Planar Shape Interpolation with Bounded Distortion

Renjie Chen[1], Ofir Weber[2], Daniel Keren[2], and Mirela Ben-Chen[1]

[1]Technion - Israel Institute of Technology
[2]University of Haifa

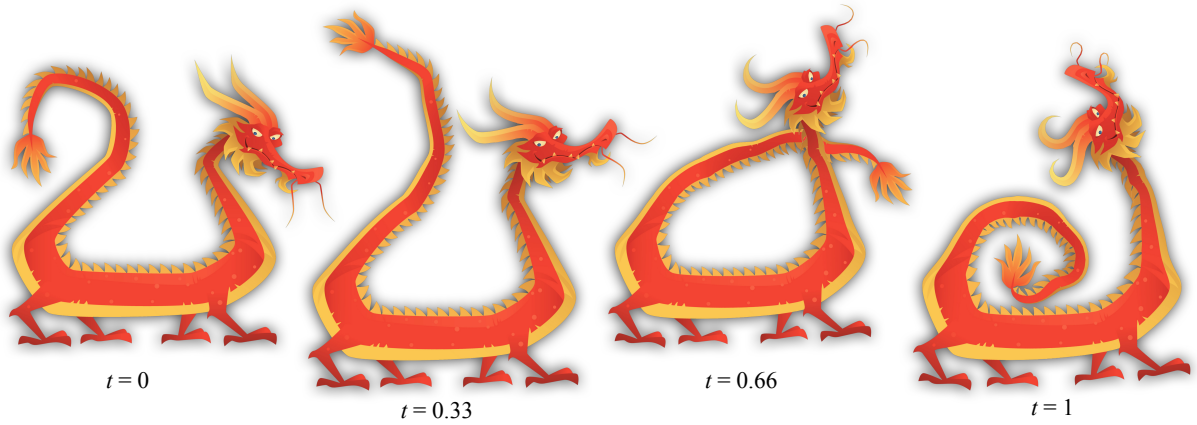

Figure 1: *Two intermediate frames interpolating between the source shape (left) and the target (right). Note the natural interpolation of both large and small scale features (the tail and the horns), seamless handling of large rotations, and small distortion in the intermediate frames.*

Abstract

Planar shape interpolation is widely used in computer graphics applications. Despite a wealth of interpolation methods, there is currently no approach that produces shapes with a bounded amount of distortion with respect to the input. As a result, existing interpolation methods may produce shapes that are significantly different than the input and can suffer from fold-overs and other visual artifacts, making them less useful in many practical scenarios. We introduce a novel shape interpolation scheme designed specifically to produce results with a bounded amount of conformal (angular) distortion. Our method is based on an elegant continuous mathematical formulation and provides several appealing properties such as existence and uniqueness of the solution as well as smoothness in space and time domains. We further present a discretization and an efficient practical algorithm to compute the interpolant and demonstrate its usability and good convergence behavior on a wide variety of input shapes. The method is simple to implement and understand. We compare our method to state-of-the-art interpolation methods and demonstrate its superiority in various cases.

CR Categories: I.3.3 [Computer Graphics]: Picture/Image Generation—Display Algorithms I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation;

Keywords: shape interpolation, bounded distortion, conformal mapping, conformal distortion, quasi-conformal, triangle meshes

Links: DL PDF

1 Introduction

Planar shape interpolation is a fundamental ingredient in many graphics and geometry processing applications. Blending two shapes is instrumental for generating in-between key-frames in computer animation sequences, whereas blending multiple shapes can be used for shape design and exploration. As different applications require different types of shape interpolation, there exists no universal standard by which one can qualitatively assess shape interpolation methods. However, in many cases it is important to preserve the underlying geometric details of the given input shapes as much as possible. Since the input shapes are different, introducing some amount of distortion is unavoidable. Nonetheless, it is somewhat surprising that despite the plurality of different interpolation schemes, none of the existing approaches enables direct control of the amount of distortion that is introduced.

Consider the following alternative formulation to the shape interpolation problem. Denote one input shape as the *source* and consider the set of maps from the source to the other *target* shapes. Augmenting this set of maps with the identity map, we can now perform an interpolation of maps rather than shapes. Clearly, the identity map is the ultimate map, having zero metric distortion. Other maps in this set, however, exhibit some amount of unavoidable angular or area distortion. The goal is to generate maps that interpolate the input while demonstrating a small amount of distortion, which is bounded pointwise by the input's distortion.

We propose a novel planar shape interpolation method designed especially to be shape preserving. We provide a continuous solution that is based on blending the pullback metrics of the input maps, which encode all the local map distortion in the planar case. The

ACM Reference Format
Chen, R., Weber, O., Keren, D., Ben-Chen, M. 2013. Planar Shape Interpolation with Bounded Distortion. ACM Trans. Graph. 32, 4, Article 108 (July 2013), 11 pages. DOI = 10.1145/2461912.2461983 http://doi.acm.org/10.1145/2461912.2461983.

DOI: http://doi.acm.org/10.1145/2461912.2461983

action of the blended metric tensor on the source shape introduces curvature and requires further processing in order to generate a planar embedding. We overcome this difficulty by applying a conformal transformation to the blended curved metric such that the two metrics agree on the boundary while the transformed metric is flat everywhere inside the domain. The choice of using a conformal transformation is motivated by the fact that it does not affect the amount of angular distortion that is introduced in the blending step. Moreover, a conformal map with such boundary conditions always exists and is unique. The actual blending is done in a way that guarantees that the conformal distortion is bounded both locally and globally by the amount of distortion that the input maps exhibit. Furthermore, algorithms for computing a high quality discrete approximation of such a conformal map exist and are relatively easy to implement. We demonstrate that using the method of [Springborn et al. 2008] for computing the conformal map produces high quality results with observed good convergence properties on a large variety of inputs.

For the special case where all the given input maps are conformal (meaning they have zero angular distortion), our method will also yield a conformal map of the *source* (also the *target*). We further show that our method is closed under affine transformations, is symmetric with respect to permuting the source and targets, can handle arbitrarily large rotations and has several other desirable properties. Finally, we demonstrate the superiority of our approach compared to state-of-the-art planar shape interpolation methods.

1.1 Previous Work

It is beyond our scope to review all existing interpolation, blending and morphing methods, and we will restrict ourselves to several notable methods and the most recent approaches. Comprehensive review of classic methods is provided in [Wolberg 1998] for image-based methods, and in [Alexa 2002] for mesh-based methods. Here, we concentrate on *planar* shape interpolation methods, although we mention a few relevant surface based approaches as well.

Every shape interpolation method boils down to three steps: Describing the geometry of the source shapes in terms of some geometric quantities, interpolating these quantities, and reconstructing the blended geometry from the interpolated quantities. The different methods differ by the quantity they choose to interpolate. The most straightforward approach is to represent the geometry using the *embedding*, namely the two- or three-dimensional vertex locations. This global approach is not invariant to isometric deformations of the input shapes and has the obvious drawback that the result is in many cases self intersecting and thus invalid.

Some interpolation methods (e.g. by Shurazky and Gotsman [2001; 2003]) were designed specifically in order to guarantee that the intermediate results are intersection free; however, they require the boundary of the input domains to be convex, which does not hold in most practical scenarios. Some efforts to overcome this limitation embed a non-convex shape in an artificial convex support structure – with possibly degraded quality of the results. Another popular approach [Alexa et al. 2000] is to represent the geometry of the target by using the element-wise Jacobian (either vertex or face based) of the source to target map. By interpolating separately the rotational and shear/skew components of the Jacobians, such methods generate pleasing blends with relatively low amount of shear in many cases. This general approach has been used in other applications such as planar shape deformation [Igarashi et al. 2005], surface modeling [Sorkine and Alexa 2007], parameterization [Liu et al. 2008] and image registration [Sỳkora et al. 2009].

Despite their usefulness, these interpolation methods fail to successfully handle maps with large rotations, and some are quite sensitive to the underlying mesh tessellation or are non-symmetrical with respect to swapping the source and target shapes. These problems were partly addressed by [Choi and Szymczak 2003; Fu et al. 2005] and more fully by [Baxter et al. 2008]. Other interpolation approaches, which are based on a differential coordinate representation, were proposed by [Sheffer and Kraevoy 2004; Lipman et al. 2005; Xu et al. 2006] and [Kircher and Garland 2008]. However, even state-of-the-art methods suffer from a severe drawback: the main ingredient in all these methods is to act element-wise, and then "combine" everything by solving a global optimization problem. Since the system is solved in the least-squares sense, there is no guarantee on the maximal local distortion introduced in the blended shapes. One of our main observations is that it is possible to compute a global embedding from the element-wise geometric properties while preserving the conformal distortion, simply by using a conformal map to the plane.

