Face Recognition

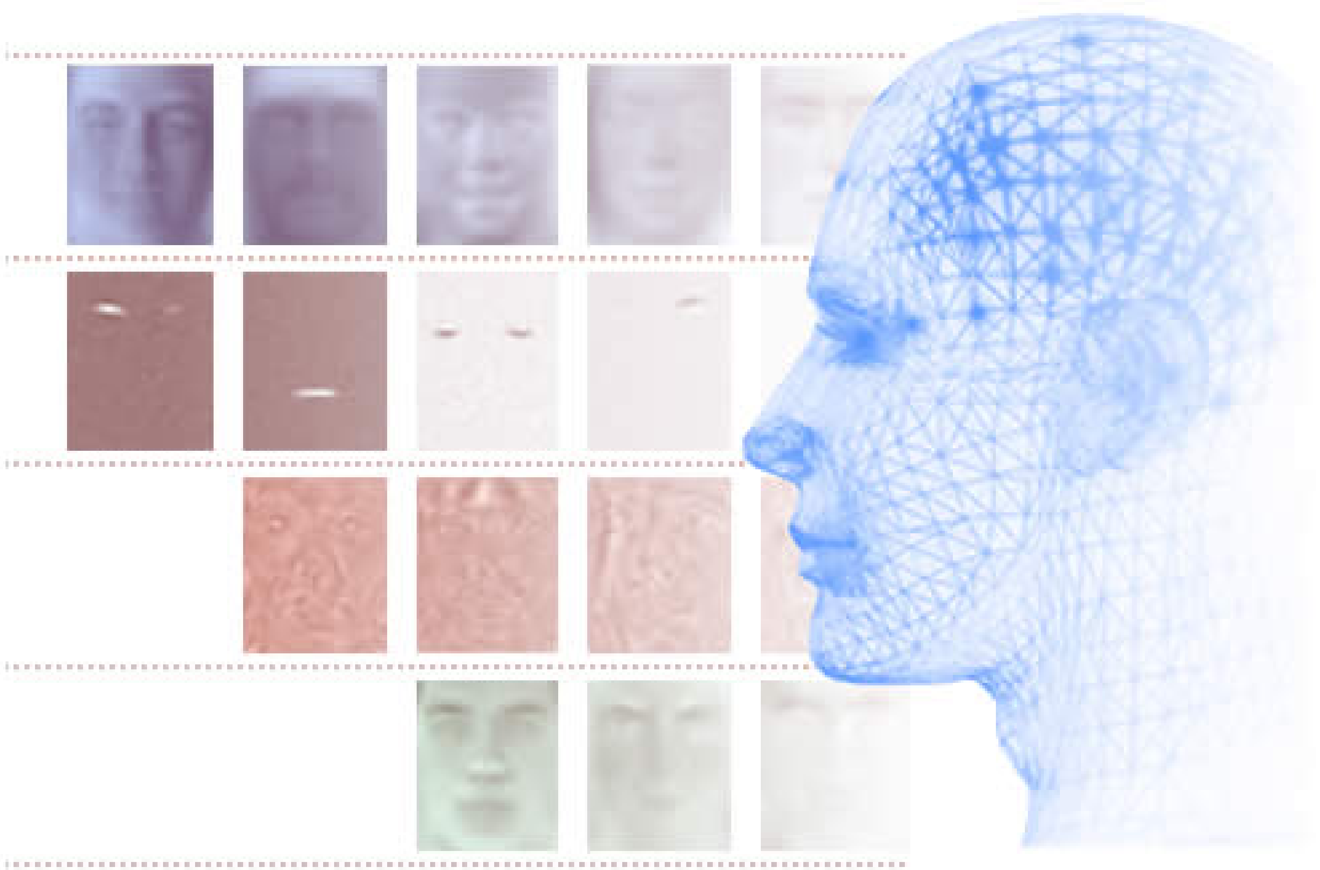

Agenda

- Face recognition by humans
- 3D – morphable models – pose and illumination
  - model creation
  - model fitting
  - results
- Face recognition vendor test 2006

Face Recognition by Humans [SBOR06]

- Result 1: Humans can recognize familiar faces in very low resolution images

32x32

[TFF07]

256x256

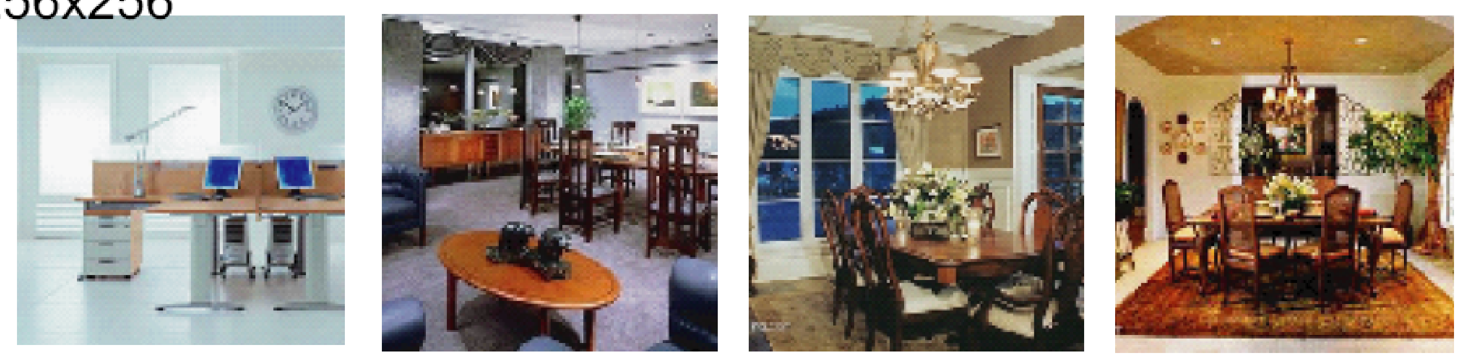

Face Recognition by Humans

- Result 2: The ability to tolerate degradations increases with familiarity

[SBOR06]

Fig. 1. *Unlike current machine-based systems, human observers are able to handle significant degradations in face images. For instance, subjects are able to recognize more than half of all familiar faces shown to them at the resolution depicted here. Individuals shown in order are: Michael Jordan, Woody Allen, Goldie Hawn, Bill Clinton, Tom Hanks, Saddam Hussein, Elvis Presley, Jay Leno, Dustin Hoffman, Prince Charles, Cher, and Richard Nixon.*

Face Recognition by Humans

- Result 3: Facial features are processed holistically

[SBOR06]

