Quadrangulation of Non-Rigid Objects using Deformation Metrics

Jiaran Zhou [1,3], Marcel Campen [2], Denis Zorin [3], Changhe Tu [1], Claudio T. Silva [3]

[1] Shandong University, China
[2] Osnabrück University, Germany
[3] New York University, USA

Abstract

We present a novel method to generate quad meshes for non-rigid objects. Our method takes into account the geometry of a collection of key poses in one-to-one correspondence or even an entire animation sequence. From this input, on a common computational domain, an extremal metric is computed that captures the local worst case behavior in terms of distortion as the object undergoes deformation. An anisotropic, non-uniformly sized quad mesh is then generated based on this metric. This mesh avoids undersampling when deformed into any of the poses specified in the input and thus reduces artifacts. Hence, in contrast to previous approaches which target static geometry, our method aims to optimize the mesh's adaptation to the shape for every pose expected during animation or deformation rather than for one specific reference state.

1 Introduction

In recent years a variety of versatile methods for the automatic quadrangular remeshing of surfaces has been developed. These methods typically optimize the properties of the quad mesh, such as element *size*, *anisotropy*, *orientation*, and mesh *connectivity*, with respect to the specific given surface geometry. Note that this is appropriate only when dealing with rigid objects with static geometric properties. At the same time, it is common for quad meshes to be used in animation or simulation. In this case, while the quad mesh might be well adapted to a particular specific state or pose (based on which it was generated and optimized), it may be a poor match for other poses or states of deformation (cf. Figures 1, 8, and 7).

We present a method that aims to produce quadrangulations *adequate for every pose* during deformation or animation, rather than *optimal for a single pose* (cf. Figures 1 and 2). If known a priori, the deformation can be given as input in form of a complete surface animation sequence. Otherwise, the expected space of deformation can be outlined by a set of extremal key poses.

Based on one-to-one correspondences between the key poses or animation frames we analyze the deformation structure and construct an extremal metric corresponding to the worst-case local metric behavior at every point of the surface. Furthermore, information about preferable mesh element orientation is determined jointly over the expected deformation. Together, this allows us to generate an adapted quad mesh with anisotropic element sizing that avoids undersampling artifacts due to being too coarse in any region for any pose when being deformed.
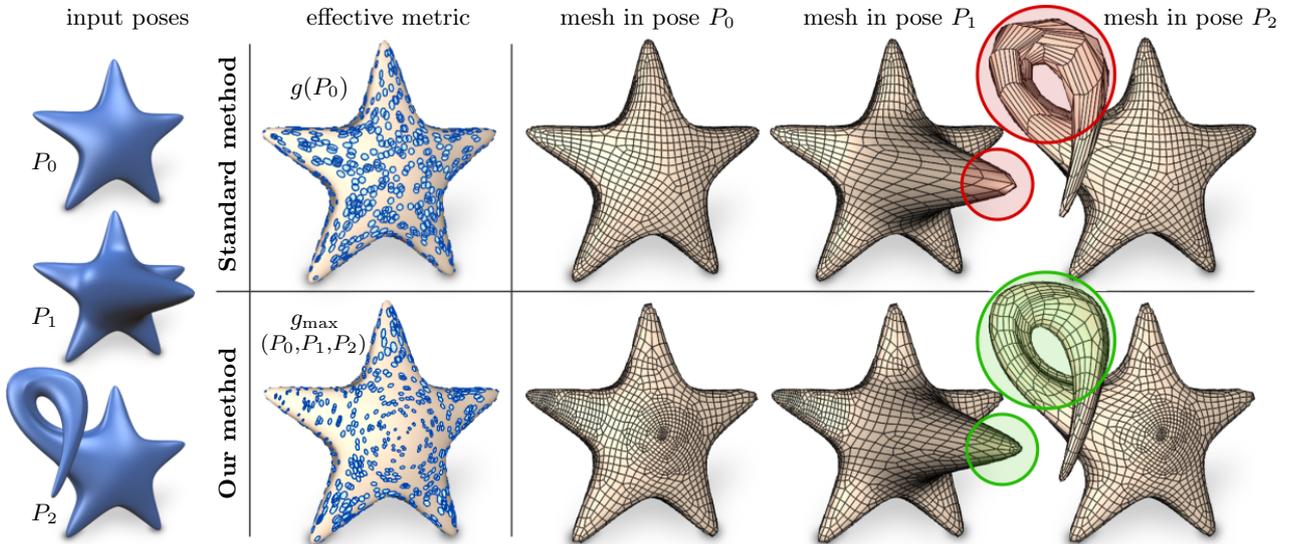
Figure 1: Illustration of our deformation-aware mesh generation method in contrast to a standard, static method [20]. Three poses (left) are taken as input. The standard method optimizes the mesh with respect to the metric of one specific pose, P_0; ours uses an extremal metric it computes over all poses. In the top row, the resulting quad mesh deformed into the poses other than P_0 exhibits artifacts, especially in areas of significant stretching and bending (red). Our quad mesh (bottom row) shows better behavior (despite even having a slightly lower total number of elements) in these poses. Notice that this comes with an increased number of irregular vertices; these are induced by the stronger local variation in element size and shape due to deformation-awareness.
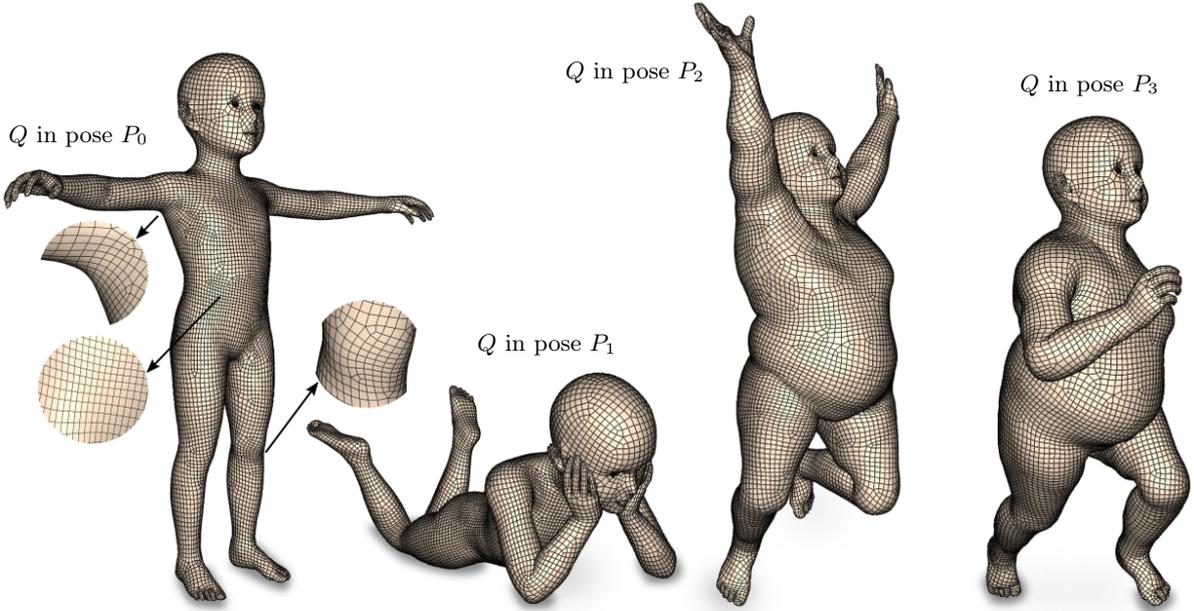
Figure 2: The same quad mesh Q in four different poses for which it was jointly optimized using our technique with extremal metric. Notice, for instance, the higher density of the mesh in the left pose on the knees (because they get bent strongly in poses P_1 and P_2), under the arm (because it gets stretched in pose P_2), and on the belly (because poses P_2 and P_3 are more corpulent).

2 Related Work

Quad Meshing   An account of modern quadrangular surface remeshing techniques is given in a recent survey [4]. Parametrization based techniques have received the most attention in recent years for their flexibility and result quality [2,5,7,9,14,23,32,37, 42,45]. These methods all optimize the mesh properties for one specific surface pose.

