On Coherent Rotation Angles for As-Rigid-As-Possible Shape Interpolation

Jaeil Choi and Andrzej Szymczak
College of Computing, Georgia Tech

Abstract

The goal of morphing algorithms is to construct visually pleasing interpolations between 2D or 3D shapes, referred to as morphs in computer graphics circles. One of the desirable properties of morphing algorithms is avoiding self-intersections of the deforming shape. This paper investigates winding number-based topological invariants of planar morphs between compatible triangulations that do allow global self-intersections but avoid local ones. Equivalently, such morphs are not allowed to make any of the triangles degenerate. We show a relationship between those invariants and the as-rigid-as-possible shape interpolation algorithm and improve it so that it allows to handle cases when a large amount of twist is required to transform the source triangulation into the target triangulation.

1 Background

Morphing is a well-established problem a number of researchers have been working on for many years. Because of lack of space, we limit ourselves to discussion of only the results that apply to the same setting as assumed in this paper, namely morphing between compatible planar triangulations.

Different variants of a method of deforming planar triangulations while avoiding self-intersections is discussed in papers [2, 3, 4, 6]. Its theoretical foundation is the theory of discrete Laplacians, which, in particular, states that if the boundary of a planar mesh is a convex polygon and every vertex of a mesh is a weighted average of the neighboring vertices with positive weights then the triangulation is intersection free. Essentially, those methods represent the locations of the vertices of the triangulations in an implicit manner as collections of the weights and construct the morphing by deforming the set of weights for the source triangulation to the set of weights for the target triangulation.

Although avoiding self-intersections is certainly important, preservation of the shape is at least equally important in practice. Contrary to the work described above, which is centered on avoiding self-intersections, the as-rigid-as-possible shape interpolation of [1] is centered around the idea of preserving the shape. It allows to construct very well looking morphs (in particular, preserve the lengths and overall shape of the parts of the shapes that need to be deformed). However, it does not come with any guarantee of no self-intersections in the morph.

2 Overview of the main results

This section presents an informal discussion and motivation for our results. Their completely precise formulations can be found in the long version of the paper.

In what follows, we will consider triangulations in the plane that do not allow local self-intersections. Formally, we shall think of the source and target triangulations as geometric realizations of a connected and oriented 2D manifold (with boundary) abstract simplicial complex. Compatible triangulations are realizations of the same abstract simplicial complex.

Let us note that, apart from planar triangulations with no self-intersections or non-manifold vertices our definition encapsulates the two triangulations shown on the right of Figure 1. Also, notice that all triangulations shown in the Figure are compatible. It is not hard to see that there exists a morph between the first two which does not make any triangle in the intermediate triangulations degenerate. However, the third one is impossible to deform to any of the other two without a degeneracy. A simple argument can be based on the concept of kink-free deformations of polygonal loops [5], which is a discrete analogue of deformations of smooth regular closed curves discussed in [7]. In the smooth case, deformations of smooth and regular loops preserve the winding number of the tangent vector. For kink-free deformations of polygonal loops (i.e. deformations that are not allowed to make any three consecutive vertices collinear), a similar statement is true. Lack of continuity of the tangent vector is circumvented by linearly interpolating between the tangent vectors of the consecutive line segments of the polygonal loop and treating the the loop obtained by joining all of these linear paths as the loop of tangent vectors. The triangulations shown in Figure 1 are circular triangle strips. By joining the midpoints of edges shared by each pair of consecutive triangles one obtains three paths shown in Figure 1. It is not hard to see that the winding number of the tangent vector to the two paths on the left is different from the winding number for the third path. Hence there is no kink-free deformation of the first or second path to the third. Therefore, there can't be a non-degenerating morph of the triangulation on the right of Figure 1 to any of the other two. If it existed, by joining the midpoints of the intervals between each pair of consecutive triangles in the intermediate stages of the morph we would obtain a forbidden kink-free deformation.

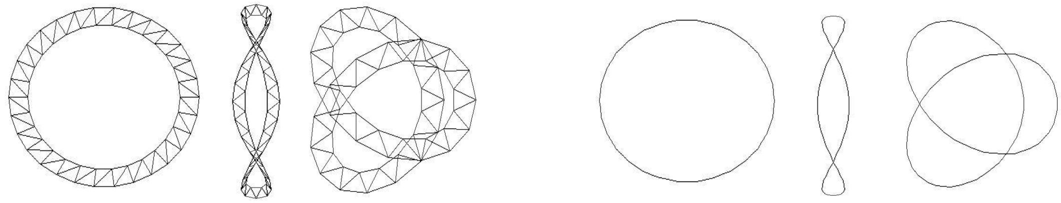

Figure 1: Three triangulations and their associated paths.

