Heterogeneous Face Recognition

Brendan F. Klare, Ph.D.

In collaboration with: Anil K. Jain, Unsang Park, Hu Han, Kate Bonnen, Zhifeng Li, Serhat Bucak, Mark Burge, Josh Klonz, Richard Vorder Bruegge, Alessandra Paulino and Tayfun Akgul

Why automated face recognition?

- Face recognition is one of the most primal tasks
- We rely on specialized processes in the brain that exclusively perform face specific tasks
- Information humans immediately derived from a face can include
  - Identity
  - Emotion
  - Age
  - Gender
  - Race
- Automation of this vital operation is of substantial benefit

Where is face recognition today?

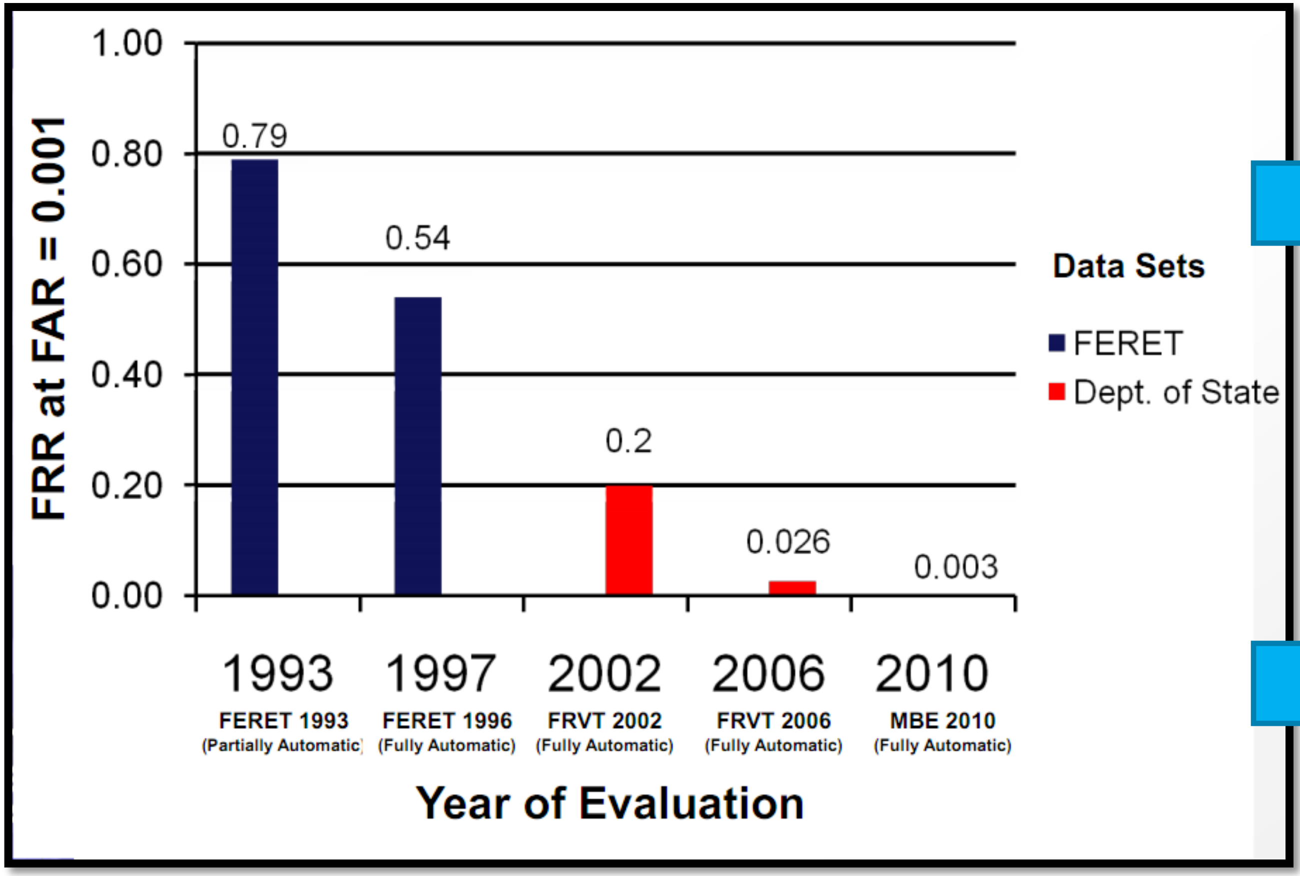

De-duplication

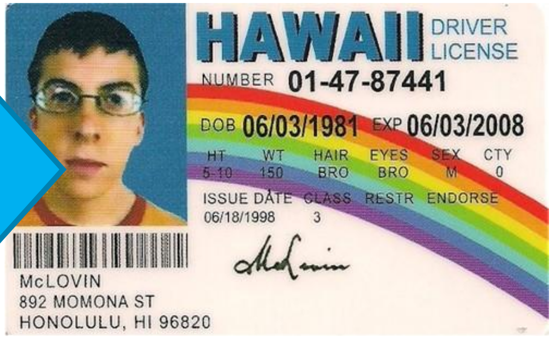

Inmate Identify Confirmation

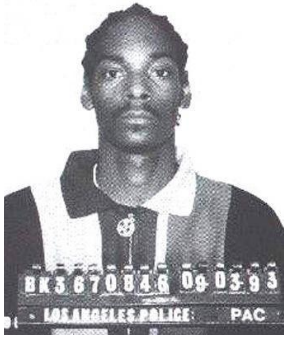

Source: NIST Interagency Report 7709 - Report on the Evaluation of 2D Still-Image Face Recognition Algorithms

Where is face recognition going?

- Face recognition technology is moving towards **ubiquity**: reducing violent, unpredictable acts, like the rioting in London

Overview of Automated Face Recognition Algorithms

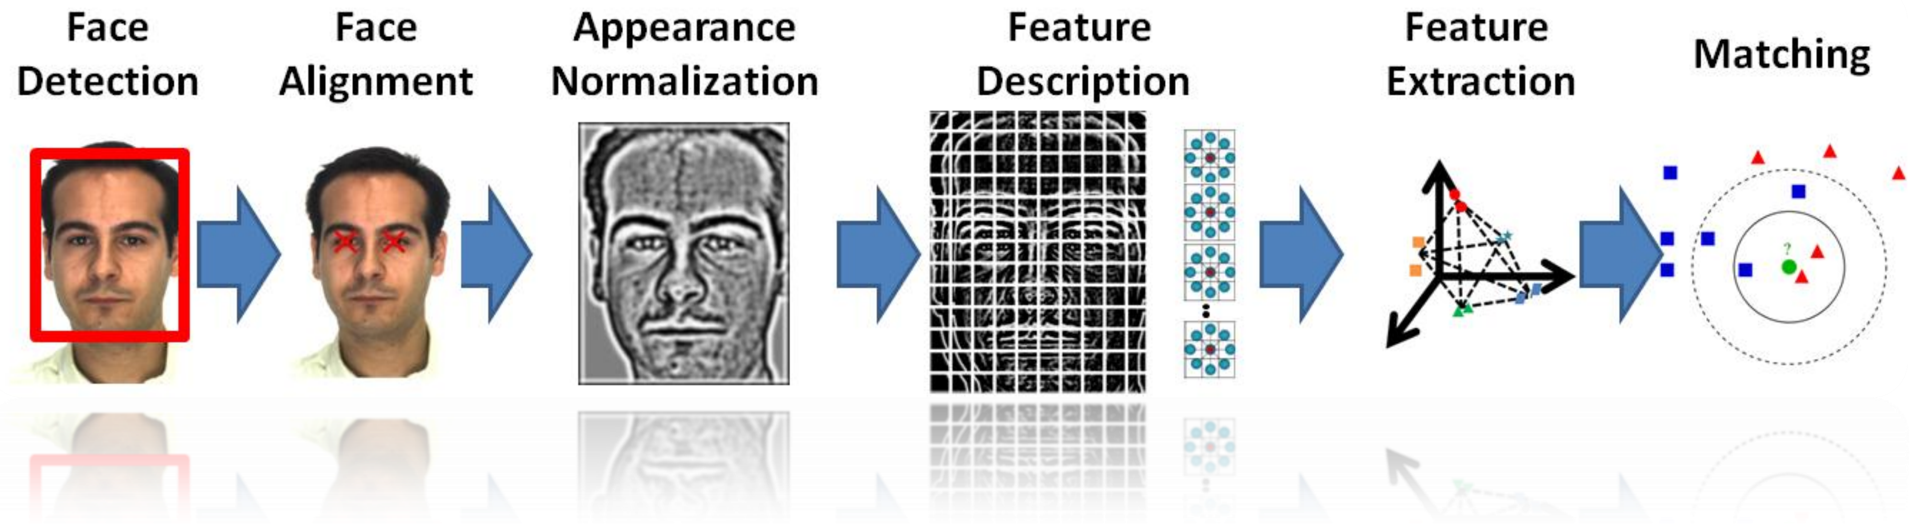

- Face recognition algorithms generally follow the same pipeline as listed above
- FR research can improve any specific stage above, or address the entire face recognition pipeline

Face Detection

- Must first find a face to perform face recognition
- Seminal approach developed by Viola and Jones [1]
- Facilitated robust face recognition in real-time
- Made available via OpenCV project

Detection = Where's Waldo?

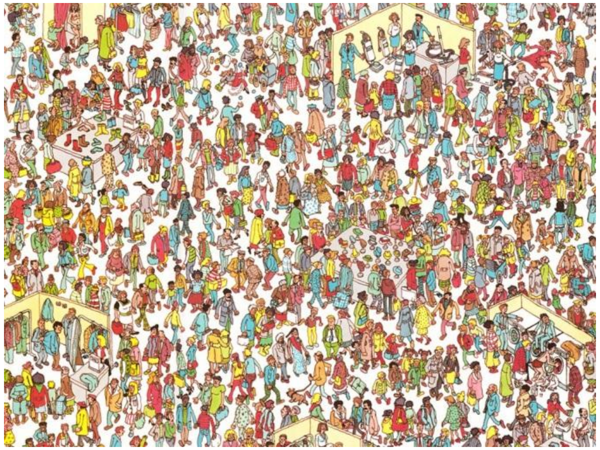

Overview of Viola Jones face detection algorithm:

Encode face images with Haar features:

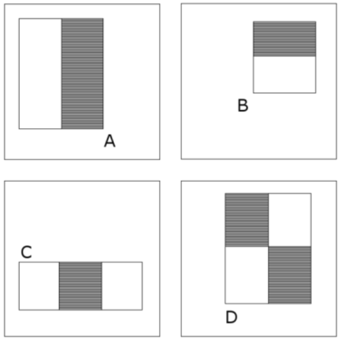

Perform cascaded detection using features learning from AdaBoost algorithm:

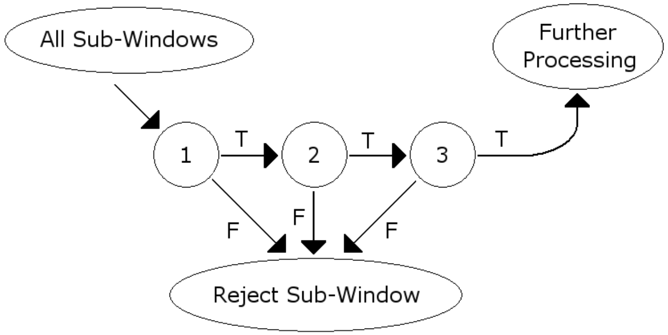

Example output:

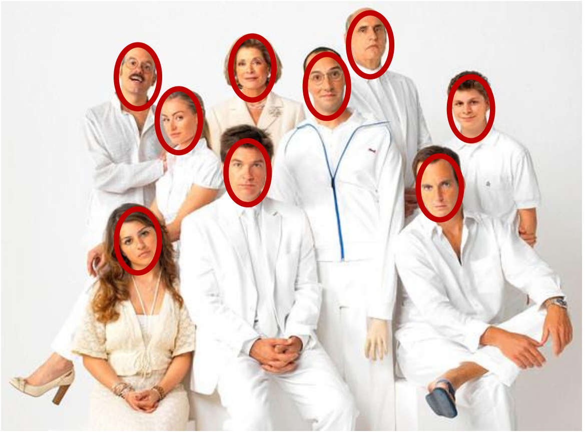

[1] Viola, Paul, and Michael J. Jones. "Robust real-time face detection." *International journal of computer vision* 57.2 (2004): 137-154.

noblis

Face Localization

- Face detector gives us rough location of the face, but where should the algorithm compute facial measurements from?
- Face has a fixed geometry (e.g., eyes above nose, mouth wider than nose)
- Learn facial geometry to aid in landmark detection (e.g. Active Shape Models, Active Appearance Models, Morphable Models)
- Landmarks can then be used to align the faces to a fixed model

