Face Recognition Based on Frontal Views Generated from Non-Frontal Images

Volker Blanz[1], Patrick Grother[2], P. Jonathon Phillips[2] and Thomas Vetter[3]

[1]Max-Planck-Institut für Informatik, Saarbrücken, Germany
[2]National Institute of Standards and Technology NIST, Gaithersburg, MD, USA
[3]University of Basel, Departement Informatik, Basel, Switzerland
blanz@mpi-sb.mpg.de, pgrother@nist.gov, jonathon@nist.gov, thomas.vetter@unibas.ch

Abstract

This paper presents a method for face recognition across large changes in viewpoint. Our method is based on a Morphable Model of 3D faces that represents face-specific information extracted from a dataset of 3D scans.

For non-frontal face recognition in 2D still images, the Morphable Model can be incorporated in two different approaches: In the first, it serves as a preprocessing step by estimating the 3D shape of novel faces from the non-frontal input images, and generating frontal views of the reconstructed faces at a standard illumination using 3D computer graphics. The transformed images are then fed into state-of-the-art face recognition systems that are optimized for frontal views. This method was shown to be extremely effective in the Face Recognition Vendor Test FRVT 2002.

In the process of estimating the 3D shape of a face from an image, a set of model coefficients are estimated. In the second method, face recognition is performed directly from these coefficients. In this paper we explain the algorithm used to preprocess the images in FRVT 2002, present additional FRVT 2002 results, and compare these results to recognition from the model coefficients.

1. Introduction

Most state-of-the-art face recognition systems are optimized for frontal views of faces only, and their performance drops significantly if the faces in the input images are shown from non-frontal viewpoints. Changes in viewpoint, however, are encountered in many real-world applications, and face recognition from non-frontal viewpoints is one of the main challenges in developing general-purpose face recognition systems.

This paper demonstrates the viability of a preprocessing approach based on a Morphable Model of 3D faces [4, 5] that corrects for changes in viewpoints between facial images. Morphable Models have been shown to provide a technique for view transformation [4, 3] and for automatic face recognition [5].

Our approach reconstructs 3D models of faces from non-frontal views. Computer Graphics techniques are used to redraw a frontal facial image. The computer-generated frontal view can then serve as input to any view-based face recognition system for person identification or verification. The 3D reconstruction of the face includes the ears and the neck, but not the hair and the shoulders of the person. For a view transformation technique to be compatible with most face recognition systems, the transformation technique must be able to place a transformed face into a "full portrait." Our algorithm automatically inserts the face into a standard image. Moreover, it automatically estimates regions of the face that are occluded in the non-frontal view, and compensates for differences in illumination between images. Along with the faces' 3D shapes, our reconstruction algorithm estimates all relevant scene parameters, such as head position, size and orientation, illumination parameters and contrast (Figure 1). We can, therefore, automatically compute the rendering parameters from the standard image and apply them for rendering standard views of novel faces (Fig. 2).

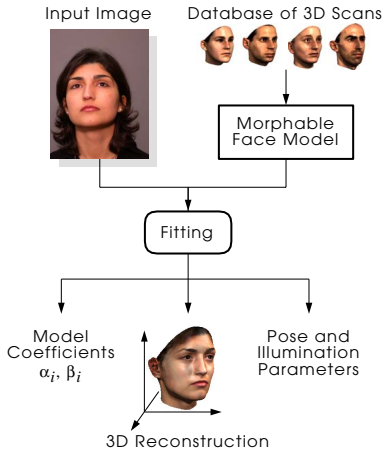

Figure 1. Fitting the Morphable Model to an Image produces not only a 3D reconstruction, but also model coefficients α_i, β_i and an estimate of head orientation, position and illumination.

The core of the 3D face reconstruction algorithm is a Morphable Model of 3D faces [4, 5] that captures general information about the natural variation of 3D shapes and textures of faces in a vector space spanned by a dataset of

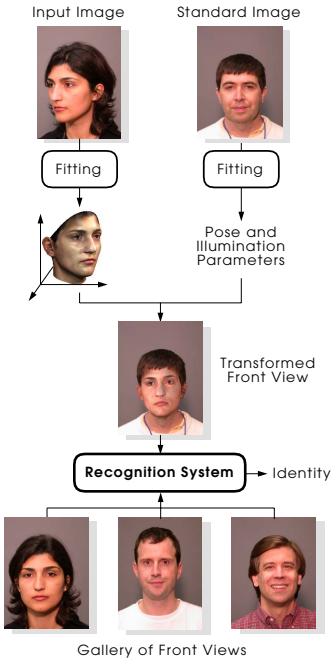

Figure 2. Viewpoint-transformed recognition: From a probe image (top left), our algorithm generates a transformed front view. This is input to a view-based face recognition system for comparison with the set of frontal gallery views.

