Face Frontalization Based on Robustly Fitting a Deformable Shape Model to 3D Landmarks

Zhiqi Kang, Mostafa Sadeghi, Radu Horaud

Face Frontalization Based on Robustly Fitting a Deformable Shape Model to 3D Landmarks

Zhiqi Kang, Mostafa Sadeghi and Radu Horaud

Abstract—Face frontalization consists of synthesizing a frontally-viewed face from an arbitrarily-viewed one. The main contribution of this paper is a robust face alignment method that enables pixel-to-pixel warping. The method simultaneously estimates the *rigid transformation* (scale, rotation, and translation) and the *non-rigid deformation* between two 3D point sets: a set of 3D landmarks extracted from an arbitrary-viewed face, and a set of 3D landmarks parameterized by a frontally-viewed deformable face model. An important merit of the proposed method is its ability to deal both with noise (small perturbations) and with outliers (large errors). We propose to model inliers and outliers with the generalized Student's t-probability distribution function – a heavy-tailed distribution that is immune to non-Gaussian errors in the data. We describe in detail the associated expectation-maximization (EM) algorithm that alternates between the estimation of (i) the rigid parameters, (ii) the deformation parameters, and (iii) the t-distribution parameters. We also propose to use the *zero-mean normalized cross-correlation*, between a frontalized face and the corresponding ground-truth frontally-viewed face, to evaluate the performance of frontalization. To this end, we use a dataset that contains pairs of profile-viewed and frontally-viewed faces. This evaluation, based on direct image-to-image comparison, stands in contrast with indirect evaluation, based on analyzing the effect of frontalization on face recognition.[1]

I. INTRODUCTION

The problem of face frontalization is the problem of synthesizing a frontal view of a face from an arbitrary view. Recent research has shown that face frontalization consistently boosts the performance of face analysis. In particular, it has been recently demonstrated that face recognition from frontal views yields better performance than face recognition from unconstrained views [1]–[3]. This observation is equally valid for other tasks, such as the analysis of facial expressions [4] or of lip reading [5], [6].

It is well established that lip movements, as well as tongue and jaw movements, are controlled by speech production and that they are correlated with word pronunciation [7]. Consequently, the capacity to properly analyse facial movements plays an important role in visual and audio-visual speech separation, speech enhancement and speech recognition, e.g. [8]–[12]. This is particularly useful when audio signals are corrupted by ambient noise and by acoustic perturbations. Nevertheless, facial-movement analysis is perturbed by inherent rigid head movements. It is therefore important to separate

non-rigid facial movements from rigid head movements, and face frontalization may well be viewed as such a rigid/non-rigid separation process.

In this paper we address face frontalization as the problem of simultaneously estimating the 3D pose and deformation parameters of an arbitrarily-viewed face, e.g. Figure 1. The main contribution is a robust point-set alignment method that enables face frontalization via pixel-to-pixel warping. In details, we propose a method that simultaneously estimates the *rigid transformation* (scale, rotation, and translation) and the *non-rigid deformation* between two 3D point sets, namely (i) a set of 3D landmarks extracted from an arbitrary-viewed face and (ii) a set of frontally-viewed 3D landmarks that are parameterized by a deformable face model. An important merit of the proposed alignment estimator is its ability to deal both with noise (small perturbations) and with outliers (large errors), i.e. robust estimation. We propose to model both inliers and outliers with the generalized Student's t-probability distribution function (pdf) – a heavy-tailed distribution that is immune to non-Gaussian errors in the data [13], [14]. We describe in detail the associated expectation-maximization (EM) algorithm that alternates between the estimation of (i) the rigid parameters, (ii) the deformation parameters, and (iii) the pdf parameters. Interestingly, the proposed method can be indifferently applied to align two rigid point sets or to *fit and align* a deformable point set to a rigid one.

We also propose to use the *zero-mean normalized cross-correlation* (ZNCC) coefficient between a frontalized face and the corresponding ground-truth frontally-viewed face, in order to empirically evaluate the performance of face frontalization. To this end, we use a dataset that contains pairs of profile-viewed and frontally-viewed faces. This direct evaluation, based on image-to-image comparison between the prediction and the ground-truth, stands in contrast with indirect evaluation, based on analyzing the effect of frontalization on face recognition, e.g. [1]–[3].

The remainder of this article is organized as follows. In Section II we review face frontalization methods that were proposed in the recent past. Section III outlines the proposed algorithm and describes in detail the proposed 3D point-set alignment method. Section IV analyzes some technical aspects of the method and provides implementation details. Section V describes the proposed evaluation metrics and shows results obtained with our method and with two state-of-the-art methods. Section VI draws some conclusions. Finally, Appendix A, Appendix B and Appendix C provide mathematical details of the proposed algorithms.

Z. Kang, M. Sadeghi and R. Horaud are with Inria Grenoble and with Université Grenoble Alpes.

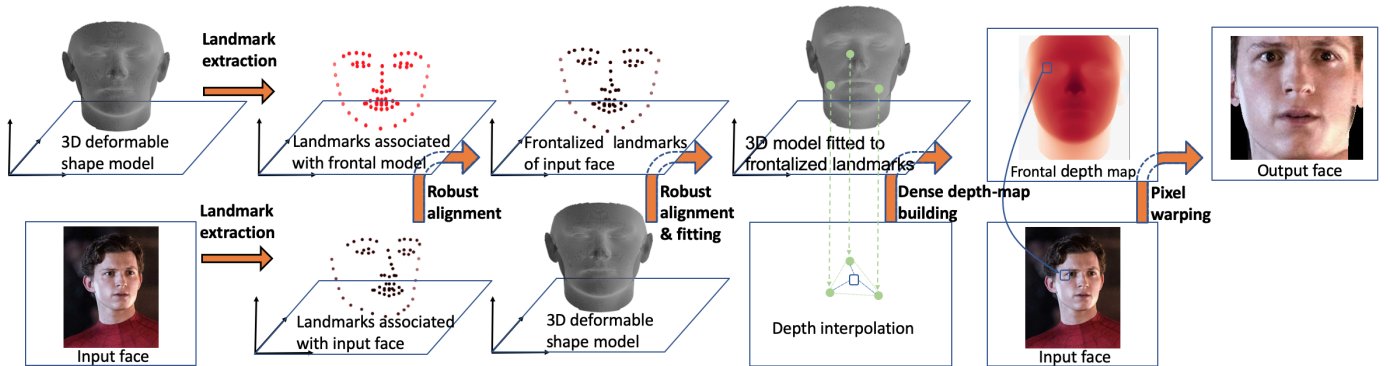[1]Supplemental material for this paper is accessible at https://team.inria.fr/perception/research/rff/.

Fig. 1: Overview of the proposed robust face frontalization method. 3D landmarks are extracted from both an arbitrarly-viewed input face and from a frontally-viewed deformable shape model. These landmarks are robustly aligned, thus enabling to compute the pose of the input face and to frontalize the input landmarks. The deformable shape model is next fitted to these landmarks, thus obtaining a frontalized shape model, followed by the computation of a frontal dense depth map; the latter is obtained by interpolation of the 3D vertices of the triangulated mesh that describes the shape model. Finally, the input-face pixels are warped onto the output-image pixels.

II. RELATED WORK

As already mentioned, face frontalization consists of synthesizing a frontally-viewed face from an arbitrarily-viewed face. A popular approach has been to train deep neural networks (DNNs) to learn a mapping from an arbitrary view to a frontal view, using a massive collection of input/target pairs of faces. It was shown that architectures based on generative adversarial networks (GANs), e.g. [3], [15]–[17] outperform architectures based on convolutional neural networks (CNNs) [1] for face recognition. The main drawback of these DNN-based methods is that they are designed to predict frontal faces that are as neutral as possible, in order to improve the performance of face recognition. Hence, there is no guarantee that the predicted output preserves non-rigid facial deformations. These methods are therefore hardly usable for facial expression recognition or for lip reading.

Another approach has been to estimate the 3D pose of an input face with respect to a frontal 3D face model, and then to use the pose parameters to warp the facial pixels from the input image onto a frontal view. Recently proposed face frontalization methods capitalize on 3D pose estimation, e.g. [18], [2]. The problem of 3D pose estimation from point-to-point correspondences has been thoroughly addressed in the past and many solutions have been proposed in the presence of various camera models, e.g. [19]. In [18] it was proposed to use a 3D generic model of a face from which a frontal image is synthesized. Next, 48 2D facial landmarks are extracted from the input face and from the generic frontal face, which provides 2D-to-3D correspondences between the input face and the generic 3D model. This amounts to estimate the projection (or camera) matrices, from the 3D face model onto the input image and onto the frontal image, respectively, and to subsequently warp pixels from one image to another. A similar method was proposed in [2] where 68 2D landmarks are extracted from both the input face and from a frontal face associated with a generic 3D model. The 3D pose parame-ters are obtained from 2D-to-3D landmark correspondences between the input face and the 3D model.

Nevertheless, the pose-based methods just cited suffer from a number of limitations. First, the 3D models that they use do not take facial deformations into account. Second, the estimation of 3D pose with an unknown projective camera model is an ill-posed problem in the case of quasi-planar shape models such as faces. Third, the pose estimator itself is not robust to the presence of large errors in landmark detection and localization, and to large non-rigid facial deformations.

The proposed method overcomes these limitations on the following grounds. First, we use a deformable 3D face model that parameterizes faces with large variabilities in appearance and in expression. Deformable shape models were initially introduced in [20], and subsequently used in [21]–[23] to model faces and facial deformations. Second, we use 3D facial landmarks rather than 2D landmarks, e.g. [24]–[26]. Consequently, there is no need to estimate the parameters of the camera model, at the price of approximating the latter with a scaled orthographic (or weak perspective) model. However, the latter is a realistic approximation of the true camera model, since human faces are quasi planar surfaces and, hence, their depth range is small relative to the camera-to-face distance [19]. Third, we propose a robust method for aligning a 3D set of landmarks with a 3D deformable face model. The associated EM solver alternates between the estimation of the rigid-transformation parameters, the model-deformation parameters, and the robust statistical-model parameters. This paper builds on a recent study showing that the use of 3D facial landmarks in combination with robust pose estimation yields excellent performance [27].

As already mentioned, face frontalization has been essentially used as a pre-processing step for face recognition and the evaluation metrics proposed so far are generally based on the recognition rate [1]–[3], [28]. This kind of evaluation lacks direct significance for standalone face frontalization meth-

Algorithm 1: Robust face frontalization (RFF).
Input: Face image and deformable 3D shape model
Initialization: Extract landmarks from the input face. Extract frontal landmarks from the shape model ;
Pose estimation: Compute the scale, rotation and translation parameters between the input landmarks and the model landmarks ;
Landmark frontalization: Apply the pose parameters to the input landmarks ;
Model fitting: Align the shape model (scale, rotation, translation and deformation parameters) with the frontalized landmarks and frontalize the shape model ;
Frontal dense depth map: Interpolate the frontalized shape model ;
Face warping: Compute pixel-to-pixel correspondences between input and output face images ;
Output: Frontalized face image

ods. Moreover, existing evaluation pipelines include facial-expression normalization [29], which biases the results and does not allow to evaluate the extent to which face frontalization preserves non-rigid facial deformations.

