

Person-specific Shape Estimation under Varying Head Pose from Single Snapshots

Fadi Dornaika^{1,2}

¹ *University of the Basque Country, San Sebastian, Spain*

² *IKERBASQUE, Basque Foundation for Science
Bilbao, Spain*

fadi_dornaika@ehu.es

Bogdan Raducanu

*Computer Vision Center
Bellaterra, Barcelona, Spain
bogdan@cvc.uab.es*

Abstract

This paper presents a new method for person-specific face shape estimation under varying head pose of a previously unseen person from a single image. We describe a featureless approach based on a deformable 3D model and a learned face subspace. The proposed approach is based on maximizing a likelihood measure associated with a learned face subspace, which is carried out by a stochastic and genetic optimizer. We conducted the experiments on a subset of Honda Video Database showing the feasibility and robustness of the proposed approach. For this reason, our approach could lend itself nicely to complex frameworks involving 3D face tracking and face gesture recognition in monocular videos.

1. Introduction

Offline or online computed 3D face shapes can be used in many applications such as face recognition [3, 5], 3D face pose tracking [10], and facial expression recognition [12]. Model-based applications exploiting monocular vision systems (the face model is given by a 3D mesh or a range model) need to personalize the face model of the person utilizing the system in order to achieve an accurate estimation. This holds true even with simple 3D models such as cylinders and ellipsoids. Recently some researchers proposed the use of special sensors such as a travelling camera or a 3-D scanner in order to build personalized facial shapes [2]. These shape models are then used for art production or for 3D face detection and recognition. Such systems suffer from several shortcomings. Some of the shortcomings can be alleviated by using stereo vision sensors. In [6], the authors propose to infer side-view shape pa-

rameters from one single frontal image using learned statistical correlation between the frontal-view parameters and the side-view parameters. The facial points (MPEG-4 points) and the frontal view parameters (relative distances) are extracted from the frontal image using some heuristics and prior knowledge.

The mainstream for shape estimation relies on extracting and matching some salient facial features such as the locations and local statistics of the eyes, nose, and mouth in one or more views. Thus, feature-based shape estimation not only require the extraction of the facial features but also requires a frontal view of the face. Feature-based approaches suffer from self-occlusions and drifting. A solution to overcome the drawbacks of feature-based approaches is given by holistic approaches (appearance-based approaches), which try to analyze the whole facial appearance [4, 8]. For example, Active Appearance Models (AAMs) were mainly used for 2D model fitting and tracking. In this paper, we present a new method for specific face shape estimation under varying head pose of a previously unseen person from a single image. Although the shape estimation is the main focus, our approach is intrinsically related with 3D head pose estimation, since they are both included in the proposed mathematical framework.

The proposed holistic approach estimates both the face shape control parameters as well as the 3D pose parameters by registering the input texture (warped region of the image) to a statistical face texture. Compared to AAMs methods our proposal has two advantages. First, there is no need to compute a Jacobian matrix neither offline nor online. Second, while AAMs merge both the inter and intra-person shape variabilities, our method separates these variabilities, and therefore the proposed method can be easily and efficiently used for initializing a real time 3D face tracker and facial expression recognizer in videos (both the personalized 3D model and its

3D pose are computed for the first frame in the video sequence). However, it is not clear how these tasks can be performed with AAMs.

Our approach does not use neither 2D AAM nor 3D AAM. The only similarity with AAMs is the use of a statistical facial texture model based on Principal Component Analysis (PCA). The remainder of the paper is organized as follows. Section 2 describes the face modelling aspects. Section 3 presents the proposed holistic approach for the simultaneous estimation of the 3D pose and shape. Section 4 presents some experimental results. Section 5 concludes the paper.