3D faces. This information helps to solve the otherwise ill-posed problem of reconstructing 3D shape from a single image. Using an iterative optimization technique, our algorithm finds the linear combination of example faces such that 3D computer graphics renders an image that is as similar as possible to the input image. For initialization, the algorithm requires the image coordinates of up to seven feature points.

Figures 2 and 3 show two alternative paradigms for face recognition with Morphable Models. In both paradigms, we assume that *gallery images* of all persons known to the system are stored, and that persons in *probe images* have to be identified by finding the most similar person in the gallery, or their claimed identity has to be verified. We address the case of frontal gallery and non-frontal probe images, and vice versa.

In the *viewpoint-transformed recognition* approach proposed in this paper, all non-frontal images are transformed into front views, and recognition is performed by a separate, view-based algorithm. In one scenario, we assume that the gallery images are frontal (e.g. mug-shots), so only the *probe images have to be transformed.* (Figure 2). This has been evaluated in the Face Recognition Vendor Test 2002 [9] with 10 different face recognition systems by commercial firms, and we present a more comprehensive report of this evaluation and additional results. We are also evaluating a scenario where the probe views are frontal, for example from a security checkpoint, and the *gallery images*

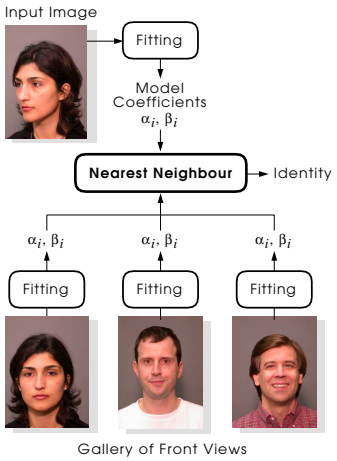

Figure 3. Coefficient-based recognition: The representation of faces in terms of model coefficients α_i, β_i for 3D shape and texture is independent of viewpoint. For recognition, all probe and gallery images are processed by the model fitting algorithm.

are transformed from given snapshots of suspects. This scenario may have even more practical relevance, since the viewpoint transformation has to be done only for the relatively small number of gallery images, and the recognition can then be done by efficient view-based methods.

Coefficient-based recognition (Figure 3) uses the model coefficients, which were computed by the model fitting algorithm, as a viewpoint-invariant, low-dimensional representation of faces [5]. Here, both the probe and the gallery images can be from any viewpoint. The drawback of this method is that each image has to be processed by the costly model fitting algorithm. To provide a control for measuring future impact of morphable models on automatic face recognition, we report performance of a coefficient-based recognition algorithm, using the same morphable model and the same fitting results as the viewpoint transformation algorithm. The model is trained and tuned on a different set of images than the test images. This allows for a meaningful comparison with the caveat that the coefficient algorithm was not tested on sequestered data.

1.1. Related Work

In an image-based approach, Beymer and Poggio have applied warping operations to generate virtual views of faces for recognition [2]. Synthesis of novel images based on two image-based linear object classes, which were constructed from images of the same set of individuals, was used for face recognition by Vetter [12]. Atick et al. presented a method for model-based shape-from-shading, using untextured 3D scans without point-to-point correspondence, and they suggested to create images at new poses for face recognition [1]. Virtual views, generated with image warping and shape–from–shading, have also been proposed for face recognition by Zhao and Chellappa [14]. Georghi-

ades et al. recovered 3D shape of faces from 3 images at a fixed pose and different lighting, and synthesized novel views that they used for face recognition [6]. For face recognition across different illuminations, but at a strictly fixed pose, Sim and Kanade presented a method that recovers surface normals with shape-from-shading, generates images at new illuminations, and trains a classifier on the extended set [11].

