The Visual Computer

Image morphing: a survey

George Wolberg

Department of Computer Science,
City College of New York, New York, NY 10031, USA
E-mail: wolberg@cs-mail.engr.ccny.cuny.edu

Image morphing has received much attention in recent years. It has proven to be a powerful tool for visual effects in film and television, enabling the fluid transformation of one digital image into another. This paper surveys the growth of this field and describes recent advances in image morphing in terms of feature specification, warp generation methods, and transition control. These areas relate to the ease of use and quality of results. We describe the role of radial basis functions, thin plate splines, energy minimization, and multilevel free-form deformations in advancing the state-of-the-art in image morphing. Recent work on a generalized framework for morphing among multiple images is described.

Key words: Image morphing – Feature specification – Warp generation – Transition control

1 Introduction

Image metamorphosis has proven to be a powerful tool for visual effects. There are now many breathtaking examples in film and television of the fluid transformation of one digital image into another. This process, commonly known as *morphing*, is realized by coupling image warping with color interpolation. Image warping applies 2D geometric transformations to the images to retain geometric alignment between their features, while color interpolation blends their colors.

Image metamorphosis between two images begins with an animator establishing their correspondence with pairs of feature primitives, e.g., mesh nodes, line segments, curves, or points. Each primitive specifies an image feature or landmark. The feature correspondence is then used to compute mapping functions that define the spatial relationship between all points in both images. Since mapping functions are central to warping, we shall refer to them as *warp functions* in this paper. They will be used to interpolate the positions of the features across the morph sequence. Once both images have been warped into alignment for intermediate feature positions, ordinary color interpolation (i.e., cross-dissolve) generates inbetween images.

Feature specification is the most tedious aspect of morphing. Although the choice of allowable primitives may vary, all morphing approaches require careful attention to the precise placement of primitives. Given feature correspondence constraints between both images, a warp function over the whole image plane must be derived. This process, which we refer to as *warp generation*, is essentially an interpolation problem. Another interesting problem in image morphing is transition control. If transition rates are allowed to vary locally across inbetween images, more interesting animations are possible.

The explosive growth of image morphing is due to the compelling and aesthetically pleasing effects made possible by warping and color blending. The extent to which artists and animators can effectively use morphing tools is directly tied to solutions to the following three problems: feature specification, warp generation, and transition control. Together, they influence the ease and effectiveness in generating high-quality metamorphosis sequences. This paper surveys recent advances in image morphing in terms of their roles in addressing these three problems. We compare various morphing techniques, including those based on

The Visual Computer (1998) 14:360–372

mesh warping (Wolberg 1990), field morphing (Beier and Neely 1992), radial basis functions (Arad et al. 1994), thin plate splines (Lee et al. 1994; Litwinowicz and Williams 1994), energy minimization (Lee et al. 1996), and multilevel free-form deformations (Lee et al.1995).
A tradeoff exists between the complexity of feature specification and warp generation. As feature specification becomes more convenient, warp generation becomes more formidable. The recent introduction of spline curves to feature specification raises a challenge to the warp generation process, making it the most critical component of morphing. It influences the smoothness of the transformation and dominates the computational cost of the morphing process. We comment on these tradeoffs and describe their role in influencing recent progress in this field.
The scope of this paper is limited to traditional 2D image morphing. We refrain from discussing related work in 3D volume morphing (Lerios et al. 1995), view interpolation (Chen and Williams 1993), or view morphing (Seitz and Dyer 1996). With the exception of 3D volume morphing, which is a direct extension of 2D image morphing, the latter two areas apply morphing to compute new views of a scene, given a set of prestored images.

2 Morphing algorithms

Before the development of morphing, image transitions were generally achieved through the use of cross-dissolves, e.g., linear interpolation to fade from one image to another. Figure 1 depicts this process applied over five frames. The result is poor, owing to the double exposure effect apparent in misaligned regions. This problem is particularly apparent in the middle frame, where both input images contribute equally to the output. Morphing achieves a fluid transformation by incorporating warping to maintain geometric alignment throughout the cross-dissolve process.
In this section, we review several morphing algorithms, including those based on mesh warping, field morphing, radial basis functions, thin plate splines, energy minimization, and multilevel free-form deformations. This review is intended to motivate the discussion of progress in feature specification, warp generation, and transition control.

2.1 Mesh warping

Mesh warping was pioneered at Industrial Light & Magic (ILM) by D. Smythe for use in the movie *Willow* in 1988 (Smythe 1990). It has been successfully used in many subsequent motion pictures. To illustrate the two-pass mesh warping algorithm, consider the image sequence shown in Fig. 2. The five frames in the middle row represent a metamorphosis (or morph) between the two faces at both ends of the row. We refer to these two images as I_S and I_T, the source and the target images, respectively. The source image is associated with mesh M_S. It specifies the coordinates of control points, or landmarks. A second mesh M_T specifies their corresponding positions in the target image. Meshes M_S and M_T are respectively shown overlaid on I_S and I_T in the upper left and lower right images of the figure. Notice that landmarks such as the eyes, nose, and lips lie below the corresponding grid lines in both meshes. Together, M_S and M_T are used to define the spatial transformation that maps all points in I_S onto I_T. The meshes are constrained to be topologically equivalent, i.e., no folding or discontinuities are permitted. Therefore, the nodes in M_T may wander as far from M_S as necessary, as long as they do not cause self-intersection. Furthermore, for simplicity, the meshes are constrained to have frozen borders.
All intermediate frames in the morph sequence are the product of a four-step process:

for each frame f **do**
  Linearly interpolate mesh M, between M_S and M_T
  Warp I_S to I_1, using meshes M_S and M
  Warp I_T to I_2, using meshes M_T and M
  Linearly interpolate image I_f, between I_1 and I_2
end