An alternative approach is to develop algorithms based on some physical view, e.g., by using continuum mechanics and elasticity theory [Hu et al. 2004; Bao et al. 2005; Chao et al. 2010]. Since such approaches are based on finite elements discretization, they are guaranteed to converge in the limit of refinement. However, they still cannot provide guarantees on the amount of local distortion of the blended shapes. The same drawback holds for the shape space method in [Kilian et al. 2007] where an additional concern is computational efficiency, due to the involved non-linear and non-convex optimization method. Most recently, [Winkler et al. 2010; Fröhlich and Botsch 2011] propose a method based on interpolating the edge lengths and dihedral angles, coupled with a non-linear reconstruction method. We show that interpolating edge lengths in 2D is somewhat problematic, as in some cases it causes dependency on the source triangulation which is not reduced under refinement.

There are blending methods [Sederberg and Greenwood 1992; Shapira and Rappoport 1995; Surazhsky and Elber 2002; Liu et al. 2004; Klassen et al. 2004; Whited and Rossignac 2011] that interpolate between boundary *curves* of planar domains, instead of between the domains themselves. This problem is significantly different than the problem we are addressing, as we are assuming that we are *given* the map between the interior of the domains, and need to take it into account when computing the blended shapes.

The method of [Weber and Gotsman 2010] interpolates smooth conformal maps by interpolating the so-called angular factor on the shape's boundary. The uniqueness of this method lies in its ability to guarantee that any intermediate map is also conformal, hence has exactly zero angular distortion. Our method can be seen as an extension of this approach that supports the wider class of quasi-conformal maps for guaranteeing bounded conformal distortion.

Recently, a few methods have been developed [Zeng et al. 2009; Gu et al. 2010; Lipman et al. 2012; Weber et al. 2012] which are based on the theory of quasi-conformal maps, but whose aims are different. These focused on either computing a quasi-conformal deformation or computing the map of the interior given the map of the boundary. Lipman [2012] provides a general framework for computing parameterizations with bounded distortion which potentially could be combined with shape interpolation methods. The method requires the user to explicitly set a global bound on the distortion as well as a reference frame field. Choosing an inappropriate bound or frames may lead to an empty solution space. Our method is conceptually simpler, leads to an efficient algorithm and requires less input from the user. Furthermore, it provides a point-wise bound on the distortion which is implicitly derived from the input map.

1.2 Contributions

Our main contributions are:

- A continuous formulation of the shape blending problem with a set of useful properties (Section 2).
- A continuous solution based on metric blending and conformal mapping, which fulfills these properties (Section 3).
- A discretization and a simple practical algorithm that is experimentally shown to converge to the continuous solution, has the required properties, and compares favorably with the state-of-the-art (Sections 4 and 6).

2 The Problem

We will first describe the problem and our proposed solution in the continuous case, and later show how to discretize it. To simplify the exposition and limit the notational clutter, we will formulate the case of interpolating between two shapes - a source and a target. This setup will be extended to the case of N shapes in Section 5.

Given source and target shapes, we compute intermediate shapes such that a specific set of properties holds which guarantee the quality of the interpolations. All properties follow the same two guiding principles: the deformation of the intermediate frames with respect to the source shape should be as "simple" as the deformation of the target, and the frames should vary smoothly from source to target.

Therefore, we require the Lagrange property, namely source and target reproduction for the appropriate values of t. Further, we require that if the target is a "simple" deformation of the source (e.g. affine or conformal), then so are the intermediate frames. As a natural generalization, we require the intermediates to be "as close to conformal" to the source as the target is. This property can be defined rigorously using the *conformal distortion* introduced by a map. Finally, following the smoothness guiding principle, we require the interpolation to be invariant to the swap of source and target, and to vary smoothly in time.

In order to formally specify these properties, we will need the following standard definitions.

2.1 Definitions

Let $S_0 \subset \mathbb{R}^2$ be a simply connected domain. A *metric* g_p on S_0 is a symmetric bilinear form defined at every point $p \in S_0$: $g_p : \mathbb{R}^2 \times \mathbb{R}^2 \to \mathbb{R}$. The metric allows us to measure lengths (using $g_p(Y, Y)$) and angles (using $g_p(Y, Z)$), where Y, Z are tangent vectors at p.

When we apply a map $\phi : S_0 \to \mathbb{R}^2$, the changes in lengths and angles induced by the map can be computed using a special metric, called the *pullback metric* of the map. The pullback metric is defined as $g_p(Y, Z) = \langle d\phi_p(Y), d\phi_p(Z) \rangle$, where $\langle , \rangle$ is the standard Euclidean inner product, and $d\phi_p$ is the differential of ϕ. Since S_0 and $S_1 = \phi(S_0)$ are both flat, $d\phi(Y)$ is simply JY, where J is the 2×2 Jacobian matrix of ϕ. It is easy to check that as a result we have $g_p = J^T J$.

The *conformal distortion* of the map ϕ at a point $p \in S_0$ is given by $K(p) = \sqrt{\lambda_1/\lambda_2}$, where λ_1 and λ_2 are the maximal and minimal eigenvalues, respectively, of g_p.

Equipped with these definitions, we can now formally specify our problem, as follows.

2.2 Problem Specification

Input: Two simply connected planar domains $S_0, S_1 \subset \mathbb{R}^2$, a locally bijective and differentiable map $\phi : S_0 \to S_1$ and a scalar $t \in [0, 1]$.

Output: A map $\hat{\phi}_t : S_0 \to \mathbb{R}^2$ and the corresponding domain $\hat{S}_t = \hat{\phi}_t(S_0)$.

2.2.1 Required Properties

Source and target reproduction. $\hat{\phi}_0 = Id$ and $\hat{\phi}_1 = \phi$, where Id is the identity map.

Pointwise bounded conformal distortion. Let $p \in S_0$ and $K(p)$ be the conformal distortion of ϕ at p. In addition, let $\hat{K}_t(p)$ be the conformal distortion of $\hat{\phi}_t$ at p. Then $\hat{K}_t(p) \leq K(p)$. As a special case, if $K = 1$, thus ϕ is a conformal map, then $\hat{K}_t = 1$, and therefore, $\hat{\phi}_t$ is also conformal.

Affine reproduction. If ϕ is an affine map, then $\hat{\phi}_t$ is affine.

Symmetry. Let $\psi : S_1 \to S_0$ be the inverse of ϕ, and let $\hat{R}_t$ be the interpolation result from S_1 to S_0 at time t. Then $\hat{R}_t = \hat{S}_{1-t}$.

Smoothness. The derivative $d\hat{S}_t/dt$ exists and is bounded.

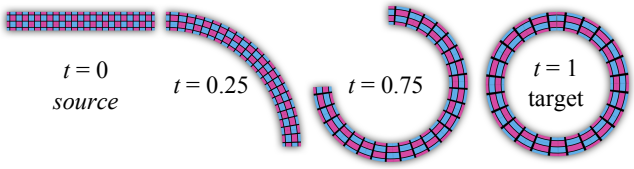

Figure 2: *An example of two intermediate frames generated by our bounded distortion interpolation algorithm. Note the natural interpolation of the large rotation.*

3 Bounded Distortion Shape Interpolation

As discussed in the Introduction, there are many possible *discrete* representations of the geometry that can be interpolated to generate the blended shapes. Our main observations are as follows: First, by choosing a geometry representation which has a continuous equivalent, we can hope for a solution which converges under refinement, and therefore is not dependent on the discretization. Second, since the blended representation is not necessarily realizable in the plane, we should choose a realization method that does not introduce additional distortion. Therefore, we propose a two-step interpolation process:

1. Compute a blended metric g_t for S_0 as the interpolation of the the identity and the pullback metric under ϕ.
2. Let $\hat{\phi}_t : S_0 \to \mathbb{C}$ be the map whose pullback metric is $\hat{g}_t = e^{2u} g_t$, with boundary conditions $u(\partial S_0) = 0$.

Our choice to interpolate the metric is motivated by the following reasons. First, the metric is easy to discretize on a triangular mesh, simply as a 2×2 positive definite matrix per triangle. The approximation of the piecewise-constant metric of the continuous metric improves when the mesh is refined, which is necessary if we expect our algorithm to be consistent for different discretizations. Second, since the metric is rotation invariant, it can simply be interpolated linearly: $g_t = (1-t)I + tg$, where I is the Euclidean metric on S_0 (the identity), and g is the pullback metric under ϕ.

Note that since our domains are planar, the metric tensors can be represented in the same global coordinate frame using 2×2 matrices, and the blending can be done by blending the matrix coefficients. Further, note that these matrices are symmetric positive definite (SPD) by definition, and since the space of SPD matrices forms a convex cone, any convex combination of SPD matrices is

an SPD matrix. Therefore, g_t is a valid metric tensor, yet, the Gaussian curvature that it implies is not necessarily zero.