2. Modelling faces

A deformable 3D mesh In our study, we use *Candide* 3D face model [1]. This common 3D deformable wireframe model accounts for person specific shape variation as well as for facial animation. The 3D shape of this wireframe model (triangular mesh) is directly recorded in coordinate form. It is given by the coordinates of its n 3D vertices. Thus, the shape up to a global scale can be fully described by the $3n$ -vector \mathbf{g} ; the concatenation of the 3D coordinates of all vertices. The vector \mathbf{g} is written as:

$$\mathbf{g} = \bar{\mathbf{g}} + \mathbf{S} \boldsymbol{\tau}_s + \mathbf{A} \boldsymbol{\tau}_a \quad (1)$$

where $\bar{\mathbf{g}}$ is the standard shape of the model, $\boldsymbol{\tau}_s$ and $\boldsymbol{\tau}_a$ are shape and animation control vectors, respectively, and the columns of \mathbf{S} and \mathbf{A} are the Shape and Animation Units. Both matrices are provided by *Candide* model package. A Shape Unit provides a means of deforming the 3D wireframe so as to be able to adapt eye width, head width, eye separation distance, etc (see Figure 1). Thus, the term $\mathbf{S} \boldsymbol{\tau}_s$ accounts for shape variability (inter-person variability) while the term $\mathbf{A} \boldsymbol{\tau}_a$ accounts for the facial animation (intra-person variability). With this model, the ideal neutral face configuration is represented by $\boldsymbol{\tau}_a = \mathbf{0}$. In this study, we assume that the images are depicting quasi-neutral faces. Thus, the expression for the deformable mesh becomes:

$$\mathbf{g} = \bar{\mathbf{g}} + \mathbf{S} \boldsymbol{\tau}_s \quad (2)$$

The shape modes were created manually to accommodate the subjectively most important changes in facial shape. In the model package, the number of modes associated with facial Shape Units matrix \mathbf{S} (inter-person variability) is twelve. However, for the purpose of our study which deals with the automatic image-based extraction of the control vector $\boldsymbol{\tau}_s$ only six components are considered as the most significant indicators of the perceived person-dependent facial shape in a given near frontal facial image. These components are:

Head height, vertical position of the eye brows, vertical position of the eye, eyes separation distance, vertical position of the nose, vertical position of the mouth. The remaining components are set to nominal values.




Figure 1. Effects of some facial shape control parameters on the deformable 3D model (standard shape, mouth width, eyes width, eyes vertical position, eye separation distance, head height).

In equation (1), the 3D shape is expressed in a local coordinate system. However, one should relate the 3D coordinates to the image coordinate system. To this end, we should add the six degrees of freedom associated with the 3D face pose. The mapping between the 3D face model and the image adopts the weak perspective projection model.

Thus, the state of the 3D wireframe model is given by the 3D face pose parameters (three rotations and three translations) and the shape control vector $\boldsymbol{\tau}_s$. This is given by the 12-dimensional vector \mathbf{b} :

$$\mathbf{b} = [\theta_x, \theta_y, \theta_z, t_x, t_y, t_z, \boldsymbol{\tau}_s^T]^T \quad (3)$$

Shape-free facial patches A facial patch is represented as a shape-free image (geometrically normalized rawbrightness image). The geometry of this image is obtained by projecting the standard shape $\bar{\mathbf{g}}$ using a centered frontal 3D pose onto an image with a given resolution. The texture of this geometrically normalized image is obtained by texture mapping from the triangular 2D mesh in the input image (see figure 2) using a piece-wise affine transform.




Figure 2. (a) an input image with correct fitting. (b) the corresponding shape-free facial patch.

3 Shape and 3D head pose parameter estimation

3.1. Face subspace

The statistical facial texture model describes the appearance variation of the shape-free facial patches \mathbf{x} (see figure 2.(b)). These patches are obtained from the training images (individual snapshots or video sequences) by fitting the 3D deformable model to the face. This fitting can be manual or automatic [1]. Using these training patches one can easily build the face subspace. For this purpose we use the Principal Component Analysis (PCA)—a well-known technique used for modeling face subspaces. We assume that we have K shape-free patches. Applying a PCA on the training patches we can compute the mean and the principal modes of variation. The use of linear subspace (PCA) can be justified by 1) the facial images are geometrically normalized, and 2) the proposed method will be carried out in a relatively constrained environment. However, currently we are investigating the use of non-linear manifold learning techniques.