Fig. 4. *Try to name the famous faces depicted in the two halves of the left image. Now try the right image. Subjects find it much more difficult to perform this task when the halves are aligned (left) compared to misaligned halves (right), presumably because holistic processing interacts (and in this case, interferes) with feature-based processing. The two individuals shown here are Woody Allen and Oprah Winfrey.*

Face Recognition Using 3D Models – Pose and Illumination

- Modeling the effects of illumination and pose
- 3D shape information and reflectance models
- Effect of illumination on the appearance of a human face

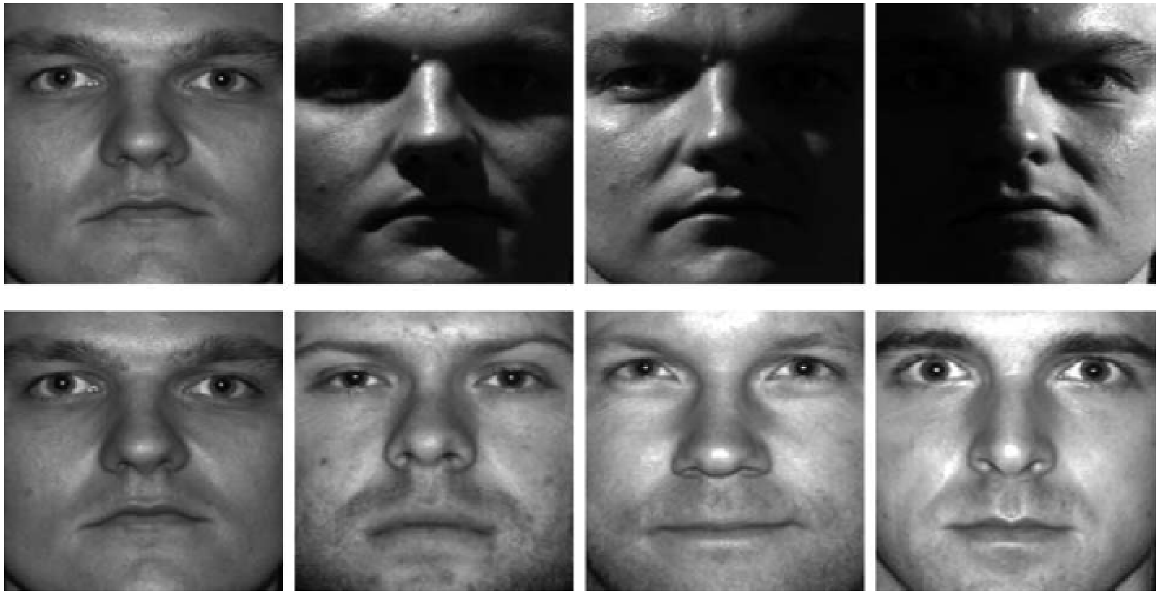

[RHVK06]

Models of Illumination and reflectance

- Lambertian model

$$I(\vec{l}) = \rho \max(\vec{l} \cdot \vec{n}, 0)$$

- Phong model

$$I(\vec{l}) = a\rho + \rho \max(\vec{l} \cdot \vec{n}, 0) + s(\vec{l}, \vec{n}, \vec{v})$$

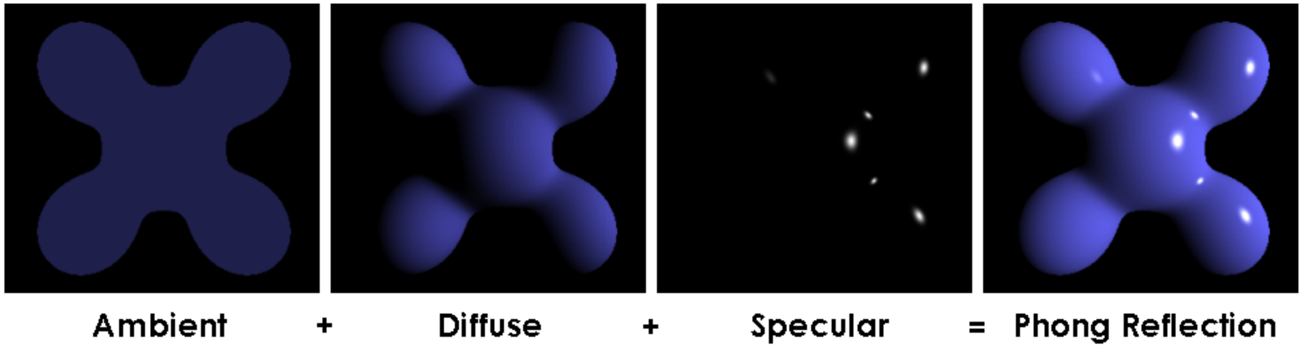

Shadows

- Attached Shadow
  - related to the local geometry
- Cast Shadow
  - related to the object's global geometry

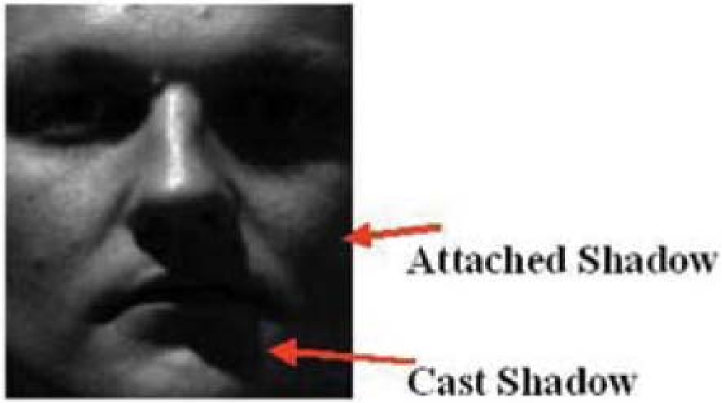

3D-Morphable Model [RHVK06]