Due to recent advances in 3D shape reconstruction [5], we can now explore a combination of high-resolution reconstructions of textured 3D faces from single images, and a variety of state-of-the-art systems for image-based recognition. It is important to note that our approach does not require 3D scans of the faces that are to be recognized (gallery or probe), but only single 2D images. Just as with our coefficient-based algorithm [5], we use the database of 3D faces only for learning general properties of human faces. None of the individuals in the 3D database is in the image database that we used for testing.

In the following two sections, we briefly summarize the Morphable Model and the algorithm for 3D shape reconstruction. Section 3 describes our procedure for rendering standardized frontal views. Section 4 gives an extensive evaluation of the combined approach with a set of different view-based algorithms, and compares this combination with coefficient-based recognition.

2. Morphable Models for Face Reconstruction

The Morphable Model of 3D faces[13, 4, 5] is a vector space of 3D shapes and surface reflectances (textures) that is spanned by a dataset of examples and that captures the variations found within this set. Our dataset contains 200 textured *Cyberware* (TM) laser scans of an equal number of males and females aged between 18 and 45 years. Except for one Asian female, all persons are Caucasian. Previous work [5] indicates that the model may well be applied to reconstruct 3D shape from images of a wider ethnic variety.

The shape and texture vectors are defined such that any linear combination of examples

$$\mathbf{S} = \sum_{i=1}^{m} a_i \mathbf{S}_i, \quad \mathbf{T} = \sum_{i=1}^{m} b_i \mathbf{T}_i. \tag{1}$$

is a realistic face, given that $\mathbf{S}$, $\mathbf{T}$ are within a few standard deviations from their averages. Each vector $\mathbf{S}_i$ stores the 3D shape in terms of x, y, z-coordinates of all vertices $k \in \{1, \ldots, n\}$ of a high-resolution 3D mesh, and textures $\mathbf{T}_i$ contain their red, green and blue color values:

$$\mathbf{S}_i = (x_1, y_1, z_1, x_2, \ldots, x_n, y_n, z_n)^T \tag{2}$$

$$\mathbf{T}_i = (R_1, G_1, B_1, R_2, \ldots, R_n, G_n, B_n)^T. \tag{3}$$

In the conversion of the laser scans into shape and texture vectors $\mathbf{S}_i$, $\mathbf{T}_i$, it is essential to establish dense point-to-point correspondence of all scans with a reference scan, to make sure that vector dimensions in $\mathbf{S}$, $\mathbf{T}$ describe the same point, such as the tip of the nose, in all faces. Dense correspondence is computed automatically with an algorithm derived from optical flow (for details, see [5]). Finally, we perform a Principal Component Analysis to estimate the probability distribution of faces around the averages $\overline{\mathbf{s}}$ and $\overline{\mathbf{t}}$ of shape and texture, and we replace the basis vectors $\mathbf{S}_i$, $\mathbf{T}_i$ in Equation (1) by orthogonal eigenvectors $\mathbf{s}_i$, $\mathbf{t}_i$:

$$\mathbf{S} = \overline{\mathbf{s}} + \sum_{i=1}^{m-1} \alpha_i \cdot \mathbf{s}_i, \qquad \mathbf{T} = \overline{\mathbf{t}} + \sum_{i=1}^{m-1} \beta_i \cdot \mathbf{t}_i \tag{4}$$

3D shape reconstruction from a single input image is achieved by fitting the Morphable Model to the image in an analysis-by-synthesis loop: At each iteration, the current model parameters define a 3D face, and computer graphics can be used to render a coloured model image with red, green and blue channels $I_{r,model}(x, y)$, $I_{g,model}(x, y)$, $I_{b,model}(x, y)$. The fitting algorithm minimizes the difference between the model image and the input image,

$$E_I = \sum_x \sum_y \sum_{c \in \{r,g,b\}} (I_{c,input}(x, y) - I_{c,model}(x, y))^2, \tag{5}$$

in a stochastic Newton optimization [5] with respect to the following model parameters: shape coefficients, texture coefficients, 3D position, 3D orientation, focal length, red, green and blue components of ambient and parallel light, direction of that parallel light, color offsets, gains and color contrast. For details, see [5]. Due to the explicit, separate parameters for pose and illumination recovered from the image, we can modify any of them independently and draw the face from any new angle and under any new illumination.

