Mapping and Manipulating Facial Expression

Barry-John Theobald[1]**, Iain Matthews**[2]**,
Michael Mangini**[3]**, Jeffrey R. Spies**[4]**,
Timothy R. Brick**[4]**, Jeffrey F. Cohn**[5]**,
Steven M. Boker**[4]

[1] *University of East Anglia, Norwich, U.K.*
[2] *Weta Digital Ltd., Miramar, New Zealand*
[3] *University of Notre Dame, Indiana, U.S.A.*
[4] *University of Virginia, Charlottesville, U.S.A.*
[5] *University of Pittsburgh, Pittsburgh, U.S.A.*

Key words

active appearance models

expression cloning

expression mapping

Abstract

Nonverbal visual cues accompany speech to supplement the meaning of spoken words, signify emotional state, indicate position in discourse, and provide back-channel feedback. This visual information includes head movements, facial expressions and body gestures. In this article we describe techniques for manipulating both verbal and nonverbal facial gestures in video sequences of people engaged in conversation. We are developing a system for use in psychological experiments, where the effects of manipulating individual components of nonverbal visual behavior during live face-to-face conversation can be studied. In particular, the techniques we describe operate in real-time at video frame-rate and the manipulation can be applied so both participants in a conversation are kept blind to the experimental conditions.

1 Mapping and manipulating facial expression

Audiovisual conversation is information rich, comprised of both the linguistic message content and non-linguistic information. This non-linguistic information might include facial expressions to convey emotional state, head nods to indicate emphatic stress, and posture and eyegaze to signal turn-taking in discourse. There is increasing interest

Language and Speech

in understanding the relationship between verbal and nonverbal conversational cues. In large part, this interest comes from the entertainment industry, which requires life-like animation of embodied characters. If interactive dialogue between animated characters is to be believable, it must seem realistic and lifelike.

In psychology and speech perception, the ability to generate experimental stimuli using computer-generated characters is attractive as individual aspects of verbal and nonverbal cues can be manipulated to arbitrary precision and in complete isolation from each other (Massaro, 1998). In this article, we describe the development of a system for automatically capturing, encoding, and manipulating facial expressions and global head movements, and transferring this information from video sequences of one person to video sequences of another. In particular, we describe techniques that allow these manipulations to be accomplished in real-time at video frame-rate during face-to-face conversations and with both participants in the conversation being kept blind to the processing and manipulation of the visual cues accompanying the speech. This system is being developed to provide experimenters with the means for controlled investigation to better understand the mechanisms of communication and social interaction.

2 Background and related work

Computer-generated models of faces are useful in studying speech perception because the experimenter has precise control over the stimuli. For example, the use of a model allows the amplitudes of expressive face movements to be exaggerated or attenuated. Likewise, the timing of facial expressions can be manipulated to create arbitrary durations. Many of these effects would be impossible even for highly accomplished actors. An effective face model can make them possible.

In some ways, the freedom of manipulation provided by face models presents a problem for researchers; models can extrapolate beyond what is plausible or acceptable. For experimental validity, it is important that the stimuli presented to the viewer be acceptably realistic. Constraining the appearance of generated facial expressions is a nontrivial challenge. The changes in the appearance of a face that we interpret as a facial expression are the product of the complex interaction of neuromuscular signals, anatomical structures (including bone, muscle, subcutaneous fat, and skin), and light (Terzopoulos & Waters, 1993). The difficulty of ensuring that only plausible facial expressions are generated is compounded by the fact that people are expert at detecting and recognizing facial expressions and are acutely sensitive to even minute discrepancies from normal behavior. As the realism of computer-generated models increases, the viewer generally becomes more reactive to subtle errors in dynamic sequences – errors are more readily detected and the face may appear "zombie-like." This occurrence often is referred to as the *Uncanny Valley* effect (Mori, 1970).

Generally, the most pleasing, natural, and lifelike facial animation is achieved using *performance-driven* techniques. To implement this approach, feature points are located on the face of an actor and the displacement of these feature points over time is used either to update the vertex locations of a polygonal model (Guenter, Grimm, Wood, Malvar, & Pighin, 1998; Noh & Neumann, 2001; Williams, 1990), or are mapped to an underlying muscle-based model (Choe, Lee, & Ko, 2001). Performance-driven

facial animation enables subtle facial movements to be captured and transferred to the animated face model, but there are several inherent limitations that make this approach unsuitable in many applications. Extensive time is required to setup the capture environment. The motion capture markers (e.g., LEDs or reflective tape or ink) must be applied to the actor, and they may inhibit ease or naturalness of movement. Finally, consistent placement of these markers must be maintained between capture sessions.

An alternative is image-based animation, in which changes in the position of the facial features observed in images or video sequences of one person are used to warp images of another person (Chuang & Bregler, 2002; Chang & Ezzat, 2005; Liu, Shan, & Zhang, 2001; Zhang, Liu, Quo, Terzopoulos, & Shum, 2006). In this way, there is no longer need for motion capture, which considerably reduces setup time. On the down side, the range of convincing facial expressions that can be transferred is limited to those that actually occurred. It also is difficult to manipulate facial expressions, such as exaggerate their intensity or qualitative appearance, as the face is represented only as raw pixel values. There is no underlying parametric model available to provide quantitative guidance in constructing realistic image manipulation.

