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

Jaeil Choi and Andrzej Szymczak
College of Computing, Georgia Tech
{jerry,andrzej}@cc.gatech.edu

Abstract

Morphing algorithms aim to construct visually pleasing interpolations (morphs) between 2D or 3D shapes. One of the desirable properties of a morph is avoiding self-intersections of the deforming shape. We investigate topological invariants (based on winding number) of planar morphs between compatible triangulations that do allow global self-intersections but avoid local ones. Equivalently, such morphs do not make any of the triangles degenerate. We discuss a variant of the as rigid as possible morphing algorithm based on these invariants that 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. We limit ourselves to discussion of 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 Laplace operator, 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 weights and construct the morph 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. In contrast to [2, 3, 4, 6], 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 guarantee that the morph does not contain self-intersections.

2 Overview of the main results

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. Such a geometric realization is required to preserve the orientation, i.e. map each oriented triangle into a triangle in the plane oriented in a counterclockwise manner. 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 (such a morph will be called non-degenerating). However, it is impossible to morph the third one to any of the other two. A simple argument can be based on the concept of kink-free deformations of polygonal loops [5]. In the smooth case [7], 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.

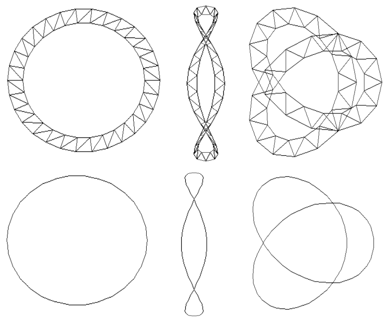

Figure 1: Three triangulations and their associated paths.

As-rigid-as-possible shape interpolation method [1] takes as input two compatible triangulations 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 $GL_2^+(\mathbb{R})$), 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 (even their homotopy classes) of this kind. The paper [1] includes experimental evidence showing 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 $\sigma_T : [0,1] \to GL_2^+(\mathbb{R})$ by

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

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 applied by the path σ_T to the triangle T, or, more precisely, the homotopy class of this path. Note that $GL_2^+(\mathbb{R})$ is homotopy equivalent to the circle. Homotopy classes of loops in that space can be distinguished as follows. For loops $\Sigma : [0,M] \to GL^+(\mathbb{R}^2)$ and $\sigma : [0,M] \to \mathbb{R}^2 \setminus \{0\}$, $\Sigma(t)(\sigma(t))$ is a loop in $\mathbb{R}^2 \setminus \{0\}$. Its winding number depends only on the homotopy types of the loops Σ and σ. By the winding number of Σ we mean the winding number of the loop $\Sigma(t)(\sigma(t))$ for a null-homotopic loop σ. One can prove that, for an arbitrary loop $\sigma : [0,M] \to \mathbb{R}^2 \setminus \{0\}$, $w(\Sigma(t)(\sigma(t))) = w(\Sigma(t)) + w(\sigma(t))$, where by $w(.)$ we denote the winding number of a loop either in $\mathbb{R} \setminus \{0\}$ or $GL_2^+(\mathbb{R})$. Two loops in $GL_2^+(\mathbb{R})$ can be proved to be homotopic if and only if they have the same winding numbers.

Intuitively, it is clear that adjacent triangles should spin in a similar manner so that their deformation paths can be gracefully compromised when global optimization step is performed. This suggests selecting rotation angles in such a way that they differ by less than π for any pair of adjacent triangles (we shall call such an assignment of rotation angles *coherent*), leading to the following greedy 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 use an open interval $(\alpha(T)-\pi, \alpha(T)+\pi)$ in the algorithm because admissible rotations of adjacent triangles T and T' cannot differ by π. This is because 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 is strictly less than π because the matrices of S_T and $S_{T'}$ are symmetric and positive definite.

The success of Algorithm 1 (i.e. whether the assignment of angles it produces is coherent or not) depends on global properties of the source and target triangulations. 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 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. The rotation angles assigned to the two triangles will be off by more than π. This is because the algorithm gradually accumulates the spin that has to be applied to triangles in the source to deform them to the corresponding triangles of the target.

