

Mesh Morphing

Marc Alexa

Interactive Graphics Systems Group, Department of Computer Science, Technische Universität Darmstadt
Rundeturmstr. 6, 64283 Darmstadt, Germany
phone: ++49 6151 155 674, fax: ++49 6151 155 669
email: alexa@gris.informatik.tu-darmstadt.de

Abstract

Meshes have become a widespread and popular representation of models in computer graphics. Morphing techniques aim at transforming a given source shape into a target shape. Morphing techniques have various applications ranging from special effects in television and movies to medical imaging and scientific visualization. Not surprisingly, morphing techniques for meshes have received a lot of interest lately.

This state of the art report sums up recent developments in the area of mesh morphing. It presents a consistent framework to classify and compare various techniques approaching the same underlying problems from different angles.

1. Introduction

Morphing techniques transform one shape into another. With the introduction in TV and movies, morphing is nowadays known to an audience beyond the computer graphics community. At the same time, morphing has established itself as interesting research area. Recently, the focus is shifting from handling representations of space (images, volumes) to using explicit boundary representations, interpolating or blending the shape of the objects. This work is concentrating on the popular piecewise linear boundary representations, namely meshes.

Blending shapes rather than the space they are embedded in can lead to better results but is also more involved, because a proper mapping between the shapes is needed. Defining such a mapping is not trivial for two main reasons. First, it requires a parametrization of the boundary representation, and second, the mapping might involve shapes with different topology.

Besides the parameterization problem, which is fundamental in many areas dealing with meshes, morphing also requires to find suitable paths for the elements of a mesh. This part has an aesthetic component, however, several reasonable conditions should be observed, i.e. the shape should not self intersect or collapse as it varies from source to target configuration.

Traditionally, morphing is applied to two shapes: a source

and a target shape. Morphing among more than two shapes can be seen as generating elements in a space of shapes. This has interesting applications for modeling, animation, and analysis. Especially analysis using well-established methods such as the principal component analysis has gained interest lately.

This report mainly explains techniques for morphing between two meshes. This avoids some clutter in the formalism. Once all methods are explained for two meshes, possible extensions to more than two meshes and their applications are discussed.

2. Terminology & Framework

Mesh morphing techniques involve computations on the geometry as well as the topology (connectivity) of meshes. For simplicity this report concentrates on triangle meshes. In the context of morphing it seems to be acceptable to triangulate polygonal meshes prior to the application of a morphing technique. To classify and understand mesh morphing techniques it is helpful to use the now widespread terminology from Spanier⁴⁹. A mesh \mathcal{M} is described by a pair (K, V) , where K is a simplicial complex representing the connectivity of vertices, edges, and faces and $V = (v_1, \dots, v_n)$ describes the geometric positions of the vertices in \mathbb{R}^d , where typically $d = 3$.

The abstract complex K describes vertices, edges, and

faces as $\{0, 1, 2\}$ -simplices, that is, edges are pairs $\{i, j\}$, and faces are triples $\{i, j, k\}$ of vertices. The *topological realization* maps K to a simplicial complex $|K|$ in \mathbb{R}^n : The vertices are identified with the canonical basis of \mathbb{R}^n and each simplex $s \in K$ is represented as the convex hull of the points $\{e_i\} \in \mathbb{R}^n, i \in s$. Thus, each 0-simplex is a point, each 1-simplex is a line segment, and each 2-simplex is a triangle in \mathbb{R}^n .

The *geometric realization* $\phi_V(|K|)$ is a linear map of the simplicial complex $|K|$ to \mathbb{R}^d , which is defined by associating the basis vectors $e_i \in \mathbb{R}^n$ with the vertex positions $v_i \in V$. The map ϕ_V is an *embedding* if ϕ_V is bijective. The importance of an embedding is that every point p on the mesh can be uniquely represented with a barycentric coordinate b , i.e. $p = \phi_V(b)$. Such barycentric coordinates have at most three non-zero components and specify the position of a point relative to a simplex. If the point is coincident with a vertex it is a canonical basis vector, if the point lies on an edge it has two non-zero components, otherwise it has three and lies on a face.

The neighborhood ring of a vertex $\{i\}$ is the set of adjacent vertices $\mathcal{N}(i) = \{j|i, j \in K\}$ and its star is the set of incident simplices $\mathcal{S}(i) = \bigcup_{i \in s, s \in K} s$.

In the classical setting of mesh morphing two meshes $\mathcal{M}_0 = (K_0, V_0)$ and $\mathcal{M}_1 = (K_1, V_1)$ are given. The goal is to generate a family of meshes $\mathcal{M}(t) = (K, V(t)), t \in [0, 1]$ so that the shape represented by the new topology together with the geometries $V(0)$ and $V(1)$ is identical with the original shapes, i.e. $\phi_{V(0)}(|K|) = \phi_{V_0}(|K_0|)$ and $\phi_{V(1)}(|K|) = \phi_{V_1}(|K_1|)$. Most of the time the paths $V(t)$ are required to be smooth. The generation of this family of shapes is typically done in three subsequent steps:

1. Finding a correspondence between the meshes. More specifically, computing coordinates W_0, W_1 that lie on the other mesh, i.e. $W_0 \in \phi_{V_1}(|K_1|)$ and $W_1 \in \phi_{V_0}(|K_0|)$. Each coordinate in W_0, W_1 is represented as a barycentric coordinate with respect to a simplex in the other mesh. Note that ϕ_{W_0} will not map $|K_0|$ to $\phi_{V_1}(|K_1|)$ (and vice versa), as only the vertices are mapped to the other mesh but not the edges and faces. Particularly important is the alignment of automatically detected or user specified features of the meshes.
2. Generating a new, consistent mesh topology K together with two geometric positions $V(0), V(1)$ for each vertex so that the shapes of the original meshes are reproduced. The traditional morphing approach to this problem is to create a superset of the simplicial complexes K_0 and K_1 . However, remeshing techniques as used in multiresolution techniques are also attractive.
3. Creating paths $V(t), t \in]0, 1[$ for the vertices. While in general this is an aesthetic problem, several constraints seem reasonable to help in the design process. For example, in most applications the shape is not expected to

collapse or self intersect and, generally, the paths are expected to be smooth.

In the following, recent work will be explained in terms of the above mentioned problem areas. This state of the art report focuses on mesh morphing, however, if believed to be instructive also techniques dealing with polygons are discussed.

3. Correspondence of shapes

In this section we aim at finding corresponding vertex positions on two or more shapes. Given two meshes \mathcal{M}_0 and \mathcal{M}_1 , the result of this procedure is a set of barycentric coordinates B_0 so that the geometry $W_0 = \phi_{V_1}(B_0)$ of the barycentric coordinates on \mathcal{M}_1 is an embedding ϕ_{W_0} of \mathcal{M}_0 on the surface of \mathcal{M}_1 , and vice versa. The idea is that this mapping of vertices from one mesh to the other accomplishes the main part of a bijective mapping between the surfaces of \mathcal{M}_0 and \mathcal{M}_1 . After this step only the edges and faces have to be adjusted accordingly.

The process is typically done by finding a common parameter domain D for the surfaces. By mapping each surface bijectively to that parameter domain, the mapping between the shapes is established. The typical parameter domains for meshes in the context of morphing are the sphere \mathbb{S}^2 (in case the meshes are topological spheres) or a collection of topological disks represented as a piecewise linear parameter domain L . In case of the disks, the meshes have to be decomposed into isomorphic structures of disks (which requires them to be homeomorphic). A major constraint is to take into account user specified or automatically generated feature correspondences (i.e. vertex-vertex correspondences). Depending on the approach chosen, this is done by reparameterization or by decomposing the meshes according to the feature correspondence.

In case of mapping to a sphere, an embedding ϕ_S with $S = \{s_0, s_1, \dots\}, s_i \in \mathbb{R}^3, |s_i| = 1$ is computed. The embeddings on the sphere are aligned according to the feature correspondence using a bijective map f that maps spheres into spheres.

$$\begin{array}{ccc} \{i\} \in K_0 & \xrightarrow{W_0} & \phi_{V_1}(B_0) \\ \phi_{S_0} \downarrow & & \uparrow \phi_{S_1}^{-1} \\ \mathbb{S}^2 & \xrightarrow{f} & \mathbb{S}^2 \end{array}$$

The main problems in this approach are to compute the vertex coordinates S_0, S_1 on the sphere and the reparameterization f .

The decomposition approach is more general and more difficult. In addition to generate embeddings of the topological disks one has to decompose the meshes in an isomorphic way, taking possible feature correspondences into account. Formally, an abstract simplicial complex L consist-

ing of a subset of the vertices in K_0, K_1 is used as coarse approximation of both meshes:

$$\phi_{v_0}(|L|) \approx \phi_{v_0}(|K_0|), \phi_{v_1}(|L|) \approx \phi_{v_1}(|K_1|)$$

Typically, L is topological minor of K_0 as well as K_1 , i.e. it is a partition of the meshes. Vertices in K_0, K_1 are identified with a face in L and all vertices belonging to a particular face are embedded in its planar shape. Thus, the common parameter domain is the topological realization $|L|$, where each vertex is represented with a barycentric coordinate with respect to a particular face in L . This requires to embed pieces of the mesh in the plane.

$$\begin{array}{ccc} \{i\} \in K_0 & \xrightarrow{w_0} & \phi_{v_1}(B_0) \\ \phi_{L_0} \downarrow & & \uparrow \phi_{L_1}^{-1} \\ |L| & \xrightarrow{f} & |L| \end{array}$$

Following, techniques to embed simply-connected bounded and unbounded meshes in the plane and on the sphere are explained. Then, approaches to dissect the meshes into isomorphic patch-networks (or, equivalently, inducing base-domains $|L|$ on $\mathcal{M}_0, \mathcal{M}_1$) are discussed. After these basic embedding steps reparameterization for feature alignment is introduced. Finally, some comments on rarely mentioned details in the correspondence problem are given.

3.1. Parameterizing topological disks

Simply-connected parts of the boundary of three dimensional shapes are homeomorphic to a disk and, therefore, called topological disks. In order to find a parameterization of such pieces we need a bijective map of a bounded, simply connected mesh to the plane.

In our application we need to find a bijective map between patches. Thus, it is necessary to constrain the boundary of the patches to a particular shape. Here, we concentrate on mapping an arbitrary bounded and simply connected mesh to a unit disk so that boundary vertices of the mesh lie on the unit circle.

In a first step the boundary vertices are fixed on the unit circle. First, the three vertices from the base domain L are fixed in an equiangular way. This is necessary to make sure that adjacent faces in the base domain have a continuous parameterization across base domain edges. The remaining boundary vertices are fixed so that the arc lengths between neighboring vertices are proportional to the original edge lengths. The remaining (interior) vertices are free and their position is determined by a relation to neighboring vertices.

