

Feature-Based Techniques for Real-Time Morphable Model Facial Image Analysis

S. Chaudhuri^a, R. K. Singh and E. Charbon.^b

^aIndian Institute of Technology, Kanpur 208016, India;

^bÉcole Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland

ABSTRACT

We present an algorithm to quickly analyse and compress facial images using a 2-dimensional morphable model. It runs in real-time on reasonable resources, and offers considerable opportunities for parallelization.

A morphable model associates a “shape vector” and a “texture vector” with each image of a sample set. The model is used to analyze a novel image by estimating the model parameters via an optimization procedure. The novel image is compressed by representing it by the set of best match parameters. For real-time performance, we separate the novel image into shape and texture components by computing correspondences between the novel image and a reference image, and match each component separately using eigenspace projection. This approach can be easily parallelized.

We improve the speed of algorithm by exploiting the fact that facial correspondence fields are smooth. By computing correspondences only at a number of “feature points” and using interpolation to approximate the dense fields, we drastically reduce the dimensionality of the vectors in the eigenspace, resulting in much smaller compression times. As an added benefit, this system reduces spurious correspondences, since weak features that may confuse the correspondence algorithm are not tracked.

Keywords: Morphable model, facial/face image, compression, features, real-time

1. INTRODUCTION

Portable and embedded devices typically need to communicate at high speed through low bandwidth channels. Frequently, images are transmitted or received by such devices, for example in a wireless security camera network or in a video phone. It is essential to compress these images to the smallest size possible, so that they can be transmitted quickly over slow connections. In such applications, lossless compression is not essential, since the display systems are usually of limited quality and size. It is however important to preserve strong visual features of the image, so that easy recognition is possible.

At the EPFL, Switzerland, the *MegaWatch* project aims to build a wristwatch-sized system with a range of functionalities, including image acquisition and transmission over very low bandwidth channels. One of the targetted applications of this feature is to capture and transmit images of faces – of the wearer and of other people. The system has considerable multi-processing power, so advanced techniques may be used for image compression and reconstruction. We present a real-time procedure designed to compress facial images on such a platform.

The method outlined in this report is extremely fast, reasonably robust and highly parallelizable. It takes advantage of the relative visual significances of image areas: strong features such as the mouth and the eyes are better preserved than flat features such as the cheeks and forehead. It has been tested on a large number of facial images and its various components are the subjects of much ongoing study, so we can expect to see considerable refinements and improvements in the near future.

Further author information: (Send correspondence to R.K.S.)

S.C.: E-mail: sidc@iitk.ac.in

R.K.S., E.C.: E-mail: {[randhirkumar.singh](mailto:randhirkumar.singh@epfl.ch), [edoardo.charbon](mailto:edoardo.charbon@epfl.ch)}@epfl.ch



Figure 1. Matching by pixel-wise overlay. On the left is the original image, on the right the best match. Images courtesy Vetter and Troje.⁸

2. BACKGROUND

Image compression may be divided into two categories: **informed** and **uninformed**. Uninformed compression assumes no knowledge of the objects represented by the image. Most general image compression standards such as GIF, PNG, JPEG and JPEG2000 fall in this category. The sequence of pixel intensities is compressed using lossless (which allow exact reconstruction) or lossy (which allow only approximate reconstruction) methods. The advantage of uninformed methods is that any image whatsoever may be compressed with adequate scope for reconstruction – they are therefore suitable for general image processing applications.

Informed algorithms deal with specific classes of images, for example facial images. Much greater compression is possible because the range of images is restricted. However, a setup to compress one class of images cannot be used to compress other classes – a facial image compression system will handle faces but not furniture. In many applications, such restriction is acceptable. For example, mugshots of criminals in police records are frontal or profile views of faces, all shot from the same viewpoint under similar lighting conditions.

Much effort in developing informed compression algorithms stemmed from research in face identification systems. The task was to identify a face from a photograph. It was assumed that another photograph of the same person was already in a large database of facial images. The new image was approximated as closely as possible by a linear combination of the images in the database. The person was identified by the database image that had the largest coefficient in this combination. This method worked reasonably well for identification, but the approximations were fuzzy and badly-defined because the faces were generally not of the same shape. Fig. 1 illustrates this.

To solve this problem, it was essential to somehow *normalize* the facial images so that they all had the same shape and could be overlaid accurately. **Morphable models** were developed to provide such a representation of image classes. In its conventional form, a morphable model associates a **shape** and a **texture** with each image in a sample set. The shape component (a correspondence function) describes how the sample image deviates from a reference image. The texture component encodes the colour/intensity information of the sample image, normalized to the reference shape. Fig. 2 shows an instance of this representation.