The algorithm does succeed when a non-degenerating morph between the source and the target exists. We

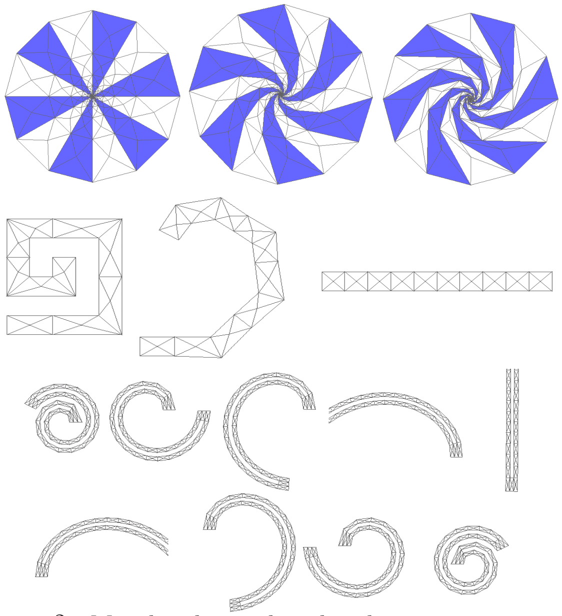

Figure 2: Morphs obtained with coherent rotation angles (some of the snapshots of the spiral have been truncated).

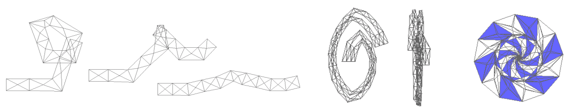

Figure 3: Intermediate stages of the morphs between the same source and target meshes as in Figure 2, but without coherent angle selection.

claim that this case is most important for applications. When finding a non-degenerating morph is not possible, then one should rather look for a different compatible triangulations of the source and target before constructing a high quality morph.

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 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 Formal definitions and proofs

In this section we prove that Algorithm 1 produces coherent rotation angles if a non-degenerating morph between the source and target triangulations exists.

3.1 Invariants of non-degenerating morphs and winding number equivalent triangulations

Let $\mathcal{C}$ be a connected 2D manifold (with boundary) oriented abstract simplicial complex. Consider a loop L in the dual graph of $\mathcal{C}$ that does not backtrack (i.e. no two consecutive edges on L are equal). Recall that the vertices of the dual graph correspond to the triangles of $\mathcal{C}$ and edges join adjacent triangles. With any such loop L we associate a planar loop, depending on a geometric realization F of $\mathcal{C}$, obtained 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 can be used to argue that the following statement holds.

Lemma 1 *The winding number along any loop in the dual graph is preserved by a non-degenerating morph.*

In particular, the winding number 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 $\tau_{F,L}$, where F denotes a geometric realization of the complex $\mathcal{C}$. The following definition specifies the class of source and target triangulations which are not precluded from being morphable in a non-degenerate manner by Lemma 1.

Definition 1 *Two geometric realizations F and G of $\mathcal{C}$ are winding number equivalent if and only if the winding numbers of the loops $\tau_{F,L}$ and $\tau_{G,L}$ are the same for any loop L in the dual graph of $\mathcal{C}$ that does not backtrack.*

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 $\Sigma_{F\to G,L}$ in $GL_2^+(\mathbb{R})$ defined as follows. For each triangle T on L, consider the linear isomorphism M_T defined as in Section 2, i.e. such that $M_T(\overrightarrow{F(a)F(b)}) = \overrightarrow{G(a)G(b)}$ and $M_T(\overrightarrow{F(a)F(c)}) = \overrightarrow{G(a)G(c)}$. The loop $\Sigma_{F\to G,L}$ is obtained by joining the isomorphisms M_T for T along L with linear paths: if the consecutive triangles on the loop L are $T_0, T_1, \ldots, T_n = T_0$ then $\Sigma_{F\to G,L} : [0,n] \to GL_2^+(\mathbb{R})$ is defined by $\Sigma_{F\to G,L}(t) = (k+1-t)M_{T_k} + (t-k)M_{T_{k+1}}$ for $t \in [k, k+1], k = 0, 1, 2, \ldots, n-1$. All transformations on the loop $\Sigma_{F\to G,L}$ can be argued to belong to $GL_2^+(\mathbb{R})$ using the following proposition.