III. ROBUST FACE FRONTALIZATION

A. Outline of the Proposed Method

The proposed robust face frontalization (RFF) method is summarized in Algorithm 1 and Figure 1. The core idea of the method is to estimate the 3D pose (scale, rotation, and translation) of an input face based on robust alignment of two 3D point sets, namely a set of 3D landmarks extracted from the input face, and a set of 3D landmarks associated with the frontal pose of a 3D shape model. This is done in the framework of maximum likelihood estimation and, in practice, an expectation-maximization algorithm is being used, i.e. (21) and Appendix A. This allows to frontalize the input landmarks, i.e. (5) and, then, to fit a deformable shape model to the frontalized landmarks, which yields a frontalized shape model of the input face, i.e. (10). Next, vertex interpolation is used to compute a dense depth map associated with the frontalized shape model, i.e. (11) and (12). Eventually, a frontalized face is computed by warping the input face onto a frontal image, i.e. (13) and (14).

B. Frontalization via Alignment of 3D Point Sets

Let I_p be an *observed* image of a face in an unknown pose.[2] A set $\mathcal{X}$ of $J = 68$ 3D landmarks is extracted from I_p with image-centered coordinates $\boldsymbol{X}_{1:J} = \{\boldsymbol{X}_j\}_{j=1}^J \subset \mathbb{R}^3$. Throughout the paper we adopt the notation $\boldsymbol{X}_j =$

---

[2] The term *face pose* refers to a rigid transformation, namely the scale factor, rotation matrix and translation vector that map a face-centered coordinate frame onto a coordinate frame that is aligned with the camera frame.

(X_{j1}, X_{j2}, X_{j3}) to designate the three coordinates of a point in $\mathbb{R}^3$. Let $\boldsymbol{Z}_{1:J} = \{\boldsymbol{Z}_j\}_{j=1}^J \subset \mathbb{R}^3$ be the 3D coordinates of a set of landmarks, $\mathcal{Z}$, that correspond to the frontal pose of a 3D deformable face model, e.g. [27], and let I_f be the frontal image corresponding to this frontal view of the model. The problem of face-pose estimation consists of finding the rigid transformation that best maps $\boldsymbol{X}_{1:J}$ onto $\boldsymbol{Z}_{1:J}$ on the premise that the landmarks are in one-to-one correspondence and up to an error (or residual) $\boldsymbol{e}_j$:

$$\boldsymbol{e}_j = \boldsymbol{Z}_j - (\rho \mathbf{R} \boldsymbol{X}_j + \boldsymbol{t}), \quad \forall j \in \{1 \ldots J\}, \qquad (1)$$

where $\rho \in \mathbb{R}^+$, $\mathbf{R} \in \mathbb{R}^{3\times3}$ and $\boldsymbol{t} \in \mathbb{R}^3$ are the scale, rotation matrix and translation vector, respectively, associated with the unknown face pose. Because the landmark locations $\boldsymbol{X}_{1:J}$ are inherently affected by detection noise and outliers, as well as by *non-rigid facial deformations*, it is suitable to use a robust alignment technique. For this purpose, we assume that the residuals $\boldsymbol{e}_{1:J}$ are samples of a random variable $\boldsymbol{e}$ drawn from a robust probability distribution function (pdf) $P(\boldsymbol{e}; \boldsymbol{\theta})$, where $\boldsymbol{\theta}$ is the set of parameters characterizing the pdf. Then, the problem is cast into maximum likelihood estimation (MLE) or, equivalently into the minimization of the following negative log-likelihood function:

$$\mathcal{L}(\boldsymbol{\theta}, \rho, \mathbf{R}, \boldsymbol{t} | \boldsymbol{X}_{1:J}, \boldsymbol{Z}_{1:J}) = -\frac{1}{2} \sum_{j=1}^J \log P(\boldsymbol{e}_j; \boldsymbol{\theta}), \quad (2)$$

The generalized Student's t-distribution [14] is a robust pdf, namely:

$$P(\boldsymbol{e}; \boldsymbol{\theta}) = \int_0^\infty \mathcal{N}(\boldsymbol{e}; 0, w^{-1}\boldsymbol{\Sigma})\mathcal{G}(w; \mu, \nu)dw \qquad (3)$$

where $\mathcal{N}()$ is the normal distribution, $\mathcal{G}()$ is the gamma distribution and $\boldsymbol{\Sigma} \in \mathbb{R}^{3\times3}$ is a covariance matrix. $w \in \mathbb{R}$ is a precision associated with a residual. The precisions $w_{1:J}$ are modeled as random variables drawn from gamma distributions.

Direct minimization of (2) is not practical. An expectation-maximization (EM) formulation is therefore adopted, namely minimization of the *expected complete-data negative log-likelihood*:

$$\min_{\boldsymbol{\theta}} \mathrm{E}_W[-\log P(\boldsymbol{e}_{1:J}, w_{1:J} | \boldsymbol{e}_{1:J}; \boldsymbol{\theta})], \qquad (4)$$

where EM alternates between the estimation of the means of the precision posteriors, $\overline{w}_{1:J}$ and the estimation of the parameters $\boldsymbol{\theta} = \{\rho, \mathbf{R}, \boldsymbol{t}, \boldsymbol{\Sigma}, \mu\}$. As it can be seen in Appendix A, the precisions are precious because they weight the relative importance of the landmarks within the optimization process, and hence they help the parameter estimation process to be robust against badly localized landmarks. Once the pose parameters are estimated, the rigid transformation is applied to the landmarks $\boldsymbol{X}_{1:J}$ in order to obtain a set $\mathcal{Y}$ of *frontalized* landmarks whose coordinates in I_f are denoted $\boldsymbol{Y}_{1:J} \subset \mathbb{R}^3$, namely:

$$\boldsymbol{Y}_j = \rho \mathbf{R} \boldsymbol{X}_j + \boldsymbol{t}, \quad \forall j \in \{1 \ldots J\}. \qquad (5)$$

The next step is to fit a deformable 3D face-shape model to this set of landmarks in order to eventually obtain a *frontal*

dense depth map of the face. For that purpose and without loss of generality, we consider a linear deformation model, e.g. the 3DMM Basel Face Model (BFM) [22]. The latter consists of a 3D mesh with a set $\hat{\mathcal{V}}$ of N vertices, whose coordinates $\hat{\boldsymbol{V}}_{1:N} = \{\hat{\boldsymbol{V}}_n\}_{n=1}^N \subset \mathbb{R}^3$ are parameterized by a *statistical linear shape-model* in the following way (please consult Appendix B):

$$\hat{\boldsymbol{V}}_n = \overline{\boldsymbol{V}}_n + \mathbf{W}_n \boldsymbol{s}, \quad \forall n \in \{1 \ldots N\}, \tag{6}$$

where $\overline{\boldsymbol{V}}_{1:N} \subset \mathbb{R}^3$ are the vertices of the mean shape model, $\mathbf{W}_{1:N} \subset \mathbb{R}^{3 \times K}$ are reconstruction matrices, and $\boldsymbol{s} \in \mathbb{R}^K$ is a low dimensional embedding of the vertex set, with $K \ll 3N$, see Appendix B. In order to fit this deformable 3D face-shape model to the frontalized landmarks $\boldsymbol{Y}_{1:J}$, we consider a subset of $J = 68$ vertices with coordinates $\hat{\boldsymbol{V}}_{1:J}$ that correspond, one-to-one, to the frontalized landmarks, namely $\{\hat{\boldsymbol{V}}_j \leftrightarrow \boldsymbol{Y}_j\}_{j=1}^J$.[3] For that purpose, the statistical shape model must be scaled, rotated, translated and deformed, such that the vertices $\hat{\boldsymbol{V}}_{1:J}$ are optimally aligned with the landmarks $\boldsymbol{Y}_{1:J}$. This yields the following negative log-likelihood function:

$$\mathcal{L}(\boldsymbol{\theta}, \sigma, \mathbf{Q}, \boldsymbol{d}, \boldsymbol{s} | \boldsymbol{Y}_{1:J}, \hat{\boldsymbol{V}}_{1:J}) = -\frac{1}{2} \sum_{j=1}^J \log P(\boldsymbol{r}_j; \boldsymbol{\theta}), \tag{7}$$

where σ, $\mathbf{Q}$ and $\boldsymbol{d}$ parameterize the rigid transformation that aligns the two point sets. The residual $\boldsymbol{r}_j$ is given by:

$$\boldsymbol{r}_j = \boldsymbol{Y}_j - (\sigma \mathbf{Q}(\overline{\boldsymbol{V}}_j + \mathbf{W}_j \boldsymbol{s}) + \boldsymbol{d}), \quad \forall j \in \{1 \ldots J\}. \tag{8}$$

As above, one can use the generalized Student's t-distribution (3) and an expectation-maximization algorithm to robustly estimate the model parameters. The notable difference between (7) and (2) is that, in addition to scale, rotation and translation, the deformable shape parameters must be estimated as well. The shape parameters are estimated with (please consult Appendix C):

$$\boldsymbol{s} = \left(\sum_{j=1}^J \overline{w}_j \mathbf{A}_j^\top \boldsymbol{\Sigma}^{-1} \mathbf{A}_j + \kappa \boldsymbol{\Lambda}^{-1} \right)^{-1} \left(\sum_{j=1}^J \overline{w}_j \mathbf{A}_j^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{b}_j \right), \tag{9}$$

where $\mathbf{A}_j = \sigma \mathbf{Q} \mathbf{W}_j$ and $\boldsymbol{b}_j = \boldsymbol{Y}_j - \sigma \mathbf{Q} \overline{\boldsymbol{V}}_j - \boldsymbol{d}$. The shape vertices can now be mapped onto the frontal view, namely $\tilde{\boldsymbol{V}}_n = (\tilde{V}_{n1}, \tilde{V}_{n2}, \tilde{V}_{n3})$, with:

$$\tilde{\boldsymbol{V}}_n = \sigma \mathbf{Q}(\overline{\boldsymbol{V}}_n + \mathbf{W}_n \boldsymbol{s}) + \boldsymbol{d}, \quad \forall n \in \{1 \ldots N\}. \tag{10}$$

A frontal dense depth map is then computed in the following way. Remember that the shape vertices form a triangulated 3D mesh; therefore the projection of $\tilde{\boldsymbol{V}}_{1:N}$ onto the image plane I_f form a 2D triangulated mesh whose vertices have $(\tilde{V}_{n1}, \tilde{V}_{n2})_{1:N}$ as 2D coordinates. Let n_1, n_2 and n_3 be the indexes of the vertices of a mesh triangle. Moreover, let a_1, a_2 and a_3 be the barycentric coordinates of a pixel (A_1, A_2) that lies inside that triangle, with $0 \leq a_1, a_2, a_3 \leq 1$ and $a_1 + a_2 + a_3 = 1$. The barycentric coordinates of pixel (A_1, A_2) are computed with:

$$\begin{pmatrix} A_1 \\ A_2 \end{pmatrix} = a_1 \begin{pmatrix} \tilde{V}_{n_1 1} \\ \tilde{V}_{n_1 2} \end{pmatrix} + a_2 \begin{pmatrix} \tilde{V}_{n_2 1} \\ \tilde{V}_{n_2 2} \end{pmatrix} + a_3 \begin{pmatrix} \tilde{V}_{n_3 1} \\ \tilde{V}_{n_3 2} \end{pmatrix}. \tag{11}$$

Once the barycentric coordinates are thus determined, the depth A_3 of the pixel is computed by interpolation, as follows:

$$A_3 = a_1 \tilde{V}_{n_1 3} + a_2 \tilde{V}_{n_2 3} + a_3 \tilde{V}_{n_3 3} \tag{12}$$

The above procedure is repeated for all the triangles and for all the points inside each triangle, thus obtaining a frontal dense depth map of the face.