A hybrid approach based on Blanz and Vetter's (1999) three-dimensional morphable model (3DMM) transfers facial expressions from images of one person's face to images of another person. This expression transfer is achieved by *copying* to the new face the changes in facial geometry and facial appearance that result from displaying a facial expression (Blanz, Basso, Poggio, & Vetter, 2003). 3DMMs have the advantage of a parameterized face model that provides control over the deformation of the facial geometry, and an image-based texture-map that provides near-photorealistic rendering of the facial appearance. 3DMMs are capable of producing highly convincing results, but there are several practical limitations that restrict their use. Firstly, this approach requires a collection of 3D laser range scan images to build the face model. Secondly, changes in face shape and appearance that result from movements of the facial features as measured in images of one person are *copied* directly to images of another person. This makes it difficult to model and apply person-specific characteristics of facial expressions. Thirdly, because the inside of the mouth is difficult to model, the same generic inner mouth must be used for all subjects. And fourthly, the algorithms are computationally expensive, which currently rules out real-time applications. Approaches that address these limitations are needed.

The lack of facial expression mapping in the system proposed by Blanz et al. (2003) was addressed recently in Vlasic, Brand, Pfister, and Popovic (2005) using multilinear face models to represent changes in images that are due to identity, age, and facial expression. The learned multilinear model can be used to predict the appearance of common facial expressions in images of new, unseen faces. However, to construct the multilinear model, a 3D laser range scanner is required to capture facial images of a number of people displaying the same prespecified facial expressions (e.g., smile, frown, surprise, and so on). While this approach remains computationally expensive, it can be approximated efficiently in 2D (Macedo, Brazil, & Velho, 2006) using Active Appearance Models (AAMs) (Cootes, Edwards, & Taylor, 2001), but this approximation still requires the used of a common set of posed facial expressions from which the multilinear model is learned.

We have developed an AAM-based system that captures, manipulates, and transfers facial expressions, visual speech gestures, and changes in head pose from video sequences, all in real-time at video frame-rates. Our approach requires no marker-based tracking, offline processing, or high-level semantic descriptors. The full face, including the eyes and inside of the mouth, can also be transferred. This makes it possible for this approach to be used for analyzing and manipulating real-time conversation in an interactive two-way videoconference experimental paradigm.

3 Linear deformable face models

The key to our approach is a linear deformable model that can both analyze and synthesize images of faces. The idea of representing an ensemble of images of faces as a compact statistical model was first proposed by Sirovich and Kirby (1987). They noted that humans store and recognize a large number of faces almost instantaneously, which suggests the information required to encode faces may be represented in only a few tens of parameters (Kirby & Sirovich, 1990). They proposed a compact linear face model of the form,

$$\mathbf{F} = \mathbf{F_0} + \sum_{i=1}^{m} \mathbf{F}_i p_i, \tag{1}$$

where $\mathbf{F}$ is an instance of a face, $\mathbf{F}_0$ is a reference face, $\mathbf{F}_i$ are basis faces that characterize the allowed deformation from the reference face, and each p_i defines the contribution of the corresponding basis face in representing $\mathbf{F}$. Thus, faces are not coded directly, but rather indirectly as displacements from the reference.

Turk and Pentland (1991) developed the idea proposed by Sirovich and Kirby (1987) into a system for recognizing faces in images. Thus, $\mathbf{F}$ is a (vectorized) face image, $\mathbf{F}_0$ is the mean (vectorized) image computed from a training set of face images, and each $\mathbf{F}_i$ is referred to as an *Eigenface*. The model is usually obtained by applying principal components analysis (PCA) to a set of training images containing faces. Linear models have since become one of the most popular techniques for modeling flexible objects, such as faces, and they have a wide range of applications, which include: low-bandwidth transmission of face images (Koufakis & Buxton, 1999; Mogghaddam & Pentland, 1995; Theobald, Kruse, Bangham, & Cawley, 2003), face recognition and verification (Blanz & Vetter, 2003; Edwards, Cootes, & Taylor, 1998; Shackleton & Welsh, 1991; Yang, Zhang, Frangi, & Yang, 2004), automated lip-reading (Luettin, Thacker, & Beet, 1996; Matthews, Cootes, Bangham, Harvey, & Cox, 2002), and face synthesis (Cosker, Marshall, Rosin, Paddock, & Rushton, 2004; Fagel, Bailly, & Elisei, 2007; Govokhina, Bailly, Bretin, & Bagshaw, 2006; Hong, Wen, Huang, & Shum, 2002; Theobald, Bangham, Matthews, & Cawley, 2004; Theobald & Wilkinson, 2007).

3.1
Active Appearance Models (AAMs)
Perhaps the most common form of linear deformable face model is the AAM (Cootes, Edwards, & Taylor, 2001), which provides a compact statistical representation of the shape *and* appearance variation of the face as measured in 2D images. A recent

extension, the so called 2D+3D AAM (Xiao, Baker, Matthews, & Kanade, 2004), provides a convenient method for constructing models of 3D shape variation from 2D images, overcoming the need for an expensive 3D laser range scanner required for a 3DMM.