Figure 2 depicts this process. In the top row of the figure, mesh M_S is shown deforming to mesh M_T, producing an intermediate mesh M for each frame f. These meshes are used to warp I_S into increasingly deformed images, thereby deforming I_S from its original state to those defined by the intermediate meshes. The identical process is shown in reverse order in the bottom row of the figure, where I_T is shown deforming from its original state. The purpose of this procedure is to maintain the align-

1

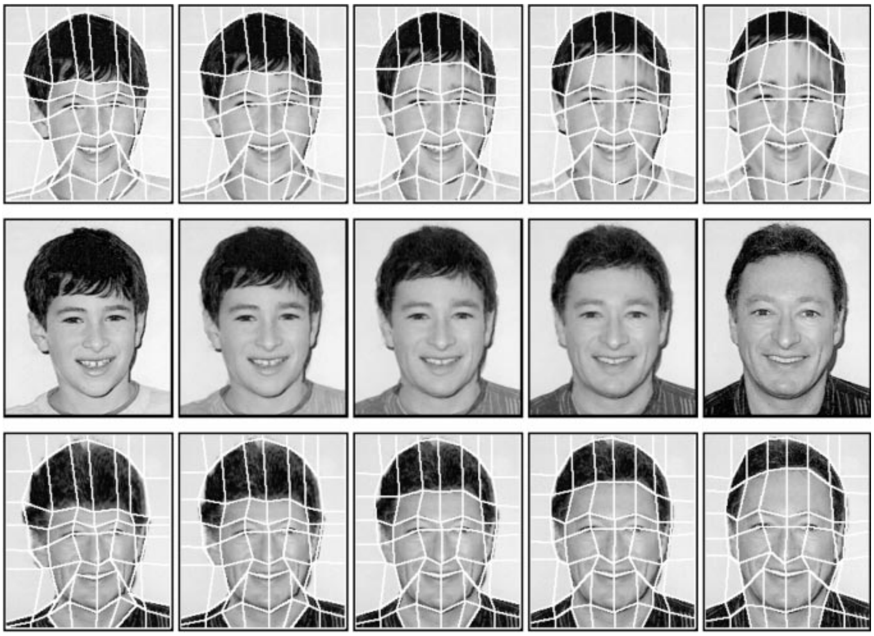

2

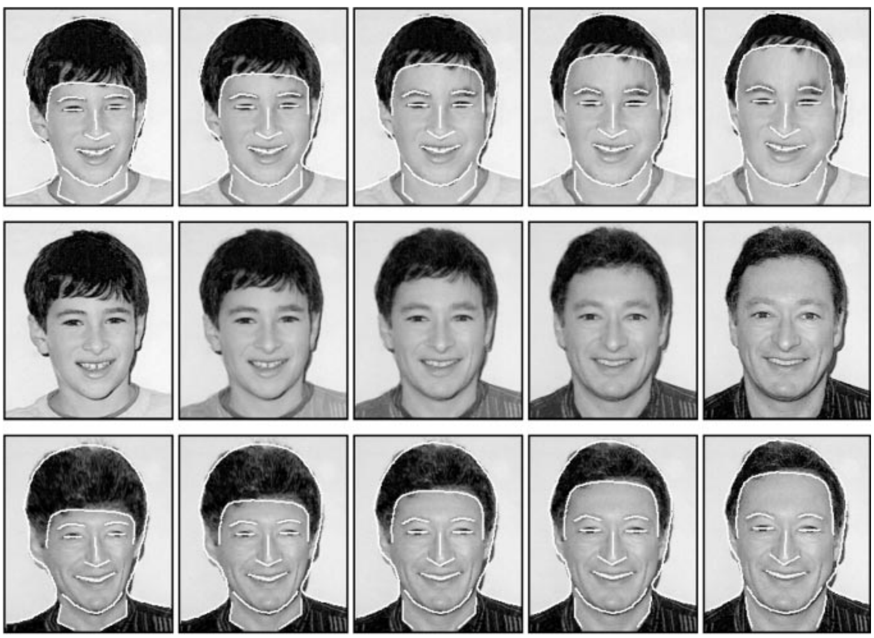

3

Fig. 1. Cross-dissolve

Fig. 2. Mesh warping

Fig. 3. Multilevel free-form deformation (MFFD) based morphing

ment of landmarks between I_S and I_T as they both deform to some intermediate state, producing the pairs of I_1 and I_2 images shown in the top and bottom rows, respectively. Only after this alignment is maintained does a cross-dissolve between successive pairs of I_1 and I_2 become meaningful, as shown in the morph sequence in the middle row. This sequence was produced by applying the weights [1, 0.75, 0.5, 0.25, 0] and [0, 0.25, 0.5, 0.75, 1] to the five images in the top and bottom rows, respectively, and adding the two sets together. This process demonstrates that morphing is simply a cross-dissolve applied to warped imagery. The important role that warping plays here is readily apparent in the comparison of the morph sequence in Fig. 2 with the cross-dissolve result in Fig. 1.

