particular, the surface itself can be used as the signal. Zayer et al. [32, 30, 31] describe a general class of parameterization methods based on solving a generalized Laplace or Poisson equation using a tensor field, which can be interpreted as a metric tensor. An elegant formulation for related quasi-conformal maps based on Beltrami factors described in [33]. The interpolation and stiffness properties of anisotropic linear triangles in finite-element context are discussed in detail in [26]. [5] derives bounds on the Hausdorff-distance approximation of manifolds using a metric closely related to the one that we use.

We show how to use a metric defined on a surface to obtain anisotropic versions of global quadrangulation algorithms, both harmonic and feature-aligned, and demonstrate the improvements in surface approximation that can be obtained in this way. To the best of our knowledge, metric-based techniques were not yet applied to quadrangulation, although [29] suggests that this is possible by altering the Laplace equation coefficients without suggesting a specific way to compute the metric.

We emphasize that we view using anisotropic metric as complementary to curvature-alignment approaches, rather than alternative to these. Curvature-alignment methods allow to obtain a geometrically meaningful set of singularities and coarse alignment with the shape; Anisotropy helps to resolve sharp features locally with fewer vertices, and allows to keep the number of parameterization singularities low.

3. ANISOTROPIC METRIC

The main idea of our approach is to define a new metric (that is, new edge lengths) on a mesh, and use an isometry-approximating parameterization based on these edge lengths for quadrangulation. The discrete metric is given by Equation (5). Our goal in this section is to explain the motivation for this choice. First, we discuss the *local error* and the choice of the best approximating quad; under the assumptions that we make, and similarly to previous work, the optimal quad is aligned with principal curvature directions, and has aspect ratio proportional to the ratio of principal curvatures.




Figure 3: Notation

Second, we discuss how local errors can be combined together to obtain equations for the parameterization of the whole surface. We show that isometry in the shape-operator corresponds to optimal *equidistributed* error.

Definitions. Important local properties of a parameterization are captured by the *metric tensor*. Suppose a surface A is defined by a function $\mathbf{f} : \mathbf{R}^2 \rightarrow A$ (Figure 3). A *surface parameterization* is the inverse map from the surface to the plane $\mathbf{g} : A \rightarrow \mathbf{R}^2$. In our exposition, it is convenient to fix a surface point \mathbf{p} and the tangent plane P at this point. In this case, $\nabla \mathbf{f}$ is a linear map from the parametric plane to P .

Recall that the metric tensor of \mathbf{f} is defined by

$$M(\mathbf{f}) = (\nabla \mathbf{f})^T \nabla \mathbf{f} \quad (1)$$

and is given by a 2×2 symmetric positive-definite matrix. For a vector $\mathbf{v} = \mathbf{q}_2 - \mathbf{q}_1$ in the parametric domain defined by a pair of close points \mathbf{q}_1 and \mathbf{q}_2 , the quadratic form $\mathbf{v}^T M(\mathbf{f}) \mathbf{v}$ is, in the limit, the squared length of the image of \mathbf{v} : $|\mathbf{f}(\mathbf{q}_2) - \mathbf{f}(\mathbf{q}_1)|^2$.

3.1 Normal approximation error and metric

The local normal approximation error measure (e.g., [6]) is similar to the gradient error measure in finite elements [9]. This error corresponds more closely to the perceived visual quality of an approximation, compared to, for example, the distance between points on the surface. For the purposes of defining a pointwise error, we consider an idealized setting: (1) The surface has well-defined curvature, with nonvanishing Gaussian curvature. (2) For a parameterization \mathbf{g} , we consider the approximation of the surface by a collection of small quads. Each quad Q is a parallelogram obtained by mapping a square Q_p of edge length h from a regular grid in the plane to the tangent plane of the surface at a point $\mathbf{g}^{-1}(\mathbf{c}) = \mathbf{f}(\mathbf{c})$, using $\nabla \mathbf{f}$. (3) We assume the surface to be well-approximated by a quadratic function over the tangent plane over each quad.

We define the error for a quad Q in the tangent plane P with normal \mathbf{n}_Q as the square of the average of the deviation of the normal on the part of the surface $A(Q)$ projected to the quad Q along \mathbf{n}_Q .

$$E_Q^2 = \frac{1}{\text{Area}(Q)} \int_{A(Q)} \|\mathbf{n}_S(\mathbf{q}) - \mathbf{n}_Q\|^2 d\mathbf{q} \quad (2)$$

It can be shown (appendix A) that up to high-order terms,

$$E_Q^2 = \frac{h^4}{12} \text{Tr} (S^2 M(\mathbf{g})^{-1}), \quad (3)$$

where S denotes the shape operator.

E_Q is highly similar to the gradient interpolation error for linear elements [26], yet there is an important distinction. As discussed in [26, 3], that error has a strong dependence on the shape of the element in the physical space (in our case, the shape of the approximating quad).

Specifically, if a square is mapped to the tangent plane using a map \mathbf{f} with metric S^{-2} , and the edges of the quad form a large angle in the tangent plane, the error, instead of being independent of curvature as suggested by (3) and (4) may be of order ah^2 , where a is the ratio of max to min curvature; so the error distribution over the surface is clearly nonuniform. The fact that the quads we consider are tangent to the surface changes this behavior. However, in this work we are primarily concerned with the case when arbitrary anisotropy is not allowed. Rather we limit it to moderate values (typically no more than 5). We also note that under our assumptions, differing from those in, e.g., [9], the error is the same for hyperbolic and elliptic points with identical principal curvatures. If the vertices of




Figure 4: Top left to right: a conformal map, a map with a small amount of anisotropy added ($\alpha = 3$), and large amount of anisotropy ($\alpha = 0.1$), where the metric tensor for the parameterization is $\alpha^2 I + S^2$. Bottom left to right: corresponding uv maps color-coded by inverse parametric triangle area.

quads are expected to interpolate the surface, optimality conditions in the hyperbolic case are different.

