Robust Face Matching Under Large Occlusions

Georgy Gimel'farb, Patrice Delmas, John Morris, and Al Shorin
Department of Computer Science, The University of Auckland, Auckland 1000, New Zealand

Abstract

Outliers due to occlusions and contrast and offset signal deviations notably hinder recognition and retrieval of facial images. We propose a new maximum likelihood matching score with "soft masking" of outliers which is robust in these conditions. Differences between two images are modelled by unknown contrast and offset deviations from an unknown template and by independent pixel-wise errors. The error distribution is a mixture of a zero-centred Gaussian noise with an unknown variance and uniformly distributed outliers. The matching score combines the maximum likelihood estimates of model parameters and the soft masks being produced by a simple iterative Expectation-Maximisation algorithm. Experiments with facial images from the MIT Face Database show the robustness of this technique in the presence of large occlusions.

1. Introduction

The main challenges in face recognition and retrieval are associated with considerable deviations between facial images of the same person. This is often due to varying observation position, orientation, and resolution, different illumination, partial occlusions (e.g. by hair, glasses, scarves), and so on. Thus robust best image matching has to allow for these deviations in the matching score and effectively eliminate them. However, most local deviations - especially caused by occlusions - cannot be easily eliminated because they result in signal outliers that do not conform to conventional formal models of image deviations.

Outliers adversely affect the accuracy of image recognition. Consider, for example, a subset of the MIT face database [8] in Appendix A. To show that classification invariant to uniform contrast and offset deviations in facial images fails or becomes unstable if the deviations are combined with large occlusions, we took an image from the MIT Face Database and formed two variants by changing the contrast and offset values: image DV1 with reduced contrast - the intensity range was reduced from 0-255 to xxx-yyy - and image DV1 with offset altered from 0 to zzz. These two images were further altered by application of a large mask (mimicking a large occlusion) to produce images DVx-L (lower mask) and DVx-U (upper mask), see Figure 1.

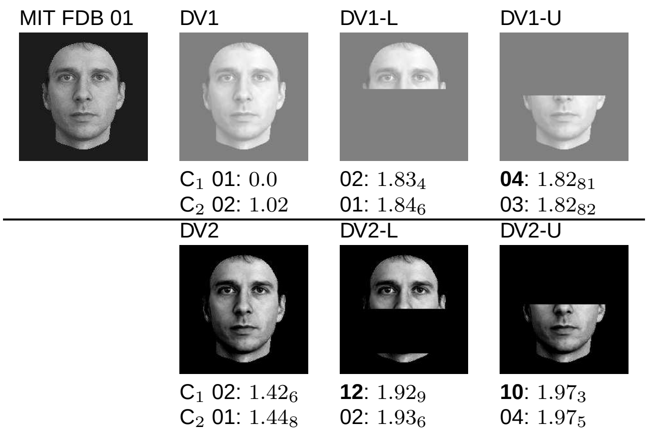

Figure 1. Best ($C_1 \times 10^{-5}$) and second best ($C_2 \times 10^{-5}$) matches to the MIT Face Database subset (App. A) based on scores D^{c}_{12} (see Section 2) for images with altered contrast and offset (DV1, DV2) and large occlusions (DVx-L, DVx-U). Matching errors are bold-faced. Note that a match to a different image of the same face is not considered an error.

Of course, outlier definition and characterisation depends on a particular image recognition or retrieval problem. For instance, not only visual occlusions can be treated as outliers. As well, it would be sensible to exclude background from matching because it is likely to be considerably different in facial images of the same person. Many face databases simply crop images to completely eliminate backgrounds as in the MIT database [8] in Fig. A6. However, backgrounds could also be considered as outliers and eliminated in the same manner as occlusions.

In order to focus on the robustness to outliers, we only considered the specific problem of pixel-to-pixel matching under relative translations, spatially uniform contrast and offset deviations and occlusions. Today's robust image matching [1, 7, 10] replaces pixel intensities (grey values or colour components) with local features (e.g. normalised signal gradients) which are more stable to contrast and offset variations and exploits mostly statistical M-estimators [6, 9] to derive a matching score. M-estimators depends less on outliers than least-square estimators as the error function has a less than than quadratic growth for large error values. However robustness of M-estimator-based scores to outliers is obtained at the cost of complex numerical optimisation for matching images under relative contrast and offset distortions compared to least-square-based scores. Because spatially uniform contrast and offset distortions cannot be analytically eliminated, M-estimator-based scores have more local optima in the parameter space that complicate the numerical search for the best match.