A step towards taking surface deformation or animation into account has been made by Marcias et al. [33]: the orientation of mesh elements is chosen based on the principal directions of local surface stretch during deformation. The work by Meng and He [35] considers combined feature determination over multiple poses, again to influence mesh element orientation. Our contribution is orthogonal and complementary to these works: we focus on the determination of appropriate element *sizing*, *anisotropy*, and *shear*; it can be used in conjunction with these techniques that determine *orientation*, as well as other, e.g. principal curvature direction based techniques, as detailed in Section 5. Interactive, user-assisted quadrangulation techniques [34, 49] are another option for mesh generation: the user can design the mesh taking knowledge about expected deformation behavior into account.

Cross Fields   A key component of quad remeshing techniques based on parametrization is the computation of 4-symmetric direction fields. These fields predetermine the parametrization's singularities and thus the quad mesh's irregular vertex structure, and provide directional guidance information for the parametrization's isolines (which correspond to the quad mesh's edges). Key contributions in this context concern the efficient handling of the 4-symmetry [18,26,41], control over the field topology [5,11,31,48], and generalization to non-orthonormal directions [13,21,42]. A recent survey [52] provides a detailed overview of the topic.

Custom Metrics   In some instances of previous work properties of cross fields or parametrizations are optimized with respect to a metric different from the standard Euclidean metric. For instance, the surface under consideration can be virtually smoothed in this way, geometrically [47] or even topologically [14]; the shape operator can be used as metric tensor to improve approximation properties [17,27]. Also user-designed metrics can be taken into account, e.g. for the creation of cross or frame fields with spatially varying magnitude and skewness [21,40].

Deformation-Aware Meshing   A number of works have addressed the problem of generating *triangle* meshes for non-rigid models. They proceed by decimating an overly dense mesh, taking an animation sequence or individual key poses into account [12,29,30,36]. This is generally accomplished by averaging the employed decimation error functions over the animation frames or poses. In Section 4.1, we discuss why direct averaging can be inadequate in our context. Several authors proposed techniques to dynamically adjust the mesh connectivity during animation [19,25,43] for application scenarios where this is appropriate. We focus here on the case of a quad mesh with static connectivity, because 1) this is more generally applicable in the context of animation and simulation and 2) the set of possible connectivity changes in quad meshes is less rich compared to the triangle case [3,44,50], even more so if not only proper connectivity but element quality is taken into consideration.

3 Overview

The input to our method is a set of N surfaces P_i, called poses, with pairwise diffeomorphisms $h_{ij} : P_i \to P_j$ mapping between them. When the poses are obtained through deformation from a reference pose, say P_0, such maps are of course immediately available.

The sequence of poses can be the frames of a complete animation sequence, or a set of key poses which ideally are extremal, in the sense that other expected deformation states are likely to lie within the convex hull, i.e. could be obtained by interpolation between some of the key poses. The input order of poses plays no role in our method.

With existing quad meshing algorithms one could generate a quad mesh optimized for an individual pose P_i (based on its Euclidean metric, or any other shape related metric, cf. Sections 2 and 4). By means of the maps h_{ij}, one can deform this quad mesh into a quad mesh for any other pose P_j. However, it may be unfit for those poses, with artifacts due to undersampling, stretch, or shear, as illustrated in Figure 1.
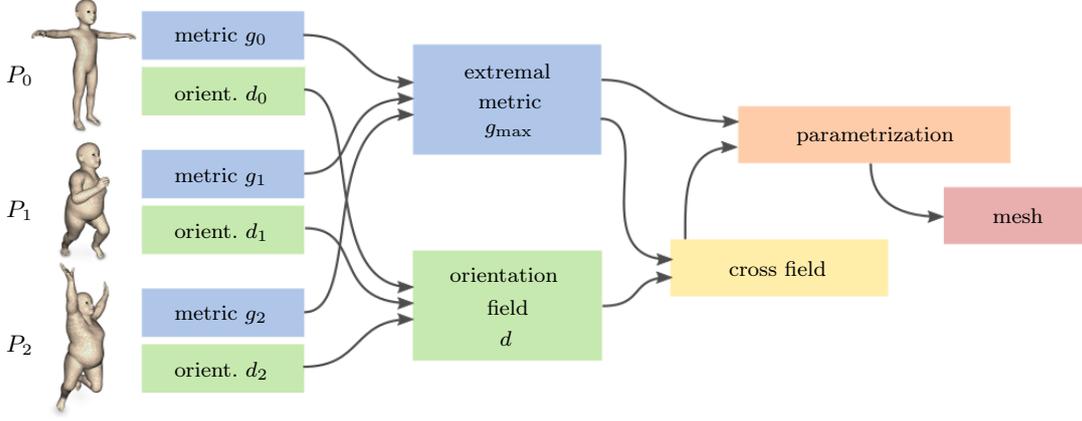
Figure 3: Overview of our method: Multiple poses or frames (P_0, P_1, P_2) are taken as input. Metrics g_i and orientation fields d_i are computed per pose. The metrics are consolidated into an extremal metric $g_{\max}$ on a common domain. Orientation fields are combined through appropriate averaging operations. Based on this a cross (or frame) field and a field guided parametrization are computed. Finally, a quad mesh is extracted from the parametrization.

In our approach we express the target metrics of the individual poses on one common computational domain and combine them into an extremal metric that expresses what quad mesh element sizing and anisotropy is appropriate for all poses, avoiding undersampling in all directions, for all poses (Section 4). Intuitively, elements are made smaller in regions that get stretched in some pose, so they do not become overly large when deformed accordingly. Similarly, information about preferred element orientation (based on shape-derived properties such as principal curvature directions, or on deformation-derived properties such as principal stretch directions [33]) is combined on the common domain (Section 5). A quad mesh can then be optimized with respect to this metric and orientation information using (variations of) state-of-the-art techniques (Section 6). We demonstrate that this strategy yields meshes which behave better when undergoing deformation (Section 7). Figure 3 shows a visual overview.

4 Deformation-Aware Metric

In order to enable control over the sizing and anisotropy of mesh elements one can approach the quad mesh generation problem as follows: let a standard quad meshing method simply aim for unit square quads, while adjusting the notion of "unit square" by means of a custom metric on the surface. For a background on the concepts of differential geometry used in the following, we refer to the introduction by Frankel [16].

Let g_i be a Riemannian metric tensor field on pose P_i. As usual, for a choice of local tangent space bases these quadratic forms can be represented by matrices, $g_i(x) \in \mathbb{R}^{2 \times 2}$, for every point x on P_i.

The following types of metrics are of particular relevance:

- $g_i = sI$ uniform, isotropic
- $g_i = s_i(x)I$ non-uniform, isotropic
- $g_i = A_i(x)$ non-uniform, anisotropic

Here $I \in \mathbb{R}^{2 \times 2}$ is the identity, $s \in \mathbb{R}$ is a global scaling factor, $s_i : P_i \to \mathbb{R}$ is a scalar sizing field on P_i, and $A_i : P_i \to \mathbb{R}^{2 \times 2}$ is an arbitrary field of positive definite quadratic forms, i.e. symmetric positive definite matrices, on P_i.

Anisotropically sized elements are particularly of interest to minimize and equalize discretization errors. A choice of $g_i = sS_i$, where S_i is the shape operator on P_i (whose eigenvalues and eigenvectors are the principal curvatures $\kappa_{\min}$ and $\kappa_{\max}$ and their directions $v_{\min}$ and $v_{\max}$, respectively) is known to optimize the geometric approximation of the surface [17] by the resulting mesh in the limit of small quads. A choice of $g_i = sS_i^T S_i$ (illustrated in the inset) is known to optimize approximation in terms of equalized surface normal error [27]. In practice, to ensure positive definiteness in regions of vanishing surface curvature, these can be regularized by adding εI with $\varepsilon > 0$. More generally, for the approximation of arbitrary functions (surface signals), $g_i = sH_i$ is L_2-optimal [38], where H is the function's Hessian.
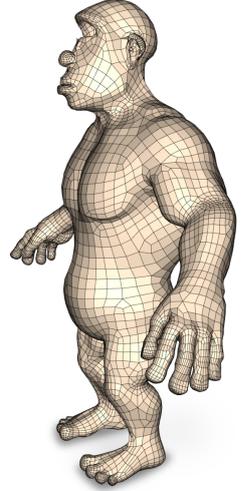

4.1 Metric Consolidation

We can "transfer" the metric g_j from P_j to P_i as the pullback $h_{ij}^* g_j$ of g_j by h_{ij}. Let J_{ij} be the Jacobian of h_{ij}, then $h_{ij}^* g_j = J_{ij}^T g_j J_{ij}$. In this way we can, in particular, obtain the set of N metric tensors $g_{0i} := h_{0i}^* g_i$ on P_0 (where $g_{00} := g_0$), cf. Figure 4. Note that computing a quad mesh on P_0 based on metric g_{0i} and mapping the result to P_i via h_{0i} would yield the same result as performing the computation on P_i based on metric g_i (in a hypothetical exact setting, i.e. up to numerical or discretization effects). We are thus now able to perform computations targeted at specific poses P_i on a common computational domain; w.l.o.g. we choose P_0.