Uniform-error parameterization and shape operator metric.

A natural approach to define an optimal parameterization given a pointwise local error is to require the error to have the same value ϵ over the whole surface, and minimize ϵ . (This is however distinct from many methods that define a global energy as an integral measure of a local error over the surface). Integrating the local error E_Q over the surface results in difficult-to-solve equations. *Equalizing* the error, however, leads to a simple condition on the error (appendix A): if we assume that the metric tensor $M(\mathbf{g})$ can be chosen arbitrarily, and the total parametric area of the surface A_{tot} is fixed, then the standard Lagrange multiplier solution of this constrained minimization problem is given by

$$M(\mathbf{g}) = cS^2, \quad (4)$$

with c independent of the point (the actual value of c is irrelevant for our purposes). In particular, the error bound is the same (under restrictive assumptions outlined below) for all parameterization differing by a rotation of the parametric plane (Figure 5).

In general, S^2 may have small or zero eigenvalues, and using it alone as a metric is not desirable, as this would result in infinitely long or thin quads. We can limit the possible quad aspect ratios by using $G(\alpha) = \alpha^2 I + S^2$ as the metric.




Figure 5: The right model shows the result of rotating the anisotropic parameterization 45 degrees. Observe that the mesh elements remain stretched along the features.

We conclude that *A uniform normal error parameterization of a surface with nonzero Gaussian curvature has a metric tensor coinciding with the square of the inverse of the shape operator up to a*

globally constant scale factor, in other words, it is isometric in the metric defined by the inverse of the shape operator.

4. ANISOTROPIC PARAMETERIZATION

The observation of the previous section reduces the problem of finding an equidistributed error parameterization to that of finding an isometric parameterization in a metric. Most currently used techniques can be regarded as approximations to the isometric parameterization in Euclidean metric, and can be naturally generalized if the shape-operator metric can be computed robustly and accurately, as discussed in Section 5.

We present anisotropic extensions for two parameterization techniques, harmonic, following [27] and feature-aligned, following [2]. As we have discussed in Section 2, the advantage of the former is more direct and explicit control over the number of singularities, while the latter yields parameterizations better aligned with mesh features, and typically closer to isometric.

We regard both harmonic parameterization and feature-aligned parameterization as two types of efficient approximations to isometric maps (in the latter case with additional condition of feature alignment) and demonstrate how these can be combined with anisotropy.

Isometric parameterization and harmonic maps. Isometric parameterizations do not exist for surfaces with nonzero Gaussian curvature: at best, we can hope to approximate an isometric parameterization. Minimizing the deviation of the metric tensor from identity leads to nonlinear systems of equations for which no robust and efficient solvers are available. For this reason, many techniques replace direct isometry optimization with various types of factorizations.

Most commonly, harmonic maps, leading to linear systems, are used to minimize the angle distortion, subject to the boundary conditions; harmonic parameterizations often result in high area distortion. The idea of a number of recent methods [14, 10, 1, 27] is to use harmonic maps with singularities to define a parameterization, and to reduce the area distortion by introducing singularities and optimizing the singularity placement.

Anisotropic harmonic maps. In case of isometry, conformal maps are defined by the condition $M(\mathbf{g}) = cI$; they preserve the ratio of the singular values of the identity tensor I exactly. *Anisotropic conformal maps* satisfying $M(\mathbf{g}) = cG(\alpha)$ have similar behavior in the shape operator metric. Intuitively, an anisotropic conformal map takes a small circle in the parametric plane to an ellipse in the tangent plane of the surface, with axes aligned with the principal curvature directions, and its aspect ratio is determined by the ratio of principal curvatures. The effects of such a map, compared to a conformal map, are illustrated in Figure 4. The anisotropic harmonic map is a least-squares approximation to the anisotropic conformal map.

Anisotropic feature-aligned maps. Feature-aligned maps [15, 2] use a feature cross-field, which locally can be regarded as a pair of orthogonal unit vectors (\mathbf{u}, \mathbf{v}) to define the target directions for the surface gradients of parametric coordinates ∇u and ∇v . If desired gradient directions for coordinate functions are fixed, finding the as-isometric-as-possible parameterization can be formulated as a linear optimization problem minimizing misalignment with the feature field and deviation of the gradient magnitude from the unit length:

$$E = \int_A (\nabla u - \mathbf{u})^2 + (\nabla v - \mathbf{v})^2 dA$$

As \mathbf{u} and \mathbf{v} are orthogonal, perfect minimization of this energy corresponds to an isometric parameterization.

To obtain the *anisotropic feature-aligned parameterization*, we replace the original surface A with \tilde{A} , which is the same surface but with the metric redefined to be G (the ways to define \tilde{A} practically are discussed in Section 5). Then we remap the feature field from A to \tilde{A} and compute a feature-aligned least-squares isometric parameterization of \tilde{A} .