The use of meshes for feature specification facilitates a straightforward solution for warp generation: bicubic spline interpolation. The example here employed Catmull-Rom spline interpolation to determine the correspondence of all pixels. Fant's algorithm was used to resample the image in a separable implementation (Fant 1986; Wolberg 1990).

2.2 Field morphing

While meshes appear to be a convenient manner of specifying pairs of feature points, they are sometimes cumbersome to use. The field morphing algorithm developed by Beier and Neely (1992) at Pacific Data Images grew out of the desire to simplify the user interface to handle correspondence by means of line pairs. A pair of corresponding lines in the source and target images defines a coordinate mapping between the two images. In addition to the straightforward correspondence provided for all points along the lines, the mapping of points in the vicinity of the line can be determined by their distance from the line. Since multiple line pairs are usually given, the displacement of a point in the source image is actually a weighted sum of the mappings due to each line pair, with the weights attributed to distance and line length.

This approach has the benefit of being more expressive than mesh warping. For example, rather than requiring all the correspondence points of Fig. 2 to lie on a mesh, line pairs can be drawn along the mouth, nose, eyes, and cheeks of the source and target images. Therefore, only key feature points need be given.

Although this approach simplifies the specification of feature correspondence, it complicates warp generation. This is due to the fact that all line pairs must be considered before the mapping of each source point is known. This global algorithm is slower than mesh warping, which uses bicubic interpolation to determine the mapping of all points not lying on the mesh. An optimization based on a piecewise linear approximation is offered by Lee et al. (1998b) to accelerate the process. A more serious difficulty, though, is that unexpected displacements may be generated after the influence of all line pairs are considered at a single point. Additional line pairs must sometimes be supplied to counter the ill effects of a previous set. In the hands of talented animators, though, both the mesh warping and field morphing algorithms have been used to produce startling visual effects.

2.3 Radial basis functions/thin plate splines

The most general form of feature specification permits the feature primitives to consist of points, lines, and curves. Since lines and curves can be point sampled, it is sufficient to consider the features on an image to be specified by a set of points. In that case, the x and y components of a warp can be derived by constructing the surfaces that interpolate scattered points. Consider, for example, M feature points labeled (u_k, v_k) in the source image and (x_k, y_k) in the target image, where $1 \leq k \leq M$. Deriving warp functions that map points from the target image to the source image is equivalent to determining two smooth surfaces: one that passes through points (x_k, y_k, u_k) and the other that passes through (x_k, y_k, v_k) for $1 \leq k \leq M$.

This formulation permits us to draw upon a large body of work on scattered data interpolation to address the warp generation problem. All subsequent morphing algorithms have facilitated general feature specification by appealing to scattered data interpolation.

Warp generation by this approach is extensively surveyed by Ruprecht and Müller (1995) and Wolberg (1990). Lee et al. (1994) and Litwinowicz and Williams (1994) independently propose two similar methods that employ the thin plate surface model. Hassanien and Nakajima (1998) describe a related Navier spline technique, which uses Navier partial differential equations to construct the warp function. Arad et al. (1994) describe another method using radial basis functions. These techniques generate smooth warps that exactly reflect the feature correspondence. Furthermore, they offer the most general form of feature specification, since any primitive (e.g., spline curves) may be sampled into a set of points. Elastic Reality, a commercial morphing package from Avid Technology (www.avid.com), uses curves to enhance feature specification. Their warp generation method, however, is unpublished.

2.4 Energy minimization

All of the methods just described do not guarantee the one-to-one property of the generated warp functions. When a warp is applied to an image, the one-to-one property prevents the warped image from folding back upon itself. Lee et al. (1996) propose an energy minimization method for deriving one-to-one warp functions. Their method allows extensive feature specification primitives such as points, polylines, and curves. Internally, all primitives are sampled and reduced to a collection of points. These points are then used to generate a warp that is interpreted as a 2D deformation of a rectangular plate. A deformation technique is provided to derive C^1-continuous and one-to-one warps from the positional constraints. The requirements for a warp are represented by energy terms and satisfied by minimizing their sum. The technique generates natural warps since it is based on physically meaningful energy terms. The performance of the method, however, is hampered by its high computational cost.