We preserve the maximum likelihood framework leading to the least-square image matching in the absence of outliers but attempt to eliminate outliers by *soft masking* of 'suspicious' pixels. To do so each pixel is weighed by a mask entry ranging from 0 (an outlier) to 1 ('pure' noise). The matching errors are modelled with a mixture of the probability distributions of random noise and outliers. The masks are estimated with an Expectation-Maximisation (EM) algorithm similar to the model identification schemes in [2, 3]. Because most of the unknown model parameters have analytical estimates, the resulting iterative maximum likelihood matching is faster and more flexible than its more conventional M-estimator-based counterparts. Moreover, this approach could be in principle extended to spatially variant relative contrast / offset deviations between the images (although the latter are eliminated in this case only numerically using an appropriate quadratic programming technique).

The paper is organised as follows. Section 2 presents a symmetric maximum likelihood matching score to measure dissimilarity between two noisy images under spatially uniform contrast and offset deviations. Then a EM-based score robust in the presence of outliers is derived from it. Experiments describing the validity of the proposed matching score in the case of face recognition or retrieval are given in Section 3.

2. Symmetric maximum likelihood matching

Matching of images distorted with different contrast (gain) and offset values and random noise leads to the following probability models of images. Let $g : \mathbf{R} \rightarrow \mathbf{Q}$ denote a digital image on a finite 2D lattice $\mathbf{R}$ with values (grey levels ot colour indices) from a signal set $\mathbf{Q}$. Each image is considered as the vector, $\mathbf{g} = [g_1, \ldots, g_p]$, where g_i is the signal intensity for the pixel i. Let $\mathbf{g}_j$ be an image derived from an (unknown) noiseless template $\mathbf{g} = (g_i : i = 1, \ldots, p)$ by arbitrary uniform contrast (a_j) and offset (b_j) deviations of signals and independent errors caused either by random noise or outliers. For the non-outliers:

$$g_{ji} = a_j g_i + b_j + \varepsilon_{ji};\ i \in \{1, \ldots, p\}, \quad (1)$$

where the errors, ε_{ji}, have a centred normal distribution with the same variance. The outliers are uniformly distributed and bear no relation to the template transformations.

Let, for simplicity, the per pixel error probability model in Eq. (1) be a mixture of the normal noise distribution $\mathcal{N}(\varepsilon|\sigma) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(\frac{-\varepsilon^2}{2\sigma^2}\right)$ with an unknown variance σ^2 and the outlier distribution $\mathcal{U}(\varepsilon)$:

$$\Pr(\varepsilon) = \rho\mathcal{N}(\varepsilon|\sigma) + (1 - \rho)\mathcal{U}(\varepsilon) \quad (2)$$

where ρ is an unknown prior probability of the noise.

Providing the errors are signed integers, $\varepsilon \in \{-Q + 1, \ldots, Q - 1\}$, where Q is the number of image grey values (typically $Q = 256$), the distribution $\mathcal{N}(\varepsilon|\sigma)$ is derived from the zero-centred normal distribution of noise:

$$\mathcal{N}(\varepsilon|\sigma) = \frac{1}{Z_\sigma} \exp\left(\frac{-\varepsilon^2}{2\sigma^2}\right) \quad (3)$$

where $Z_\sigma = \sum_{\delta=-Q+1}^{Q-1} \exp\left(\frac{-\delta^2}{2\sigma^2}\right)$. Generally, the distributions of both noise and outliers depend on the images of interest. In some cases, the error distribution for outliers can be associated with an empirical marginal distribution of all differences between the image signals [4]. However this model does not hold for many image types, in particular, for the facial images from the MIT database in Fig. A6 where the artificial black ($g_i = 0$) background unduly overemphasises zero matching errors for outliers.

2.1. Uniform contrast and offset deviations

If the model in Eq. (1) accounts for only the noise with no outliers, the maximum likelihood matching score, *leading to the minimum log-likelihood score* $D^{\mathrm{c}}_{12} = \min_{\boldsymbol{\theta}} (-\ln \Pr(\mathbf{g}_1, \mathbf{g}_2))$, measures the signal dissimilarity after replacing the unknown parameters, $\boldsymbol{\theta} = (a_1, b_1, a_2, b_2, \mathbf{g}, \sigma^2)$, with the most likely estimates (denoted by "hat" below):

$$D^{\mathrm{c}}_{12} = -\sum_{i=1}^{p} \ln \mathcal{N}(\hat{\varepsilon}_i) = p + p \ln \hat{Z}^{\mathrm{c}}_\sigma \quad (4)$$