The as-rigid-as-possible shape interpolation [1] method takes as input two compatible triangulations, (which we shall call the source and target), realizing an abstract simplicial complex $\mathcal{C}$ and proceeds in two steps. First, for every triangle T in $\mathcal{C}$, it selects a deformation of the 'shape' of the triangle corresponding to T in the source to that of the corresponding triangle in the target. This deformation is described by a path in the space of all orientation-preserving automorphisms of the plane (which we denote by $GL_2^+(\mathbf{R})$), starting at the identity and ending at the transformation M_T, which takes the vectors running along the edges of the triangle corresponding to T in the source into the vectors running along the corresponding edges the target. Clearly, there are infinitely many loops of this kind (there are even infinitely many their homotopy classes). The paper [1] includes experimental evidence that paths constructed in the following way produce naturally looking deformations. Decompose M_T into a product of a rotation matrix, $R_{\alpha(T)}$ (where $\alpha(T)$ is the rotation angle) and a positive definite symmetric matrix S_T. Then, define the path $\sigma_T : [0,1] \to$ by

$$\sigma_T(t) = R_{t\alpha(T)} \circ ((1-t)I + tS_T).$$

The triangulation interpolating between the source ($t = 0$) and the target ($t = 1$) corresponding to parameter $t \in [0,1]$ is computed by minimizing (in the least squares sense) the difference between the entries of the matrices of $\sigma_T(t)$ and the matrix of the transformation that takes vectors running along edges of the triangle corresponding to T in the source to vectors running along corresponding edges in the unknown interpolating triangulation. Since all entries of the matrix of the latter transformation can be expressed as linear combinations of the unknowns (coordinates of vertices of the interpolating triangulation), this boils down to solving a sparse global quadratic optimization problem. The reader is referred to [1] for details.

The angle $\alpha(T)$ can be selected in infinitely many ways, the admissible choices differing by a multiple of 2π from each other. Careful choice of $\alpha(T)$ allows to control the amount of spin performed by the path $\sigma(T)$ on the triangle T, or, more precisely, the homotopy class of this path. Note that $GL_2^+(\mathbf{R})$ is homotopy equivalent to the circle, which means that distinguishing between homotopy classes of loops in that space is relatively simple. The proofs of our statements heavily depend on the topological structure of $GL_2^+(\mathbf{R})$. Intuitively, it is clear that the adjacent triangles of the triangulations should spin in a similar manner so that their deformation paths can be gracefully compromised when global optimization step is performed. This suggests that selecting rotation angles in such a way that they differ by less than π for adjacent triangles might be a good idea, leading to the following recursive algorithm.

Algorithm 1 *For an arbitrarily selected triangle T_0, select any admissible rotation angle $\alpha(T_0)$. Any time a triangle T has a rotation angle assigned, assign rotation angles in the interval $(\alpha(T) - \pi, \alpha(T) + \pi)$ to its adjacent triangles (which have not got a rotation angle yet), until all triangles are exhausted.*

Note that we used an open interval $(\alpha(T) - \pi, \alpha(T) + \pi)$ in the statement of the algorithm because admissible rotations of adjacent triangles T and T' cannot differ by π. The vector $\vec{v}$ running along the edge separating the two triangles in the source triangulation is mapped to the vector $\vec{v}'$ along the edge separating the corresponding triangles in the target triangulation by *both* M_T and $M_{T'}$. The difference of the rotation angles $\alpha(T)$ and $\alpha(T')$ (modulo 2π) has to be equal to the angle between $S_T(\vec{v})$ and $S_{T'}(\vec{v})$, which has to be strictly less than π because the matrices of S_* are symmetric and positive definite.

The success of Algorithm 1 turns out to depend on global properties of the source and target triangulations. For example, it is not going to work if the source and the target triangulations are the triangulations on the left and right of Figure 1. Assume it starts at some triangle of the mesh on the left and assigns the rotation angles to triangles in the counterclockwise order. Just before terminating, it is going to assign a rotation angle to the clockwise neighbor of the initial triangle. It turns out that the rotation angles assigned to the two triangles will be off by more than π. The reason is that the algorithm gradually accumulates the spin that has to be applied to triangles in the source to deform them to corresponding triangles of the target.

