DOI: 10.1111/cgf.12246
Pacific Graphics 2013
B. Levy, X. Tong, and K. Yin
(Guest Editors)

Volume 32 (2013), Number 7

As-Rigid-As-Possible Distance Field Metamorphosis

Yanlin Weng[1] Menglei Chai[1] Weiwei Xu[2] Yiying Tong[3] Kun Zhou[1]

[1]State Key Lab of CAD&CG, Zhejiang University, China
[2]Hangzhou Normal University, China
[3]Michigan State University, USA

Abstract

Widely used for morphing between objects with arbitrary topology, distance field interpolation (DFI) handles topological transition naturally without the need for correspondence or remeshing, unlike surface-based interpolation approaches. However, lack of correspondence in DFI also leads to ineffective control over the morphing process. In particular, unless the user specifies a dense set of landmarks, it is not even possible to measure the distortion of intermediate shapes during interpolation, let alone control it. To remedy such issues, we introduce an approach for establishing correspondence between the interior of two arbitrary objects, formulated as an optimal mass transport problem with a sparse set of landmarks. This correspondence enables us to compute non-rigid warping functions that better align the source and target objects as well as to incorporate local rigidity constraints to perform as-rigid-as-possible DFI. We demonstrate how our approach helps achieve flexible morphing results with a small number of landmarks.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric algorithms, languages, and systems

1. Introduction

Shape interpolation/morphing between 2D or 3D objects with arbitrary topology is of great interest for various applications. As a useful tool for smooth transition from a source object to a target object, it has long been used in animation and special effects in the film industry. There are two main approaches for the morphing problem, surface morphing and distance field interpolation (DFI). Surface morphing usually establishes one-to-one correspondence between the source and target surfaces via cross-parameterization and remeshing, and achieves as-rigid-as-possible interpolation results. It, however, has difficulty in handling objects with different topologies. DFI approaches, on the other hand, represent the shape as an isosurface, which does not require dense surface correspondences and allows topological change without special treatment.

However, without dense correspondences, DFI approaches also lack effective control over the morphing process. One solution [COSL98] is to specify sparse anchor points in the source and target objects, compute a warping function associated with the anchors, and then blend the two warped distance fields to create intermediate objects. In this manner, the morphing process is roughly controlled by the sparse anchor points. Another solution [TPG01] approximates the volume enclosed by the surface with spheres, using correspondences between source and target spheres to provide control. Unfortunately, the sparse correspondences used in these approaches are not enough to precisely align the features during the morphing process. There is no way to perform as-rigid-as-possible interpolation, which needs neighborhood correspondences.

In order to improve the controllability and quality of the DFI process, we propose a novel approach to as-rigid-as-possible distance field morphing. We first construct shape interior correspondence, a per-voxel fuzzy correspondence between two distance fields. Such a correspondence allows maintaining shape alignment throughout the morphing process, without an excessive number of landmarks. Making use of this correspondence, we then warp the source and target distance fields to an intermediate time and blend them to generate the intermediate distance field, taking into account the local rigidity constraints implied in the correspondence. The result is an as-rigid-as-possible distance field metamorphosis algorithm that can generate superior morphing sequences to those produced by previous DFI methods.

Our shape interior correspondence is computed through

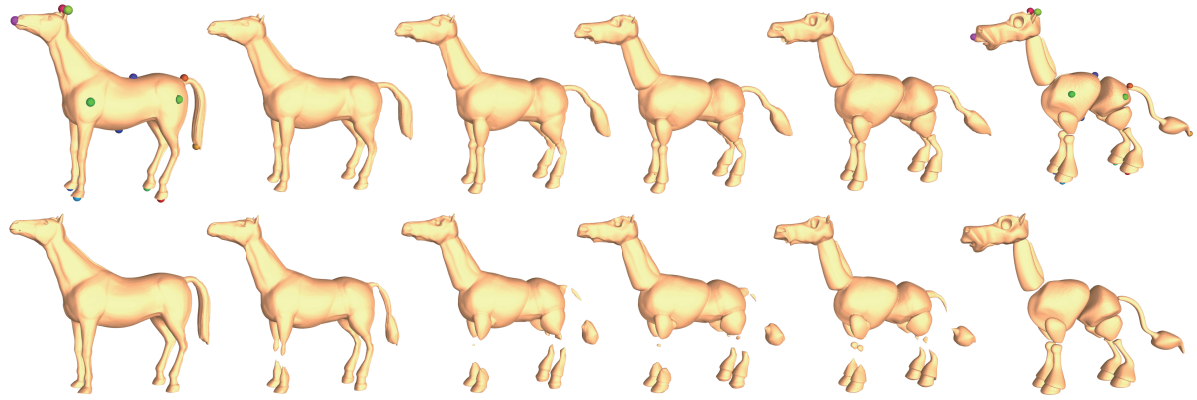

Figure 1: *Top row: a morphing sequence generated by our algorithm. Bottom row: the corresponding sequence generated using [COSL98]. The user-specified anchor points are also shown in the top row.*

minimizing earth mover's distance [Kan06], based on sparse corresponding anchor points. Thus, our initial correspondence is a fuzzy many-to-many correspondence, represented by a function storing the portion of a source voxel transported to each target voxel. With a simple weighted average, each target voxel center corresponds to a point in the source domain. This procedure avoids the difficulty in directly establishing a one-to-one correspondence between voxels, while reducing the shape distortion in the morphing as well as providing flexible control.

We tested our approach on several 2D and 3D objects with complicated shapes and topologies. In these tests, our voxel-level as-rigid-as possible interpolation scheme produced smooth morphing results with proper feature alignment, as shown in Fig. 1, Fig. 5 and Fig. 7.

1.1. Related Work

We briefly review both the surface-based and the volumetric approaches for morphing. We also discuss the mass transport problem, the technique used in our fuzzy correspondence algorithm, in the context of related areas.