where $\hat{Z}^{\mathrm{c}}_{\sigma} = \frac{\pi}{p}\Phi_{12}$ and $\Phi_{12} = \min\limits_{\boldsymbol{\theta}} \sum\limits_{i=1}^{p} \left(\varepsilon_{1i}^2 + \varepsilon_{2i}^2\right)$, that is,

$$\begin{aligned}\Phi_{12} &= \min_{\boldsymbol{\theta}} \sum_{i=1}^{p} \left[\left(\widetilde{g_1}_i - a_1 g_i - b_1\right)^2 + \left(\widetilde{g_2}_i - a_2 g_i - b_2\right)^2 \right] \\ &= \sum_{i=1}^{p} \hat{\varepsilon}_i^2 = \frac{1}{2}\left[S_{11} + S_{22} - \sqrt{\left(S_{11} - S_{22}\right)^2 + 4S_{12}^2} \right]\end{aligned}$$

where $S_{jk} = \sum_{i=1}^{p}(g_{ji} - \mu_j)(g_{ki} - \mu_k)$; $j,k = 1,2$, is the non-normalised signal covariance, $\mu_j = \frac{1}{\nu}\sum_{i=1}^{p} g_{ji}$; j=1,2, the mean signal value, and $\hat{\varepsilon}_i$ is the estimated mutual residual noise:

$$\begin{aligned}\hat{\varepsilon}_i &= \alpha_2 \left(g_{1i} - \mu_1\right) - \alpha_1 \left(g_{2i} - \mu_2\right); \\ \alpha_1^2 &= \frac{1}{2}\left(1 + \frac{S_{11} - S_{22}}{\left[(S_{11} - S_{22})^2 + 4S_{12}^2\right]^{\frac{1}{2}}} \right); \\ \alpha_2^2 &= \frac{1}{2}\left(1 - \frac{S_{11} - S_{22}}{\left[(S_{11} - S_{22})^2 + 4S_{12}^2\right]^{\frac{1}{2}}} \right)\end{aligned} \quad (5)$$

The score in Eq. (4) is obtained using the maximum likelihood estimate, $\hat{\sigma}^2 = \frac{1}{2p}\Phi_{12}$, of the noise variance σ^2. The matching score of Eq. (4) holds also for the discrete noise distribution in Eq. (3) provided that the latter noise variance estimate, $\hat{\sigma}^2$, is used in this case, too.

2.2. Expectation Maximisation Outlier Masking

To derive the robust version of the score in Eq. (4), let a soft noise mask, $\boldsymbol{\Gamma} = (\gamma_i : i = 1, \ldots, p)$; $\gamma_i \in [0,1]$, be applied to the signals in the images, $\boldsymbol{g_1}$ and $\boldsymbol{g_2}$, to be matched. The maximum likelihood signal dissimilarity, $D_{12} = \min_{\boldsymbol{\theta},\boldsymbol{\gamma}}\left(-\ln \Pr(\mathbf{g}_1, \mathbf{g}_2)\right)$, combines non-outliers and outliers. Let ε_i be the residual matching error, or difference between the signals for the pixel, i, after their contrast and offset deviations are eliminated. Then

$$\begin{aligned}D_{12} &= -\sum_{i=1}^{p} \left[\gamma_i \left(\ln \mathcal{N}(\varepsilon_i)\right) + (1 - \gamma_i)\ln \mathcal{U}(\varepsilon_i)\right] \\ &= \nu + \nu \ln(\hat{Z}_{\sigma}) + \Psi_{12}\end{aligned} \quad (6)$$

where $\nu = \sum_{i=1}^{p}\gamma_i$, $\Psi_{12} = \sum_{i=1}^{p}(1 - \gamma_i)\ln \mathcal{U}(\varepsilon_i)$, $\hat{Z}_{\sigma} = \frac{\pi}{\nu}\Phi_{12}$, and Φ_{12} is the minimum total squared error with respect to the model parameters $\boldsymbol{\theta}' = \{a_1, a_2, b_1, b_2, \mathbf{g}\}$:

$$\begin{aligned}\Phi_{12} &= \min_{\boldsymbol{\theta}} \sum_{i=1}^{p} \gamma_i \left(\varepsilon_{1i}^2 + \varepsilon_{2i}^2\right) \\ &= \min_{\boldsymbol{\theta}} \sum_{i=1}^{p} \gamma_i \left[\left(g_{1i} - a_1 g_i - b_1\right)^2 + \left(g_{2i} - a_2 g_i - b_2\right)^2 \right] \\ &= \sum_{i=1}^{p} \gamma_i \hat{\varepsilon}_i^2 = \frac{1}{2}\left[S_{11} + S_{22} - \sqrt{\left(S_{11} - S_{22}\right)^2 + 4S_{12}^2} \right]\end{aligned}$$