Most of the publicized approaches to solve this task boil down to solving a system of linear equations (an exception is the approach Gregory et al.²⁰).

More specifically, let $\{v_i\}$ be the vertices to be mapped

to the disk so that the free interior vertices have indices $0 \leq i < n$ and the fixed boundary vertices have indices $n \leq i < N$. We aim at finding positions w_i in the plane with $|w_i| = 1, n \leq i < N$. The mapping is bijective if and only if no edges cross. In the following we discuss three ways to define a linear system, whose solution yields positions for the vertices. In addition, the hierarchical parameterization used in MAPS³⁴ is explained.

3.1.1. Barycentric mapping

Tutte⁵³ has shown how to embed planar graphs in the plane using barycentric mapping. In our restricted setting, the idea is simply to place every interior vertex at the centroid of its neighbors:

$$w_i = \sum_{j \in \mathcal{N}(i)} \frac{1}{d_i} w_j \quad (1)$$

Setting $\Lambda = \{\lambda_{i,j}\}$ with

$$\lambda_{i,j} = \begin{cases} d_i^{-1} & \{i,j\} \in K \\ 0 & \{i,j\} \notin K \end{cases} \quad (2)$$

this can be written as the mentioned system of linear equations

$$(I - \Lambda) \begin{pmatrix} w_0 \\ w_1 \\ \dots \\ w_{n-1} \end{pmatrix} = \begin{pmatrix} \sum_{i=n}^{N-1} \lambda_{0,i} w_i \\ \sum_{i=n}^{N-1} \lambda_{1,i} w_i \\ \dots \\ \sum_{i=n}^{N-1} \lambda_{n-1,i} w_i \end{pmatrix} \quad (3)$$

The matrix $(I - \Lambda)$ has full rank and, thus, there is exactly one solution. It is easy to see that this embedding has to be valid if the fixed vertices are placed on a convex boundary: Every vertex is placed at the centroid of its neighbors, i.e. it is inside the convex hulls of all convex subgraphs. From this it follows that all neighbor rings are convex and every vertex is, indeed, placed inside the face consisting of its neighbors.

Note, that the shape of the mesh has no effect on the placement of vertices in the plane. All information for the embedding comes from K and it is clear that the embedding cannot reflect geometric properties contained in V of the mesh. In the following we try to incorporate information about the original shape.

3.1.2. Shape preserving parameterization

In the barycentric mapping the weights λ contain only topological information. Floater¹⁴ determines weights that reflect the local shape of the mesh. More precisely, the λ are so chosen that the angles and lengths of edges around a vertex are taken into account.

To compute the weights for a particular vertex v_i this vertex is placed in the origin and incident edges are laid out in the plane using the original edge lengths and angles proportional to the original angles. This is assumed to be the ideal parameterization w'_i of the mesh with respect to v_i .

The weights are computed in way that would result in placing \mathbf{w}_i in the origin if the neighbors \mathbf{w}'_j were fixed and the system of equations had to be solved. Thus, we have

$$\mathbf{w}_i = \begin{pmatrix} 0 \\ 0 \end{pmatrix} = \sum_{j \in \mathcal{N}(i)} \lambda_{i,j} \mathbf{w}'_j \quad (4)$$

and

$$1 = \sum_{j \in \mathcal{N}(i)} \lambda_{i,j}. \quad (5)$$

If \mathbf{v}_i has only three neighbors this exactly determines the positive weights, for more than three neighbors a positive solution has to be chosen from the space of possible solutions. Note that positivity results in convex combinations, which are necessary to assure a valid embedding. Floater presents a method to compute reasonable weights, which are guaranteed to be positive: Take the cyclically ordered set of neighbors $j_k \in \mathcal{N}(i)$, $k \in \mathbb{Z}_{|\mathcal{N}(i)|}$. Determine sets of weights $\lambda_{i,j}(k)$ with respect to three subsequent neighbors j_k, j_{k+1}, j_{k+2} . This yields non-negative $\lambda_{i,j}(k)$ for each k . These weights are averaged to yield the final weights:

$$\lambda_{i,j} = \frac{1}{|\mathcal{N}(i)|} \sum_k \lambda_{i,j}(k) \quad (6)$$

The positions w_i are computed by solving (3).

3.1.3. Discrete harmonic mappings

Harmonic mappings are a concept found in several fields in mathematics using differentials. Harmonic maps are often described as the function u among all functions mapping to a given domain Ω that minimize the Dirichlet energy

$$E_D(u) = \frac{1}{2} \int_{\Omega} |\nabla u|^2. \quad (7)$$

Pinkall and Polthier³⁸ show how to discretize this problem for triangles, so that weights are derived per vertex and neighbor leading to a system of linear equations of the form of Eq. (3). A somewhat clearer derivation can be found in a more recent work of Polthier³⁹. There, it is shown that the discrete Dirichlet energy is

$$E_D(u) = \frac{1}{4} \sum_{i,j \in \{i,j\} \in K} (\cot \alpha_{i,j} + \cot \beta_{i,j}) |v_i - v_j|^2, \quad (8)$$

$$\alpha_{i,j} = \angle(i, k_0, j), \beta_{i,j} = \angle(i, k_1, j), \{i, j, k_c\} \in K$$

and that the minimizer solves

$$0 = \frac{1}{2} \sum_{j \in \mathcal{N}(i)} (\cot \alpha_{i,j} + \cot \beta_{i,j}) (\mathbf{v}_i - \mathbf{v}_j) \quad (9)$$

at each vertex i . This leads to weights

$$\lambda_{i,j} = \begin{cases} \frac{\cot \alpha_{i,j} + \cot \beta_{i,j}}{\sum_{j \in \mathcal{N}(i)} (\cot \alpha_{i,j} + \cot \beta_{i,j})} & \{i, j\} \in K \\ 0 & \{i, j\} \notin K \end{cases} \quad (10)$$

which are used to obtain an embedding by solving Eq. (3).

Another formulation, which is probably better known in the graphics community, is given by Eck et al.¹².

3.1.4. Hierarchical conformal maps

The hierarchical parametrization approach in the context of MAPS³⁴ could be adapted to work on given bounded meshes. The basic idea is to simplify the mesh using vertex removal and parameterize each removed vertex. If a vertex $\{i\}$ is removed, its star is embedded in the plane using a conformal map, i.e. the edge lengths $|\mathbf{v}_i - \mathbf{v}_j|$ and relative angles at \mathbf{v}_i are preserved (i.e. as in Floater's approach). The resulting hole is triangulated in the plane and the triangle containing \mathbf{v}_i is used to yield a barycentric coordinate for $\{i\}$.

During the removal of vertices the barycentric coordinates have to be updated when the triangle containing the vertex is altered. Such changes can only occur during a vertex removal and the embedding used for that vertex removal is used to update the barycentric coordinates: All vertices parameterized with respect to a triangle in the embedding are also embedded using their barycentric coordinate. The conformal planar setting is used to retriangulate and to compute new barycentric coordinates for all embedded vertices.

3.1.5. Comparison and Conclusion

We have embedded parts of a mesh using the three approaches presented above. We have not implemented the hierarchical conformal approach since it is not particular suited if the base domain is known prior to simplification. Specifically, the point location problem in each conformal embedding adds a large constant factor to the computational complexity. Note that the solution of matrix equation (3) is computationally not more expensive than the hierarchical MPAS approach if it is solved using multigrid methods or exploiting its sparse structure.

The results of the comparison are shown in Figure 1. It is apparent that the general structure of larger and smaller triangles is very similar in all embeddings. This suggests that topology is the major factor in these type of embeddings. Changing the weights used to compute the embedding only changes the local behavior of the embedding. In tests we have found the major problem of embeddings in the plane to be the effect called area compression: Inner triangles have much less area than outer triangles. In fact, this can make parameterizations unusable due to the high ratio of areas and the limited precision of floating point numbers. It has been observed that the base domain should have enough "skin" to allow for a reasonable parametrization of the mesh.

The small differences in local shape do not seem to have much influence on the resulting correspondence of the shapes. This is even more true when local features of the shapes are aligned by reparameterization (see Section 3.4).




Figure 1: A mesh parameterized on the unit disk using different mapping techniques. The embedding in (a) is a barycentric mapping (see Section 3.1.1), which does not reflect the geometry of the mesh. The embedding in (b) tries to capture the local shape of the mesh (see Section 3.1.2) and (c) represents a discretized harmonic embedding (see Section 3.1.3).

3.2. Parameterizing topological spheres

Unbounded simply-connected 2-manifolds are called topological spheres because they are homeomorphic to spheres. A natural parameter domain for such shapes is, therefore, a unit sphere.

3.2.1. Star shapes

Kent et al.^{28, 29} were the first to present techniques to map certain classes of genus 0 meshes to a sphere. A particularly simple class of objects are convex shapes. A convex shape has the property that a straight line connecting any two boundary points of the shape lies completely inside the model. Thus, all points are visible from any interior point of the shape and a projection through an interior point onto an enclosing sphere is necessarily bijective.

A generalization of this idea extends the class of shapes to star shapes. Such shapes have at least one interior point so that straight lines connecting this interior point with boundary points lie completely inside the shape. Interior points with this property are called star points. Obviously, projecting the boundary points of a shape through a star point onto an enclosing sphere is a bijective mapping. Specifically, if point O is visible from all vertices of the mesh then translate all points so that O coincides with the origin. Then normalize all vertex coordinates. These vertex coordinates are the parameterization of the mesh vertices on a unit sphere. An illustration is given in Figure 2.

The only problem is to determine whether a shape is star shaped and if so to find a star point. For piecewise linear shapes (meshes) this can be done by intersecting halfspaces induced by the face elements of the mesh. The intersection of all halfspaces is called kernel. If the kernel is non-empty the mesh is star shaped and every point inside the (convex) kernel is a suitable star point. The kernel of a mesh in 3D can be computed in $O(n \log n)$ using standard techniques⁴¹.




Figure 2: A polygonal star shape and its projection to a circle. The kernel of a star shape is the intersection of all open half spaces over the edges (faces in case of a polyhedron). Every point in the kernel induces a bijective mapping to the circle by projection.

3.2.2. Simplification

Shapiro and Tal⁴⁷ seem to be the first to present a reliable scheme that turns arbitrary genus 0 polyhedra into convex shapes. They first simplify the shape using vertex removal until the simplified shape is a tetrahedron. Only vertices with valence 3, 4, and 5 are removed. Since the mesh is triangular such vertices always exist: It follows easily from the Euler-Poincare formulas that the average degree in any triangular (surface) mesh is less than 6. Thus, at least one vertex with degree strictly less than 6 has to exist.

Once the shape is simplified to a tetrahedron, vertices are reattach making sure that the shape stays convex. More specifically, it is shown how to attach vertices with degree 3, 4 and 5 to a convex shape so that the shape stays convex. More specifically, if a vertex $\{i\}$ has to be added to a face f , its position has to be outside the convex hull of the current mesh but inside the kernel of faces adjacent to f .