2.5 Multilevel free-form deformation

Lee et al. (1995) present a new warp generation method that is much simpler and faster than the related energy minimization method (Lee et al. 1996). Large performance gains are achieved by applying multilevel free-form deformation (MFFD) across a hierarchy of control lattices to generate one-to-one and C^2-continuous warp function. In particular, warps are derived from positional constraints by introducing the MFFD as an extension to free-form deformation (FFD) (Sederberg and Parry 1986). In their paper, the bivariate cubic B-spline tensor product is used to define the FFD function. A new direct manipulation technique for FFDs, based on 2D B-spline approximation, is applied to a hierarchy of control lattices to exactly satisfy the positional constraints. To guarantee the one-to-one property of a warp, a sufficient condition for a 2D cubic B-spline surface to be one-to-one is presented. The MFFD generates C^2-continuous and one-to-one warps that yield fluid image distortions. The MFFD algorithm is combined with the energy minimization method (Lee et al. 1996) in a hybrid approach.

An example of MFFD-based morphing is given in Fig. 3. Notice that the morph sequence shown in the middle row of the figure is virtually identical to that produced by mesh warping in Fig. 2. The benefit of this approach, however, is that feature specification is more expressive and less cumbersome. Rather than editing a mesh, a small set of feature primitives are specified. To further assist the user, snakes are introduced to reduce the burden of feature specification. Snakes are energy-minimizing splines that move under the influence of image and constraint forces. They were first adopted in computer vision as an active contour model (Kass et al. 1988). Snakes streamline feature specification because primitives must only be positioned near the features. Image forces push snakes toward salient edges, thereby refining their final positions and making it possible to capture the exact position of a feature easily and precisely.

2.6 Work minimization

Given that two images are sufficiently similar, image registration techniques (Brown 1992) offer a potential solution for automatically determining the correspondence among all points across the source and target images. Optic flow techniques, for instance, are regularly used to track the frame-to-frame motion of video frames for use in compression and digital video processing (Tekalp 1995). Such methods exploit the fact that the two frames represent the same scene imaged only moments apart, and therefore only small displacements must be recovered. This is generally not true in morphing applications, where the source and target images may be vastly different. Nevertheless, it is reasonable to consider the extent to which the feature specification problem can be automated for similar input images. This is precisely the problem considered by Gao and Sederberg (1998).

Gao and Sederberg (1998) present a new algorithm for determining the warp function between two similar images with little or no human interaction. This approach is consistent with the trend towards minimizing the amount of user interaction in specifying feature correspondence. The algorithm is based on a work minimization approach that derives its cost directly from the image intensities, and not from user-specified constraints. Motivated from related work conducted by the authors for solving the polygon shape blending problem (Sederberg and Greenwood 1992), they attempt to minimize the cost associated with warping and recoloring the image. They apply a hierarchical warp using a grid of bilinear B-spline patches and associate a cost function to it. Even after the warp is applied, there will be differences in the intensity values and a cost function is associated with recoloration. The morph generation process is then reduced to that of finding the global minimum cost function among all possible warps. The authors offer a method that is simple and fast, albeit it does not guarantee a global minimum.

The premise of the work minimization algorithm is that the most visually pleasing morphs are associated with those images that satisfy the minimal work constraints. This assumption is sometimes false, and therefore some user interaction is necessary to specify features and initialize a warp. In spite of any initial warp that may need to be specified, the rationale behind this approach is that user interaction can be significantly reduced by exploiting work minimization constraints. Although this work is still in its early stages, it has shown promising results for morphing among similar images.

2.7 Discussion

The progression of morphing algorithms has been marked by more expressive and less cumbersome tools for feature specification. A significant step beyond meshes was made possible by the specification of line pairs in field morphing. The complications that this brought to warp generation, however, sometimes undermined the usefulness of the approach. For instance, the method sometimes produced undesirable artifacts, referred to as *ghosts*, due to the computed warp function (Beier and Neely 1992). To counter these problems, the user was required to specify additional line pairs beyond the minimal set that would otherwise be warranted. All subsequent algorithms, including those based on radial basis functions, thin plate splines, and energy minimization, formulated warp generation as a scattered data interpolation problem and sought to improve the quality (smoothness) of the computed warp function. They did so at a relatively high computational cost. The newest approach, based on the MFFD algorithm, significantly improves matters by accelerating warp generation. The use of snakes further assists the user in reducing the burden of feature specification. The ultimate goal of eliminating user interaction from feature specification remains elusive, but the work by Gao and Sederberg (1998) demonstrates the feasibility of an algorithm approaching this goal when the source and target images are similar.

3 Transition control

Transition control determines the rate of warping and color blending across the morph sequence. If transition rates differ from part to part in inbetween images, more interesting animations are possible. Such nonuniform transition functions can dramatically improve the visual content. Note that the examples shown thus far all use a uniform transition function, whereby the positions of the source features steadily move to their corresponding target positions at a constant rate.