Since the metric g_t is not flat, it is not possible to directly embed it in the plane. This problem arises in other shape interpolation methods, where the interpolated quantity is no longer consistent - either there is no embedding that realizes it, or such an embedding is not flat. Many shape interpolation method address this problem by solving a Poisson system, namely, finding an embedding which best realizes the interpolated geometric quantity in the least squares sense. Unfortunately, such an approach cannot provide guarantees on properties of the result, such as bounded distortion. We, therefore, chose instead to find a flat metric in the space of metrics *conformal* to g_t. This guarantees that the flattening process does not increase the angular distortion introduced by the interpolation.

In the continuous case, although it is hard to solve analytically for $\hat{g}_t$, it is possible to show that it exists and is unique, as follows. The metric $g_t(p), p \in S_0$ uniquely determines the Gaussian curvature $\kappa_t(p)$, and the Laplace-Beltrami operator $\Delta_t(p)$. If $\hat{g}_t$ is a metric conformally equivalent to g_t, namely $\exists u : S_0 \to \mathbb{R}$, such that $\hat{g}_t = e^{2u} g_t$, and if the Gaussian curvature under $\hat{g}_t$ is 0, then the following partial differential equation holds

$$\kappa_t = \Delta_t u, \tag{1}$$

(see, e.g., [Gu et al. 2010], equation (7.8)). For specific boundary conditions, the solution u exists and is unique, and defines the conformal flat metric $\hat{g}_t$. Note that in 2D, the metric uniquely defines the embedding up to global translation and rotation. These three global degrees of freedom can be fixed automatically by linearly interpolating their values between source and target. Alternatively, manually setting them allows for some artistic freedom and control. For example, an experienced user may draw a curve describing the path that a fixed anchor point should follow.

Of course, this procedure does not provide us with an algorithm for computing $\hat{\phi}_t$ directly, because, in general, equation (1) does not have an analytic solution. We, therefore, justify this model, by proving the properties it has in the continuous case, and then approximate it using a discrete algorithm. Figure 2 shows an example of two interpolated frames computed using the discrete algorithm. We will show experimentally that in the discrete case, the conformal distortion $\hat{K}_t(p)$ converges to the continuous conformal distortion, which we can compute analytically in some simple cases.

3.1 Properties

We will now explain why the continuous solution we defined satisfies all the properties we require.

Source and target reproduction. This property trivially holds. By definition we have that g_0 is the identity, which is a flat metric. Since our boundary conditions are $u = 0$ on the boundary of S_0, the trivial solution $u = 0$ on *all* the domain S_0 solves equation (1). Therefore, the flat metric $\hat{g}_0$ equals g_0. Given the additional degrees of freedom provided by the user (a single point and orientation), we get that $\hat{\phi}_0 = Id$ namely, source reproduction. Using similar arguments, we have that $g_1 = g$, which is also flat and therefore $\hat{g}_1 = g_1 = g$. Combined with interpolating the position and orientation degrees of freedom provided by the user, we get that $\hat{\phi}_1 = \phi$ namely, target reproduction.

Pointwise bounded conformal distortion. The conformal distortion K of a map ϕ is defined as $\sqrt{\lambda_1/\lambda_2}$, where λ_1 and λ_2 are the largest and smallest eigenvalues, respectively, of the pullback metric under ϕ. First, let us consider the conformal distortion K_t of the blended metric g_t. Using the variational definition of the eigenvalues of a matrix, we can show (see Appendix A) that if M_1 and M_2 are positive definite matrices, and $M = tM_1 + (1-t)M_2, t \in [0,1]$, then $\exists \alpha \in [0,1]$ such that $c(M) \leq \alpha c(M_1) + (1-\alpha)c(M_2)$, where $c(M) = \lambda_1/\lambda_2$ is the l_2 matrix condition number of M. Therefore, $c(M) \leq \max(c(M_1), c(M_2))$. Since the conformal distortion K and the condition number of the metric g are related by a monotonic function, we have that $K_t \leq K$. As the flat metric $\hat{g}_t$ is conformal to the metric g_t, the conformal distortion $\hat{K}_t = K_t$. This is easy to see: since $\hat{g}_t = e^{2u} g_t$, the eigenvalues of $\hat{g}_t$ are scaled uniformly, and therefore, their ratio does not change. This computation holds for any point $p \in S_0$, and accordingly we have $\hat{K}_t(p) \leq K(p)$ as required.

Figure 3 shows an example of the interpolation results (based on the discrete algorithm described in Section 4). The figure shows the source, target, some representative intermediate frames, as well as a graph showing the conformal distortion for a few triangles as a function of t. It is evident from the graphs that the pointwise bounded conformal distortion property holds for this discrete case.

Affine reproduction. ϕ is a global affine map if and only if $g(p)$ is constant, independent of p. Therefore, the blended metric g_t is also a constant. As a consequence, g_t is flat, since Gaussian curvature is a result of variation in the metric (this is is easy to see by considering the Brioschi formula, see, e.g., [Gray et al. 2006] Theorem 17.3). Finally, as we have seen previously, if g_t is flat, then we reproduce it; namely, $\hat{g}_t = g_t$. Therefore, $\hat{g}_t$ is also constant and $\hat{\phi}_t$ is a global affine map. Figure 4 demonstrates this property.

Symmetry. Let $h = g^{-1}$ be the pullback metric under ψ, and let $h_s = (1-s)I + sh$. In addition, let $\psi_s : S_1 \to M$ be a map that realizes the metric h_s. The pullback metric of the composed map $\psi_s \circ \phi : S_0 \to M$ is a metric on S_0 given by gh_s (because the pull-back of a composition is the compositions of the pull-backs in reverse order. See e.g. [Weintraub 1997], Proposition 3.61 (b)). To show symmetry, it is enough to show that $gh_{1-t} = g_t$, since if the metrics before the conformal map are identical, the metrics after the conformal map will also be. A simple calculation shows that $gh_{1-t} = g((1-(1-t))I + (1-t)h) = g(tI + (1-t)g^{-1}) = tg + (1-t)I = g_t$.

Smoothness. If the map ϕ is continuously differentiable, the metric g exists and is differentiable everywhere on S_0. Furthermore, g_t is a smooth function of t, and therefore, its derivative with respect to t exists (and is, in fact, constant). If we additionally assume that ϕ has bounded angular distortion, then the determinant of the metric g, as well as that of g_t is bounded away from 0. This implies that κ_t and Δ_t depend smoothly on t, and as a consequence, u (the solution to (1)) also depends smoothly on t. Hence, the metric $\hat{g}_t$, the map $\hat{\phi}_t$ and the domain $\hat{S}_t$ vary smoothly with t. Figure 5 demonstrates this property.

4 Discretization

Now that we have seen that the continuous solution has various desirable properties, we need to discretize it in order to be able to compute the blended shapes in practice. We proceed by formulating the discrete equivalent of the interpolation problem, as follows.

Notations. We use $S = (X, F)$ to represent a triangle mesh, where X is the 2D embedding of the vertices, and F is the connectivity. Variables with subscript t correspond to the intermediate frame t, and we use $\hat{x}$ to represent variables after the conformal map. For example, M_t is the interpolated metric at time t, and $\hat{M}_t$ is the resulting metric after the conformal map.

Input: Two planar triangle meshes with identical connectivity $S_0 = (X_0, F)$ and $S_1 = (X_1, F)$ such that no triangle in S_1 is flipped or collapsed, and a scalar $t \in [0,1]$.

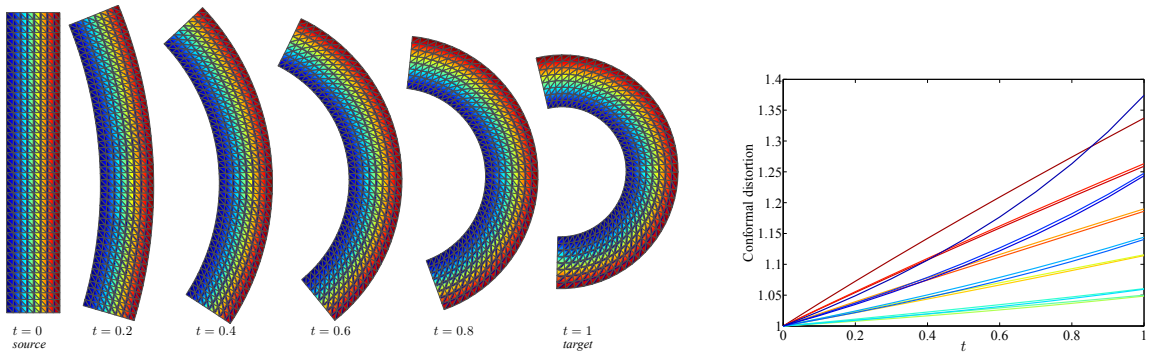

Figure 3: *Shape interpolation with bounded conformal distortion. Every plotted line represents one triangle in the mesh (the lines are colored the same as their corresponding triangles), and the graph plots the conformal distortion of this triangle as a function of t. Note that the graphs are monotonic, and in no case is the intermediate distortion bigger than the distortion at $t = 1$ for* any *triangle.*