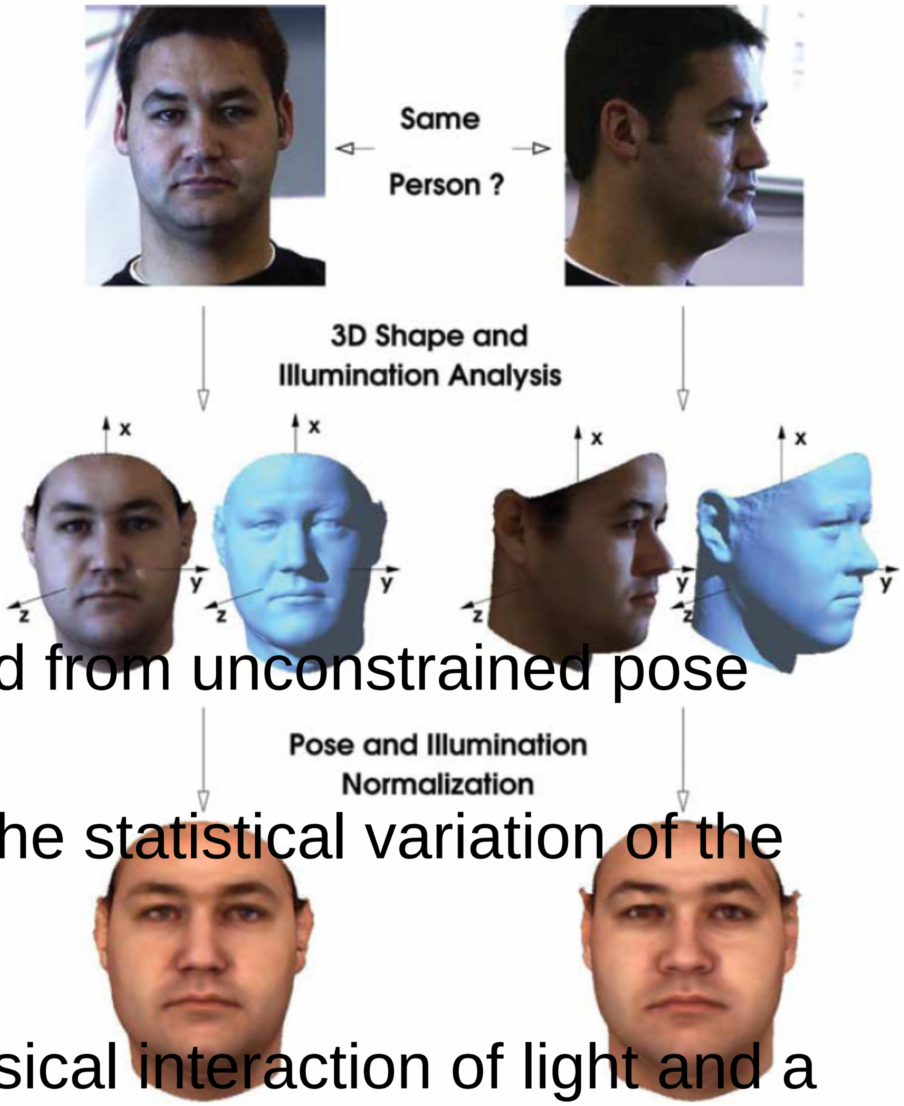

- Estimating
  - 3D-shape
  - albedo
  - pose
  - illumination
- From a single image viewed from unconstrained pose and illumination conditions
- Strong prior knowledge of the statistical variation of the shape and texture
  and
  a model simulating the physical interaction of light and a face's surface

3D-Morphable Model

- In computer graphics, faces can be represented by a mesh that includes a dense set of vertices (N_v) and a triangle list, that specifies the connectivity between the vertices.
  - global illumination effects (e.g. cast shadows) can be used in the image analysis algorithm
  - pose can be estimated

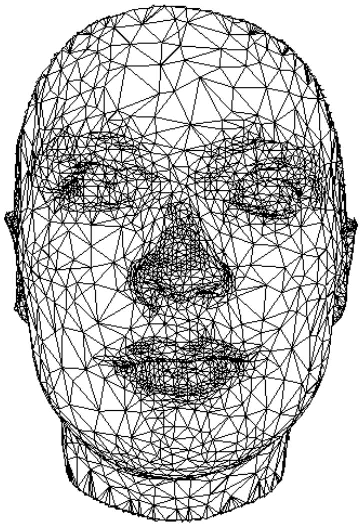

- Shape matrix

$$\mathbf{S} = \begin{pmatrix} x_1 & x_2 & \cdots & x_{N_v} \\ y_1 & y_2 & \cdots & y_{N_v} \\ z_1 & z_2 & \cdots & z_{N_v} \end{pmatrix}$$

- Texture matrix

$$\mathbf{T} = \begin{pmatrix} r_1 & r_2 & \cdots & r_{N_v} \\ g_1 & g_2 & \cdots & g_{N_v} \\ b_1 & b_2 & \cdots & b_{N_v} \end{pmatrix}$$

3D-Morphable Model – Image Creation

- Produce an image of the shape object
  - rotation and translation of the object

$$\mathbf{W} = \mathbf{R}_\gamma \mathbf{R}_\theta \mathbf{R}_\phi \mathbf{S} + \tau_w \mathbf{1}_{1\times N_v}$$

  - perspective projection

$$x_i = f\frac{\mathbf{W}_{1,i}}{\mathbf{W}_{3,i}} \quad y_i = f\frac{\mathbf{W}_{2,i}}{\mathbf{W}_{3,i}}$$

3D-Morphable Model – Image Creation

- Produce an image of the texture model
  - Phong reflectance model
  - considers diffuse and specular reflections

$$\mathbf{t}_i^I = \begin{pmatrix} l_r^a & 0 & 0 \\ 0 & l_g^a & 0 \\ 0 & 0 & l_b^a \end{pmatrix} \cdot \mathbf{t}_i + \begin{pmatrix} l_r^d & 0 & 0 \\ 0 & l_g^d & 0 \\ 0 & 0 & l_b^d \end{pmatrix} \cdot \Big((\vec{n}_i \cdot \vec{l}) \mathbf{t}_i + k_s (\vec{v}_i \cdot \vec{r}_i)^{\nu} \mathbf{1}_{3\times 1} \Big)$$

3D-Morphable Model

- Extent the mesh model
  - parameterize the morphable model
  - probability distribution is learned from the training images
  - PCA is applied to the registered shape and texture exemplars

- The face of any individual can be obtained by

$$\mathbf{S} = \overline{\mathbf{S}} + \sum_{i=1}^{N_S} \alpha_i \cdot \mathbf{S}^i, \quad \mathbf{T} = \overline{\mathbf{T}} + \sum_{i=1}^{N_T} \beta_i \cdot \mathbf{T}^i$$

3D-Morphable Model – Face Analysis

- The task is to find the model parameters
  - => the face image rendered from those parameters has to be as close as possible to the input image
  - => minimize the error of a cost function
    or in our case
    maximizing the posterior probability of the model parameters

$$\min_{\theta} \sum_{i} \frac{1}{\sigma_I^2} \left\| I(x_i(\theta), y_i(\theta)) - \mathbf{t}_i^I(\theta) \right\|^2 + \sum_{i=1}^{N_S} \frac{\alpha_i^2}{\sigma_{S,i}^2} + \sum_{i=1}^{N_T} \frac{\beta_i^2}{\sigma_{T,i}^2}$$