Detected…

But not localized

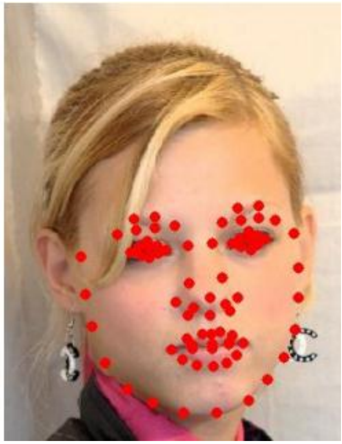

Automated Landmark Detection Results

Landmark Detection and Face Alignment

Landmark Detection

Active Shape Model (global) [1][2]:

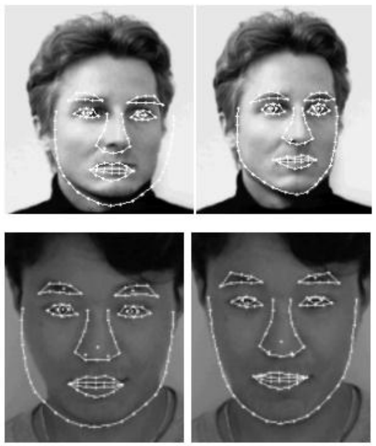

Local + global approach [3]:

Face Alignment

Procustes-based:

Morphable Model [4]:

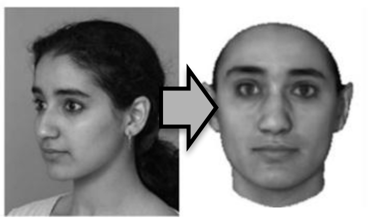

Structure from Motion [5]:

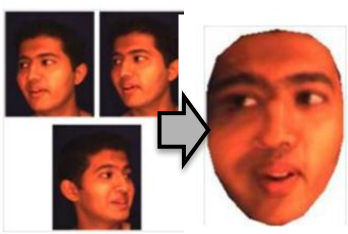

Component-based [6]:

[1] Cootes, Timothy F., et al. "Active shape models-their training and application."*Computer vision and image understanding* 61.1 (1995): 38-59.
[2]Wang, Wei, et al. "An improved active shape model for face alignment."*Proceedings of the 4th IEEE International Conference on Multimodal Interfaces*. IEEE Computer Society, 2002.
[3] Belhumeur, Peter N., et al. "Localizing parts of faces using a consensus of exemplars." *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*. IEEE, 2011.
[4] Blanz, Volker, and Thomas Vetter. "Face recognition based on fitting a 3D morphable model." *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 25.9 (2003): 1063-1074.
[5] Park, Unsang, and Anil Jain. "3D model-based face recognition in video."*Advances in Biometrics* (2007): 1085-1094.
[6] K. Bonnen, B. Klare, and A. K. Jain, "Component-Based Representation in Automated Face Recognition", IEEE Transactions on Information Forensics and Security, Vol. 8, No. 1, pp. 239-253, January 2013

Face Representation

- Face detected and then aligned: can we measure their similarity yet?
- **Yes** – but with respect to what measurement?
- Primitive technique #1:
  - Feature vector from facial measurements such as the distance between eyes, nose, mouth (i.e, anthropometric measurements)
- Primitive technique #2:
  - Using the image pixels values as your feature vector

Anthropometric features contain little information:

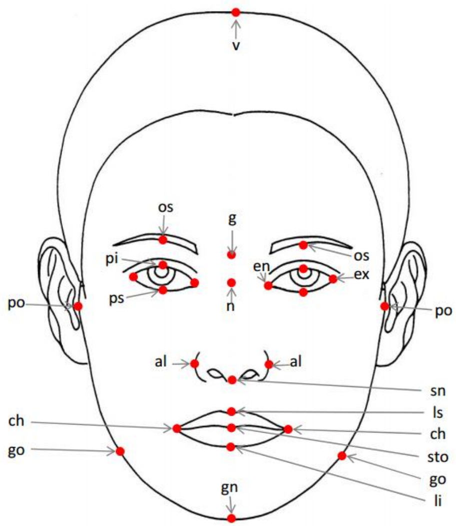

Raw pixels values redundant, sensitive to variates:

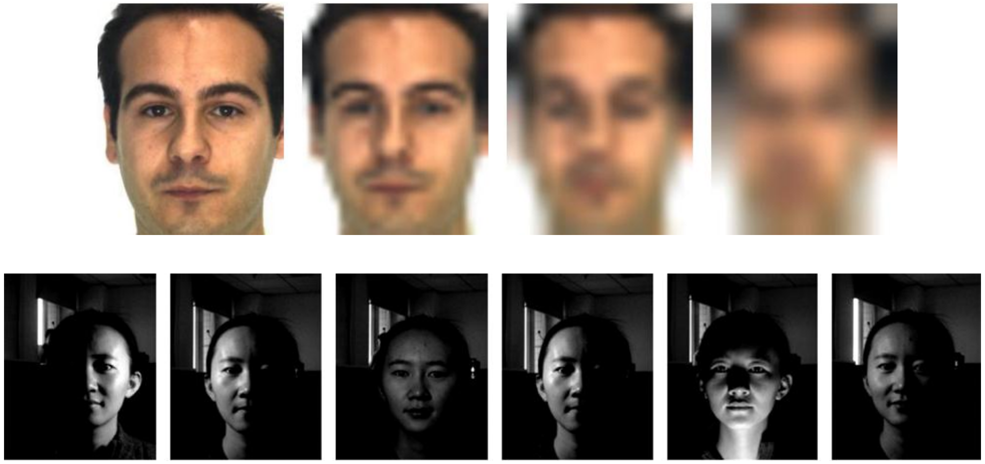

Facial Feature Taxonomy

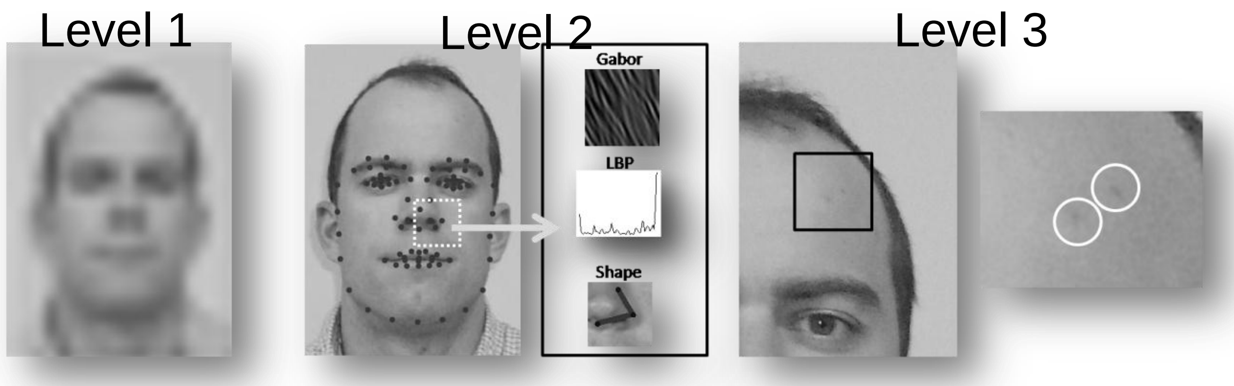

- Facial feature taxonomy provided to organize facial feature representations [1]
- Facial individuality models needed for legal admissibility of FR evidence
- Organized in a similar manner as fingerprints

	Source	
	Humans and machine	Machine Only
Level 1	gender, race, age	appearance-based methods (PCA, LDA, etc.)
Level 2	anthropometric features	distribution-based feature descriptors (LBP, SIFT, etc.), shape distribution models, texture descriptors
Level 3	moles, scars, freckles, birth marks	high spatial frequency

[1] B. Klare and A. K. Jain, "On a Taxonomy of Facial Features," Proc. of IEEE Conf. on Biometrics: Theory, Applications and Systems (BTAS), 2010.

Local Binary Patterns

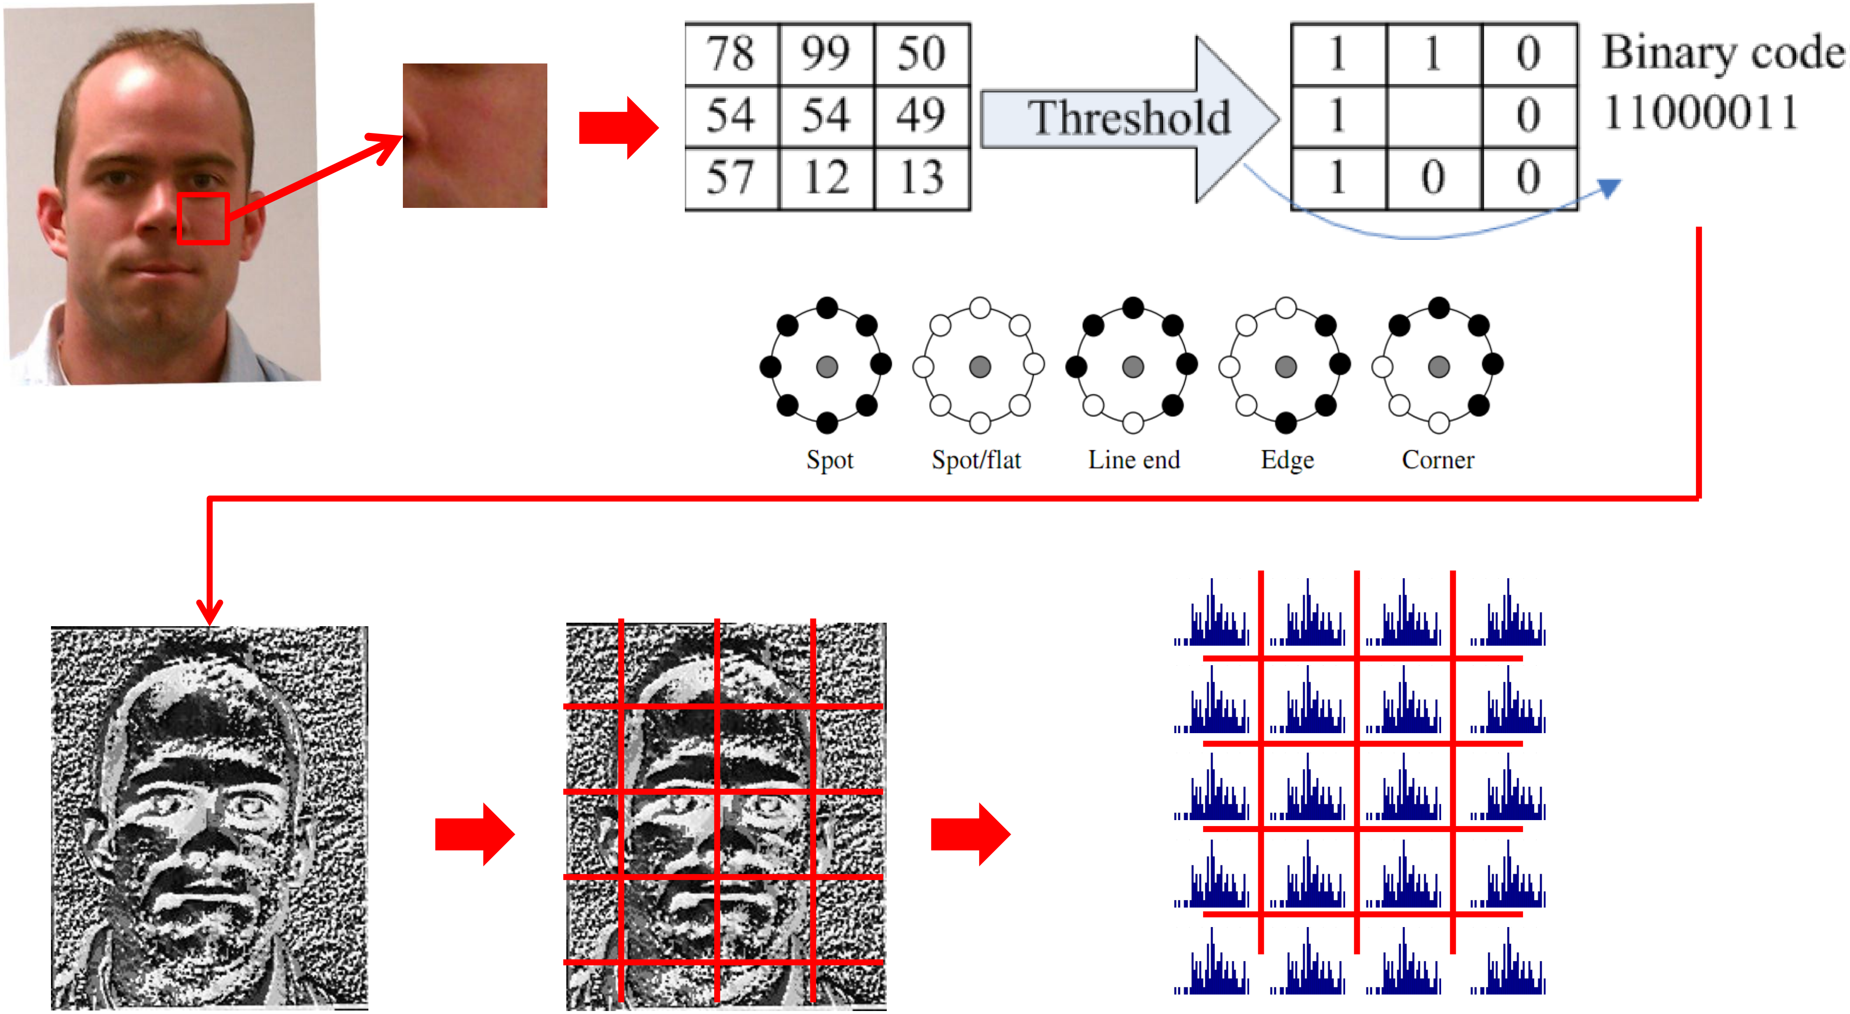

T. Ojala, et al. “Multiresolution Gray-Scale and Rotation Invariant Texture Classification with Local Binary Patterns,” TPAMI, 2002