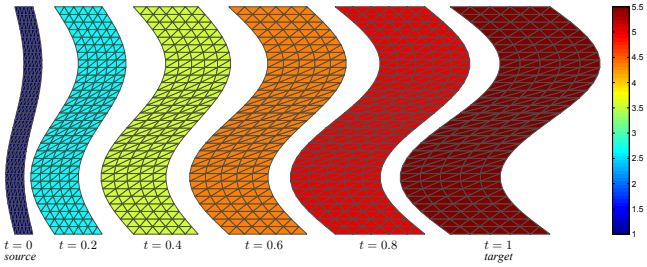

Figure 4: *Affine reproduction. The color shows the conformal distortion with respect to the source. Note, that since the target is a global affine map of the source (with uniform conformal distortion), so are the intermediate frames.*

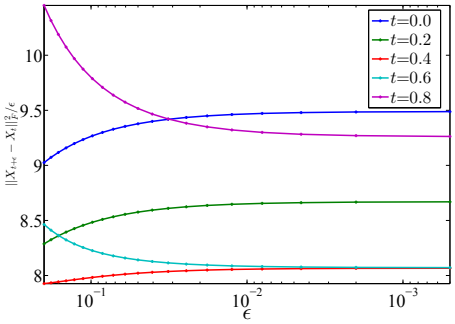

Figure 5: *Smoothness. We show the value $||\hat{X}_{t+\epsilon} - \hat{X}_t||_{Fro}/\epsilon$ for decreasing ϵ. We use the same source and target as in Figure 3, and show different values of t. Note that these values converge as ϵ goes to* 0.

Output: An intermediate planar mesh $\hat{S}_t = (\hat{X}_t, F)$.

We choose the following discretizations for the two main steps of our algorithm:

1. We discretize the metric to be piecewise constant, namely we have a function $M_t : F \to \mathbb{R}^{2\times 2}$.
2. We discretize conformal maps using CETM, the approach in [Springborn et al. 2008].

The first choice is a standard finite element approach. It guarantees that the discrete representation of the metric will become a better approximation of the continuous metric as the triangulation is refined. The second choice does not guarantee convergence to the continuous conformal map; however, it is shown to be well-behaved in practice. Additionally, CETM is well-behaved not only in the limit of refinement, but also for coarse meshes. It is worth noting that any *intrinsic* method, requiring only the metric and not the embedding in order to compute the conformal map, can be used.

Our full algorithm is described in Algorithm 1. In the next section we explain why our algorithm is valid and discuss its properties.

Algorithm Notations. For each triangle $f \in F$, $E_t(f,i), i \in (1,2,3)$ is the length of the edge in the triangle f opposite to the vertex i, $M_t(f)$ is the interpolated metric on f, and e_0 is an edge vector of f.

4.1 Algorithm Validity

For our algorithm to be valid we need to show the following. First, that $e_0^T M_t e_0$ is always positive, and that the triangle inequality holds for E_t. Second, the edge length $E_t(f,i)$ is computed twice, for both of its neighboring triangles - we need to show that these computations are consistent.

Positivity and triangle inequality follow from the fact that the matrix M_t is symmetric positive definite, and therefore the inner product it defines is a distance metric. Consistency of E_t follows from the following simple lemma, which shows that blending metrics is equivalent to blending *squared* edge lengths.

Lemma 4.1 *Let M be the metric for a triangle $f \in F$, as defined in lines (2,3) in Algorithm 1. Further, let e_0 be one edge vector in f in S_0, and e_1 the corresponding edge vector in S_1. Then, $e_0^T M_t e_0 = (1-t)||e_0||^2 + t||e_1||^2$.*

Proof First, note that $||e_1||^2 = e_0^T M e_0$ by definition of M. The result follows by linearity: $e_0^T((1-t)I + tM)e_0 = (1-t)e_0^T e_0 + te_0^T M e_0 = (1-t)||e_0||^2 + t||e_1||^2$.

This simple observation guarantees that the edge lengths we compute are consistent in its adjacent triangles. In fact, our algorithm can be thus simplified to Algorithm 2.

In the following section we investigate which of the properties of the continuous algorithm holds also for its discrete counterpart. In addition, we demonstrate the consistency of the algorithm under a change of triangulation, and the convergence of the discrete conformal distortion to the continuous one.

Algorithm 1: Bounded Distortion Shape Interpolation Algorithm

Input: Triangle meshes $S_0 = (X_0, F)$, $S_1 = (X_1, F)$, a scalar $t \in [0, 1]$

Output: Intermediate mesh $\hat{S}_t = (\hat{X}_t, F)$

foreach *Triangle* $f \in F$ **do**
  J = Jacobian of affine transformation from $X_0(f)$ to $X_1(f)$;
  $M(f) = J^T J$;
  $M_t(f) = (1-t)I + tM(f)$;
  for $i = 1$ **to** 3 **do**
    $e_0 = X_0(f, i+1) - X_0(f, i+2)$;
    `// Here, + is computed modulus 3`
    $E_t(f, i) = \sqrt{e_0^T M_t(f) e_0}$;
  end
end
$\hat{E}_t$ = CETM(E_t, $u = 0$ on the boundary);
$\hat{X}_t$ = embed $\hat{E}_t$ using single position and orientation constraints supplied by user;

Algorithm 2: Simpler BD Shape Interpolation Algorithm

Input: Triangle meshes $S_0 = (X_0, F)$, $S_1 = (X_1, F)$, a scalar $t \in [0, 1]$

Output: Intermediate mesh $\hat{S}_t = (\hat{X}_t, F)$

foreach *edge* $e \in F$ **do**
  Let $||e_0||, ||e_1||$ be the edge lengths of e in S_0 and S_1 respectively;
  $E_t(e) = \sqrt{(1-t)||e_0||^2 + t||e_1||^2}$;
end
$\hat{E}_t$ = CETM(E_t, $u = 0$ on the boundary);
$\hat{X}_t$ = embed $\hat{E}_t$ using single position and orientation constraints supplied by user;

4.2 Properties

Whenever the intermediate edge lengths E_t are embeddable in the plane (namely, the angle deficit at every vertex is 0), then CETM returns $\hat{E}_t = E_t$, and we have all the continuous properties. This shows source, target and affine reproduction. Symmetry also trivially holds, because we are blending the squared edge lengths.

The conformal distortion implied by the edge lengths E_t – namely, before the conformal map – is pointwise bounded using the same arguments we used in the continuous case (simply considering the metric of a triangle, instead of a point). However, CETM returns only a discrete conformal map, and therefore the conformal distortion at every triangle is bounded only in the limit of refinement.

We will now demonstrate experimentally the convergence and consistency properties of our algorithm. First, we show an example for which we can compute the conformal distortion induced by our continuous algorithm explicitly. We demonstrate that the discrete conformal distortion resulting from the discrete algorithm converges under refinement to the continuous conformal distortion. Second, we discuss a scenario where our algorithm is invariant to the choice of triangulation, but a similar algorithm is not.

4.3 Experimental Convergence

In order to show convergence to the continuous case, we need to compute analytically the continuous solution, which is infeasible except in trivial cases. We will, therefore, show convergence of the *conformal distortion* to the continuous conformal distortion, which we can compute analytically if we know explicitly the map ϕ. In the following we refer to ϕ as a map to $\mathbb{C}$, instead of as a map to $\mathbb{R}^2$, as it is easier to compute analytically the conformal distortion if the map is represented using complex functions.

We set up our experiment as follows. Let $\phi : \mathbb{C} \to \mathbb{C}$ be a complex function from the plane to itself, and let $\phi_z(z) = \partial\phi/\partial z$ and $\phi_{\bar{z}}(z) = \partial\phi/\partial\bar{z}$. Then, a simple calculation shows that the metric tensor $M(z), z \in \mathbb{C}$ is given by:

$$M(z) = |\phi_z|^2(|\mu|^2+1)\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} + 2|\phi_z|^2\begin{bmatrix} \Re(\mu) & \Im(\mu) \\ \Im(\mu) & -\Re(\mu) \end{bmatrix},$$

where $\Re(\mu)$ and $\Im(\mu)$ are the real and imaginary components of μ, and $\mu(z) = \phi_{\bar{z}}(z)/\phi_z(z)$.