Now $S_{jk} = \sum_{i=1}^{p}\gamma_i \left(g_{ji} - \mu_j\right)\left(g_{ki} - \mu_k\right)$; $j,k = 1,2$, $\mu_j = \frac{1}{\nu}\sum_{i=1}^{p}\gamma_i g_{ji}$, and the estimated residual noise $\hat{\varepsilon}_i$ and factors α_1 and α_2 are computed as in Eq. (6).

The estimated noise variance is $\hat{\sigma}^2 = \frac{1}{2\nu}\Phi_{12}$. As before, with this estimate, the matching score of Eq. (6) holds also for the discrete noise distribution in Eq. (3). The local minimum of D_{12} by the parameters $\boldsymbol{\theta}, \boldsymbol{\gamma}$ is obtained using the following EM-based iterative re-evaluation of both the soft masks and model parameters.

EM-based image matching:

Input: two images $\mathbf{g}_1$ and $\mathbf{g}_2$.

Initial step $t = 0$: match the images with the mask $\boldsymbol{\Gamma}^{[0]} = [\gamma_i^{[0]} = 1 : i = 1, \ldots, p]$ in order to find $\Phi_{12}^{[0]}$, $\hat{\sigma}^2_{[0]} = \frac{1}{2p}\Phi_{12}^{[0]}$, $D_{12}^{[0]} = D_{12}^{\mathrm{c}}$, $\alpha_1^{[0]}$, $\alpha_2^{[0]}$, $\mu_1^{[0]}$, $\mu_2^{[0]}$ (the conventional matching score and its components in Section 2.1). Set the prior $\rho_{[0]} = 0.5$.

Iteration $t = 1, 2, \ldots$: reset the mask and prior for the current residual errors $\hat{\varepsilon}_i^{[t]} = \alpha_2^{[t-1]}\left(g_{1i} - \mu_1^{[t-1]}\right) - \alpha_1^{[t-1]}\left(g_{i2} - \mu_2^{[t-1]}\right)$:

$$\gamma_i^{[t]} = \frac{\rho_{[t-1]}\mathcal{N}\left(\varepsilon_i^{[t]}|\sigma_{[t-1]}\right)}{\rho_{[t-1]}\mathcal{N}\left(\varepsilon_i^{[t]}|\sigma_{[t-1]}\right) + (1 - \rho_{[t-1]})\mathcal{U}\left(\varepsilon_i^{[t]}\right)};$$

$$\nu_{[t]} = \sum_{i=1}^{p}\gamma_i^{[t]}; \; \rho_{[t]} = \frac{\nu_{[t]}}{p}$$

and update $S_{jk}^{[t]}$; $j,k = 1,2$, $\Phi_{12}^{[t]}$, $\sigma^2_{[t]} = \frac{1}{2\nu_{[t]}}\Phi_{12}^{[t]}$, $D_{12}^{[t]}$, $\alpha_1^{[t]}$, $\alpha_2^{[t]}$, $\mu_1^{[t]}$, and $\mu_2^{[t]}$.

Stopping rule: terminate if $|D_{12}^{[t]} - D_{12}^{[t-1]}| \leq \theta_{\mathrm{r}} D_{12}^{[t]}$ or $t > \theta_{\mathrm{i}}$ where θ_{r} and θ_{i} are the fixed thresholds.

This algorithm provides additional cues for robust matching: if the prior ρ is less than a given threshold (specifying roughly an admissible level of outliers) or one of the relative scaling factors, α_1 or α_2 is close to zero, the matching obviously fails (i.e. the images are too dissimilar overall).

3. Experimental results

To evaluate the proposed robust matching, we considered the simplest case of rigid 2D matching under only contrast and offset deviations, mutual translations, random noise, and occlusions. Figure 2 shows that the EM-based symmetric dissimilarity score is more accurate and stable than its more conventional counterpart using Eq. (4) for the cases shown in Fig. 1.

