

# Supplementary Material for CVPR’16 Paper: “A 3D Morphable Model learnt from 10,000 faces”

James Booth      Anastasios Roussos      Stefanos Zafeiriou  
Allan Ponniah      David Dunaway

## 1 Visualization of LSFM-global model (video)

As an extension of Figure 4 in the main paper, we present a short video showcasing the variation present in the first 5 principal components of the shape variation of LSFM-global model: please refer to the video file “LSFM\_components.mp4”.

## 2 Visualizations of LSFM-bespoke models

Figure 1 shows additional visualizations of the first 5 components of the 5 different bespoke models produced. We note that we isolate the inner facial regions in this experiment in order to showcase changes in facial identity more clearly.

On these results we note two observations. Firstly, ethnic traits can be seen to differ between the ethnically trained models. Whilst this is perhaps unsurprising, this is the first time it has been shown that ethnic traits are statistically dominant in facial appearance, appearing in the first 5 modes of variation in population-specific models.

The second observation is made with regards to age. Despite the fact that the four LSFM-bespoke models for White ethnicity of different age groups are built from completely disjoint datasets, the principal components present in each model are clearly interrelated. Of great interest is the fact that these individual principal components can be seen to age between the LSFM-bespoke models of different ages.

## 3 LSFM-global out-of-sample reconstructions

To showcase the power of the LSFM-global model, faces from the BU3D-FE [3] database were placed into correspondence with the model and then projected onto the subspace of the model. Figure 2 shows two original subjects from BU3D-FE and their LSFM-global reconstructions. As can be seen, LSFM-global achieves excellent detail retention, and does a commendable job of representing these out-of-sample subjects from a lower resolution dataset.

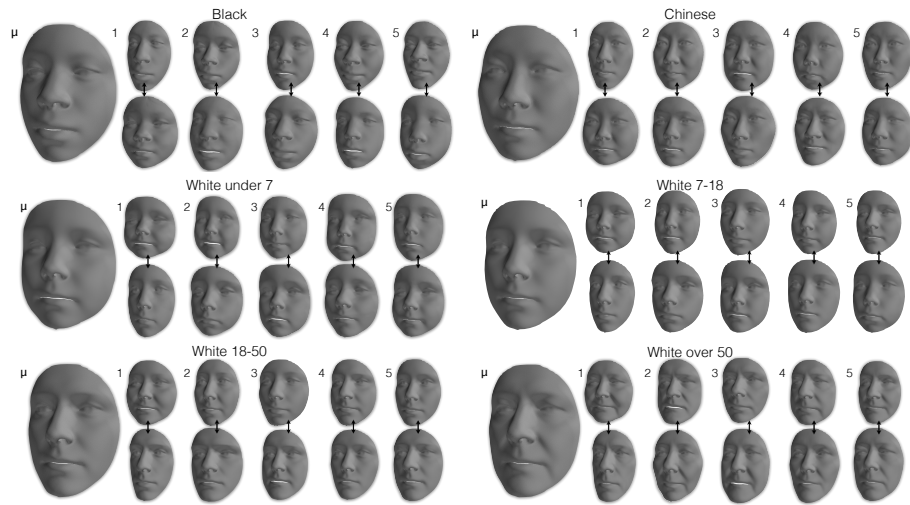


Figure 1: Visualization of the mean and first 5 principal components for the 5 LSFM-bespoke models.




Figure 2: Two subjects from the BU3D-FE dataset, along with their reconstructions inside the LSFM-global subspace.

## 4 Automatic error pruning

In this section, we provide more details about the automatic error pruning (Section 5.2 of the main paper) of our 3DMM construction pipeline.

Adopting a commonly-used probabilistic interpretation of the PCA model of shape variation (Eq. (1) of the main submission) [1], we assume that the shape parameters  $\alpha_1, \dots, \alpha_d$  are independent random variables and that each  $\alpha_i$  follows a Gaussian distribution with zero mean and variance  $\lambda_i$ , where  $\lambda_i$  is the  $i$ -th PCA eigenvalue (i.e. the  $i$ -th eigenvalue of the training data covariance matrix).

Therefore, the normalized shape parameters  $\frac{\alpha_1}{\sqrt{\lambda_1}}, \dots, \frac{\alpha_d}{\sqrt{\lambda_d}}$  are independent and identically distributed following a zero-mean and unit-variance Gaussian distribution and their squared sum, which can be written as:

$$F(\alpha) = \sum_{i=1}^d \frac{\alpha_i^2}{\lambda_i} \quad (1)$$

follows a chi-square distribution with  $d$  degrees of freedom [2]. The above sum is actually a weighted norm of the shape vector  $\alpha$  and yields a squared Mahalanobis distance between the current shape and the mean shape. This can be used as a measure of plausibility of the shape with shape parameters  $\alpha$ , under the current PCA model.

Based on the aforementioned remarks, for every training face mesh that has been put in correspondence using NICP and afterwards subjected in Procrustes alignment, we find its shape parameters  $\alpha$  by projecting on the initial global PCA model. Then, we use the squared norm  $F(\alpha)$  as the criterion to detect failures of the dense correspondence estimation. This is due to the fact that these failures behave as outliers of the Gaussian distribution.

We classify as outliers all shape vectors  $\alpha$  with a squared norm  $F(\alpha)$  above a threshold  $\theta_f$ . This threshold is selected so that only a small portion (e.g. 1.5% as used in the reported experiments) of the training data yields a value  $F(\alpha)$  that is greater than  $\theta_f$ . This means that  $F(\alpha)$  is expected to be less than  $\theta_f$  with a very high probability. Note that the set of shape vectors  $\alpha$  with  $F(\alpha) < \theta_f$  corresponds to a hyper-ellipse in the  $d$ -dimensional space of shape parameters.

## References

- [1] R. Davies, C. Taylor, et al. *Statistical models of shape: Optimisation and evaluation*. Springer Science & Business Media, 2008. 3
- [2] A. Patel and W. A. Smith. 3d morphable face models revisited. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 1327–1334. IEEE, 2009. 3
- [3] L. Yin, X. Wei, Y. Sun, J. Wang, and M. J. Rosato. A 3d facial expression database for facial behavior research. In *Automatic face and gesture recognition, 2006. FGR 2006. 7th international conference on*, pages 211–216. IEEE, 2006. 1