Remapping the cross-fields. Conceptually, parametrizing the surface \tilde{A} embedded in six dimensions is not different from parametrizing a surface in three dimensions. One could remap the salient points on A to \tilde{A} , using the natural map $\mathbf{p} \rightarrow (\mathbf{p}, \mathbf{n})$, and then compute the feature cross-field directly on \tilde{A} . However, in practice we observe that the surface \tilde{A} is much “bumpier”, i.e., has greater oscillations of the Gaussian curvature, due to higher variation of the shape operator included in the metric. The cross-field optimization procedure of [2] tends to place cones at Gaussian curvature extrema, which results in large numbers of cones. Instead, we perform cross field optimization in three dimensions as before, and remap the resulting cross field to \tilde{A} .

To remap the cross-field, we regard it as a bilinear form on tangent vectors. If for a triangle T , the linear transform from T to \tilde{T} in some two-dimensional local coordinate systems is C , and the principal directions of the cross-field are \mathbf{u} and $\mathbf{v} = \mathbf{u}^\perp$, the cross-field, can be associated with a unique (up to a sign) tensor B with singular values -1 and 1 , and singular directions \mathbf{v} and \mathbf{u} $B = \mathbf{u} \otimes \mathbf{u} - \mathbf{v} \otimes \mathbf{v}$. We obtain the new cross-field by transforming B using the standard transformation for symmetric tensors

$$\tilde{B} = C^T B C,$$

and obtain the cross-field directions $\pm \tilde{\mathbf{u}}$ and $\pm \tilde{\mathbf{v}}$ as singular directions of this tensor. One can show that as long as C is nonsingular, there is a one-to-one correspondence between these four directions and the original cross-field directions¹

5. DISCRETE METRIC

To complete our construction, it remains to define a discrete metric $G(\alpha)$ by assigning new lengths to each edge (5). While a variety of techniques can be used, we found that the results can be quite sensitive to the choice of technique.

The most obvious approach is to estimate the shape operator S either per vertex [7, 21, 16] or per triangle face [23, 13]. Using these estimates of the shape operator makes it possible to use the quasi-conformal map setting of [30] to compute the parameterization.

However, it proves to be fundamentally difficult to transition from a metric tensor defined in either way to a discrete metric defined per edge length, which satisfies the triangle inequality for general meshes. The reason for this can be seen from Figure 6. Suppose a triangle has bad alignment (long edge along principal direction with larger curvature). If the metric length of each edge e is determined as the average of two lengths $(\sqrt{e^T M_1 e} + \sqrt{e^T M_2 e})/2$, and the singular values of M_i are 1 and k^2 , except M_3 for which they are 1 and $(1+a)k^2$, then for large l , a can be at most $4/(lk)$ before the triangle inequality is violated. So any averaging method is likely to fail even for small curvature variation: for $k = 10$ and $l = 10$, for instance, only 4% variation is possible across an edge.

¹In [19], it was observed that cross-fields are most naturally interpreted as symmetric 4-tensors, rather than 2-tensors, which eliminates the sign ambiguity; one could use this correspondence to remap the fields.

An attractive alternative is to define an embedding of the surface such that the Euclidean metric on the surface for this embedding yields an approximation to the desired metric tensor [3]. For the shape operator, the relevant embedding is the *Gauss map*: $f(\mathbf{p}) = \mathbf{n}(\mathbf{p}) \in \mathbf{R}^3$, because $S = \nabla \mathbf{n}^T$, i.e. S^2 is exactly the metric tensor of the Gauss map.

The shape operator satisfies $S\mathbf{v} = \nabla_{\mathbf{v}}\mathbf{n}$ for a tangent vector \mathbf{v} . Applied to edge vectors $\mathbf{e}_{ij} = p_i - p_j$ on a triangle mesh, it can be discretized by $S\mathbf{e}_{ij} = \mathbf{n}_j - \mathbf{n}_i = \Delta \mathbf{n}_{ij}$. Then the squared shape operator metric S^2 is given by

$$\mathbf{e}_{ij}^T S^2 \mathbf{e}_{ij} = |\Delta \mathbf{n}_{ij}|^2$$

In other words, *the optimal metric edge length is simply the distance between endpoints of the edge in the Gauss map image of the mesh*. Note that so far, we only considered the embedding of the mesh into a two-dimensional sphere given by mapping each vertex v_i to its normal \mathbf{n}_i . This, however, is not sufficient to obtain the metric $G(\alpha) = \alpha^2 I + S^2$.

We therefore embed the mesh into \mathbf{R}^6 , with a vertex v_i mapped to the point $(\alpha \mathbf{p}_i, \mathbf{n}_i)$, where α is a scale factor controlling the aspect ratios. In this case, the Euclidean metric in \mathbf{R}^6 yields

$$\begin{aligned} l_{ij}^2 &= (\alpha \mathbf{p}_i - \alpha \mathbf{p}_j)^2 + (\mathbf{n}_i - \mathbf{n}_j)^2 \\ &= \mathbf{e}_{ij}^T (\alpha^2 I + S^2) \mathbf{e}_{ij} = \mathbf{e}_{ij}^T G \mathbf{e}_{ij} \end{aligned} \quad (5)$$

i.e., it corresponds to a linear combination of isometry and normal error metrics. This defines the metric tensor G in terms of metric edge lengths l_{ij} . Since every mesh triangle is embedded in Euclidean space, the metric edge lengths satisfy the triangle inequality *by construction*.