The *shape* of an AAM is defined by a 2D triangulated mesh and in particular the vertex locations of the mesh. Mathematically, the shape, $\mathbf{s}$, of an AAM is defined as the concatenation of the x and y-coordinates of n vertices that make up the mesh:

$$\mathbf{s} = (x_1, y_1, x_2, y_2 \ldots, x_n, y_n)^{\mathrm{T}}.$$

For face models, vertices typically are defined to delineate the eyes, eyebrows, nose, mouth and the outline of the face. A compact model that allows a linear variation in the shape is given by,

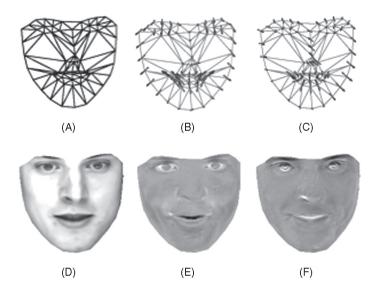$$\mathbf{s} = \mathbf{s}_0 + \sum_{i=1}^{m} \mathbf{s}_i p_i, \tag{2}$$

where the coefficients p_i are the shape parameters that allow the mesh to undergo both rigid (head pose) and non-rigid (expression) variation. Such a model is usually computed by applying PCA to a set of meshes that have been aligned with a corresponding set of training images. Typically this alignment is done manually, although automated approaches have shown some success (Wang, Lucey, & Cohn, 2008). The base shape $\mathbf{s}_0$ is the mean shape, and the vectors $\mathbf{s}_i$ are the (reshaped) eigenvectors corresponding to the m largest eigenvalues. Care is required when hand-labeling images to ensure that all example shapes have the same number of vertices, and vertices are labeled in the same order in all images. Hand-labeling errors can result in significant sources of variation in the shape component of the AAM.

The *appearance*, A($\mathbf{x}$), of an AAM is an image that exhibits variation under the control of the model parameters and is defined as the pixels, $\mathbf{x} = (x, y)^{\mathrm{T}}$, that lie within the base mesh, $\mathbf{x} \in \mathbf{s}_0$. AAMs allow linear appearance variation. That is, the appearance can be expressed as a base appearance, $\mathbf{x} \in \mathbf{s}_0$ plus a linear combination of l appearance images, $A_i(\mathbf{x})$:

$$A(\mathbf{x}) = A_0(\mathbf{x}) + \sum_{i=1}^{l} A_i(\mathbf{x})\lambda_i \quad \forall \mathbf{x} \in \mathbf{s}_0, \tag{3}$$

where the coefficients λ_i are the appearance parameters. The appearance component of the AAM is usually computed by first warping the manually annotated training images to the base shape, then applying PCA to the shape-normalized images (Cootes et al., 2001). The base appearance $A_0(\mathbf{x})$ is the mean shape normalized image, and the vectors $A_i(\mathbf{x})$ are the (reshaped) eigenvectors corresponding to the l largest eigenvalues. The advantage of this two-stage approach to modeling the face is that the resulting modes of variation are more compact and cleaner, in the sense that they model changes in the facial features rather than change in the images, as is the case with *Eigenfaces*. An example AAM is shown in Figure 1.

Figure 1

An example AAM derived from 15 hand-labelled images from a longer sequence of a single individual. Shown are the mean shape and appearance (A and D respectively), and the first two modes of variation of the shape (B–C) and appearance (E–F) components of the model. Note the AAM includes the region of the face only to the eyebrows. There are no defining features that are easily identifiable on the forehead that can consistently be placed manually across individuals. Work is ongoing to overcome this limitation

(A)   (B)   (C)

(D)   (E)   (F)

3.1.1
Facial image synthesis

Near-photorealistic images of faces are synthesized using an AAM by first applying the shape parameters, Equation (2), to generate the shape of the AAM, then applying the appearance parameters, Equation (3), to generate the AAM image, $A(\mathbf{x})$. The synthesized image is created by warping $A(\mathbf{x})$ from the base shape, s_0, to the model-synthesized shape using a piece-wise affine warp. Example images synthesized using an AAM are shown in Figure 2. We note here that the AAMs used in this work require a common structure. Each person-specific model requires the equivalent number of landmarks and pixels and the ordering of the landmarks must be consistent across all models. The shape of the model is constrained to include information in the region of the face only to the eyebrows. The forehead region is not included in our current work as there are no identifying features that can consistently be located across all individuals. Occlusion of the forehead by hair, for instance, is a problem. Solutions have been proposed (e.g., Gross, Matthews, & Baker, 2006); however, these solutions significantly increase the computational expense.

3.1.2
Multi-segment AAMs

The standard approach to constructing the appearance component of the AAM is to warp the images onto s_0 and concatenate *all* pixels bound by the mesh before applying

Figure 2

Example face images synthesized using AAMs trained on three different individuals

PCA. The assumption is that the probability distribution of the pixel intensities is Gaussian. However, this assumption often is not met for face images. Consequently, some important information might be considered noise and discarded. In addition, if the distribution of the data is multimodal or nonlinear, then linearly interpolating the basis vectors to re-synthesize face images can result in parameters that appear valid, in that they are within the limits of the training examples, but lie in invalid regions of the global parameter space (Bregler & Omohundro, 1995). When assumptions about the distribution of the data are not met, some blurring in the rendered images may occur. This consequence is most striking in the eyes and inner mouth, in particular the teeth. These are important areas of the face as they are ones on which viewers tend to focus their attention.

To address this issue, we apply a relatively straightforward extension to improve the quality of rendering by constructing a piece-wise PCA model. Independent appearance models are constructed for each region of the face (skin, eyes, nose, and inner-mouth). This can be done in the coordinate frame of $\mathbf{s}_0$, so the pixel indices for the different regions of the face are constant across all images. The appearance for individual segments then can be regenerated and copied into the appearance vector $A(\mathbf{x})$ before warping to the shape $\mathbf{s}$. This notion could be taken further by clustering the regions of the face according to the shape parameters, for example building separate models of appearance to account for an open and closed mouth. This also allows different model segments to be encoded with more or less resolution, giving more significance to regions of the face to which a viewer is likely to pay particular attention.