For convergence of the optimization, the system is initialized by providing 2D image positions of some feature points, such as the tip of the nose or the corners of the eyes. The algorithm then converges automatically in about 4.5 minutes on a 2GHz Pentium 4 processor.

Since the linear combination of textures $\mathbf{T}_i$ cannot reproduce all local characteristics of the novel face, such as moles or scars, which in fact may be highly relevant for recognition, we extract the person's true texture from the image wherever it is visible. This is done by an illumination-corrected texture extraction algorithm [4]. The boundary between the extracted texture and the predicted regions (where the true texture is occluded in the input image) was still visible in some of the images used in the figures and evaluations reported in this paper. We have recently improved our algorithm to achieve a smooth transition, based on a reliability criterion for texture extraction.

In the database of images used for evaluation, some individuals had facial hair or wore eyeglasses, unlike the 200

Figure 4. From the original images (top row), we recover 3D shape (second row), by finding an optimal linear combination of example faces, and redraw the faces in frontal pose into a standard background (bottom row). In Section 4.1, the original front view (top, center) was used in the gallery, and the four non-frontal originals (top) and all five transformed views (bottom) were in the nine probe sets, respectively. The frontal–to–frontal mapping served as a baseline test (bottom row, center).

3D faces in the Morphable Model. Texture extraction captures the beard in a texture map and reproduces it on the frontal view, even if the beard's thickness is neglected in this transformation. In the same way, eyeglasses are mapped on the facial surface by the algorithm, which distorts the shape of the glasses in rotated viewpoints. The estimate of face shape, however, is not affected much by hair or glasses (see [5] for examples).

3. Hair and Portrait Processing

The goal of viewpoint transformation is to render frontviews images that are optimal for the subsequent, view-based recognition system in terms of imaging conditions. Rendering only the part of the head covered by our model, i.e. face, ears and neck, in front of a uniform background might affect algorithms that expect complete face images. Our approach, therefore, is to draw the rotated face into a standard portrait of a short-haired person at a frontal pose (Figure 4), so all transformed images have the same hair–style, shoulders, and background. With additional classification of gender and skin-complexion, which may be based on the 3D face reconstruction, it would be possible to select from a choice of different standard images automatically.

In order to determine head size, position, orientation and illumination of the standard image, we fitted the Morphable Model to the image, as described in the previous section. To be able to draw the standard portrait's hair in front of the forehead, we manually defined a transparency map that is opaque on the hair and transparent everywhere else [3]. If the transformed face is smaller than the face in the standard portraits, part of the original face would be visible in the background. We apply a simple background continuation method [3] to extend the background pattern into the background face region along the facial contour.

Given an input image, we perform the following steps:

1. Manually define feature points, such as the tip of the nose, the corners of the eyes, or any points along the occluding contours on the cheeks. For the results presented in this paper, we defined an average number of 11 feature points to ensure optimum quality. However, 6 points are often sufficient. In future systems, these points can be found by automated feature detection.
2. Run the optimization algorithm to estimate the 3D shape of the face.
3. Render the 3D face in front of the standard image, using the parameters for position, orientation, size, and illumination of the standard image.
4. Draw the hair in front of the forehead, using the transparency map which has been defined once on the standard image.

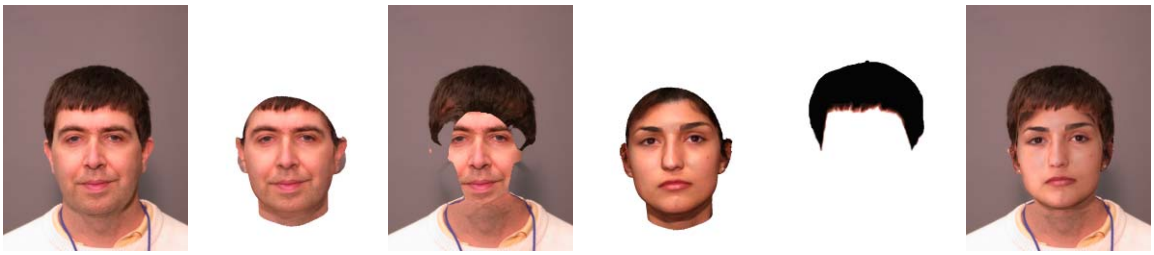

Standard Portrait Reconstruction Background Novel Face Hair Layer Result

Figure 5. Compositing. From the standard image (left), we reconstruct the face in order to obtain estimates of the head's pose and illumination. Along the edge of the reconstructed head, we extend the background colors into the face region. Then, we render the novel face with the standard imaging parameters and composite the hair layer in the foreground to obtain the final image (right).