Figures 4 and 5 show examples of the use of uniform and nonuniform transition functions, respectively. The upper left and lower right images of Fig. 4 are the source and target images, respectively. The features used to define the warp functions are shown overlaid on the two images. The top and bottom rows depict a uniform transition rate applied to the warping of the source and target images, respectively. Notice, for instance, that all points

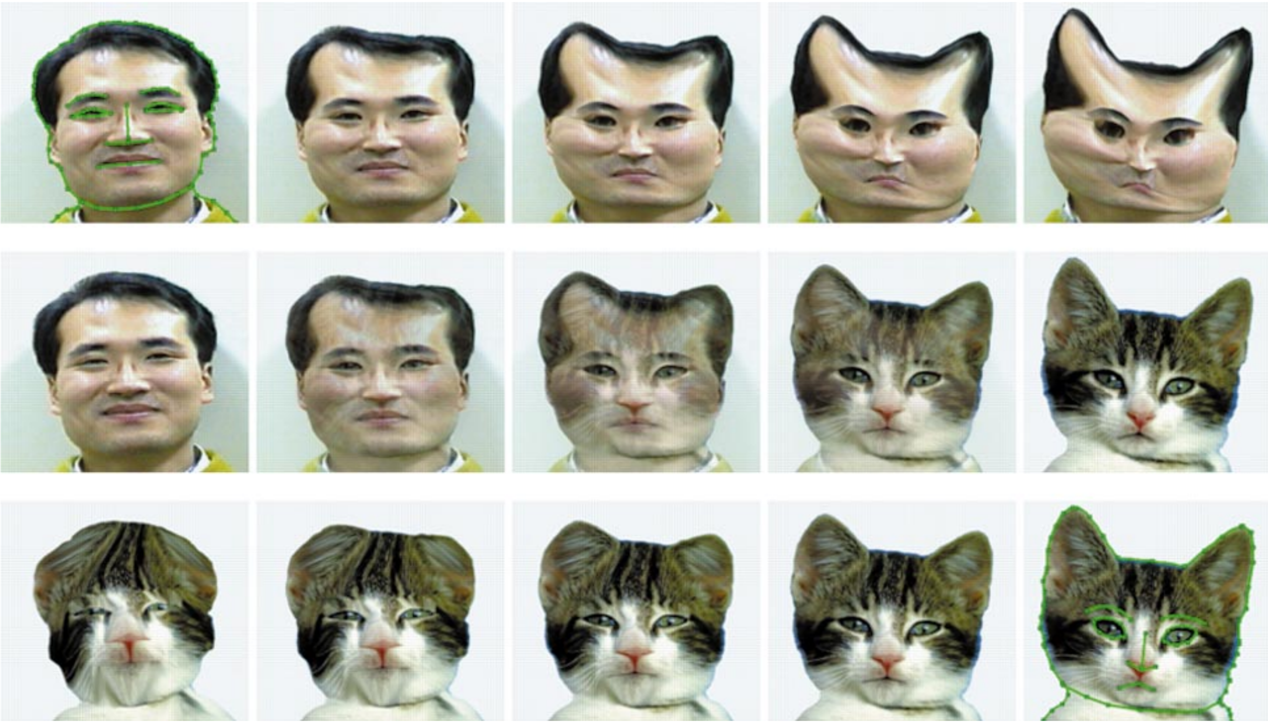

4

5

Fig. 4. Uniform metamorphosis

Fig. 5. Nonuniform metamorphosis

Fig. 6. Procedural transformation

in the source and target images are moving at a uniform rate to their final positions. Those two rows of warped imagery are attenuated by the same transition functions and added together to yield the middle row of inbetween images. Note that geometric alignment is maintained among the two sets of warped inbetween images before color blending merges them into the final morph sequence.

The example in Fig. 5 demonstrates the effects of a nonuniform transition function applied to the same source and target images. In this example, a transition function has been defined that accelerates the deformation of the nose of the source image, while leaving the shape of the head intact for the first half of the sequence. The deformation of the head begins in the middle of the sequence and continues linearly to the end. The same transition function is used for the bottom row. Notice that this use of nonuniform transition functions is responsible for the dramatic improvement in the morph sequence.

Transition control in mesh-based techniques is achieved by assigning a transition curve to each mesh node. This proves tedious when complicated meshes are used to specify the features. Nishita et al. (1993) mention that the transition behavior can be controlled by a Bézier function defined on the mesh.

In the energy minimization method, transition functions are obtained by selecting a set of points on a given image and specifying a transition curve for each point. Although earlier morphing algorithms generally couple the feature specification and transition control primitives, this method permits them to be decoupled. That is, the location of transition control primitives does not necessarily coincide with those of the features. The transition curves determine the transition behavior of the selected points over time. For a given time, transition functions must have the values assigned by the transition curves at the selected points. Considering a transition rate as the vertical distance from a plane, transition functions are reduced to smooth surfaces that interpolate a set of scattered points. The thin plate surface model (Terzopoulos 1983) is employed to obtain C^1-continuous surfaces for transition functions.

In the MFFD-based approach, the MFFD technique for warp generation is simplified and applied to efficiently generate a C^2-continuous surface for deriving transition functions. The examples in Figs. 4 and 5 were generated with the MFFD-based morphing algorithm.

Transition curves can be replaced with procedural transition functions (Lee et al.1995, 1996). An example is depicted in Fig. 6, where a linear function varying in the vertical direction is applied to two input images. The result is a convincing transformation in which one input image varies into the other from top to bottom.