A novel image is approximated as a warped combination of the components in the morphable model. More precisely, for n sample images, if the shapes are represented by the set $\{S_i, i = 1 \dots n\}$ and the textures by the set $\{T_i, i = 1 \dots n\}$, then the novel image I_{novel} is approximated as:

$$I_{novel}(x, y) \approx T_{model} \circ S_{model}^{-1}(x, y)$$

where

$$T_{model} = \sum_{i=1}^n b_i T_i$$

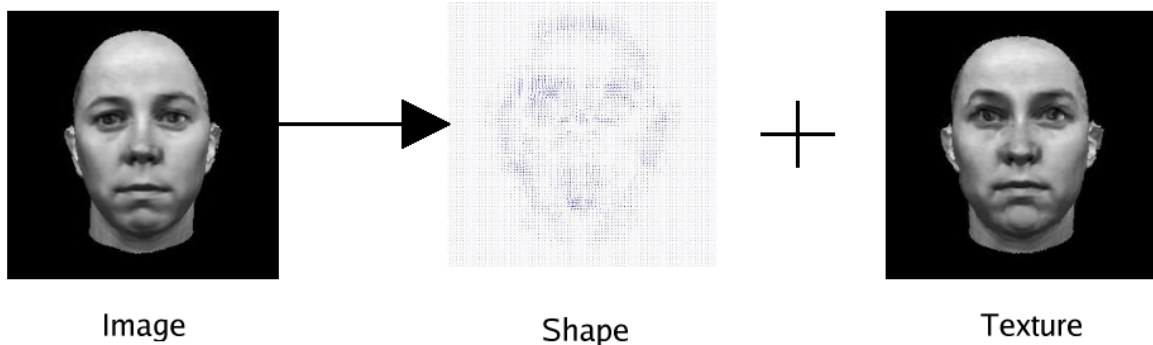


Figure 2. An image split into shape and texture components.

$$S_{model} = \sum_{i=1}^n c_i S_i$$

Intuitively, the novel texture is approximated by a linear combination of the sample textures, and is then warped to the approximate shape of the novel image by a linear combination of sample shapes. The scalars $\{b_i\}$ and $\{c_i\}$ are called the *best match coefficients*.

We have assumed here that our shape components S_i are defined in the direction of a “backward warp”. That is, $S_i(x, y)$ is the location of the point in the i th sample image that corresponds to the point (x, y) in the reference image. To obtain the novel image from the novel texture, we must apply a “forward warp”, so we use the inverse of the model shape S_{model} .

Assuming the approximation is good, drastic compression is achieved by representing the novel image only by the best match coefficients. It is easy to see that if the model is known, these coefficients are sufficient to reconstruct the closest match to the novel image. Also, the size of the compressed image depends only on the number of components in the database of sample images. A database of 50 sample images yields 50 shape and 50 texture components: if each coefficient is stored as a 4-byte floating-point number, then the compressed image is only $(50 + 50) \times 4 = 400$ bytes. (In our implementation, each shape component is represented as two independent components (x and y), so the actual size is 600 bytes.)

Vetter, Jones, Poggio, Beymer and others have done a considerable amount of work in developing two and three dimensional morphable models. Their work has mainly focussed on facial images. This is influenced by the fact that the morphable model description lends itself very well to faces: a dense and smooth correspondence field can be established between features on two different faces, something which is difficult to do with, say, buildings, in which external features present in one are frequently not present in another. Of course, we assume that the two faces have similar hair growth, accessories (such as spectacles), and no missing features (van Gogh, minus one ear, might be difficult to represent with the model).

In the approach of Jones, Poggio, Vetter et al,⁵ the mismatch between the morphable model and novel image is represented by a single sum-of-squared-errors function, which may be written as:

$$E(\mathbf{b}, \mathbf{c}) = \sum_{x,y} \left(I_{novel}(x, y) - T_{model} \circ S_{model}^{-1}(x, y) \right)^2$$

where \mathbf{b} represents the vector of b_i 's and \mathbf{c} represents the vector of c_i 's. A numerical minimization method such as stochastic gradient descent is used to minimize the error and thus obtain the best match coefficients \mathbf{b} and \mathbf{c} .

3. A REAL-TIME APPROACH

Our objective was to design a real-time system that would match a morphable two-dimensional face model to a novel image. We experimented with the stochastic gradient descent approach to minimize the matching error⁵

but ruled it out as it took too much time, partly because it is an iterative process that must loop many times before it converges to a solution and partly because the evaluation of the error (and its derivatives) at each step requires a large number of operations.

A faster solution was to first separate the novel image into shape and texture components as described by Vetter and Troje⁸ – the construction of the model requires that sample images be split in this way, so code could be reused. The components are then matched individually: the novel shape is approximated as a linear combination of sample shapes and the novel texture as a linear combination of sample textures.