Our goal is to compute a quad mesh that is not optimized for a specific pose, but for all poses. We thus create a new metric on P_0 that takes all the metrics g_{0i} into account. A simple choice is the *average metric*

$$g_{\text{avg}} = \frac{1}{N} \sum_{i=0}^{N-1} g_{0i} \qquad (1)$$
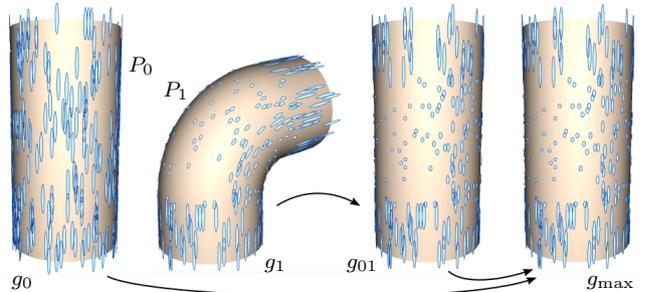
Figure 4: Visualization of metric tensors at random points as ellipses which are unit circles under the respective metric. The metric shown here is based on the shape operator S. Metric g_1 on pose P_1 is mapped to pose P_0 as g_{01}. Then g_{00} (=g_0) and g_{01} are combined to $g_{\max}$ on P_0 (which in this particular example is essentially equal to g_{01}).

defined on P_0. This is akin to the average (or sum) error functions used in pose-aware triangle mesh decimation techniques (cf. Section 2). The problem with this averaging is that it implicitly weights local deformation states by how common they are in the pose set. For instance, consider a large set of poses of a character model where each pose is derived from a rest pose by exercising one joint to one of its extremes. In this case the average metric does hardly differ from the metric of the rest pose, because every surface region is in its rest state in all but one of the many poses.

Extremal Metric We thus propose the use of an *extremal metric* rather than an average metric. For the above example this means that the metric locally reflects the surface stretch experienced in the poses that deform a region the most, rather than the rest pose (cf. Figures 4 and 5).

One needs to decide whether the metric should be extremal in a maximum or a minimum sense; intuitively, whether it should be sensitive to the largest stretching or the largest squeezing of the surface. For the purpose of meshing, the former is relevant, because stretching leads to undersampling, whereas squeezing only leads to unnecessarily fine sampling (which is only an issue if mesh simplicity is of highest priority). We thus define

$$g'_{\max}(v,v) = \max_{i=0}^{N-1} g_{0i}(v,v) \tag{2}$$

for any tangent vector v. Notice that $g'_{\max}$ defined in this way is not a quadratic form, i.e. not a Riemannian metric, in general. We thus consider instead a tight bounding quadratic form $g_{\max}$ defined via

$$\det(g_{\max}) \to \min \ \ \text{s.t.} \ \ g_{\max}(v,v) \geq g'_{\max}(v,v) \text{ for all } v. \tag{3}$$

Figure 5: The effect of the extremal metric $g_{\max}$ on a simple plane model, bent into different poses. Depending on whether only pose P_0 (left), poses P_0 and P_1 (middle), or poses P_0, P_1, and P_2 (right) are taken as input, the resulting quad mesh (shown deformed into the different poses below) will have its element sizing and anisotropy adapted accordingly. A metric based on shape operator S was used here, so element shapes are adapted to curvature: the most extreme curvature encountered locally in any of the poses determines the local element shape and size.

4.2 Implementation Details

The poses P_i are given as triangle meshes. We compute the necessary metric-related values per triangle or vertex of P_0 as follows.

Computing the Jacobians J_{0i} Let each triangle in each of the poses be equipped with a local 2D coordinate system. Let a, b, c be the coordinates of the three vertices of a triangle in P_0. The Jacobian of h_{0i} in this triangle is then computed as

$$J_{0i} = \begin{bmatrix} h_{0i}(b) - h_{0i}(a) & h_{0i}(c) - h_{0i}(a) \end{bmatrix} \begin{bmatrix} b - a & c - a \end{bmatrix}^{-1}.$$

Depending on the mesh generation technique to be used, a metric defined on a per-triangle or per-vertex basis might be required. In the per-vertex case, per-vertex Jacobians can be interpolated from those of the incident triangles (after rotations about the respective axes $n_t \times n_v$, aligning a triangle tangent plane, with normal n_t, to the vertex tangent plane, with normal n_v).

Computing the shape operator S We estimate principal curvatures $\kappa_{\min,i}$, $\kappa_{\max,i}$ and their directions $d_{\min,i}$, $d_{\max,i}$ on a pose P_i [10] (again with respect to local coordinate systems). Then $S_i = \kappa_{\min,i} d_{\min,i} d_{\min,i}^\mathsf{T} + \kappa_{\max,i} d_{\max,i} d_{\max,i}^\mathsf{T}$.

Computing the extremal metric g_{max} While it is long known that the optimization problem (3) has a unique solution [1,22], no simple technique to find it if $N > 2$ is available. Even if we just require the inequality to hold for a sample of directions v, non-trivial iterative solution techniques are necessary [28]. As we need to perform this for every single triangle, we use an approximation instead:

Each metric g_{0i} can be written as $g_{0i} = \lambda_{ai} d_{ai} d_{ai}^\mathsf{T} + \lambda_{bi} d_{bi} d_{bi}^\mathsf{T}$, where λ_{ai}, λ_{bi} are the eigenvalues, and d_{ai}, d_{bi} the eigenvectors of g_{0i}. We construct $g_{\max} = \lambda_a d_a d_a^\mathsf{T} + \lambda_b d_b d_b^\mathsf{T}$, where the eigenvectors d_a and d_b are taken from g_{avg}, and the eigenvalues λ_a, λ_b are computed as:

$$\lambda_{a/b} = \max_{i=0}^{N-1} \lambda_{ai}(d_{a/b}^\mathsf{T} d_{ai})^2 + \lambda_{bi}(d_{a/b}^\mathsf{T} d_{bi})^2. \tag{4}$$

A simple calculation shows that we then have

$$g_{\max}(d_{a/b}, d_{a/b}) = g'_{\max}(d_{a/b}, d_{a/b}) = \max_{i=0}^{N-1} g_{0i}(d_{a/b}, d_{a/b}),$$
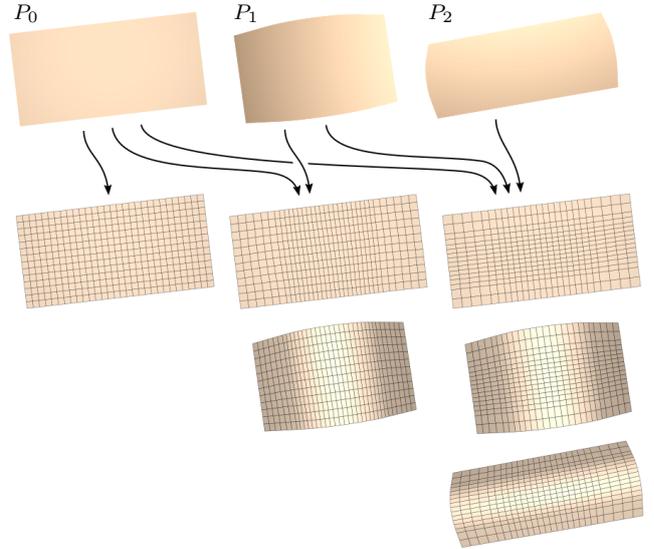
i.e. the approximation is exact in the principal directions d_a and d_b. One can furthermore show (cf. Appendix A) that for other directions v, the approximation error is bounded: $g'_{\max}(v,v)/g_{\max}(v,v) < 2$, i.e. the underestimation of lengths is bounded everywhere and in any direction by a factor $\sqrt{2}$. The approximation can thus easily be made conservative by scaling with a corresponding factor. This global scaling, of course, is without effect if the quadrangulation is performed with respect to a target mesh complexity rather than a specific target sizing, and might as well be omitted.