Given a triangle mesh $S_0 = (X_0, F)$, we compute $S_1 = (X_1, F)$ using $X_1(z) = \phi(z)$. For every triangle $f \in F$, take z_f to be the barycenter of the triangle, compute $M_t(z_f) = (1-t)I + tM(z_f)$. Given the metric $M_t(z_f)$ we can compute $K_t(z_f) = \sqrt{\lambda_1/\lambda_2}$, where λ_1 and λ_2 are the largest and smallest eigenvalues of $M_t(z_f)$, respectively. We compute the discrete conformal factor $\hat{K}_t(f)$ similarly, using the Jacobian of the affine map that takes the triangle f in S_0 to the triangle f in $\hat{S}_t$. Finally we compute the following quantities:

$$\begin{aligned} E_2 &= \sum_{f\in F} A_f ||K_t(f) - \hat{K}_t(f)||^2 / \sum_{f\in F} A_f, \\ E_\infty &= \max_{f\in F} |K_t(f) - \hat{K}_t(f)|, \end{aligned} \tag{2}$$

where A_f is the area of the triangle f in $\hat{S}_t$. E_2 and E_∞ represent the approximation error of the discrete solution. As the size of the triangles decreases, we expect the discrete solution to provide a better approximation, and these values to decrease as well.

Figure 6 (left) shows a plot of E_2 and E_∞ as a function of the mesh resolution (number of vertices) for the map $\phi = m_1 \circ A \circ m_2$, where m_1 and m_2 are Möbius maps of the form $m(z) = \frac{az+b}{cz+d}$. $A(z) = (1+K)\Re(z)+(1-K)\Im(z)i$ is a simple affine stretch. The map ϕ is a Teichmüller map (a conformal map composed with an affine stretch, composed with a conformal map), one of the simplest examples of a quasi-conformal map. The results for the following three Teichmüller maps:

$$\begin{aligned} \phi_1(z) &= \frac{2z}{z-5} + \frac{2}{3}\frac{\bar{z}}{\bar{z}-5}, \\ \phi_2(z) &= \frac{20}{2z-5} + \frac{12}{2\bar{z}-5}, \\ \phi_3(z) &= \frac{3(4z+5)}{3z+1} + \frac{4\bar{z}+5}{3\bar{z}+1} \end{aligned}$$

appear in Figure 6. Note, that both E_2 and E_∞ converge to 0.

4.4 Why *squared* edge lengths?

A few previous approaches [Lipman et al. 2005; Winkler et al. 2010] based the interpolation on blending the edge lengths of the triangle mesh (in addition to some extra properties). Although the edge lengths of a mesh are often regarded as the "discrete metric", we show that directly interpolating edge lengths (as opposed

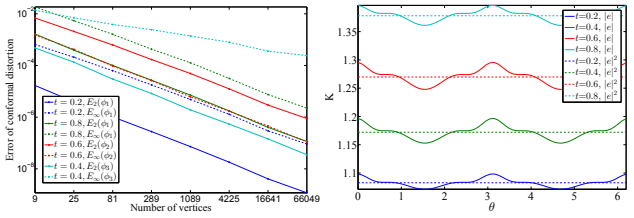

Figure 6: *(left) Experimental convergence of the discrete conformal distortion to the continuous value, for three Teichmüller maps, using various values of t. The plot shows the values of E_2 and E_∞ defined in* (2) *as a function of the mesh resolution. Note that both measures converge to 0. (right) Blending squared edge lengths vs. blending edge lengths. We show the discrete conformal distortion $K(\theta)$ as a function of the rotation angle θ of a source triangle. Note that blending squared edge lengths results in a constant function $K(\theta)$, whereas blending edge lengths does not.*

to *squared* edge lengths) is not consistent, in the sense that the resulting conformal distortion depends on the triangulation.

Consider the following experiment: Let $\phi(z) = (1+k)\Re(z) + i(1-k)\Im(z)$, for some real $k \in \mathbb{R}$, $0 \leq k < 1$. We choose a triangle f in the plane, apply ϕ to the vertices of f, and compute the interpolated result at some t, once by blending edge lengths, and once by blending squared edge lengths. Given the new triangles, we can compute the conformal distortion for both cases. We then rotate f by an angle θ, apply ϕ, and repeat the process. If the interpolation scheme is consistent, the result should be invariant to θ.

Figure 6 (right) shows that this is the case for blending squared edge lengths, but is not the case for blending edge lengths. The figure shows the resulting discrete K as a function of θ. Note that the result when blending edge-lengths varies, whereas the result when blending squared edge lengths remains the same. Note further that since the map ϕ has constant conformal distortion, taking a smaller triangle will not improve the result. Hence, blending edge lengths cannot be invariant to the triangulation of the source domain.

5 Multiple Targets

Thus far we have focused on the simple case, with a single source and a single target shape. However, our formulation can be easily extended to multiple targets. The problem definition in the continuous case changes as follows:

Input: N simply connected planar domains $S_i, i \in [1, N]$ and a base domain S_0, N locally bijective and differentiable maps $\phi_i : S_0 \to S_i$ and weights $w_i \geq 0, \sum_{i=1}^N w_i = 1$.

Output: A map $\hat{\phi}_w : S_0 \to \mathbb{R}^2$ and the corresponding domain $\hat{S}_w = \hat{\phi}_w(S_0)$, such that the following properties hold:

Input reproduction. $\hat{\phi}_{w^i} = \phi_i$ where w^i are weights that are 0 for all shapes except S_i.

Pointwise bounded conformal distortion. Let $p \in S_0$ and $K_i(p)$ be the conformal distortion of ϕ_i at p. In addition, let $\hat{K}_w(p)$ be the conformal distortion of $\hat{\phi}_w$ at p. Then $\hat{K}_w(p) \leq \max_i K_i(p)$. As a special case, if $\forall i, K_i = 1$, and thus all ϕ_i are conformal maps, then $\hat{K}_w = 1$, and therefore, $\hat{\phi}_w$ is also conformal.

Affine reproduction. If all ϕ_i are affine maps, then $\hat{\phi}_w$ is affine.

Invariance to the choice of S_0. Let $R_0 \neq S_0$ and let $\psi_i : R_0 \to S_i$, such that $\phi_j \circ \phi_i^{-1} = \psi_j \circ \psi_i^{-1}, \forall i, j \in [1, N]$. Let $\hat{S}_w$ be the interpolation result using S_0 and weights w, and $\hat{R}_w$ the interpolation result using R_0 and weights w. Then $\hat{S}_w = \hat{R}_w$.

Smoothness. The derivatives $\partial \hat{S}_w / \partial w_i$ exist and are bounded.

Our two-step solution can be easily extended to this case:

1. Compute $g_w = \sum_{i=1}^N w_i g_i$, where g_i are the pullback metrics under ϕ_i. Note that g_i are defined on S_0, and therefore, can be safely interpolated.

2. Let $\hat{\phi}_w : S_0 \to \mathbb{C}$ be the map whose pullback metric is $\hat{g}_w = e^{2u} g_w$, with boundary conditions $u(\partial S_0) = 0$.

The proofs for most properties are analogous to the single-source single-target case. First, g_w is a valid metric, since it is a convex combination of SPD matrices. Input and affine reproduction are a result of reproducing flat metrics. The bounded distortion property also carries over to the general case, since all we need is a convex combination of SPD matrices. Smoothness also follows from similar considerations.

The only property that is not trivial is the invariance to the choice of S_0. However, the proof is similar to the symmetry proof. Let h_i be the pullback metrics under ψ_i. Let $h_w = \sum_{i=1}^N w_i h_i$. We cannot compare g_w and h_w directly, given that they are defined on different domains. We will, therefore, pull them both to S_1. Accordingly, we need to show that $g_1^{-1} g_w = h_1^{-1} h_w$. Since we assumed that $\phi_j \circ \phi_i^{-1} = \psi_j \circ \psi_i^{-1}, \forall i, j \in [1, N]$, we have $g_1^{-1} g_i = h_1^{-1} h_i \forall i \in [1, N]$, and the result follows.

Figure 7 shows how multiple target blends with different weights can be used for generating novel poses, given a collection of input poses. See also the accompanying video.