Scale Invariant Feature Transform (SIFT) Feature Descriptor

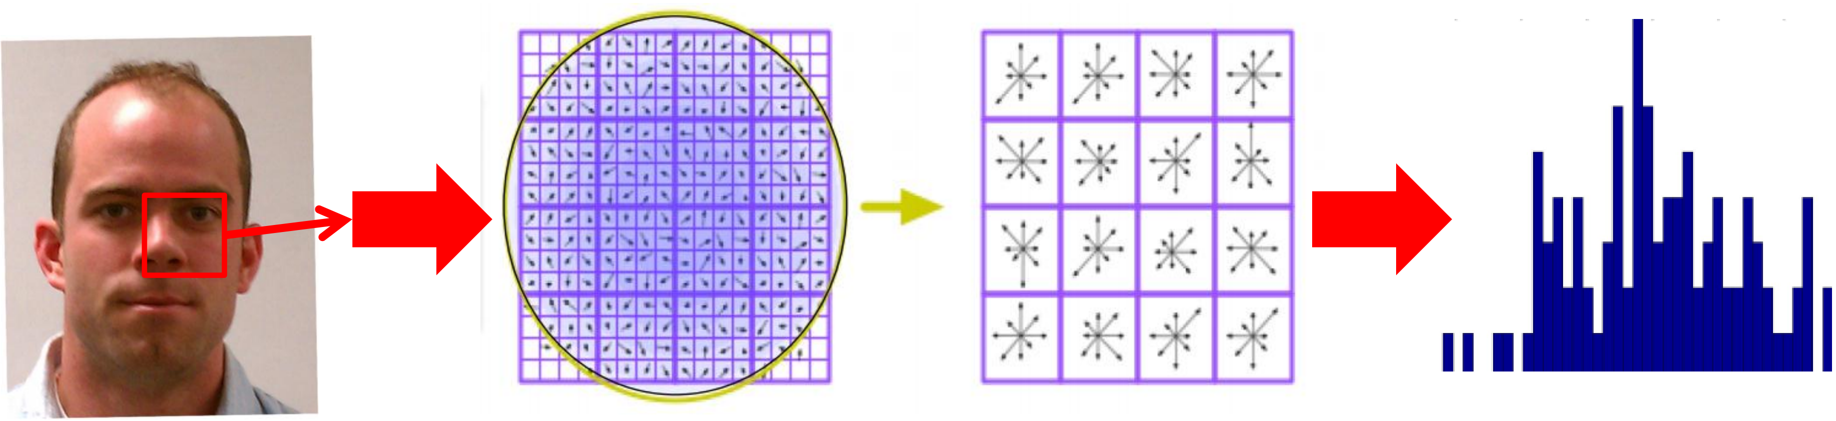

Image Gradients   Quantized Gradients   Histogram Descripto

D. Lowe, “Distinctive Image Features from Scale-Invariant Keypoints,” IJCV, 2004

Statistical Learning / Feature Extraction

- Feature extraction is a critical stage of face recognition which performs statistical learning to (ideally) discover optimal feature weighting
- Earliest approach was PCA (Eigenfaces) [1]
- Used PCA to significantly reduce feature dimensionality

Geometric Interpretation of PCA:

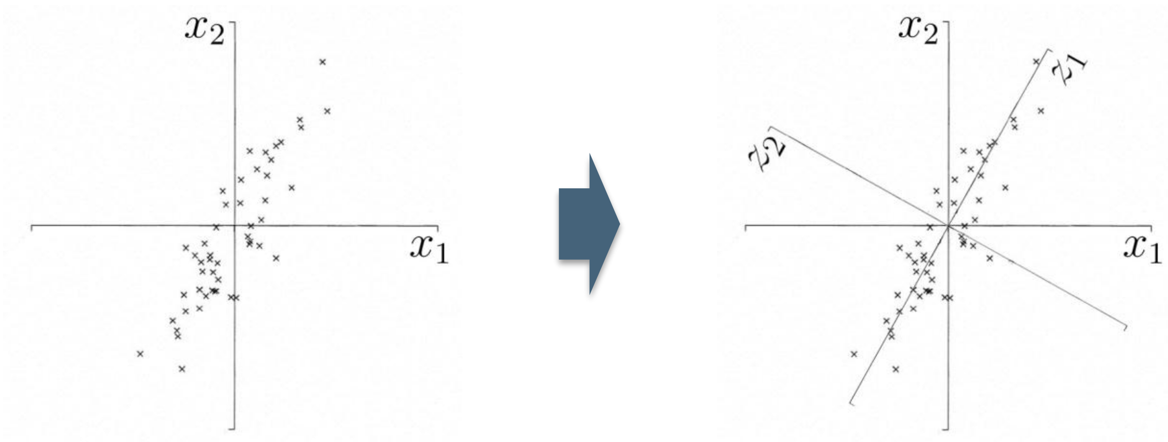

[1] M. Turk and A. Pentland, "Eigenfaces for recognition," Journal of Cognitive Neuroscience, Vol. 3, No. 1, pp. 71-86, 1991.

How Eigenface Method Works:

- Let's consider a set of N face images $\{x^1, x^2, \ldots, x^N\}$ where each x_k is a n-dimensional vector.
- Compute the total scatter matrix S_T as
- Or, if we consider a **n**-by-**N** matrix $\mathbf{X=[(x^1-\mu)\,(x^2-\mu)\,\ldots\,(x^N-\mu)]}$, $\mathbf{S_T}$ can be calculated as $\mathbf{S_T=XX^T}$
- Solve eigen decomposition: $\mathbf{S_T W = \lambda W}$
- Where **W** is matrix of eigenvectors, and **λ** is diagonal matrix of eigenvalues
- Keep top **m** eigenvectors, such that a predetermined amount of variance (e.g 98%) is retained (determined from eigenvalues)

$$S_T = \sum_{k=1}^{N} (x^k - \mu)(x^k - \mu)^T \qquad \mu = 1/N \sum_{k=1}^{N} x^k$$

How Eigenface Method Works:

Input image (size 250x200 -> 50,000 –d vector)

Project input image along eigenvectors shown

Coefficients from eigenvector projections becomes new image representation (Here, top 8 eigenvectors retained -> 8-d vector)

[56.4 38.6 -19.7 9.8 -5.9 1.6 -2.4 0.2]

Distance-based matching can be performed in this subspace

Benefit of PCA/Eigenface Method

- May significantly reduce feature vector size without lose of discriminative information
- Allows compact data storage, improved matching speeds, and linearly independent features
- Unsupervised method -> **does not improve matching accuracy**

Accuracy of PCA vs original feature representation (pixels):

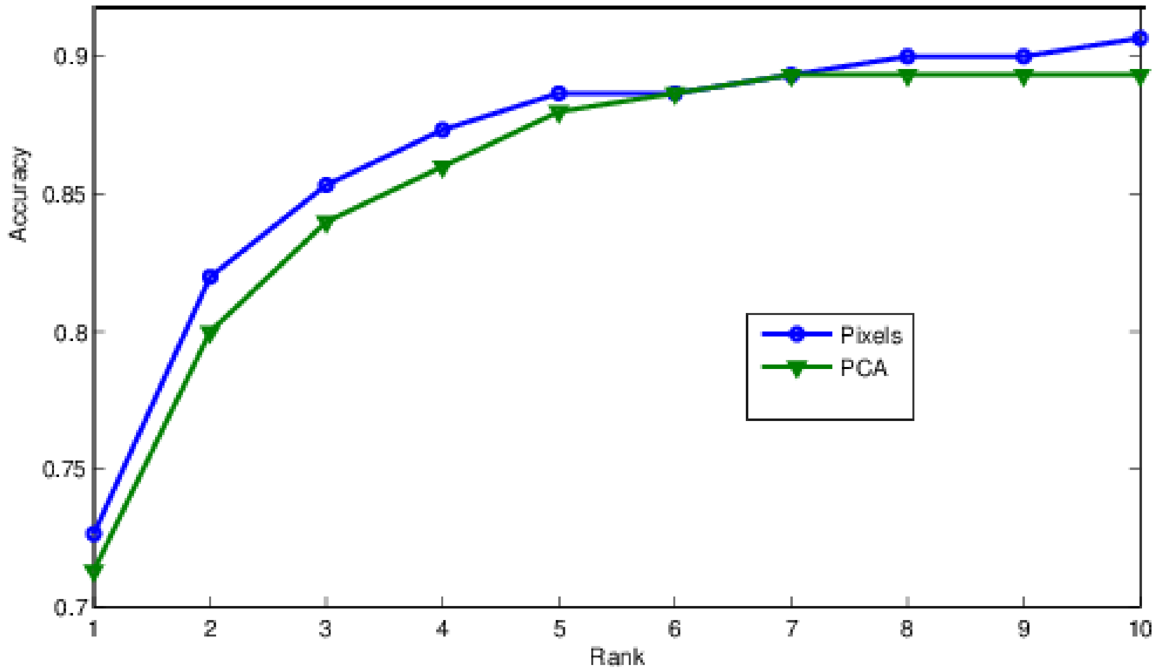

noblis

Problem with PCA

- PCA objective function is to retain maximal energy/variations
- But variations between face images may be due to environmental (illumination) or intrinsic (pose, expression) variations
  - i.e. they may be unwanted
- PCA is unsupervised, and will learn projections that maximizes this variation

Certain variations, such as illumination, should not be retained

FisherFaces

- FisherFaces was seminal approach that used Fisher's Linear Discriminant Analysis (LDA) technique for feature extraction [1]
- Supervised learning method
- Seeks learn subspace W that maximizes Fisher criterion:
  - det($\mathbf{W}^T\mathbf{S}_b\mathbf{W}$) / det ($\mathbf{W}^T\mathbf{S}_w\mathbf{W}$)
- $\mathbf{S}_b$ is between-class scatter, $\mathbf{S}_w$ and is within-class scatter

Geometric Interpretation of LDA:

LDA projection perfectly separates classes/subjects in this example

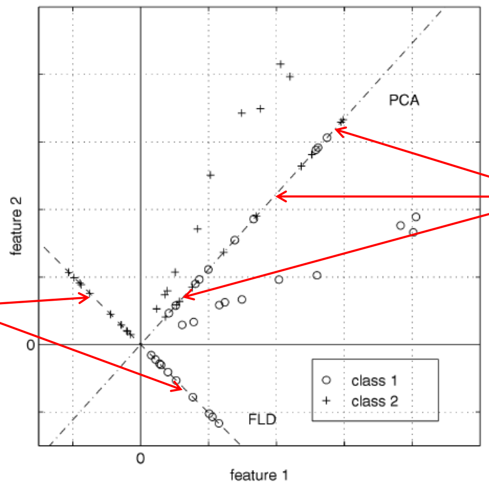

Projection onto principal component (PCA) prevents recognition of subjects

[1] P. Belhumeur, J. Hespanha, and D. Kriegman, "Eigenfaces vs. Fisherfaces: Recognition using class specific linear projection," IEEE Trans. Pattern Analysis and Machine Intelligence, Vol. 19, pp. 711-720, 1997.