4 Manipulating facial expressions

In this section we illustrate the usefulness of AAMs for studies of social interaction. Individual verbal and nonverbal conversational cues obtained from the face in a video sequence can be manipulated independently so that their separate and combined effects can be studied. For example, one might wish to decouple the apparent sex of an individual from his/her behavior to determine each factor's influence during conversation. Or, one might wish to attenuate or exaggerate the degree of expressiveness to examine its influence on perceived personality traits. Exaggerated expressions and head movements are believed to increase a sense of liveliness and lead to greater rapport (Grahe & Bernieri, 1999). Attenuating expressiveness and head movement

may create a sense of disinterest, sadness, or depressed affect (Gehricke & Shapiro, 2000; Schwartz, Fair, Salt, Mandel, & Klerman, 1976). Using AAMs, the behavior and identity of each person can be manipulated during actual conversation without the individual's awareness that a change has occurred.

The following descriptions of mapping and manipulating facial expressions consider only the *intensity* of those expressions. Focus to date has been on developing techniques that operate in real-time at video frame-rate, allowing conversational cues to be manipulated during live face-to-face interaction in a videoconferencing environment. Manipulating other factors, such as the frequency of occurrence of expressions, and the temporal properties of facial expressions will be the focus of future work. This is difficult to achieve during live conversation, as *a priori* knowledge of the temporal properties of the facial expressions over the conversation is required.

4.1
Manipulating the intensity of facial gestures

Facial expressions can be represented in terms of the parameters of an AAM, which in turn represent displacements from the origin in the space spanned by the model. It follows that multiplying the parameters by a scalar greater than unity increases the distance from the mean, thereby exaggerating the facial expression represented by the parameters. Conversely, scaling by a value less than unity decreases the distance from the origin, thereby attenuating the intensity of the facial expression. Equations (2) and (3) can easily be extended to allow the intensity of facial expressions to be manipulated, as follows:

$$\mathbf{s} = \mathbf{s}_0 + \sum_{i=1}^{m} \mathbf{s}_i p_i \alpha, \tag{4}$$

$$A(\mathbf{x}) = A_0(\mathbf{x}) + \sum_{i=1}^{l} A_i(\mathbf{x}) \lambda_i \beta \quad \forall \mathbf{x} \in s_0, \tag{5}$$

To illustrate the effect of manipulating the intensity of facial expressions in this way, an AAM was constructed and used to track a face during a conversation in a video sequence. The tracking algorithm used is the *inverse compositional* algorithm described by Matthews and Baker (2004), which runs far in excess of video frame-rate (> 200 fps). The goal is to minimize the sum of squares difference between a model instance $A(\mathbf{x}) = A_0(\mathbf{x}) + \sum_{i=1}^{m} \lambda_i A_i(\mathbf{x})$ and the input image warped back onto the base mesh $I(\mathbf{W}(\mathbf{x};\mathbf{p}))$. Following (Gross, Matthews, & Baker, 2006), this is achieved as follows:

Pre-compute:

- Compute the gradient of base appearance $\nabla A_0(\mathbf{x})$.
- Compute the Jacobian of the warp at $(\mathbf{x};0), \dfrac{\partial \mathbf{W}}{\partial \mathbf{p}}$.
- Compute steepest descent images $S(\mathbf{x}) = \nabla A_0(\mathbf{x}) \dfrac{\partial \mathbf{W}}{\partial \mathbf{p}}$.
- Project out appearance variation $S_p(\mathbf{x}) = S(\mathbf{x}) - \sum_{i=1}^{m} \left[\sum_{x \in s_0} A_i(\mathbf{x}) S(\mathbf{x}) \right] A_i(\mathbf{x})$.

- Compute the Hessian $\mathbf{H} = \sum_{\mathbf{x}} \left[\nabla A_0 \frac{\partial \mathbf{W}}{\partial \mathbf{p}} \right]^T \left[\nabla A_0 \frac{\partial \mathbf{W}}{\partial \mathbf{p}} \right]$.

Iterate:

- Warp the image onto the base shape $I(\mathbf{W}(\mathbf{x}; \mathbf{p}))$.
- Compute the error image $E(\mathbf{x}) = I(\mathbf{W}(\mathbf{x}; \mathbf{p})) - A_0(\mathbf{x})$.
- Compute the warp update $\nabla \mathbf{p} = -H^{-1} \left[\nabla A_0 \frac{\partial \mathbf{W}}{\partial \mathbf{p}} \right]^T E(\mathbf{x})$.
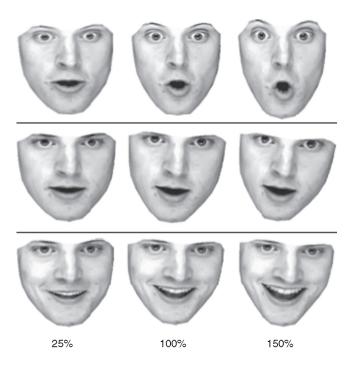- Update the warp $\mathbf{W}(\mathbf{x}; \mathbf{p}) \leftarrow \mathbf{W}(\mathbf{x}; \mathbf{p}) \circ \mathbf{W}(\mathbf{x}; \Delta\mathbf{p})^{-1}$.