Surface morphing: Surface morphing techniques usually start with building dense vertex-to-vertex correspondences and then compute vertex paths to create intermediate shapes. Since the source and target shapes are usually of different mesh connectivity, a cross parameterization method is required to remesh the shapes and build correspondence [KS04, SAPH04]. The vertex correspondences can also be established through functional correspondences on the mesh, such as harmonic coordinates [ZRKS05] and functional maps [OBCS*12].

The challenge in vertex path computation is how to reduce the shape distortion in the morphing. Sederberg et al. [SGWM93] developed an intrinsic coordinate for 2D shape blending, using edge lengths and interior angles. This method has been extended to polyhedral meshes in [SWC97]. To reduce the non-rigid distortion in surface shape blending, the as-rigid-as possible interpolation technique was introduced in [ACOL00, IMH05]. The technique proposed by Alexa et al. [ACOL00] can handle volumetric object interpolation but requires that the source and target objects have the same number of tetrahedra and the same connectivity. For 3D meshes, low-distortion surface morphing results can be achieved by changing the underlying surface representation. Laplacian coordinates [SCOL*04], rotation-invariant coordinates [LSLCO05] and Poisson shape representation [XZWB05] are popular techniques to enforce the local rigidity constraint during morphing. A recent contribution [KMP07] realized as-isometric-as-possible shape morphing results via finding a geodesic path in shape space with the designated metric.

DFI morphing: Signed distance field is an implicit representation of a surface. One of its advantages in shape morphing is that it can handle the topological change during the morphing between two surfaces with different topology without any additional processing, while a triangle-based representation requires special topological modifications for surface split and merge operations and may have to remesh some regions to maintain reasonable triangle shapes.

The early work in DFI morphing simply cross-dissolves two distance fields and reconstructs the intermediate surface from the blended distance field [PT92]. However, this simple scheme cannot preserve the essential features in the source and target objects. To improve feature preservation, a wavelet volumetric morphing technique is developed in [HWK94]. Fine control of feature alignment for two distance fields is achieved in [LGL95] by using various types of feature handles. Sphere correspondences are adopted in [TPG01] to control the DFI morphing process. The most related DFI method to our work is the 3D distance field metamorphosis by Cohen-Or et al. [COSL98]. It separates the transformation between source and target objects into rigid and non-rigid parts according to the sparse correspondence points on the surface, and the intermediate

distance field is generated by blending the warped distance fields.

As described earlier, the advantage of our algorithm is to compute a dense shape *interior* correspondence, which allows the shape alignment to be consistent with the 3D objects enclosed by the surfaces. To our knowledge, ours is the first algorithm that incorporates local rigidity constraints into DFI morhping, leading to reduced distortion.

Mass transport: We model the morphing problem as a mass transport problem, a shape transformation with a mass preservation property, where the weighted 2D or 3D volume is treated as the mass to be transported. The optimal mass transport of this sort was first considered in 1781 by Gaspard Monge, who referred to it as the "Earth Mover's Distance" (EMD). Its modern formulation first appeared in [Kan06]. Optimal mass transport has been applied to image registration [HZTA04], content-based image retrieval [RTG00] and feature similarity measure [GD05]. In [MY11], EMD is also applied to topology free 2D image morphing. However, explicit non-rigid distortion control is not directly available in this method, which hinders its application to shape morphing with the feature preservation requirement. In [LPD11, SNB*12], EMD is used to compute soft correspondences between surfaces.

1.2. Overview

We present the overall framework here, before elaborating on the algorithm in the following sections.

The two given 3D solid objects, a source Ω_0 and a target Ω_1, are represented as signed distance fields to their respective boundary surfaces, D_0 and D_1, stored on a volumetric grid. More precisely, for any point q in the domain, the value of $D_0(q)$ (or $D_1(q)$) is defined as the signed shortest Euclidean distance between q and the boundary of Ω_0 (or Ω_1), negative for points inside the objects and positive for the outside points. Our goal is to continuously deform Ω_0 to Ω_1 and produce the in-between objects $\{\Omega_t, 0 < t < 1\}$ (i.e., $\{D_t, 0 < t < 1\}$). The user can specify a set of anchor point pairs $\{(p_0^i, p_1^i), 1 \le i \le K\}$, such that point p_0^i in the source domain corresponds to point p_1^i in the target domain.

Our algorithm first establishes a fuzzy correspondence between the interiors of Ω_0 and Ω_1 (Sec. 2). Treating morphing as a process of transporting the mass of Ω_0 to Ω_1, we formulate the correspondence as the solution to an optimal mass transport problem from the interior voxels of Ω_0 to those of Ω_1. The underlying physical optimality of this correspondence makes the morphing results reasonable, in the sense that it moves the voxels with the minimum total transport cost. To produce correspondences consistent with the user-specified anchor pairs, we define a mass transportation cost induced by the anchor points.

Making use of the above correspondence, we use thin plate spline interpolation to compute warping functions $W_{0\to1}$ and $W_{1\to0}$ that can align the shapes of Ω_0 and Ω_1 as well as possible (Sec. 3.1), i.e.,

$$W_{0\to1}(\Omega_0) \approx \Omega_1, \quad W_{1\to0}(\Omega_1) \approx \Omega_0. \tag{1}$$

Next, given a morphing parameter $t \in [0,1]$, we find warping functions $W_{0\to t}$ ($W_{1\to t}$) that can transform Ω_0 (Ω_1) to time t (Sec. 3.2) as an as-rigid-as-possible interpolation between the identity transformation and $W_{0\to1}$ ($W_{1\to0}$). We then compute their inverses, the backward mapping functions $B_{t\to0}$ and $B_{t\to1}$. Finally, the distance field D_t at t can be evaluated as the interpolation of D_0 and D_1 guided by the backward mappings:

$$D_t(v) = (1-t)D_0(B_{t\to0}(v)) + tD_1(B_{t\to1}(v)), \tag{2}$$

where v is an arbitrary voxel in the domain.