PCA vs. LDA

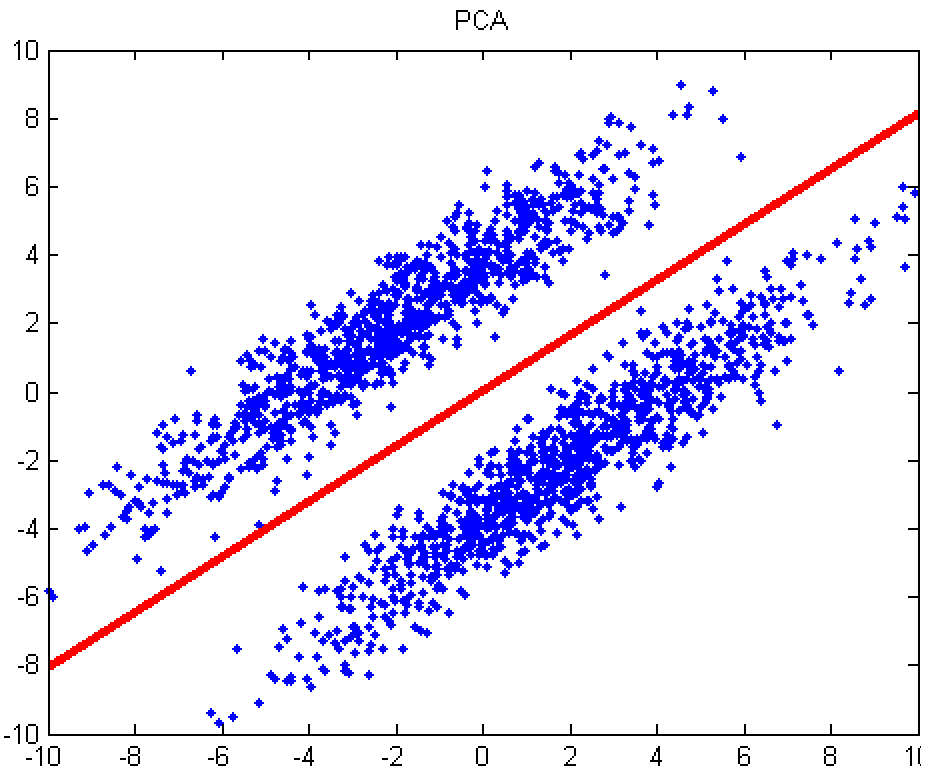

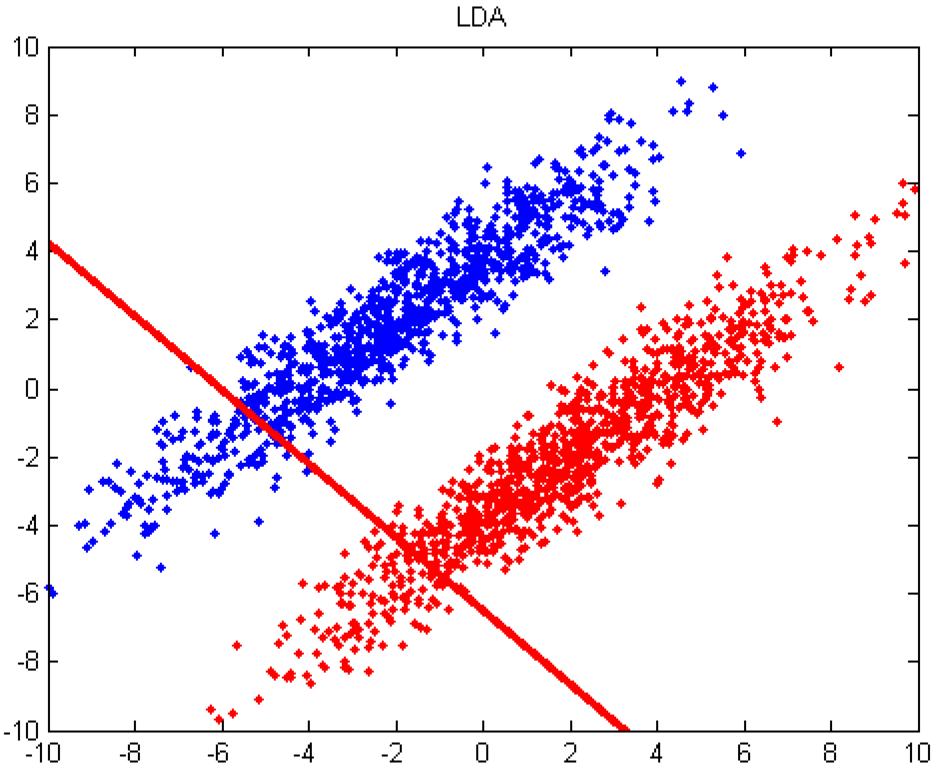

PCA: Find a transformation *w*, such that the $w^T x$ is dispersed the most (maximum distribution)

LDA: Find a transformation w, such that the $w^T X_1$ and $w^T X_2$ are maximally separated & each class is minimally dispersed (maximum separation)

Face Recognition Basics

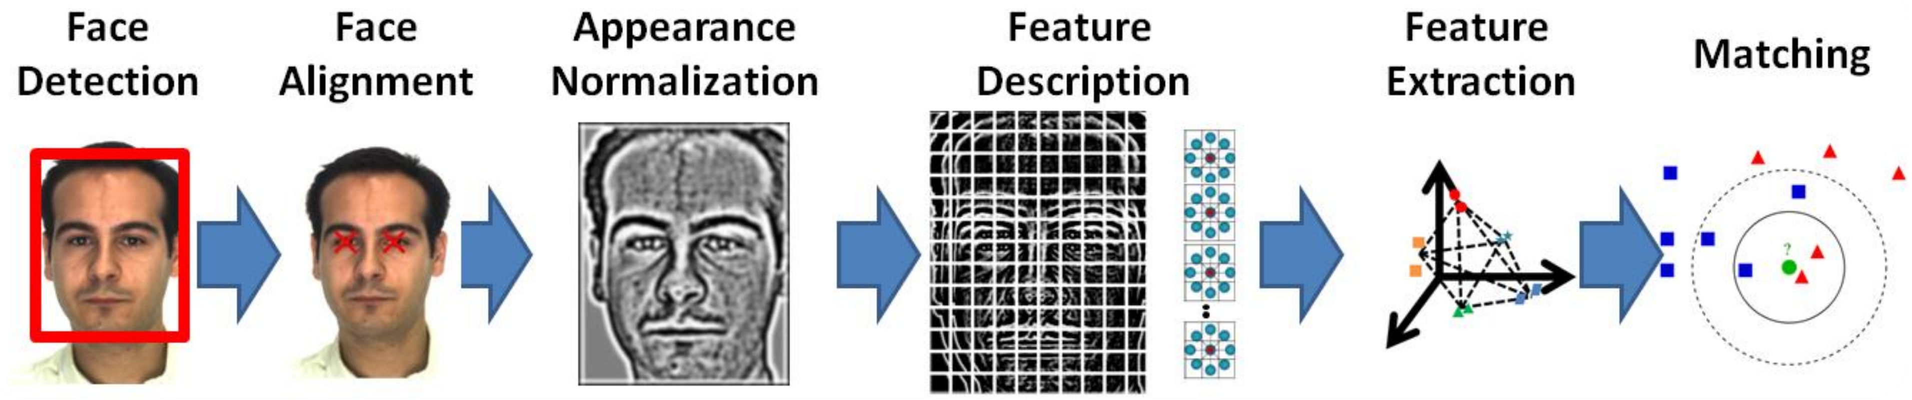

- Discussed many of the basic building blocks of face recognition algorithms
- 100's of papers on each individual topic
- Great challenge in integrating each component as well
- Changes in early stages of the pipeline may effect the later stages
  - Sometimes a matter of different parameters needed at later stages
  - Other times requires entirely new algorithms at later stages

noblis

Seminal Advances in Face Recognition

EigenFaces [1]

1990

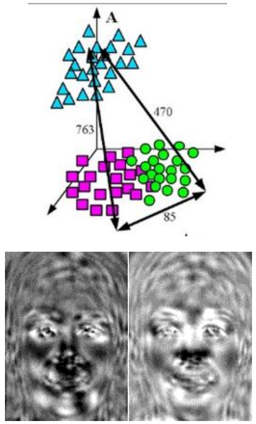

Fisherfaces [2]

1995

Active Appearance Models [3]

2000

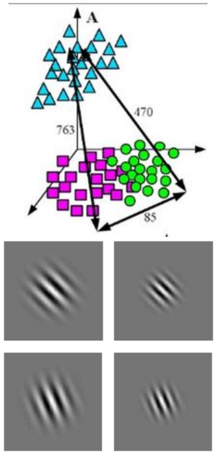

Gabor + LDA [4]

2002

Local Binary Patterns [5]

2005

[1] Turk, Matthew, and Alex Pentland. "Eigenfaces for recognition." *Journal of cognitive neuroscience* 3.1 (1991): 71-86.
[2] Belhumeur, Peter N., Joao P. Hespanha, and David J. Kriegman. "Eigenfaces vs. fisherfaces: Recognition using class specific linear projection." *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 19.7 (1997): 711-720.
[3] Cootes, Timothy F., Gareth J. Edwards, and Christopher J. Taylor. "Active appearance models." Pattern Analysis and Machine Intelligence, IEEE Transactions on 23.6 (2001): 681-685.
[4] Liu, Chengjun, and Harry Wechsler. "Gabor feature based classification using the enhanced fisher linear discriminant model for face recognition." *Image processing, IEEE Transactions on* 11.4 (2002): 467-476.
[5] Ahonen, Timo, Abdenour Hadid, and Matti Pietikainen. "Face description with local binary patterns: Application to face recognition." *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 28.12 (2006): 2037-2041.

Heterogeneous Face Recognition (HFR)

A frontal photograph image exists for the majority of the population

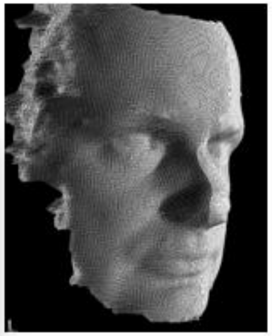

Many security and intelligence scenarios necessitate identification from different modalities of face images (e.g. forensic sketch, infrared image)

Matching non-photograph face images (probe images) to large databases of frontal photographs (gallery images) is called ***heterogeneous face recognition (HFR)***. Current technology does not support this scenario.

HFR Use Cases

- HFR one of the most challenging problems in face recognition due to high intra-class variability due to change in modality
- Successful solutions greatly expand the opportunities to apply face recognition technology
- Common modalities:
  - Sketch – faciliates FR when no face image exists
  - NIR – nighttime and controlled condition face capture, close to visible spectrum
  - Thermal – passive sensing method, highly covert

Examples of heterogeneous face images:

Forensic sketches:

Near infrared:

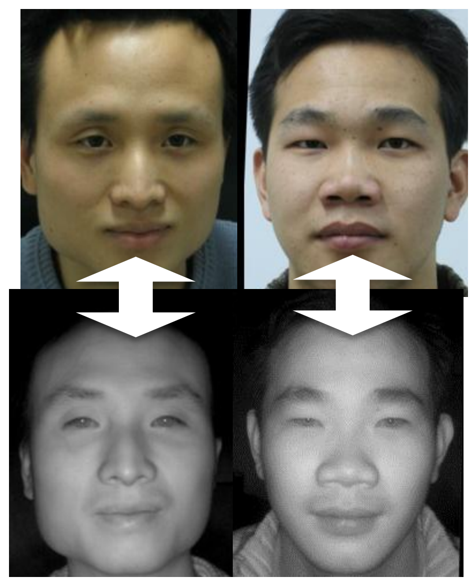

Thermal infrared:

Matching Forensic Sketches to Mug Shot Photographs

- Forensic sketches are drawn by a police artist based on verbal description provided by witness/victim
- Useful when no surveillance video or other biometric data available
- FR engines do not perform well in matching sketch to photo
- FR capabilities need to be enhanced to identify these high value targets
- Early discovery was the invariance of SIFT feature descriptors between sketch and photo [1]
- Prior research only focused on "viewed sketches" [2]

Forensic Sketch Examples

Sketches: *Mugshots:*

Viewed Sketch Examples:

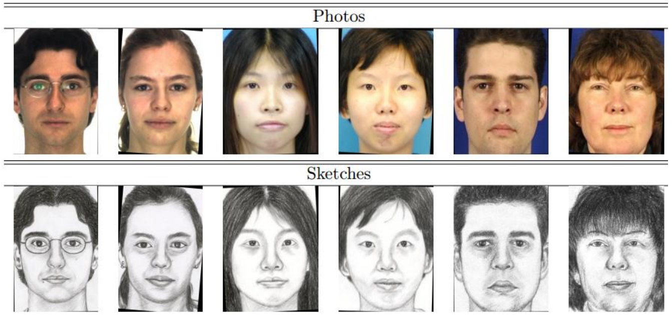

[1] B. Klare and A. K. Jain, "Sketch to Photo Matching: A Feature-based Approach",Proc of SPIE, Biometric Technology for Human Identification VII, April 2010.
[2] Wang, X. and Tang, X., "Face photo-sketch synthesis and recognition," IEEE Trans. Pattern Analysis & Machine Intelligence 31, 1955{1967 (Nov. 2009).