However, the algorithm does succeed when a non-degenerating morph between the source and the target exists. It also succeeds when both source and target do not have self intersections. We claim that these two cases are most important for applications. When finding a non-degenerating morph is not

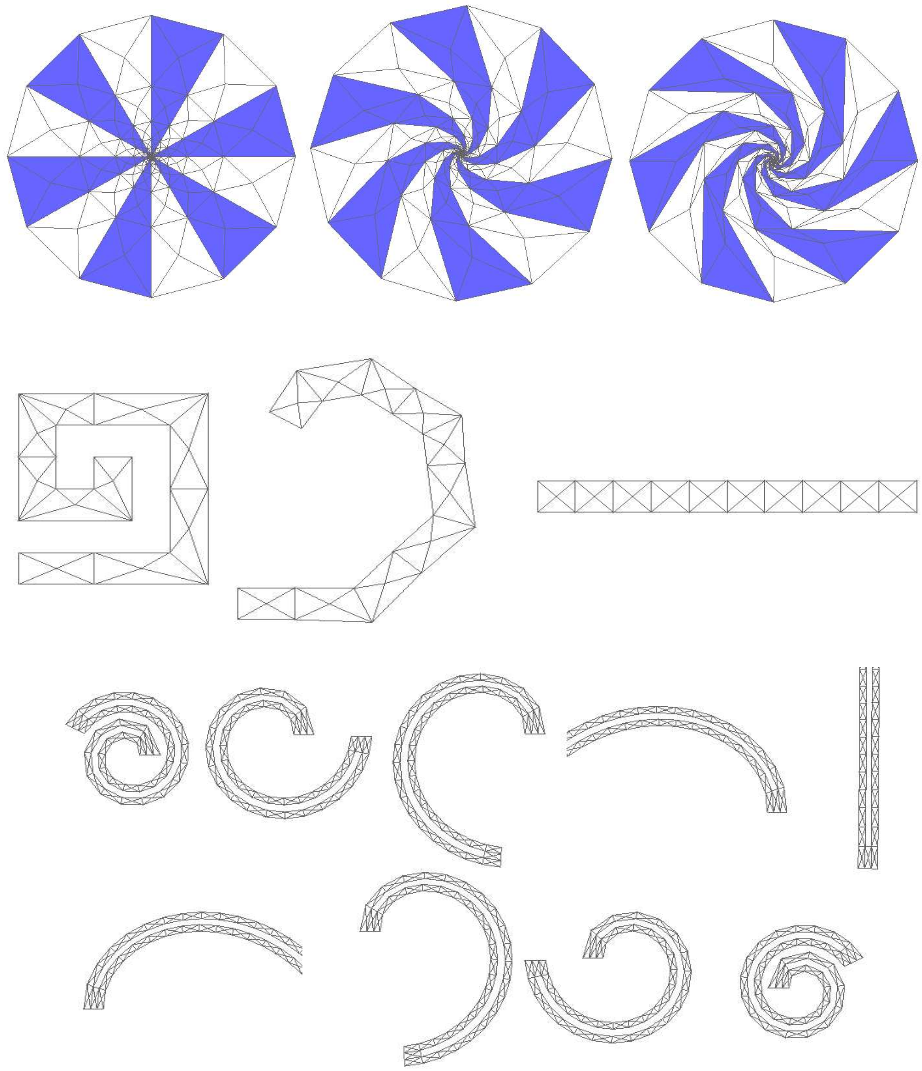

Figure 2: Samples of morphs obtained with coherent rotation angles (some of the snapshots of the hollow spiral have been truncated).

possible, then one should probably look for a different compatible triangulations of the source and target before constructing a high quality morph. The most natural of triangulations of planar shapes will be non self-intersecting.

As was mentioned before, selecting $\alpha(T)$ means selecting a homotopy class for the path σ_T. If a non-degenerating morph of the source to the target exists, then one can prove that it is possible to find one in which, for any triangle T, the path τ_T in $GL_2^+(\mathbf{R})$ defined so that $\tau_T(t)$ maps the vectors running along T in the source into the vectors running along the corresponding edges in the interpolating triangulation for parameter t is homotopic to the path σ_T constructed based on the angle produced by the algorithm. This means that our algorithm can be interpreted as computing the *homotopy classes of deformations of individual triangles in a non-degenerating morph between the source and the target meshes*.

Figure 2 show examples of morphs obtained using as-rigid-as-possible morphing with coherent rotation angles. Figure 3 shows some of the intermediate stages of the morphs with the same source and target where all rotation angles have simply been selected from the interval $[-\pi, \pi]$ (leading to degeneracies in the morph). Note that in all our examples, the differences between rotation angles for some triangles exceed 2π. Choosing rotation angles from a different interval of length 2π would not allow to avoid degeneracies.

3 Sketch of proof

We limit ourselves to the proof that Algorithm 1 produces coherent rotation angles provided the source and target triangulations can be deformed in a non-degenerating fashion.