Once we have the distance field, we can use the marching cubes method [LC87] to extract the zero set surface, which is represented as a triangular mesh. Note that the extracted meshes at contiguous time frames may have different number of vertices/triangles and connectivity, and it is difficult to enforce any temporal coherence between the two meshes. Fortunately, in practice, we found that the generated morphing sequences exhibit smooth transition as demonstrated in our video demo. It is also possible to use the recent surface tracking technique [BHLW12] to track the correspondence for the mesh sequence and improve the mesh coherence.

2. Shape Interior Correspondence

We formulate the desired shape interior correspondence between Ω_0 and Ω_1 as the optimal mass transport. Suppose that the total mass of a solid object is a constant value 1, evenly distributed among its interior voxels. The problem is then to find a mapping f that transports all the mass from the interior voxels of Ω_0 to those of Ω_1 with the minimal transportation cost.

Specifically, let Ω_0 have N_0 interior voxels and Ω_1 have N_1 interior voxels. We construct a complete bipartite graph, with N_0 source nodes each of which corresponds to an interior voxel of Ω_0 and N_1 target nodes each of which corresponds to an interior voxel of Ω_1. We can then regard the mapping f as an assignment of edge weights, with $f(i,j)$ representing the amount of mass transported from source voxel i to target voxel j (see Fig. 2). The goal is thus to solve for the mapping f that minimizes the total transportation cost,

$$\arg\min_f \sum_{i,j} d(i,j) \cdot f(i,j), \tag{3}$$

subject to $f(i,j) \ge 0$, $\sum_i f(i,j) = 1/N_1$ and $\sum_j f(i,j) = 1/N_0$, where $0 < i \le N_0$ and $0 < j \le N_1$. Here $d(i,j)$ is some distance between the i-th source node and the j-th target node. We solve Eq. (3) using the network simplex algorithm proposed by Bonneel et al. [BPPH11].

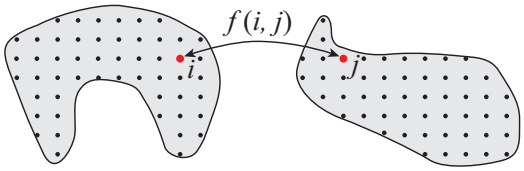

Figure 2: *Illustration of shape interior correspondence. For every source node i and target node j, $f(i,j)$ represents the amount of mass transported from source voxel i to target voxel j.*

The distance $d(i,j)$ between source nodes and target nodes should be defined with the correspondence of user-specified anchor point pairs taken into account. Intuitively, if (i,j) is an anchor point pair, $d(i,j)$ should be a minimal value to favor the mass transportation between the two nodes. Inspired by Zayer et al.'s work [ZRKS05], we compute a K-dimensional vector field in either of the source or target domains with respect to the K anchor point pairs. Specifically, for each source anchor point i, we compute a harmonic field h_0^i in the source domain with Dirichlet boundary conditions by setting its value to 1 at this anchor point and to 0 at all other source anchor points. This is equivalent to solving the following linear system in the volumetric grid

$$\Delta h_0^i = 0, \ \ with \ \ h_0^i(p_0^i) = 1, \ h_0^i(p_0^j) = 0, \ j \neq i, \qquad (4)$$

where $\Delta = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2}$ is the Laplacian operator. In our current implementation, the Laplacian at voxel v, $\Delta h_0^i(v)$, is calculated as $h_0^i(v) - \frac{1}{6}\sum_{u\in\Psi(v)} h_0^i(u)$, where $\Psi(v)$ represents v's 6-connected neighboring voxels.

Solving the above equation for all K source anchor points gives us a K-dimensional vector field $\mathbf{h}_0 = (h_0^1, ..., h_0^K)$. Similarly we can compute a K-dimensional vector field $\mathbf{h}_1 = (h_1^1, ..., h_1^K)$ in the target domain. We then define the distance between source node i and target node j as

$$d(i,j) = \|\mathbf{h}_0(v_0^i) - \mathbf{h}_1(v_1^j)\|, \qquad (5)$$

where v_0^i and v_1^j are the positions of the i-th source node and the j-th target node respectively.

Note that unlike [ZRKS05], we cannot directly use the harmonic fields to construct the correspondence as this may make exterior regions of the source correspond to interior regions of the target or vice versa. We only use harmonic fields to calculate the distance between source and target nodes. The final correspondence is computed by minimizing the total mass transportation cost.

The minimizer f of the total transportation cost gives a many-to-many correspondence: for each target node j, any source node i that has a nonzero $f(i,j)$ could partially correspond to j. To perform DFI, we need to compute a unique corresponding position in the source domain for each target node j. For this purpose, we first find the source node with the largest $f(i,j)$ with given j, i.e., $l = \arg\max_i f(i,j)$. The source corresponding position of the j-th target node, q_0^j, is then computed through

$$q_0^j = \frac{\sum_{k\in\Upsilon(l)} f(k,j) v_0^k}{\sum_{k\in\Upsilon(l)} f(k,j)}, \qquad (6)$$

where $\Upsilon(l)$ is the $3\times3\times3$ volume grid centered at source node l. We can compute the target corresponding position q_1^i of every source node i in a similar way.

Although we compute a unique corresponding position in the source for each target node, such a correspondence is not a one-to-one mapping between the source and target domains. Our intuition is that although the corresponding position of each individual node may not be optimal, the corresponding positions of all nodes can be used as a good guidance to aligning the source and target shapes. Note that as we only need a fuzzy correspondence, we can also compute q_0^j as the the average position of all source nodes having nonzero $f(i,j)$ values. According to our experiments, this hardly affects the final morphing results.

3. As-Rigid-As-Possible Distance Field Interpolation

We first describe how to compute warping functions that align the source and target shapes as well as possible, then introduce an algorithm to perform as-rigid-as-possible interpolation for these warping functions and generate the intermediate shape.