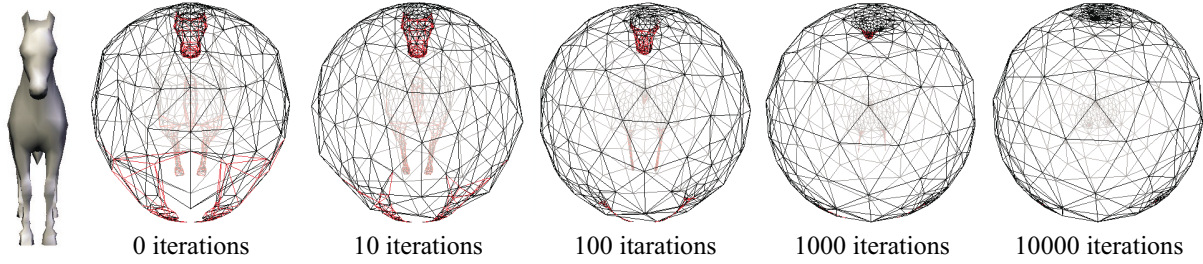


Figure 3: Embedding a polyhedral object on a sphere using relaxation. Initially, the vertices are projected through an interior point of the model onto a unit sphere. The relaxation is finished when all faces are oriented correctly. Incorrectly oriented faces are surrounded by red edges.

3.2.3. Spring embedding

Alexa¹ introduced a variation of the methods presented for planar embeddings to embed polyhedra on the unit sphere. The basic idea is the same as in barycentric mappings: Place each vertex in the centroid of its neighbors. On the sphere, however, two conditions of the planar case are violated. First, convex combinations of the neighbors' positions are not part of the domain (the sphere) and, second, no peripheral cycle is given to support the embedding.

The approach is to use a relaxation algorithm to compute the solution to the barycentric constraints. The starting configuration is generated by computing an interior point of the solid model represented by the mesh and then projecting all vertices to a sphere, which is centered at the interior point. The relaxation algorithm repeatedly places each vertex at the centroid of its neighbors. Since the centroid is not on the sphere the coordinate is normalized:

$$\mathbf{w}_i^{l+1} = \frac{\sum_{j \in \mathcal{N}(i)} \mathbf{w}_j^l}{\left\| \sum_{j \in \mathcal{N}(i)} \mathbf{w}_j^l \right\|} \quad (11)$$

The main problem of this approach is the missing of fixed vertices. In the planar case the fixed vertices avoid that all vertices collapse to one point, which is the trivial solution to (11). On the sphere, a local minimum exists such that the points are distributed over the sphere, however, the naive relaxation algorithm tends to find the global minimum, i.e. all vertices coincide.

Alexa¹ proposes to penalize long edges with a quadratic weight on the edge length. Because the collapse of the vertices into one point has to pull at least one triangle over the equator, penalizing long edges effectively prevents the vertices from collapsing. However, a simpler solution has been reported by Gumhold²¹: The sphere is recentered after each relaxation round, i.e.

$$\mathbf{w}_i^{l+1} = \frac{1}{n} \left(\sum_{j \in K} \mathbf{w}_j^l \right) - \mathbf{w}_i^l \quad (12)$$

If we want to guarantee the topological correctness of the

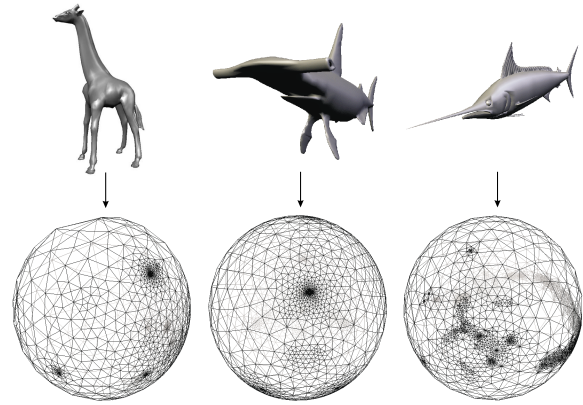


Figure 4: Sphere embeddings of the models of a giraffe, a hammerhead shark, and a swordfish.

embedding, an epsilon bound of any kind is inadequate as the only termination criterion. Instead, the process is finished only when a valid embedding is found. The embedding is valid, if and only if all faces are oriented the same, i.e. the side that was on the outside of the model is on the outside of the sphere (obviously, the surface cannot fold back upon itself without at least one triangle being upside down).

We can check this condition by testing the orientation of three consecutive vertices along the boundary of each face. Here, orientation refers to whether the three vertices make a clockwise turn on the surface of the sphere. This can be computed by evaluating $sgn((\mathbf{v}_0 \times \mathbf{v}_1) \cdot \mathbf{v}_2)$. So, the relaxation is not terminated until the orientation of each face is the same as in the original model. After the embedding is valid a conventional epsilon bound is used as the final termination criterion.

A relaxation process for the polyhedral model of a horse is depicted in Figure 3 and resulting embeddings for several models are shown in Figure 4.

3.3. Isomorphic dissection

The more general approach to establish correspondence between meshes is to dissect them into pieces. Each piece is a topological disk and can be mapped to a to the plane using one of the techniques discussed in Section 3.1. Of course, the shapes have to be split in such a way that the graphs representing the dissections have equivalent topologies.

This approach is not limited to a particular topology of the shapes, since the dissection results in a set of topological disks. However, the shapes need to be homeomorphic so that their dissections could be topologically equivalent. With extra conditions it is possible to deal also with topologically different shapes.

3.3.1. Automatic dissection of shapes

Ideally, the dissection process would not require the user to assist. However, the fully automatic dissection of two meshes into isomorphic structures seems to be a hard problem. The approach of Kanai et al.^{26, 24} uses a single patch and, thus, automatically decomposes into isomorphic structures. However, the approach is limited to genus 0 meshes and suffers from the already mentioned area compression problems in the embedding.

Several techniques exist for the dissection of a single mesh. In the context of multi resolution models several approaches require the mesh to be broken into patches. This problem is known as mesh partitioning and naturally related to graph theory. Some algorithms try to balance the size of patches (e.g., Eck et al.¹², Karypis & Kumar²⁷).

In many multi resolution methods, however, the base domain (the structure of large patches) is found by simplifying the mesh using vertex removal^{44, 43, 30, 32} or edge collapse^{22, 17, 35}.

These techniques might help in deriving a single base domain for two meshes. Lee et al.³³ use two independently established base domain to generate one base domain for both meshes. They employ their MAPS scheme³⁴ to build independent parameterizations over different base domains. These base domains are merged (see Section 4) so that the resulting merged base domain contains both independent base domains as subgraphs. Note, that in general the correspondence problem had to be solved for the geometry of the base domains. Lee et al., however, assume that the geometry of the base domains is so similar that this problem could be solved with simple heuristics (e.g. projecting in normal direction).

3.3.2. User specification of isomorphic dissections

The underlying idea of all works in this section is that the user specifies the topology of the base domain and the location of the base domain vertices on the original meshes. Tracing the edges of the base domain on the mesh is more or less done automatically.

DeCarlo and Gallier¹¹ do not assist the user specifying the edges. While this way of defining the dissection gives a lot of freedom to the user it is very time consuming.

Gregory et al.²⁰ assist the user in defining the edges. The base domain is developed while intersecting the surfaces. The user defines a pair of vertices on a mesh and the system finds a shortest path of mesh vertices connecting the defined vertices. Subsequently, feature vertices can be connected to existing feature vertices using shortest paths along the mesh. By picking corresponding vertices in the input meshes the system will construct the same graph in the input meshes. A problem could arise from the fact that only mesh vertices are used to find shortest path.

The works of Bao and Peng⁶ and Zöckler et al.⁵⁵ are similar in spirit. However, it seems that they allow to use more points to define the boundary of a patch. Points are connected with the shortest paths in the vertex-edge graph as in the work of Gregory et al.²⁰

In the approach of Kanai et al.²⁵ the user first defines a set of corresponding feature vertices. Aware of the problems resulting from using a shortest path consisting of mesh vertices the authors compute the shortest path on the piecewise linear surface connecting the feature vertices. This path may or may not coincide with vertices and edges. Since computing exact shortest path on polyhedral surfaces is difficult and time consuming they employ an approximate method that refines the original mesh and uses Dijkstra's algorithm²³.

However, even using the exact shortest path can lead to problems. Praun et al.⁴⁰ illustrate the problem and propose better solutions: If a shortest path would cross an already established edge of the base domain, the shortest possible connection avoiding the intersection is computed using a wavefront algorithm. However, also the order of vertices being connected is important, because several edges might enclose an unconnected vertex. This problem can be avoided by traversing the vertices along a spanning tree.

In our view, the underlying problem is that on non-convex and unbounded shapes more than one geodesic between two points exists on the surface. We believe that a set of these geodesics is sufficient to trace out the given topology of the base domain. To implement this, first all geodesics between connected vertices of the base domain would be computed. Then, these edges would be inspected for possible intersection. The intersection-free subset yields the decomposition of the original mesh.

3.4. Feature alignment

The necessity for aligning prominent features becomes evident even in very simple examples. Figure 5 shows two morphs between models of a young pig and a grown-up pig. In the upper sequence, no features were aligned and the resulting morph is unacceptable. The lower sequence of Figure




Figure 5: Morphs between the models of a young pig and a grown-up pig. In the upper row, no feature alignment is used, which leads to unpleasant effects (e.g., eight legs in the intermediate models). In the lower row, the eyes, ears, hoofs, and the tail are aligned (a total of 17 vertex-vertex correspondences), yielding a smooth transformation.

5 shows a morph produced with some features (ears, eyes, hoofs, and the tail) aligned. The result is obviously more pleasing. Surprisingly, the need of user guidance becomes more obvious when the shapes are similar. This is because we can envision a transformation, i.e. we expect common features of the models (head, legs, etc.) to be preserved. But this does not happen, of course, due to the different mesh topology of the models (in this example, the different mesh topologies are obvious from the different vertex counts of the models).

3.4.1. User-selected vs. shape features

A difficult task is to identify common features in several shapes. It seems impossible to automatically find such common features as they are mostly defined in a semantic and not necessarily in a geometric way. The user can identify these features and provide information about their location and correspondence (for instance as vertex-vertex correspondence of a few vertices). The algorithm should exploit this information as much as possible.

All dissection type methods explained above offer this way of user-control. Since the user explicitly chooses corresponding patches (and, therefore, corresponding edges and vertices) they can specify which parts of the meshes correspond. However, the user is also involved in other tasks, which can make the process complicated and lengthy.