Two Difficulties in Sketch Matching

- **Inaccurate Sketches:**
  - Sketches are drawn from human memory
  - May cause inaccurate description of the suspect
    - i.e. sketch may not even look like the same person

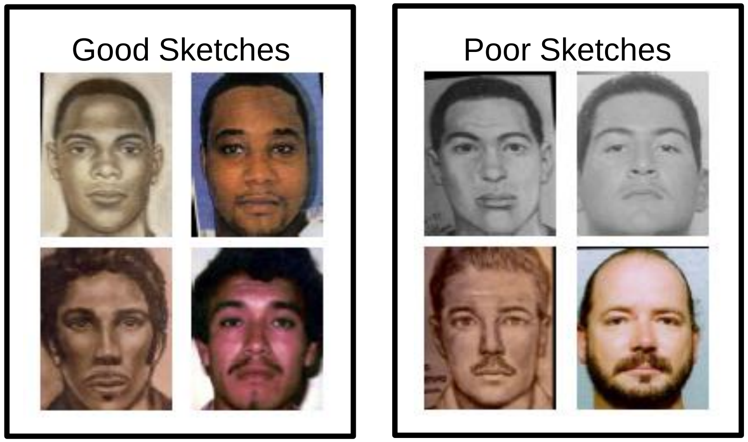

- **Different image modalities:**
  - Cannot directly compare a sketch to a photograph
  - Though accurate, the sketch has a different “appearance”

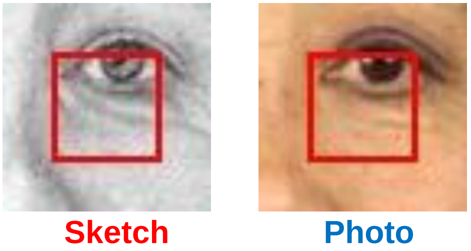

Framework for Matching Forensic Sketches

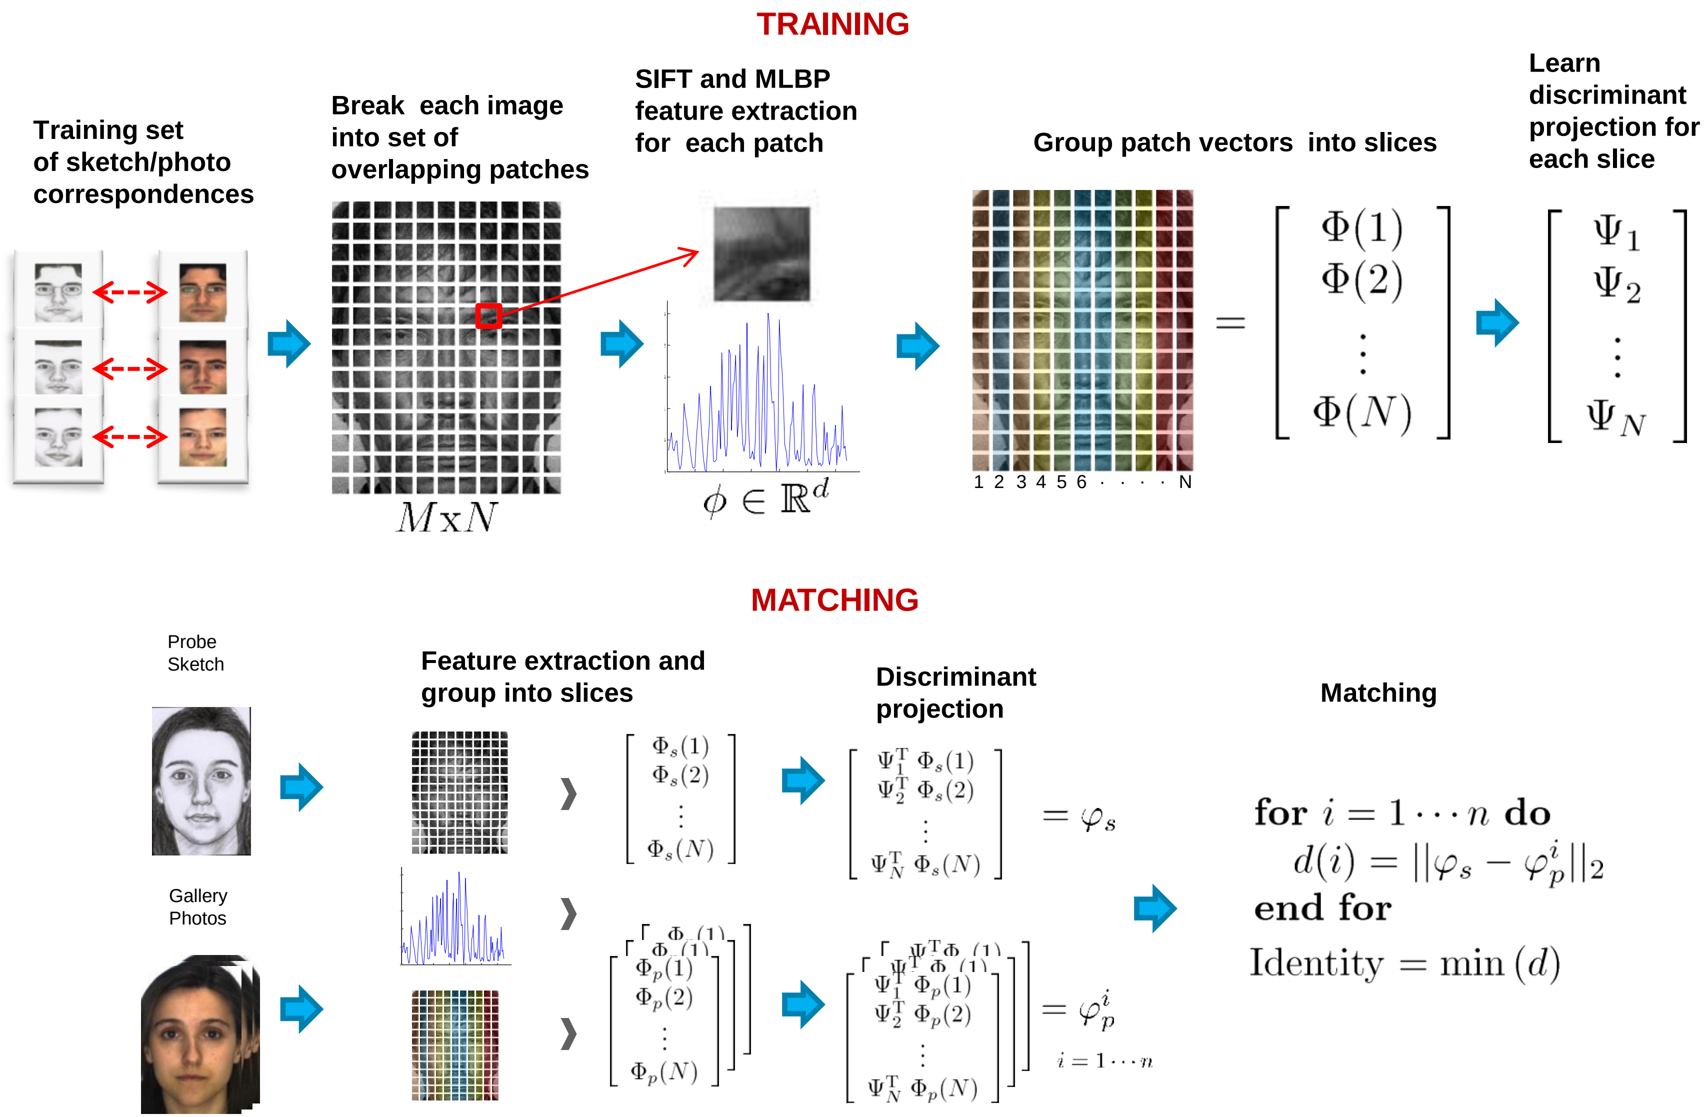

[1] B. Klare, Z. Li, and A. K. Jain, "Matching Forensic Sketches to Mugshot Photos", IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 33, No. 3, pp. 639-646, March 2011.

Sketch Recognition Experiments

- Sketch database:
  - 159 total pairs of mated sketches and photos
  - Labeled as "Good" (49 pairs) and "Poor" sketches (100 pairs) based on resemblance to photograph
- Matched against an additional 10,000 mugshot images provided by the MSP
- Baselined against a leading commercial face recognition system:
  - Cognitec's FaceVACS (a top performer in NIST MBE)

Good vs Poor

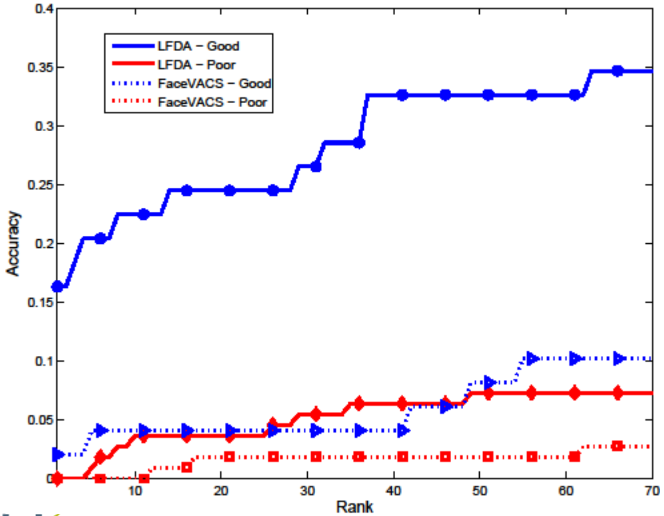

Demographic Filtering

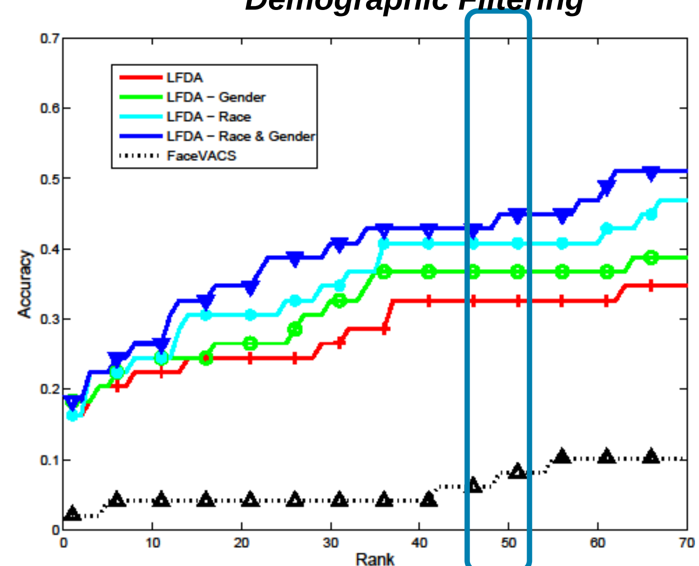

Failed Examples

- Most failed matches were due to poorly drawn sketches with little resemblance to the true photo:

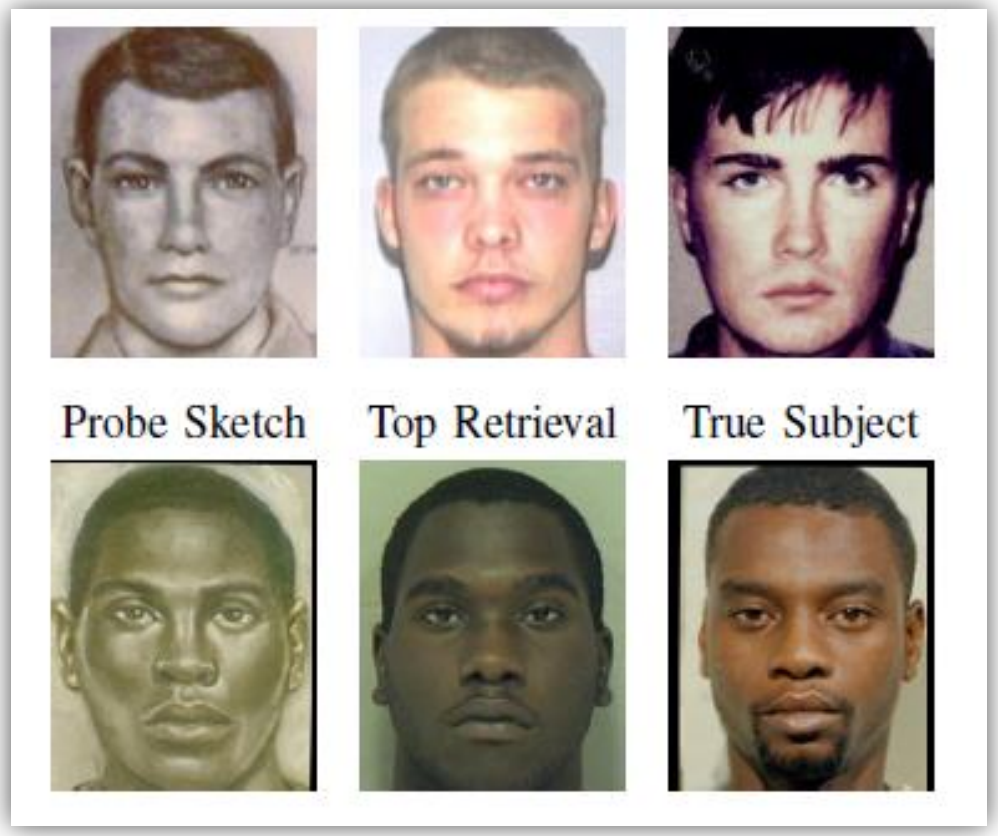

Sketch Recognition From Video

The New York Times

Los Angeles Officials Identify Video Assault Suspects

*“**Composite drawings** of four of the suspects have been made **based upon video images**”*