3.1. Computing Warping Functions

We employ thin plate spline (TPS) interpolation [Boo89] to compute the warping functions from the dense correspondence computed in the previous section.

Specifically, to compute the warping function $W_{1\to0}$ from the target shape to the source shape, we solve the following problem

$$W_{1\to0} = \arg\min_W \sum_j \|W(v_1^j) - q_0^j\|^2 + \lambda \int |D^2 W|^2, \qquad (7)$$

where D^2W is the matrix of second-order partial derivatives of W and the matrix norm used in $|D^2W|^2$ is the Frobenius norm. The first term preserves the interior correspondence, and the second term ensures the smoothness of the warping. The weight λ is set to 0.5 in all of our tests.

The above problem has a closed form solution when restricted to the following form:

$$W_{1\to0}(p) = a_1 + a_2 p_x + a_3 p_y + a_4 p_z + \sum_j c_i \varphi(\|p - v_1^j\|), \qquad (8)$$

where p_x, p_y and p_z indicate coordinate components of p, $\varphi(r) = r^2 \log r$ is the kernel function, and a_1, a_2, a_3, a_4 and c_j are mapping coefficients ($\in \mathbb{R}^3$) solved from Eq. (7),

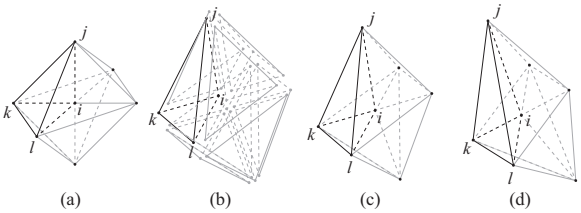

Figure 3: *Illustration of the eight trirectangular tetrahedra surrounding voxel i. For each tetrahedron, an affine transformation can be computed from its node positions in the source (a) and target (d). Interpolating the transformations of the eight tetrahedra independently will disconnect them (b). Our global optimization stitches them together again (c).*

which is simply a quadratic energy minimization with respect to these coefficients. See [Boo89] for details on computing the mapping coefficients.

The warping function from the source to the target, $W_{0\to1}$, can be computed in a similar way.

3.2. As-Rigid-As-Possible Interpolation

Given the warping functions $W_{0\to1}$ and $W_{1\to0}$ computed above, we can now warp both the source Ω_0 and the target Ω_1 to time t and compute the intermediate shape Ω_t. First, we need to compute the warping functions $W_{0\to t}$ and $W_{1\to t}$, and their backward mappings $B_{t\to0}$ and $B_{t\to1}$. Then, Ω_t (represented by D_t) is generated according to Eq. (2).

Note that $W_{0\to t}$ should be smoothly changing with t, and be the identity transformation at $t=0$ and $W_{0\to1}$ at $t=1$. At time t, each interior voxel v_i in the source domain will be transformed to $v_i^t = W_{0\to t}(v_i)$. In the following we describe an approach for computing $W_{0\to t}$ as an as-rigid-as-possible interpolation between the identity transformation and $W_{0\to1}$. We first explain how to solve for the optimal positions of all interior voxels $\mathbf{v}^t = (v_1^t, \ldots, v_{N_0}^t)$ at time t.

For each interior voxel v_i and its 6-connected neighbors in the source domain (see Fig. 3), we can construct eight trirectangular tetrahedra with voxel center v_i as their right angle. Let (v_i, v_j, v_k, v_l) be one of the tetrahedra, we can find its corresponding tetrahedra (u_i, u_j, u_k, u_l) in the target domain through $W_{0\to1}$,

$$u_m = W_{0\to1}(v_m), \quad m \in \{i,j,k,l\}. \tag{9}$$

An affine transformation defined by a 3×3 matrix M and a displacement vector b can transform the source tetrahedron to the target tetrahedron

$$u_m = Mv_m + b, \quad m \in \{i,j,k,l\}. \tag{10}$$

M can be easily solved through

$$M = [u_j - u_i \;\; u_k - u_i \;\; u_l - u_i][v_j - v_i \;\; v_k - v_i \;\; v_l - v_i]^{-1}. \tag{11}$$

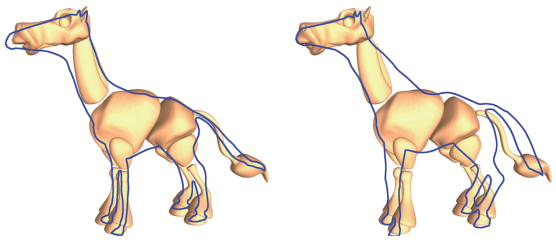

Figure 4: *A horse is warped to align with a mechanical-like horse. Left: our alignment result. Right: the alignment result by [COSL98]. We only show the silhouette of the horse for better illustration.*

Following [ACOL00, SZGP05], we decompose M (deformation gradient) through polar decomposition into a single rotation and a symmetric matrix (right stretch tensor): $M = R_r S$. Based on this decomposition, the transformation for the tetrahedron at time t can be constructed by linearly interpolating the free parameters in the factorization

$$M(t) = R_t((1-t)I + tS), \tag{12}$$

where I is an identity matrix.

Under this transformation and ignoring the translation, the positions of the tetrahedron at time t should be

$$g_m^t = M(t)v_m, \quad m \in \{i,j,k,l\}. \tag{13}$$

Now we solve for the optimal $\mathbf{v}^t$ that can match the shape of tetrahedron (v_i, v_j, v_k, v_l) at time t to the shape of the tetrahedron computed with Eq. (13) as much as possible. Specifically, we measure the similarity using the relative positions of v_j, v_k and v_l to v_i, and propose the following energy

$$E_a(v_i^t, v_j^t, v_k^t, v_l^t) = \sum_{m\in\{j,k,l\}} \|v_m^t - v_i^t - (g_m^t - g_i^t)\|^2. \tag{14}$$

The optimal positions of all interior voxels can be calculated by summing the energies for all possible tetrahedra together

$$\arg\min_{\mathbf{v}^t} \sum E_a(v_i^t, v_j^t, v_k^t, v_l^t), \tag{15}$$

where the summation is over all tetrahedra whose vertices are all interior voxels.