3.2. Optimization

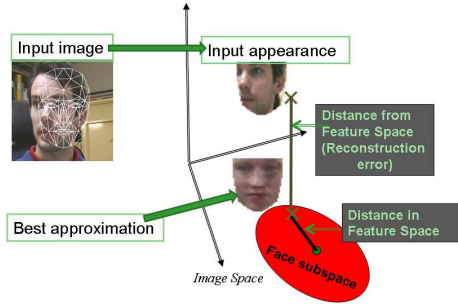


Figure 3. The unknown parameters are estimated by maximizing a likelihood measure taking into account the reconstruction error and the distance in feature space. The face subspace is linear.

The basic idea is to estimate the 3D face pose and shape parameters, i.e. the vector \mathbf{b} , such that the associated shape-free patch will be as close as possible to the facial sub-space. This can be carried out by maximizing a certain likelihood measure. For this purpose, we use the likelihood measure proposed in [9]:

$$p(\mathbf{x}|\mathbf{b}) \propto \exp\left(-\frac{1}{2} \sum_{i=1}^M \frac{\xi_i^2}{\lambda_i}\right) \exp\left(-\frac{e}{2\rho^*}\right) \quad (4)$$

where e is the reconstruction error, λ_i s are the M largest eigenvalues given by the PCA, ξ_i s represent the texture projection onto the corresponding M eigenvectors, and ρ^* is the arithmetic average of the remaining eigenvalues (in the complementary subspace). The reconstruction error is the distance between the original shape-free texture \mathbf{x} and its projection onto the PCA subspace. The above likelihood measure takes into account two distances (i) the distance-from-feature-space, and (ii) the distance-in-feature-space. These two distances are illustrated in Figure 3. Maximizing this likelihood is equivalent to minimizing the *Mahalanobis* distance over the original textures. The unknown 3D face pose and shape parameters (the vector \mathbf{b}) can be estimated by seeking the maximum of the likelihood (4):

$$\mathbf{b} = \arg \max_{\mathbf{b}} p(\mathbf{x}|\mathbf{b}) \quad (5)$$

To this end, we use the Differential Evolution (DE) algorithm [11] in order to maximize (4) with respect to the 3D face pose and shape parameters. The DE algorithm is a practical approach to global numerical optimization that is easy to implement, reliable and fast. The crucial idea behind DE is a scheme for generating trial parameter vectors. Basically, DE adds the weighted difference between two population vectors to a third vector. In our case, the initial population is randomly selected between the lower and upper bounds defined for each variable using uniform distributions. The distributions associated with the translational part of the 3D face pose are centered on the output of the 2D face detector [13].

4 Experimental results

Experiments were conducted to evaluate the performance of the proposed fitting algorithm in image snapshots extracted from several video sequences of the Honda video database [7]. The video sequences were recorded in realistic conditions, with persons featuring unconstrained in-plane and out-of-plane head movements. The sequences are at least 15 seconds long and are recorded at 15 frames per second. A subset of 20 video sequences (corresponding to 20 different persons) has been retrieved for our experiments.

However, since our 3D model is based on Candide model not all snapshots extracted from the videos database can be used. Only those depicting 3D head poses belonging to the interval $[-40^\circ, +40^\circ]$ for the pitch and yaw angles were considered. In total, we selected for our experiments about 900 images. A PCA model is built from a set of 500 shape-free templates, belonging to 5 persons. The remaining persons were used for test. We have found that PCA models with 20

principal components are usually enough for representing the face sub-space. More precisely, we found that the retained variance is above 95% of the total variance. Figure 4 illustrates the fitting results obtained with four unseen persons. As can be appreciated, the face pose and shape parameters (relative positions of eyebrows, eyes, nose, and mouth) are correctly fitted on the face.