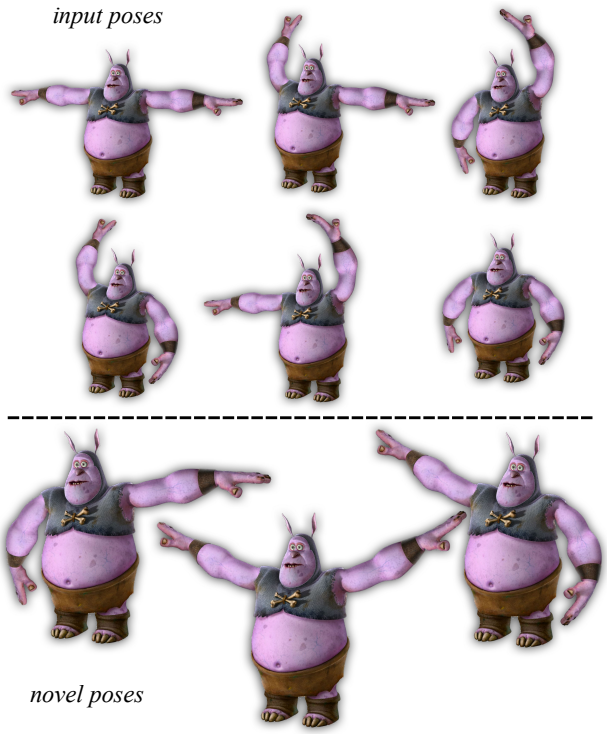

Figure 7: *Blending 6 input poses (top) to generate novel poses (bottom). The deformed textured images look very natural, with a three-dimensional feel even though they are completely planar.*

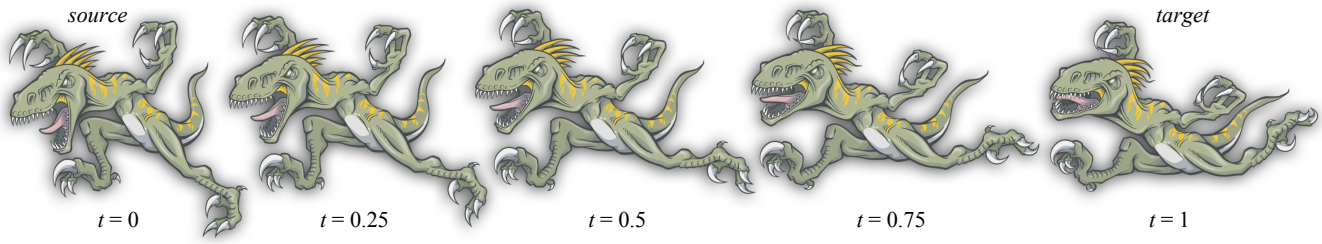

Figure 8: *Bounded distortion interpolation results.*

6 Experimental Results

6.1 Implementation Details

We implemented our algorithm as a plugin to Autodesk Maya®. For implementing CETM we implemented a Newton solver (essentially iterating equation (1)). In almost all cases the solver converged after $2-3$ iterations. For each Newton step a linear system was solved by computing a Cholesky decomposition using the Intel MKL® package. Although this process is computationally heavier than existing methods, we achieve real-time frame rates (30 fps) for meshes with up to 5000 vertices on a PC with i7-3770 CPU, as demonstrated in the accompanying video.

6.2 Limitations

One limitation of our method is that it is designed to bound only the *angular* distortion and not the *area* distortion. However, we have noticed experimentally that the area distortion of the intermediate shape is bounded, as is demonstrated by all our examples, as well as the accompanying video. Note, though, that the bound depends on both the area and the conformal distortions of the target. Specifically, if the conformal distortion of the target is very large, the intermediate might have a larger area than either source or target.

In addition, our method only guarantees bounded angular distortion in the limit of refinement. Although we showed fast convergence experimentally, it would be desirable to design a method that can guarantee bounded distortion discretely by defining a discrete quasi-conformal map, analogously to the definition of discrete conformality in CETM.

6.3 Results and Comparisons

We compared our approach to ARAP [Alexa et al. 2000], triangle based ARAP local/global [Liu et al. 2008], and FFMP [Kircher and Garland 2008]. For local/global ARAP, the rigid energies of the intermediate shape with respect to the source/target are interpolated and minimized, as in [Chao et al. 2010]. The results are summarized in Figure 13, where for each method we show the interpolation result for one t value ($3/4$ for the spirals and $1/2$ for the other examples), and the conformal distortion (normalized to $[0,1]$) graph of 100 most distorted triangles, which are found separately for each t. The first two comparisons are also shown in the accompanying video. The source and target are shown in the first and last columns, respectively, and they are either textured or color coded to show the correspondence between them.

The method with results closest to ours is FFMP, which exhibits nice interpolations on all the examples, except for the "spirals" example, where it introduces an unnecessary global distortion. And, as is evident from the distortion graphs, it always introduces unnecessary distortion in some regions, which can be found by carefully inspecting the results, for example, on one of the blue monster's fingers. The large rotation required in the first 3 examples causes the ARAP method to generate fold-overs and collapsed triangles. Finally, due to the non-isometric input mapping between the source and the target, ARAP local/global generates unnatural results for the dragon, horse and blue monster examples (note e.g. the leg of the dragon).

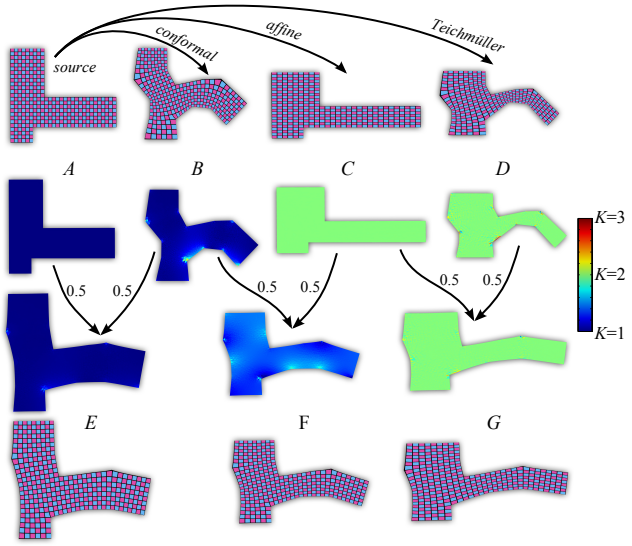

Figure 9: *This figure demonstrates how our method behaves for conformal and Teichmüller maps. The figure shows (top) a source shape A, and three targets B,C,D, along with their conformal distortion. In addition, we show three combinations E = (A+B)/2, F = (B+C)/2 and G = (C+D)/2. See the text for details.*

Figure 9 demonstrates how our method behaves for conformal and Teichmüller maps. The figure shows (top) a source shape A, and three targets B,C,D. In addition to the shapes, we show the conformal distortion with respect to A. B is conformal to A, as is evident from the conformal distortion K_B which is close to 1. C is an affine map of A, as can be seen from the uniform conformal distortion $K_C = 2$. Finally, D is a conformal map of C, as is demonstrated by the fact that $K_D \approx K_C$. We generate the new shapes E = (A + B)/2, F = (B + C)/2 and G = (C + D)/2. Since B is conformal to A, so is E, as can be seen from the fact that the shapes of the quads are preserved and $K_E \approx 1$. Furthermore, we see that in this case blending Teichmüller maps results in an almost Teichmüller map G, demonstrated by $K_G \approx K_C$, which is uniform. In fact, since D is conformal to C, G is also conformal to C, and therefore by definition a Teichmüller of A.

Figure 10 demonstrates the behavior of our algorithm when invalid inputs (e.g. with fold-overs) are provided. Dealing with such cases is sometimes inevitable in practical scenarios. Our algorithm manages to produce high quality, well behaved results. Figure 11

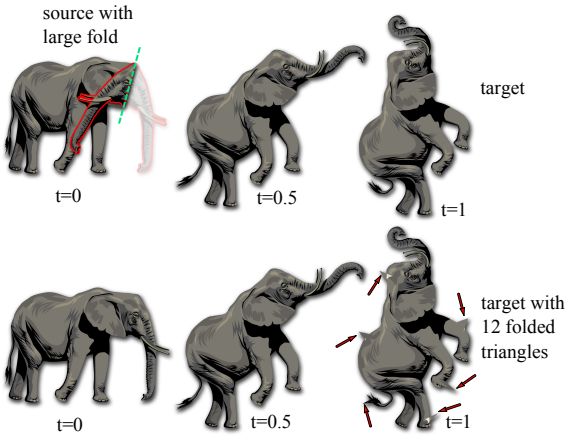

Figure 10: *Robustness to invalid input: (top row) part of the elephant's head in the source is folded onto itself. (bottom row) 12 individual triangles in the target shape are folded. In both cases, the source to target map is not locally bijective, yet the interpolated results are well behaved.*

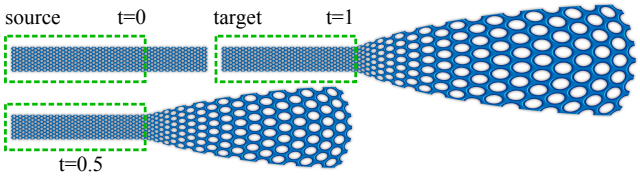

Figure 11: *Local support behavior: the right part of the bar shape is being deformed, whereas the left part remains intact (marked green). The intermediate (t=0.5) exhibits similar behavior, leaving the left part of the shape completely static.*

demonstrates the natural local support behavior of our algorithm. In this figure, only a part of the source shape is being deformed while the rest of the shape remains intact. The same behavior is observed for intermediate frames where only the relevant part of the shape is actively deformed. Figures 1, 8 and 12, as well as the accompanying video, show some additional examples of interpolated results.