The shapes' geometry also contain information useful to exploit. Several functions over the parameter domain of the meshes seem to be worth looking at. It is important that these functions are independent of the parameterization, i.e. are intrinsic to the shape and do not change if the description of the

shape is changed. Such functions are especially considered in *differential geometry*, which could be seen as exploring a shape on the shape, i.e. without a distant view. The most prominent assets for describing shapes in differential geometry are

- normals, which are independent of translation and scaling but sensitive to rotation and
- curvature (principal curvatures, mean or gaussian curvature), which is independent of translation and rotation but sensitive to scaling.

The parameterization of the shape's boundary allows to represent these quantities as a function in two variables, i.e. the normal $n : \mathbb{R}^2 \rightarrow \mathcal{S}^2$ or the gaussian curvature $c : \mathbb{R}^2 \rightarrow \mathbb{R}$.

It is clear that this information about the meshes does not lead to point to point correspondences such as user selected features. Instead the quality of the match of two shapes is quantified as a function of the distance of the shape descriptors. For example, Surahzky and Elber⁵¹ use the integral over the inner products of normals:

$$R = \int_D \langle n_1, n_2 \rangle dD \quad (13)$$

Here, the inner product between normals and the integration over the surface represent particular choices. One might choose another metric for the difference of normals as well as another method to take into account the set of differences (e.g. the maximum of the angles between normals). In order to match shapes based on such criteria the parameterization is changed so that the functional is minimized. Note that no point has an a priori optimum placement making this prob-

lem much harder to solve than aligning specified point to point correspondences.

3.4.2. Transforming to align features

As a first step in an alignment procedure the parameter domains should be transformed using affine transform to roughly align the features. Note that this is not possible for parameterizations resulting from dissection as the orientation of each patch is determined by neighboring patches.

Alexa¹ aligns a set of point to point correspondences by rotating the spherical embeddings of the mesh. The objective function to be minimized is the squared distance of corresponding points. The minimization problem can be solved using the techniques explained in Section 5.2.

3.4.3. Warping parameterizations to align features

In general, one could generate any parameterization of the meshes as a first step to establish correspondence. After this, the parameterization domain can be used to align user selected features or automatically generated features in terms of a re-parameterization of one or more of the initial parameterizations.

Alexa¹ and Zöckler et al⁵⁵ explicitly allow the user to select a set of point to point correspondences. Warping techniques similar to those used in image morphing (e.g., see the overview works of Ruprecht⁴² or Wolberg⁵⁴) are used to deform the parameterization so that corresponding points coincide. Whether the parameter domain is a disk⁵⁵ or a sphere¹ does not make a difference for the general approach.

In contrast to image warping, it is absolutely necessary that the warp does not introduce incorrectly oriented faces. This would be less of a problem if vertices as well as edges were warped. But since the algorithm later might require edge-edge intersection tests, warping the edges is impractical. Instead, edges should be (still) defined as the shortest path between vertices. That is, we warp the vertices only. Thus, even injective warping functions might introduce foldover.

Two solutions have been proposed: Alexa warps only as much as is possible with the given triangulation. If the mapping starts to introduce foldover in the triangulation the warp is made more local by adjusting the radius of influence. However, the features are not guaranteed to coincide after this process.

Zöckler et al use the foldover free warping scheme of Fujimura and Makarov¹⁶. They also warp in small steps. However, if foldover occurs they change the mesh topology to assure that the embedding stays valid. In particular, they use edge flips for this task. This changes the original triangulation of the meshes.

3.5. Conclusions

The ideal algorithm for finding a parameterization of a mesh has not been found. In general, coarse simplifications of the original meshes are accepted as useful parameter domains. In the context of morphing they are not ideal for two reasons:

- For seemingly different shapes a common base domain might be hard to find and the decomposition of the original mesh *forces* the user to interact.
- The alignment of features (e.g. shape features) is restricted to corresponding patches of the base domain.

In view of these limitations the simple solution to embed topological spheres on a unit sphere has some appeal. However, embedding complex shapes on a sphere might result in a distorted parameterization because the local ratio of surface area between sphere and original shape differs.

It seems that finding a common base domain is the method of choice. For applications, in which one base domain is needed for more than one shape, techniques should be developed that include geometric features in the decomposition process.

We still search for a reliable method that works on arbitrary input, takes any number of user-constraints into account, optimizes a reasonable resemblance of the shapes, and is sufficiently fast.

4. Representation mesh

Given two embeddings W_0, W_1 of meshes $(V_0, K_0), (V_1, K_1)$ on a common domain D we aim at generating one mesh topology K with vertex positions $V(0), V(1)$ so that the original shapes are reproduced, i.e.

$$\phi_V(0)(|K|) = \phi_{V_0}(|K_0|), \phi_V(1)(|K|) = \phi_{V_1}(|K_1|). \quad (14)$$

Note that the vertex positions $V(0), V(1)$ are already available using the barycentric coordinates of each vertex w.r.t. the base domain. These barycentric coordinates allow to map each vertex from one mesh to the other. However, the exact mapping of vertices onto the piecewise linear surface might lead to bad results. The next subsection discusses better alternatives for the absolute position of vertices.

The main point of this section is to establish the common topology K . The typical approach found in the morphing literature is to generate a supergraph of the connectivities K_0, K_1 , i.e. one that contains the simplices of both plus additional vertices if edges cross. This graph is found by *map overlay*. Here, we distinguish two cases:

1. Bounded meshes embedded in a disk.
2. Unbounded meshes, assuming the geometries of several meshes are sufficiently close.

These cases stem from the parameterization methods presented in the previous chapter.

Looking at multiresolution techniques for meshes an alternative way of generating a common topology is remeshing. In particular, the parameterization is exploited to map planar coordinates of refinement operators to coordinates on the surface of the shapes. Provided the base domain accurately represents sharp features of the meshes this approach has the advantage that it is much easier to scale. The size can be easily adapted to the desired precision. For the same reason this approach is easier to extend to more than two meshes.

4.1. Mapping parameter values to the surface

After the meshes have been parameterized it is easy to find the position of a particular vertex on the surface of a mesh. Assume we want to find the position of vertex v_{1_i} of the first mesh on the second mesh. We determine the vertices $\{w_{2_j}\}$ comprising the face in the parameterization in which the parameter domain position w_{1_i} lies. Then, w_{1_i} is represented in barycentric coordinates with respect to $\{w_{2_j}\}$:

$$w_{1_i} = \sum b_k w_{2_j(k)} \quad (15)$$

The position of v_{1_i} in the other mesh is found as

$$v_{1_i}' = \sum b_k v_{2_j(k)} \quad (16)$$

This is the exact position on the piecewise linear shape and the way used in most of the morphing literature.

However, this does not take into account the idea that piecewise linear shapes are (in most cases) just approximations of smooth shapes. Particular practical problems occur when normals have to be rebuild from these new geometric positions: Vertices inside a face get the face's normal. If standard rendering methods are used (vertex normals and Gouraud shading) this results in deteriorated shading.

It would be advantageous to find positions which result in a smooth surface. More specifically, we would like to use the barycentric coordinates to find positions *over* a triangular face and not necessarily on the face. This calls for methods defining a smooth surface from a coarse mesh. An obvious choice for such a method would be subdivision (e.g., Loop subdivision³⁶ or Kobbelts $\sqrt{3}$ -scheme³¹).

4.2. Map overlay data structure

We need a data structure to store the meshes, which allows to add and remove edges, gives quick access to topological information (e.g., the ordering of edges around a vertex), and is not too heavy in terms of storage. We choose the doubly connected edge list³⁷ (sometimes called twin-edge data structure). The basic data type of this data structure is the edge. Edges are stored as two directed half edges. More specifically, the following information is stored:

Face The face record contains a pointer to an arbitrary half edge on its boundary.




Figure 6: The doubly connected edge list.

Edge Each edge record contains pointers to

- its originating vertex,
- the face it bounds,
- the half edge connecting the same vertices but in the opposite direction (its *twin*),
- the next half edge along the boundary of the bounded face.

Vertex The vertex record contains a pointer to an arbitrary half edge originating from this vertex as well as location in space and other attributes (e.g., normal, color, texture coordinate).

Figure 6 illustrates the data structure. Note that it is particular easy to iterate along the boundaries of faces (next pointers) or through all edges incident upon a vertex in their circular order (twin \rightarrow next). A good description of the doubly connected edge list can be found in Berg et al¹⁰.

4.3. Open meshes embedded in a disk

Several algorithms were proposed for the problem of overlaying planar graphs - see a textbook¹⁰. In general, the planar map overlay has the complexity $O(n \log n + k)$, where n is the number of edges and k is the number of intersections. If the two subdivisions are connected (as in our case) the planar overlay can be computed in $O(n + k)$ ¹³.

The general paradigm for planar overlay is *plane sweep*. Sweep algorithms process the input with a virtual line moving along its normal direction. Whenever a vertex intersects the sweep line the corresponding edge is added (the vertex is the starting point of this edge) or removed (the vertex is the endpoint) from the list of active edges. The list of active edges is tested for intersection with added edges. To further reduce the number of necessary intersection tests the active edges are stored in their order along the sweep line. This is done by inserting edges in the correct position. In addition, the order has to be updated at intersection points. Using the ordering, only neighboring edges have to be tested for intersection. This processing leads to an algorithm with complexity $O(n \log n + k)$. By exploiting that two connected graphs are intersected the complexity can be reduced to $O(n + k)$.

In the case that meshes are embedded on the disk special

care has to be taken for the boundaries of the meshes. While we assume that the embedding is surjective (i.e. fills the disk), the boundary in fact is a polygon leaving small empty regions between the disk and the polygon. However, it is clear that the boundaries of the meshes to overlay should be mapped onto each other. So in order to avoid that the boundary polygons intersect with inner edges of the other mesh the boundaries have to be merged first. This is done by simply connecting the vertices of all meshes on the disk along the linear order given by the disks boundary. After this boundary polygon has been established the planar mesh overlay procedure can be computed.

4.4. Closed meshes in arbitrary position

There seem to be only a few publications about the overlay of meshes in general position. Note that plane sweep solutions are not applicable for meshes in general position. Few publications deal with overlaying two subdivisions of the sphere. Kent et al.²⁹ give an algorithm for the sphere overlay problem, which needs $O(n + k \log k)$ time. Alexa¹ has presented a solution to this particular problem, which reports the intersection of two spherical subdivisions in the optimum time of $O(n + k)$. Also, both algorithms exploit the topological properties of both subdivisions, which are used to guarantee the correct order of intersections. Here, we generalize these algorithms to work on two arbitrary shaped meshes, which are assumed to be sufficiently close to each other. We also alleviate the problem that the published version¹ had a worst case complexity of $O(n + k \log n)$ for the construction of the merged mesh using the already reported intersections.

The algorithm consists of two main parts: First, finding all intersections, and second, constructing a representation for the merged model.