As E_a is a quadratic function, the above equation can be solved via least-squares optimization (using a conjugate gradient solver in our implementation). To set the boundary condition, we ask the user to specify an anchor point whose position is directly computed via linear interpolation. Alternatively, the center of mass can be forced to be static. This forms the boundary condition for Eq. (15).

Having solved $\mathbf{v}^t$, we directly apply the TPS interpolation described in Sec. 3.1 to calculate $W_{0\to t}$ and its backward mapping $B_{t\to0}$. $W_{1\to t}$ and $B_{t\to1}$ can be computed in a similar way.

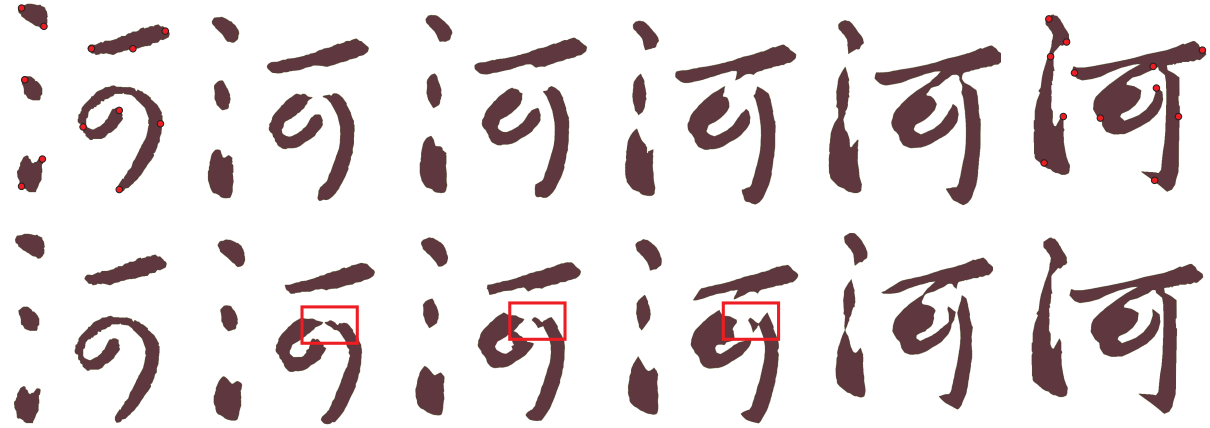

Figure 5: *Top row: a morphing sequence generated by our algorithm. Bottom row: the corresponding sequence generated using [COSL98]. The user-specified anchor points are also shown in the top row.*

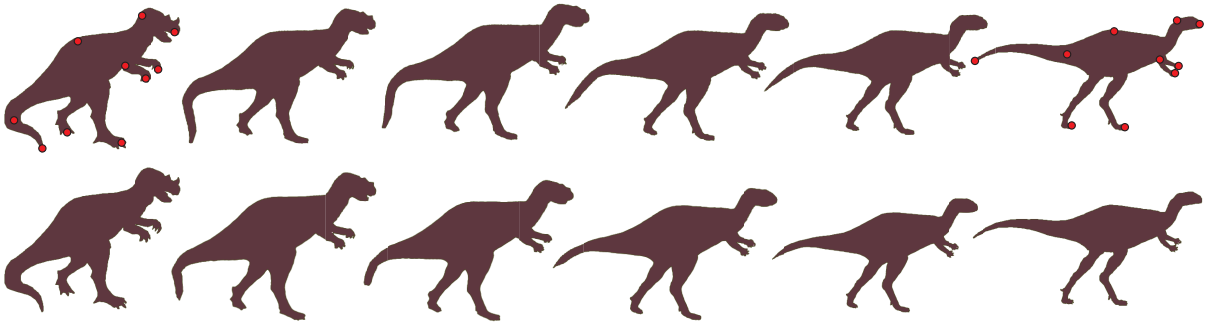

Figure 6: *Top row: a morphing sequence generated by our as-rigid-as-possible DFI. Bottom row: the corresponding sequence generated using linearly interpolated warping functions.*

We note that many tetrahedra constructed in our approach overlap with each other and do not form a tetrahedralization of the object. They are only used to impose local transformation constraints and finally produce global as-rigid-as-possible interpolation. This distinguishes our approach from [ACOL00], where an isomorphic tetrahedralization needs to be constructed for two objects with the correspondence on the boundaries already established.

4. Experimental Results

We have implemented the correspondence and DFI morphing algorithms on an Intel Xeon E5620 workstation. We provide statistics for the models presented in this paper in Table 1, including the timings and the numbers of anchor points used. We selected several examples to demonstrate the effects of the resulting algorithm. See also the accompanying video for morphing animations.

As shown in Table 1, among the stages of our approach, constructing correspondence takes the least amount of time as it is performed on grids of smaller sizes. TPS interpolation and as-rigid-as-possible DFI spend comparable time.

Example	Grid Size	#Achors	Corresp.	TPS	Interp.
Fig. 1	$180^3/40^3$	12	10	58	24
Fig. 5	$500^2/50^2$	12	1	14	1
Fig. 7	$180^3/30^3$	5	9	20	22
Fig. 9	$180^3/40^3$	13	10	62	18
Fig. 10	$180^3/40^3$	14	9	21	33
Fig. 11	$180^3/40^3$	12	10	19	27

Table 1: *The volume grid sizes for distance fields/mass transport, the number of anchor points and the timing results (in seconds) for building correspondence, TPS interpolation, and as-rigid-as-possible DFI for a single frame.*

The zero set surface extraction (i.e., marching cube) takes less than two seconds for a single frame in all examples. Overall, we can generate an animation frame in less than two minutes, with much room for acceleration as the current implementation is not well optimized.

