

Actual Morphing: A Physics-Based Approach to Blending

Shi-Min Hu^a, Chen-Feng Li^b, Hui Zhang^a

^aTsinghua University, Beijing, People's Republic of China

^bUniversity of Wales Swansea, UK

Abstract

When two topologically identical shapes are blended, various possible transformation paths exist from the source shape to the target shape. Which one is the most plausible? Here we propose that the transformation process should obey a quasi-physical law. This paper combines morphing with deformation theory from continuum mechanics. By using strain energy, which reflects the magnitude of deformation, as an objective function, we convert the problem of path interpolation into an unconstrained optimization problem. To reduce the number of variables in the optimization we adopt shape functions, as used in the finite element method (FEM). A point-to-point correspondence between the source and target shapes is naturally established using these polynomial functions plus a distance map.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Animation

1. Introduction

Morphing, the process of transforming one shape into another via intermediate shapes, plays an important role in computer animation, industrial design, and simulation. Extensive investigations have been conducted on this subject and existing methods can be classified into four categories according to how the objects are represented: image morphing [BN92, LCSW95], 2D shape blending [SG92, SGWM93], volume-based morphing [Hug92, LGL95, TO99, TPG01] and boundary-based morphing [KCP92].

There are two main aspects to a morphing problem: vertex correspondence and path interpolation. For image morphing, the main problem is to find a pixel-to-pixel correspondence. For 2D shape blending, there are algorithms that automatically determine vertex correspondence between the polygons [Shi99]. The principal task is to identify the right path. Sederberg and Greenwood [SG92] presented an algorithm whereby the first shape is bent and/or stretched into the second with the minimum amount of work. Sederberg et al. [SGWM93] subsequently proposed a solution that works by interpolating the intrinsic parameters—angles and edge lengths—of the initial and final shapes. In 3D cases, as in volume-based morphing and boundary-based morphing, both vertex correspondence and path interpolation are non-trivial. The main techniques used to set up 3D vertex correspondence are harmonic mapping [KSK00], MAPS [LDSS99] and consistent mesh parameterization [PSS01]. The path problem has also been comprehensively studied: Cohen et al. [CLS98] presented a warp transformation method of decomposing the transformation into a rigid component and an elastic component, based on distance field interpolation. The interpolation technique of Alexa et al. [ACL00] aims to make body elements, i.e. triangles for 2D or tetrahedrons for 3D, as rigid as possible dur-

ing shape interpolation. Breen et al. [BW01] proposed a level-set approach for 3D morphing which has the advantage of being able to cope with shapes of different topological genus.

Although some efficient techniques have been proposed for both vertex correspondence and path interpolation, most concentrate on only one of these two aspects. For example, Lee et al. [LDSS99] used the MAPS algorithm to parameterize source and target meshes over simple base domains and an additional harmonic map, and to bring the latter into correspondence, while only using simple interpolation for path determination. Alexa et al. [ACL00] considered only the path problem, and assumed the availability of a boundary vertex correspondence between source and target shapes. It is generally accepted, however, that setting up boundary vertex correspondence for 3D objects with different numbers of vertices is non-trivial.

In this paper, we present a novel method which uses a quasi-physical law to determine the ideal transformation path when blending two topologically identical shapes. The path is then found by solving an unconstrained optimization problem, and a point-to-point correspondence between the source and target shapes is set up by using the same discrete formulation used in tracing the path, i.e. shape functions, plus a distance map.

The paper is set out as follows. Section 2 outlines the underlying mathematical theory of the algorithm from two aspects, i.e. quasi-physical law with path optimization, and numerical solution with vertex correspondence. In Section 3 we explain the algorithm in detail for the 3D case. Section 4 gives some examples in both 2D and 3D. Finally we give some conclusions in Section 5.

2. Mathematical basis

2.1. Quasi-physical law and path optimization

As shown in Fig. 1, given a body B 's source configuration B^0 and target configuration B^M , the objective of morphing is to find intermediate configurations B^i , where $i = 1, 2, \dots, M-1$, which transform the object from B^0 to B^M smoothly.