3D-Morphable Model – Face Analysis

- Nonlinear optimization algorithm => initial estimate
  - mean shape
  - mean texture
  - pose and light direction is taken to be frontal
  - seven landmark points

Original Features

[BV03]

3D-Morphable Model – Face Analysis

- Prior models (shape, texture) are not strong enough to obtain an accurate estimate of the 3D-shape, when only those few anchor points are used

- => add other features
  - edges
  - specular highlights
  - texture constraint

3D-Morphable Model – Face Analysis

- Multiple feature fitting algorithm

$$\min_{\theta} \tau^c C^c + \tau^e C^e + \tau^s C^s + \tau^p C^p + \tau^t C^t$$

  - pixel intensity feature cost function

$$C^c = \sum_i (1/\sigma_I^2) \| I(x_i(\theta), y_i(\theta)) - \mathbf{t}_i^I(\theta) \|^2$$

  - prior feature cost function

$$C^p = \sum_{i=1}^{N_S} (\alpha_i^2 / \sigma_{S,i}^2) + \sum_{i=1}^{N_T} (\beta_i^2 / \sigma_{T,i}^2)$$

  - edge-, specular highlight- and texture constraint cost function C^e, C^s, and C^t

3D-Morphable Model – Face Analysis

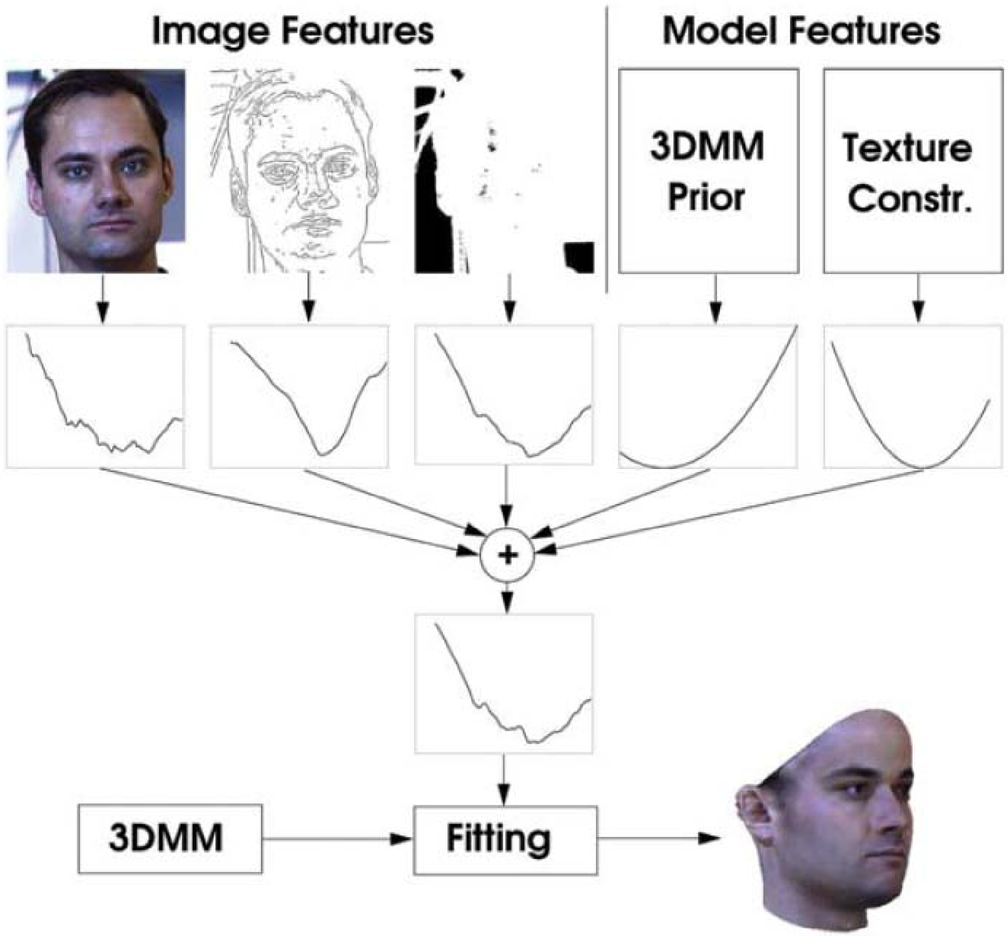

3D-Morphable Model – Face Analysis

[BV03]

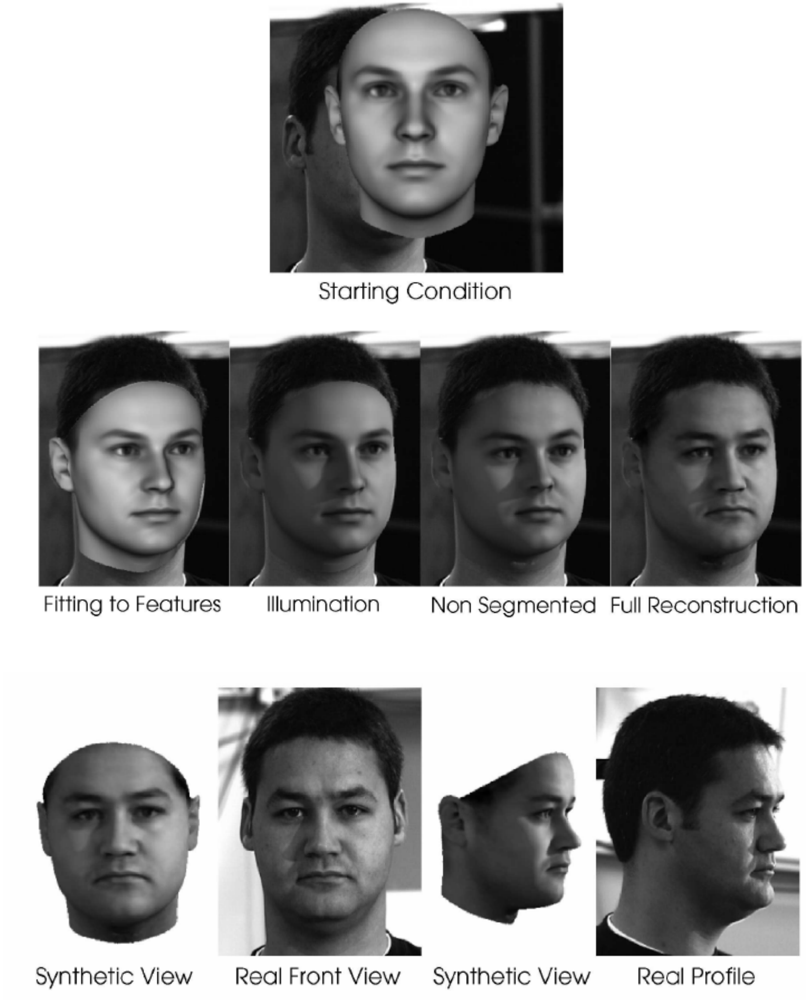

3D-Morphable Model Demonstration

- http://www.mpi-inf.mpg.de/~blanz/

Pose Normalization Using 3D-Morphable Models

[RHVK06]