Controlling aspect ratios. The parameter α can be used to control the maximal distortion either globally or locally. We found that the method is stable even for very small values of α , which allow quads to stretch a lot. The singular values of the tensor are $\alpha + \kappa_1^2$ and $\alpha + \kappa_2^2$, and the aspect ratio of the images of infinitesimal quads is $\sqrt{(\alpha + \kappa_1^2)/(\alpha + \kappa_2^2)}$, where we assume $|\kappa_1| > |\kappa_2|$. By choosing

$$\alpha = \sqrt{\frac{r \max \kappa^2 - \min \kappa^2}{r - 1}} \quad (6)$$

globally, we can keep the aspect ratio below r . This is, however, a very conservative choice, which may eliminate the advantages of the method for surfaces with very nonuniform curvature.

6. IMPLEMENTATION

The idea of using a shape-operator metric can be integrated with any quadrangulation approach that only relies on the surface metric: the main change required is to modify the metric-dependent quantities to use (5); for methods using vector or tensor fields on surfaces, these need to be remapped accordingly.

The details of both harmonic and feature-aligned mixed-integer parameterization can be found in [1, 27] and [15, 2] respectively. Here we present only a brief overview, to point out the aspect of algorithms that were modified.




Figure 6: A triangle with aspect ratio l , with 3 metric tensors M_i at vertices

For both methods, we start with computing a normal field (we use the robust method of [16]) and compute and smooth the scaling function α , followed by evaluating the metric lengths l_{ij} using (5). Once the global parameterization is computed, we generate a quad mesh by tracing parametric lines $u = i$, and $v = i$ where i is an integer, and determine quad vertex positions at integer u/v locations by linearly interpolating the original mesh vertices.

Harmonic parameterization. The main steps in this case are: – iteratively optimize cone locations solving the Laplace equation for the scale factors using metric edge length l_{ij} , or specify singularity locations manually; – cut the mesh into a disk; – quantize singularity indices to $k\pi/2$ (if not specified by hand), and singularity positions to integer locations; – use harmonic parameterization with cotangent weights computed from l_{ij} to obtain a global mesh parameterization matching across the seams of the cut. The main distinction compared to the original method is computing all metric quantities (cotangent weights in particular) using lengths obtained in (5).

Feature-aligned parameterization. In this case, we start with constructing the 3D feature cross-field: – identify the *salient* locations on the mesh for which the cross-field directions are fixed; – compute a global smooth feature cross-field using the quadratic mixed integer optimization of [2]; – detected singularities, and cut the mesh into a disk, so that the cut passes through all singularities; – label globally consistent \mathbf{u} and \mathbf{v} directions on the cut mesh; – minimize the fit energy for parameterization gradients to \mathbf{u} and \mathbf{v} , enforcing constraints along the cuts and constraining the changes in coordinates across cuts to be integer.

The last step may be repeated multiple times with increasing weights in the energy to eliminate inverted triangles in the parameterization (stiffening).

For anisotropic feature-aligned parameterization, we remap the cross-fields on each triangle to the new metric to the using the approach described in Section 4. Each triangle T of the mesh for surface A corresponds to a triangle \tilde{T} on the mesh for surface \tilde{A} , with rescaled edge lengths \tilde{l}_{ij} . The linear transformation C is uniquely determined by the affine transformation mapping \tilde{T} to T .

7. RESULTS

Comparison of different metric discretizations. First, we demonstrate the robustness and feature sensitivity of our technique (Figure 7). We compare to an approach similar to that of [30], rewriting the shape-metric Dirichlet energy as $\int_A (\nabla_G u)^2 dA$, where the gradient ∇_G is computed using a per-triangle estimate of the shape operator; effectively this is similar to averaging metric lengths of an edge with respect to the operators on two incident triangles. This method results in significant smoothing of the metric, and, as a consequence, sharper features are not captured (Figure 7b.)

We attempt to set the scaled edge lengths again by averaging the lengths computed using per-vertex shape operators at two endpoints (Figure 7d,e). We observe that even for modest anisotropy, for a large number of facets the triangle inequality is violated; refining the mesh in most cases eliminates the triangle inequality violations, but a large number of iterations may be needed and resulting quadrangulation suffers from metric smoothing similar to the per-triangle case (Figure 7e).

Figure 5 shows the effects of rotating parametric axes for anisotropic




Figure 7: Comparison of different ways of specifying metric lengths (a) the original face mesh. (b) face-tensor-based (c) our method (d) vertex-tensor averaging, triangles not satisfying metric inequality (e) after refinement, metric inequality is satisfied, but quadrangulation misses some features.

harmonic parameterization of a shape which does not require adding cones or cuts. Note that the parameterization automatically squeezes quads to the lines of high curvature: the mesh elements appear to preserve their orientation, while rotating in the parametric domain.

Comparisons with isotropic quadrangulation. Our primary comparison is to the mixed-integer quadrangulation of [2] with no anisotropy. Figure 8, Figure 11, and Figure 12 show feature-aligned quadrangulations for a number of models.




Figure 8: Periodic global parameterization and (unaligned) harmonic anisotropic parameterization. Normal error distribution is shown in pseudocolor.

For two models, we also compare to the periodic global parameterization (PGP) (Figure 8 and Figure 9). We observe that under some conditions, *unaligned* anisotropic harmonic quadrangulation produces better results compared to aligned but isotropic quadrangulation.