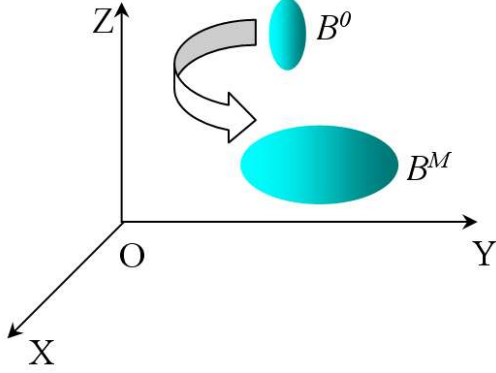


Figure 1: Schematic illustration of morphing.

Using visual inspection, a user can judge the quality of the morphing sequence. But how can quality be *measured*? A practical criterion is that the object should undergo the least distortion and travel along the simplest path. Using ideas from deformation theory from continuum mechanics [Fun65], material properties should be assigned to a geometric object, and its strain energy should be used to measure its distortion during a morphing process. Using these ideas, we can now suggest that the ideal path should be the one with the minimum strain energy. We call this path interpolation criterion a *quasi-physical law*.

Going from B^0 to B^M , each new configuration serves as an equilibrium position from which the next deformation is computed, and the strain energy is calculated for a succession of deformations taking place at each animation step. Let $\mathbf{X} = [x, y, z]^T$ be a vector in \mathfrak{R}^3 that denotes the location of a point. Thus,

$$\mathbf{X}^i = \mathbf{X}^{i-1} + \mathbf{u}^i, i = 1, 2, \dots, M \quad (1)$$

defines the path of the point during morphing, $\mathbf{X}^0 \xrightarrow{\mathbf{u}^1} \mathbf{X}^1 \xrightarrow{\mathbf{u}^2} \mathbf{X}^2, \dots, \mathbf{X}^{M-2} \xrightarrow{\mathbf{u}^{M-1}} \mathbf{X}^{M-1} \xrightarrow{\mathbf{u}^M} \mathbf{X}^M$, where \mathbf{u}^i denotes the relative displacement of the point between configurations B^i and B^{i-1} . Here, \mathbf{X}^0 and \mathbf{X}^M are pre-defined by source and target configurations B^0 and B^M of the body respectively. From these displacements, the strain energy U which reflects the magnitude of deformation in a morphing process can be obtained by

$$\boldsymbol{\varepsilon}^i = \frac{1}{2}(\mathbf{u}^i \nabla + \nabla \mathbf{u}^i) \quad (2)$$

$$\boldsymbol{\sigma}^i = C : \boldsymbol{\varepsilon}^i \quad (3)$$

$$U^i = \int_{\Omega^i} \frac{1}{2} \boldsymbol{\sigma}^i : \boldsymbol{\varepsilon}^i d\Omega \quad (4)$$

$$U = \sum_{i=1}^M U^i(\mathbf{X}^i) \quad (5)$$

where the strain tensor $\boldsymbol{\varepsilon}^i$ describes the deformation of the object, the stress tensor $\boldsymbol{\sigma}^i$ denotes the internal “force”, C is the elastic stiffness tensor—a rank four constant tensor with only two independent parameters for isotropic elastic material, $:$ represents tensor contraction, and ∇ is the gradient operator. (For concepts of tensor analysis refer to [Flü72]).

Using our quasi-physical law, intermediate configurations B^i ($i = 1, 2, \dots, M-1$) are determined by solving an unconstrained optimization problem:

$$\min_{\mathbf{X}^i} (U) = \min_{\mathbf{X}^i} \sum_{i=1}^M U^i(\mathbf{X}^i) \quad (6)$$