4.4.1. Finding the intersections

In the algorithm two geometric functions are needed: One to decide if and where two edges intersect on the sphere, and a second to decide whether a point lies inside a face. Both geometric properties can be checked in a projection to the tangent plane of the surfaces. Since the meshes are supposed to be close in space their tangent planes should not differ to much. A suitable way of finding a common tangent plane is to take the cross product of two edges (i.e. the two edges to intersect, or two edges of the face to check).

The basic idea is to traverse the graphs breadth first, keeping information about the face that contains the current working edge and exploiting face-to-face neighbor information. Choose an arbitrary vertex $\{i\} \in K_0$ and search the 2-simplex $f = \{f_1, f_2, f_3\} \in K_1$ that contains it in under the bijective mapping. Start with an edge $e \in \mathcal{S}(i)$. Store e together with f on a stack. In general, the stack will always contain a directed edge together with the face in the other mesh containing the origin of this edge. The basic idea of the




Figure 7: Edge-edge intersections are determined by following an edge (blue in this illustration) over the faces of the other triangulation (red). After finding an intersection the face-to-face coherence exploited and only the edges of the next face are tested.

traversal is to walk over the faces following an edge (see Figure 7). Each edge $e = \{e_1, e_2\}$ is intersected first with the three edges $\{f_1, f_2\}, \{f_2, f_3\}, \{f_3, f_1\}$ bounding f , which contains $\phi_{W_0}(e_1)$. When an intersection is found the working edge e is emanating to the next face, i.e. the one that shares the intersected edge. This face is set to be f and is inspected in turn. Each edge is tested against three edges plus two additional intersection tests for each intersection being found. Thus, the algorithm has constant costs per edge and per intersection and the complexity is $O(n + k)$.

4.4.2. Generating the data structures

An appropriate data structure for storing the intersections is needed. Information about an intersection should be accessible from both intersection edges at constant costs. We use a hashtable with edge indices as key values. When edge-edge intersections are found and stored in the intersection lists a pointer to the entry in the hashtable is stored. This means, both edges point to the same data structure containing information about the intersection (the intersecting edges in the beginning). The hashtable is only needed to access the entry when the intersections are generated. After reporting all intersections it is discarded.

The following two step algorithm constructs the merged mesh: First, edges in K_0 are cutted. We iterate through the intersection list of an edge and cut the edge at each intersection point. Thus, a new edge (two half edges) are generated for each intersection. The new edge represents the part of the edge that has to be processed. At each intersection the data structure containing the respective information is updated to now contains the two parts of the edge incident upon the intersection point. At this point only the twin pointers of the half edges are updated. The next pointers are left empty.

Second, edges in K_1 are processed. As in the first step edges are cut into two pieces at each intersection point. However, this time also the next pointers are updated. This is done

by using the information stored in the intersection data structure, which now contains both edges of the already cut edge in K_1 .

After all intersections are processed in this way we have a valid vertex and edge lists of the embedding. It remains to compute the records for the faces. Note that faces created from intersecting triangles are convex polygons with 3 to 6 sides, which should be triangulated. This is another subtlety, which is more involved as it may seem: While the polygon resulting from the intersection is convex it is not clear what shape it has in other geometric configurations, e.g. those of the source meshes. In principle one should find a triangulation that is admissible in all source geometries. This might be difficult and could lead to the need for additional vertices. The problem is known as *compatible triangulation* and discussed in detail in another context in Section 5.4.1.

4.5. Remeshing

A mesh is typically just an approximation of a shape. We have already seen that the mesh overlay process together with using coordinates lying exactly on the mesh might introduce artifacts into the source meshes (see Section 4.1). Thus, even if the original mesh topologies are available as subsets of K the reproduction of the original shapes though exact is not ideal. It seems that the perfect reconstruction of the source shapes is impossible and we could as well use any mesh topology to approximate both given shapes.

Remeshing techniques have been used to construct semi-regular meshes from irregular input³⁴. The irregular mesh is reduced to an irregular base domain. The base domain is refined inserting only regular vertices. The idea is to use refinement operators as known from subdivision surfaces, however, without using the geometric rules attached to the refinement. Instead, geometric positions are found by exploiting the bijection between original surface geometry and the parameterization. For example, using the 1-4 split the parameter domain positions of inserted vertices are given as edge bisectors. This parameter leads to the coordinate on the surface of the mesh.

In the context of morphing each parameter value would lead to two coordinates. After several refinement steps a semi-regular mesh topology K is constructed together with coordinates $V(0), V(1)$ as desired. Because the refined topology is defined by the rules of the refinement used, only the base domain connectivity has to be stored explicitly.

To achieve a desired approximation accuracy, the number of refinement steps should be adapted to the geometric complexity of the meshes. Note that refinement could be done adaptively depending on the viewing conditions without necessarily computing and storing all coordinates of the refinement levels.

4.6. Comments

The remeshing approach is appealing because it allows to scale the size of the representation mesh. Its only limitation is the accurate representation of sharp features in the original shapes. In conventional multiresolution models this problem is alleviated by fitting the base domain to these features. In the context of morphing the base domain has to represent the features of several meshes, which do not necessarily coincide. This, again, incurs extra burden on the user, because a more complex base domain has to be induced on the input meshes. In addition, a more complex base domain limits the possibilities of automatic feature alignment methods. However, the flexible and lean representation mesh seem worth it.

5. Vertex paths

After the computation of one mesh topology K and two mesh geometries represented by vertex coordinates $V(0)$ and $V(1)$ it remains to compute vertex coordinates for the blended shapes. For a typical morphing animation, a set of vertex coordinates $V(t), t \in]0, 1[$ has to be generated.

A simple choice is linear interpolation^{29, 20, 55}. A rigid^{8, 9} or affine¹ transform prior to linear vertex interpolation yields better results. More complex behavior during the transform calls for more elaborate methods. Such methods decompose the shape in to linear pieces and treat these pieces separately.

5.1. Linear interpolation of vertices

The easiest way to produce blends of corresponding shapes is to interpolate the coordinates of vertices. Given a transition parameter t the coordinates of an interpolated shape are computed by

$$V(t) = (1 - t)V(0) + tV(1) \quad (17)$$

This type of interpolation produces good results if the shapes have the same orientation and are somewhat similar.

Different orientation could lead to displeasing results. Imagine two squares that are rotated by 180 degrees against each other. If simple vertex interpolation is applied in this configuration, the interpolated shapes will shrink until the shape is collapsed to one point and then grow again. This is not the desired result in most applications. It is advisable to interpolate the orientation separately from the vertex coordinates.

5.2. Interpolation of Orientation

Several ways exist to compute a relative orientation of two shapes. Note that it is difficult to interpolate the orientation of more than two shapes in 3D so the following discussion will be restricted to two shapes.

As a first step, the shapes are usually translated so that

their centers of mass coincide with the origin. Then, a rotation^{8,9} or an affine transform¹ is computed separating the rigid/affine part from the elastic part of the morph. A way of defining the rigid/affine part is to minimize the squared distances of corresponding vertices using the corresponding transform. The minimization problem of finding an affine transform can be solved using the SVD of the coordinate vector. Let the vertex vectors be arranged as a $n \times 3$ matrix

$$V = \begin{pmatrix} v_{1x} & v_{1y} & 1_z \\ v_{2x} & v_{2y} & 2_z \\ v_{3x} & v_{3y} & 3_z \\ \dots & & \end{pmatrix}.$$

Then the squared distance of coordinates under an affine transform A is

$$(V(0)A - V(1))^2 \quad (18)$$

and has to be minimized. This leads to linear system of equations. Alternatively, the least squares solution to $V(0)A \approx V(1)$ can be computed using the SVD. Let $V(0) = UDW$, where U and W are orthogonal and $D = \{d_i\}$ diagonal. Set

$$D' = \{d'_i\}, d'_i = \begin{cases} d_i^{-1} & d_i > \epsilon \\ 0 & \text{else} \end{cases} \quad (19)$$

and compute $A = V^T D' W^T V(1)$. This is the desired least squares solution¹⁸.

Intermediate shapes $V(t) = \{\mathbf{v}_1(t), \mathbf{v}_2(t), \dots\}$ are described as $V(t) = A(t)V(0)$. The question is how to define $A(t)$ reasonably? The simplest solution would be: $A(t) = (1-t)I + tA$. However, some properties of $A(t)$ seem to be desirable, calling for a more elaborate approach:

- The transformation should be symmetric.
- The rotational angle(s) and scale should change monotonic.
- The transform should not reflect.
- The resulting paths should be simple.

The basic idea is to factor A into rotations (orthogonal matrices) and scale-shear parts with positive scaling components. Alexa et al.³ have examined several decompositions. Through experimentation, they have found a decomposition into a single rotation and a symmetric matrix, to yield the visually-best transformations. This result is supported by Shoemake⁴⁸ for mathematical, as well as psychological, reasons. The decomposition can be deduced from the SVD as follows:

$$\begin{aligned} A &= R_\alpha D R_\beta = R_\alpha (R_\beta R_\beta^T) D R_\beta = \\ &= (R_\alpha R_\beta) (R_\beta^T D R_\beta) = R_\gamma S \end{aligned} \quad (20)$$

with $\det(R_\gamma) = 1$ and $s_{ii} > 0$. Based on the decomposition, $A(t)$ is computed by linearly interpolating the free parameters in the factorizations in (20), i.e.

$$A_\gamma(t) = R_{t\gamma}((1-t)I + tS). \quad (21)$$




Figure 8: Transformations of a single triangle. (a) Linear vertex interpolation. (b-d) An affine map from the source to the target triangle is computed and factored into rotational and scale-shear parts. Intermediate triangles are constructed by linearly interpolating the angle(s) of rotation, the scaling factors, and the shear parameter. (b) is generated using the SVD; (c) shows the results of reducing the overall angle of (b) by subtracting 2π from one of the angles; (d) corresponds to Equation 21 and represents the method of our choice. The last column in all rows shows plots of the vertex paths.

Figure 8 illustrates the resulting transformations for a triangle. For comparison, 8(a) shows linear interpolation of vertex coordinates. The transformation resulting from a singular value decomposition and linear interpolation $A_{\alpha,\beta}(t)$ is depicted in 8(b). Note that the result is symmetric and linear in the rotation angle but still unsatisfactory, since a rotation of more than π is unnecessary. However, if we subtract 2π from one of the angles (depicted in 8(c)) the result is even more displeasing. We have found that decomposing A into one rotation and a symmetric matrix and using $A_\gamma(t)$ yields the best results (Figure 8(d)). It avoids unnecessary rotation or shear compared to the SVD and is usually more symmetric than a QR decomposition-based approach.