Figure 3 and Table 1 compare the performance of the matching scores based on Eqs. (4) and (6) for faces from the MIT database in Appendix A distorted only by large occlusions and matched to the same database. The occlusions

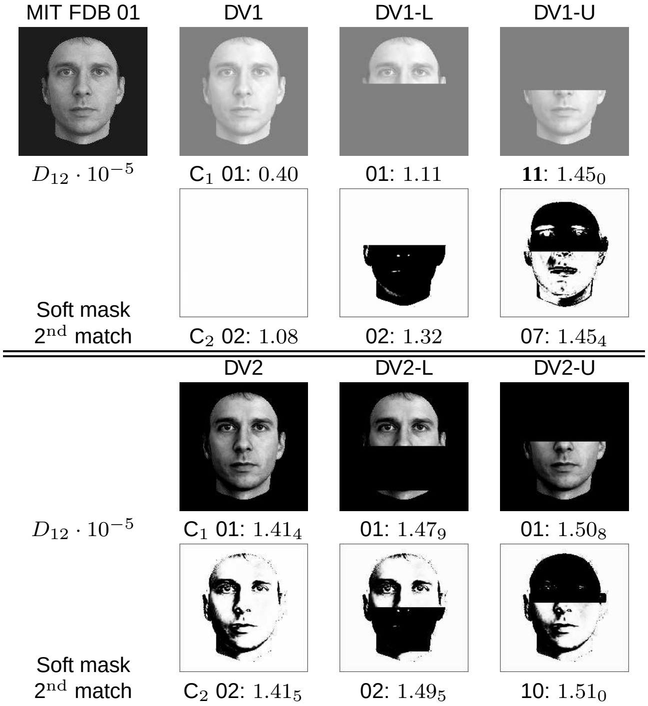

Figure 2. Robust matching of the images from Fig. 1: best (C_1) and second best (C_2) scores D_{12} using Eq. (6) and grey-coded soft masks (black - 0; white - 1).

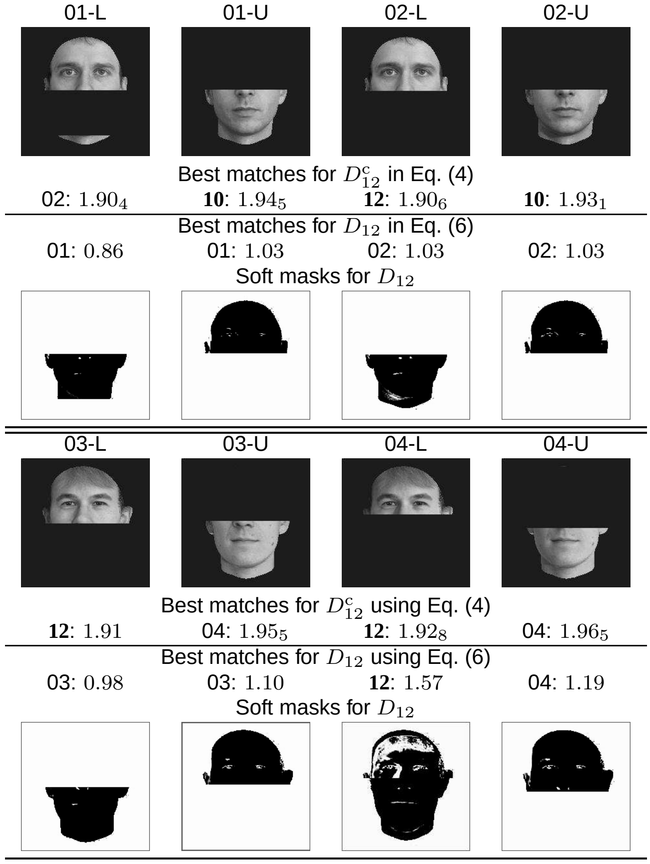

Figure 3. Matching of the half-occluded images from the MIT Face Database in Fig. 6: the robust score D_{12} in Eq. (6) vs. the score D^c_{12} using Eq. (4) (errors are boldfaced**).**

cover either the upper (U) or lower (L) half of each face. In these experiments, $\rho \leq 0.45$ and $\alpha_j \leq 0.10$. Note that images in the MIT database occur in pairs, images 1 and 2 are of the same subject as are images 3 and 4, etc. A match to another image of the same subject was considered a correct match.

The experiments show that the EM-based matching score based on Eq. (6) results in considerably more accurate recognition of the occluded faces - 95% correct matches compared to less than 50% using Eq. (4). It is also more stable with respect to noise: differences between best match and second best match were typically $\sim 25\%$ of the best EM-based match score compared to differences of $\sim 2\%$ for D^c_{12}. as regarding score differences between the best match and second best match.