4 Generalized morphing framework

The traditional formulation for image morphing considers only two input images at a time, i.e., the source and target images. In this case, morphing among multiple images is understood to mean a series of transformations from one image to another. This limits any morphed image to take on the features and colors blended from just two input images. Given the success of morphing using this paradigm, it is reasonable to consider the benefits possible from a blend of more than two images at a time. For instance, consider the generation of a facial image that is to have its eyes, ears, nose, and profile derived from four different input images. In this case, morphing among multiple images is understood to mean a seamless blend of several images at once. We refer to this process as *polymorph*.

Rowland and Perrett (1995) consider a special case of polymorph to obtain a prototype face, in terms of gender and age, from several tens of sample faces. They superimpose feature points on input images to specify the various positions of features in sample faces. The shape of a prototype face is determined by averaging the specified feature positions. A prototype face is generated by applying image warping to each input image and then performing a cross-dissolve operation among the warped images. The shape and color differences between prototypes from different genders and ages are used to manipulate a facial image to perform predictive gender and age transformations.

Despite the explosive growth of morphing in recent years, the subject of morphing among multiple images has been neglected. In recent work conducted by the author and his colleagues, a general framework is developed that extends the traditional image morphing paradigm applied to two images (Lee et al.1998a). We formulate each input image to be a vertex of a regular convex polyhedron in (n–1)-dimensional space, where n is the number of input images. An inbetween (morphed) image is considered to be a point in the convex polyhedron. The barycentric coordinates of the point determine the weights used to blend the input images into the inbetween image.

In morphing among two images, nonuniform blending was introduced to derive an inbetween image in which blending rates differ across the image (Lee et al.1995, 1996). This allows us to generate more interesting animations such as transformation of the source image to the target from top to bottom. Nonuniform blending has also been considered in volume metamorphosis to control blending schedules (Hughes 1992; Lerios et al.1995). In (Lee et al. 1998a), the framework for polymorph includes nonuniform blending of features in several input images. For instance, a facial image can be generated to have its eyes, nose, mouth, and ears derived from four different input faces.

An example is shown in Fig. 7. The top row of Fig. 7 shows input images I_0, I_1, and I_2. We selected three groups of features in these images and assigned them unity blending values to generate inbetween images. The feature groups consist of the hair, eyes and nose, and mouth and chin. Each feature group was selected in a different input image. For instance, Fig. 8 shows the feature groups selected to generate the leftmost inbetween image in the middle row of Fig. 7. The lower rightmost image in the figure shows their projections onto a single image. The middle and bottom rows of Fig. 7 show the inbetween images resulting from all possible combinations in selecting those feature groups.

Polymorph is ideally suited for image composition applications in which elements are seamlessly blended from two or more images. A composite image is treated as a metamorphosis of selected regions in several input images. The regions seamlessly blend together with respect to geometry and color. In future work, we will determine the extent to which the technique produces high-quality composites with considerably less effort than conventional image composition techniques. In this regard, the technique can bring to image composition what image warping has brought to cross-dissolve in deriving morphing: a richer and more sophisticated class of visual effects that are achieved with intuitive and minimal user interaction.

5 Applications

Image morphing has traditionally been associated with visual effects for entertainment. Visually compelling fluid transformations are created by synthesizing intermediate images between supplied image pairs. The basis for these results is attributed to the geometric alignment that is maintained throughout the image sequence. This same result applies to other domains where image interpolation can benefit from supplied geometric correspondence. One such application includes medical visualization.

In medical imaging, CT or MRI scans are available at a fixed interslice resolution. Although intermediate slices can be computed with conventional linear, cubic, or higher-degree interpolation functions, this traditional approach does not consider the underlying structure of the imaged organs. Superior results are made possible by establishing geometric correspondence of features among successive pairs of scans. Features may include organ

7

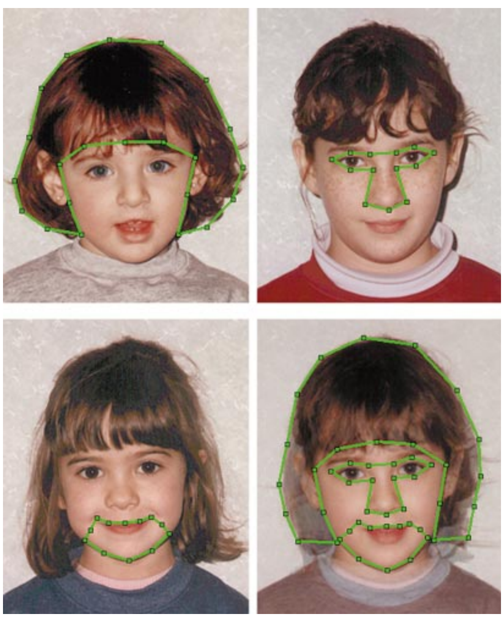

8

Fig. 7. Input images (*top row*) from which combinations of the hair, eyes-and-nose, and mouth-and-chin (*middle and bottom rows*) have been taken

Fig. 8. Selected parts on input images and their projections onto the central image

boundaries and anatomical landmarks. The computed morphed images constitute the intermediate images that conform to the geometric correspondence provided by the user. Ruprecht and Müller (1994) describe work in this area. Automatic feature selection remains an active area of research, and the feature selection process remains largely manual.