5 Element Orientation

Often, orienting quad elements such that edges follow principal curvature directions, at least in significantly curved regions, is of interest [5,8]. Alternatively, the user may want to specify the orientation manually, usually in a sparse manner.

We can model such directional guidance objectives in a general way as follows. Let d_i be a tangent vector field on P_i whose direction specifies the desired edge orientation (e.g. the minimum principal curvature direction or another direction set by the user) and whose magnitude specifies the importance (e.g. it may be related to the local anisotropy of the surface, or be 0 in regions where there is no preference in terms of element orientation).

Let $d_{0i} := J_{i0} d_i$ be the pullback of d_i by h_{0i} on P_0. Note that these directions can obviously be contradictory: d_{0i} may differ from d_{0j} for $i \neq j$. This, for instance, is the case if d_i and d_j are principal curvature directions and the surface is bent differently in poses P_i and P_j. In such cases we strive to align to these directions as well as possible by striking a balance.

4

To this end, we compute a weighted average of the directions specified by d_{0i}, where the weights are dictated by the magnitudes $\|d_{0i}\|$. Here it is important to note that a cross field as well as a regular quad mesh, away from extraordinary vertices, has a 4-symmetric structure; whether we enforce element orientation according to a direction d, $d^\perp$, $-d$, or $-d^\perp$ is irrelevant. Here $d^\perp$ denotes a $\frac{\pi}{2}$-rotated (in the tangent plane) version of d. This invariance to rotations by multiples of $\frac{\pi}{2}$ must be considered to obtain a proper average. For instance, averaging d and $d^\perp$ should not result in (some multiple of) $d + d^\perp$, but in d (or any $k\frac{\pi}{2}$-rotation thereof, $k \in \mathbb{Z}$).

To achieve this invariance, we map the vector fields d_{0i} to 4th-order tensor fields [41] which are inherently invariant to rotations by multiples of $\frac{\pi}{2}$. In form of the so-called representation vector [41], they can be averaged by simple vector addition:

$$\bar{d}_{\text{avg}} := \frac{1}{N} \sum_{i}^{N-1} \bar{d}_{0i},$$

where $\bar{d}$ denotes the 4th-order tensor representation of d. The result is then mapped back to a vector field d_{avg}. Note that this, due to the involvement of the magnitude of d_{0i}, is a weighted average. A very similar $\frac{\pi}{2}$-symmetric averaging operation is performed in [33]; some difference lies in that that work uses as final weight (i.e. magnitude of d_{avg}) the (weighted) average of the magnitudes of the d_{0i}, regardless of how compatible or contradictory the individual directions are. With the formulation here, the weights (encoded in the magnitude of the $\bar{d}_{0i}$) are attenuated depending on how strongly the individual directions differ.

5.1 Implementation Details

Curvature Guidance   Some binary feature filters have been proposed in the past [5, 8, 39], that can be used to identify salient regions on meshes. We can apply them per pose P_i and set the directional guidance field d_i to a unit vector in minimum curvature direction in salient regions, and to zero everywhere else. Alternatively, a continuously weighted alignment term can be used by setting the local magnitude of d_i based on a curvedness measure on P_i, e.g., as in [23, 26, 46].

Stretch Guidance   As proposed in [33], one can orient quad mesh elements according to the principal directions of stretch occurring when deforming P_0 to P_i. Here this corresponds to setting d_i to $(|\lambda_0|/|\lambda_1| - 1)s_0$, where $|\lambda_0| > |\lambda_1|$ are the singular values and s_0 is an eigenvector of J_{0i}. Such a choice of directional guidance puts a focus on element shear reduction.

Tensor Conversion   Let d be a tangent vector represented using coordinates in a local 2D tangent coordinate system. Using the standard identification of $\mathbb{R}^2$ and $\mathbb{C}$, we can write $d = re^{i\theta}$, which is the complex representation of d. The representation vector of the 4th-order tensor then simply is $\bar{d} = re^{i4\theta}$.

6   Mesh Generation

Having constructed a metric g (with the case of particular interest here being $g = g_{\max}$) and a guidance field $d = d_{\text{avg}}$, both on $P = P_0$, we strive to generate a quad mesh that is as uniform as possible in g (i.e. with unit length edges and right face angles, measured in g), while edges tend to align to d. The meshing approaches based on cross field guided parametrization (cf. Section 2) are particularly suited for this task. There are two main categories: methods based on *global* parametrization (e.g. [5, 23]) or *local* parametrization (e.g. [20, 46]). The former approach generates a globally consistent integer grid map that implies a pure quad mesh, but requires rather involved measures to achieve robustness [2, 7]. The latter approach is simpler but generates a quad-dominant mesh only—which can be turned into a pure quad mesh by subdivision (introducing an additional irregular vertex for each non-quad) [20].

Most of the meshing methods from these two categories were proposed in a formulation that assumes the Euclidean metric $g_{\text{Euclid}} = I$. Global parametrization methods can be adapted to a different metric g (in particular our $g_{\max}$) either through 'metric uniformization by surface deformation' as proposed by [42], or through the use of g in the definition of parallel transport for the cross field generation and in the definition of norms in the parametrization energy.

The local parametrization based approach has the characteristic property of automatically introducing additional extraordinary vertices (besides those already imposed by the guiding cross field) where necessary to keep the mesh elements very close to their desired local size and shape. According to our experiments, this is particularly beneficial when the metric varies significantly across the surface, e.g. due to strong deformations between poses. The global parametrization approach, by contrast, leads to a smaller number of extraordinary vertices in such cases, at the cost of higher element distortion (cf. Figure 6). As the optimal trade-off is application dependent, we do not intend to advocate one approach over the other and note that either one can be used with our technique. As it more clearly conveys the effects of metric variations, we choose to use a local parametrization approach [20] for demonstrations herein. We adapt it to custom metrics as detailed in the following.
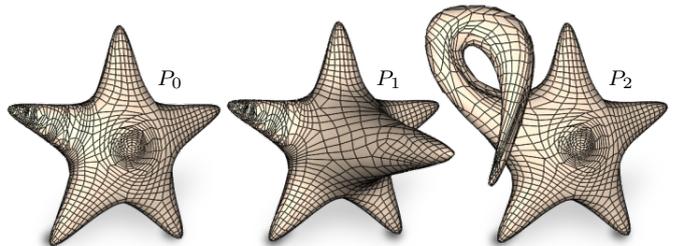
Figure 6: The mesh of Figure 1 bottom as generated when using [2] instead of [20] with our technique. The number of non-regular vertices is lower, while quads are shaped worse.

6.1   Metric-Aware Local Parametrization

Like most related methods, the recent local parametrization approach [20] was described for g_{Euclid}. The authors state that it is straightforward to use custom metrics with this technique. While this is indeed true from a high level perspective, a number of non-trivial issues and details need to be addressed when taking a closer look on a lower level. We describe in the following how this method can be generalized to Riemannian metrics such as our $g_{\max}$. For brevity, we do not repeat the overall algorithm here but refer the reader to the original paper [20].

Symmetry Groups   The method relies on the sets $\mathcal{R}$ and $\mathcal{T}$, which represent the cross field's rotational symmetry group and the integer grid's translational symmetry group, respectively. Under the Euclidean metric, the elements of $\mathcal{R}$ can be obtained from one representative element (unit vector) $\mathbf{o}$ via $\mathcal{R}(\mathbf{o}, \mathbf{n}, k) = \text{rot}(\mathbf{n}, k\frac{\pi}{2})\mathbf{o}$, where $\text{rot}(\mathbf{n}, \theta)$ is a rotation matrix around axis $\mathbf{n}$ by angle θ.