The final face frontalization step consists of *warping* the input face onto a frontal face. The rigid transformation from the frontal view to the profile view is the inverse of the pose, namely the inverse of (5): $\rho' = \rho^{-1}$, $\mathbf{R}' = \mathbf{R}^\top$, and $\boldsymbol{t}' = -\rho^{-1}\mathbf{R}^\top \boldsymbol{t}$. Assuming scaled orthographic projection, a one-to-one correspondence between a pixel $(A_1, A_2) \in I_f$ with depth A_3, and a pixel $(B_1, B_2) \in I_p$ is obtained with the following formula:

$$\begin{pmatrix} B_1 \\ B_2 \end{pmatrix} = \rho' \begin{pmatrix} R'_{11} & {}'R'_{12} & {}'R'_{13} \\ R'_{21} & {}'R'_{32} & {}'R'_{33} \end{pmatrix} \begin{pmatrix} A_1 \\ A_2 \\ A_3 \end{pmatrix} + \begin{pmatrix} t'_1 \\ t'_2 \end{pmatrix} \tag{13}$$

Finally, face frontalization consists of building an image whose intensities (or colors) are provided by:

$$I_f(A_1, A_2) = I_p([B_1], [B_2]). \tag{14}$$

where $[\cdot]$ returns the integer part of a real number.

IV. ALGORITHM ANALYSIS AND IMPLEMENTATION DETAILS

The proposed method is summarized in Algorithm 1. The method starts with 3D facial landmark extraction which is achieved with [24] as it is one of the best-performing methods [27]. Moreover, the pose estimation and model fitting steps of Algorithm 1 are achieved with Algorithm 2 (Appendix A) and Algorithm 3 (Appendix C), respectively. Both these two algorithms are EM procedures and hence they have good convergence properties.

All the computations inside these two algorithms are in closed-form, with the notable exception of the estimation of the rotation matrix. The latter is parameterized with a unit quaternion [30], which allows us to reduce the number of parameters, from nine to four, and to express the orthogonality constraints of the rotation matrix in a much simpler way. The minimization (27) is solved using a sequential quadratic problem [31]. More precisely, a sequential least squares pro-gramming (SLSQP) solver[4] is used in combination with a root-finding software package [32]. The SLSQP minimizer found at the previous EM iteration is used to initialize the current EM iteration. At the start of EM, the closed-form method of [30] is used to initialize the rotation.

The proposed method also requires a deformable shape model, i.e. Appendix B. We consider a set $\mathcal{S}^I = \{\boldsymbol{S}_1^I, \ldots, \boldsymbol{S}_m^I, \ldots, \boldsymbol{S}_M^I\} \subset \mathbb{R}^{3N}$ of M shapes, where $\boldsymbol{S}_m^I$ is a concatenation of N 3D coordinates that corresponds to a different face identity, and where each face was scanned in a

---

[3]For the sake of simplifying the notations, we use $\{1 \ldots J\} \subset \{1 \ldots N\}$.

[4]https://docs.scipy.org/doc/scipy/reference/optimize.html

frontal view and with a neutral expression. We also consider a set $\mathcal{S}^E = \{\boldsymbol{S}_1^E, \ldots, \boldsymbol{S}_m^E, \ldots, \boldsymbol{S}_M^E\} \subset \mathbb{R}^{3N}$ that contains M scans of the same faces but with facial expressions, associated one-to-one with the M identities. Let $\boldsymbol{S}_m^\triangle = \boldsymbol{S}_m^E - \boldsymbol{S}_m^I$ be the expressive-neutral difference of face identity m, namely the expressive offset. A face $\boldsymbol{S}$ can be reconstructed from its identity and expression embeddings, i.e. Appendix B:

$$\hat{\boldsymbol{S}} = \overline{\boldsymbol{S}}^I + \overline{\boldsymbol{S}}^\triangle + \widetilde{\mathbf{U}}^I \boldsymbol{s}^I + \widetilde{\mathbf{U}}^\triangle \boldsymbol{s}^\triangle, \qquad (15)$$

where $\overline{\boldsymbol{S}}^I$ and $\overline{\boldsymbol{S}}^\triangle$ are the means associated with the identity- and with the expressive-neutral difference, matrices $\widetilde{\mathbf{U}}^I$ and $\widetilde{\mathbf{U}}^\triangle$ contain the K principle eigenvectors, and $\boldsymbol{s}^I$ and $\boldsymbol{s}^\triangle$ are the corresponding embeddings. Note that we use $\hat{\boldsymbol{S}}$ to emphasize that the reconstruction is an approximation of $\boldsymbol{S}$.

In practice we use the publicly available Basel Shape Model (BSM) [22] augmented with [23]. This provides $M = 200$ registered face scans in a frontal view and with neutral expressions, corresponding to M different identities, as well as M expressive scans of the same identities. Each scan is described by a triangulated mesh with an identical number of vertices, namely $N = 53490$. The dimension of the embedding is $K = 200$ and hence we have $K \ll 3N$. We use the landmark locations associated with the mean identity $\overline{\boldsymbol{S}}^I$ to compute the pose, i.e $\boldsymbol{Z}_{1:J}$.

V. Experiments

As with any methodological development, performance evaluation is extremely important. In this paper we make recourse to empirical evaluation based on a dataset with associated ground truth. Such a dataset should contain pairs of frontal and profile views of faces for a large number of subjects. Quantitative performance evaluation consists of computing a metric between an image obtained by face frontalization of a profile view of a subject, with an image containing a frontally-viewed face of the same subject. It is desirable that the profile and frontal images are simultaneously recorded with two synchronized cameras. Therefore, the proposed evaluation is based on image-to-image comparison. Several metrics were developed in the past for comparing two images, e.g. feature-based and pixel-based metrics. In this work we use the *zero-mean normalized cross correlation* (ZNCC) coefficient between two image regions, a measure that has successfully been used for stereo matching, e.g. [33], because it is invariant to differences in brightness and contrast between the two images, due to the normalization with respect to mean and standard deviation.

Let $R_f(h, v) \subset I_f$ be a region of size $H \times V$ whose center coincides with pixel location (h, v) of a frontalized image I_f. Similarly, let $R_t(h, v) \subset I_t$ be a region of the same size and whose center coincides with pixel location (h, v) of a ground-truth image I_t. The ZNCC coefficient between these two regions writes:

$$C(h, v, h', v') = \qquad (16)$$
$$\max_{\delta h, \delta v} \left\{ \frac{\text{Cov}\left[R_f(h, v), R_t(h + \delta h, v + \delta v)\right]}{\text{Var}\left[R_f(h, v)\right]^{1/2} \text{Var}\left[R_t(h + \delta h, v + \delta v)\right]^{1/2}} \right\},$$

TABLE I: Mean ZNCC coefficients for 15 participants. The best score is in **bold** and the second best is in *slanted bold*.

Method	Post-processing	ZNCC
Hassner et al. [18]	-	0.723
Hassner et al. [18]	Soft symmetry	0.780
Hassner et al. [18]	Hard symmetry	0.722
Banerjee et al. [2]	-	0.739
Banerjee et al. [2]	Soft symmetry	*0.788*
Proposed	-	**0.824**

TABLE II: Results for nine participants as a function of estimated yaw angle (in degrees) that corresponds to the horizontal head orientation computed with the proposed 3D head-pose estimator. Both [18] and [2] make use of symmetry to fill in the occluded face areas. The best scores are in **bold** and the second best are in *slanted bold*.

Participant	Yaw	[18]	[2]	Proposed
#31	19.1	*0.905*	0.856	**0.924**
#01	23.5	**0.915**	0.893	*0.913*
#02	24.9	*0.888*	0.878	**0.948**
#10	29.0	0.805	**0.812**	*0.806*
#23	30.0	0.810	**0.857**	*0.842*
#27	32.9	0.685	**0.852**	*0.732*
#19	37.8	*0.752*	0.650	**0.754**
#12	38.5	*0.731*	0.713	**0.774**
#21	40.6	0.632	**0.743**	*0.735*
Mean & std.		0.791±0.100	*0.806*±0.084	**0.825**±0.085

where $\text{Cov}\left[\cdot, \cdot\right]$ is the centered covariance between the two regions, $\text{Var}\left[\cdot\right]$ is the centered variance of a region, δh and δv are horizontal and vertical shifts, and h' and v' are the horizontal and vertical shifts that maximize the ZNCC coefficient.

In order to evaluate the performance of the proposed frontalization method and to compare it with state-of-the-art methods, we used a publicly available dataset, namely the OulouVS2 dataset [34]. This dataset targets the understanding of speech perception, more precisely, the analysis of non-rigid lip motions that are associated with speech production. The dataset was recorded in an office with ordinary (artificial and natural) lighting conditions. The recording setup consists of five synchronized cameras (2 MP, 30 FPS) placed in five different points of view, namely $0°$, $30°$, $45°$, $60°$ and $90°$.

The dataset contains 5×780 videos recorded with 53 participants. Each participant was instructed to read loudly several text sequences displayed on a computer monitor placed slightly to the left and behind the $0°$ (frontal) camera. The displayed text consists of digit sequences, e.g. "one, seven, three, zero, two, nine", of phrases, e.g. "thank you", "have a good time", and "you are welcome", as well as of sequences from the TIMIT dataset, e.g. "agricultural products are unevenly distributed". While participants were asked to keep their head still, natural uncontrolled head movements and body position changes were inevitable. As a consequence the actual head pose varies from one participant to another and there is no exact match between the head and camera orientations.