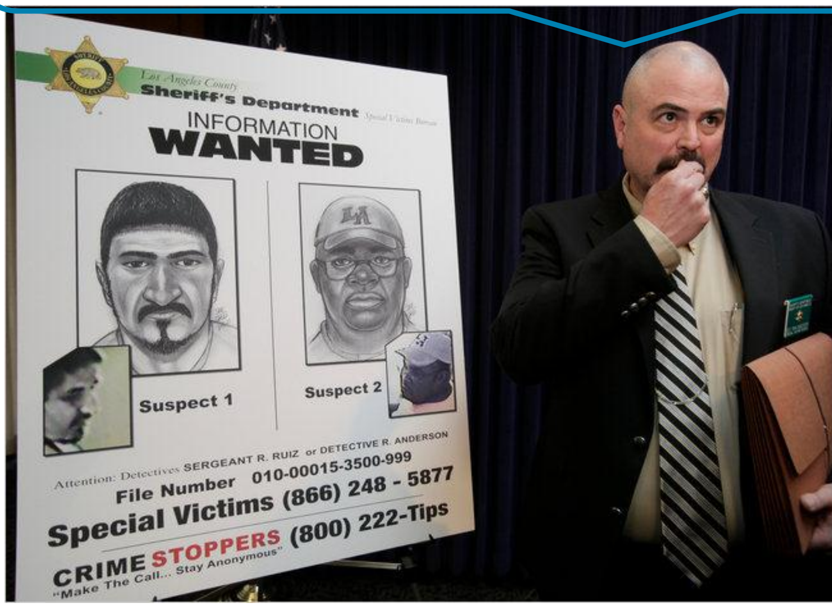

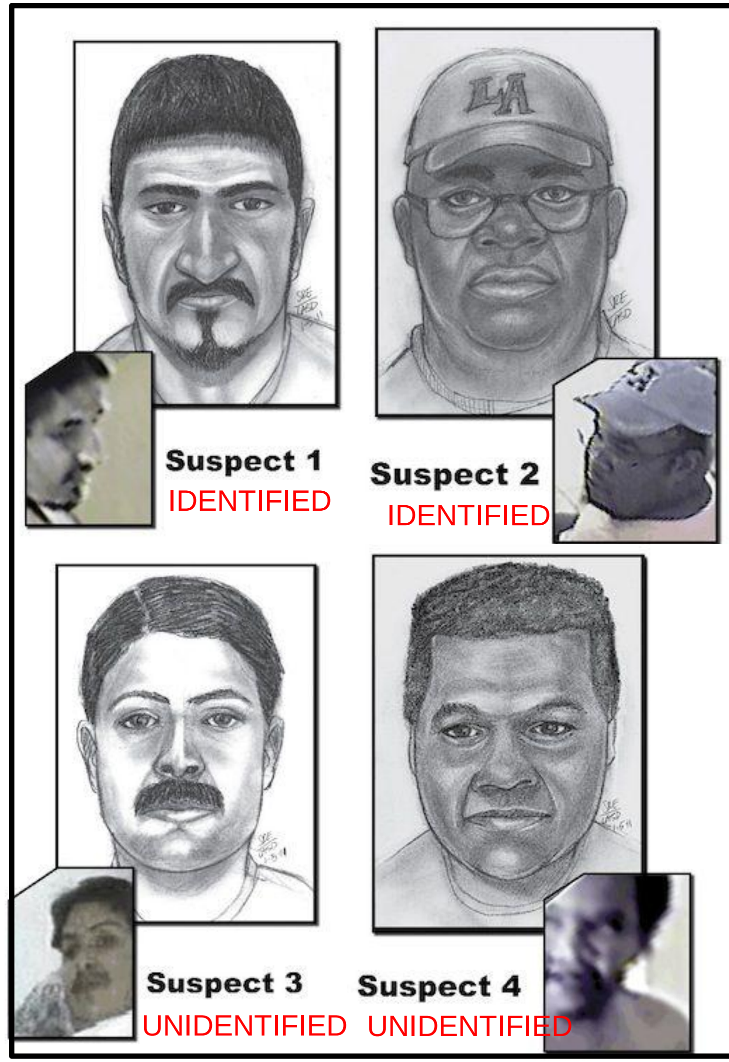

http://www.nytimes.com/2011/01/08/us/08disabled.html

The sketches shown were drawn by Sandra Enslow, LA Sheriff's epartment

Prototype-based Approach to Heterogeneous Face Recognition

- Feature-based methods have demonstrated high accuracy on sketch and NIR matching
- However, other scenarios (e.g. thermal, 3D) do not have invariant feature descriptors
- We seek a generic method for HFR that is not specific to any specific modality
- Proposed prototype representation achieves this goal

Prototype-based approach to HFR:

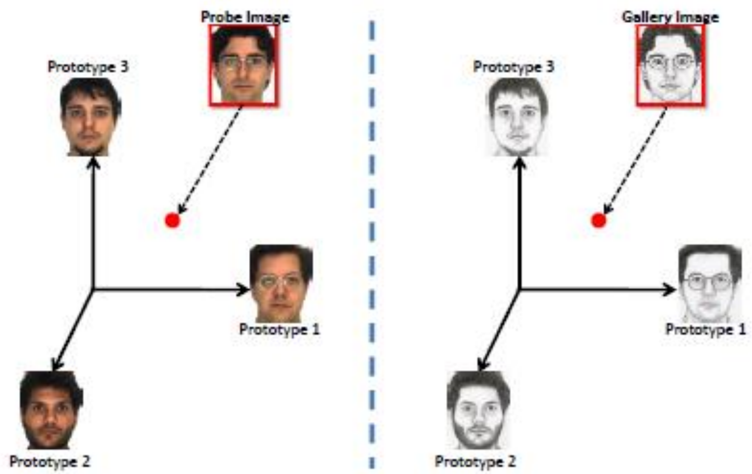

HFR using Kernel Prototypes

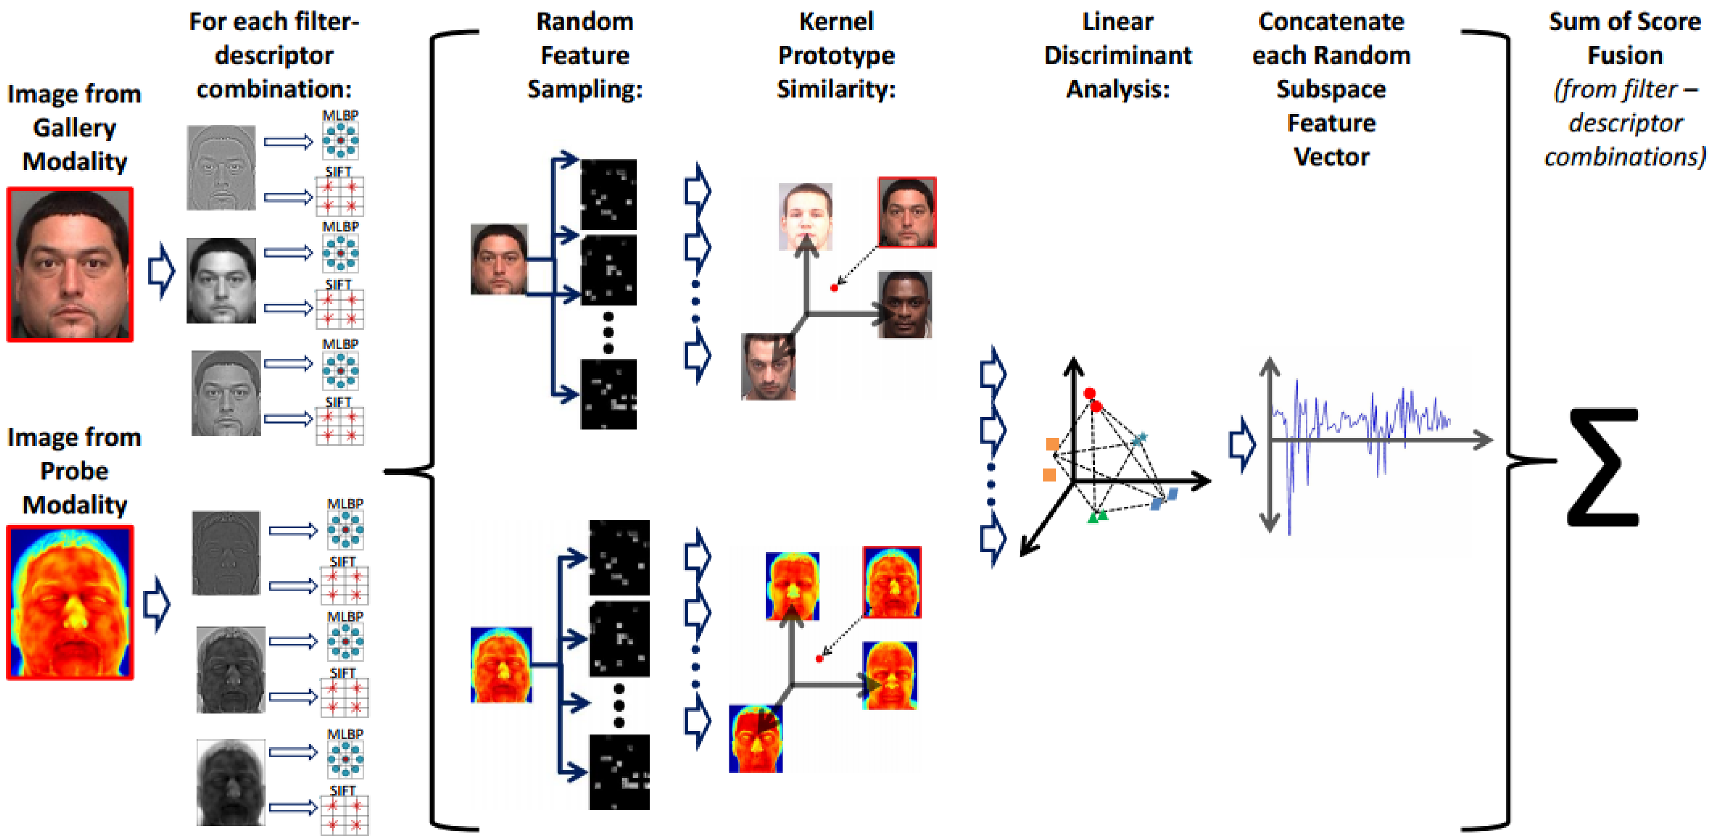

B. Klare and A. K. Jain, "Heterogeneous Face Recognition using Kernel Prototype Similarities", IEEE Transactions on Pattern Analysis and Machine Intelligence, 2013 (to appear)

HFR Experiments

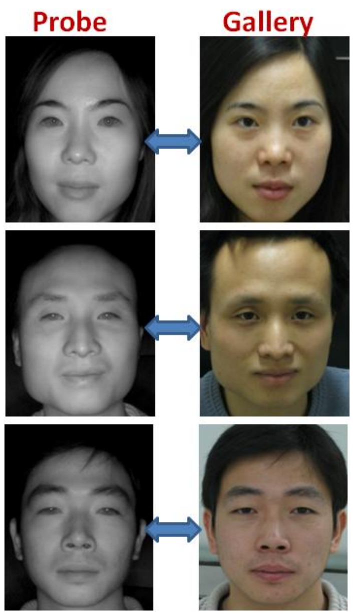

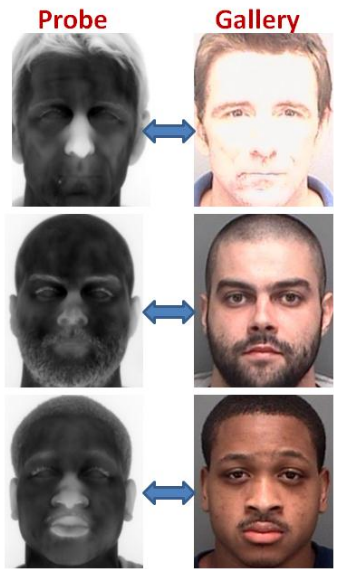

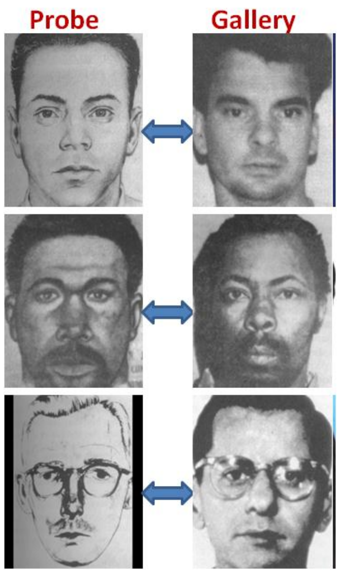

CASIA HFB Dataset

- 200 Probe (NIR) & Gallery (photo) images; 5 splits of 133 training, 67 testing
- Background Gallery: 10,000 mugshot images