4. Results

The data in the FRVT 2002 MCInt morphable model experiment was designed to examine the effects of pose variation and morphable models on performance [9]. The data consisted of five images of 87 subjects who were not in the dataset of 3D scans used for the morphable model. All images were taken taken indoors under studio conditions with a single incandescent photo floodlight, and all images of a subject were taken within five minutes. The five images of each subject consist of a frontal view and images with subjects looking left and right 45 degrees and up and down 30 degrees. The pose angles of the faces are nominal because pose was controlled by having subjects look at an object.

4.1. View Transformation of Probe Images

For the FRVT 2002 MCInt morphable model experiment the gallery consisted of the frontal image of the 87 subjects. Nine probe sets were constructed to measure the impact of pose variation and morphable models on performance. Figure 4 shows examples of the nine probe sets for one individual. The *45 left* and *45 right* probe sets contained facial images facing 45 degrees to the left and right of center respectively. The 45 L and 45 R columns in Figures 6 and 7 report verification and identification results for the 45 left and 45 right probe sets. Line segments are drawn between original and corresponding transformed probe sets for visualizing the effect of transformation. The *30 up* and *30 down* probe sets contain facial images facing 30 degrees up and down respectively. The performance results for these two probe sets are reported in the 30 U and 30 L columns. In the remaining five probe sets, a three-dimensional morphable model has been applied to the probes according to the paradigm shown in Figure 2.

The *frontal morph* probe set provides a baseline for how replacing the hairstyle and background affects a system if the viewpoint transformation is close to 0 degrees: The probe views are transformed versions of the original front views that form the gallery. In Figures 6 and 7, the results for the frontal morph probe set are in column frontal (morph). If a system were insensitive to the artifacts introduced by the morphable model, then the verification and identification rates would be 1.0. In Figure 6, sensitivity to morphable models range from 0.98 down to 0.45.

To investigate the effects of morphable models, performance was computed for four probe sets: *45 left morphed*, *45 right morphed*, *30 up morphed*, and *30 down morphed*. These probe sets were produced by applying the morphable model to the 45 left, 45 right, 30 up, and 30 down probe sets respectively. The results for the morphed probe sets are in columns 45 L (morph), 45 R (morph), 30 U (morph), and 30 D (morph). The results show that with the exception of Iconquest, morphable models significantly improved performance.

4.2. View Transformation of Gallery Images

We now examine the scenario of using one or multiple non-frontal images in the gallery and a single frontal image as the probe. This situation is representative of applications where non-frontal surveillance images are compared with standardized probe images such as those captured in an airport lane.

Table 1 lists the verification results for viewpoint-transformed gallery against un-transformed frontal probe images with one of the leading algorithms from the previous tests. The correct acceptance rate is in the same range as in the previous scenario. Table 2 summarizes verification results for all combinations of probe and gallery views. In the pairs of numbers listed, the first refers to the condition with both gallery and probe images unchanged, and the second to both gallery and probe images view-transformed. The data show a dramatic increase in correct acceptance rate in all conditions.

An additional question is whether, together, four non-frontal images offer superior performance than any one alone. We perform post-matching score-level fusion on the results obtained with each probe view. This involves replacing the four scores that result from the comparison of an unknown probe with the four non-frontals with a single