In practice, we evaluated the performance of the proposed method and we compared it with two state-of-the-art methods
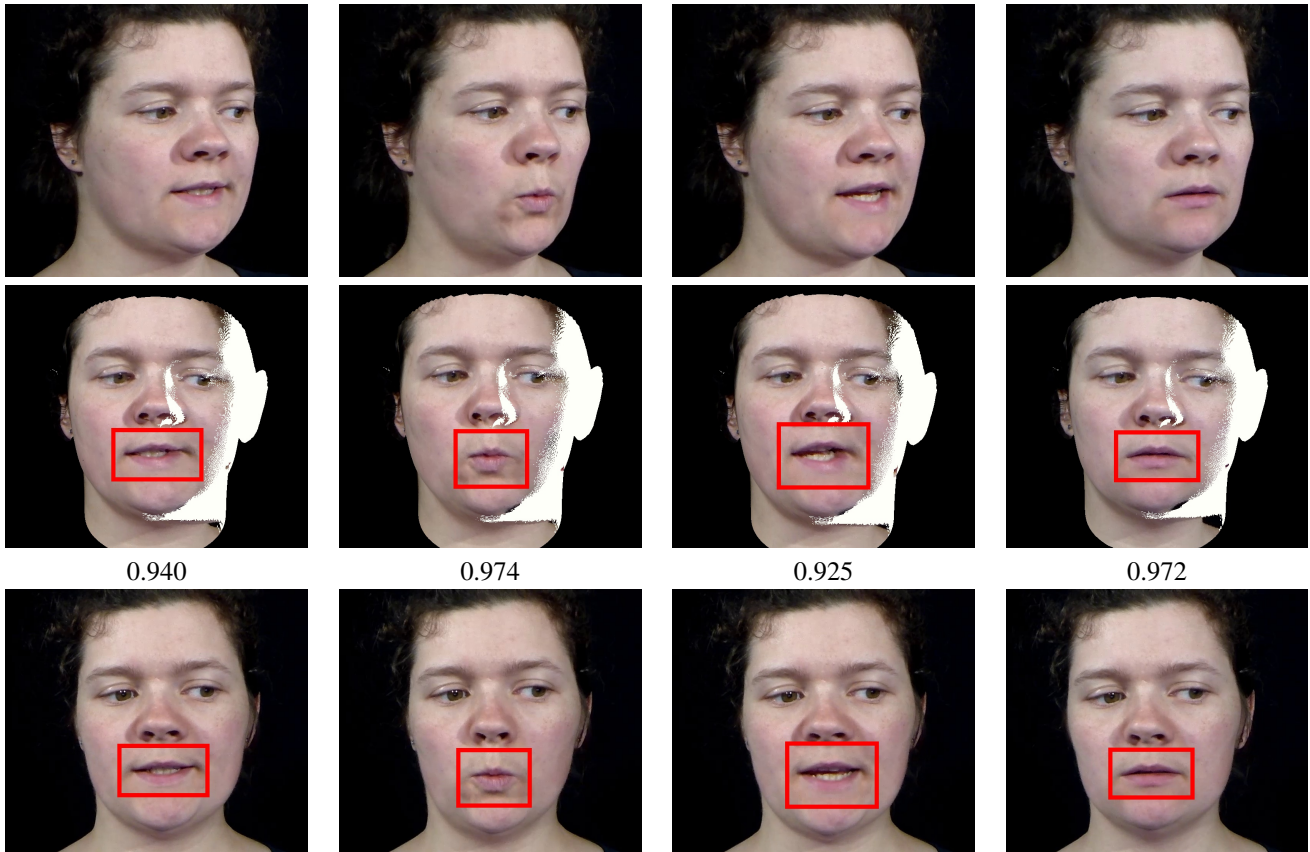
Fig. 2: Face frontalization exemples and ZNCC coefficients obtained with our method for participant #02 from the OulouVS2 dataset. Top row: input face recorded with the 30° camera. The estimated horizontal head orientation (yaw angle) is of 25° in this case. Middle row: frontalization results with self occluded facial regions displayed in white. Bottom row: ground-truth frontal face recorded with the 0° camera. The ZNCC coefficients correspond to the mouth bounding boxes shown in red.
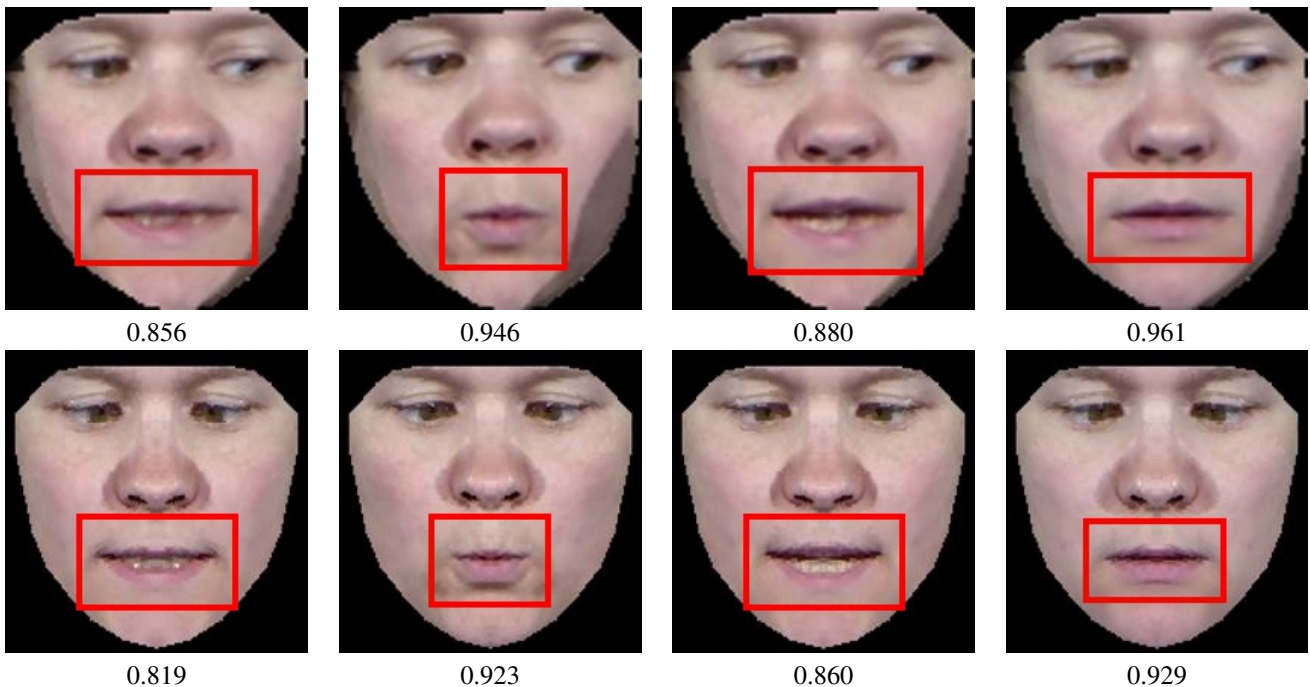

Fig. 3: Face frontalization results and ZNCC coefficients for participant #02. Top row: Hassner et al [18]. Bottom row: Banerjee et al. [2]. Both these methods enforce symmetry to fill in the gaps caused by self occlusions.

0.707 0.778 0.818 0.663

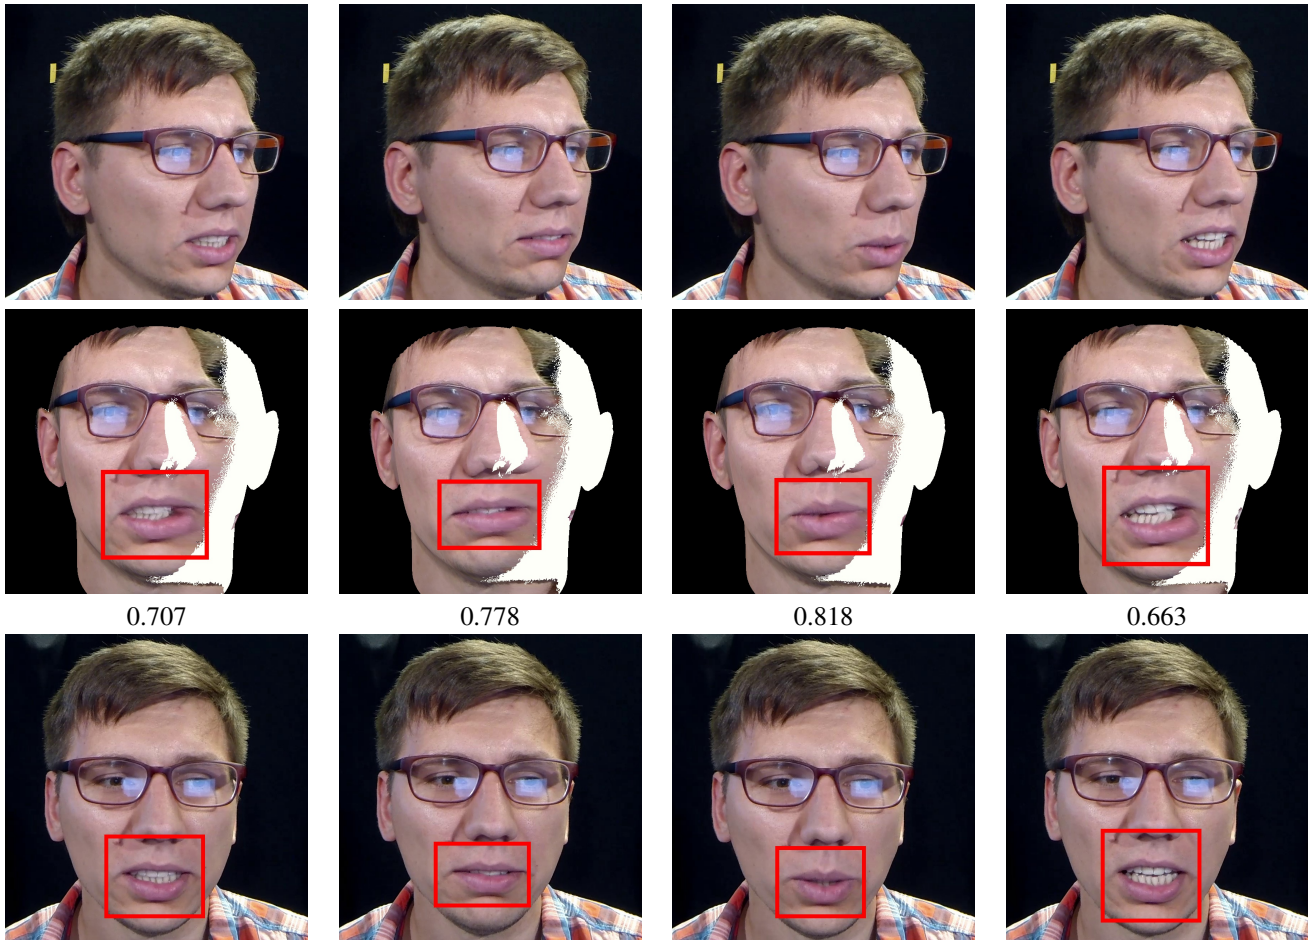Fig. 4: Same as Figure 2 for participant #21. The estimated horizontal head orientation (yaw angle) is of $40°$ in this case. Notice that in spite of the large yaw angle, the mouth area correctly frontalized with no self occluded regions.
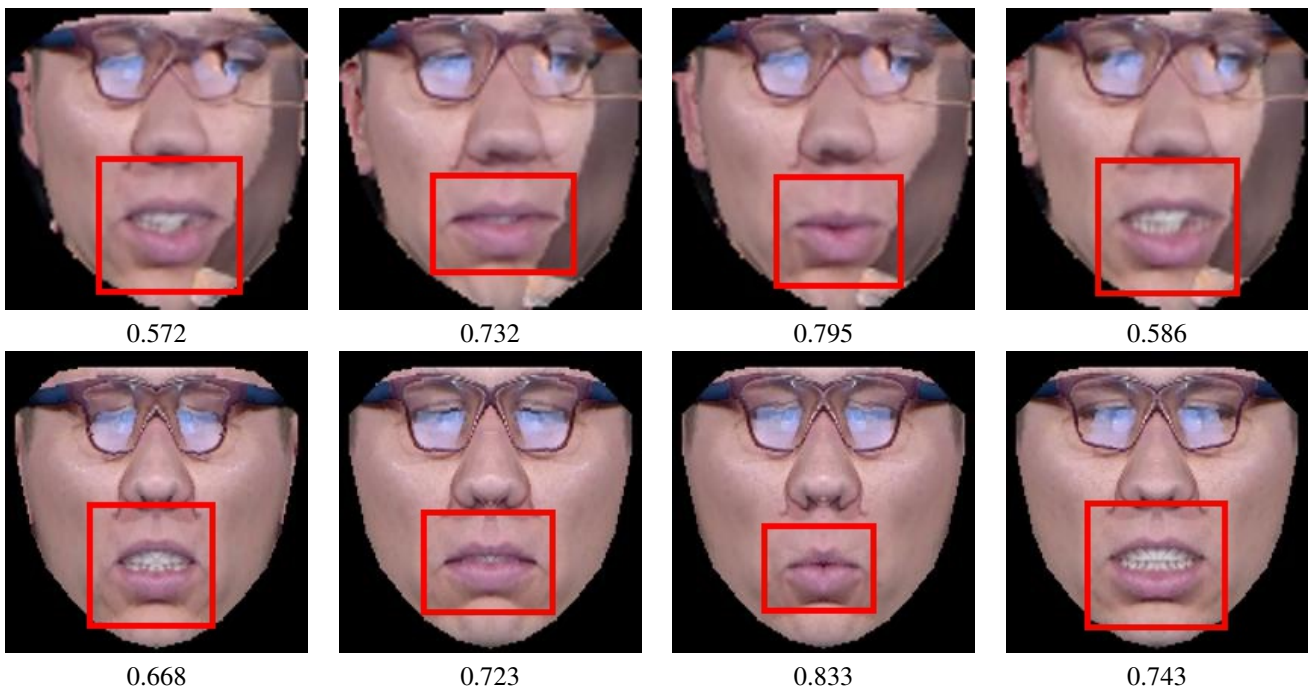

0.572 0.732 0.795 0.586

0.668 0.723 0.833 0.743

Fig. 5: Same as Figure 3 for participant #21. Note that in this example symmetry post-processing introduces important facial distorsions.