As our main target application is approximating the original meshes with semiregular meshes with good visual quality, the ultimate criterion in this case (vs., for example, remeshing for finite element simulation) is the appearance of the resulting models. For this reason, we present smoothly shaded images of the remeshed models in Figure 12, along with a pseudocolor rendering of the pointwise normal error (dark red corresponds to maximal error, dark blue to no error). We choose relatively coarse quadrangulations to make the

errors more apparent. The number of facets in the original models, the number of quads as the fraction of the original model size, and the number of singularities are summarized in the following table.

model	facets	reduced to	cones
lion head	16674	17%	41
Julius	39168	28%	25
screwdriver	54300	3%	20
Stanford bunny	111364	3.5%	32
rocker arm	20088	8%	26
Omotondo	10000	25%	36
Max Planck	50790	35%	15

We emphasize that our technique aims to make the error distribution more uniform, *not* to minimize an integral error measure, hence it is difficult to quantify the relative quality of the result by a single number. In pseudocolor visualizations in Figure 12, one can observe greater uniformity in pointwise error. A consistent increase in uniformity is also confirmed by the plots of the pointwise error distribution: these plots show, for a given abscissa β , (in percent of the max possible error in normal), the fraction of vertices with error above β in log scale. Plots for anisotropic models are in red and for isotropic in blue. Higher slope corresponds to more even error distribution.

Sharp features. As Figure 9 demonstrates, anisotropic harmonic quadrangulation can handle models with sharp features robustly, even with no feature alignment. The mesh for the fan disk model has only 8 singularities, i.e., the whole surface is mapped to the surface of the cube. Although for noise-free the quality of the result is inferior to the one that can be obtained by explicitly constraining the parameterization to be aligned with sharp edges as described in [2], for scanned meshes similar to the screwdriver example (Figure 12) when the edges of the mesh are not aligned with sharp features of the underlying geometry.




Figure 9: Quadrangulation of a model with sharp features. From left to right: the original model, remeshing using PGP, remeshing using anisotropic harmonic map. Both remeshed models retain approximately 20% of faces of the original model. 8 singularities are used for the anisotropic map, i.e., the model is parametrized over the surface of a cube.

For certain types of models, it may be highly desirable to preserve sharp features. For feature-based parameterization, one can explicitly integrate perfectly sharp feature edges into the process, by forcing the field to be aligned with these edges and forcing one of parametric coordinates to be constant along these edges. This typically requires introducing a sufficient number of singular vertices.

In the context of our method, one can introduce parameterization discontinuities along sharp edges *without* introducing extraordinary vertices, at the expense of introducing collapsed quads on a regular mesh. Figure 10 shows a case where sharp features were tagged along the connection of the model to the plane, and degenerate triangles were inserted along these creases. For normal calculations the creases were treated as internal boundaries.