The fitting algorithm iterates until convergence, that is, until the change in the parameters between iterations falls below some threshold. The key to the efficiency is the error is computed in the coordinate frame of model and not the image. Thus the Hessian matrix required to minimize the error function is constant and can be pre-computed. Also, since the appearance variation is *projected-out* from the steepest-descent images, the appearance parameters can be ignored during the fitting process. The AAM search is local and so requires a reasonable starting estimate for the parameters **p**. To track the face throughout a video sequence, the first frame can be manually hand-labeled, and then each subsequent frame can use the fit from the previous frame as the starting estimate. The parameters measured from the video can then be scaled using Equations (4) and (5), and the resultant facial expressions re-synthesized by applying the parameters to the model. Figure 3 shows example frames from a video sequence re-rendered using the AAM both with and without adjustment to the intensity of the facial expressions.

Varying degrees of exaggeration/attenuation can be applied to the different parameters of the AAM. However, to ensure the resultant facial expressions are valid, we generally apply the same degree of scaling to all parameters. The scaling parameter in Equations (4) and (5) is different to indicate the shape and appearance components of the model can be manipulated independently. To ensure any manipulated facial expression is valid, in the sense that it appears genuine, constraints are placed on the parameter values to ensure that they lie within chosen limits of the original training data. Typically the parameters are constrained to lie within ±3 standard deviations from the mean (Cootes et al., 2001). This prevents artifacts such as the mouth opening abnormally wide, or other violations of facial topology, for example the upper and lower lips intersecting when the mouth closes.

An advantage of using an AAM for this task is that it takes only in the order of 150 ms to capture a video frame, extract the AAM parameters automatically, apply the manipulation, re-render the face, and display the resultant image. Thus, AAMs allow face images to be manipulated during face-to-face interaction within a video-conferencing environment. Thus, two people can communicate with one another over the videoconference, but the face(s) displayed are re-rendered using the AAM, and not the image captured directly from the camera. The manipulation of the parameters can be applied prior to re-synthesis, thus both interlocutors are free to converse naturally and are kept blind to any manipulation of the facial expression. A further advantage

Figure 3

Facial expressions of varying intensity rendered using an AAM. Left column shows the expressions scaled to 25% intensity, the middle column shows the expressions as measured in original video, and the right column shows expressions exaggerated to 150% of the intensity measured in the video. The effect of scaling the parameters is much more pronounced in more extreme expressions

25%	100%	150%

is the scaling of the parameters allows facial expressions to be manipulated *precisely* (e.g., smile exactly half as much), which is impossible using an actor.

4.2
Copying facial expressions between faces

The previous section described how the *intensity* of facial expressions can be manipulated using AAMs. However, it might also be useful to present facial expressions *independently of identity*. This would be particularly beneficial in experiments contrasting the influence of social expectation and behavior on natural conversations. This *could* be achieved using an actor, whereby the actor behaves differently in different conversations, thereby effectively dissociating identity from behavior. However, the conversations are no longer natural, and the reliability is largely dependent on the skill of the actor.

Alternatively, this can be achieved by transferring facial expressions from images of the face of one person (the source) to images of the face of another (the target). This is achieved by measuring the difference between an image containing a reference facial expression (say, neutral) and an expressive image, then *imposing* this difference

onto an image of the same reference facial expression for the target identity. This difference can be measured in terms of shape and appearance information (Blanz, Basso, Poggio, & Vetter, 2003) or in terms of raw pixel intensity changes (Liu, Shan, & Zhang, 2001).

This form of facial expression mimicry fits naturally within our AAM framework. The parameters of an AAM are, by definition, offsets from the mean. Provided that the base shape and appearance for two models represent approximately the same facial expression, then swapping the mean shape and appearance from one AAM into another will effectively recreate the same facial expression (behavior) on a different face (identity). Equation (2) for reconstructing the shape can be easily modified to take this transfer into account as follows:

$$p_i = \mathbf{s}_i^T (\mathbf{s_0} - \mathbf{s}) \rightarrow \mathbf{t} = \mathbf{t_0} + \sum_{i=1}^{m} \mathbf{s}_i p_i \alpha \qquad (6)$$

where the left hand side of Equation (6) uses the source model to compute the parameters from a known shape, and the right-hand side applies the parameters to the source modes of variation and the target mean shape. The structure of the equation for transferring appearance is exactly the same. Thus facial expressions can be simultaneously copied to an image of the target face and manipulated in terms of the intensity. The computational cost of this mimicry is exactly as outlined in the previous sections, so again it can be used to manipulate live face-to-face conversations over a videoconference without either conversant being aware of the manipulation. Example facial expressions rendered using AAMs both before and after substituting the mean shape and appearance from a target model into a source model are shown in Figure 4, where the case for both within-gender and between-gender mimicry are shown.

The requirement that the mean shape and appearance for both the source and the target AAM represent similar facial expressions is generally not a problem. The mean shape and appearance are computed from a number of images displaying a range of facial expressions, so averaging over many examples usually results in the means converging to approximately the same facial expression, even for different identities and gender—see Figure 5.

4.3
Mapping facial expressions between faces

The method for copying facial expressions between faces described previously substitutes the mean (shape and appearance) from one AAM to another. The modes of variation (the basis shape and appearance vectors) of the source model still are used to generate the expression itself. Hence *the cloned expression appears as the source expression imposed onto the target identity.* See, for example, the teeth in Figure 4. Whilst this is useful for presenting facial expressions independently of identity, we might wish to *map* gestures from facial images of one person to facial images of another such that the characteristics of the target person are retained.