Recent work in face recognition has also drawn upon morphing. Bichsel (1996) presents a new method for generating an optimum mapping of one image into another. The mapping is calculated by maximizing the probability of the morph field in a Bayesian framework. The approach uses an iterative approximation where, in a neighborhood around the current morph field solution, the probability distribution is approximated by a Gaussian probability distribution for which the most likely solution is evaluated by linear algebra techniques. This method has been used to interpolate between different views of a single face and between images of different persons. The technique was also demonstrated on face recognition under varying view and illumination conditions.

An automated technique to recover facial features for automatic image morphing has been explored by Covell (1996). That work has been demonstrated for use in lip syncing (Bregler et al. 1997). In an application that the authors refer to as video rewrite, the system automatically synthesizes faces with proper lip sync. It is a facial animation system that is driven by audio input, and it can be used for dubbing movies, teleconferencing, and special effects.

In related work on facial analysis and synthesis, Ezzat and Poggio (1996) explore an image-based facial modeling approach, in which they model facial movements using example images. They exploit the viability of a view interpolation approach to image synthesis. A learning network is trained on example images, each of which is associated with a position in a high-level, multidimensional parameter space denoting pose and expression. The authors demonstrate their analysis-by-synthesis technique by estimating pose, eye, and mouth movements, as well as head translations, rotations, and scales.

6 Conclusions

This paper has reviewed the growth of image morphing and described recent advances in the field. Morphing algorithms all share the following components: feature specification, warp generation, and transition control. The ease with which an artist can effectively use morphing tools is determined by the manner in which these components are addressed. We surveyed various morphing techniques, including those based on mesh warping, field morphing, radial basis functions, thin plate splines, energy minimization, and multilevel free-form deformations.

The earliest morphing approach was based on mesh warping. It was motivated by a reasonably straightforward interface requiring meshes to mark features and bicubic spline interpolation to compute warp functions. The field morphing approach attempted to simplify feature specification with the use of line pairs to select landmarks. This added benefit demanded a more computationally expensive warp generation stage. Subsequent morphing algorithms have sought to maintain the use of curves and polylines to select features. Warp generation has consequently become formulated as a scattered data interpolation problem. Standard solutions such as radial basis functions and thin plate splines have been presented.

The newest approach based on multilevel free-form deformations has further accelerated warp generation. This approach uses snakes to assist the user in placing feature primitives, thereby reducing the burden in feature specification. Snakes are particularly useful when features lie along large intensity gradients.

Transition control determines the rate of warping and color blending across the morph sequence. If transition rates differ from part to part in inbetween images, more interesting animations are possible. The techniques used to compute warp functions can also be applied for transition control functions, thereby propagating that information everywhere across the image. Although early morphing algorithms generally coupled the feature specification and transition control primitives, more recent algorithms have permitted them to be decoupled.

For those cases where the input images are sufficiently similar, feature specification can potentially be automated. The algorithm proposed by Gao and Sederberg (1998) addresses this problem and offers some promising results. The attempt to attain feature specification with no user interaction will likely remain elusive for morphing applications. Aside from the standard registration problem, a key difficulty with fully automating feature correspondence is the aesthetic issue that must be considered in generating a visually pleasing morph. High-level understanding of the input images is necessary to specify the correspondence of points in areas where feature information is either partially occluded or missing.

Future morphing work includes further work in morphing among multiple images, and improving automatic morphing among a limited class of images and video sequences. In the limited, but common, class of facial images, researchers have explored the use of computer vision techniques to automatically register features between two images. As the examples given in Figs. 2 and 3 show, facial images require feature primitives to be specified along the eyes, nose, mouth, hair, and profile. Model-based vision has been investigated to exploit knowledge about the relative position of these features and automatically locate them for feature specification (Brunelli and Poggio 1993). Currently, this is an active area of research, particularly for compression schemes designed for videoconference applications. The same automation applies to morphing among two video sequences, where time-varying features must be tracked. Progress in this area is tied to that of the motion estimation research community.

Acknowledgements. This work was supported in part by a NASA Faculty Award for Research grant (NAG-7129) and a PSC-CUNY grant (RF-668373). The author thanks S. Lee for valuable discussions and help with some of the figures.