Proposition 1 *For any two adjacent triangles T and T' in $\mathcal{C}$ all convex combinations of M_T and $M_{T'}$ are isomorphisms.*

Proof. Let p and q be the two vertices shared by T and T' and r and s - the remaining vertices of T and T' (respectively). Let $P = F(p)$, $Q = F(q)$, $R = F(r)$, $S = F(s)$, $P' = G(p)$, $Q' = G(q)$, $R' = G(r)$, $S' = G(s)$. Notice that both M_T and $M_{T'}$ map the vector $\overrightarrow{PQ}$ into the same nonzero vector, $\overrightarrow{P'Q'}$. Therefore, so does any their convex combination. Moreover, they both map an arbitrarily chosen nonzero vector $\vec{w}$ perpendicular to

$\overrightarrow{PQ}$ to a vector pointing to the same side of the line through the origin parallel to $\overrightarrow{P'Q'}$ (since they preserve orientation). Therefore, convex combinations of M_T and $M_{T'}$ cannot map $\vec{w}$ to a vector parallel to $\overrightarrow{P'Q'}$ and therefore they are isomorphisms. □

We finish this section with the following result.

Proposition 2 *If F and G are winding number equivalent then the loop $\Sigma_{F\to G,L}$ is contractible in $GL_2^+(\mathbb{R})$ for any loop L in the dual graph of $\mathcal{C}$.*

Proof. Let $T_0, T_1, \ldots, T_n = T_0$ be the consecutive triangles on L. Let P_i and P_i' be the midpoints of the interval joining the points corresponding to vertices shared by the triangles T_i and T_{i+1} in the source and target triangulations for $i = 0, 1, \ldots, n-1$. Define loops $\vec{v}, \vec{v'} : [0, n] \to \mathbb{R}^2 \setminus \{0\}$ (which, in fact, are the loops $\tau_{F,L}$ and $\tau_{G,L}$ defined in Section 3.1) by: $\vec{v}(t) = (k+1-t)\overrightarrow{P_{k-1}P_k} + (t-k)\overrightarrow{P_kP_{k+1}}$ and $\vec{v'}(t) = (k+1-t)\overrightarrow{P_{k-1}'P_k'} + (t-k)\overrightarrow{P_k'P_{k+1}'}$ for $t \in [k, k+1], k = 0, 1, \ldots, n-1$ (where $P_{-1} = P_{n-1}$ and $P_{-1}' = P_{n-1}'$). The loop $\theta(t) = \Sigma_{F\to G,L}(t)(\vec{v}(t))$ is homotopic to $\vec{v'}$. This is because, by an elementary calculation, for $t \in [k, k+1]$ with $k \in \{0, 1, \ldots, n-1\}$, $\theta(t)$ is a convex combination of vectors $M_{T_k}(\overrightarrow{P_kP_{k+1}})$, $M_{T_{k+1}}(\overrightarrow{P_{k-1}P_k})$, $M_{T_k}(\overrightarrow{P_{k-1}P_k})$ and $M_{T_{k+1}}(\overrightarrow{P_kP_{k+1}})$. All of those vectors belong to the same open half-plane bounded by the line passing through the origin and parallel to the edge corresponding to that shared by T_k and T_{k+1} in G. Since $\vec{v'}$ is a convex combination of vectors $\overrightarrow{P_k'P_{k+1}'}$ and $\overrightarrow{P_{k-1}'P_k'}$ which both belong to the same half-plane, $\angle(\theta(t), \vec{v'}(t)) < \pi$, and therefore the loops θ and $\vec{v'}$ are indeed homotopic in $\mathbb{R}^2 \setminus \{0\}$.

Since the winding numbers of the loops $\vec{v}$ and $\vec{v'}$ are the same, this means that the winding number of $\Sigma_{F\to G,L}$ is zero and therefore it is contractible. □

3.2 Algorithm 1 for winding number equivalent triangulations

In this section, we shall prove the following theorem.