The advantage of matching components separately is that the error functions are much simpler. They do not have the implicit warping induced by the $T_{model} \circ S_{model}^{-1}$ term. For example, the texture error in the new formulation is simply:

$$E_{texture}(\mathbf{b}) = \sum_{x,y} \left(T_{novel}(x,y) - T_{model}(x,y) \right)^2$$

where T_{novel} is the novel texture and T_{model} is $\sum_{i=1}^n b_i T_i$ as before.

3.1. Shape and Texture Representation

Following the standard practice, we represent both shapes and textures as vectors indexed by image coordinates. Each element of a vector corresponds to an image pixel. The texture vector stores intensity values, and the shape vector stores the relative coordinates of corresponding points (i.e. we store the values of the shape and texture functions at each image pixel). We note that each element of a shape vector has two components, in the x and y directions. In our implementation, we simplify this further with two separate vectors, one storing the correspondence values in the x direction and the other in the y direction. Our three error terms are, therefore:

$$\begin{aligned} E_{x,shape}(\mathbf{c}_x) &= \sum_{x,y} \left(S_{x,novel}[x,y] - \sum_{i=1}^n c_{x,i} S_{x,i}[x,y] \right)^2 \\ E_{y,shape}(\mathbf{c}_y) &= \sum_{x,y} \left(S_{y,novel}[x,y] - \sum_{i=1}^n c_{y,i} S_{y,i}[x,y] \right)^2 \\ E_{texture}(\mathbf{b}) &= \sum_{x,y} \left(T_{novel}[x,y] - \sum_{i=1}^n b_i T_i[x,y] \right)^2 \end{aligned}$$

where square brackets represent vector (or array) lookup instead of function evaluation. This representation facilitates a uniform treatment of shape and texture. The cost of the extra coefficients (\mathbf{c}_x and \mathbf{c}_y instead of just \mathbf{c}) is outweighed by the ease of implementation. It is possible to work with single shape vectors in which each element has an x and a y component, but the implementation is messier. We will assume separate x and y vectors for shape in the rest of this paper: the interested reader may modify the algorithms to work with single vectors.

3.2. Matching the Components

We observe that each of the errors above is a simple quadratic function of the appropriate coefficients. Unconstrained quadratic programming can be used to minimize each error. However, by preconditioning the sample set during model construction we can have much faster matching. The trick is to transform the set of sample vectors into an orthonormal basis. Adding up the projections of the novel vector onto the elements of the basis, we obtain the best match to the novel vector, measured in terms of the sum of squared errors (which is nothing but the squared L_2 distance between the vectors). The following theorem shows this:

THEOREM 3.1. *Let V be an n -dimensional subspace of \mathbb{R}^m , $n < m$. Let $\{\mathbf{B}_i\}$ be an orthonormal basis of V . Given a vector $\mathbf{X} \in \mathbb{R}^m$, the element of V closest to \mathbf{X} is*

$$\sum_{i=1}^n (\mathbf{X} \cdot \mathbf{B}_i) \mathbf{B}_i$$

where the “closest” element \mathbf{C} is defined as that which minimizes the squared L_2 distance $\|\mathbf{X} - \mathbf{C}\|^2$.

Each projection takes time linear in the size of the vectors. If each vector has m dimensions (which is, in our case, the number of pixels in each image), then the matching is performed in $O(mn)$ time.

3.3. Principal Component Analysis and Eigenvectors

In practice, the transformation to an orthonormal basis is not performed on the sample vectors themselves, but on the *deviations* of the sample vectors from their mean. This complication is introduced by **Principal Component Analysis** (PCA), which has the benefit of reducing the size of the sample space by discarding “less significant” elements from the basis. PCA transforms a set of linearly independent vectors $\{\mathbf{X}_i, i = 1 \dots n\}$ as follows: Let $\bar{\mathbf{X}}$ be the mean of the \mathbf{X}_i ’s. Consider the set of vectors $\{\mathbf{X}'_i \mid \mathbf{X}'_i = \mathbf{X}_i - \bar{\mathbf{X}}\}$. This set has rank $n - 1$, since $\sum_{i=1}^n (\mathbf{X}_i - \bar{\mathbf{X}}) = 0$. PCA gives an orthonormal basis $\{\mathbf{E}_i\}$ of size $n - 1$ for the space spanned by the set $\{\mathbf{X}'_i\}$. The \mathbf{E}_i ’s are conventionally called *eigenvectors*. PCA additionally reduces the size of the basis by discarding “less significant” eigenvectors, giving a final approximate basis of $p \leq n - 1$ elements.