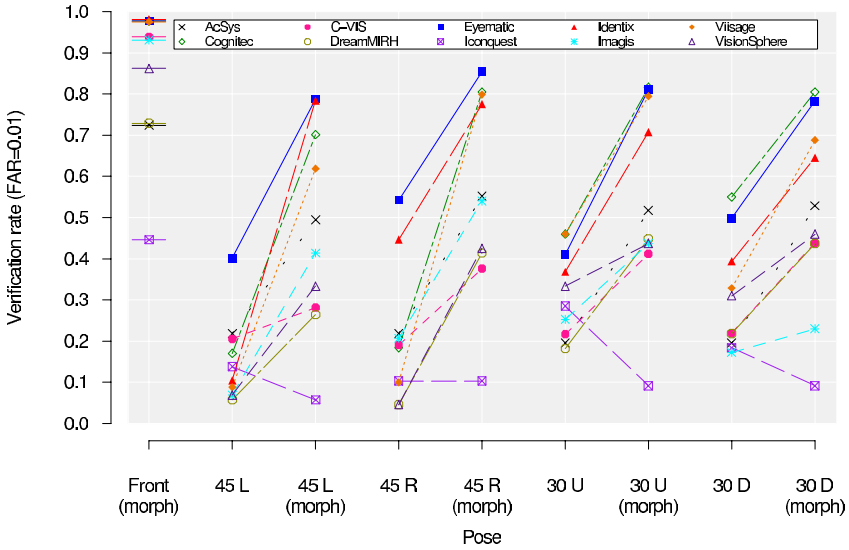

Figure 6. Verification rate for 10 different face recognition systems from FRVT02 at a 1% false accept rate. The diagram shows results for original probe images taken from different viewpoints (45L, 45R, 30U, 30D) and viewpoint-transformed probe images ("morph", see Section 4.1). Gallery images were untransformed frontal views.

View	Without View Transform	With View Transform
30 deg up view	0.45	0.81
30 deg down view	0.54	0.80
45 deg left view	0.16	0.70
45 deg right view	0.18	0.80
fused views (sum)	0.44	0.88
fused views (max)	0.55	0.90

Table 1. Verification performance for non-frontal gallery images with and without viewpoint-transformation, tested with untransformed, frontal probe images. The bottom rows are obtained with fusion of outputs from multiple views, according to two different schemes. The figures give the rate of correct acceptance at 1% false accept rate.

score. The standard approach [8, 7] is to sum the scores, but the use of the maximum, which corresponds to the assignment of the most similar score, is equivalent to regarding each of the four views as separate gallery identities. Note that the option of fitting the Morphable Model to multiple images simultaneously [4] is not used here.

The results shown in Table 1 form a lower bound on the benefits attainable with more sophisticated fusion schemes. The table suggests two conclusions. First that the viewpoint transformation is effective, especially in conjunction with even an elementary fusion scheme: The last number in the table, max-rule fusion applied to viewpoint-transformed images, is a 47% reduction in error (the correct acceptance rate increased to 0.90 from 0.81) over the best single transformed view. Secondly neither fusion rule is particularly effective on untransformed views.

4.3. Viewpoint Transformation versus Coefficient Based Recognition

From fitting the model's 99 most relevant principal components to images, we obtain coefficients α_i and β_i, $i = 1, ..., 99$ for each image. The algorithm does not only fit the entire face model, but also separate regions around the eyes, nose, mouth and the surrounding part, which yields four more sets of model coefficients. We scale all coefficients by their standard deviations according to PCA, and concatenate them to a vector $\mathbf{c}$. This vector can be used for identification using a nearest neighbour search. For verification, we compare the difference with a threshold value that is set such that false accept rate is at 0.01. In both recognition tasks, we use as a measure of similarity [5]

$$d_W = \frac{\langle \mathbf{c}_1, \mathbf{c}_2 \rangle_W}{\|\mathbf{c}_1\|_W \cdot \|\mathbf{c}_2\|_W} \tag{6}$$

between the coefficient vectors $\mathbf{c}_1$, $\mathbf{c}_2$ in a scalar product $\langle \mathbf{c}_1, \mathbf{c}_2 \rangle_W = \langle \mathbf{c}_1, \mathbf{C}_W^{-1} \mathbf{c}_2 \rangle$. $\mathbf{C}_W$ is the covariance matrix of the coefficients obtained when separately fitting the model to multiple images of the same persons, and it captures the within-class variation of $\mathbf{c}$ [5]. We use a matrix $\mathbf{C}_W$ obtained from a portion of the FERET database [10]. None of the individuals of the FERET database are in the gallery or probe sets used for our performance measurements.