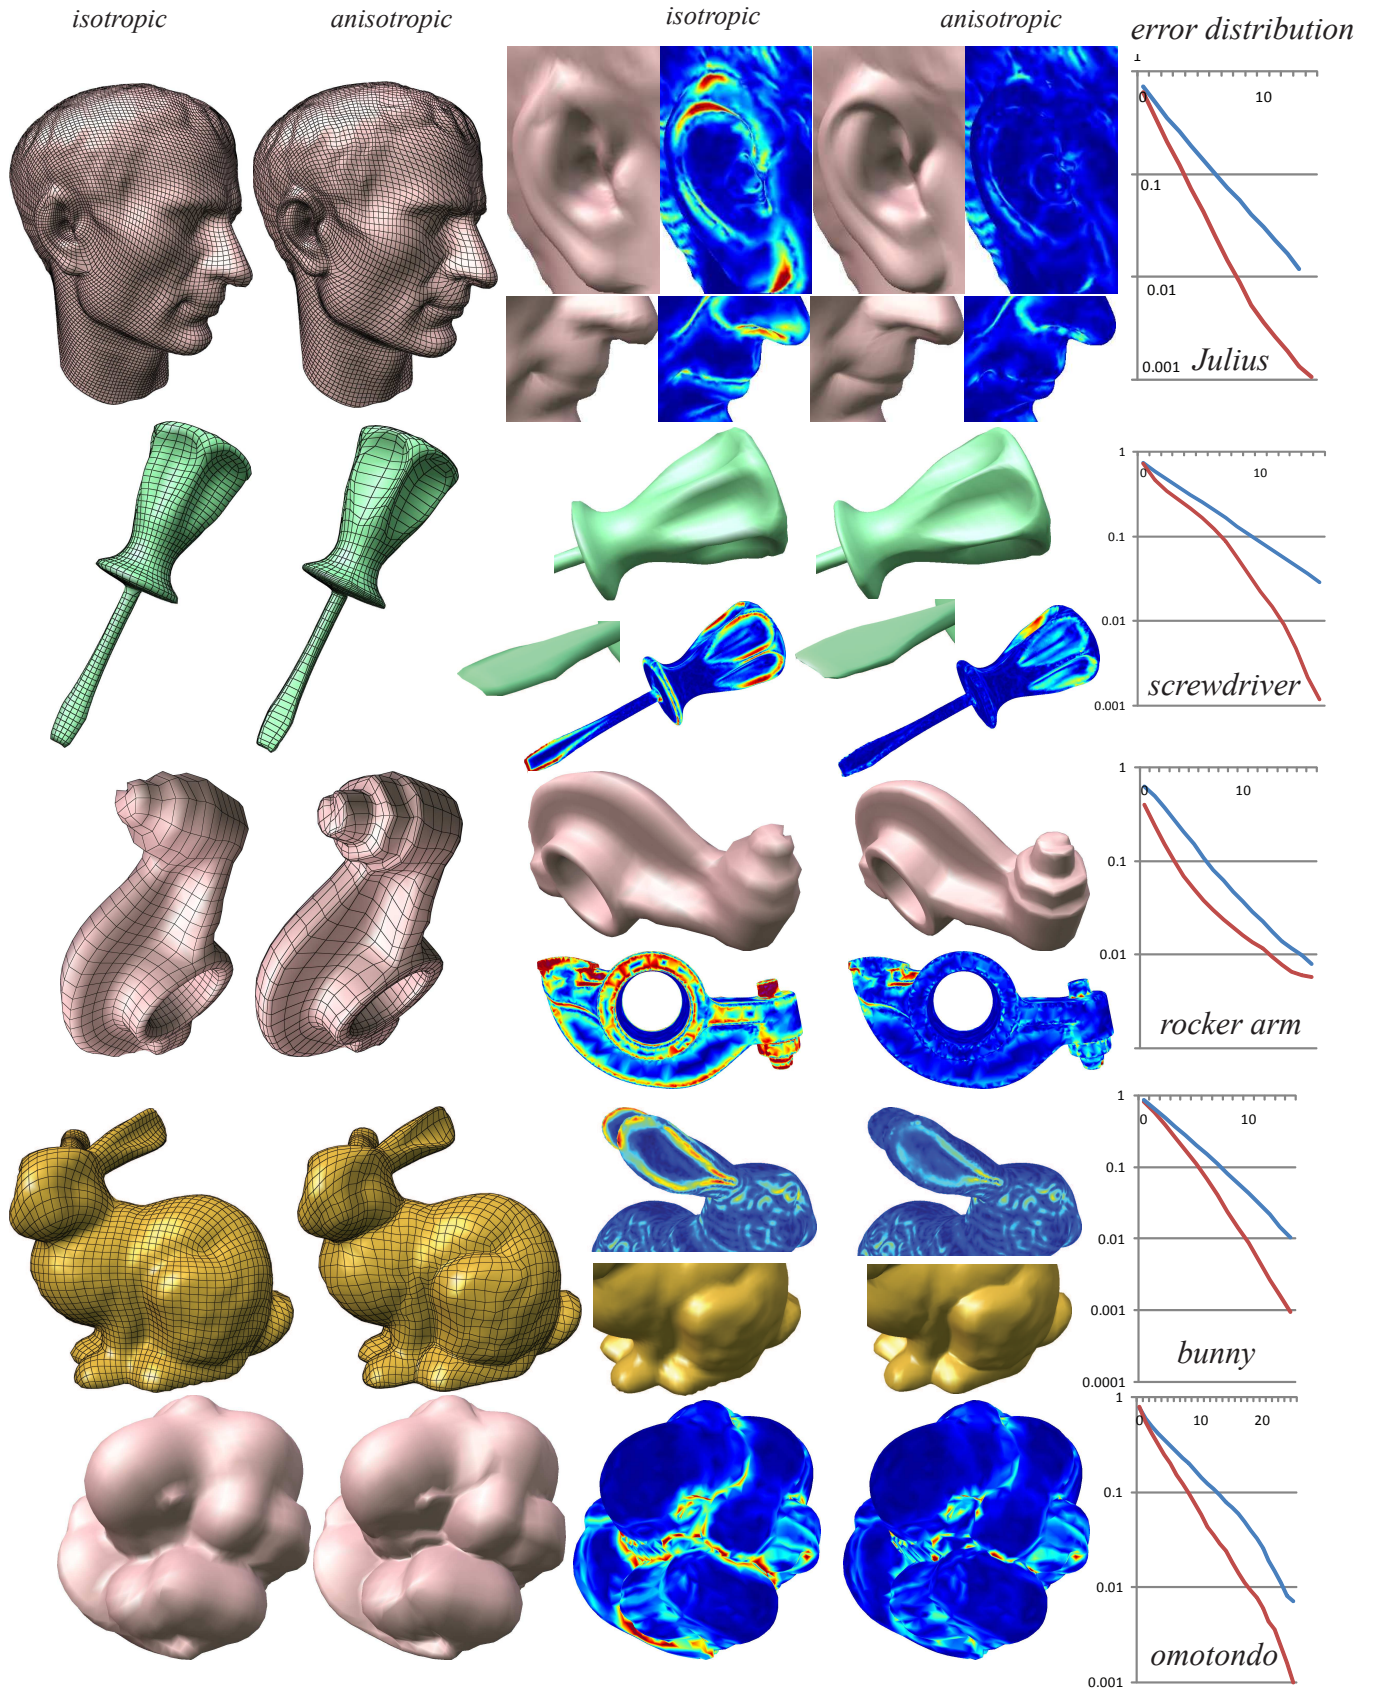


Figure 12: Isotropic and anisotropic feature-aligned quadrangulations and error visualization. Error plots show in log scale for a given abscissa β (in % of max. normal error) the fraction of vertices with error above β .




Figure 10: Quadrangulation of a model with sharp features, with additional edges inserted at creases; no singularities are used. Top: the original model, and our quadrangulation with 25% of faces. Bottom left: a harmonic map quadrangulation with the same number of faces; Bottom right: the original mesh in the parametric domain. Note the extremely stretched bands of triangles: these are thin triangles inserted along the sharp feature.

8. CONCLUSION

The most appealing features of the proposed method are its remarkable robustness (even for high aspect ratios, compared to other metric-based techniques we have experimented with), its simplicity and its compatibility with a number of other approaches.