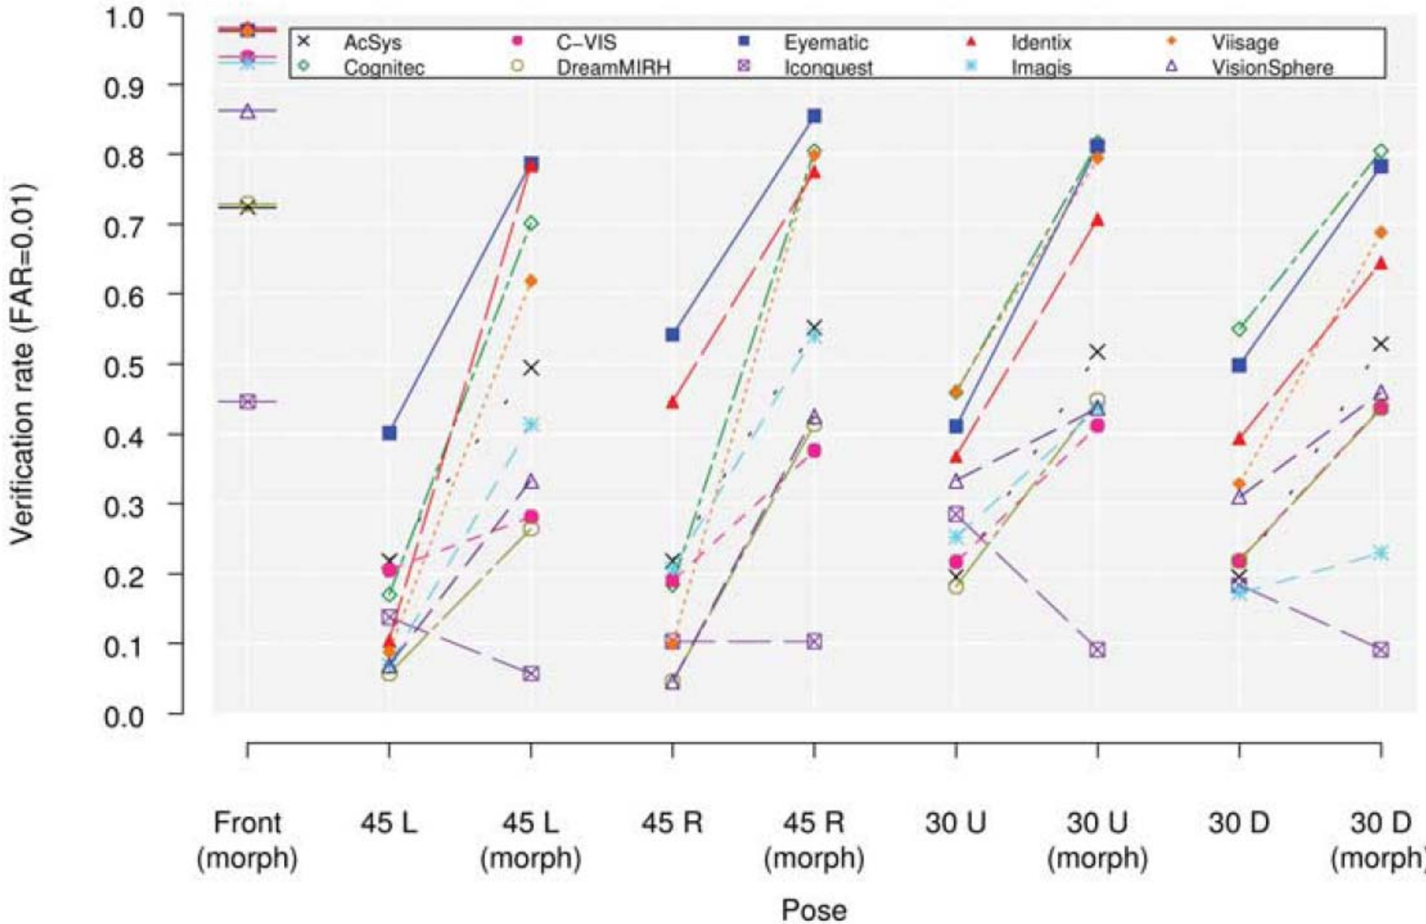

The effect on the verification performance of the original images versus normalized images using the 3-D Morphable Model. The verification rate at a false alarm rate of 1% is plotted. This figure was first published in the face Recognition Vendor Test 2002

Comparison of Recognition Methods – Variable Pose and Lighting

[RHVK06]

COMPARISON OF RECOGNITION METHODS			
Gallery View	Probe View		
	front (0°)	side (16.5°)	profile (62.1°)
3DMM [38]:			
front	0.1%	1.6%	24.4%
side	3.6%	0.7%	16.3%
profile	23.7%	14.0%	10.6%
Spherical-basis Morphable Model of Zhang & Samaras [54]:			
front	3.5%	5.4%	21.3%
side	6.1%	3.3%	21.4%
profile	18.2%	18.5%	10.4%
Zhou & Chellappa [59]:			
front	3%	12%	48%
Pose-encoded spherical harmonics of Yue *et al.* [50]:			
front	NA	1.2%	NA

Table 2 Mean Identification Error Percentage for Different Methods Obtained for the PIE Data Set, Averaged Over All Lighting Conditions for Front, Side, and Profile View Galleries. All the Experiments Included a Single Image per Individual in the Gallery Set With a Front Illumination Condition

Face Recognition Vendor Test – FRVT 2006 [PSO+07]

- Recognition from
  - high-resolution still-images
  - 3D images
- Under controlled and uncontrolled illumination