If the basis vectors between two AAM spaces could be interpreted to have the same meaning, say the first vector opens and closes the mouth, the second the eyes,

Figure 4

Dissociating expression and identity by *copying* changes in the features of the face that characterize expressions. Here expressions on (A), a reference face, are copied to other faces of (B) the same gender and (C) the opposite gender. As expressions are copied, the characteristics of the expression are the same on all faces. For example, the nasiolabial furrows, the shape of the smile, the mole on the cheek and the teeth are all characteristic of the person in (A)
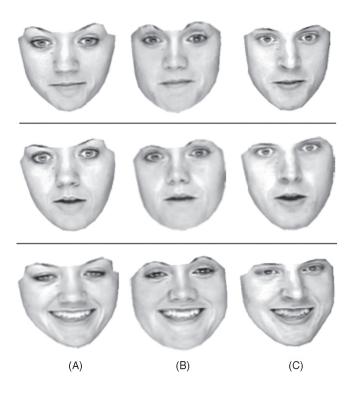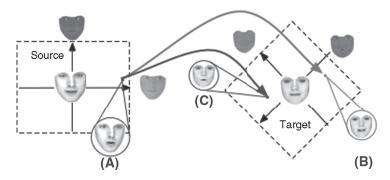

(A)   (B)   (C)

Figure 5

The base appearance of three AAMs each trained on 20 images of the respective individual. The mean appearances converge to approximately the same expression

(A)   (B)   (C)

Figure 6

Illustration of the mismatch in alignment between model-spaces. The parameters (A) in the source space representing the expression to be mapped can be copied directly to (B) the same location on the 2D plane. However, the resulting misalignment means that this point is not in the same location relative to the target basis vectors. Thus, the expression generated does not match the original. Instead, the parameters should be mapped by considering the rotation of the basis vectors between the model-spaces such that (A) is mapped to the point (C) in the same relative location between model-spaces

and so on, the entire model for the target person could be substituted for the source, rather than substituting only the mean shape/appearance. The source model could be used to analyze the video and determine the AAM parameters, which could then be applied directly to the target model. However, basis vectors computed using PCA simply maximize the projected variance of the input data—the source and target basis vectors do not correspond in any physically meaningful way. Simply applying the parameters directly from one model to another will produce incorrect facial expressions. The problem is illustrated in Figure 6.

The manner in which the features of the face can vary is highly constrained. Individual features generally do not move independently of one another (in a global sense), and features generally move in the same way between faces—the mouth can open and close, the lips can round, the eyes can blink, and so on. It follows therefore that the information contained within one AAM space should be embedded in some way within another. The question is: how does one map the meaning of a basis vector of one AAM to the space spanned by another? One approach to solving this problem is to capture many individuals displaying the same facial expressions, then use multi-linear models to separate expression from identity (and any other source of variation of interest). Expressions can then be predicted and synthesized on new faces (Macedo et al., 2006; Vlasic et al., 2005). Our approach to *mapping* between AAM spaces avoids the need for this complex, controlled, data capture.

The basis of our approach is not to think about the problem as mapping facial expressions between images of different people, but rather to consider it a geometric problem. A change in the facial features captured by a basis vector in one model-space ought to be captured by a combination of the basis vectors in another model-space (because the movement of the facial features is highly constrained, as stated). The basis

vectors themselves are computed when the model is trained and remain constant thereafter, so we can pre-compute the alignment of a basis vector from one model (the source model) with the basis vectors of another model (the target model) by computing the inner products

$$\langle \mathbf{s}_i, \mathbf{t}_j \rangle \; j = 1 \ldots m.$$

These inner products are then used to weight the parameter values computed using the source model so they correctly represent, as far as possible, the same deformation on the target face. Thus, Equation (2) rewritten in matrix form and extended to incorporate the mapping can be written as:

$$t = t_0 + T(R\mathbf{p}_s), \tag{7}$$

where the columns of **T** are the basis vectors spanning the *target* space, R is a $q \times r$ matrix of inner products (the target space is of dimension q and the source of dimension r), and p_s are the parameters representing the expression in the *source* space. Note: because **R** does not depend on expression and can be pre-computed, the cost of mapping an expression is only a matrix-vector product. Equation (3), which is used to generate the appearance images, takes exactly the same form as Equation (7). Cloning facial expressions in this way has the advantage that we do not explicitly need to define anything about the facial expressions. This information is implicit from the basis vectors in the source and target spaces. Also, we are not concerned with the direction of the eigenvectors. For example, an increasingly positive value for a source parameter might, say, open the mouth, while the same action could be defined by an increasingly negative value in the target space. In this instance the inner product for that combination of vectors is negative (the displacements from the mean are largely in opposite directions), so the sign of the parameter value is flipped when the parameter is mapped. The parameters are also constrained to lie within the chosen ±3 standard deviation limits of the *target* model, so the appearance of the resulting synthesized faces are valid given the original training data of the target person. This ensures the characteristics of the target face are preserved after mapping. The parameters can also be scaled after mapping to manipulate the intensity of the facial expression, as described previously.

Example images illustrating the mapping between various facial expressions are shown in Figure 7. Although there is some degradation in image quality, images can be generated at video frame-rate. If greater image quality is desired, multisegment modeling can be used to improve the rendering and reduce the effects of blurring (see Figure 8) at the additional cost of a precomputed look-up from pixels within a segment to pixels within the face image. Further effort can be applied to improve the quality of static images; however, this would increase the computational cost and rendering time. In a previous study (Theobald et al., 2003) we found that the image degradation such as that introduced by an AAM was not statistically significant in terms of the perceived naturalness of the *speech dynamics* in comparison with, say,