After constructing the basis, we match a novel vector \mathbf{X}_{novel} as follows:

1. Compute $\mathbf{X}'_{novel} = \mathbf{X}_{novel} - \bar{\mathbf{X}}$
2. Project \mathbf{X}'_{novel} onto each eigenvector \mathbf{E}_i , obtaining scalars $e_i = \mathbf{E}_i \cdot \mathbf{X}_{novel}$. Note that all eigenvectors are of unit length, since the basis is orthonormal.
3. Approximate \mathbf{X}_{novel} as $\bar{\mathbf{X}} + \sum_{i=1}^p e_i \mathbf{E}_i$.

THEOREM 3.2. *Approximation by projection onto the full set of $n - 1$ eigenvectors is equivalent to minimizing the L_2 distance between a novel vector \mathbf{X}_{novel} and the vector space $V_1 = \{\sum_{i=1}^n b_i \mathbf{X}_i, \sum b_i = 1\}$.*

In our experiments, we observed that except in cases where lighting conditions, shape or overall intensity of the novel image deviated drastically from the images of the sample set, the constraint that the sum of the coefficients must be 1 was not very restrictive and the best matches were virtually identical to those without the constraint (Sect. 3.2). Further, by retaining only the first few principal components, we reduced the number of projections required, thus speeding up the matching and generating excellent approximations to the matches produced with the full eigenspace. Fig. 3 demonstrates this.

4. FEATURE-BASED OPTIMIZATIONS

The eigenvector-projection approach is fast, but has a significant associated challenge: the separation of the novel image into shape and texture components must be done as accurately as possible. This calls for a robust algorithm to set points in two facial images in correspondence.

The simplest approach is to compute the optical flow between the reference image and the novel image – the flow field is taken as the correspondence field. This approach is conceptually flawed, because optical flow algorithms are designed to track points on the *same* object across multiple frames, not locate corresponding points in images of *different* objects. However, since two faces are superficially similar, this approach works reasonably well in practice.

We experimented with the Bergen-Hingorani algorithm¹ for dense optical flow. It gave acceptable results on the whole, but spurious correspondences were frequently generated. Larger window sizes gave noticeably better results.

More sophisticated algorithms designed specifically for computing facial correspondences exist, but they are usually also more computation-intensive and not suitable for real-time processing.

To develop a suitable algorithm for our application which gave few spurious correspondences and ran in real-time, we took note of the following facts:

- Certain points in facial images are easier to track than others. Finding correspondences for a point at the corner of the mouth where contrast and texture are strong is, for instance, easier than for a point in the middle of the cheek, where the surface is smooth and unbroken.




Figure 3. Matches to the novel image (a) obtained by retaining b) all 49, c) 30, d) 15 and e) 7 eigenvectors.

- Correspondence fields between two facial images are smooth. This is to be expected since human faces have elastic, organic structure. This suggests that the field can be approximated fairly accurately by interpolation from a set of “defining points”.
- Human vision is sensitive to strong features. When we observe a face, recognition is triggered more by the intensity variation in sharply-defined regions such as the overall outline, the eyes, the nose and the mouth than in undistinguished ones such as the cheek and forehead. Therefore, in compressing a facial image, it is important to preserve the structure of strong features using accurate correspondences at these locations, but other regions need not be rendered so exactly.

A suitable correspondence algorithm could therefore select a set of easily trackable *feature points*, compute correspondences at these points only, and interpolate from the resulting values to obtain the correspondences at every pixel in the source image. By not tracking weakly-defined points (which resemble their neighbours), spurious correspondences may be reduced.

We implemented this method in the following manner:

Feature selection: We used a Harris-type corner detector⁴ to select well-defined feature points in the source image. The Harris detector is sensitive to strong gradients in the image. We ensured that the points corresponded to local maxima of the image gradient and were well-separated.³ Typically, about 300 feature points were selected. We were pleased to observe that these points typically coincided with dominant facial features.

Sparse correspondences: To calculate the correspondences at the feature points, we sampled a dense optical flow field. This field was generated using the Bergen-Hingorani algorithm.¹ Sparse sampling of a dense flow field has the advantage that the smoothing influence of neighbouring pixels on the value at each pixel is retained in the sparse set of flows. This helps prevent spurious correspondences.




Figure 4. Barycentric coordinates in an irregular convex polygon. Image courtesy Meyer et al.⁷

Interpolation: We triangulated the set of feature points using Delaunay triangulation. To compute the flow at a pixel \mathbf{p} , we determined which Delaunay facet f it lay in. The flow at \mathbf{p} was computed by interpolation from the vertices of f . We used generalized barycentric coordinates as interpolation coefficients: the coefficient w_j corresponding to the j th vertex was calculated as⁷:

$$w_j = \frac{\cot(\gamma_j) + \cot(\delta_j)}{\|\mathbf{p} - \mathbf{q}_j\|^2}$$

where the vertex \mathbf{q}_j and the angles γ_j and δ_j are as in Fig. 4. In practice, the coefficients at each pixel were computed during model construction and stored in a lookup table. The costs of triangulation, point location and barycentric coordinate computation were therefore not incurred during matching.