As we generate quads with large angles, the resulting meshes are in general not suitable for solving equations on surfaces, unless the aspect ratio is limited to a moderate value; even with this restriction we can still expect a reduction in the number of quads needed for a given approximation quality.

While we do provide control over maximal aspect ratios, it is far from a complete solution, especially in cases of rapid edge length variation.

The method takes advantage of the possibility of discretizing the shape operator metric using a high-dimensional embedding. We would like to extend this to approximate embedding discretizations for arbitrary metric tensors.

9. REFERENCES

- [1] M. Ben-Chen, C. Gotsman, and G. Bunin. Conformal Flattening by Curvature Prescription and Metric Scaling. In *Computer Graphics Forum*, volume 27, pages 449–458. Blackwell Synergy, 2008.
- [2] D. Bommes, H. Zimmer, and L. Kobbelt. Mixed-integer quadrangulation. *ACM Transactions on Graphics (TOG)*, 28(3):77, 2009.
- [3] G. Cañas and S. Gortler. On Asymptotically Optimal Meshes by Coordinate Transformation. *Proceedings of 15th International Meshing Roundtable*, 2006.
- [4] G. Cañas and S. Gortler. Surface remeshing in arbitrary codimensions. *The Visual Computer*, 22(9):885–895, 2006.
- [5] K. Clarkson. Building triangulations using ε -nets. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 326–335. ACM New York, NY, USA, 2006.
- [6] D. Cohen-Steiner, P. Alliez, and M. Desbrun. Variational shape approximation. *ACM Transactions on Graphics (TOG)*, 23(3):905–914, 2004.
- [7] D. Cohen-Steiner and J. Morvan. Restricted delaunay triangulations and normal cycle. In *Proceedings of the nineteenth annual symposium on Computational geometry*, pages 312–321. ACM New York, NY, USA, 2003.
- [8] J. Daniels, C. Silva, and E. Cohen. Semi-regular Quadrilateral-only Remeshing from Simplified Base Domains. In *Computer Graphics Forum*, volume 28, pages 1427–1435. Blackwell Publishing Ltd, 2009.
- [9] E. D’Azevedo and R. Simpson. On optimal triangular meshes for minimizing the gradient error. *Numerische Mathematik*, 59(1):321–348, 1991.
- [10] S. Dong, P. Bremer, M. Garland, V. Pascucci, and J. Hart. Spectral surface quadrangulation. *ACM Transactions on Graphics (TOG)*, 25(3):1057–1066, 2006.
- [11] Q. Du and D. Wang. Anisotropic centroidal voronoi tessellations and their applications. *SIAM J. Sci. Comput.*, 26(3):737–761, 2005.
- [12] M. Eck, T. DeRose, T. Duchamp, H. Hoppe, M. Lounsbery, and W. Stuetzle. Multiresolution analysis of arbitrary meshes. *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pages 173–182, 1995.
- [13] E. Grinspun, Y. Gingold, J. Reisman, and D. Zorin. Computing discrete shape operators on general meshes. In *Computer Graphics Forum*, volume 25, pages 547–556. Blackwell Synergy, 2006.
- [14] X. Gu and S. Yau. Global conformal surface parameterization. *Proceedings of the 2003 Eurographics/ACM SIGGRAPH symposium on Geometry processing*, pages 127–137, 2003.
- [15] F. Kälberer, M. Nieser, and K. Polthier. QuadCover: Surface Parameterization using Branched Coverings. *Computer Graphics Forum*, 26(3):375–384, 2007.
- [16] E. Kalogerakis, P. Simari, D. Nowrouzezahrai, and K. Singh.




Figure 11: Anisotropic quadrangulation preserves essential features even for extreme simplification (3%) The lower row shows the normal error distribution.