2.2. Numerical solution and vertex correspondence

To solve the optimization problem in Eq. 6 efficiently, a finite element mesh is generated for the object and the coordinates of points on the i^{th} configuration B^i , $i = 0, 1, \dots, M$ are approximated by

$$\mathbf{X}^i = \mathbf{N}\mathbf{X}_{(e)}^i = [\mathbf{N}_{j_1}, \mathbf{N}_{j_2}, \dots] \begin{Bmatrix} \mathbf{X}_{j_1} \\ \mathbf{X}_{j_2} \\ \vdots \end{Bmatrix}_{(e)} \quad (7)$$

where $\mathbf{X}_{(e)}$ denotes the node coordinates of the mesh and $\mathbf{N}_{j_i}(\xi_i)$ is the shape function at node j_i . A shape function $\mathbf{N}_{j_i}(\xi_i)$ [ZT00] can take various polynomial forms; the formulation used here produces exact coordinates on nodes. Substitution of Eq. 7 into Eqs. 1–6 yields the minimization problem

$$\min_{\mathbf{X}^i} \sum_{i=1}^M \sum_{e \in \Omega} \mathbf{u}_{(e)}^i{}^T \mathbf{K}_{(e)} \mathbf{u}_{(e)}^i \quad (8)$$

where the stiffness matrix $\mathbf{K}_{(e)}$ is calculated by Gaussian numerical integration. Given the predefined finite element meshes on configurations B^0 and B^M , Eq. 8 is the final numerical form of path optimization, from which a succession of intermediate meshes representing configuration B^i , $i = 1, 2, \dots, M-1$ can be determined. Since $\mathbf{u}_{(e)}^i{}^T \mathbf{K}_{(e)} \mathbf{u}_{(e)}^i$ is a positive definite quadratic form with good properties (it is unimodal and adequately smooth), the above optimization can be solved efficiently with any classical optimization algorithm such as the gradient method or the quasi-Newton method.

Shape functions within a finite element mesh grid define a locally smooth patch with regular parametric range, i.e., $[-1, 1]$ for each coordinate variable, and combination of these patches produces a geometrically continuous interpolated surface for the approximate configuration B^i , $i = 0, 1, \dots, M$. Since the topology of the mesh stays unchanged during a morphing procedure, the mesh surface provides a grid-to-grid correspondence. In order to compute the accurate configuration B^i , $i = 0, 1, \dots, M$ and the point-to-point correspondence, the interpolated surface is taken as one component of the original configuration B^0 or final configuration B^M . The other component, i.e. the elevation map (the distance between the interpolated surface and the original shape surface), is calculated by re-sampling configuration B^0 or B^M at the appropriate resolution. The elevation map for configuration B^i , $i = 1, 2, \dots, M-1$ is calculated

by linear interpolation between the elevation maps for B^0 and B^M . Then, the configuration $B^i, i = 0, 1, \dots, M$ can be represented as a superimposition of the interpolated surface and the corresponding elevation map. An example is shown in Fig. 2, where four nodes (yellow) and corresponding shape functions define the parametric patch (yellow), with a local coordinate system and regular parametric range $[-1, 1]$ for each coordinate valuable. Re-sampling is done on the original shape surface (blue), with a vector from the parametric anchor point (red dot) which is perpendicular to the parametric patch (red dashed segment).




Figure 2: Schematic illustration of re-sampling.

Finite element meshes over the source and target objects aid the task of shape correspondence from grid to grid; incorporating shape functions and elevation maps, with the same resolution, into each mesh grid, enables the creation of a point-to-point correspondence between the source and target shapes.

3. Algorithm

In order to illustrate the morphing technique in detail, 3D morphing is used as an example. From the known configurations B^0 and B^M , configurations $B^i, i = 1, 2, \dots, M - 1$ are calculated with our algorithm using the following steps.

3.1. Generating FE meshes and calculating elevation maps

To reduce the number of unknown optimization variables, a surface finite element model is used rather than a volumetric model. First, a succession of point pairs is selected by inspection of the source and target shapes. Each of these point pairs is given a unique serial number and defines a pair of corresponding characteristic points on the source and target shapes. Now, by connecting these points in accordance with their serial numbers, two finite element meshes with the same topology are obtained, one for the source shape and the other for the target. Sometimes, these two meshes may not be fine enough to capture the source and target shapes in detail. It is trivial to refine them, and this is done by adding nodes and edges to the initial mesh grids. Elevation maps for configuration B^0 and B^M are calculated according to the method described in Section 2.2. Using these finite element meshes, a grid-to-grid correspondence is defined on the object; with the elevation maps re-sampled at identical resolutions, a point-to-point correspondence is further set up.