We developed this approach from original ideas, but discovered later that Kardouchi et al⁶ had done similar work. Fig. 5 illustrates the entire process.

4.1. Features for Speed

Although our feature-based approach was initially designed to generate good correspondences, we observed that it could also be used to drastically speed up the matching process. The shape components were completely defined by their values at the feature points. Hence, we could restrict the shape vectors to these values only, reducing their dimensionality from m , the total number of pixels in each image, to m_f , the number of feature points. In our experiments with 256×256 images and 300 feature points, this implied a reduction from 65536 dimensions to 300.




Figure 5. The feature-based interpolation process. For clarity, background triangles are not shown in the Delaunay triangulation. The interpolation map is colour-coded: each feature point is assigned an unique intensity and the intensity at each pixel is obtained by interpolation from the appropriate feature points.

We compute the PCA space of the restricted shape vectors and match by projection onto this space. The projection-based approach executes in time linear in the dimensionality of the vectors, so execution time is reduced from $O(mn)$ (see Section 3.2) to $O(m_f n)$. The full (dense) shape component is approximated *after* matching, by interpolating from the best match restricted shape.

We would like to show that matching with restricted vectors *first* and interpolating *afterwards* is equivalent to interpolating *first* and matching with full vectors *afterwards*. We must mention that in practice, very little difference was observed between the two methods.

4.1.1. Investigating Equivalence

We are given a set X of n vectors $\{\mathbf{X}_i, i = 1 \dots n\}$ in \mathbb{R}^m , $n \leq m$. We are also given a set Y of “restricted vectors” $\{\mathbf{Y}_i, i = 1 \dots n\}$ in \mathbb{R}^{m_f} , $m_f < m$. These sets have the following property: for any $j \in \{1, 2 \dots m\}$:

$$\begin{aligned} X_{1j} &= \sum_{k=1}^{m_f} c_{jk} Y_{1k} \\ X_{2j} &= \sum_{k=1}^{m_f} c_{jk} Y_{2k} \\ &\vdots \\ X_{nj} &= \sum_{k=1}^{m_f} c_{jk} Y_{nk} \end{aligned}$$

where X_{ij} denotes the j th element of \mathbf{X}_i and Y_{ij} denotes the j th element of \mathbf{Y}_i . The c_{jk} ’s are a set of scalar *interpolation coefficients* – there is one such set for each j . In simple terms, the set X may be obtained by interpolation from the set Y .

We are now given a novel vector $\mathbf{X}_{novel} \in \mathbb{R}^m$ and the corresponding restricted vector $\mathbf{Y}_{novel} \in \mathbb{R}^{m_f}$, related by the same sets of interpolation coefficients as above: for any $j \in \{1, 2 \dots m\}$:

$$X_{novel,j} = \sum_{k=1}^{m_f} c_{jk} Y_{novel,k}$$

We obtain two approximations to X_{novel} as follows:

1. Minimize the L_2 distance from the space spanned by the elements of X to \mathbf{X}_{novel} , obtaining the closest match \mathbf{X}_{approx}^1 .
2. Minimize the L_2 distance from the space spanned by the elements of Y to \mathbf{Y}_{novel} , obtaining the closest match \mathbf{Y}_{approx} . Now interpolate from \mathbf{Y}_{approx} using the coefficients $\{c_{jk}\}$ to obtain \mathbf{X}_{approx}^2 .

We ask if and under what conditions \mathbf{X}_{approx}^1 and \mathbf{X}_{approx}^2 are equal.

THEOREM 4.1. \mathbf{X}_{approx}^1 and \mathbf{X}_{approx}^2 are identically equal for all novel vectors if and only if for any orthonormal basis $\mathcal{A} = \{\mathbf{A}_1, \mathbf{A}_2 \dots \mathbf{A}_\alpha\}$ for X and any orthonormal basis $\mathcal{B} = \{\mathbf{B}_1, \mathbf{B}_2 \dots \mathbf{B}_\beta\}$ for Y ,

$$A^T A C = C B^T B$$

where

$$A = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \\ \vdots \\ \mathbf{A}_\alpha \end{bmatrix}, \quad B = \begin{bmatrix} \mathbf{B}_1 \\ \mathbf{B}_2 \\ \vdots \\ \mathbf{B}_\beta \end{bmatrix}, \quad C = [c_{jk}]_{m \times m_f}$$