We have evaluated the coefficient-based approach with the original images and the recognition tasks described in Section 4.1, based on the model coefficients α_i, β_i that were used to generate the transformed views. Even though the coefficient-based recognition was done as a post-hoc analysis, i.e. knowing the images, we did not fine-tune parame-

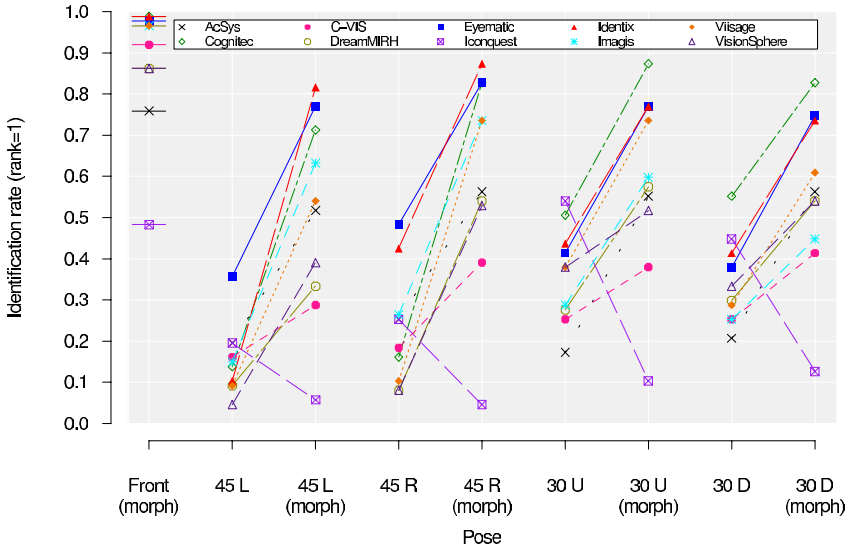

Figure 7. Rank 1 identification rates for original versus viewpoint-transformed probe images from different viewpoints. Gallery images were untransformed frontal views.

View	30 up	45 left	frontal	45 right	30 down
30 up	- -	0.07 0.36	0.45 0.83	0.08 0.43	0.14 0.36
45 left	0.07 0.36	- -	0.16 0.74	0.05 0.51	0.08 0.43
frontal	0.45 0.83	0.16 0.74	- -	0.18 0.85	0.54 0.80
45 right	0.08 0.43	0.05 0.51	0.18 0.85	- -	0.12 0.52
30 down	0.14 0.36	0.08 0.43	0.54 0.80	0.12 0.52	- -

Table 2. Cross viewpoint performance: Each pair of numbers gives the rate of correct acceptance (at 1% false accept rate) before and after viewpoint transformation of both gallery and probe images.

probe view	Coefficient Based	View Transformation
	gallery view front	gallery view front
up	81.6	81
down	65.5	80
right	79.3	80
left	79.3	77

Table 3. Percentages of correct acceptance for verification with coefficient-based recognition, compared to the best view-based system at each angle in the view-tranformation paradigm at 1% false accept rate.

ters of the system, but used the settings from previous studies [5]. The results are given in Tables 3 and 4 in a comparison with the performance of the best view-based system from Section 4.1. Since all evaluation conditions, including the fitting results, are identical, we can directly assess the appropriateness of the two approaches.

The results in Tables 3 and 4 indicate that the coefficient-based comparison and the recognition on transformed images perform equally well. The high rates of correct identification and verification indicate that diagnostic information about identity is preserved in the viewpoint transformation despite the fact that a projection from 3D to 2D is involved. For some viewing angles, view transformation slightly outperforms coefficient-based recognition, which may be due to the fact that not all structures in the textures, such as scars or moles, can be captured by the model coefficients, but are still transferred to the novel viewpoint, as described in Section 2.

5. Conclusions

We have presented a combined approach of a 3D Morphable Model with state–of–the–art face recognition systems for recognizing faces in images taken from arbitrary viewpoints. In a wide range of application scenarios, the approach can be applied both to probe and to gallery images, transforming faces into any given standard imaging conditions.