FRVT 2006 – Error Rates

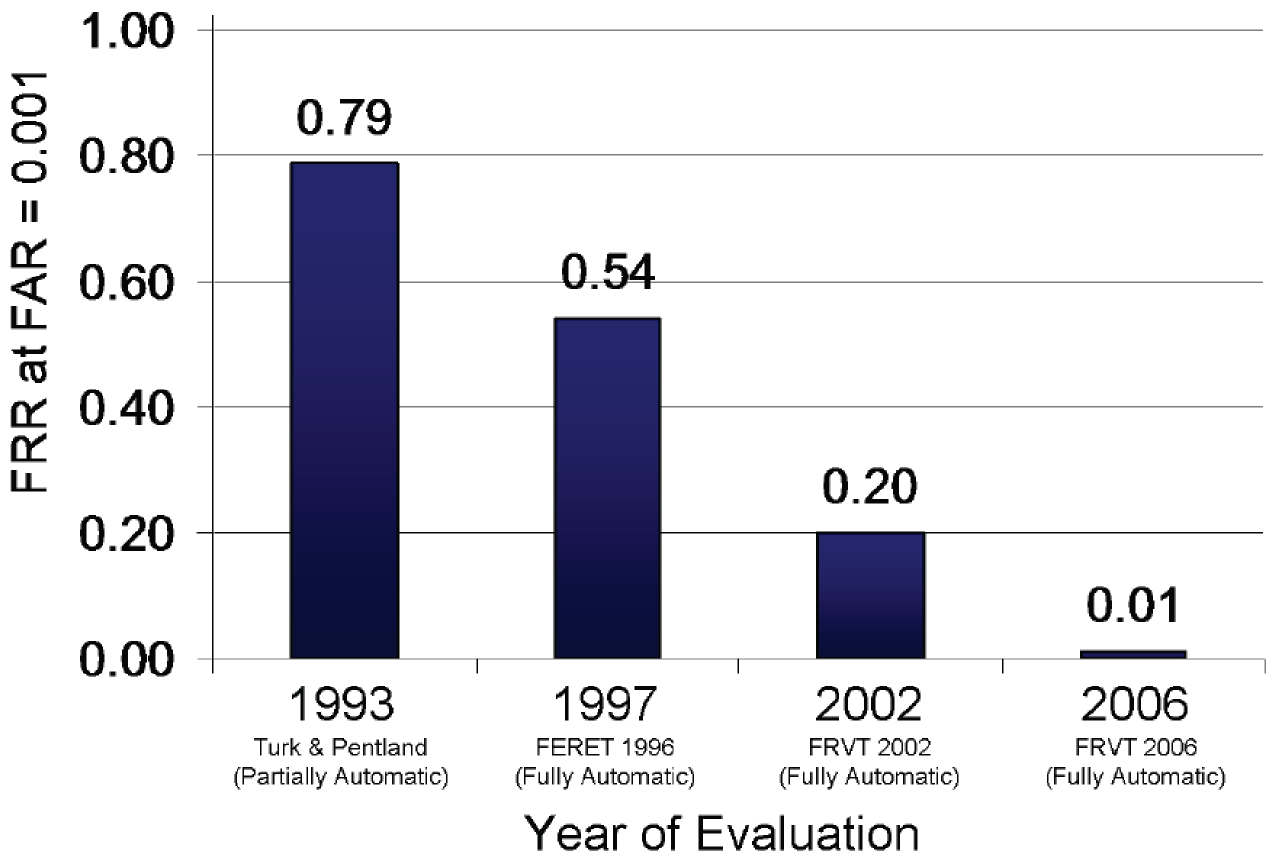

[PSO+07]

- The best performing face recognition algorithms were more accurate than humans !!