Proof. Consider any two orthonormal bases \mathcal{A} for X and \mathcal{B} for Y . By Thm. 3.1, a closest match to \mathbf{X}_{novel} may be obtained as follows (as usual, V_{ab} or $V_{a,b}$ denotes the b th element of vector \mathbf{V}_a):

$$\begin{aligned}
\mathbf{X}_{approx}^1 &= \sum_{i=1}^{\alpha} (\mathbf{X}_{novel} \cdot \mathbf{A}_i) \mathbf{A}_i \\
\Rightarrow X_{approx,j}^1 &= \sum_{i=1}^{\alpha} (\mathbf{X}_{novel} \cdot \mathbf{A}_i) A_{ij} \\
&= \sum_{i=1}^{\alpha} \sum_{l=1}^m X_{novel,l} A_{il} A_{ij} \\
&= \sum_{i=1}^{\alpha} \sum_{l=1}^m \sum_{k=1}^{m_f} c_{lk} Y_{novel,k} A_{il} A_{ij} \\
&= \sum_{k=1}^{m_f} Y_{novel,k} \left(\sum_{l=1}^m \sum_{i=1}^{\alpha} A_{ij} A_{il} c_{lk} \right)
\end{aligned} \tag{1}$$

Also,

$$\mathbf{Y}_{approx} = \sum_{i=1}^{\beta} (\mathbf{Y}_{novel} \cdot \mathbf{B}_i) \mathbf{B}_i$$

Interpolating to obtain \mathbf{X}_{approx}^2 , we have

$$\begin{aligned}
X_{approx,j}^2 &= \sum_{l=1}^{m_f} c_{jl} \sum_{i=1}^{\beta} (\mathbf{Y}_{novel} \cdot \mathbf{B}_i) B_{il} \\
&= \sum_{l=1}^{m_f} c_{jl} \sum_{i=1}^{\beta} \sum_{k=1}^{m_f} Y_{novel,k} B_{ik} B_{il} \\
&= \sum_{k=1}^{m_f} Y_{novel,k} \left(\sum_{i=1}^{\beta} \sum_{l=1}^{m_f} c_{jl} B_{il} B_{ik} \right)
\end{aligned} \tag{2}$$

We observe that the right hand sides of Eqns. 1 and 2 are identically equal for all novel vectors and all $j \in \{1, 2, \dots, m\}$ if and only if the corresponding bracketed expressions are equal. But the bracketed expression in Eqn. 1 is $(A^T A C)_{jk}$ and the bracketed expression in Eqn. 2 is $(C B^T B)_{jk}$. Therefore $A^T A C = C B^T B$ is a necessary and sufficient condition for equivalence. \square

4.1.2. Theoretical Speed Gain

We conclude this section with some explicit calculations for the speed gain from the feature-based approach, compared to the case when full vectors are projected for all components. We assume that features are used to restrict only the two shape components (x and y), not the texture component. We also assume that the platform is serial-processing, not parallel-processing.

As mentioned previously, each shape component restricted to the feature points is projected onto the set of eigenvectors in $O(m_f n)$ time, compared to $O(mn)$ for projecting full vectors. If a sequence of one multiplication and one addition (a unit operation in projection) takes time τ , then we obtain a net speed gain in matching by a factor of g :

$$g = \frac{3mn\tau}{mn\tau + 2m_f n\tau} = \frac{3m}{m + 2m_f} = 3 - \frac{6m_f}{m + 2m_f}$$

For practical purposes, $m_f \ll m$, so $g \approx 3$: we expect the feature-based approach to be three times faster in projection. (Matching in PCA space requires us to add the mean, but this is linear in the dimensions of

the vectors and does not change g .) The results satisfy our expectations: as noted in Section 5, feature-based projection takes ~ 20 ms while full-vector projection takes ~ 60 ms. We exclude the (more or less identical) times taken to compute correspondences in the two approaches.

If the texture vectors could somehow be restricted as well, i.e. if we could predict an entire texture map from values at feature points, then the speedup is $g = m/m_f$, typically of the order of a few hundred.

5. IMPLEMENTATION AND RESULTS

We implemented our approach on a dual 2.7 GHz Pentium 4 system with 1GB RAM, running RedHat Linux 9 with the symmetric multi-processing (SMP) kernel 2.4.20-8smp. The coding language was C. OpenCV-0.9.5 was used as the image-processing library. No explicit multi-processing or threading instructions were included in the source code.

In our experiments, we obtained a minimum average matching time of 120ms. A slightly slower implementation which was easier to time took 140ms, of which 120ms was for computing sparse correspondences (pruning a Bergen-Hingorani optical flow field) and 20ms for projecting feature-based restricted vectors. We quote the times from the latter implementation since we are more confident that the timing was accurate.

We present the execution times for various approaches in tabular form (all measurements on the above platform).

Method	Approximate average matching time (ms)
Iterative error-minimization (Conjugate Gradient Descent)	3×10^4
Full-image projection with Black ² correspondences	5×10^3
Full-image projection with Bergen-Hingorani ¹ correspondences	180 (Correspondence: 120 + Projection: 60)
Feature-based projection with pruned Bergen-Hingorani ¹ correspondences	140 (Correspondence: 120 + Projection: 20)

Reconstruction typically took 40-80ms.