In our case, the four vectors of a metric-dependent cross are supposed to be in $\frac{\pi}{2}$ increments *with respect to g*. This is achieved using the definition $\mathcal{R}_g(\mathbf{o}, \mathbf{n}, k) = \sqrt{g}^{-1}\text{rot}(\mathbf{n}, k\frac{\pi}{2})\sqrt{g}\,\mathbf{o}$, where $\sqrt{g}$ is computed via singular value decomposition: $\sqrt{g} = U\sqrt{\Sigma}\,U^T$

where $U\Sigma U^T = g$. Note that if $\mathbf{o}$ is a unit vector *under* g, i.e. $\mathbf{o}^T g\, \mathbf{o} = 1$, so are the other resulting elements of $\mathcal{R}_g$. Hence $\mathcal{R}_g$ properly represents crosses that are orthonormal under g.

$\mathcal{T}_g$ in our case is supposed to represent a regular grid under g, i.e. a general oblique lattice in Euclidean space. It is simply spanned by the vectors of $\mathcal{R}_g$.

Metric Transport   In order to compare vectors from different tangent spaces, they must be transported to a common space. This is accomplished by a definition of parallel transport. A common choice in the discrete setting is *unfolding* via the rotation matrix $R_{ji} = \mathrm{rot}(\mathbf{n}_j \times \mathbf{n}_i, \angle(\mathbf{n}_j, \mathbf{n}_i))$, where $\mathbf{n}_i$ and $\mathbf{n}_j$ are the normals of two tangent spaces on P. The vector $\mathbf{o}_j$ from tangent space j transported to tangent space i then is computed as $\mathbf{o}_{ji} = R_{ji}\mathbf{o}_j$.

In our setting we also need to be able to transport metric tensors. We define g_{ji} (the tensor g_j from tangent space j transported to tangent space i) as $g_{ji} = R_{ij}^T g_j\, R_{ij}$, where $R_{ij} = R_{ji}^{-1}$. A simple calculation shows that $o_{ji}^T g_{ji}\, o'_{ji} = o_j^T g_j\, o'_j$ holds, i.e. this definition of metric transport is consistent with the above vector transport.

One situation where this metric transport is to be used is in the initial construction of a multiresolution hierarchy of P. When two vertices i and j are merged to a vertex v on a coarser level, a metric tensor for this vertex is interpolated; the tensors of i and j are transported to v and averaged: $g_v = \frac{1}{2}(g_{iv} + g_{jv})$. Further uses of metric transport occur in the following.

Period Jumps   When comparing and averaging crosses in neighboring tangent spaces i and j, the *period jump* k_{ij} needs to be determined. For orthogonal crosses this can be done as originally described: $k_{ij} = \mathrm{argmin}_k \angle(\mathbf{o}_i, \mathcal{R}(\mathbf{o}_{ji}, k))$, i.e. by just considering one pair of representative vectors. The result is independent of which vector $\mathbf{o}_i$ is used as representative for the cross at i.

With our metric g, the crosses are not orthogonal, as described above. The choice must thus be made differently in order to determine the period jump that best aligns the crosses, not just individual representative vectors. This is achieved by:

$$k_{ij} = \underset{k}{\mathrm{argmin}} \sum_{0 \le l < 4} \angle(\mathcal{R}_{g_i}(\mathbf{o}_i, l), \mathcal{R}_{g_{ji}}(\mathbf{o}_{ji}, l + k)).$$

In this way, the combined alignment of all four vector pairs is measured for each choice of period jump to determine the optimal one. By symmetry, the sum can actually be restricted to $0 \le j < 2$ without changing the result.

Finally, it bears noting that the extrinsic smoothness measures which are demonstrated as being beneficial in [20] are not appropriate for our case of a deformable shape (or a shape in multiple poses): there is no static embedding based on which extrinsic measures could be evaluated. We make use of the intrinsic formulation instead, and use our field d as directional guiding input to this method. As suggested, one step of midpoint subdivision is performed on the resulting quad-dominant mesh in order to yield a quad-only mesh.

7   Results

When performing comparisons of different approaches based on different metrics g_a and g_b in the following, we rescale these metrics with a globally constant factor such that $\int_P \det g_a\, dA = \int_P \det g_b\, dA$. This leads to quad meshes of approximately equal complexity (number of elements), allowing for meaningful assessment of the improvements in terms of approximation quality.