FRVT 2006 – Face Recognition, Controlled Environment

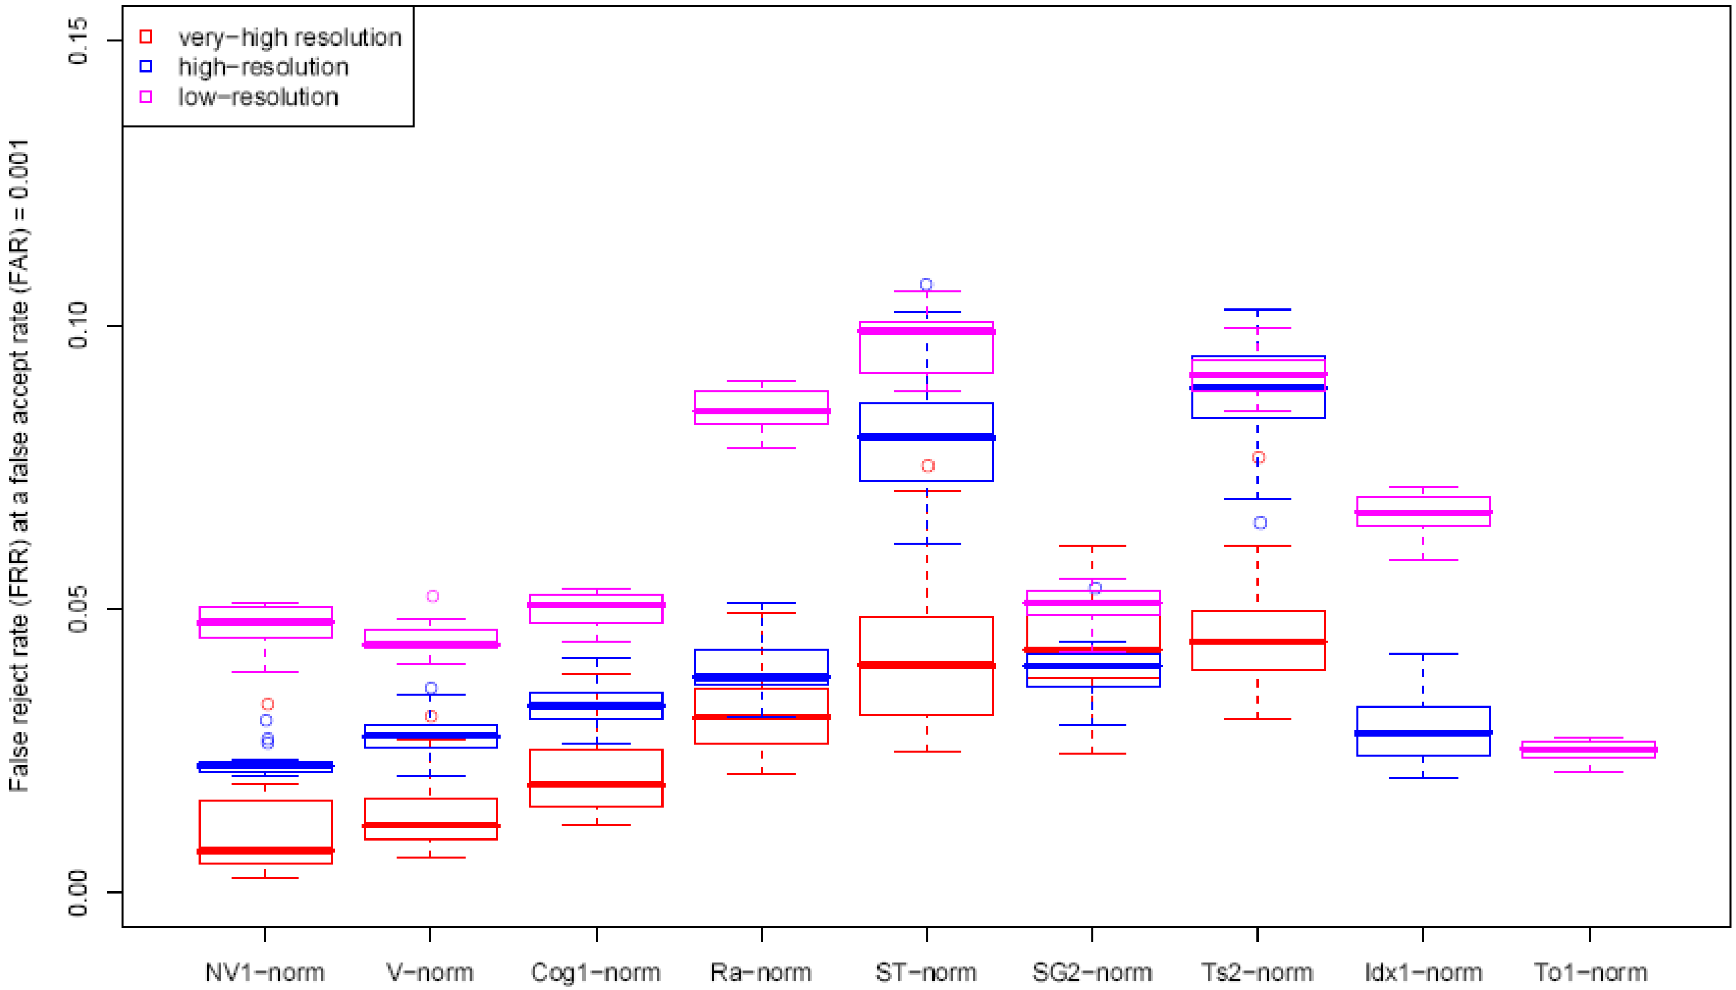

FRVT 2006 – Face Recognition, Uncontrolled Environment

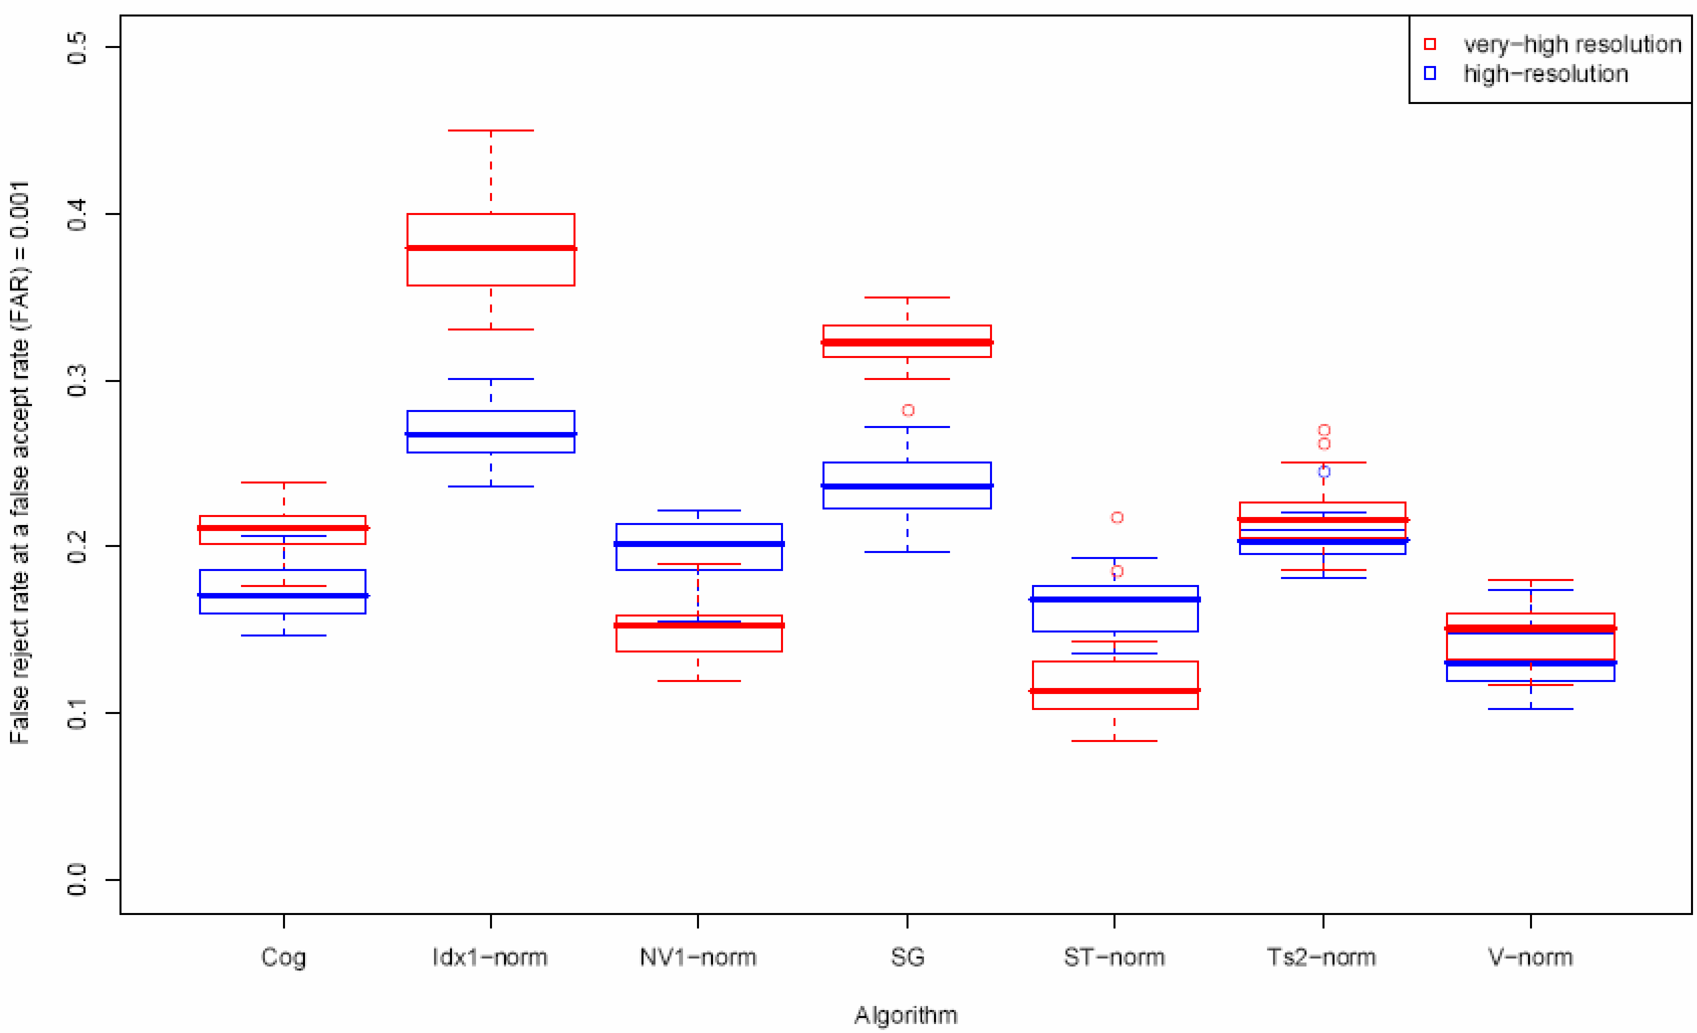

FRVT 2006 – Results

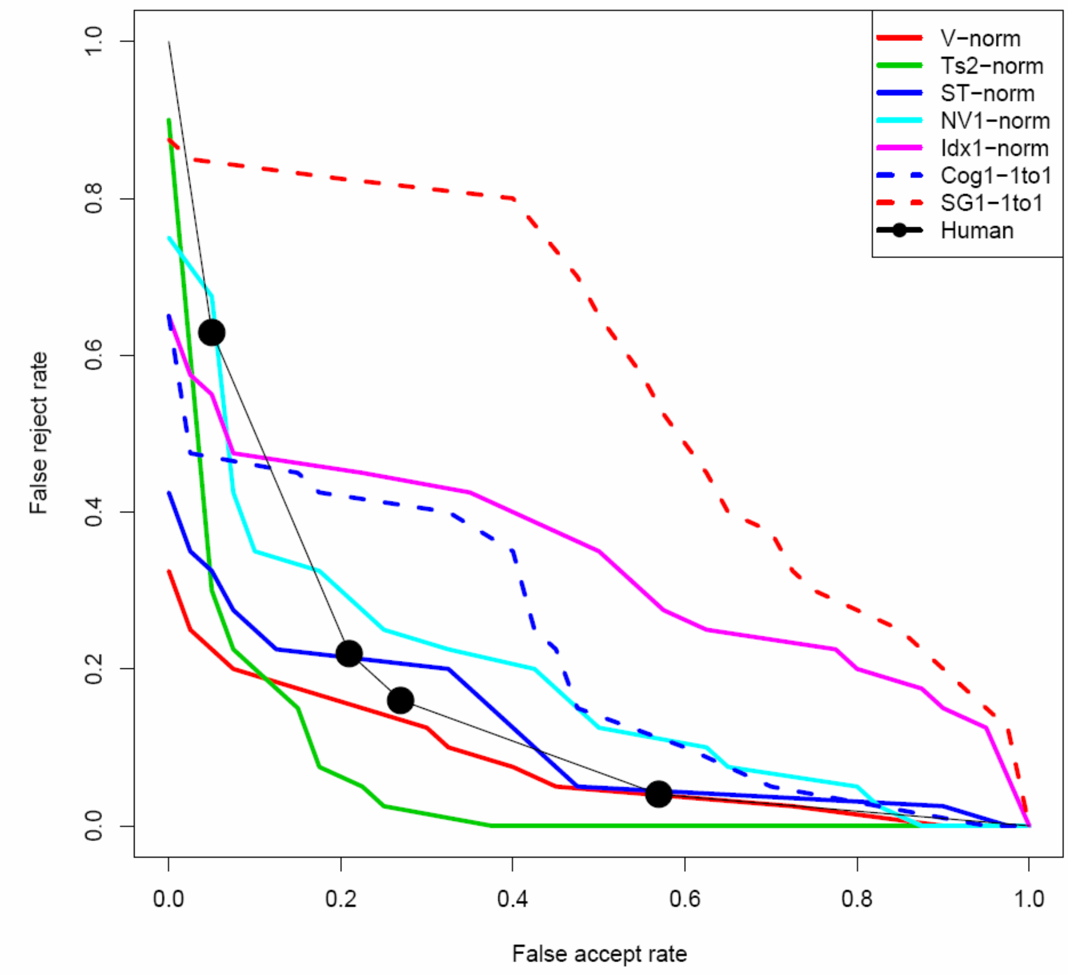

[PSO+07]

Literature

[BV03] Volker Blanz and Thomas Vetter. Face recognition based on fitting a 3d morphable model. *Transactions on Pattern Analysis and Machine Intelligence*, 25(9):1063–1074, September 2003.

[PSO+07] P. Jonathon Phillips, W. Todd Scruggs, Alice J. O'Toole, Patrick J. Flynn, Kevin W. Bowyer, Cathy L. Schott, and Matthew Sharpe. Frvt 2006 and ice 2006 large-scale results. Technical report, NISTIR 7408 - National Institute of Standards and Technology, Gaithersburg, MD 20899, March 2007.

[RHVK06] Sami Romdhani, Jeffrey Ho, Thomas Vetter, and David J. Kriegman. Face recognition using 3d models - pose and illumination. *Proceedings of the IEEE*, 94(11):1977–1999, November 2006.

[SBOR06] Pawan Sinha, Benjamin Balas, Yuri Ostrovsky, and Richard Russel. Face recognition by humans - nineteen results all computer vision researchers should know about. *Proceedings of the IEEE*, 94(11):1948–1962, November 2006.

[TFF07] Antonio Torralba, Rob Fergus, and William T. Freeman. 80 million tiny images: a large dataset for non-parametric object and scene recognition. Technical Report MIT-CSAIL-TR-2007-024, Computer Science and Artificial Intelligence Lab, Massachusetts Institute of Technology, 2007.

Thanks for your Attention