The reconstructed images in general closely resembled the original novel images. However, there were some spurious deformations in the results that are a cause for concern. It seems a better feature-tracking / correspondence / interpolation scheme is required. Incidentally, we note that interpolation with barycentric coordinates only ensures C^0 continuity, i.e. the interpolation surface is continuous but not necessarily smooth across edges of the subdivision. Spline-based interpolation schemes give C^1 (continuous first derivatives) and higher continuity, and are more suitable.

We present a selection of matches obtained, both good and bad, in Fig. 6. 50 images were present in the sample set. 40 principal components (out of a possible maximum of 49) were used for both shape and texture. The novel images were rendered from textured 3D models at the Max Planck Institute, Germany.

REFERENCES

1. J. Bergen and R. Hingorani, "Hierarchical Motion-Based Frame Rate Conversion," Technical Report, David Sarnoff Research Center, Princeton, 1990.
2. M. J. Black, "Robust Incremental Optical Flow," Ph.D. Thesis, Yale, 1992.
3. J.-Y. Bouget, "Pyramidal Implementation of the Lucas Kanade Feature Tracker: Description of the Algorithm," Microprocessor Research Labs, Intel Corp., 2000.

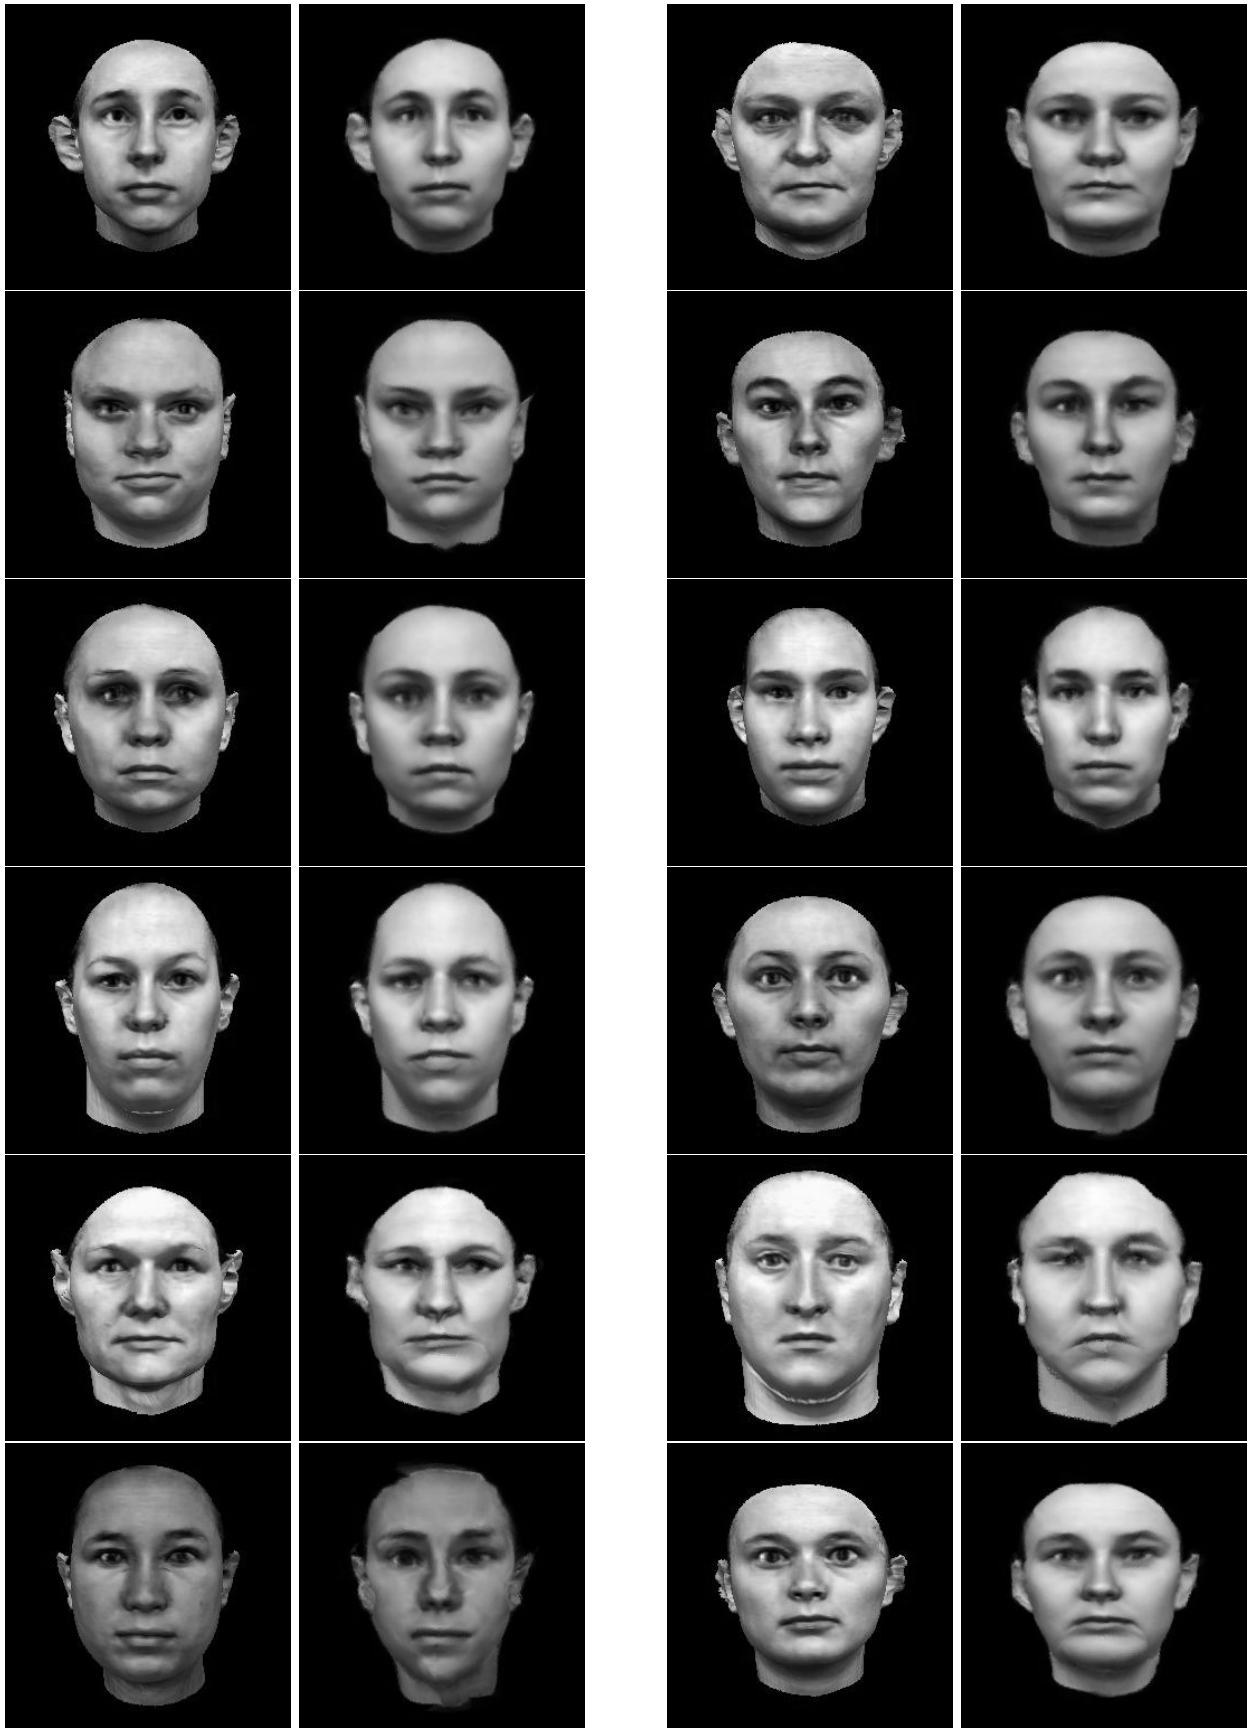


Figure 6. A selection of novel images (first and third columns) and the corresponding feature-based matches (second and fourth columns): some good, some not-so-good.

4. C. Harris and M. Stephens, "A combined corner and edge detector," *Proc. Alvey Vision Conf.*, pp.147-151, 1988.
5. M. J. Jones and T. Poggio, "Model-Based Matching by Linear Combinations of Prototypes," unpublished AI memo, MIT, 1996.
6. M. Kardouchi, J. Konrad and C. Vázquez, "Estimation of large-amplitude motion and disparity fields: Application to intermediate view reconstruction," *Proc. Visual Communications and Image Processing, IS&T/SPIE Symp. on Elec. Imaging*, San Jose, 2000.
7. M. Meyer, Haeyoung L., A. Barr, M. Desbrun, "Generalized Barycentric Coordinates on Irregular Polygons," *Journal of Graphics Tools*, 7(1):pp.13-22, 2002.
8. T. Vetter and N. F. Troje, "Separation of texture and shape in images of faces for image coding and synthesis," *JOSA, A* 14:9 pp.2152-2161, 1997.