Key Pose Sets   In Figures 8 and 7 we compare quad meshes optimized for a static geometry (pose P_0 of the input) with quad meshes optimized jointly for all given poses (two in Figure 7, three in Figure 8) with our method. The S-based metric was used in these examples. The local shape approximation quality in terms of pointwise input mesh to quad mesh distance is visualized over the surface.
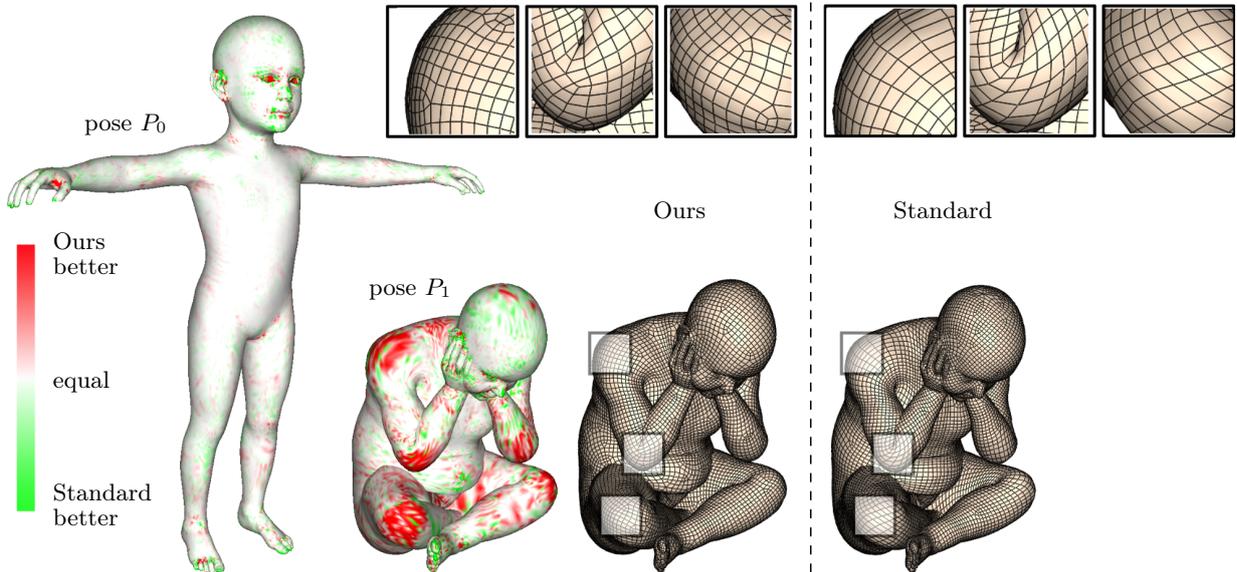
Figure 7: Left: visualization of the *difference* in local shape approximation quality (pointwise input surface to quad mesh distance) when comparing our method (metric $g_{\max}$) and a standard, static method (effectively using metric g_0 only). In green regions, the mesh generated by the standard approach is better, in red regions, it is worse than our result. On P_0 the differences are minor, as expected. On P_1, the mesh optimized for P_0 using the standard, static approach (and then deformed to pose P_1) has significantly larger approximation error (red regions, especially on stretched and bent parts, such as knees, shoulders, elbows) than the mesh optimized over both poses using our method. Right: the two quad meshes in pose P_1; notice that long, stretched quad elements (in particular in direction of strong curvature) are avoided by our method, thereby reducing approximation error due to undersampling. It can also be observed how additional extraordinary vertices are introduced in order to facilitate the required mesh density transitions.
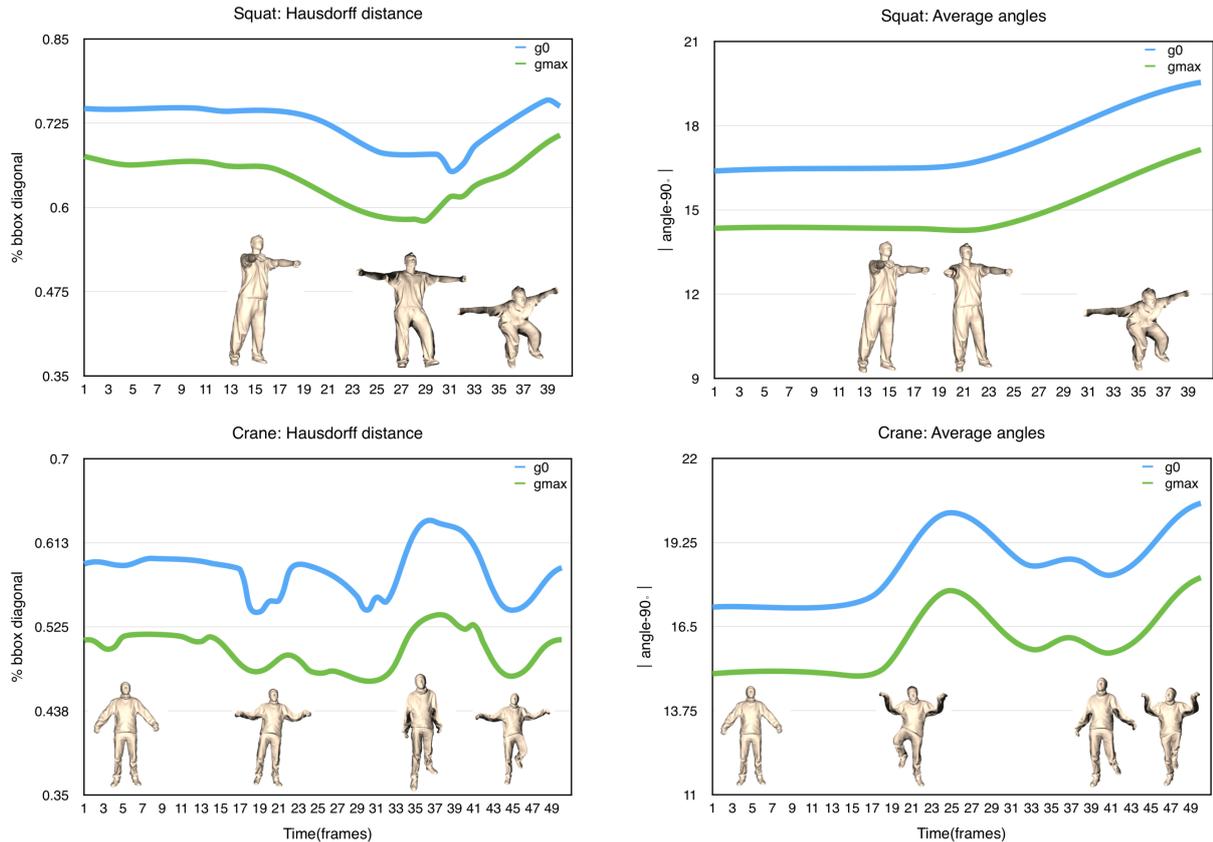
Figure 9: The shape approximation quality (in terms of Hausdorff distance, left column) and element quality (in terms of quad corner angles, right column) of quad meshes during deformation according to given animation sequences (Squat, Crane), some frames of which are depicted below the graphs. We compare our method (green) with a standard approach (blue) that optimizes the mesh for frame 0 of the animation. See the text for more details on the algorithmic choices.

Animation Sequences In Figure 9 we consider the quality of quad meshes undergoing deformation according to given animation sequences. We compare meshes generated by our proposed method, optimized for the entire input animation sequences (using g_{max}), with quad meshes generated for static pose P_0 (as with traditional quad mesh generation methods, using g_0). We used the S-based metric in this comparison, and made use of orientation fields d computed from the principal curvatures as described in [39]. As can be observed, the Hausdorff distance from the input model to the quad mesh is lower when using our method, and the average element shape is better (element corner angles closer to $90°$). Interestingly, in these examples this even holds for the frame P_0; we attribute this to the joint determination of the curvature-based orientation constraints over the entire sequence (cf. Section 5) in our method, which provides more stable, reliable guidance for the model. Possibly, the typically higher number of extraordinary vertices due to g_{max} also contributes to this effect.

Figure 8: Visualization of difference in local shape approximation quality (pointwise input surface to quad mesh distance) between our method, optimizing for all poses, and a method optimizing for a static reference pose (P_0) only. The corresponding quad meshes are depicted in Figure 1. Color coding as in Figure 7, i.e. in red regions our method is better.
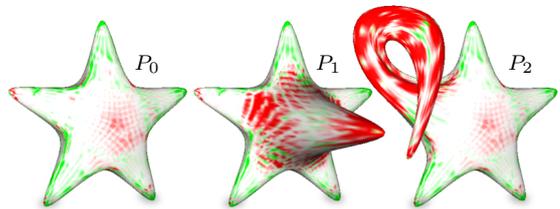
Figure 10 shows a similar comparison based on animation sequences. However, for this experiment we used our method in combination with the orientation guidance based on principal stretch directions proposed by [33], as detailed in Section 5.1. The further improvement due to the use of our g_{max} in addition to this animation-aware orientation guidance becomes clear in this experiment.

8 Limitations & Future Work

A high quality solution to the problem of determining suitable guidance fields d is desirable; those discussed in Section 5.1 are not always satisfactory. When working with a single static model, manual input (such as stroke guidance [15]) can be reasonable, but for multiple poses and even sequences this is impractical; automatic solutions are important in this context.

Finding an exact solution to the tight bounding metric problem (3) would certainly be of value; though we expect the practical implications to be minor in the context at hand. On a related note, metrics that consider stretching *and* squeezing could be of value in certain contexts; averages of g_{max} and an analogously defined g_{min} could provide benefits over a traditional g_{avg} (1) here.