3.2. Material properties

Two independent material parameters for isotropic elastic material are used in this paper: Young's modulus E and Poisson's ratio ν [Fun65]. Since these material properties are at the user's disposal, it is not necessary to choose these two parameters appropriate to real materials. Actually, one can freely select values for Young's modulus E and Poisson's ratio ν within their individual definition domains, i.e. $(0, +\infty)$ for E and $[1/2, 1]$ for ν . Furthermore, our experimental tests show that the morphing results obtained are not particularly sensitive to choice of these two parameters.

3.3. Tracing the morphing path

Initial values for iterative solution of the optimization problem in Eq. 8 i.e., a family of coarse intermediate nodal coordinates, are obtained from linear interpolation between the nodal coordinates for configurations B^0 and B^M . A family of optimal nodal coordinates is then obtained using the gradient method. We use a gradual refinement procedure, where the number of animation frames is doubled at each step. For example, starting at $n = 1$, we (i) generate $2n + 1$ animation frames by interpolation and optimization, then (2) take these optimized animation frames as input values and update n by replacing it by $2n$. The procedure is then repeated as necessary.

3.4. Rebuilding the shape at each animation step

We compute the local parametric patches defined in Eq. 7 using the nodal coordinates calculated above together with the elevation maps found for configurations $B^i, i = 1, 2, \dots, M - 1$ by means of linear interpolation between configurations B^0 and B^M . Thus, the i^{th} intermediate shape is rebuilt by superimposing the parametric patches and the corresponding elevation maps.

4. Examples

4.1. Morphing a colored bar from straight to arched

A simple morphing sequence for a colored bar going from a straight shape to an arched shape is shown in Figs. 3.1–3.3. The shape functions used in this example are listed as follows:


$$N_i = \frac{1}{4}(1 + \xi_i \xi)(1 + \eta_i \eta) \quad i = 1, \dots, 4 \quad (9)$$

where $(\xi_i, \eta_i), i = 1, \dots, 4$ take values $(-1, -1), (1, -1), (1, 1), (-1, 1)$ respectively.

Figure 3.1 shows the finite element meshes of the source and target shapes. Figures 3.2 and 3.3 show morphing results produced using our proposed method and linear interpolation, respectively. Comparison of Figs 3.2 and 3.3 shows that the length and the width of the bar remain unchanged during morphing, unlike the result produced by linear interpolation, which causes drastic shape distortion.

4.2. Morphing of gymnastic patterns

Figures 4.1 and 4.2 show a morphing example involving three 2D gymnastic patterns in which the shape functions in Eq. 9 are again applied. Figure 4.1 shows the finite element meshes of the three gymnastic patterns, and Fig. 4.2 shows the morphing sequence produced; the three key patterns are circled. The arms and legs of the



(a) Straight source bar

(b) Arched target bar

3.1: FE meshes for the two bars




3.2: Animation frames by the proposed method



3.3: Animation frames by linear interpolation.

Figure 3: Morphing a colored bar

gymnastic patterns undergo large and different translations and rotations, which cannot be handled using traditional methods for determining the transformation path. Note that purely rigid rotation consumes zero strain energy; by minimizing the strain energy a satisfactory morphing result that has the least distortion is achieved.

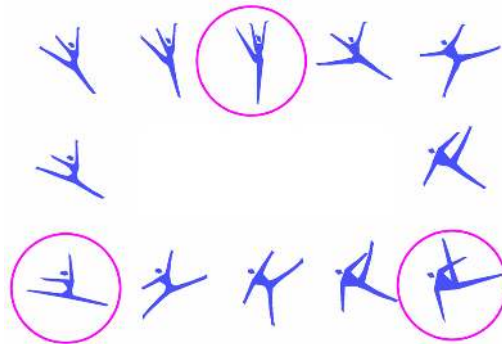


(a) Key pattern 1

(b) Key pattern 2

(c) Key pattern 3

4.1: FE meshes of gymnastic patterns



4.2: Animation frames by the proposed method

Figure 4: Morphing gymnastic patterns

4.3. Morphing between two 3D human heads


A classical morphing test case used in previous research [LDSS99] involves two 3D human heads from a man and a woman. We use it here to test the new morphing algorithm. The man's head has 13,408 triangle facets and the woman's 100,000; the number of vertices is 6,738 for the man's head and 50,002 for the woman's head. In this case we let