(a) Input image   (b) 0.403   (c) 0.739 (hard symmetry)   (d) ground truth
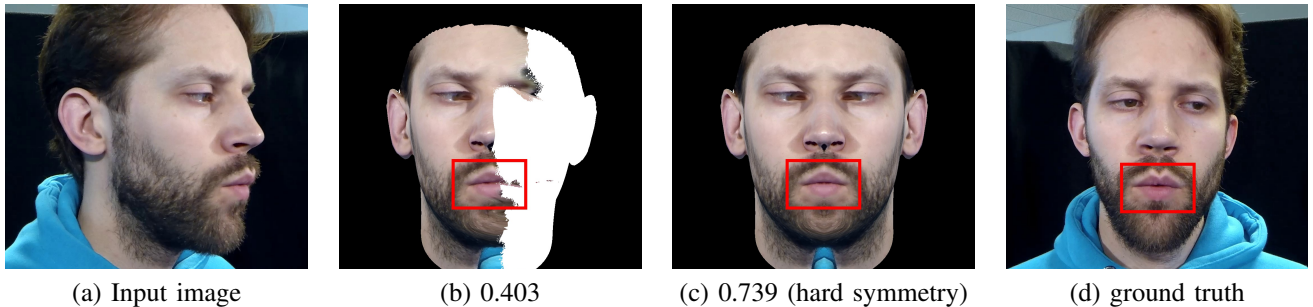
Fig. 6: An example of face frontalization result obtained with the $45°$ camera. (a) The estimated yaw angle is $57°$ in this case. (b) Because the lip area is half occluded, the ZNCC coefficient is below 0.5. (c) Exploiting hard symmetry (the occluded part is filled in by flipping the visible part around the face's mirror-symmetric axis). (d) Ground-truth provided by the $0°$ camera.

for which the code is publicly available, namely [2], [18]. We applied frontalization to images extracted from the videos recorded with the $30°$ camera and compared the results with the "ground-truth", namely the corresponding images extracted from the videos recorded with the $0°$ camera. Notice that videos recorded with higher viewing angles, i.e. $45°$, $60°$ and $90°$, can be hardly exploited by a frontalization algorithm because half of the face is occluded, e.g. Figure 6. For each frontalized image I_f we extract the mouth region R_f and we search in the ground-truth image I_t for the best-matching region R_t. This provides a ZNCC coefficient (16) for each test image. Notice that (16) only cares about the horizontal and vertical shifts in the image plane and assumes that the frontalized face and the corresponding ground-truth frontal face share the same scale. In practice, different frontalization algorithms output faces at different scales. For this reason and for the sake of fairness, prior to applying (16), we extract facial landmarks from both the frontalized and ground-truth faces and we use a subset of this set of landmarks to estimate the scale factor between the two faces.

We randomly selected 30 video pairs, recorded with the $30°$ and $0°$ cameras, of 15 participants from the OulouVS2 dataset. Each video contains approximatively 160 images, therefore we used a total of 4800 images in our benchmark. The mean ZNCC coefficients obtained with two state-of-the-art methods and with the proposed method are displayed in Table I. Both [18] and [2] exploit facial symmetry to fill in the occluded areas, which improves the ZNCC metrics. Indeed, facial symmetry can be used, either by replacing occluded pixels with the mirror-symmetric ones (soft symmetry), or by flipping the visible half of the face (hard symmetry).

We noticed that there were important discrepancies in method performance across participants. In order to better understand this phenomenon, we computed the mean ZNCC coefficient for nine participants and displayed these means and standard deviations as a function of the yaw angle (in degrees), i.e. the horizontal head orientation associated with the 3D head-pose estimated with the proposed method, please refer to Table II. One may notice that there is a wide range of yaw angles, from $19°$ and up to $40°$, and that the performance gracefully decreases as the yaw angle increases. Nevertheless, the proposed method yields the best scores for participants

#19 and #12 and the second best score for participant #21, and this is without using mirror-symmetric information.

Examples of face frontalization obtained with our method and with the methods of [18] and [2] are shown on Figure 2, Figure 2, Figure 3 and Figure 5. As already mentioned, both [18] and [2] enforce facial symmetry as a post-processing frontalization step to compensate for the gaps caused by self occlusions. We also show an example of applying frontalization to the $45°$ camera, Figure 6. In this example the estimated yaw angle is $57°$ and only half of the lip area is visible. Consequently, the ZNCC coefficient is below 0.5 in this case. Notice that when the hard-symmetry strategy of [18] is applied, the ZNCC coefficient increases considerably.

Finally, Figure 7 and Figure 8 show a few examples of frontalization results obtained with faces from the LFW dataset [35]. It is worthwhile to notice that our method yields frontal images of a better quality than the methods of [18] and [2], and that these two methods produce many artefacts and unrealistic facial deformations.

VI. CONCLUSIONS

In this paper we introduced a robust face frontalization (RFF) method that is based on the simultaneous estimation of the rigid transformation between two 3D point sets and the non-rigid deformation of a 3D face model. This is combined with pixel-to-pixel warping, between an input image of an arbitrarily-viewed face and a synthesized frontal view of the face. The proposed method yields state-of-the-art results, both quantitatively and qualitatively, when compared with two recently proposed methods. Up to now, the performance of face frontalization has been evaluated using face recognition benchmarks. We advocate that direct image-to-image comparison between the predicted output and the associated ground-truth yields a better assessment criterion that is not biased by another task, e.g. face recognition.

It is worthwhile to notice that the performance of face recognition can be boosted by combining face frontalization with the use of facial symmetry. However, the latter is likely to introduce undesired artefacts and unrealistic facial deformations, which are quite damaging for other face analysis tasks, such as expression recognition or lip reading.
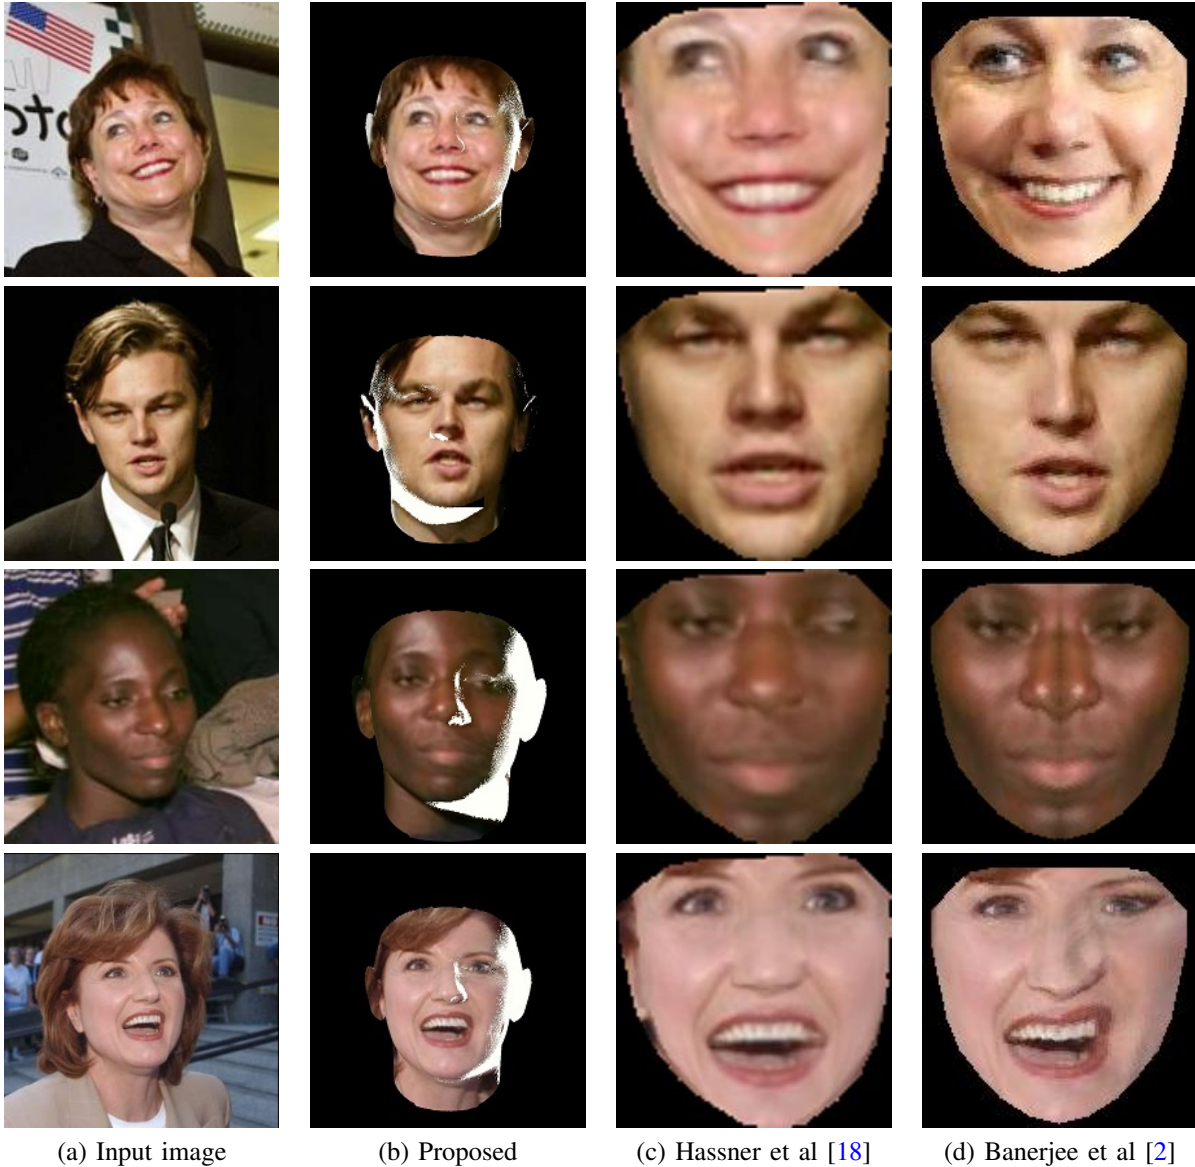
(a) Input image   (b) Proposed   (c) Hassner et al [18]   (d) Banerjee et al [2]

Fig. 7: Face frontalization examples from the Labeled Face in the Wild (LFW) dataset [35]. Both [18] and [2] make use of mirror symmetry, which introduces unrealistic facial deformations.

In the future, we plan to extend the proposed method to image sequences in order to allow robust temporal analysis of faces. In particular, we are interested in combining face frontalization with audio-visual speech enhancement. As already outlined, face frontalization may well be viewed as a process of discriminating between rigid head movements and non-rigid facial deformations, which can be used to eliminate head motions that naturally accompany speech production, and hence leverage the performance of visual speech reading and of audio-visual speech enhancement.