As discussed in Section 6, the introduction of additional extraordinary vertices is of benefit for individual mesh element quality. Depending on the use case it is, however, likewise beneficial if the number of extraordinary vertices is small [4]. For use cases that have a strong preference for meshes with a low number of extraordinary vertices, our method in combination with the chosen meshing strategy [20] is likely not a viable choice, and so far no alternative automatic quad mesh generation method enables explicit and practical continuous control over the trade-off between element quality and extraordinary vertex count. Progress in this direction would thus certainly be of interest. Likewise no fully robust global parametrization based mesh generation technique
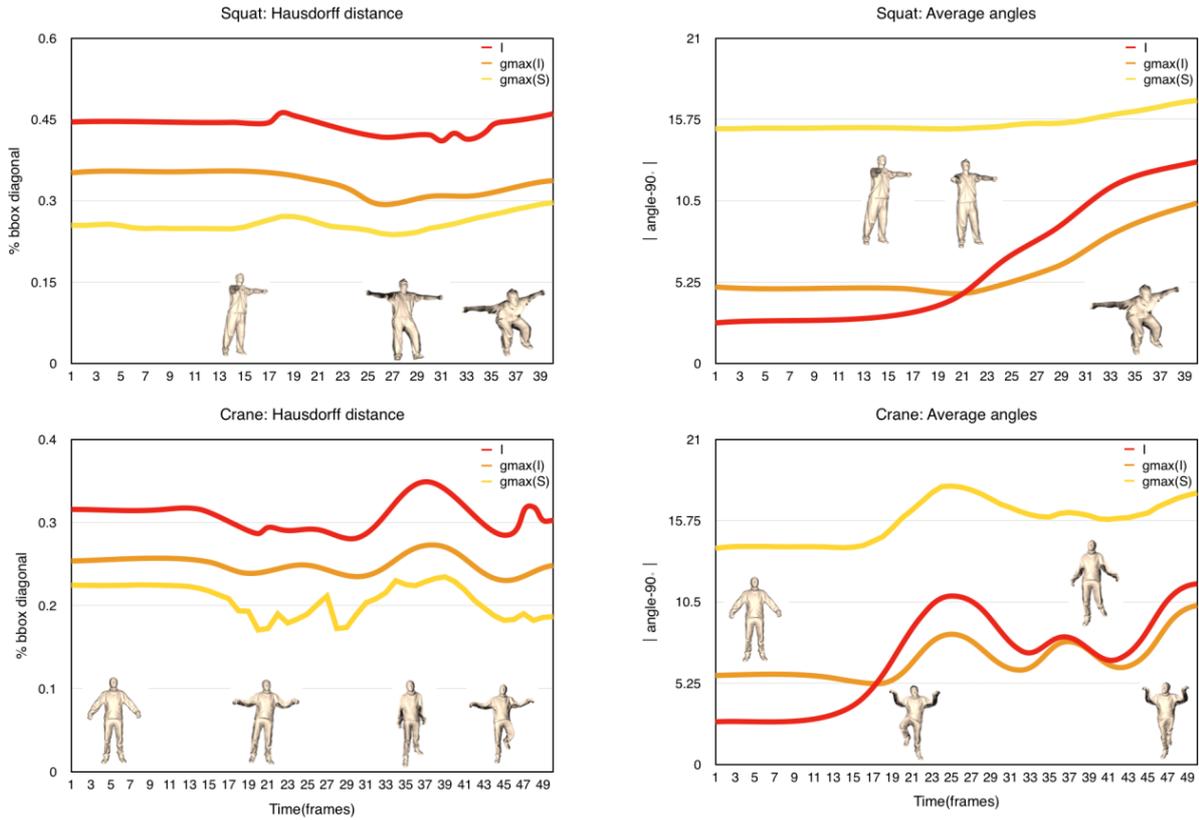
Figure 10: The shape approximation quality (in terms of Hausdorff distance, left column) and element quality (in terms of quad corner angles, right column) of quad meshes during deformation according to given animation sequences (Squat, Crane), some frames of which are depicted. We compare meshes optimized for the standard Euclidean metric I of frame 1 (red graph) with meshes optimized for this metric extremally combined over all frames using our method (orange graph). It can be observed that in these examples the extremal metric reduces approximation error (left). As expected, the use of the S (shape-operator) metric in combination with our approach (yellow graph) leads to further error reduction. The element quality (in terms of corner angles) is best in the beginning of the sequences with the mesh optimized for frame 1 (red), while later the benefit of the extremal metric (orange) shows. The use of S (yellow) naturally leads to less square elements by design; it rather targets low approximation error. Note that the method leading to the red graph is conceptually equal to [33]—with the exception of the local technique [20] rather than the global technique [5] being used for mesh generation in the end.

(that does not introduce additional extraordinary vertices) is available, but would be of interest to reliably achieve low numbers of extraordinary vertices even with highly variable metrics.

It will be interesting to further explore the use of our method for different models of the same category (as opposed to different poses of the same model) to obtain high quality generic quad meshes for shape categories, essentially by a form of *co-quadrangulation*. Depending on the specific scenario, methods for non-rigid shape correspondence finding [24, 51], including recent machine learning based variants [6], could play a central role in this for determining the diffeomorphisms between models.

References

[1] F. Behrend. Über die kleinste umbeschriebene und die größte einbeschriebene Ellipse eines konvexen Bereichs. *Math. Ann.*, 115(1):379–411, 1938.

[2] D. Bommes, M. Campen, H.-C. Ebke, P. Alliez, and L. Kobbelt. Integer-grid maps for reliable quad meshing. *ACM Trans. Graph.*, 32(4):98:1–98:12, 2013.

[3] D. Bommes, T. Lempfer, and L. Kobbelt. Global structure optimization of quadrilateral meshes. *Comp. Graph. Forum*, 30(2):375–384, 2011.

[4] D. Bommes, B. Lévy, N. Pietroni, E. Puppo, C. Silva, M. Tarini, and D. Zorin. Quad-mesh generation and processing: A survey. *Comp. Graph. Forum*, 32(6):51–76, 2013.

[5] D. Bommes, H. Zimmer, and L. Kobbelt. Mixed-integer quadrangulation. *ACM Trans. Graph.*, 28(3):77, 2009.

[6] D. Boscaini, J. Masci, E. Rodolà, and M. Bronstein. Learning shape correspondence with anisotropic convolutional neural networks. In *Advances in Neural Information Processing Systems 29*, pages 3189–3197. 2016.

[7] M. Campen, D. Bommes, and L. Kobbelt. Quantized global parametrization. *ACM Trans. Graph.*, 34(6), 2015.

[8] M. Campen, M. Ibing, H.-C. Ebke, D. Zorin, and L. Kobbelt. Scale-invariant directional alignment of surface parametrizations. *Comp. Graph. Forum*, 35(5), 2016.

[9] M. Campen and D. Zorin. Similarity maps and field-guided T-splines: a perfect couple. *ACM Trans. Graph.*, 36(4), 2017.

[10] F. Cazals and M. Pouget. Estimating differential quantities using polynomial fitting of osculating jets. In *Proc. SGP '03*, pages 177–187, 2003.

[11] K. Crane, M. Desbrun, and P. Schröder. Trivial connections on discrete surfaces. *Comp. Graph. Forum*, 29(5), 2010.

[12] C. DeCoro and S. Rusinkiewicz. Pose-independent simplification of articulated meshes. In *SI3D*, pages 17–24. ACM, 2005.

[13] O. Diamanti, A. Vaxman, D. Panozzo, and O. Sorkine-Hornung. Designing N-PolyVector fields with complex polynomials. *Comp. Graph. Forum*, 33(5):1–11, 2014.

[14] H.-C. Ebke, M. Campen, D. Bommes, and L. Kobbelt. Level-of-detail quad meshing. *ACM Trans. Graph.*, 33(6), 2014.

[15] H.-C. Ebke, P. Schmidt, M. Campen, and L. Kobbelt. Interactively controlled quad remeshing of high resolution 3d models. *ACM Trans. Graph.*, 35(6), 2016.

[16] T. Frankel. *The Geometry of Physics: An Introduction*. Cambridge University Press, 2011.

[17] P. S. Heckbert and M. Garland. Optimal triangulation and quadric-based surface simplification. *Comput. Geom.*, 14(1-3):49–65, 1999.

[18] A. Hertzmann and D. Zorin. Illustrating smooth surfaces. In *Proc. SIGGRAPH '00*, pages 517–526, 2000.

[19] F.-C. Huang, B.-Y. Chen, and Y.-Y. Chuang. Progressive deforming meshes based on deformation oriented decimation and dynamic connectivity updating. In *Proc. SCA '06*, pages 53–62, 2006.

[20] W. Jakob, M. Tarini, D. Panozzo, and O. Sorkine-Hornung. Instant field-aligned meshes. *ACM Trans. on Graphics*, 34(6):189:1–189:15, 2015.

[21] T. Jiang, X. Fang, J. Huang, H. Bao, Y. Tong, and M. Desbrun. Frame field generation through metric customization. *ACM Trans. Graph.*, 34(4):40:1–40:11, 2015.

[22] F. John. Extremum problems with inequalities as subsidiary conditions. In *Studies and Essays Presented to R. Courant on his 60th Birthday*, pages 187–204. Interscience Publishers, New York, 1948.

[23] F. Kälberer, M. Nieser, and K. Polthier. Quadcover - surface parameterization using branched coverings. *Comp. Graph. Forum*, 26(3):375–384, 2007.

[24] V. G. Kim, Y. Lipman, and T. Funkhouser. Blended intrinsic maps. *ACM Trans. Graph.*, 30(4), 2011.

[25] S. Kircher and M. Garland. Progressive multiresolution meshes for deforming surfaces. In *Proc. SCA '05*, pages 191–200, 2005.

[26] F. Knöppel, K. Crane, U. Pinkall, and P. Schröder. Globally optimal direction fields. *ACM Trans. Graph.*, 32(4), 2013.

[27] D. Kovacs, A. Myles, and D. Zorin. Anisotropic quadrangulation. *Computer Aided Geometric Design*, 28(8), 2011.