$$\bar{N}_i = \frac{1}{8} (1 + \xi \xi_i) (1 + \eta \eta_i) (1 + \zeta \zeta_i) \quad i = 1, 2, \dots, 8 \quad (10)$$

where (ξ_i, η_i, ζ_i) equal $(1, 1, 1)$, $(-1, 1, 1)$, $(-1, -1, 1)$, $(1, -1, 1)$, $(1, 1, -1)$, $(-1, 1, -1)$, $(-1, -1, -1)$, and $(1, -1, -1)$ in turn, with i ranging from 1 to 8. The shape functions [Par95] can be written as:

$$\begin{cases} N_i = \bar{N}_i - \bar{N}_{i+4} \\ N_{i+4} = \bar{N}_i + \bar{N}_{i+4} \end{cases} \quad i = 1, 2, 3, 4. \quad (11)$$

Fig. 5.1 shows the finite element meshes of the two human heads, while Fig. 5.2 a few typical animation frames. This example illustrates that a complex 3D model can be accurately represented by shape functions and elevation maps. The proposed morphing technique demonstrates good performance in such 3D cases as well as in the 2D cases.



(a) Male source head

(b) Female target head

5.1: FE meshes of human heads



5.2: Animation frames of human heads

Figure 5: Morphing heads

4.4. Morphing from a disk to a 3D vase

This example shows morphing from a circular disk to a vase using the shape functions in Eq. 11. Fig. 6.1 illustrates the finite element meshes of the source disk and the target vase, and Fig. 6.2 shows the morphing results. The results show that smooth shapes can be well captured by only a few mesh grids.

4.5. Performance of the algorithm

All the computations were done using a Pentium 4 1.7GHz PC. The total animation frames, the total mesh grids and the CPU time for the optimization are listed in Table 1.




Figure 6: *Morphing a disk into a vase*

Example	Frames	Grids	Time
Colored bar	65	14	1 minute
Gymnastic pattern	65	36	8 minutes
Human head	65	51	15 minutes
Disk to vase	65	52	15 minutes

Table 1: *Computation times*

5. Conclusions

This work has described a simple morphing method in which, an object is described by a superimposition of a shape-function-interpolated surface with an elevation map. A point-to-point correspondence between the source shape and the target shape is built up by re-sampling the models at the same resolution, and both the motion and deformation path are found by minimizing strain energy accumulated during the morphing process (using a so-called quasi-physical law). Several animations were presented, showing that the algorithm generates good quality morphing results closely maintain much of the original shape and volume of the two models.

The approach is philosophically similar to that described in [TM91] but differs in the details of the model representation and purpose of the analysis. In particular, morphing can be seen as a boundary value problem without any predefined forces, whereas Terzopoulos [TM91] considers shape registration as an initial value problem in which external forces must be defined in advance. While Terzopoulos' outstanding work has achieved great success in medical image registration, it is well known that deriving external forces from data constraints is non-trivial. This is overcome here by only including coordinate variables in the algorithm.

Some areas of future research we intend to investigate are as follows:

- To reduce the number of unknown optimization variables and time cost involved in generating finite element meshes, a surface elastic model has been used for blending 3D shapes rather than a volumetric one. A volumetric model is better suited to minimiz-

ing volume variations. More work needs to be done to see how these conflicting issues can best be resolved.

- There are presently few research results on morphing control. Since nodal coordinates at each animation step have been taken as the independent variables for optimization, constraints can easily be incorporated into the morphing, especially location constraints. This will certainly be useful in manipulating animation.
- The finite element meshes used in the algorithm are defined manually, which is cumbersome. More work is needed on automatically generating simple meshes based on the characteristics of the source and the target objects.

Acknowledgements

We thank Mr. Wei Hong and Miss. Bin Guo for their helpful assistance in coding, as well as Prof. Chiew-Lan Tai for her kind help. This work was supported by the Natural Science Foundation of China (Project Number 60225016), the Specialized Research Fund for the Doctoral Program of Higher Education (Project Number 20020003051) and the National Basic Research Project of China (Project Number 2002CB312100).