Figure 7

Dissociating expression and identity by *mapping* changes in the features of the face that characterize expressions. Here expressions on (A), a reference face, are mapped to another face (B) of the same gender and (C) of the opposite gender. As expressions are mapped the characteristics of the expression are, as far as possible, consistent with expression changes on the respective people. These are the same expressions as those shown in Figure 5. Notice the mole on the cheek no longer appears on the cloned faces, and the teeth for each person appear as their own. In Figure 5, the same teeth as are visible in (A) are copied into the inner mouth for each cloned person
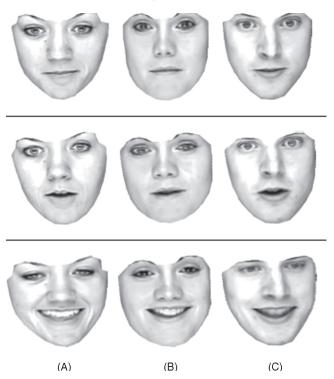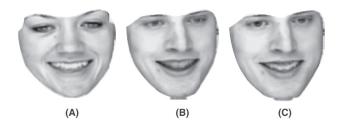

(A)   (B)   (C)

Figure 8

Multi-segment modelling and re-rendering. (A) An expression on a source confederate as measured in a video sequence and re-rendered on their model, and the expression mapped to a target face and rendered using (B) a standard AAM and (C) a multi-segment AAM. The multi-segment model reduces some of the blurring artefacts apparent in the standard AAM rendered image

(A)   (B)   (C)

temporally distorted video. The temporal properties of facial gestures are unaffected by the manipulation we have described in this article.

5 Summary and discussion

In this article we have described techniques for manipulating and transferring facial expressions from video sequences containing one face to video sequences containing other faces. AAMs are constructed for a source face and one or more target faces. The source face is tracked in a video sequence using the inverse compositional AAM fitting algorithm and the model parameters representing the facial expressions are mapped to the target face(s) and re-rendered, all in real-time at video frame-rate. The advantages of using AAMs are as follows: (1) The mapping is simple and intuitive and requires no manual specification of the type or degree of facial gestures. (2) The model can account for a high degree of variability in the images, which offers more flexibility than purely image-based approaches. (3) Near-videorealistic avatars for new subjects can be created without the cost of recording and processing a complete training corpus for each—visual speech and expression information can be synthesized for one face and transferred to another. (4) The model allows the mapped expressions to be constrained, so that they best match the appearance of the target producing the expression. (5) The model parameters can be manipulated easily to exaggerate or attenuate the mapped expression, which would be difficult using comparable image-based approaches.

The methods described here will be used in perceptual studies of dyadic conversation in a videoconferencing environment. The facial expressions of one conversant will be manipulated as described and the effect of these adjustments on behavior during real-time natural conversations will be investigated. The advantage of this approach is both interlocutors are blind to any manipulation, either in terms of the expressiveness of their actions, or the identity or apparent gender of the person they appear to be.

References

BLANZ, V., BASSO, C., POGGIO, T., & VETTER, T. (2003). Reanimating faces in images and video. *Computer Graphics Forum*, **22**(3), 641–650.

BLANZ, V., & VETTER, T. (1999). A morphable model for the synthesis of 3D faces. In W. Waggenspack (Ed.), *Proceedings of SIGGRAPH* (pp.187–194). Los Angeles, CA: ACM Press.

BLANZ, V., & VETTER, T. (2003). Face recognition based on fitting a 3D morphable model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **25**(9), 1063–1074.

BREGLER, C., & OMOHUNDRO, S. (1995). Nonlinear manifold learning for visual speech recognition. In E. Grimson (Ed.), *Proceedings of the International Conference on Computer Vision* (pp.494–499). Cambridge, MA: IEEE Computer Society Press.

CHANG, Y., & EZZAT, T. (2005). Transferable videorealistic speech animation. In *ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (pp.29–31).

CHOE, B., LEE, H., & KO, H. (2001). Performance-driven muscle-based facial Animation. *Journal of Visualization and Computer Animation*, **11**(2), 67–79.

CHUANG, E., & BREGLER, C. (2002). *Performance driven facial animation using blendshape interpolation*. Technical Report CS-TR-2002–02, Stanford University, U.S.A.

COOTES, T., EDWARDS, G., & TAYLOR. C. (2001). Active appearance models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **23**(6), 681–685.

COSKER, D., MARSHALL, D., ROSIN, S., PADDOCK, P., & RUSHTON, S. (2004). Towards perceptually realistic talking heads: Models, metrics, and McGurk. In V. Interrante & A. McNamara (Eds.), *Proceedings of Symposium on Applied Perception in Graphics and Visualization* (pp.2474–2477). Los Angeles, CA: ACM Press.

EDWARDS, G., COOTES, T., & TAYLOR, C. (1998). Face recognition using active appearance models. In H. Burkhardt & B. Neumann (Eds.), *Proceedings of the European Conference on Computer Vision* (pp.581–595). Freiburg: Springer.

FAGEL, S., BAILLY, G., & ELISEI, F. (2007). Intelligibility of natural and 3D-cloned German speech. In *Proceedings of Auditory Visual Speech Processing* (pp.126–131).

GEHRICKE, J., & SHAPIRO, D. (2000). Reduced facial expression and social context in major depression: Discrepancies between facial muscle activity and self-reported emotion. *Psychiatry Research*, **95**(2), 157–167.

GOVOKHINA, O., BAILLY, G., BRETIN, G., & BAGSHAW, P. (2006). TDA: A new trainable trajectory formation system for facial animation. In R. Stern, A. Black, & T. Schultz (Eds.), *Proceedings of Interspeech* (pp.2474–2477). Pittsburgh, PA: International Speech Communication Association.