[28] P. Kumar and E. A. Yildirim. Minimum-volume enclosing ellipsoids and core sets. *Journal of Optimization Theory and Applications*, 126(1):1–21, 2005.

[29] E. Landreneau and S. Schaefer. Simplification of articulated meshes. *Comp. Graph. Forum*, 28(2):347–353, 2009.

[30] H. Lee, M. Ahn, and S. Lee. Displaced subdivision surfaces of animated meshes. *Computers & Graphics*, 35(3):532–541, 2011.

[31] W.-C. Li, B. Vallet, N. Ray, and B. Lévy. Representing higher-order singularities in vector fields on piecewise linear surfaces. *IEEE TVCG*, 12(5):1315–1322, 2006.

[32] Y. Liu, W. Xu, J. Wang, L. Zhu, B. Guo, F. Chen, and G. Wang. General planar quadrilateral mesh design using conjugate direction field. *ACM Trans. Graph.*, 30(6), 2011.

[33] G. Marcias, N. Pietroni, D. Panozzo, E. Puppo, and O. Sorkine-Hornung. Animation-aware quadrangulation. *Comp. Graph. Forum*, 32(5):167–175, 2013.

[34] G. Marcias, K. Takayama, N. Pietroni, D. Panozzo, O. Sorkine-Hornung, E. Puppo, and P. Cignoni. Data-driven interactive quadrangulation. *ACM Trans. Graph.*, 34(4):65:1–65:10, 2015.

[35] M. Meng and Y. He. Consistent quadrangulation for shape collections via feature line co-extraction. *Computer-Aided Design*, 70:78–88, 2016.

[36] A. Mohr and M. Gleicher. Deformation sensitive decimation. Technical report, University of Wisconsin, 2003.

[37] A. Myles and D. Zorin. Controlled-distortion constrained global parametrization. *ACM Trans. Graph.*, 32(4), 2013.

[38] E. Nadler. Piecewise linear best L_2 approximation on triangulations. *Approximation Theory*, V:499–502, 1986.

[39] M. Nieser, J. Palacios, K. Polthier, and E. Zhang. Hexagonal global parameterization of arbitrary surfaces. *IEEE Trans. Vis. Comput. Graph.*, 18(6), 2012.

[40] K. Pal, C. Schüller, D. Panozzo, O. Sorkine-Hornung, and T. Weyrich. Content-aware surface parameterization for interactive restoration of historical documents. *Comput. Graph. Forum*, 33(2):401–409, 2014.

[41] J. Palacios and E. Zhang. Rotational symmetry field design on surfaces. *ACM Trans. Graph.*, 26(3), 2007.

[42] D. Panozzo, E. Puppo, M. Tarini, and O. Sorkine-Hornung. Frame fields: Anisotropic and non-orthogonal cross fields. *ACM Trans. Graph.*, 33(4):134, 2014.

[43] F. Payan, S. Hahmann, and G. P. Bonneau. Deforming surface simplification based on dynamic geometry sampling. In *Proc. Shape Modeling and Applications (SMI '07)*, pages 71–80, 2007.

[44] C.-H. Peng, E. Zhang, Y. Kobayashi, and P. Wonka. Connectivity editing for quadrilateral meshes. *ACM Trans. Graph.*, 30(6), 2011.

[45] N. Pietroni, M. Tarini, O. Sorkine, and D. Zorin. Global parametrization of range image sets. *ACM Trans. Graph.*, 30(6), 2011.

[46] N. Ray, W. C. Li, B. Lévy, A. Sheffer, and P. Alliez. Periodic global parameterization. *ACM Trans. Graph.*, 25(4):1460–1485, 2006.

[47] N. Ray, B. Vallet, L. Alonso, and B. Levy. Geometry-aware direction field processing. *ACM Trans. Graph.*, 29(1), 2009.

[48] N. Ray, B. Vallet, W. C. Li, and B. Lévy. N-symmetry direction field design. *ACM Trans. Graph.*, 27(2):10:1–10:13, 2008.

[49] K. Takayama, D. Panozzo, A. Sorkine-Hornung, and O. Sorkine-Hornung. Sketch-based generation and editing of quad meshes. *ACM Trans. Graph.*, 32(4):97:1–97:8, 2013.

[50] M. Tarini, N. Pietroni, P. Cignoni, D. Panozzo, and E. Puppo. Practical quad mesh simplification. *Comp. Graph. Forum*, 29(2):407–418, 2010.

[51] O. van Kaick, H. Zhang, G. Hamarneh, and D. Cohen-Or. A survey on shape correspondence. *Comp. Graph. Forum*, 30(6):1681–1707, 2011.

[52] A. Vaxman, M. Campen, O. Diamanti, D. Panozzo, D. Bommes, K. Hildebrandt, and M. Ben-Chen. Directional field synthesis, design, and processing. *Comp. Graph. Forum*, 35(2), 2016.

A   Extremal Metric Approximation Bound

For a given tangent vector v, let $i = \text{argmax}_{i=0}^{N-1} g_{0i}(v, v)$. Let

$$g_{0i} = \mu_a e_a e_a^\intercal + \mu_b e_b e_b^\intercal, \quad g_{\max} = \lambda_a d_a d_a^\intercal + \lambda_b d_b d_b^\intercal$$

be expressions of the metrics in terms of their eigenvalues and eigenvectors. Then

$$g_{\max}(v, v) = (d_a^\intercal v)^2 \lambda_a + (d_b^\intercal v)^2 \lambda_b.$$

Noticing that $\lambda_{a/b} = g_{0i}(d_{a/b}, d_{a/b})$ (cf. Equation (4)), we obtain

$$g_{\max}(v, v) = (d_a^\intercal v)^2 g_{0i}(d_a, d_a) + (d_b^\intercal v)^2 g_{0i}(d_b, d_b)$$

$$= (d_a^\intercal v)^2 \big((e_a^\intercal d_a)^2 \mu_a + (e_b^\intercal d_a)^2 \mu_b\big) + (d_b^\intercal v)^2 \big((e_a^\intercal d_b)^2 \mu_a + (e_b^\intercal d_b)^2 \mu_b\big)$$

$$= (e_a^\intercal d_a d_a^\intercal v)^2 \mu_a + (e_b^\intercal d_a d_a^\intercal v)^2 \mu_b + (e_a^\intercal d_b d_b^\intercal v)^2 \mu_a + (e_b^\intercal d_b d_b^\intercal v)^2 \mu_b$$

$$= \big((e_a^\intercal d_a d_a^\intercal v)^2 + (e_a^\intercal d_b d_b^\intercal v)^2\big) \mu_a + \big((e_b^\intercal d_a d_a^\intercal v)^2 + (e_b^\intercal d_b d_b^\intercal v)^2\big) \mu_b.$$

We compare this to $g'_{\max}(v, v) = g_{0i}(v, v) = (e_a^\intercal v)^2 \mu_a + (e_b^\intercal v)^2 \mu_b$, by considering the coefficients of μ_a and μ_b. In particular, we show in the following that these are related by

$$\big((e_a^\intercal d_a d_a^\intercal v)^2 + (e_a^\intercal d_b d_b^\intercal v)^2\big) \geq \frac{1}{2}(e_a^\intercal v)^2$$

(and analogously for the coefficient of μ_b). From this it immediately follows that $g'_{\max}(v, v)/g_{\max}(v, v) < 2$.

Due to orthonormality of d_a, d_b ($d_a^\intercal d_a = d_b^\intercal d_b = 1$ and $d_a^\intercal d_b = 0$), we have $e_a^\intercal d_a d_a^\intercal v + e_a^\intercal d_b d_b^\intercal v = e_a^\intercal (d_a d_a^\intercal + d_b d_b^\intercal) v = e_a^\intercal v$. Defining $p = e_a^\intercal d_a d_a^\intercal v$ and $q = e_a^\intercal d_b d_b^\intercal v$, we thus only need to show that $p^2 + q^2 \geq \frac{1}{2}(p+q)^2 = \frac{1}{2}(p^2 + q^2) + pq$ to prove the above inequality, or equivalently that $p^2 + q^2 \geq 2pq$. One easily verifies that this holds for any $p, q \in \mathbb{R}$.