Theorem 1 *If the source and target triangulations are winding number equivalent then Algorithm 1 produces a coherent assignment of rotation angles.*

In what follows, by $\star$ we shall denote the operation of concatenating paths. For two paths σ and τ in a topological space X such that the starting point of τ coincides with the endpoint of σ, $\sigma \star \tau$ is the path that first follows the path σ and then follows τ. For two paths η and θ that have the same starting points and the same endpoints we shall write $\eta \equiv \theta$ if they are homotopic with the starting points and endpoints fixed.

We start with a proof of the following proposition.

Proposition 3 *Let T and T' be two adjacent triangles in $\mathcal{C}$ and $L(T,T')$ be the linear path joining M_T and $M_{T'}$. Let σ_T be defined as in Equation (1). Then*

$$\sigma_T \star L(T,T') \equiv \sigma_{T'} \qquad (2)$$

if and only if the rotation angles $\alpha(T)$ and $\alpha(T')$ are chosen in such a way that $|\alpha(T) - \alpha(T')| < \pi$.

Proof. We concentrate on the 'if' part. The 'only if' part holds since validity of Equation (2) depends only on the difference of the two rotation angles.

Let p and q be the vertices of the edge shared by T and T', P and Q their corresponding vertices in the source triangulation and P' and Q' - their corresponding vertices in the target triangulation.

We are going to evaluate the paths on both sides of Equation 2 on the vector $\overrightarrow{PQ}$, proving that this leads to homotopic paths in $\mathbb{R}^2 \setminus \{0\}$. Notice that $\angle(S_T(\overrightarrow{PQ}), \overrightarrow{P'Q'}) = \alpha(T)$, $\angle(S_{T'}(\overrightarrow{PQ}), \overrightarrow{P'Q'}) = \alpha(T')$ (modulo 2π; both here and below, the angles are signed, i.e. they denote the angle needed to rotate the first vector into one pointing in the direction of the second vector) and, since $M_T(\overrightarrow{PQ}) = M_{T'}(\overrightarrow{PQ}) = \overrightarrow{P'Q'}$, $L(T,T')(t)(\overrightarrow{PQ}) = \overrightarrow{P'Q'}$. Since S_T and $S_{T'}$ are symmetric positive definite, $\angle(\overrightarrow{PQ}, S_T(\overrightarrow{PQ})) < \pi/2$, and $\angle(\overrightarrow{PQ}, S_{T'}(\overrightarrow{PQ})) < \pi/2$. Because $\alpha(T)$ and $\alpha(T')$ differ by less than π, $\alpha(T) + \angle(\overrightarrow{PQ}, S_T(\overrightarrow{PQ})) = \alpha(T') + \angle(\overrightarrow{PQ}, S_{T'}(\overrightarrow{PQ}))$. Thus both paths spin the vector $\overrightarrow{PQ}$ by the same amount. Hence they are homotopy equivalent. □

Proof of Theorem 1. Recall that Algorithm 1 first assigns a rotation angle to a triangle T_0 and then recursively to neighbors of triangles which have an angle assigned. This means that there are paths τ and τ' in the dual graph joining T_0 with T and T' (respectively) along which the angles are assigned (meaning that the angles for any two consecutive triangles on any of the two paths differ by less than π). By first following the path σ backward and then the path σ' we obtain a path $\bar{\tau}$ in the dual graph starting at T and ending at T'. Build a path $\bar{\tau}_*$ in $GL_2^+(\mathbb{R})$ by joining transformations M_T on the path $\bar{\tau}$ with linear paths. By Proposition 2, $\bar{\tau}_*$ is homotopic to the linear path $L(T,T')$. A simple inductive proof based on Proposition 3 shows that $\sigma_T \star \bar{\tau}_* \equiv \sigma_{T'}$, where the rotation angles for σ_T and $\sigma_{T'}$ are produced by the Algorithm. We conclude that $\sigma_T \star L(T,T') \equiv \sigma_{T'}$. Using Proposition 3 again, we have $|\alpha(T) - \alpha(T')| < \pi$. □

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 Verdiére, 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.