GRAHE, J., & BERNIERI, F. (1999). The importance of nonverbal cues in judging rapport. *Journal of Nonverbal Behavior*, **23**(4), 253–269.

GROSS, R., MATTHEWS, I., & BAKER, S. (2006). Active appearance models with occlusion. *Image and Vision Computing*, **24**(6), 593–604.

GUENTER, B., GRIMM, C., WOOD, D., MALVAR, H., & PIGHIN. F. (1998). Making faces. In *Proceedings of SIGGRAPH* (pp.55–66).

HONG, P., WEN, Z., HUANG, T., & SHUM, H. (2002). Real-time speech-driven 3D animation. In *Proceedings of 3D Data Processing, Visualization and Transmission* (pp.713–716).

KIRBY, M., & SIROVICH, L. (1990). Application of the Karhunen-Loeve procedure for the characterization of human faces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **12**(1), 103–108.

KOUFAKIS, I., & BUXTON, B. (1999). Very low bit-rate face video compression using linear combination of 2D face views and principal components analysis. *Image and Vision Computing*, **17**, 1031–1051.

LIU, Z., SHAN, Y., & ZHANG, Z. (2001). Expression mapping with expressive ratio images. In L. Pockock (Ed.), *SIGGRAPH* (pp.271–276). Los Angeles, CA: ACM Press.

LUETTIN, J., THACKER, N., & BEET, S. (1996). Speechreading using shape and intensity information. In *Proceedings of the International Conference on Spoken Language Processing* (pp.58–61).

MACEDO, I., BRAZIL, E., & VELHO, L. (2006). Expression transfer between photographs through multilinear AAMs. *Brazilian Symposium on Computer Graphics and Image Processing* (pp.239–246).

MASSARO, D. (1998). *Perceiving talking faces*. Cambridge, MA: MIT Press.

MATTHEWS, I., & BAKER, S. (2004). Active Appearance Models revisited. *International Journal of Computer Vision*, **60**(20), 135–164.

MATTHEWS, I., COOTES, T., BANGHAM, A., HARVEY, R., & COX, S. (2002). Extraction of visual features for lipreading. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **24**(2), 198–213.

MORI, M. (1970). The uncanny valley. *Energy*, **7**(4), 33–35.

MOGGHADDAM, B., & PENTLAND, A. (1995). An automatic system for model-based coding of faces. In J. Storer & M. Cohn (Eds.), *Proceedings of the Data Compression Conferences* (pp.362–370). Snowbird, UT: IEEE Computer Society Press.

NOH, J., & NEUMANN, U. (2001). Expression cloning. In L. Pocock (Ed.), *SIGGRAPH* (pp.277–288). Los Angeles, CA: ACM Press.

SCHWARTZ, G., FAIR, P., SALT, P., MANDEL, M., & KLERMAN, G. (1976). Facial expression and imagery in depression: An electromyographic study. *Psychosomoatic Medicine*, **28**(5), 337–347.

SHACKLETON, M., & WELSH, W. (1991). Classification of facial features for recognition. In S. Negahdaripour (Ed.), *Proceedings of Computer Vision and Pattern Recognition* (pp.573–579). Lhanina, Maui: IEEE Computer Society Press.

SIROVICH, L., & KIRBY, M. (1987). Low-dimensional procedure for the characterization of human faces. *Journal of the Optical Society of America (A)*, **4**, 519–524.

TERZOPOULOS, D., & WATERS, K. (1993). Analysis and synthesis of facial image sequences using physical and anatomical models. *IEEE Transactions of Pattern Analysis and Machine Intelligence*, **15**(6), 569–579.

THEOBALD, B., BANGHAM, A., MATTHEWS, I., & CAWLEY, G. (2004). Near-videorealistic synthetic talking faces: Implementation and evaluation. *Speech Communication*, **44**, 127–140.

THEOBALD, B., KRUSE, S., BANGHAM, A., & CAWLEY G. (2003). Towards a low-bandwidth talking face using appearance models. *Image and Vision Computing*, **21**(13–14), 1117–1124.

THEOBALD, B., & WILKINSON, N. (2007). A real-time speech-driven talking head using appearance models. In *Proceedings of Auditory Visual Speech Processing* (pp.264–269).

TURK, M., & PENTLAND, A. (1991). Eigenfaces for recognition. *Journal of Cognitive Neuro-Science*, **3**(1), 71–86.

VLASIC, D., BRAND, M., PFISTER, H., & POPOVIC, J. (2005). Face transfer with multilinear models. *ACM Transactions on Graphics*, **24**(3), 426–433.

WANG, Y., LUCEY, S., & COHN, J. (2008). Enforcing convexity for improved alignment with constrained local model. In N. Ahuja & L. Shapiro (Eds.), *Proceedings of Computer Vision and Pattern Recognition* (pp.1–8). Anchorage, AK: IEEE Computer Society Press.

WILLIAMS, L. (1990). Performance driven facial animation. *Computer Graphics*, **24**(2), 235–242.

XIAO, J., BAKER, S., MATTHEWS, I., & KANADE, T. (2004). Real-time combined 2D+3D active appearance models. *In Proceedings of Computer Vision and Pattern Recognition* (pp.535–542).

YANG, J., ZHANG, D., FRANGI, A., & YANG, J. (2004). Two-dimensional PCA: A new approach to appearance-based face representation and recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **26**(1), 131–131.

ZHANG, Q., LIU, Z., QUO, G., TERZOPOULOS, D., & SHUM, H. (2006). Geometry-driven photorealistic facial expression synthesis. *IEEE Transactions on Visualization and Computer Graphics*, **12**(1), 48–60.