7 Conclusion and Discussion

We presented a novel shape interpolation method that produces provably high quality results with a bounded amount of distortion. We formulated the problem in the continuous case, suggested a solution, and proved the properties it satisfies. We further suggested a discretization which experimentally converges to the continuous solution, and showed why a popular alternative solution will not converge. We demonstrated that our method out-performs in practice state-of-the-art shape interpolation methods, which are not guaranteed to have bounded distortion.

Our approach provides a novel model for maps between domains - conformal through an alternative metric with natural boundary conditions. This suggests a natural extension of our algorithm to interpolation between surfaces, using surface based conformal maps, as in [Crane et al. 2011]. Additionally, we believe that this model can be beneficial for other applications such as computing surface parameterization and maps between curved surfaces.

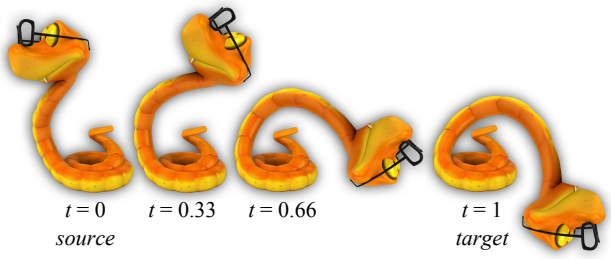

Figure 12: *Bounded distortion interpolation results. Note the natural way our method handles large rotations.*

Acknowledgements

The authors acknowledge the support of Aly Kaufmann postdoctoral fellowship at the Technion, ISF grant 699/12, ISF equipment grant, Marie Curie CIG 303511 and the state of Lower-Saxony and the Volkswagen Foundation, Hannover, Germany.

References

ALEXA, M., COHEN-OR, D., AND LEVIN, D. 2000. As-rigid-as-possible shape interpolation. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, ACM Press/Addison-Wesley Publishing Co., 157–164.

ALEXA, M. 2002. Recent advances in mesh morphing. In *Computer graphics forum*, vol. 21, Wiley Online Library, 173–198.

BAO, Y., GUO, X., AND QIN, H. 2005. Physically based morphing of point-sampled surfaces. *Computer Animation and Virtual Worlds 16*, 3-4, 509–518.

BAXTER, W., BARLA, P., AND ANJYO, K. 2008. Rigid shape interpolation using normal equations. In *Proceedings of the 6th international symposium on Non-photorealistic animation and rendering*, ACM, 59–64.

CHAO, I., PINKALL, U., SANAN, P., AND SCHRÖDER, P. 2010. A simple geometric model for elastic deformations. *ACM Transactions on Graphics (TOG) 29*, 4, 38.

CHOI, J., AND SZYMCZAK, A. 2003. On coherent rotation angles for as-rigid-as-possible shape interpolation.

CRANE, K., PINKALL, U., AND SCHRÖDER, P. 2011. Spin transformations of discrete surfaces. *ACM Trans. Graph. 40*.

FRÖHLICH, S., AND BOTSCH, M. 2011. Example-driven deformations based on discrete shells. In *Computer Graphics Forum*, vol. 30, Wiley Online Library, 2246–2257.

FU, H., TAI, C., AND AU, O. 2005. Morphing with laplacian coordinates and spatial-temporal texture. In *Proceedings of Pacific Graphics*, 100–102.

GRAY, A., ABBENA, E., AND SALAMON, S. 2006. *Modern differential geometry of curves and surfaces with Mathematica*. Chapman & Hall/CRC.

GU, D., LUO, F., AND YAU, S. 2010. Fundamentals of computational conformal geometry. *Mathematics in Computer Science 4*, 4, 389–429.

HU, S., LI, C., AND ZHANG, H. 2004. Actual morphing: a physics-based approach to blending. In *Proceedings of the*

ninth ACM symposium on Solid modeling and applications, Eurographics Association, 309–314.

IGARASHI, T., MOSCOVICH, T., AND HUGHES, J. 2005. As-rigid-as-possible shape manipulation. In *ACM Transactions on Graphics (TOG)*, vol. 24, ACM, 1134–1141.

KILIAN, M., MITRA, N., AND POTTMANN, H. 2007. Geometric modeling in shape space. In *ACM Transactions on Graphics (TOG)*, vol. 26, ACM, 64.

KIRCHER, S., AND GARLAND, M. 2008. Free-form motion processing. *ACM Transactions on Graphics (TOG) 27*, 2, 12.

KLASSEN, E., SRIVASTAVA, A., MIO, M., AND JOSHI, S. 2004. Analysis of planar shapes using geodesic paths on shape spaces. *Pattern Analysis and Machine Intelligence, IEEE Transactions on 26*, 3, 372–383.

LIPMAN, Y., SORKINE, O., LEVIN, D., AND COHEN-OR, D. 2005. Linear rotation-invariant coordinates for meshes. *ACM Transactions on Graphics (TOG) 24*, 3, 479–487.

LIPMAN, Y., KIM, V., AND FUNKHOUSER, T. 2012. Simple formulas for quasiconformal plane deformations. *ACM Transactions on Graphics (TOG) 31*, 5, 124.

LIPMAN, Y. 2012. Bounded distortion mapping spaces for triangular meshes. *ACM Transactions on Graphics (TOG) 31*, 4, 108.

LIU, L., WANG, G., ZHANG, B., GUO, B., AND SHUM, H. 2004. Perceptually based approach for planar shape morphing. In *Computer Graphics and Applications, 2004. PG 2004. Proceedings. 12th Pacific Conference on*, IEEE, 111–120.

LIU, L., ZHANG, L., XU, Y., GOTSMAN, C., AND GORTLER, S. J. 2008. A local/global approach to mesh parameterization. In *Computer Graphics Forum*, vol. 27, Wiley Online Library, 1495–1504.

SEDERBERG, T., AND GREENWOOD, E. 1992. A physically based approach to 2–d shape blending. In *ACM SIGGRAPH Computer Graphics*, vol. 26, ACM, 25–34.

SHAPIRA, M., AND RAPPOPORT, A. 1995. Shape blending using the star-skeleton representation. *Computer Graphics and Applications, IEEE 15*, 2, 44–50.

SHEFFER, A., AND KRAEVOY, V. 2004. Pyramid coordinates for morphing and deformation. In *3D Data Processing, Visualization and Transmission, 2004. 3DPVT 2004. Proceedings. 2nd International Symposium on*, IEEE, 68–75.

SORKINE, O., AND ALEXA, M. 2007. As-rigid-as-possible surface modeling. In *ACM International Conference Proceeding Series*, vol. 257, Citeseer, 109–116.

SPRINGBORN, B., SCHRÖDER, P., AND PINKALL, U. 2008. Conformal equivalence of triangle meshes. *ACM Trans. Graph. 27*, 3 (Aug.), 77:1–77:11.

SURAZHSKY, T., AND ELBER, G. 2002. Metamorphosis of planar parametric curves via curvature interpolation. *International Journal of Shape Modeling 8*, 02, 201–216.

SURAZHSKY, V., AND GOTSMAN, C. 2001. Controllable morphing of compatible planar triangulations. *ACM Transactions on Graphics 20*, 4, 203–231.

SURAZHSKY, V., AND GOTSMAN, C. 2003. Intrinsic morphing of compatible triangulations. *International Journal of Shape Modeling 9*, 02, 191–201.

SỲKORA, D., DINGLIANA, J., AND COLLINS, S. 2009. As-rigid-as-possible image registration for hand-drawn cartoon animations. In *Proceedings of the 7th International Symposium on Non-Photorealistic Animation and Rendering*, ACM, 25–33.

WEBER, O., AND GOTSMAN, C. 2010. Controllable conformal maps for shape deformation and interpolation. *ACM Transactions on Graphics (TOG) 29*, 4, 78.

WEBER, O., MYLES, A., AND ZORIN, D. 2012. Computing extremal quasiconformal maps. In *Computer Graphics Forum*, vol. 31, Wiley Online Library, 1679–1689.

WEINTRAUB, S. 1997. *Differential forms : a complement to vector calculus*. Academic Press, San Diego.

WHITED, B., AND ROSSIGNAC, J. 2011. Ball-morph: Definition, implementation, and comparative evaluation. *Visualization and Computer Graphics, IEEE Transactions on 17*, 6, 757–769.

WINKLER, T., DRIESEBERG, J., ALEXA, M., AND HORMANN, K. 2010. Multi-scale geometry interpolation. In *Computer Graphics Forum*, vol. 29, Wiley Online Library, 309–318.

WOLBERG, G. 1998. Image morphing: a survey. *The visual computer 14*, 8, 360–372.

XU, D., ZHANG, H., WANG, Q., AND BAO, H. 2006. Poisson shape interpolation. *Graphical models 68*, 3, 268–281.

ZENG, W., LUO, F., YAU, S., AND GU, X. 2009. Surface quasiconformal mapping by solving beltrami equations. *Mathematics of Surfaces XIII*, 391–408.