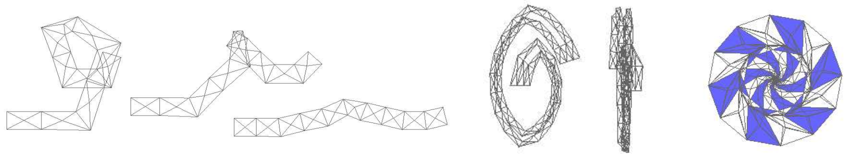

Figure 3: Samples of intermediate stages of the morphs between the same source and target meshes as in Figure 2, but without coherent angle selection.

By a triangulation we mean a planar geometric realization of connected, manifold (with boundary) oriented abstract 2D simplicial complexes (in what follows, briefly called admissible complexes). Te realization must preserve the orientation, i.e. map vertices of a triangle ordered consistently with the orientation into a counterclockwise ordering of vertices of the corresponding 2D triangle. Formally, a geometric realization is a function F that specifies 2D location of each of the vertices of the complex.

Let $\mathcal{C}$ be an admissible complex. With a loop L in the dual graph of $\mathcal{C}$ (the vertices of the dual graph correspond to the triangles of $\mathcal{C}$ and edges join adjacent triangles), we associate a planar loop depending on a geometric realization F of $\mathcal{C}$ by joining the midpoints of edges shared by consecutive pairs of triangles on L. Examples of such planar loops are shown in Figure 1. By the *winding number along* L we mean the winding number of the tangent vector of that planar loop (interpolated to form a continuous curve as described in the previous section). The argument based on kink-free deformations we have used in Section 2 proves that *the winding numbers along any loop is preserved by a non-degenerating morph.* In particular, the winding numbers along any loop L in the dual graph must be the same for any two compatible triangulations that are possible to morph in a non-degenerating fashion. In what follows, the loop obtained by interpolating the consecutive tangent vectors will play an important role. We denote it by $\tau_{F,L}$.

Let F and G be geometric realizations of the same complex $\mathcal{C}$. For any loop L in the dual graph of $\mathcal{C}$ we will associate a loop Σ_L in $GL_2^+(\mathbf{R})$ defined as follows. For each triangle T on L, consider the linear isomorphism M_T defined as in Section 2. The loop Σ_L is obtained by joining the isomorphisms M_T for T along L with linear paths (all the transformations on these paths can be proved to belong to $GL_2^+(\mathbf{R})$). One can prove that the loop in $\mathbf{R}^2 \setminus \{0\}$ obtained by evaluating $\Sigma_L(t)$ on $\tau_{F,L}(t)$ for $t \in [0,1]$ is homotopy equivalent to the loop $\tau_{G,L}$ (if the parametrization of the two loops is chosen carefully, the angles between vectors on both paths are always strictly less than π). Since we assumed the two realizations are possible to morph in a non-degenerate manner, the loops $\tau_{F,L}$ and $\tau_{G,L}$ are homotopic. Based on the structure of $GL_2^+(\mathbf{R})$ we conclude that the loop Σ_L is contractible in that space.

Algorithm 1 assigns rotation angles along paths originating from the initial triangle T_0. One can prove that for any triangle T_1 the homotopy class of the path σ_{T_1} is the same as the homotopy class of the path which first follows the path σ_{T_0} and then follows the piecewise linear path joining the M_T transformations for consecutive triangles T on the path from T_0 to T_1 along which the algorithm has constructed the rotation angles. The algorithm fails if the rotation angles (or homotopy classes of σ_{T_1}) computed for T_1 using different paths are different for some triangle T_1. However, this would mean that the loop Σ_L for the loop L formed by the two paths is not contractible in $GL_2^+(\mathbf{R})$ and contradict the statement proved in the previous paragraph.

References

[1] M.Alexa D.Cohen-Or and D.Levin, *As-Rigid-As-Possible Shape Interpolation.* Proc. SIGGRAPH 2000, pp.157–164.

[2] M.Floater and C.Gotsman, *How to Morph Tilings Injectively.* Journal of Computational and Applied Mathematics, 101:117-129, 1999.

[3] C.Gotsman and V.Surazhsky, *Guaranteed Intersection-Free Polygon Morphing.* Computers and Graphics, 25(1):67-75, 2001.

[4] V.Surazhsky and C.Gotsman, *Controllable Morphing of Compatible Planar Triangulations.* ACM Transactions on Graphics, 20(4):203-231, 2001.

[5] G.Vegter, *Kink-free deformations of polygons.* Proc. Annual Symposium on Computational Geometry 1989, pp. 61-68.

[6] E.C.de Verdire, M.Pocchiola and G.Vegter, *Tutte's Barycenter Method applied to Isotopies.* Proceedings CCCG 2001, p. 57-60.

[7] H. Whitney, *On regular closed curves in the plane.* Compos. Math. 4, pp. 276–284, 1937.