Figures 4 and 5 show results from additional experiments in which subject's aversions to razors, barbers and the sun (typical sources of occlusions in real face images) were simulated. EM-based scores produced correct matches in all the examples shown whereas scores based on Eq. (4). As might be expected, the masks derived in the matching process correspond closely to the shapes 'added' to the images to simulate beards, glasses and additional hair.

4. Conclusions

It is clear that a requirement for effective face matching with templates in a database is an ability to handle both photometric deviations (variations in contrast and offset) as well as occlusions generated by artefacts covering regions of the face. By modelling both contrast and offset deviations as well as the possibility of outliers caused by large changes in the image being matched, we have shown that it is possible to match effectively even when large regions are modified.

References

[1] J.-H. Chen, C.-S. Chen, and Y.-S. Chen, Fast algorithm for robust template matching with M-estimators, *IEEE Trans. Signal Processing*, vol. 51, pp. 230–243, 2003.

Table 1. Best (C_1) and second best (C_2) match for the half-occluded images from the MIT Face Database in Appendix A using scores based on Eq. (4) and Eq. (6) (Matching errors are in bold: a match to another template of the same person is considered correct **).**

Face	Eq. (4): $D^c_{12} \cdot 10^{-5}$ C$_1$: D^c_{12}	C$_2$: D^c_{12}	Eq. (6): $D_{12} \cdot 10^{-5}$ C$_1$: D_{12}	C$_2$: D_{12}
01-L	02: 1.90$_4$	12: 1.91$_2$	01: 0.86	02: 1.28
02-L	**12**: 1.90$_6$	08: 1.92$_4$	02: 1.03	01: 1.27
03-L	**12**: 1.91$_0$	08: 1.92$_0$	03: 0.98	04: 1.45
04-L	**12**: 1.92$_8$	02: 1.94$_0$	**12**: 1.57	06: 1.60
05-L	06: 1.91$_5$	12: 1.92$_1$	05: 1.13	09: 1.50
06-L	06: 1.85$_5$	12: 1.87$_7$	06: 1.14	05: 1.53
07-L	08: 1.88$_9$	12: 1.91$_4$	07: 1.01	08: 1.53
08-L	08: 1.48$_0$	12: 1‘.89$_6$	08: 1.08	07: 1.86
09-L	**12**: 1.91$_0$	09: 1.93$_0$	09: 0.82	10: 1.38
10-L	**12**: 1.89$_4$	08: 1.92$_0$	10: 0.94	09: 1.39
11-L	12: 1.90$_2$	02: 1.93$_7$	11: 0.90	09: 1.50
12-L	12: 1.83$_0$	11: 1.87$_6$	12: 0.88	10: 1.57
13-L	14: 1.96$_4$	16: 1.96$_5$	13: 1.10	14: 1.49
14-L	**08**: 1.93$_6$	12: 1.93$_8$	14: 1.17	13: 1.50
15-L	16: 1.92$_0$	08: 1.95$_8$	15: 1.12	11: 1.53
16-L	16: 1.88$_5$	08: 1.93$_9$	16: 1.18	07: 1.55
17-L	**16**: 1.93$_6$	08: 1.94$_8$	17: 1.13	18: 1.50
18-L	18: 1.92$_5$	17: 1.92$_9$	18: 1.04	17: 1.48
19-L	**12**: 1.94$_9$	08: 1.94$_9$	19: 1.04	20: 1.48
20-L	**12**: 1.94$_6$	08: 1.95$_0$	20: 1.05	19: 1.49
01-U	**10**: 1.94$_5$	04: 1.94$_5$	01: 1.03	02: 1.32
02-U	**10**: 1.93$_1$	04: 1.93$_1$	02: 1.03	01: 1.32
03-U	04: 1.96$_5$	03: 1.95$_6$	03: 1.10	09: 1.47
04-U	04: 1.96$_5$	03: 1.97$_1$	04: 1.19	10: 1.49
05-U	05: 1.94$_8$	06: 1.95$_2$	05: 1.18	09: 1.48
06-U	06: 1.88$_5$	05: 1.89$_7$	06: 1.21	10: 1.52
07-U	**04**: 1.95$_1$	03: 1.95$_2$	07: 1.11	01: 1.48
08-U	**04**: 1.88$_4$	03: 1.88$_4$	08: 1.14	09: 1.51
09-U	10: 1.92$_7$	09: 1.93$_0$	09: 1.01	10: 1.42
10-U	10: 1.92$_8$	09: 1.93$_7$	10: 1.06	12: 1.51
11-U	**03**: 1.95$_1$	04: 1.95$_4$	11: 1.08	09: 1.49
12-U	**04**: 1.85$_5$	03: 1.85$_6$	12: 1.02	10: 1.49
13-U	13: 2.00$_6$	14: 2.00$_8$	13: 1.17	14: 1.54
14-U	13: 1.98$_9$	14: 1.99$_3$	14: 1.18	13: 1.54
15-U	15: 2.02$_4$	13: 2.03$_6$	15: 1.11	10: 1.56
16-U	15: 1.93$_7$	13: 1.95$_4$	**12**: 1.58$_5$	10: 1.59$_0$
17-U	**03**: 1.98$_9$	04: 1.99$_0$	17: 1.02	18: 1.51
18-U	**04**: 1.98$_2$	03: 1.98$_5$	18: 1.13	17: 1.51
19-U	**03**: 1.98$_7$	04: 1.98$_8$	19: 1.14	10: 1.51
20-U	**04**: 1.97$_5$	03: 1.97$_8$	20: 1.16	09: 1.48
	Error rate: 47%		Error rate: 5%	