APPENDIX A
ROBUST ESTIMATION OF THE RIGID ALIGNMENT BETWEEN TWO 3D POINT SETS

In this appendix, we address the problem of robust alignement between two 3D point sets, $\mathcal{X}$ with coordinates

$\boldsymbol{X}_{1:J} = \{\boldsymbol{X}_j\}_{j=1}^J \subset \mathbb{R}^{3 \times J}$ and $\mathcal{Z}$ with coordinates $\boldsymbol{Z}_{1:J} = \{\boldsymbol{Z}_j\}_{j=1}^J \subset \mathbb{R}^{3 \times J}$, respectively. As mentioned in Section III the residuals (1) are samples of a random variable drawn from the generalized Student t-distribution [14]:

$$P(\boldsymbol{e}; \boldsymbol{\Sigma}, \mu, \nu) = \int_0^\infty \mathcal{N}(\boldsymbol{e}; 0, w^{-1}\boldsymbol{\Sigma})\mathcal{G}(w, \mu, \nu)dw$$

$$= \frac{\Gamma(\mu + \frac{3}{2})}{|\boldsymbol{\Sigma}|^{\frac{1}{2}}\Gamma(\mu)(2\pi\nu)^{\frac{3}{2}}}\left(1 + \frac{\|\boldsymbol{e}\|_{\mathcal{M}}^2}{2\nu}\right)^{-(\mu + \frac{3}{2})}, \quad (17)$$

where μ and ν are the parameters of the *prior gamma distribution* of the precision variable w, and $\Gamma(\cdot)$ is the gamma function. Without loss of generality, we set $\nu = 1$ [14]. Notice that using the conjugate prior property, the posterior distribution of w is also a gamma distribution, namely the

(a) Input image   (b) Proposed   (c) Proposed + soft symmetry   (d) Frontal image
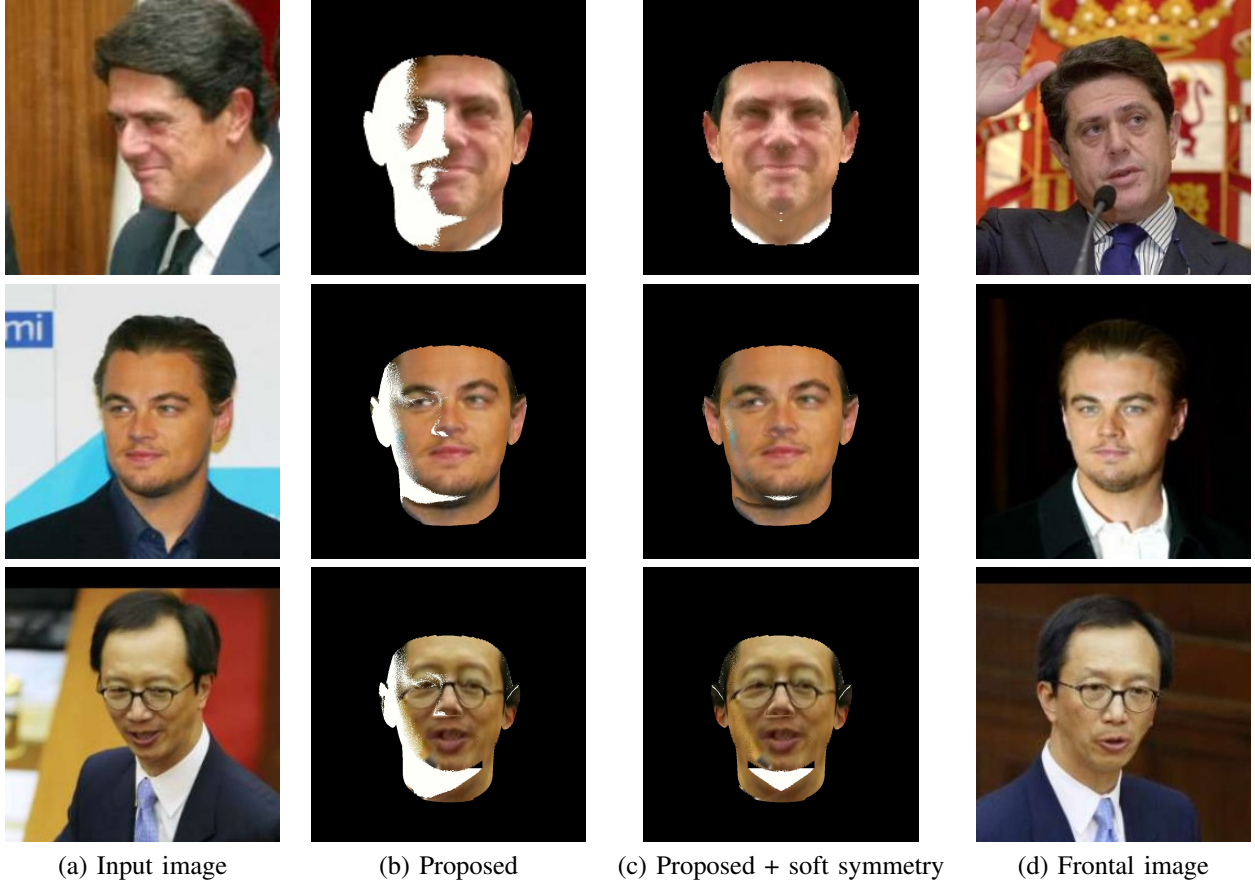
Fig. 8: Additional examples from LFW showing the visual impact of using soft symmetry in conjunction with the proposed method. The frontal image (d) is shown for qualitative visual comparison and it is not used as ground truth.

posterior gamma distribution:

$$P(w|\boldsymbol{e}; \boldsymbol{\Sigma}, \mu, \nu) = \mathcal{N}(\boldsymbol{e}; 0, w^{-1}\boldsymbol{\Sigma})\mathcal{G}(w, \mu, 1)$$
$$= \mathcal{G}(w; a, b), \quad (18)$$

with parameters:

$$a = \mu + \frac{3}{2}, \quad b = 1 + \frac{1}{2}\|\boldsymbol{e}\|_{\boldsymbol{\Sigma}}^2. \quad (19)$$

The posterior mean of the precision is:

$$\overline{w} = \mathrm{E}[w|\boldsymbol{e}] = \frac{a}{b}. \quad (20)$$

Direct minimization of the likelihood function is not practical. An expectation-maximization formulation is therefore adopted, namely the minimization of the *expected complete-data negative log-likelihood* (4), $\mathrm{E}_W[-\log P(\boldsymbol{e}_{1:J}, w_{1:J}|\boldsymbol{e}_{1:J}; \boldsymbol{\theta})]$ and in this case the parameter vector is a concatenation of the alignment and pdf parameters, namely $\boldsymbol{\theta} = \{\rho, \mathbf{R}, \boldsymbol{t}, \boldsymbol{\Sigma}, \mu\}$ since we set $\nu = 1$. This yields the minimization of:

$$\mathcal{Q}(\boldsymbol{\theta}) = \frac{1}{2} \sum_{j=1}^{J} \left(\overline{w}_j \|\boldsymbol{Y}_j - \rho\mathbf{R}\boldsymbol{X}_j - \boldsymbol{t}\|_{\boldsymbol{\Sigma}}^2 + \log|\boldsymbol{\Sigma}| \right), \quad (21)$$

By taking the derivatives of (21) with respect to the parameters and equating to zero, we obtain the following analytical expressions for optimal values of the following parameters:

$$\boldsymbol{t}^\star = \overline{\boldsymbol{Y}} - \rho\mathbf{R}\overline{\boldsymbol{X}}, \quad (22)$$

$$\rho^\star = \left(\frac{\sum_{j=1}^{J} \overline{w}_j \boldsymbol{Y}_j'^{\top} \boldsymbol{\Sigma}^{-1} \boldsymbol{Y}_j'}{\sum_{j=1}^{J} \overline{w}_j (\mathbf{R}\boldsymbol{X}_j')^{\top} \boldsymbol{\Sigma}^{-1}(\mathbf{R}\boldsymbol{X}_j')} \right)^{\frac{1}{2}}, \quad (23)$$

$$\boldsymbol{\Sigma}^\star = \frac{1}{N} \sum_{j=1}^{J} \overline{w}_j (\boldsymbol{Y}_j' - \rho\mathbf{R}\boldsymbol{X}_j')(\boldsymbol{Y}_j' - \rho\mathbf{R}\boldsymbol{X}_j')^{\top}, \quad (24)$$

with the following subsitutions:

$$\overline{\boldsymbol{X}} = \frac{\sum_{j=1}^{J} \overline{w}_j \boldsymbol{X}_j}{\sum_{j=1}^{N} \overline{w}_j}, \quad \boldsymbol{X}_j' = \boldsymbol{X}_j - \overline{\boldsymbol{X}}, \quad (25)$$

$$\overline{\boldsymbol{Y}} = \frac{\sum_{j=1}^{J} \overline{w}_j \boldsymbol{Y}_j}{\sum_{j=1}^{N} \overline{w}_j}, \quad \boldsymbol{Y}_j' = \boldsymbol{Y}_j - \overline{\boldsymbol{Y}}. \quad (26)$$

By substituting (22) in (21) and using the notations above, the rotation is estimated by minimization of:

$$\mathbf{R}^\star = \underset{\mathbf{R}}{\operatorname{argmin}} \ \frac{1}{2} \sum_{j=1}^{J} (\overline{w}_j \|\boldsymbol{Y}_j' - \rho\mathbf{R}\boldsymbol{X}_j'\|_{\boldsymbol{\Sigma}}^2). \quad (27)$$

The minimizer (27) yields a closed-form solution for an isotropic covariance, i.e. $\boldsymbol{\Sigma} = \sigma\mathbf{I}$. In the general case, however, one has to make recourse to nonlinear minimization. It is practical to represent the rotation with a unit quaternion and

to minimize (27) using a sequential quadratic programming method [31]. The parameter μ is updated by solving the following equation, where $\Psi(\cdot)$ is the digamma function:

$$\mu = \Psi^{-1}\left(\Psi(a) - \frac{1}{n}\sum_{j=1}^{J}\log b_j\right). \tag{28}$$

This yields Algorithm 2:

Algorithm 2: Robust estimation of the rigid transformation between two 3D point sets.

Data: Centered point coordinates $\boldsymbol{X}'_{1:J}$ and $\boldsymbol{Y}'_{1:J}$
Initialization of $\boldsymbol{\theta}^{\text{old}}$: Use a closed-form solution for the rotation to evaluate $\mathbf{R}^{\text{old}}$ and ρ^{old}; evaluate $\boldsymbol{\Sigma}^{\text{old}}$ using (24). Provide μ^{old} ;
while $\|\boldsymbol{\theta}^{\text{new}} - \boldsymbol{\theta}^{\text{old}}\| > \epsilon$ **do**
  E-step: evaluate a^{new} and $\boldsymbol{b}^{\text{new}}_{1:N}$ using (19) with $\boldsymbol{\theta}^{\text{old}}$, then evaluate $\overline{w}^{\text{new}}_{1:N}$ using (20) ;
  Update the centered coordinates ;
  M-scale-step: Evaluate ρ^{new} using (23);
  M-rotation-step: Estimate $\mathbf{R}^{\text{new}}$ via minimization of (27) ;
  M-covariance-step: Evaluate $\boldsymbol{\Sigma}^{\text{new}}$ using (24) ;
  M-mu-step: Evaluate μ^{new} using (28) ;
  $\boldsymbol{\theta}^{\text{old}} \leftarrow \boldsymbol{\theta}^{\text{new}}$;
end
Optimal translation: Evaluate the translation vector using (22);
Result: Optimal scale $\rho^{\star}$, rotation $\mathbf{R}^{\star}$, translation $\boldsymbol{t}^{\star}$, covariance $\boldsymbol{\Sigma}^{\star}$, and precisions $\boldsymbol{w}_{1:N}$.