References

- [ACL00] ALEXA M., COHEN-OR D., LEVIN D.: As-rigid-as-possible shape interpolation. In *Proc. SIGGRAPH '00* (2000), pp. 157–164. 1
- [BN92] BEIER T., NEELY S.: Feature-based image metamorphosis. In *Proc. SIGGRAPH '92* (1992), pp. 35–42. 1
- [BW01] BREEN D. E., WHITAKER R. T.: A level-set approach for the metamorphosis of solid models. *IEEE Transactions on Visualization and Computer Graphics* 7, 2 (2001), 173–192. 1
- [CLS98] COHEN-OR D., LEVIN D., SOLOMOVICI A.: Three-dimensional distance field metamorphosis. *ACM TOG* 17, 2 (1998), 116–141. 1
- [Flü72] FLÜGGE W.: *Tensor analysis and continuum mechanics*. Springer-Verlag, Berlin, 1972. 2
- [Fun65] FUNG Y. C.: *Foundation of solid mechanics*. Prentice-Hall, Englewood Cliffs, 1965. 2, 3
- [Hug92] HUGHES J. F.: Scheduled fourier volume morphing. In *Proc. SIGGRAPH '92* (1992), pp. 35–42. 1
- [KCP92] KENT J. R., CARLSON W. E., PARENT R. E.: Shape transformation for polyhedral objects. In *Proc. SIGGRAPH '92* (1992), pp. 47–54. 1
- [KSK00] KANAI T., SUZUKI H., KIMURA F.: Metamorphosis of arbitrary triangular meshes. *IEEE Computer Graphics and Applications* 20, 2 (2000), 62–75. 1
- [LCSW95] LEE S. Y., CHWA K. Y., SHIN S. Y., WOLBERG G.: Image metamorphosis using snakes and free-form deformations. In *Proc. SIGGRAPH '95* (1995), pp. 439–448. 1

- [LDSS99] LEE A. W. F., DOBKIN D., SWELDENS W., SCHRODER P.: Multiresolution mesh morphing. In *Proc. SIGGRAPH '99* (1999), pp. 343–350. [1](#), [4](#)
- [LGL95] LERIOS A., GARFINKLE C. D., LEVOY M.: Feature-based volume metamorphosis. In *Proc. SIGGRAPH '95* (1995), pp. 449–456. [1](#)
- [Par95] PARISCH H.: A continuum-based shell theory for non-linear applications. *International Journal for Numerical Methods in Engineering* 38 (1995), 855–1883. [4](#)
- [PSS01] PRAUN E., SWELDENS W., SCHRÖDER P.: Consistent mesh parameterization. In *Proc. SIGGRAPH '01* (2001), pp. 179–184. [1](#)
- [SG92] SEDERBERG T. W., GREENWOOD E.: A physically based approach to 2-D shape blending. In *Proc. SIGGRAPH '92* (1992), pp. 25–34. [1](#)
- [SGWM93] SEDERBERG T. W., GAO P., WANG G., MU H.: 2D shape blending: an intrinsic solution to the vertex path problem. In *Proc. SIGGRAPH '93* (1993), pp. 15–18. [1](#)
- [Shi99] SHI J. Y.: *Visualization in scientific computing*. Chinese Science Press, Beijing, 1999. [1](#)
- [TM91] TERZOPOULOS D., METAXAS D.: Dynamic 3D models with local and global deformations: deformable superquadrics. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 13, 7 (1991), 703–714. [5](#)
- [TO99] TURK, O'BRIEN J. F.: Shape transformation using variational implicit functions. In *Proc. SIGGRAPH '99* (1999), pp. 335–342. [1](#)
- [TPG01] TREECE G. M., PRAGER R. W., GEE A. H.: Volume-based three-dimensional metamorphosis using automatic region correspondence. *The Visual Computer* 17, 7 (2001), 397–414. [1](#)
- [ZT00] ZIENKIEWICZ O. C., TAYLOR R. L.: *The Finite Element Method*, 5 ed., vol. 1-2. Oxford, Butterworth-Heinemann, 2000. [2](#)