Fig. 1 demonstrates the advantages of our dense interior correspondence. As only twelve anchors are specified by the user, some shape features of the source and target are not well aligned by the warping function computed in [COSL98] (see Fig. 4, right), producing artifacts (e.g., small blobs) dur-

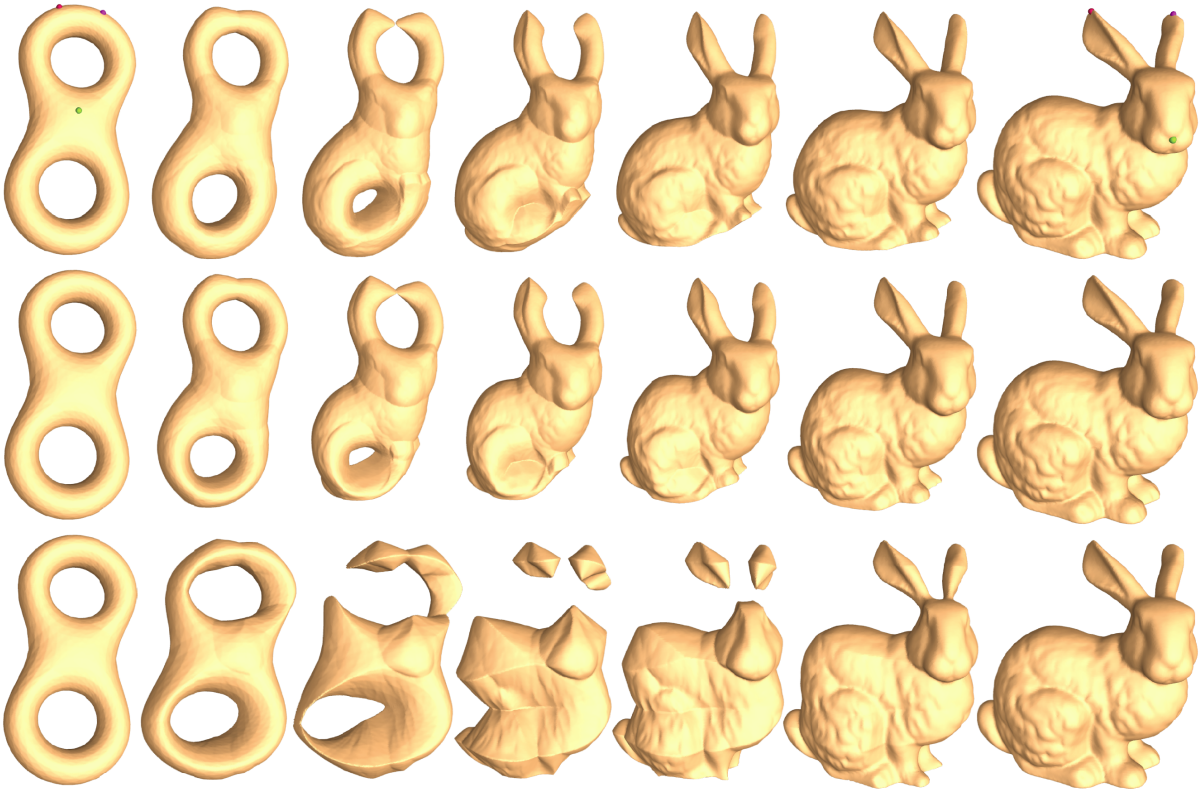

Figure 7: *Morphing a simple object of genus 2 to the Bunny model. Top row: a morphing sequence generated by our as-rigid-as-possible DFI. Middle row: the corresponding sequence generated using linearly interpolated warping functions with our interior correspondence. Bottom row: the corresponding sequence generated using [COSL98].*

ing morphing as observed in the bottom row of Fig. 1. While increasing the number of anchor points can help alleviate this problem, it becomes a tedious trial-and-error process. In contrast, using the same sparse anchor points, our algorithm can compute a dense interior correspondence that better aligns shape features such as the horse leg (see Fig. 4, left), generating superior results to those produced by [COSL98]. Fig. 5 compares our algorithm with [COSL98] in a 2D example. While judging the quality of the morphing results in this example could be subjective and different people may have different opinions, we would like to point out that our interior correspondence helps to avoid creating unpleasant ghosting features during morphing that are not exhibited in the source and target, as highlighted in the the bottom row of Fig. 5.

Our as-rigid-as-possible DFI can greatly improve the quality of the morphing results. Fig. 6 shows a 2D example. The results in both the first and second rows are generated using our dense interior correspondence, which is employed to compute the warping functions between the source and target. In the first row, we perform as-rigid-as-possible DFI, while in the second row we linearly interpolate the warping functions. A shown, as-rigid-as-possible DFI can generate more natural intermediate objects with less distortion – linear interpolation makes the dinosaur's tail shrink in the morphing process.

In Fig. 7, an object of genus two is changing smoothly to the Bunny model. Again, the sparse correspondence of five landmarks cannot align the source and target well, resulting in disconnected components in the morphing results produced by [COSL98] (see the bottom row). While linear interpolation using our computed dense interior correspondence greatly improves the results (the middle row), it introduces significant volume shrinkage in the bunny's body part, leading to more distortion than our as-rigid-as-possible DFI results (the top row).

Limitations and Discussion. Although our technique can produce visually plausible morphing animations with a small number of landmarks, we note that in order to get satisfactory results these landmarks need to be carefully placed in meaningful positions that correspond to geometric features of the source and target objects. And if the number of landmarks is too low, our approach may generate unwanted results (see Fig. 8, top row). Adding more landmarks can certainly improve the results (see Fig. 8, bottom row).