References

Arad N, Dyn N, Reisfeld D, Yeshurun Y (1994) Image warping by radial basis functions: applications to facial expressions. CVGIP: Graph Models Image Processing 56:161–172

Beier T, Neely S (1992) Feature-based image metamorphosis. (Proceedings of SIGGRAPH'92) Comput Graph 26:35–42

Bichsel M (1996) Automatic interpolation and recognition of face images by morphing. The 2nd International on Automatic Face and Gesture Recognition. Killington, Vermont, IEEE Computer Society Press, Los Alamitos, CA, pp 128–135

Bregler C, Covell M, Slaney M (1997) Video rewrite: driving visual speech with audio. (Proceedings of SIGGRAPH'97) Comput Graph, pp 353–360

Brown LG (1992) A survey of image registration techniques. ACM Comput Surveys 24:325–376

Brunelli R, Poggio T (1993) Face recognition: Features versus templates. IEEE Trans Patt Anal Machine Intell 15:1042–1052

Chen SE, Williams L (1993) View interpolation for image synthesis. (Proceedings of SIGGRAPH'93) Comput Graph pp 279–288

Covell M (1996) Eigen-points: control-point location using principal component analysis. The 2nd International on Automatic Face and Gesture Recognition, Killington, IEEE Computer Society Press, Vermont, pp 122–127

Ezzat T, Poggio T (1996) Facial analysis and synthesis using image-based models. The 2nd International on Automatic Face and Gesture Recognition, Killington, IEEE Computer Society Press, Los Alamitos, CA, pp 116–121

Fant KM (1986) A nonaliasing, real-time spatial transform technique. IEEE Comput Graph Appl 6:71–80

Gao P, Sederberg TW (1998) A work minimization approach to image morphing. Visual Comput (this volume)

Hassanien AE, Nakajima M (1998) Image morphing of the facial image transformation based on Navier elastic body splines. Proceedings of Computer Animation'98, Geneva, Switzerland, IEEE Computer Society Press, Los Alamitos, CA, pp 119–125

Hughes JF (1992) Scheduled fourier volume morphing. (Proceedings of SIGGRAPH'92) Comput Graph 26:43–46

Kass M, Witkin A, Terzopoulos D (1988) Snakes: active contour models. Int J Comput Vision 1(4):321–331

Lee S, Chwa K-Y, Hahn J, Shin SY (1994) Image morphing using deformable surfaces. Proceedings of Computer Animation'94, IEEE Computer Society Press, Geneva, Switzerland, pp 31–39

Lee S, Chwa K-Y, Shin SY, Wolberg G (1995) Image metamorphosis using snakes and free-form deformations. (Proceedings of SIGGRAPH'95) Comput Graph, pp 439–448

Lee S, Chwa K-Y, Hahn J, Shin SY (1996) Image morphing using deformation techniques. J Visualization Comput Anim 7:3–23

Lee S, Wolberg G, Shin SY (1998a) Polymorph: morphing among multiple images. IEEE Comput Graph Appl 18:58–71

Lee T, Lin Y-C, Lin L, Sun Y (1998b) Fast feature-based metamorphosis and operator design. (Proceedings of Eurographics'98), Computer Graphics Forum, 17(3):15–22

Lerios A, Garfinkle CD, Levoy M (1995) Feature-based volume metamorphosis. (Proceedings of SIGGRAPH'95) Comput Graph, pp 449–456

Litwinowicz P, Williams L (1994) Animating images with drawings. (Proceedings of SIGGRAPH'94) Comput Graph, pp 409–412

Nishita T, Fujii T, Nakamae E (1993) Metamorphosis using Bézier clipping. Proceedings of the 1st Pacific Conference on Computer Graphics and Applications, Seoul, Korea, World Scientific, Singapore, pp 162–173

Rowland DA, Perrett DI (1995) Manipulating facial appearance through shape and color. IEEE Comput Graph Appl 15:70–76

Ruprecht D, Müller H (1994) Deformed cross-dissolves for interpolation in scientific visualization. J Visualization Comput Anim 5:167–181

Ruprecht D, Müller H (1995) Image warping with scattered data interpolation. IEEE Comput Graph Appl 15:37–43

Sederberg TW, Greenwood E (1992) A physically based approach to 2D shape blending. (Proceedings of SIGGRAPH'92) Comput Graph 26:25–34

Sederberg TW, Parry SR (1986) Free-form deformation of solid geometric models. (Proceedings of SIGGRAPH'86) Comput Graph 20:151–160

Seitz SM, Dyer CR (1996) View morphing. (Proceedings of SIGGRAPH'96) Comput Graph, pp 21–30

Smythe DB (1990) A two-pass mesh warping alogrithm for object transformation and image interpolation. Technical Report 1030, ILM Computer Graphics Department, Lucasfilm, San Rafael, Calif

Tekalp AM (1995) Digital video processing. Prentice Hall PTR, Upper Saddle River, NJ

Terzopoulos D (1983) Multilevel computational processes for visual surface reconstruction. Comput Vision Graph Image Processing 24:52–96

Wolberg G (1990) Digital image warping. IEEE Computer Society Press, Los Alamitos, Calif

GEORGE WOLBERG is a Full Professor of Computer Science at the City College of New York (CCNY). He received his BS and MS degrees in Electrical Engineering from Cooper Union in 1985, and his PhD degree in Computer Science from Columbia University in 1990. His research interests include image processing, computer graphics, and computer vision. Prof. Wolberg is the recipient of a 1991 National Science Foundation (NSF) Presidential Young Investigator Award and the 1997 CCNY Outstanding Teaching Award. He is the author of *Digital Image Warping*, the first comprehensive monograph on warping and morphing.