5.3. Interpolation of intrinsic boundary representation

Linear interpolation of vertices can lead to undesirable effects such as shortening of parts of the boundary during the transition. To avoid such problems, Sederberg et al.⁴⁵ propose to interpolate an intrinsic representation of the boundary. For polygons, such an intrinsic representation are edge length and interior angles. Unfortunately, there is no simple analogon in 3D. An attempt was made to extend the ideas of to polyhedra⁵⁰ but the methods are computationally expensive and unreliable.

5.4. Interpolation of the interior of shapes

Shapira and Rappaport⁴⁶ suggest that a proper morph cannot be expressed merely as a boundary interpolation, but as a smooth blend of the interior of the objects. To achieve such

an interior interpolation, they represent the interior of the 2D shapes by compatible skeletons and apply the blend to the parametric description of the skeletons. An extension of this approach to meshes - though theoretically possible - has not been presented so far.




Figure 9: Contour blends of an elephant and a giraffe. Interpolation of the boundary (a) vs. a decomposition based approach (b).

Another way to represent the interior of the shapes is to decompose the shape into linear pieces or, more specifically, into simplices. The works of Floater, Gotsman, and Surahzhky^{15, 19} and Alexa et al.³ use this type of decomposition mainly for polygons, however, the extension to meshes has been demonstrated. The main difficulty in extending these approaches lies in the reliable computation of isomorphic dissections of meshes into simplicial complexes.

We will first discuss ways of generating such isomorphic complexes and then explain the possibilities they open for computing vertex paths.

5.4.1. Isomorphic Simplicial Complexes of Shapes

Simplicial complexes allow the local deformation of the shapes to be analyzed and controled. Here, we explain how to construct isomorphic dissections given two shapes with identical boundary topology.

The problem was first discussed by Aronov et al.⁵ for polygons. They offer two general approaches: The first approach is to triangulate the polygons independently and then use a piecewise linear bijective map to compute a planar overlay of the triangulations. This is somewhat similar to the planar embeddings explained in Section 3.1 together with the overlay procedures in Section 4.3. The second approach of Aronov et al. is a universal triangulation that fits every n -sided polygon. This approach is extended by Gotsman and Surahzky to generate triangulations with few interior vertices. However, it is unclear how to extend the general triangulation to meshes because of the more complex boundary topology. For this reason, we concentrate on the first approach.

It seems that the mesh-version of compatible triangulations has not been discussed in the literature. However, the procedure is conceptually the same. The meshes are tetrahedralized independently using common techniques⁷. Then, a piecewise linear bijective map is computed between the

shapes, typically using a common parameter domain. This parameter domain is used to compute an overlay of the simplicial complexes.

In case the common parameter domain for the meshes is a sphere, the interior of the sphere could be used as the parameter domain for the tetrahedra. If a piecewise linear parameter domain is used it seems more difficult to find a mapping of the interiors. If the parameter domain is given and then induced onto the meshes one could as well prescribe a simplicial base domain and induce this simplicial complex onto the (independent) simplicial complexes of the original meshes. The resulting structures are merged using a plane sweep algorithm similar to the line sweep algorithms discussed in Section 4.3.

The resulting simplicial complex might contain many, ill-shaped simplices, which cause the following determination of the vertex paths to deteriorate. For that reason, both, Gotsman and Surahzky as well as Alexa et al. try to improve the simplicial complex while preserving the isomorphy. Gotsman and Surahzky try to minimize the number of resulting simplices. Alexa et al. employ an approach motivated by meshing techniques, however, adapted to the situation that one topology has to work for two shapes.

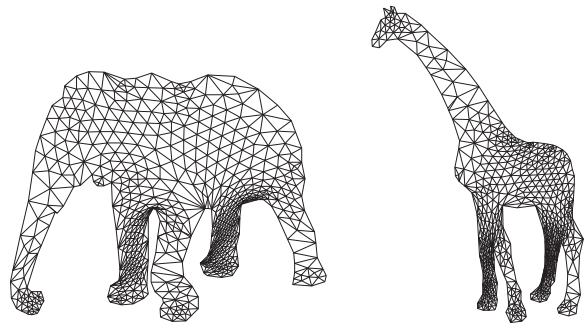


Figure 10: The homeomorphic dissections of the shapes in the elephant-giraffe example

It seems advantageous to start with Delaunay triangulations because they avoid the generation of skinny simplices, which would be inherited in the merged complex, and because the same topology is produced for similar regions, which reduces the number of extra simplices generated by the overlay. The following smoothing strategy tries to maximize the minimum angle (as the Delaunay triangulation does) by independently moving interior vertices and concurrently flipping edges. This procedure is called *compatible mesh smoothing*. If the result needs to be improved further, the vertex count is increased by means of splitting edges. The split operation is well-defined in terms of topology, if it is applied to both triangulations simultaneously, the isomorphy remains. The idea is to split long edges to avoid long skinny triangles. Figure 10 shows a result achieved with this approach.

5.4.2. Morphing barycentric coordinates

The approach of Gotsman et al. requires not only the interior of the shapes to be decomposed but also the exterior. The exterior is bounded by a common fixed convex shape. This fixed convex shape allows to represent the interior vertices with barycentric coordinates.

The idea of the approach is to linearly interpolate these barycentric coordinates. As was shown in Section 3.1.1 a convex boundary together with barycentric coordinates for the interior vertices results in a valid embedding of the complex (this is obviously not restricted to the two-dimensional case). A convex combination of the barycentric coordinate results in another barycentric coordinate so that the resulting vertex coordinates along the paths describe only valid (i.e. non-intersecting) shapes. This is a unique feature of this approach.

5.5. Composing local ideal transforms

The general idea of Alexa et al. is to find a transformation which is locally as similar as possible to the optimal transformation between each pair of corresponding simplices. The optimal simplex transform is found by factoring the affine transform defined by the pair of corresponding simplices as explained in Section 5.2. This defines ideal trajectories for each simplex.

Now consider the simplicial complex rather than a single triangle. We define an error functional for a candidate vertex configuration $V(t) = (\mathbf{v}_1, \mathbf{v}_2, \dots)$

$$E_{V(t)} = \sum_{s \in K} \|A_s(t) - B_s(t)\|^2, \quad (22)$$

where $A_s(t)$ is the desired ideal transform computed for simplex s , $B_s(t)$ is the affine transform induced from $V(0)$ to $V(t)$, and $\|\cdot\|$ is the Frobenius norm.

We define an intermediate shape $V(t)$ as the vertex configuration which minimizes this error between the desired coordinates for each individual simplex and the space of admissible coordinates.

Note that the coefficients of $B_s(t)$ are linear in $V(t)$ and that the $A_s(t)$ are known for a fixed time t . Thus, $E_{V(t)}$ is a positive quadratic form in the elements of $V(t)$. The functional $E_{V(t)}$ can be expressed in matrix form as

$$E_{V(t)} = \mathbf{v}^T \begin{pmatrix} c & G^T \\ G & H \end{pmatrix} \mathbf{v}, \quad (23)$$

where $c \in \mathbb{R}$ represents the constant, $G \in \mathbb{R}^{2n \times 1}$ the linear, and $H \in \mathbb{R}^{2n \times 2n}$ the mixed and pure quadratic coefficients of the quadratic form $E_{V(t)}$. The minimization problem is solved by setting the gradient $\nabla E_{V(t)}$ over the free variables to zero. Note that H is independent of t . In practice, we compute the LU decomposition of H and find $V(t)$ by back substitution. Furthermore, the computations are separable and are performed independently for the dimensions

of space. Note that only G depends on the dimension, while H is the same for all components. Thus, H is effectively of size $n - 1 \times n - 1$, which means the dominating factor of the computation is independent of the dimension.

The above definition has the following notable properties:

- For a given t , the solution is unique.
- The solution requires only one matrix inversion for a specific source and target shape. Every intermediate shape is found by multiplying the inverted matrix by a vector.
- The vertex path is infinitely smooth, starts exactly in the source shape, and ends exactly in the target shape. These are properties typically difficult to achieve in physically-based simulations.

Figure 11 shows transformations of some simple shapes produced with the described method.

5.6. Non-uniform interpolation

So far we have always morphed the whole mesh, i.e. the transition has been described by a scalar transition parameter t . Now, we want to locally morph certain features or regions of interest, i.e. the transition parameters are different for different vertices. We will call the set of transition parameters for vertices the *transition state*. A major problem when morphing only locally arises from the fact that corresponding features might not have the same position in space and, thus, interpolation of absolute coordinates could lead to undesirable effects. This problem is illustrated in Figure 12. The shapes in a) and b) are source and target geometry of one mesh. The idea is to locally change the geometry of the baby's face so that the nose takes the shape of the boy's. Locally interpolating vertex coordinates leads to the shape depicted in c), which is clearly not usable. Note that the faces are overall aligned in space and that the misalignment of the noses results from different relative positions in the faces.

We could ease the problem of misalignment by assigning an affine transform to a local morph. This can be done in case the mesh is represented over a base domain. The coordinates could be represented relative to the base domain. However, the fixed base domain limits the flexibility and might introduce continuity problems at the base domain edges. More generally, a shape should be defined by the transition state of its vertices. In that way, the transition states is representative for the shape of a morphable object. This could be a very compact way of representing deforming or animated objects.

The main idea to overcome the limitations is to represent vertex coordinates with respect to their neighbors in the mesh (and not with respect to some larger structure). Given a vertex and its one-neighborhood ring (see Figure 13 a), the position should be described relative to the positions of vertices in the neighborhood. Further, the representation of a vertex should be linear in the absolute coordinates. Non-linear functions tend to be numerically difficult to handle