One problem with our current correspondence is that the

Figure 8: *Morphing results with different numbers of landmarks. Top row: the results of three landmarks. Bottom row: the results of five landmarks.*

computation of harmonic fields is performed in the ambient space and does not consider the shapes of the source and target objects. This may cause correspondences not in accordance with the user's perception of the shapes and generates unsatisfactory morphing results (see Fig. 9, top row). To remedy this problem, we can restrict the harmonic field computation in Sec. 2 to be within the shape interior. Specifically, we only compute the harmonic field values for interior voxels. And when calculating the Laplacian at a voxel, we only consider those interior voxels in its 6-connected neighbors. The harmonic fields computed this way lead to a better correspondence and significant improve the morphing results (see Fig. 9, bottom row). Fig. 10 and Fig. 11 show two more morphing results generated using this approach. Note that this approach works only for shapes containing a single connected component (i.e., all interior voxels of the shape are connected). For shapes with multiple disconnected components such as those shown in Fig. 1 and Fig. 5, we still have to compute harmonic fields in the ambient space.

Another problem of our work is that the computed correspondence is not a one-to-one correspondence between shapes. It is thus difficult to evaluate the quality of the correspondence as in surface parameterization where the distortion of the correspondence can be rigorously defined. While the morphing results based on this correspondence are better than previous DFI methods, we found it hard to further improve the results without using more landmarks. Furthermore, lack of one-to-one correspondence also makes our method problematic when handling colored or textured shapes. The correspondence of each individual voxel is fuzzy, which may cause unpleasant ghosting artifacts if it is directly used for color interpolation. A possible solution is to firstly generate the morphing surfaces using our algorithm, track the correspondence information through time for the surfaces using the recent method introduced in [BHLW12], and use the tracked correspondence for color interpolation.

Finally, if an object contains multiple sematic components, we have to treat the object as a single distance field and ignore the semantic information. This may generate morphing results that do not capture structure-aware effects, which are critical for some manmade objects (e.g., furniture). We would like to explore these problems in the future.

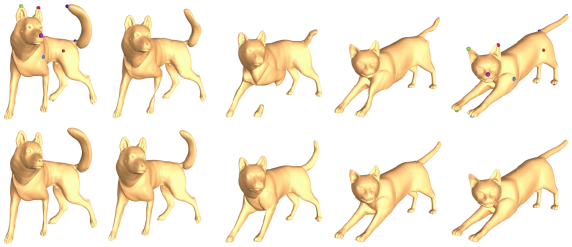

Figure 9: *Top row: a morphing sequence generated using harmonic fields computed in the whole volume. Bottom row: the corresponding sequence generated using harmonic fields computed in the shape interior.*

5. Conclusion

We have introduced a novel approach to performing as-rigid-as-possible interpolation between two distance fields. It establishes dense correspondence between the interior of two arbitrary objects based on the theory of optimal mass transportation. The resulting correspondence is consistent with the user-specified anchor pairs, enabling us to compute non-rigid warping functions that better align the source and target objects, and thus makes it possible to incorporate local rigidity constraints to perform as-rigid-as-possible DFI. We found in our tests that our approach greatly improves the quality and flexibility of DFI morphing techniques.

Acknowledgements: We would like to thank Stanford University and AIM@SHAPE for sharing their 3D models, and the anonymous reviewers for their constructive comments. The work is partially supported by NSF of China (No. 61003145, 61272305 and 61272392) and the Open Project of State Key Lab of CAD&CG.

References

[ACOL00] ALEXA M., COHEN-OR D., LEVIN D.: As-rigid-as-possible shape interpolation. In *ACM SIGGRAPH '00* (2000), pp. 157–164. 2, 5, 6

[BHLW12] BOJSEN-HANSEN M., LI H., WOJTAN C.: Tracking surfaces with evolving topology. *ACM Trans. Graph. 31*, 4 (July 2012). 3, 8

[Boo89] BOOKSTEIN F. L.: Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Trans. Pattern Anal. Mach. Intell. 11*, 6 (1989), 567–585. 4, 5

[BPPH11] BONNEEL N., PANNE M. V. D., PARIS S., HEIDRICH W.: Displacement interpolation using Lagrangian mass transport. *ACM Trans. Graph. 30*, 6 (Dec. 2011). 3

[COSL98] COHEN-OR D., SOLOMOVIC A., LEVIN D.: Three-dimensional distance field metamorphosis. *ACM Trans. Graph. 17*, 2 (Apr. 1998), 116–141. 1, 2, 5, 6, 7

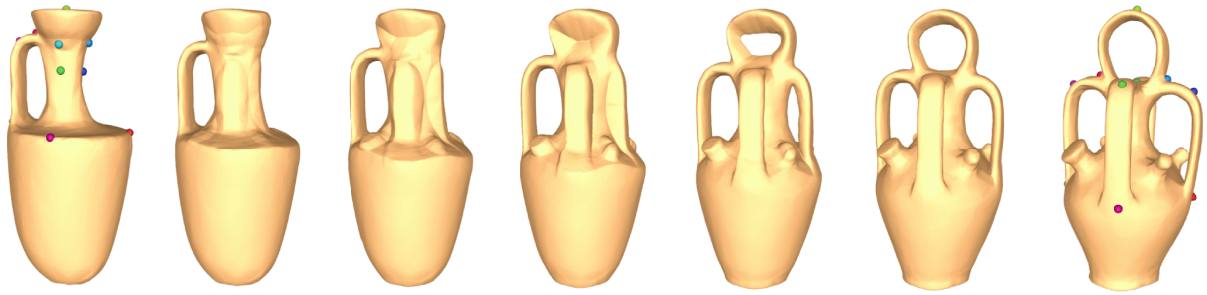

Figure 10: *Morphing a jar of genus one to another jar of genus five.*

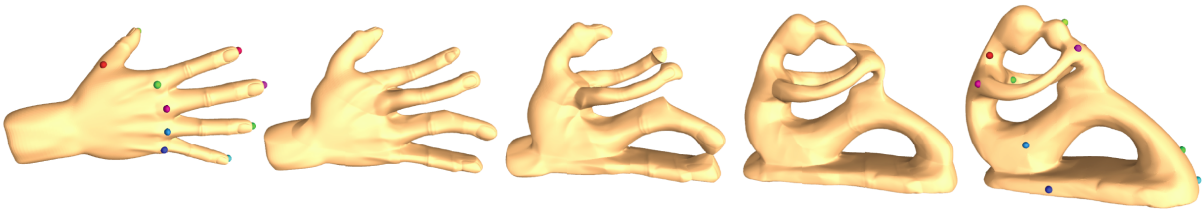

Figure 11: *Morphing a hand model of genus zero to a statue model of genus five.*