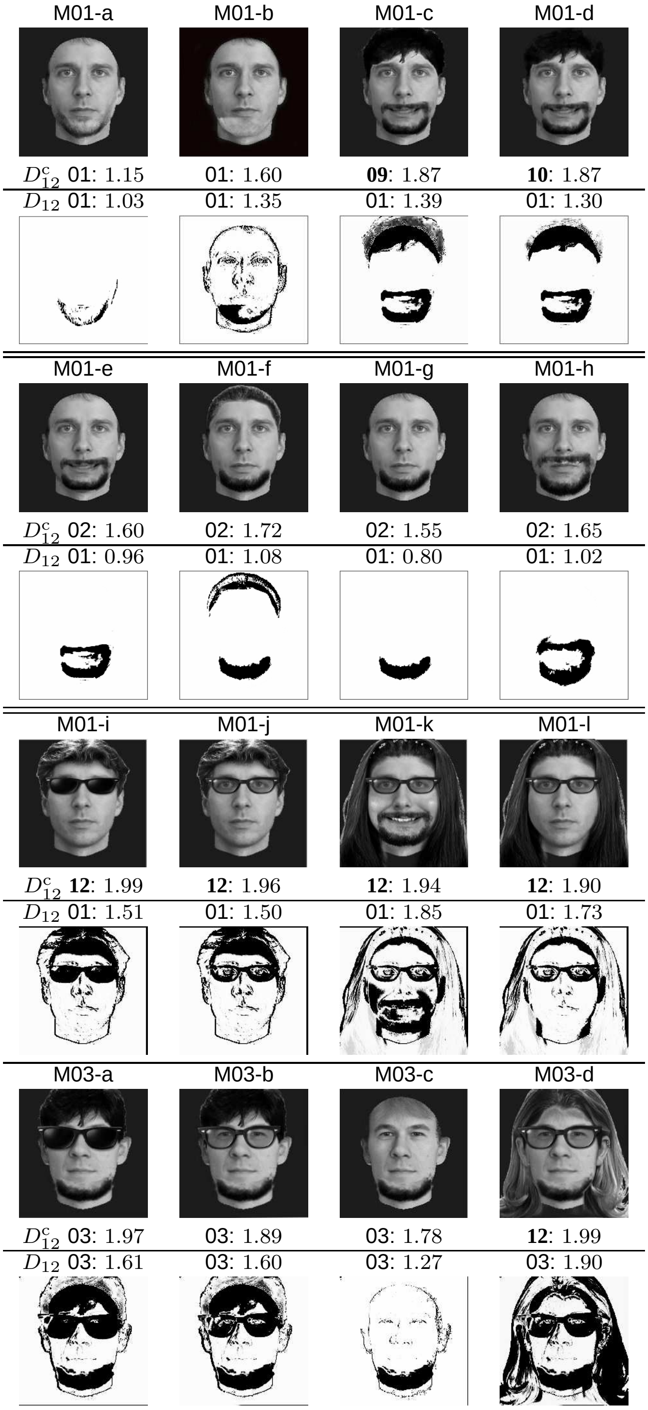

Figure 4. Best match for the faces synthesized from faces '01' and '03' of the MIT database in Fig. 6: robust score D_{12} using Eq. (6) vs. D^c_{12} using Eq. (4) and soft masks (no errors vs. 50% errors (boldfaced), respectively**).**