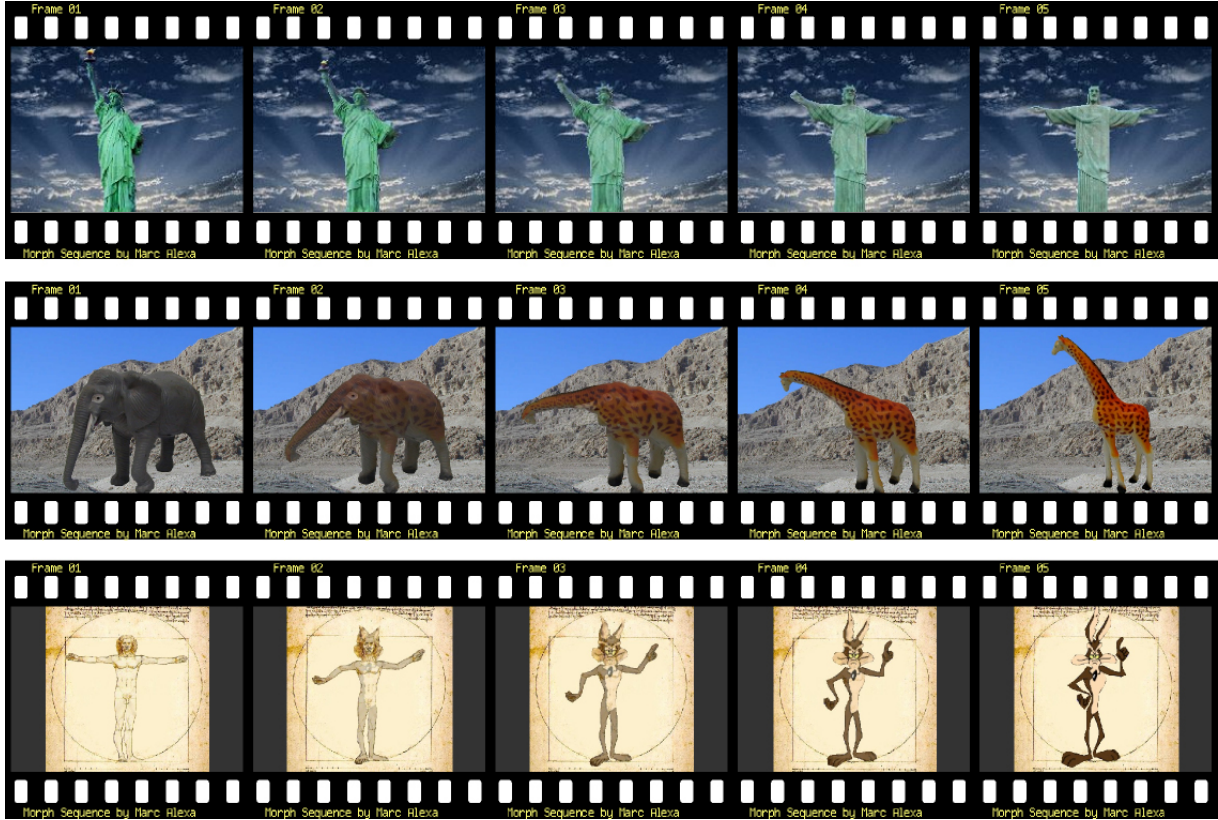


Figure 11: Transformations of different shapes representing solid objects. Note that parts of the shapes transform rigidly whenever possible.

and many morphable meshes have sliver triangles, which, together, leads to unpredictable results.

The relative representation aims at making the shape of the mesh invariant to translation or, ideally, invariant under affine transforms. If a vertex were represented in the affine space of its neighbors invariance under affine transforms would trivially follow.

5.6.1. Laplacian Representation

Alexa² uses a rather simple scheme, which is not invariant under rotation, scaling, and shearing. Assume we want to represent the position of vertex $\{i\}$. Compute the center of mass of the neighbors

$$\bar{\mathbf{v}}_i = \frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} \mathbf{v}_j \quad (24)$$

and let the new representation be the difference of this center of mass to the original position:

$$\tilde{\mathbf{v}}_i = \mathbf{v}_i - \bar{\mathbf{v}} \quad (25)$$

For an illustration see Figure 13. If we write all vertices as a

vector the forward transformation (from absolute to relative coordinates) can be represented in matrix form. Let A be the adjacency matrix of the mesh and D be a diagonal matrix with $d_{ii} = 1/|\mathcal{N}(i)|$. The transform is represented by $L = I - DA$. Note that L is a Laplacian of the mesh⁵². This is an important observation as it generalizes the approach to shape representations other than meshes, e.g. parametric or implicit functions.

The backward transformation (from relative to absolute coordinates) is, by construction, not unique. It should be uniquely determined up to a translation. This means, $L \in \mathbb{R}^{m \times m}$ should have rank $m - 1$, which is indeed so².

The main idea of this approach is to morph by linearly interpolating Laplacian coordinates rather than absolute coordinates. Since Laplacian coordinates are linear in absolute coordinates morphing the whole shape (i.e. all vertices have the same transition state) will be the same in absolute and Laplacian coordinates. Yet, if the desired transitions are different for subsets of vertices interpolating Laplacian coordinates yields more reasonable results.




Figure 12: Given a mesh with two geometries a) and b) so that corresponding features (eyes, ears, nose, mouth, etc.) are represented by the same vertices in both geometries. If one feature (in this example the nose) is morphed towards the target geometry in absolute coordinates, different positions in space lead to undesirable effects shown in c). The shape in d) shows a more pleasing result achieved by interpolating a differential encoding of the vertices.




Figure 13: A vertex (black) and its neighborhood ring (white) in a). In Laplacian coordinates a vertex is represented by the difference to the centroid of its neighbors (b).

Alexa² gives details on how to generate and represent transition states and how to solve the resulting linear system.

6. More than two meshes

The conceptual extension of the framework to more shapes is straightforward. Given meshes $\mathcal{M}_i = (V_i, K_i)$ a common topology K together with vertex sets $V(\mathbf{e}_i)$ is established. The vertex sets form a base of a space, which is reflected by using canonical base vectors \mathbf{e}_i as indices. A morphed shape $(V(\mathbf{s}), K)$ is represented by a vector $\mathbf{s} = (s_0, s_1, \dots)$ reflecting the shares of the meshes $\mathcal{M}_0, \mathcal{M}_1, \dots$

Not all techniques presented in this framework are equally suited to be extended to more meshes. The correspondence problem discussed in Section 3 seems to be relatively easy

to extend. All meshes are embedded in the given parameter domain, which leads to barycentric representation of the original vertices. If each set of original vertices V_i needs to be mapped to all other meshes $\mathcal{M}_j, i \neq j$ the complexity would grow quadratically with the number of meshes. However, this is not necessary if a remeshing strategy is used to generate a consistent mesh topology (see Section 4.5). This procedure generates the same set of vertices over all shapes, thus, the complexity is linear in the number of meshes times the number of vertices used in the remesh, which is the best we can expect. Concluding, the best way to generate the set $\{(V(\mathbf{e}_i), K)\}$ is to embed all meshes in a common parameter domain (spherical or piecewise linear) and then remesh to the desired accuracy.

The vertex path problem now extends to compute combinations of several vertex vectors. Linear vertex combination is easily extended:

$$V(\mathbf{s}) = \sum_i s_i V(\mathbf{e}_i) \quad (26)$$

Surprisingly, any technique involving rotations such as the ones explained in Sections 5.2 and 5.4 seem to be difficult to extend. Instead of interpolating the orientation one could compute the principal components (moments) of the shapes and align them with the canonical axes of the coordinate system. To extend the local morph approach explained in Section 5.5 the linear combination has to be applied to the Laplacian coordinates.

Applications of such spaces of meshes range from modeling and analysis of shapes to animation. Praun et al. have termed the synthesis-analysis part digital geometry processing (DGP)⁴⁰. Modeling could be achieved by combining several shape (features) to yield the desired result. Using techniques such as the principal component analysis, spectral properties for the mesh family can be explored.

The space of meshes $(V(\mathbf{e}_i), K)$ allows to represent animations as a curve $\mathbf{s}(t)$. A classical key frame animation with k key frames could be simply modeled as a k -dimensional space, where the curve linearly interpolates subsequent key frames. Alexa and Müller⁴ use the PCA together with rigid motion detection to find a more compact space (see an illustration in Figure 14). In this space the main part of the animation is stored in the rigid motion of the first base vector. The additional base vectors are sorted according to their energy in the spectrum of the animation. This allows to progressively store and stream mesh animations, where the progressiveness is with respect to movements and not model fidelity. The understanding of certain features of the animation as bases of a linear space gives this representation semantics. It is possible to identify e.g. the smile in a facial animation with a particular basis and, thus, to modify only the smile without the need to work on all key frames.

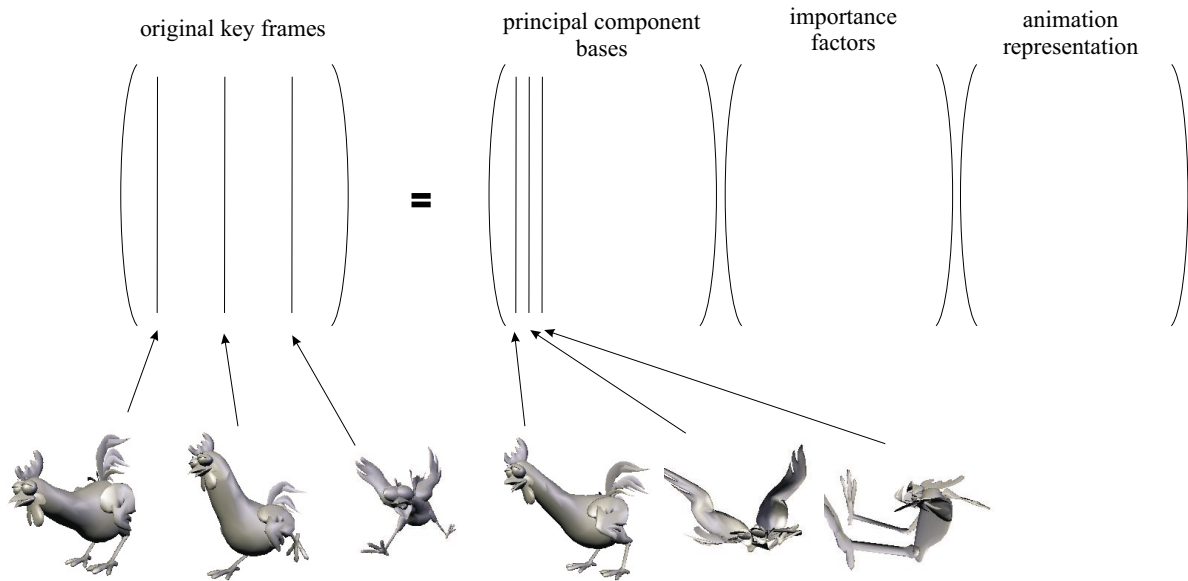


Figure 14: The SVD applied to a space of meshes. Here the original space represents a key frame animation, the result allows to represent the same animation with less geometries containing most of the original animation's energy.

7. Conclusions

Mesh morphing has reached a state where basic problems are solved, yet, a practically working system is not available. The correspondence and representation problems can be seen as the application of several techniques now common in multiresolution representations and modeling. The vertex path problem is specific to morphing applications and leaves room for improvement.

What is still missing is a robust implementation of current techniques. As with other geometric techniques, many of the approaches suffer from numerical problems. Many "detail" problems such as normal and texture coordinate interpolation can cause trouble in practice.

With the extension of mesh morphing to linear spaces of meshes several interesting avenues for future research arise. However, one might ask the question whether meshes (with fixed connectivity) are the right representation for deforming shapes, anyway.

Acknowledgements

This work contains (rephrased) parts of the author's prior works. The encouraging and insightful discussion with many knowledgeable people, namely Daniel Cohen-Or, Jose Encarnação, Craig Gotsman, David Levin, Wolfgang Müller, are gratefully acknowledged.