[GD05] GRAUMAN K., DARRELL T.: The pyramid match kernel: discriminative classification with sets of image features. In *ICCV '05* (2005), vol. 2, pp. 1458–1465 Vol. 2. 3

[HWK94] HE T., WANG S., KAUFMAN A.: Wavelet-based volume morphing. In *IEEE Visualization '94* (1994), pp. 85–92, CP8. 2

[HZTA04] HAKER S., ZHU L., TANNENBAUM A., ANGENENT S.: Optimal mass transport for registration and warping. *International Journal on Computer Vision 60* (2004), 225–240. 3

[IMH05] IGARASHI T., MOSCOVICH T., HUGHES J. F.: As-rigid-as-possible shape manipulation. *ACM Trans. Graph. 24*, 3 (July 2005), 1134–1141. 2

[Kan06] KANTOROVICH L.: On a problem of Monge. *Journal of Mathematical Sciences 133*, 4 (2006), 1383–1383, translated from Uspekhi Mat. Nauk, 3(2), 225ÍC226 (1948). 2, 3

[KMP07] KILIAN M., MITRA N. J., POTTMANN H.: Geometric modeling in shape space. *ACM Trans. Graph. 26*, 3 (July 2007). 2

[KS04] KRAEVOY V., SHEFFER A.: Cross-parameterization and compatible remeshing of 3d models. *ACM Trans. Graph. 23*, 3 (Aug. 2004), 861–869. 2

[LC87] LORENSEN W. E., CLINE H. E.: Marching cubes: A high resolution 3d surface construction algorithm. In *ACM SIGGRAPH '87* (1987), pp. 163–169. 3

[LGL95] LERIOS A., GARFINKLE C. D., LEVOY M.: Feature-based volume metamorphosis. In *ACM SIGGRAPH '95* (1995), pp. 449–456. 2

[LPD11] LIPMAN Y., PUENTE J., DAUBECHIES I.: Conformal wasserstein distance: Ii. computational aspects and extensions. *arXiv preprint arXiv:1103.4681* (2011). 3

[LSLCO05] LIPMAN Y., SORKINE O., LEVIN D., COHEN-OR D.: Linear rotation-invariant coordinates for meshes. *ACM Trans. Graph. 24*, 3 (July 2005), 479–487. 2

[MY11] MAKIHARA Y., YAGI Y.: Earth mover's morphing: topology-free shape morphing using cluster-based emd flows. In *ACCV '10, Volume Part IV* (2011), pp. 202–215. 3

[OBCS*12] OVSJANIKOV M., BEN-CHEN M., SOLOMON J., BUTSCHER A., GUIBAS L.: Functional maps: a flexible representation of maps between shapes. *ACM Trans. Graph. 31*, 4 (July 2012), 30:1–30:11. 2

[PT92] PAYNE B. A., TOGA A. W.: Distance field manipulation of surface models. *IEEE Comput. Graph. Appl. 12*, 1 (Jan. 1992), 65–71. 2

[RTG00] RUBNER Y., TOMASI C., GUIBAS L.: The earth mover's distance as a metric for image retrieval. *International Journal of Computer Vision 40*, 2 (2000), 99–121. 3

[SAPH04] SCHREINER J., ASIRVATHAM A., PRAUN E., HOPPE H.: Inter-surface mapping. *ACM Trans. Graph. 23*, 3 (Aug. 2004), 870–877. 2

[SCOL*04] SORKINE O., COHEN-OR D., LIPMAN Y., ALEXA M., RÖSSL C., SEIDEL H.-P.: Laplacian surface editing. In *SGP '04* (2004), pp. 175–184. 2

[SGWM93] SEDERBERG T. W., GAO P., WANG G., MU H.: 2-D shape blending: an intrinsic solution to the vertex path problem. In *ACM SIGGRAPH '93* (1993), pp. 15–18. 2

[SNB*12] SOLOMON J., NGUYEN A., BUTSCHER A., BEN-CHEN M., GUIBAS L.: Soft maps between surfaces. *Comp. Graph. Forum 31*, 5 (Aug. 2012), 1617–1626. 3

[SWC97] SUN Y. M., WANG W., CHIN F. Y. L.: Interpolating polyhedral models using intrinsic shape parameters. *The Journal of Visualization and Computer Animation 8*, 2 (1997), 81–96. 2

[SZGP05] SUMNER R. W., ZWICKER M., GOTSMAN C., POPOVIĆ J.: Mesh-based inverse kinematics. *ACM Trans. Graph. 24*, 3 (Aug. 2005), 488–495. 5

[TPG01] TREECE G., PRAGER R., GEE A.: Volume-based three-dimensional metamorphosis using region correspondence. *Visual Computer 17* (2001), 397–414. 1, 2

[XZWB05] XU D., ZHANG H., WANG Q., BAO H.: Poisson shape interpolation. In *SPM '05* (2005), pp. 267–274. 2

[ZRKS05] ZAYER R., RÖSSL C., KARNI Z., SEIDEL H.-P.: Harmonic guidance for surface deformation. *Comput. Graph. Forum 24*, 3 (2005), 601–609. 2, 4