Our evaluation, part of which was included in the Face Recognition Vendor Test FRVT2002, has demonstrated that the viewpoint transformation by the Morphable Model reduced the effect of pose changes considerably for 9 out of 10 systems tested, and achieved promising recognition results on a challenging set of test images. In the FRVT 2002, participants were not aware prior to the evaluation that performance on viewpoint-transformed images would be com-

probe view	Coefficient Based					View Transformation
	gallery view					gallery view
	front	up	down	right	left	front
front	**–**	82.8	67.8	87.4	83.9	–
up	**89.7**	–	36.8	58.6	59.8	87
down	**80.5**	47.1	–	64.4	58.6	83
right	**82.8**	69.0	71.3	–	88.5	87
left	**89.7**	66.7	66.7	88.5	–	82
total	**85.6**	66.4	60.6	74.7	72.7	84.75

Table 4. Percentages of correct identification with coefficient-based recognition (left columns), and results in the viewpoint transformation paradigm with the best view-based system at each viewing angle (right column).

puted. Still, the Morphable Model significantly improved performance of recognition from non-frontal views without the systems adapting to the preprocessing. This suggests that the impact will be even greater when the systems have been tuned to the specific properties of the processed data, or the output of the Morphable Model's viewpoint transformation is optimized for their requirements.

Unlike other methods that involve 3D reconstruction, our system requires only a single image per person for the gallery, and a single probe image. Additional images can be incorporated by fusing separate results, which has further increased the correct acceptance rate in our experiments, or by fitting the model to multiple images simultaneously. Our method currently requires manual labelling of a small number of feature points in the image. We do not expect this labelling to be a major problem in many applications, for example if a relatively small number of suspects' images have to be added to a gallery that is used for search. Combined with a feature detection algorithm, our system can be fully automated.

The viewpoint transformation approach combines the efficiency of view-based methods with the versatility of the 3D Morphable Model. We have compared the performance with an algorithm that uses the model coefficients of the 3D Morphable Model directly, and found comparable recognition rates as those of the best systems in the combined approach. This indicates that the combined system does not involve a significant loss of diagnostic identity information when the transformed image is rendered and subsequently analyzed in view-based recognition. As a consequence, we may conclude that with the techniques presented in this paper, transformed images may be used as a natural interface between different face recognition algorithms.

References

[1] J. J. Atick, P. A. Griffin, and A. N. Redlich. Statistical approach to shape from shading: Reconstruction of 3D face surfaces from single 2d images. *Neural Computation* 8, pages 1321–1340, 1996.

[2] D. Beymer and T. Poggio. Face recognition from one model view. In *Proceedings of the 5th International Conference on Computer Vision*, 1995.

[3] V. Blanz, K. Scherbaum, T. Vetter, and H.-P. Seidel. Exchanging faces in images. In M.-P. Cani and M. Slater, editors, *Computer Graphics Forum, Vol. 23, No. 3 EUROGRAPHICS 2004*, pages 669–676, Grenoble, France, 2004.

[4] V. Blanz and T. Vetter. A morphable model for the synthesis of 3D faces. In *Computer Graphics Proc. SIGGRAPH'99*, pages 187–194, 1999.

[5] V. Blanz and T. Vetter. Face recognition based on fitting a 3d morphable model. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 25(9):1063– 1074, 2003.

[6] A. Georghiades, P. Belhumeur, and D. Kriegman. From few to many: Illumination cone models for face recognition under variable lighting and pose. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 23(6):643–660, 2001.

[7] P. Grother. Face recognition vendor test 2002: Supplemental Report. NISTIR 7083, Nat. Inst. of Standards and Technology, 2004.

[8] J. Kittler, M. Hatef, R. Duin, and J. Matas. On Combining Classifiers *IEEE Trans. on Pattern Analysis and Machine Intelligence* 20(3):226–239, 1998.

[9] P. J. Phillips, P. Grother, R. Michaels, D. Blackburn, E. Tabassi, and M. Bone. Face recognition vendor test 2002: Evaluation report. NISTIR 6965, Nat. Inst. of Standards and Tech., 2003.

[10] P. J. Phillips, H. Wechsler, J. Huang, and P. Rauss. The feret database and evaluation procedure for face recognition algorithms. *Image and Vision Computing J*, 16(5):295–306, 1998.

[11] T. Sim and T. Kanade. Combining models and exemplars for face recognition: An illuminating example. *Proc. Workshop on Models versus Exemplars in Computer Vision*, 2001.

[12] T. Vetter. Synthestis of novel views from a single face image. *Int. Journal of Comp. Vision* 28(2), pages 103–116, 1998.

[13] T. Vetter and T. Poggio. Linear object classes and image synthesis from a single example image. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 19(7):733–742, 1997.

[14] W. Zhao and R. Chellappa. SFS based view synthesis for robust face recognition. In *Int. Conf. on Autom. Face and Gesture Recognition*, pages 285–292, 2000.