References

1. Marc Alexa. Merging polyhedral shapes with scattered features. *The Visual Computer*, 16(1):26–37, 2000. ISSN 0178-2789.
2. Marc Alexa. Local control for mesh morphing. *Shape Modeling International '01 Proceedings*, 2001.
3. Marc Alexa, Daniel Cohen-Or, and David Levin. As-rigid-as-possible shape interpolation. *Proceedings of SIGGRAPH 2000*, pages 157–164, July 2000. ISBN 1-58113-208-5.
4. Marc Alexa and Wolfgang Müller. Representing animations by principal components. *Computer Graphics Forum*, 19(3):411–418, August 2000. ISSN 1067-7055.
5. Boris Aronov, Raimund Seidel, and Diane Souvaine. On compatible triangulations of simple polygons. *Computational Geometry: Theory and Applications*, 3:27–35, 1993.
6. Hujun Bao and Qunsheng Peng. Interactive 3d morphing. *Computer Graphics Forum*, 17(3):23–30, 1998. ISSN 1067-7055.
7. B. Chazelle and L. Palios. Triangulating a non-convex polytope. In Kurt Mehlhorn, editor, *Proceedings of the 5th Annual Symposium on Computational Geometry (SCG '89)*, page 393, Saarbrücken, FRG, June 1989. ACM Press.
8. Daniel Cohen-Or and Eyal Carmel. Warp-guided

- object-space morphing. *The Visual Computer*, 13(9-10):465–478, 1998. ISSN 0178-2789.
9. Daniel Cohen-Or, Amira Solomovici, and David Levin. Three-dimensional distance field metamorphosis. *ACM Transactions on Graphics*, 17(2):116–141, April 1998. ISSN 0730-0301.
 10. Mark de Berg, Mark van Kreveld, Mark Overmars, and Otfried Schwarzkopf. *Computational Geometry – Algorithms and Applications*. Springer-Verlag, Berlin Heidelberg, 1997.
 11. Douglas DeCarlo and Jean Gallier. Topological evolution of surfaces. *Graphics Interface '96*, pages 194–203, May 1996. ISBN 0-9695338-5-3.
 12. Matthias Eck, Tony DeRose, Tom Duchamp, Hugues Hoppe, Michael Lounsbery, and Werner Stuetzle. Multiresolution analysis of arbitrary meshes. *Proceedings of SIGGRAPH 95*, pages 173–182, August 1995. ISBN 0-201-84776-0. Held in Los Angeles, California.
 13. U. Finke and A. Hinrichs. Overlaying simply connected planar subdivisions in linear time. In *Proceedings of the 11th Annual Symposium on Computational Geometry*, pages 119–126, New York, NY, USA, June 1995. ACM Press.
 14. Michael S. Floater. Parametrization and smooth approximation of surface triangulations. *Computer Aided Geometric Design*, 14(3):231–250, 1997. ISSN 0167-8396.
 15. Michael S. Floater and Craig Gotsman. How to morph tilings injectively. *Journal of Computational and Applied Mathematics*, 101:117–129, 1999.
 16. Kikuo Fujimura and Mihail Makarov. Folder-free image warping. *Graphical Models and Image Processing*, 60(2):100–111, March 1998.
 17. Michael Garland and Paul S. Heckbert. Surface simplification using quadric error metrics. *Proceedings of SIGGRAPH 97*, pages 209–216, August 1997. ISBN 0-89791-896-7. Held in Los Angeles, California.
 18. Gene H. Golub and Charles F. Van Loan. *Matrix Computations*, volume 3 of *Johns Hopkins Series in the Mathematical Sciences*. The Johns Hopkins University Press, Baltimore, MD, USA, second edition, 1989. Second edition.
 19. Craig Gotsman and Vitaly Surazhsky. Guaranteed intersection-free polygon morphing. *Computers & Graphics*, 25(1):67–75, February 2001. ISSN 0097-8493.
 20. A. Gregory, A. State, M. Lin, D. Manocha, and M. Livingston. Feature-based surface decomposition for correspondence and morphing between polyhedra. *Computer Animation '98*, June 1998. Held in Philadelphia, Pennsylvania, USA.
 21. Stefan Gumhold. Personal communication on embedding meshes, 2000.
 22. Hugues Hoppe. Progressive meshes. *Proceedings of SIGGRAPH 96*, pages 99–108, August 1996. ISBN 0-201-94800-1. Held in New Orleans, Louisiana.
 23. Takashi Kanai and Hiromasa Suzuki. Approximate shortest path on a polyhedral surface based on selective refinement of the discrete graph and its applications. *Proc. Geometric Modeling and Processing 2000*, pages 241–250, 2000.
 24. Takashi Kanai, Hiromasa Suzuki, and Fumihiko Kimura. Three-dimensional geometric metamorphosis based on harmonic maps. *The Visual Computer*, 14(4):166–176, 1998. ISSN 0178-2789.
 25. Takashi Kanai, Hiromasa Suzuki, and Fumihiko Kimura. Metamorphosis of arbitrary triangular meshes. *IEEE Computer Graphics & Applications*, 20(2):62–75, March/April 2000. ISSN 0272-1716.
 26. Takashi Kanai, Hiroshima Suzuki, and Fumihiko Kimura. 3d geometric metamorphosis based on harmonic map. *Pacific Graphics '97*, October 1997. Held in Seoul, Korea.
 27. G. Karypis and V. Kumar. Multilevel k -way hypergraph partitioning. Technical Report 98–036, Department of Computer Science and Engineering, University of Minnesota, Minneapolis, MN 55455, 1998.
 28. James Kent, Richard Parent, and Wayne E. Carlson. Establishing correspondences by topological merging: A new approach to 3-d shape transformation. *Graphics Interface '91*, pages 271–278, June 1991.
 29. James R. Kent, Wayne E. Carlson, and Richard E. Parent. Shape transformation for polyhedral objects. *Computer Graphics (Proceedings of SIGGRAPH 92)*, 26(2):47–54, July 1992. ISBN 0-201-51585-7. Held in Chicago, Illinois.
 30. Reinhard Klein. Multiresolution representations for surfaces meshes based on the vertex decimation method. *Computers & Graphics*, 22(1):13–26, February 1998. ISSN 0097-8493.
 31. Leif Kobbelt. sqrt(3) subdivision. *Proceedings of SIGGRAPH 2000*, pages 103–112, July 2000. ISBN 1-58113-208-5.
 32. Leif Kobbelt, Swen Campagna, and Hans-Peter Seidel. A general framework for mesh decimation. *Graphics Interface '98*, pages 43–50, June 1998. ISBN 0-9695338-6-1.
 33. Aaron Lee, David Dobkin, Wim Sweldens, and Peter Schröder. Multiresolution mesh morphing. *Proceedings of SIGGRAPH 99*, pages 343–350, August 1999. ISBN 0-20148-560-5. Held in Los Angeles, California.

34. Aaron W. F. Lee, Wim Sweldens, Peter Schröder, Lawrence Cowsar, and David Dobkin. Maps: Multiresolution adaptive parameterization of surfaces. *Proceedings of SIGGRAPH 98*, pages 95–104, July 1998. ISBN 0-89791-999-8. Held in Orlando, Florida.
35. Peter Lindstrom and Greg Turk. Fast and memory efficient polygonal simplification. *IEEE Visualization '98*, pages 279–286, October 1998. ISBN 0-8186-9176-X.
36. Charles Loop and Tony DeRose. Generalized b-spline surfaces of arbitrary topology. *Computer Graphics (Proceedings of SIGGRAPH 90)*, 24(4):347–356, August 1990. ISBN 0-201-50933-4. Held in Dallas, Texas.
37. David E. Muller and Franco P. Preparata. Finding the intersection of two convex polyhedra. *Theoretical Computer Science*, 7(2):217–236, 1978.
38. Ulrich Pinkall and Konrad Polthier. Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics*, 2(1):15–36, 1993.
39. Konrad Polthier. Conjugate harmonic maps and minimal surfaces. Technical Report Preprint No. 446, TU Berlin, SFB 288, 2000.
40. Emil Praun, Peter Schröder, and Wim Sweldens. Consistent mesh parameterizations. *to be published in Proceedings of SIGGRAPH 2001*, 2001.
41. Franco P. Preparata and Michael Ian Shamos. *Computational Geometry: An Introduction*. Texts and Monographs in Computer Science. Springer-Verlag, Berlin, Germany, 1985.
42. D. Ruprecht and H. Muller. Image warping with scattered data interpolation. *IEEE Computer Graphics & Applications*, 15(2):37–43, March 1995.
43. William J. Schroeder. A topology modifying progressive decimation algorithm. *IEEE Visualization '97*, pages 205–212, November 1997. ISBN 0-58113-011-2.
44. William J. Schroeder, Jonathan A. Zarge, and William E. Lorensen. Decimation of triangle meshes. *Computer Graphics (Proceedings of SIGGRAPH 92)*, 26(2):65–70, July 1992. ISBN 0-201-51585-7. Held in Chicago, Illinois.
45. Thomas W. Sederberg, Peisheng Gao, Guojin Wang, and Hong Mu. 2d shape blending: An intrinsic solution to the vertex path problem. *Proceedings of SIGGRAPH 93*, pages 15–18, August 1993. ISBN 0-201-58889-7. Held in Anaheim, California.
46. M. Shapira and A. Rappoport. Shape blending using the star-skeleton representation. *IEEE Computer Graphics & Applications*, 15(2):44–50, March 1995.
47. Avner Shapiro and Ayellet Tal. Polyhedron realization for shape transformation. *The Visual Computer*, 14(8-9):429–444, 1998. ISSN 0178-2789.
48. Ken Shoemake and Tom Duff. Matrix animation and polar decomposition. *Graphics Interface '92*, pages 258–264, May 1992.
49. Edwin H. Spanier. *Algebraic Topology*. McGraw-Hill, New York, 1966.
50. Yue Man Sun, Wenping Wang, and Francis Y. L. Chin. Interpolating polyhedral models using intrinsic shape parameters. *The Journal of Visualization and Computer Animation*, 8(2):81–96, April-June 1997. ISSN 1049-8907.
51. Tatiana Surazhsky and Gershon Elber. Matching free form surfaces. *Computers & Graphics*, 26(1):??–??, 2001. ISSN 0097-8493.
52. Gabriel Taubin. A signal processing approach to fair surface design. *Proceedings of SIGGRAPH 95*, pages 351–358, August 1995. ISBN 0-201-84776-0. Held in Los Angeles, California.
53. W. T. Tutte. How to draw a graph. *Proc. London Mathematical Society*, 13:743–768, 1963.
54. George Wolberg. Image morphing: a survey. *The Visual Computer*, 14(8-9):360–372, 1998. ISSN 0178-2789.
55. Malte Zöckler, Detlev Stalling, and Hans-Christian Hege. Fast and intuitive generation of geometric shape transitions. *The Visual Computer*, 16(5):241–253, 2000. ISSN 0178-2789.

Marc Alexa / Mesh Morphing