Appendix A

Theorem 7.1 *Let $M_i, i \in [1, N]$ be symmetric positive definite (SPD) matrices, $w_i, i \in [1, N]$ convex weights $w_i \geq 0, \sum_i w_i = 1$, and let $M = \sum_{i=1}^{N} w_i M_i$. Then $c(M) \leq max_{i\in[1,N]} c(M_i)$, where $c(M)$ is the l_2 matrix condition number of M.*

Proof The l_2 matrix condition number of an SPD matrix M is given by $c(M) = \lambda_1(M)/\lambda_2(M)$, where $\lambda_1(M)$ and $\lambda_2(M)$ are the largest and smallest eigenvalues of M respectively. From the variational characterization of eigenvalues, we have that $\lambda_1(M) = \max_{||x||=1} xMx^t$ and $\lambda_2(M) = \min_{||x||=1} xMx^t$. Therefore,

$$\lambda_1(M) = \max_{||x||=1} x(\sum_{i=1}^{N} w_i M_i)x^t = \\ = \max_{||x||=1} \sum_{i=1}^{N} w_i(xM_ix^t) \leq \sum_{i=1}^{N} w_i \max_{||x||=1} xM_ix^t.$$

Hence, $\lambda_1(M) \leq \sum_i w_i\lambda_1(M_i)$, and similarly $\lambda_2(M) \geq \sum_i w_i\lambda_2(M_i)$.

Since all the numbers involved are positive (as each M_i is SPD), we have:

$$c(M) \leq \frac{\sum_{i=1}^{N} w_i\lambda_1(M_i)}{\sum_{i=1}^{N} w_i\lambda_2(M_i)} \leq \sum_{i=1}^{N} u_i \frac{\lambda_1(M_i)}{\lambda_2(M_i)} = \sum_{i=1}^{N} u_i c(M_i),$$

where $u_i = w_i\lambda_2(M_i)/\sum_{j=1}^{N} w_j\lambda_2(M_j)$. Note that, $\sum_{i=1}^{N} u_i = 1$, and since M_i are SPD, $u_i \geq 0$. Therefore, the u_i are a set of convex weights, implying $c(M) \leq \max_i c(M_i)$, as required.

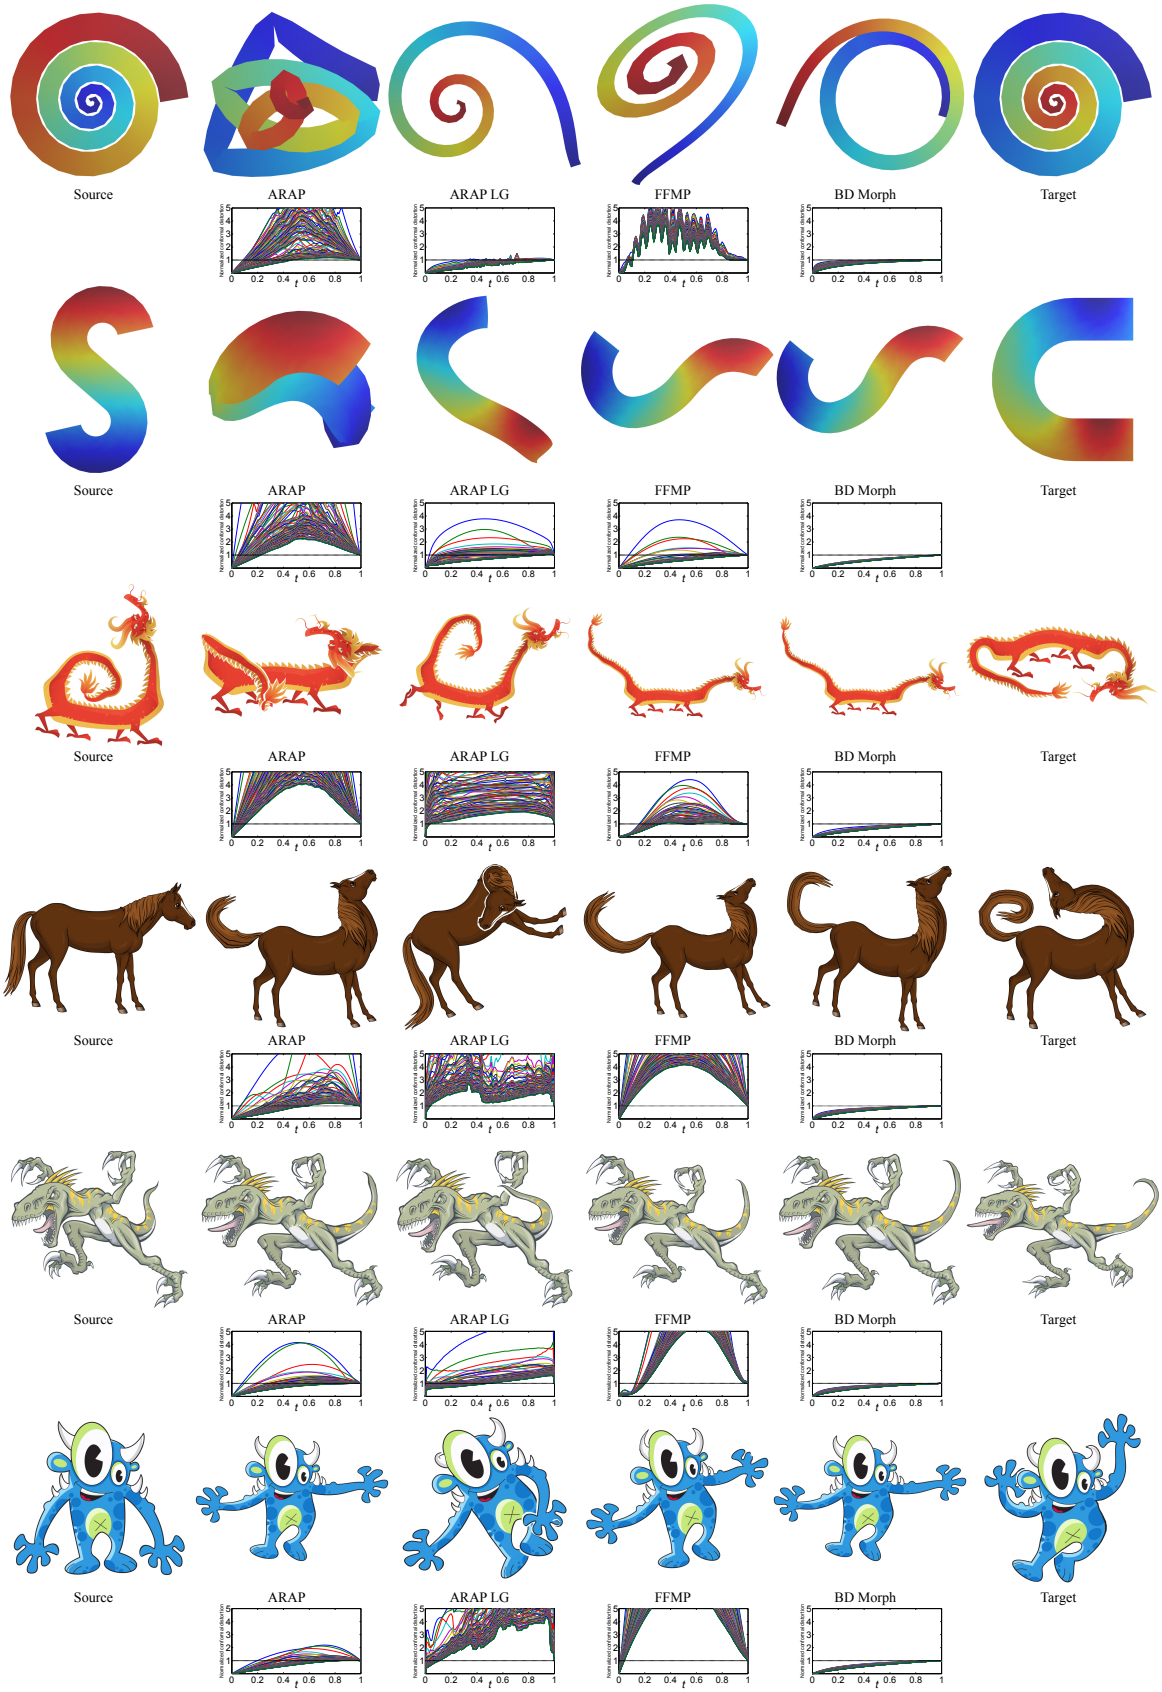

Figure 13: *Comparisons. In the first and second rows, the color coding shows the correspondence between the source and the target. For the first example we show the interpolated result for $t = 0.75$, and for all the rest $t = 0.5$. The plot shows the angular distortion (normalized to $[0, 1]$) of 100 most distorted triangles, which are found separately for each sample of t.*