APPENDIX B
DEFORMABLE SHAPE MODEL

In this appendix we summarize the deformable shape model that we use. The model is based on a training set $\mathcal{S} = \{\boldsymbol{S}_1, \ldots, \boldsymbol{S}_m, \ldots, \boldsymbol{S}_M\}$ of M 3D face scans. Each 3D face scan is described with a triangulated mesh composed of N 3D points, or vertices, $\boldsymbol{V}_{1:N} = \{\boldsymbol{V}_1, \ldots, \boldsymbol{V}_N\}$, where the coordinates of each vertex are $\boldsymbol{V}_n = (V_{n1}, V_{n2}, V_{n3})^\top$. Hence, a shape in $\mathcal{S}$ is described by the vector $\boldsymbol{S}$:

$$\boldsymbol{S} = (V_{11}, V_{12}, V_{13}, \ldots, V_{N1}, V_{N2}, V_{N3})^\top \in \mathcal{S} \subset \mathbb{R}^{3N}.$$

It is assumed that all the shapes are represented in the same coordinate frame, that they share the same number of vertices and that there is a one-to-one correspondence between their vertices, namely $\{\boldsymbol{V}_{nm} \leftrightarrow \boldsymbol{V}_{nm'}\}_{n=1}^N$, for any two shapes in the training set, $\boldsymbol{S}_m$ and $\boldsymbol{S}_{m'}$ and for each vertex index n. Notice, however, that the problem of finding a rigid alignment between M shapes based on point-to-point correspondences is not trivial because of the presence of non-rigid shape deformations and of large shape variabilities. In what follows we will treat each shape/face as a point set. The covariance matrix associated with these M registered shapes, $\{\boldsymbol{S}_m\}_{m=1}^M$, is:

$$\boldsymbol{\Sigma} = \frac{1}{M}\sum_{m=1}^{M}(\boldsymbol{S}_m - \overline{\boldsymbol{S}})(\boldsymbol{S}_m - \overline{\boldsymbol{S}})^\top, \tag{29}$$

where $\overline{\boldsymbol{S}} = (\overline{\boldsymbol{V}}_1^\top, \ldots, \overline{\boldsymbol{V}}_N^\top)^\top \in \mathbb{R}^{3N}$ is the mean shape with, $\overline{\boldsymbol{V}}_n = 1/M\sum_{m=1}^M \boldsymbol{V}_{nm}$. The covariance $\boldsymbol{\Sigma} \in \mathbb{R}^{3N \times 3N}$ is a symmetric semi-definite positive matrix, hence its eigendecomposition writes $\boldsymbol{\Sigma} = \mathbf{U}\boldsymbol{\Lambda}\mathbf{U}^\top$, where $\mathbf{U} = \begin{pmatrix} \boldsymbol{U}_1 & \ldots & \boldsymbol{U}_{3N} \end{pmatrix} \in \mathbb{R}^{3N \times 3N}$ is an orthogonal matrix and $\boldsymbol{\Lambda} = \text{Diag}\begin{pmatrix} \lambda_1, \ldots, \lambda_{3N} \end{pmatrix}$, with $\lambda_1 \geq \cdots \geq \lambda_{3N} \geq 0$. It is common practice to select the number K of *principal components*, or *modes*, such that the sum of the K largest eigenvalues represents 95% of the total variance, $\sum_{k=1}^K \lambda_k = 0.95 \sum_{n=1}^{3N} \lambda_n$. We introduce the following *truncated* matrices:

$$\widetilde{\mathbf{U}} = \begin{pmatrix} \boldsymbol{U}_1 & \ldots & \boldsymbol{U}_K \end{pmatrix} \in \mathbb{R}^{3N \times K},$$
$$\widetilde{\boldsymbol{\Lambda}} = \text{Diag}\begin{pmatrix} \lambda_1, \ldots, \lambda_K \end{pmatrix} \in \mathbb{R}^{K \times K},$$

with the property $\widetilde{\mathbf{U}}^\top \widetilde{\mathbf{U}} = \mathbf{I}_K$. Matrix $\widetilde{\mathbf{U}}^\top$ projects the centered shapes from $\mathbb{R}^{3N}$ onto $\mathbb{R}^K$:

$$\boldsymbol{s}_m = \widetilde{\mathbf{U}}^\top(\boldsymbol{S}_m - \overline{\boldsymbol{S}}), \tag{30}$$

where $\boldsymbol{s}_m \in \mathbb{R}^K$ is a low-dimensional embedding of $\boldsymbol{S}_m$. Conversely, it is possible to reconstruct an approximation of $\boldsymbol{S}_m$, denoted $\hat{\boldsymbol{S}}_m$, from its low-dimensional embedding:

$$\hat{\boldsymbol{S}}_m = \overline{\boldsymbol{S}} + \widetilde{\mathbf{U}}\boldsymbol{s}_m, \tag{31}$$

which is a concatenation of:

$$\hat{\boldsymbol{V}}_{nm} = \overline{\boldsymbol{V}}_n + \mathbf{W}_n \boldsymbol{s}_m, \tag{32}$$

where the N block-matrices $\mathbf{W}_n \in \mathbb{R}^{3 \times K}$ are such that:

$$\widetilde{\mathbf{U}} = \begin{pmatrix} \mathbf{W}_1 \\ \vdots \\ \mathbf{W}_n \\ \vdots \\ \mathbf{W}_N \end{pmatrix} \tag{33}$$

The concept of *statistical shape model* is summarized in Fig. 9. A training set of M shapes is embedded into a low-dimensional space spanned by the K principal eigenvectors of the covariance matrix (29) associated with the training set and is characterized by the K principal (largest) eigenvalues, where an eigenvalue λ_k corresponds to the variance along direction $\boldsymbol{u}_k$. Importantly, a *valid* shape $\hat{\boldsymbol{S}}$ is a shape reconstructed from an embedding $\boldsymbol{s}$ that lies inside an ellipsoid whose half axes are equal to $\sqrt{\lambda_k}$, or:

$$\hat{\boldsymbol{S}} = \overline{\boldsymbol{S}} + \widetilde{\mathbf{U}}\boldsymbol{s}, \text{ s.t. } \boldsymbol{s}^\top \widetilde{\boldsymbol{\Lambda}}^{-1}\boldsymbol{s} \leq 1. \tag{34}$$

This guarantees that the reconstructed shape belongs to the family of training shapes with 99% confidence.
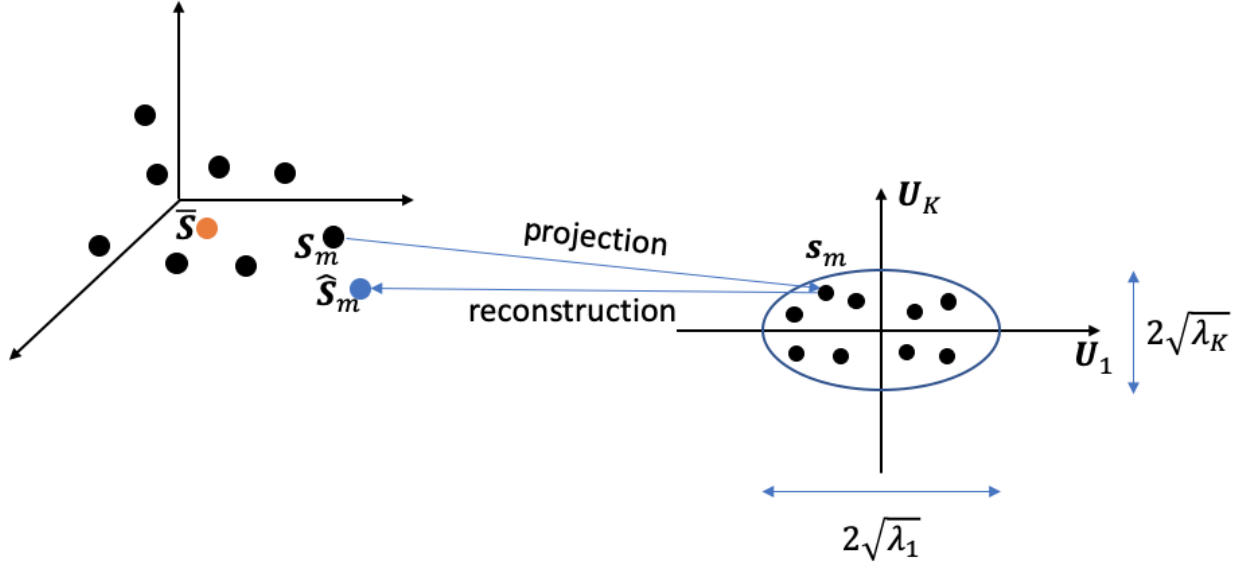
Fig. 9: A training set of M shapes (black dots), characterized by their mean (orange dot) and covariance, is projected onto the space spanned by the the principal eigenvectors of the shape covariance matrix, $\boldsymbol{U}_1$ to $\boldsymbol{U}_K$. The number K of principal eigenvectors is selected such that the sum of the associated eigenvalues represents 95% of the total variance. A shape $\boldsymbol{S}_m$ projects onto a shape embedding $\boldsymbol{s}_m$. Conversely, shape $\hat{\boldsymbol{S}}_m$ (blue dot), reconstructed from $\boldsymbol{s}_m$, approximates shape $\boldsymbol{S}_m$.

APPENDIX C
ROBUST ESTIMATION OF THE ALIGNMENT BETWEEN A DEFORMABLE SHAPE AND A 3D POINT SET

In this appendix we extend the robust alignment method of Appendix A to deformable shapes. The expression (21) can be re-written as:

$$\mathcal{Q}(\boldsymbol{\theta}) = \frac{1}{2}\sum_{j=1}^{J}\left(\overline{w}_j\|\boldsymbol{Y}_j - \sigma\mathbf{Q}(\overline{\boldsymbol{V}}_j + \mathbf{W}_j\boldsymbol{s}) - \boldsymbol{d}\|_{\boldsymbol{\Sigma}}^2 + \log|\boldsymbol{\Sigma}|\right) + \frac{\eta}{2}\boldsymbol{s}^\top\widetilde{\boldsymbol{\Lambda}}^{-1}\boldsymbol{s}, \tag{35}$$

where the statistical shape model is scaled, rotated, translated and deformed in order to be aligned with a set of landmarks, and where the last term constrains the embedding $\boldsymbol{s}$ to lie inside a 99% confidence ellipsoid, i.e. (34).