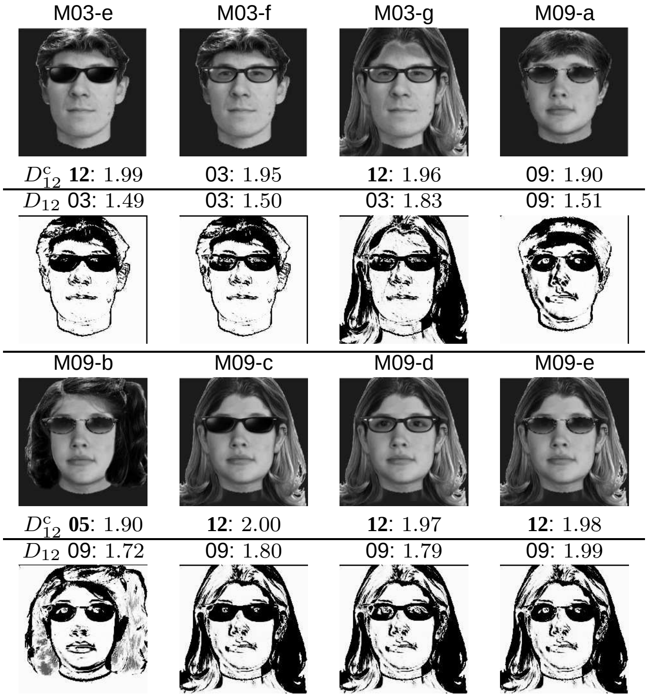

Figure 5. Best match for faces synthesized from faces '03' and '09' from the MIT database in Fig. 6: the robust score D_{12} using Eq. (6) vs. the score D^c_{12} using Eq. (4) and grey-coded soft masks (no errors vs. 50% errors (boldfaced), respectively**).**

[2] B. J. Frey and N. Jojic, "A comparison of algorithms for inference and learning in probabilistic graphical models", *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 27, pp. 1392–1416, 2005.

[3] G. Gimel'farb, A. A. Farag, and A. El-Baz, "Expectation-Maximization for a linear combination of Gaussians", in *Proc. 17th IAPR Int. Conf. on Pattern Recognition, August 22 27, 2004, Cambridge, UK*, IEEE CS Press: Los Alamitos, vol. 3, pp. 422–425, 2004.

[4] D. Hasler, L. Sbaiz, S. Süsstrunk, and M. Vetterli, Outlier modeling in image matching, *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 25, pp. 301–315, 2003.

[5] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, New York, 2001.

[6] P. J. Huber, *Robust Statistics*. John Wiley & Sons, New York, 1981.

[7] S.-H. Lai, Robust image matching under partial occlusion and spatially varying illumination change, *Computer Vision and Image Understanding*, vol.78, pp. 84–98, 2000.

[8] MIT Face Database [on-line] (accessed 24 Aug 2006) http://vismod.media.mit.edu/pub/images/.

[9] P. J. Rousseeuw and A. M. Leroy, *Robust Regression and Outlier Detection*, John Wiley & Sons, New York, 1987.

[10] C.-H. T. Yang, S.-H. Lai, and L.-W.Chang, Robust face image matching under illumination variations, *EURASIP Journal on Applied Signal Processing*, vol. 2004, pp. 2533–2543, 2004.

A The MIT Face Database

The above experiments use a subset of the MIT database [8] containing two differently illuminated images for each of the ten persons: in total, 20 normalised 200×200 grayscale images shown in Fig. 6.

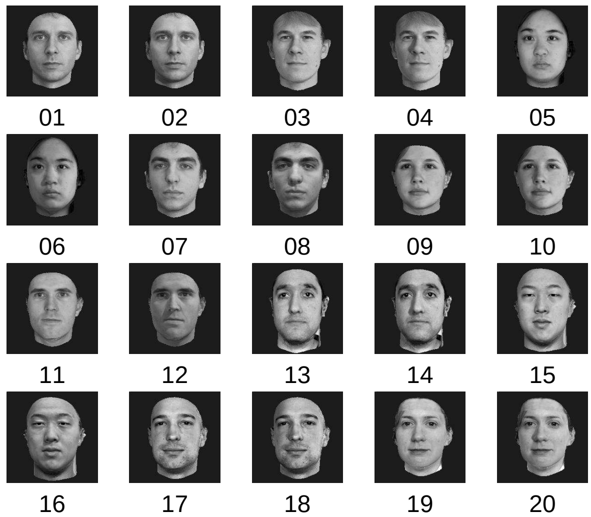

Figure 6. The MIT face database [8] subset.