- Robust statistical estimation of curvature on discretized surfaces. In *Proceedings of the fifth Eurographics symposium on Geometry processing*, pages 13–22. Eurographics Association Aire-la-Ville, Switzerland, Switzerland, 2007.
- [17] A. Khodakovsky, N. Litke, and P. Schröder. Globally smooth parameterizations with low distortion. *ACM Transactions on Graphics (TOG)*, 22(3):350–357, 2003.
- [18] A. Lee, W. Sweldens, P. Schröder, L. Cowsar, and D. Dobkin. MAPS: multiresolution adaptive parameterization of surfaces. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, pages 95–104. ACM New York, NY, USA, 1998.
- [19] J. Palacios and E. Zhang. Rotational symmetry field design on surfaces. *ACM Transactions on Graphics (TOG)*, 26(3):55, 2007.
- [20] H. Pottmann, T. Steiner, M. Hofer, C. Haider, and A. Hanbury. The isophotic metric and its application to feature sensitive morphology on surfaces. *Lecture Notes in Computer Science*, pages 560–572, 2004.
- [21] H. Pottmann, J. Wallner, Q. Huang, and Y. Yang. Integral invariants for robust geometry processing. *Computer Aided Geometric Design*, 2008.
- [22] N. Ray, W. Li, B. Lévy, A. Sheffer, and P. Alliez. Periodic global parameterization. *ACM Transactions on Graphics (TOG)*, 25(4):1460–1485, 2006.
- [23] S. Rusinkiewicz. Estimating curvatures and their derivatives on triangle meshes. In *3D Data Processing, Visualization and Transmission, 2004. 3DPVT 2004. Proceedings. 2nd International Symposium on*, pages 486–493, 2004.
- [24] P. Sander, S. Gortler, J. Snyder, and H. Hoppe. Signal-specialized parameterization. *Eurographics Workshop on Rendering*, 2002, 2002.
- [25] A. Sheffer and E. de Sturler. Parameterization of Faceted Surfaces for Meshing using Angle-Based Flattening. *Engineering with Computers*, 17(3):326–337, 2001.
- [26] J. Shewchuk. What is a good linear element? interpolation, conditioning, and quality measures. *11th International Meshing Roundtable*, pages 115–126, 2002.
- [27] B. Springborn, P. Schröder, and U. Pinkall. Conformal equivalence of triangle meshes. In *SIGGRAPH '08: ACM SIGGRAPH 2008 papers*, pages 1–11, New York, NY, USA, 2008. ACM.
- [28] G. Tewari, J. Snyder, P. Sander, S. Gortler, and H. Hoppe. Signal-specialized parameterization for piecewise linear reconstruction. *Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing*, pages 55–64, 2004.
- [29] Y. Tong, P. Alliez, D. Cohen-Steiner, and M. Desbrun. Designing quadrangulations with discrete harmonic forms. *Symposium on Geometry Processing*, pages 201–210, 2006.
- [30] R. Zayer, C. Rössl, and H. Seidel. Discrete Tensorial Quasi-Harmonic Maps. *Proceedings of Shape Modeling and Applications*, pages 276–285, 2005.
- [31] R. Zayer, C. Rössl, and H. Seidel. Setting the boundary free: A composite approach to surface parameterization. 2005.
- [32] R. Zayer, C. Rössl, and H.-P. Seidel. r-Adaptive parameterization of surfaces. Technical report, 2004.
- [33] W. Zeng, F. L. 0002, S.-T. Yau, and X. D. Gu. Surface quasi-conformal mapping by solving beltrami equations. In *IMA Conference on the Mathematics of Surfaces*, pages 391–408, 2009.

APPENDIX

A. E_Q AND OPTIMAL UNIFORM ERROR

Let (u, v) be local coordinates in the parametric plane. The linear approximation to the surface normal over $A(Q)$ is $\mathbf{n}_0 + \nabla \mathbf{n} \mathbf{p}$, where \mathbf{p} is a point in the parametric plane, $\nabla \mathbf{n}$ is the Jacobian of $\mathbf{n} = \mathbf{n}(u, v)$, and \mathbf{n}_0 is the normal at the origin of parametric coordinate system; we assume that the origin is chosen to be at the point where the quad Q is tangent to the surface, i.e. $\mathbf{n}_0 = \mathbf{n}_Q$. By definition of the shape operator S , $S\mathbf{f}_u = -\mathbf{n}_u$ and $S\mathbf{f}_v = -\mathbf{n}_v$, in other words, $S\nabla \mathbf{f} = -\nabla \mathbf{n}$. We rewrite the expression for the normal as $\mathbf{n}_0 - S\nabla \mathbf{f} \mathbf{p}$. Then the pointwise squared error is given by

$$E_{pt}^2 = (\mathbf{n} - \mathbf{n}_0)^2 = \mathbf{p}^T \nabla \mathbf{f}^T S^T S \nabla \mathbf{f} \mathbf{p} \quad (7)$$

We assume that the surface is tangent to the quad at the center, (we need to expand the quad in two directions to make this true for an arbitrary tangent point), integrating E_{pt} over the quad Q in the tangent plane, we obtain

$$\begin{aligned} E_Q^2 &= \frac{1}{\text{Area}(Q)} \int_Q (\mathbf{n} - \mathbf{n}_0)^2 \det \nabla \mathbf{f} \, du dv \\ &= \frac{h^4}{12} \text{Tr}(\nabla \mathbf{f}^T S^T S \nabla \mathbf{f}) = \frac{h^4}{12} \text{Tr}(S^T S \nabla \mathbf{f} \nabla \mathbf{f}^T) \\ &= \frac{h^4}{12} \text{Tr}(S^2 M(\mathbf{g})^{-1}), \end{aligned} \quad (8)$$

where we use $\det \nabla \mathbf{f} = \text{Area}(Q)$ and $\nabla \mathbf{f} = \nabla \mathbf{g}^{-1}$. This quantity approximates the integral of previously defined quad error up to $O(h^5)$ for each quad. Let $H = M(\mathbf{g})^{-1}$.

In general, H cannot be chosen arbitrarily (i.e., for an arbitrary choice of H corresponding \mathbf{g} may not exist). One can define an “ideal” H , solving the minimal uniform-error density optimization problem with H as a free variable, without the constraint $H = M(\mathbf{g})^{-1}$.

We require the error density (3) to be constant everywhere on a surface, $\text{Tr} S^2 H = \epsilon$, and minimize ϵ . We constrain the total area the image of the surface has in the parametric plane (with overlaps counted multiple times) to be fixed:

$$\int_A \det \nabla \mathbf{g} \, dA = \int_A \det H^{-\frac{1}{2}} \, dA.$$

Then the Lagrange function with multipliers λ and μ for the constrained minimization of ϵ is

$$\epsilon + \int_A \lambda \text{Tr}(S^2 M) + \mu \det H^{-\frac{1}{2}} \, dA.$$

Computing the L^2 -gradient of this expression with respect to H yields $\lambda S^2 + \frac{1}{2} \mu H \det H^{-\frac{3}{2}} = 0$, i.e., $H = k S^{-2}$. Substituting into $\text{Tr} S^2 H = \epsilon$, we get $k = \epsilon/2$, i.e., the scale factor is independent of H .