Figure 4. 3D face pose and person-specific shape estimation associated with four unseen persons.

Quantitatively speaking, we performed an evaluation process taking into account both intra-person and inter-person estimation accuracy. Ground-truth parameters have been obtained manually. Table 1 depicts the average error of shape parameters in the first case. Due to lack of space, we present the results of only 5 persons from the total 15. The shape parameters are normalized, i.e., each parameter belongs to the interval $[-1, 1]$.

	eyebrow	eye	eyes separa.	nose	mouth
#1	0.7%	6.2%	3.2%	6.4%	1.4%
#7	2.8%	0.4%	0.3%	2.8%	1.1%
#10	3.1%	4.3%	1.1%	5.9%	0.7%
#13	1.9%	1.6%	2.3%	3.6%	1.6%
#15	2.5%	3.5%	2.1%	5.3%	1.7%

Table 1. Average error for five individuals.

Table 2 depicts the average error between the manually obtained parameters and the automatically estimated ones for all the persons considered.

5. Conclusions

This paper presented a featureless method that fits a generic deformable 3D face model to a single facial image where the face is not required to be frontal. The fitted parameters are some salient shape control parameters as well as the 3D face pose parameters. The proposed method has several

advantages that make it attractive, being useful for 3D face pose tracking and facial expression recognition in real-time.

	eyebrow	eye	eyes separa.	nose	mouth
All	2.4%	3.6%	1.6%	4.2%	1.4%

Table 2. Global average error.

Acknowledgments

This work has been supported in part by the projects TIN2009-14404-C02-00 and CONSOLIDER-INGENIO CSD 2007-00018, Ministerio de Educación y Ciencia, Spain.

References

- [1] J. Ahlberg. An active model for facial feature tracking. *EURASIP Journal on Applied Signal Processing*, 2002(6):566–571, June 2002.
- [2] M. D. Breitenstein, D. Kuettel, T. Weise, L. Gool, and H. Pfister. Real-time face pose estimation from single range images. In *IEEE International Conference on Computer Vision and Pattern Recognition*, 2008.
- [3] A. Bronstein, M. Bronstein, and R. Kimmel. Three-dimensional face recognition. *International Journal of Computer Vision*, 64:5–30, 2005.
- [4] T. Cootes, G. Edwards, and C. Taylor. Active appearance models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(6):681–684, 2001.
- [5] F. Haar and R. Velkamp. 3D face model fitting for recognition. In *European Conference on Computer Vision*, 2008.
- [6] C. J. Kuo, R. Huang, and T. Lin. 3D facial model estimation from single front-view facial image. *IEEE Trans. on Circuits and Systems for Video Technology*, 12(3):183–192, 2002.
- [7] K. Lee, J. Ho, M. Yang, and D. Kriegman. Video-based face recognition using probabilistic appearance manifolds. *IEEE Conf. On Computer Vision and Pattern Recognition*, 1:313–320, 2003.
- [8] I. Matthews and S. Baker. Active appearance models revisited. *International Journal of Computer Vision*, 60(2):135–164, 2004.
- [9] B. Moghaddam and A. Pentland. Probabilistic visual learning for object representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(7):696–710, 1997.
- [10] E. Murphy-Chutorian and M. Trivedi. Head pose estimation in computer vision: a survey. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 31(4):607–626, 2009.
- [11] K. V. Price, J. A. Lampinen, and R. M. Storn. *Differential Evolution: A Practical Approach To Global Optimization*. Springer, 2005.
- [12] Y. Sun and L. Yin. Facial expression recognition based on 3D dynamic range model sequences. In D. Forsyth, P. Torr, and A. Zisserman, editors, *Computer Vision – ECCV 2008*, volume 5303 of *LNCS*, 2008.
- [13] P. Viola and M. Jones. Robust real-time object detection. *International Journal of Computer Vision*, 57(2):137–154, 2004.