By taking the derivatives of (35) with respect to the parameters, i.e. $\boldsymbol{\theta} = \{\sigma, \mathbf{Q}, \boldsymbol{d}, \boldsymbol{\Sigma}, \boldsymbol{s}, \mu\}$, and equating to zero, we obtain the following analytical expressions for the optimal values of the following parameters:

$$\boldsymbol{d}^\star = \overline{\boldsymbol{Y}} - \sigma\mathbf{Q}\overline{\overline{\boldsymbol{V}}}, \tag{36}$$

$$\sigma^\star = \left(\frac{\sum_{j=1}^{J}\overline{w}_j\boldsymbol{Y}_j'^\top\boldsymbol{\Sigma}^{-1}\boldsymbol{Y}_j'}{\sum_{j=1}^{J}\overline{w}_j(\mathbf{Q}\hat{V}'_j)^\top\boldsymbol{\Sigma}^{-1}(\mathbf{Q}\hat{V}'_j)}\right)^{\frac{1}{2}}, \tag{37}$$

$$\boldsymbol{\Sigma}^\star = \frac{1}{N}\sum_{j=1}^{J}\overline{w}_j(\boldsymbol{Y}_j' - \sigma\mathbf{Q}\hat{V}'_j)(\boldsymbol{Y}_j' - \sigma\mathbf{Q}\hat{V}'_j)^\top, \tag{38}$$

where similar notations as in Appendix A and Appendix B were used, i.e (25), (26) and (31). An identical nonlinear optimization problem for the rotation is obtained as well:

$$\mathbf{Q}^\star = \underset{\mathbf{Q}}{\operatorname{argmin}}\ \frac{1}{2}\sum_{j=1}^{J}(\overline{w}_j\|\boldsymbol{Y}_j' - \rho\mathbf{Q}\hat{V}'_j\|_{\boldsymbol{\Sigma}}^2). \tag{39}$$

Additionally, there is a closed-form solution for the optimal embedding:

$$\boldsymbol{s}^\star = \left(\sum_{j=1}^{J}\overline{w}_j\mathbf{A}_j^\top\boldsymbol{\Sigma}^{-1}\mathbf{A}_j + \eta\widetilde{\boldsymbol{\Lambda}}^{-1}\right)^{-1}\left(\sum_{j=1}^{J}\overline{w}_j\mathbf{A}_j^\top\boldsymbol{\Sigma}^{-1}\boldsymbol{b}_j\right), \tag{40}$$

where the following notations were used: $\mathbf{A}_j = \sigma\mathbf{Q}\mathbf{W}_j$ and $\boldsymbol{b}_j = \boldsymbol{Y}_j - \sigma\mathbf{Q}\overline{\boldsymbol{V}}_j - \boldsymbol{d}$. The associated EM algorithm is similar to Algorithm 2; the most notable difference is that the M-step alternates between the estimation of the rigid parameters and the shape parameters, namely Algorithm 3.

REFERENCES

[1] J. Yim, H. Jung, B. Yoo, C. Choi, D. Park, and J. Kim, "Rotating your face using multi-task deep neural network," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 676–684.

[2] S. Banerjee, J. Brogan, J. Krizaj, A. Bharati, B. R. Webster, V. Struc, P. J. Flynn, and W. J. Scheirer, "To frontalize or not to frontalize: Do we really need elaborate pre-processing to improve face recognition?" in *IEEE Winter Conference on Applications of Computer Vision*, 2018, pp. 20–29.

[3] J. Zhao, Y. Cheng, Y. Xu, L. Xiong, J. Li, F. Zhao, K. Jayashree, S. Pranata, S. Shen, J. Xing *et al.*, "Towards pose invariant face recognition in the wild," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 2207–2216.

Algorithm 3: Robust fitting of a deformable shape model to a 3D point set.

Data: Centered landmark coordinates $\boldsymbol{Y}'_{1:J}$, mean shape $\overline{\boldsymbol{V}}_{1:J}$ and shape reconstruction matrices $\mathbf{W}_{1:J}$

Initialization of $\boldsymbol{\theta}^{\text{old}}$: Set $\boldsymbol{s}^{\text{old}} = \mathbf{0}$, use the closed-form solution [30] to evaluate $\mathbf{Q}^{\text{old}}$ and σ^{old}; evaluate $\boldsymbol{\Sigma}^{\text{old}}$ using (38). Provide μ^{old}. Compute the centered vertex coordinates $\hat{\boldsymbol{V}}'_{1:J}$;

while $\|\boldsymbol{\theta}^{\text{new}} - \boldsymbol{\theta}^{\text{old}}\| > \epsilon$ **do**

> **E-step**: evaluate a^{new} and $\boldsymbol{b}^{\text{new}}_{1:N}$ using (19) with $\boldsymbol{\theta}^{\text{old}}$, then evaluate $\overline{\boldsymbol{w}}^{\text{new}}_{1:N}$ using (20) ;
> Update the centered coordinates ;
> **M-scale-step**: Evaluate σ^{new} using (37);
> **M-rotation-step**: Estimate $\mathbf{Q}^{\text{new}}$ via minimization of (39) ;
> **M-covariance-step**: Evaluate $\boldsymbol{\Sigma}^{\text{new}}$ using (38) ;
> **M-shape-step**: Evaluate $\boldsymbol{s}^{\text{new}}$ using (40) ;
> **M-mu-step**: Evaluate μ^{new} using (28) ;
> $\boldsymbol{\theta}^{\text{old}} \leftarrow \boldsymbol{\theta}^{\text{new}}$;

end

Optimal translation: Evaluate the translation vector using (36);

Result: Optimal scale $\sigma^\star$, rotation $\mathbf{Q}^\star$, translation $\boldsymbol{d}^\star$, covariance $\boldsymbol{\Sigma}^\star$, and precisions $\boldsymbol{w}_{1:N}$.

---

[4] E. Pei, M. C. Oveneke, Y. Zhao, D. Jiang, and H. Sahli, "Monocular 3d facial expression features for continuous affect recognition," *IEEE Transactions on Multimedia*, 2020.

[5] A. Fernandez-Lopez and F. M. Sukno, "Survey on automatic lip-reading in the era of deep learning," *Image and Vision Computing*, vol. 78, pp. 53–72, 2018.

[6] A. Adeel, M. Gogate, A. Hussain, and W. M. Whitmer, "Lip-reading driven deep learning approach for speech enhancement," *IEEE Transactions on Emerging Topics in Computational Intelligence*, 2019.

[7] T. Schultz, M. Wand, T. Hueber, D. J. Krusienski, C. Herff, and J. S. Brumberg, "Biosignal-based spoken communication: A survey," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 12, pp. 2257–2271, 2017.

[8] S. Dupont and J. Luettin, "Audio-visual speech modeling for continuous speech recognition," *IEEE Transactions on Multimedia*, vol. 2, no. 3, pp. 141–151, 2000.

[9] B. Rivet, L. Girin, and C. Jutten, "Mixing audiovisual speech processing and blind source separation for the extraction of speech signals from convolutive mixtures," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 1, pp. 96–108, 2006.

[10] P. Wu, H. Liu, X. Li, T. Fan, and X. Zhang, "A novel lip descriptor for audio-visual keyword spotting based on adaptive decision fusion," *IEEE Transactions on Multimedia*, vol. 18, no. 3, pp. 326–338, 2016.

[11] M. Sadeghi, S. Leglaive, X. Alameda-Pineda, L. Girin, and R. Horaud, "Audio-visual speech enhancement using conditional variational autoencoders," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 1788–1800, 2020.

[12] F. Tao and C. Busso, "End-to-end audiovisual speech recognition system with multitask learning," *IEEE Transactions on Multimedia*, 2020.

[13] J. Sun, A. Kabán, and J. M. Garibaldi, "Robust mixture clustering using Pearson type VII distribution," *Pattern Recognition Letters*, vol. 31, no. 16, pp. 2447–2454, 2010.

[14] F. Forbes and D. Wraith, "A new family of multivariate heavy-tailed distributions with variable marginal amounts of tailweight: application to robust clustering," *Statistics and Computing*, vol. 24, no. 6, pp. 971–984, 2014.

[15] R. Huang, S. Zhang, T. Li, and R. He, "Beyond face rotation: Global and local perception GAN for photorealistic and identity preserving frontal view synthesis," in *IEEE International Conference on Computer Vision*, 2017, pp. 2439–2448.

[16] L. Tran, X. Yin, and X. Liu, "Disentangled representation learning gan for pose-invariant face recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 1415–1424.

[17] C. Rong, X. Zhang, and Y. Lin, "Feature-improving generative adversarial network for face frontalization," *IEEE Access*, vol. 8, pp. 68842–68851, 2020.

[18] T. Hassner, S. Harel, E. Paz, and R. Enbar, "Effective face frontalization in unconstrained images," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 4295–4304.

[19] R. Horaud, F. Dornaika, S. Christy, and B. Lamiroy, "Object pose: The link between weak perspective, paraperspective, and full perspective," *International Journal of Computer Vision*, vol. 22, no. 2, pp. 173–189, 1997.

[20] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham, "Active shape models – their training and application," *Computer Vision and Image Understanding*, vol. 61, no. 1, pp. 38–59, 1995.

[21] V. Blanz and T. Vetter, "A morphable model for the synthesis of 3d faces." in *ACM SIGGRAPH*, vol. 99, 1999, pp. 187–194.

[22] P. Paysan, R. Knothe, B. Amberg, S. Romdhani, and T. Vetter, "A 3d face model for pose and illumination invariant face recognition," in *IEEE International Conference on Advanced Video and Signal Based Surveillance*, 2009, pp. 296–301.

[23] C. Cao, Y. Weng, S. Zhou, Y. Tong, and K. Zhou, "Facewarehouse: A 3d facial expression database for visual computing," *IEEE Transactions on Visualization and Computer Graphics*, vol. 20, no. 3, pp. 413–425, 2013.

[24] A. Bulat and G. Tzimiropoulos, "How far are we from solving the 2D & 3D face alignment problem? (and a dataset of 230,000 3D facial landmarks)," in *IEEE International Conference on Computer Vision*, 2017, pp. 1021–1030.

[25] Y. Feng, F. Wu, X. Shao, Y. Wang, and X. Zhou, "Joint 3D face reconstruction and dense alignment with position map regression network," in *European Conference on Computer Vision*, 2018, pp. 534–551.

[26] X. Zhu, X. Liu, Z. Lei, and S. Z. Li, "Face alignment in full pose range: A 3d total solution," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 41, no. 1, pp. 78–92, 2019.

[27] M. Sadeghi, S. Guy, A. Raison, X. Alameda-Pineda, and R. Horaud, "Unsupervised performance analysis of 3d face alignment," *arXiv preprint arXiv:2004.06550*, 2020.

[28] Y. Wang, H. Yu, J. Dong, B. Stevens, and H. Liu, "Facial expression-aware face frontalization," in *Asian Conference on Computer Vision*. Springer, 2016, pp. 375–388.

[29] X. Zhu, Z. Lei, J. Yan, D. Yi, and S. Z. Li, "High-fidelity pose and expression normalization for face recognition in the wild," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 787–796.

[30] B. K. Horn, "Closed-form solution of absolute orientation using unit quaternions," *Journal of the Optical Society of America A*, vol. 4, no. 4, pp. 629–642, 1987.

[31] J.-F. Bonnans, J. C. Gilbert, C. Lemaréchal, and C. A. Sagastizábal, *Numerical optimization: theoretical and practical aspects*. Springer Science & Business Media, 2006.

[32] D. Kraft, "A software package for sequential quadratic programming," DLR German Aerospace Center – Institute for Flight Mechanics, Koln, Germany, Tech. Rep. DFVLR-FB 88-28, 1988.

[33] C. Sun, "Fast stereo matching using rectangular subregioning and 3D maximum-surface techniques," *International Journal of Computer Vision*, vol. 47, no. 1-3, pp. 99–117, 2002.

[34] I. Anina, Z. Zhou, G. Zhao, and M. Pietikäinen, "OuluVS2: a multi-view audiovisual database for non-rigid mouth motion analysis," in *International Conference on Automatic Face and Gesture Recognition*, vol. 1. IEEE, 2015, pp. 1–5.

[35] G. B. Huang, M. Mattar, T. Berg, and E. Learned-Miller, "Labeled faces in the wild: A database for studying face recognition in unconstrained environments," 2008.