Thermal FR

- Dataset collected at PCSO
- 1,000 Probe (thermal) & Gallery (photo); 5 splits of 667 training and 333 testing
- Background Gallery: 10,000 mugshot images

CUHK Sketch Dataset

- 606 Probe (viewed sketch) & Gallery (photo); 5 splits of 404 training and 202 testing
- Background Gallery: 10,000 mugshot images

Forensic sketches

- 159 Probe (forensic sketch) & Gallery (photo); 5 splits of 106 training and 53 testing
- Background Gallery: 10,000 mugshot images

noblis

Recognition Results

Rank-1 Accuracy (%)

Method	NIR	Thermal	Sketch	Forensic*	Standard
P-RS	87.8 ± 4.53	46.7 ± 2.41	74.6 ± 5.42	14.7 ± 1.69	92.5 ± 1.91
D-RS	66.6 ± 6.97	41.5 ± 0.98	96.4 ± 1.54	17.4 ± 3.10	93.7 ± 0.20
(P-RS)+(D-RS)	86.6 ± 4.35	49.2 ± 1.90	92.5 ± 3.52	20.8 ± 2.07	93.0 ± 1.05
FaceVACS	87.8 ± 4.14	21.5 ± 0.83	84.8 ± 2.05	1.9 ± 1.03	98.7 ± 0.40

* *Results for forensic sketch are the Rank-50 accuracy.*

(a)

TAR @ FAR = 1.0%

Method	NIR	Thermal	Sketch	Forensic	Standard
P-RS	98.2 ± 1.63	76.4 ± 2.55	99.5 ± 0.35	14.7 ± 3.38	96.8 ± 0.52
D-RS	94.0 ± 3.50	77.5 ± 1.22	99.6 ± 0.41	17.7 ± 6.88	97.9 ± 0.34
(P-RS)+(D-RS)	97.0 ± 2.36	78.2 ± 0.13	99.7 ± 0.27	18.9 ± 2.31	97.6 ± 0.34
FaceVACS	93.7 ± 1.63	47.5 ± 2.49	92.2 ± 1.50	2.6 ± 1.03	99.5 ± 0.40

(b)

TAR @ FAR = 0.1%

Method	NIR	Thermal	Sketch	Forensic	Standard
P-RS	95.8 ± 6.15	71.2 ± 12.94	99.0 ± 1.25	12.9 ± 5.39	95.5 ± 2.55
D-RS	90.8 ± 8.52	72.5 ± 12.25	99.5 ± 0.44	15.7 ± 7.89	96.7 ± 2.41
(P-RS)+(D-RS)	94.5 ± 6.45	72.7 ± 13.47	99.4 ± 0.73	16.0 ± 7.23	96.4 ± 2.41
FaceVACS	92.0 ± 4.39	44.4 ± 7.85	89.6 ± 6.45	2.5 ± 0.97	99.1 ± 1.02

(c)

OpenBR and the 4SF Algorithm

- Open source biometrics recognition project, OpenBR released (http://openbiometrics.org/)
- Offers suite of image processing, computer vision, and machine learning algorithms used to perform face recognition
- Can perform roughly 3.8 million facial comparisons, per second (per CPU thread)
- Currently participating in NIST FRVT 2012
- Highest accuracy algorithm based on Spectrally Sampled Structural Subspace Features (4SF) algorithm [1]
  - Has been used in other face recognition research projects [2][3][4]

Overview of the 4SF algorithm:

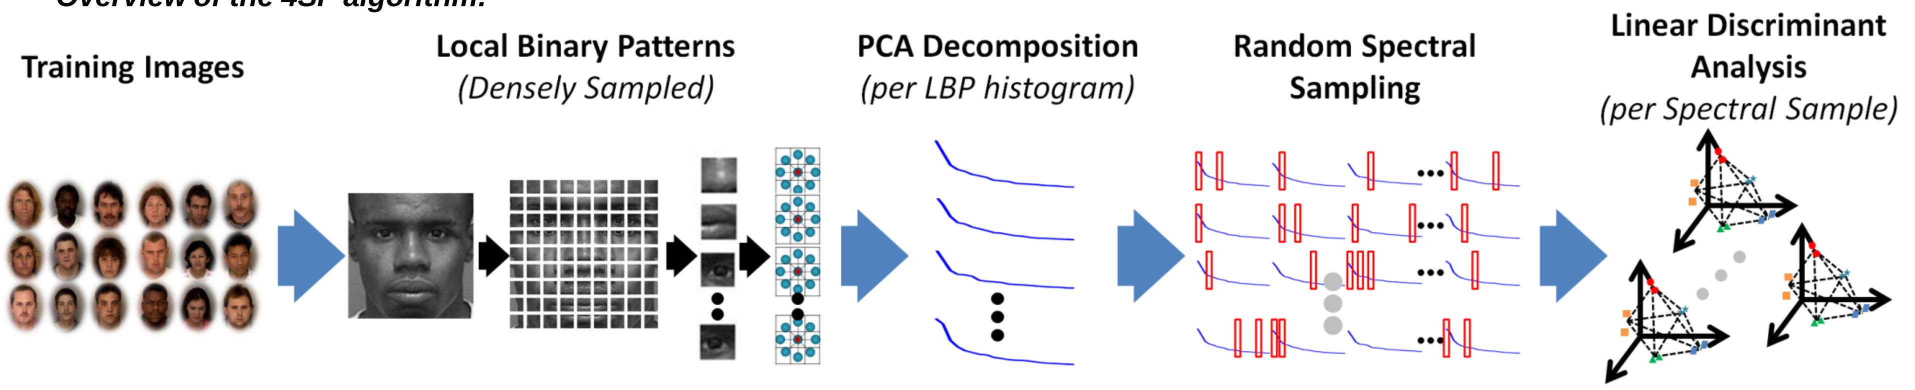

[1] B. F. Klare, “Spectrally sampled structural subspace features (4SF),” in Michigan State University Technical Report, MSU-CSE-11-16, 2011.
[2] B. F. Klare, M. Burge, J. Klontz, R. W. Vorder Bruegge, and A. K. Jain, "Face Recognition Performance: Role of Demographic Information", IEEE Transactions on Information Forensics and Security, Vol. 7, No. 6, pp. 1789-1801, December 2012.
[3] B. F. Klare and A. K. Jain, "Face Recognition: Impostor-based Measures of Uniqueness and Quality", Proceedings of the IEEE Conference on Biometrics: Theory, Applications, and Systems (BTAS), 2012.
[4] B. F. Klare and A. K. Jain, "Face Recognition Across Time Lapse: On Learning Feature Subspaces", Proceedings of the International Joint Conference on Biometrics, 2011.

Aging Invariant Face Recognition

- Invariance to facial aging is a significant challenge in face recognition
- **Aging-invariant FR algorithms** must **learn features** and/or synthesize appearances that offset facial variations over time
- State-of-the-art approaches rely heavily of training data
- Using 200,000 mug shot face images of 64,000 subjects, trained five version of 4SF algorithm [1]

Experimental Design:

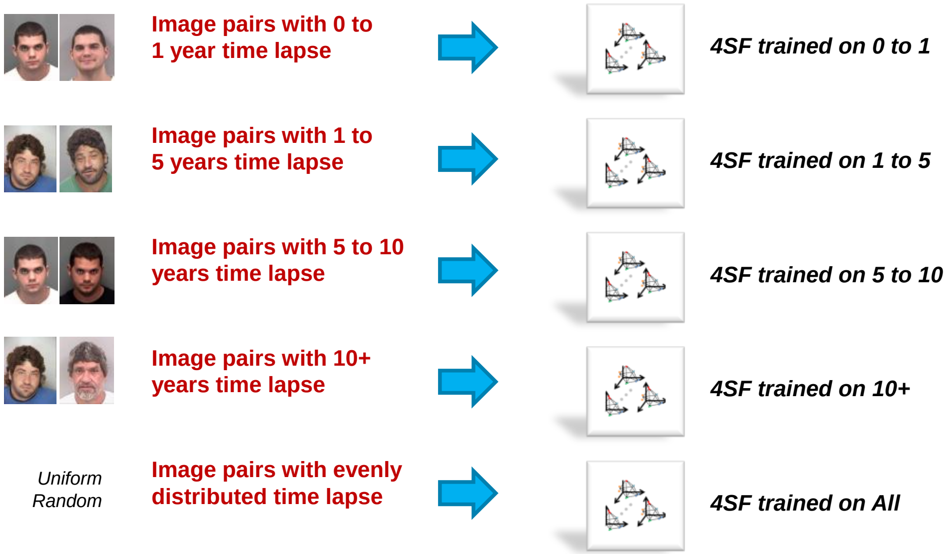

Example of the aging process:

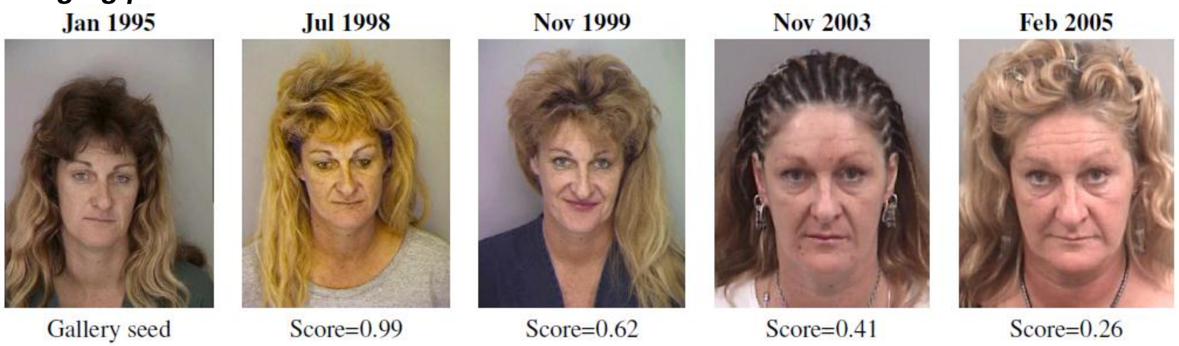

[1] B. Klare and A. K. Jain, "Face Recognition Across Time Lapse: On Learning Feature Subspaces", Proc. of IEEE *Int Joint Conference on Biometrics (IJCB)* 2011.

noblis

Experimental Results

Test set: 0 to 1 year time lapse

RS-LDA trained on (time lapse in years):					Baselines:		
(0-1)	(1-5)	(5-10)	(10+)	(All)	MLBP Only	COTS1	COTS2
94.5%	94.1%	93.1%	91.8%	94.1%	71.2%	96.3%	89.8%

# of Match Comparions: *19,996*
# of Non-Match Comparions: *239,572,034*

Test set: 1 to 5 year time lapse

RS-LDA trained on (time lapse in years):					Baselines:		
(0-1)	(1-5)	(5-10)	(10+)	(All)	MLBP Only	COTS1	COTS2
90.3%	**90.5%**	89.1%	87.7%	90.2%	62.9%	94.3%	84.6%

# of Match Comparions: *33,443*
# of Non-Match Comparions: *401,282,557*

Test set: 5 to 10 year time lapse

RS-LDA trained on (time lapse in years):					Baselines:		
(0-1)	(1-5)	(5-10)	(10+)	(All)	MLBP Only	COTS1	COTS2
75.2%	81.2%	**82.0%**	80.4%	81.3%	46.7%	88.6%	75.5%

# of Match Comparions: *24,036*
# of Non-Match Comparions: *215,795,208*

Template update per-aging time lapse:

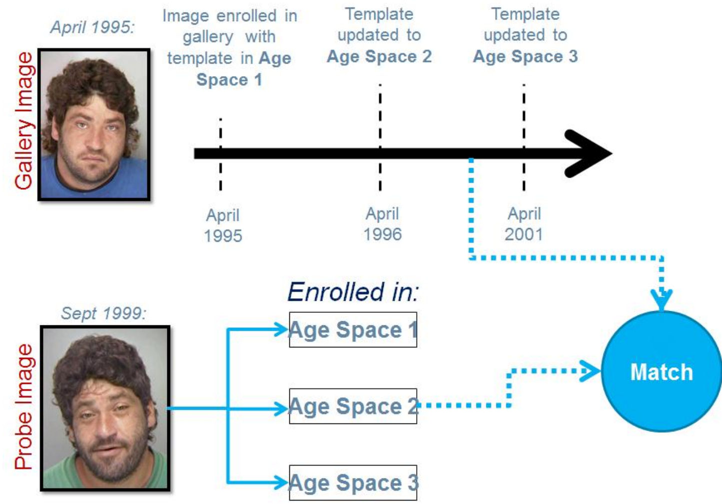

Face Recognition Across Demographics

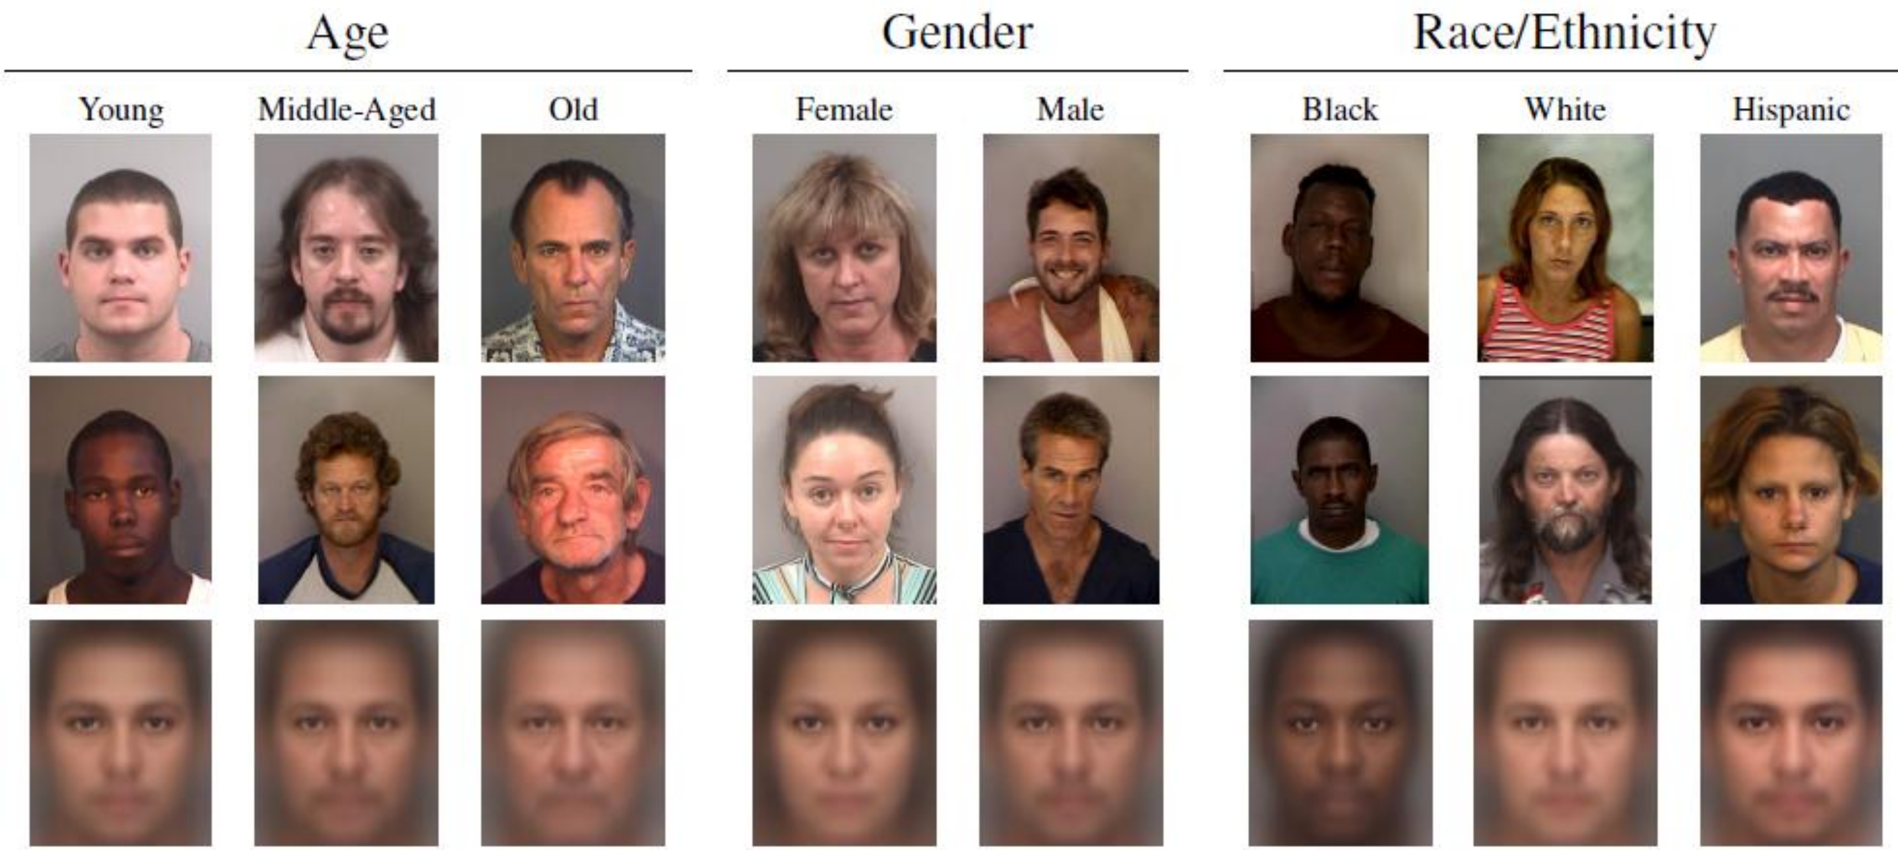

- Different demographics have been shown to be more difficult to recognize [1]

[1] P. J. Grother, G. W. Quinn, and P. J. Phillips, "MBE 2010: Report on the evaluation of 2D still-image face recognition algorithms," National Institute of Standards and Technology, NISTIR, vol. 7709, 2010.

Facial Demographics

- Is unbalanced recognition performance on different cohorts a function of:
  - Unbalanced training?
  - Inherent difficulty of the demographic cohort?
- Answered by studying whether we can improve face recognition performance by training exclusively on a cohort [1]
- Analyzed face image from the Pinellas County Sherriff's Office
- Mug shot data with subject demographics
- Dataset partitioned to consists entirely of specific demographic

Number of subjects available for testing and training for each cohort:

Demographic	Cohort	# Training	# Testing
Gender	Female	7995	7996
	Male	7996	7998
Race	Black	7993	7992
	White	7997	8000
	Hispanic	1384	1425
Age	18 to 30	7998	7999
	30 to 50	7995	7997
	50 to 70	2801	2853

Are black subjects more difficult to match, or are matchers not properly trained on black subjects?

What about females?

Algorithms Studied:

- **Commercial of the shelf (COTS):**
  - Cognitec's FaceVACS v8.2
  - PittPatt v5.2.2
  - Neurotechnology's MegaMatcher v3.1
- **Non-trainable:**
  - Local binary patterns (LBP)
  - Gabor-based
- **Trainable:**
  - Spectrally Sample Structural Subspace Features (4SF)

[1] B. F. Klare, M. Burge, J. Klontz, R. W. Vorder Bruegge, and A. K. Jain, "Face Recognition Performance: Role of Demographic Information", IEEE Transactions on Information Forensics and Security, Vol. 7, No. 6, pp. 1789-1801, December 2012.

Experimental Results: Key Findings

Gender:

	Females	Males
COTS-A	89.5	94.4
COTS-B	81.6	89.3
COTS-C	70.3	80.9
LBP	54.4	74.0
Gabor	56.0	68.2
4SF trained on All	73.0	86.2
4SF trained on Females	71.5	85.0
4SF trained on Males	69.0	86.3

Females inherently more difficult to recognize:

- All matchers (three commercial ,LBP and Gabor) the worst on females (with respect to males)
- Training exclusively on females did not improve accuracy
  - i.e. cannot improve on females through training

Race:

	Black	White
COTS-A	88.7	94.4
COTS-B	81.3	89.0
COTS-C	74.0	79.8
LBP	65.3	70.5
Gabor	61.6	63.7
4SF trained on All	78.4	83.0
4SF trained on Black	80.2	81.0
4SF trained on White	75.4	84.5
4SF trained on Hispanic	74.5	80.2

Blacks more difficult, but can be improved:

- All matchers (three commercial ,LBP and Gabor) the worst on blacks (with respect to whites and Hispanics)
- Can improve recognition performance by training on black subjects
- Can also improve on whites by training on whites

Age:

	18 to 30 y.o.	30 to 50 y.o.
COTS-A	91.7	94.6
COTS-B	86.1	89.1
COTS-C	76.5	80.7
LBP	69.4	74.7
Gabor	61.7	68.2
4SF trained on All	81.5	85.6
4SF trained on 18 to 30 y.o.	83.3	85.9
4SF trained on 30 to 50 y.o.	82.1	86.0
4SF trained on 50 to 70 y.o.	78.7	84.5

Young more difficult, but can be improved

- All matchers (three commercial, LBP and Gabor) the worst on younger subjects (with respect to middle-age and old)
- Can improve recognition performance by training on young subjects

Dynamic Face Matcher Selection

- Ability to improve performance on race and age suggest dynamic face matcher selection
- Particularly useful in forensic scenarios where first pass does not yield a successful match

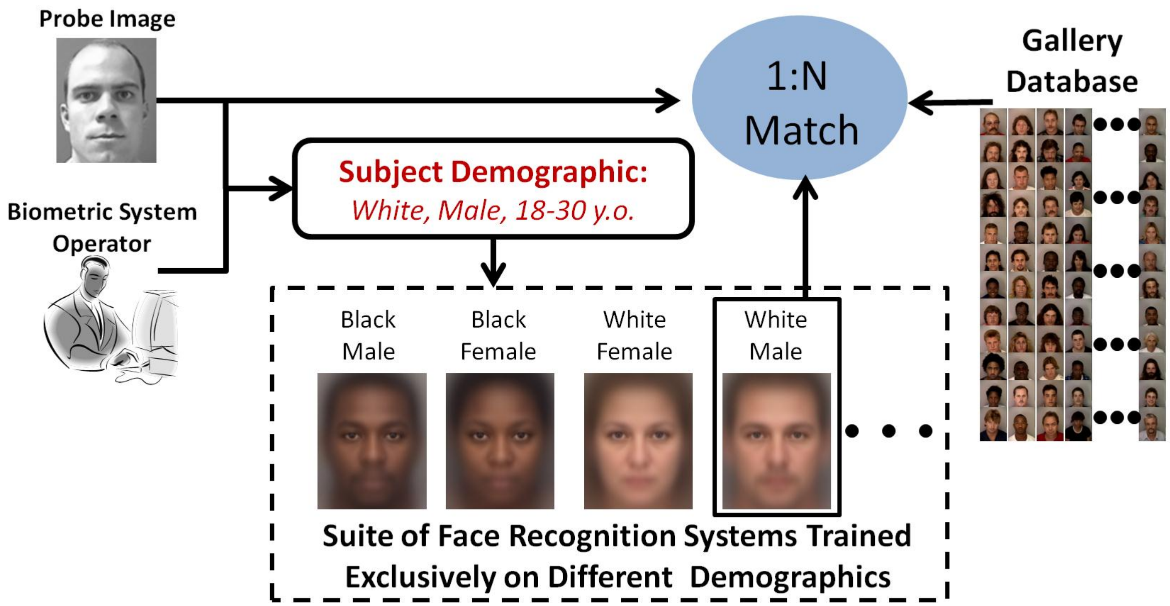

Who am I?

Image drawn by: Grant Pominville

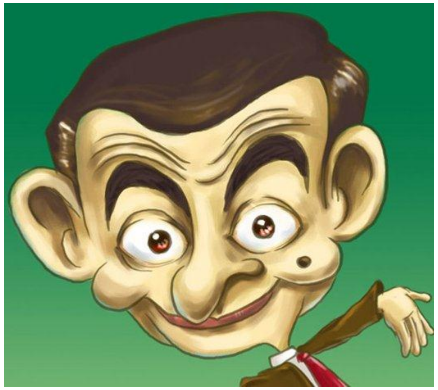

Image drawn by: "Hikari"

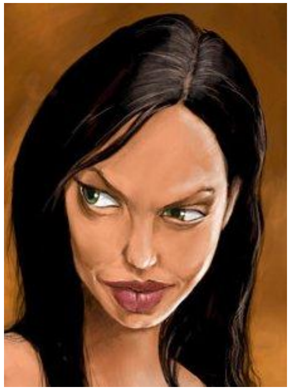

Image drawn by: Rok Dovecar

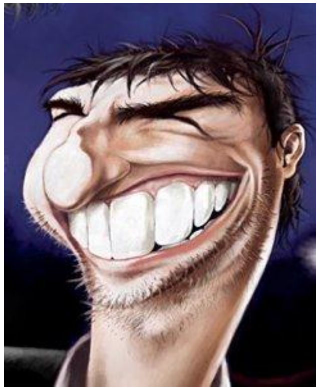

Conceptualizing a Caricature

- Humans found to recognize caricature sketches better than realistic sketches [1] [2]
- A caricature can be thought as an extrapolation between the mean face and the subject's face [3]
- Studies suggest we encode face images as deviations for prototypical face images
- Through exaggeration, caricatures exploit our internal face representation

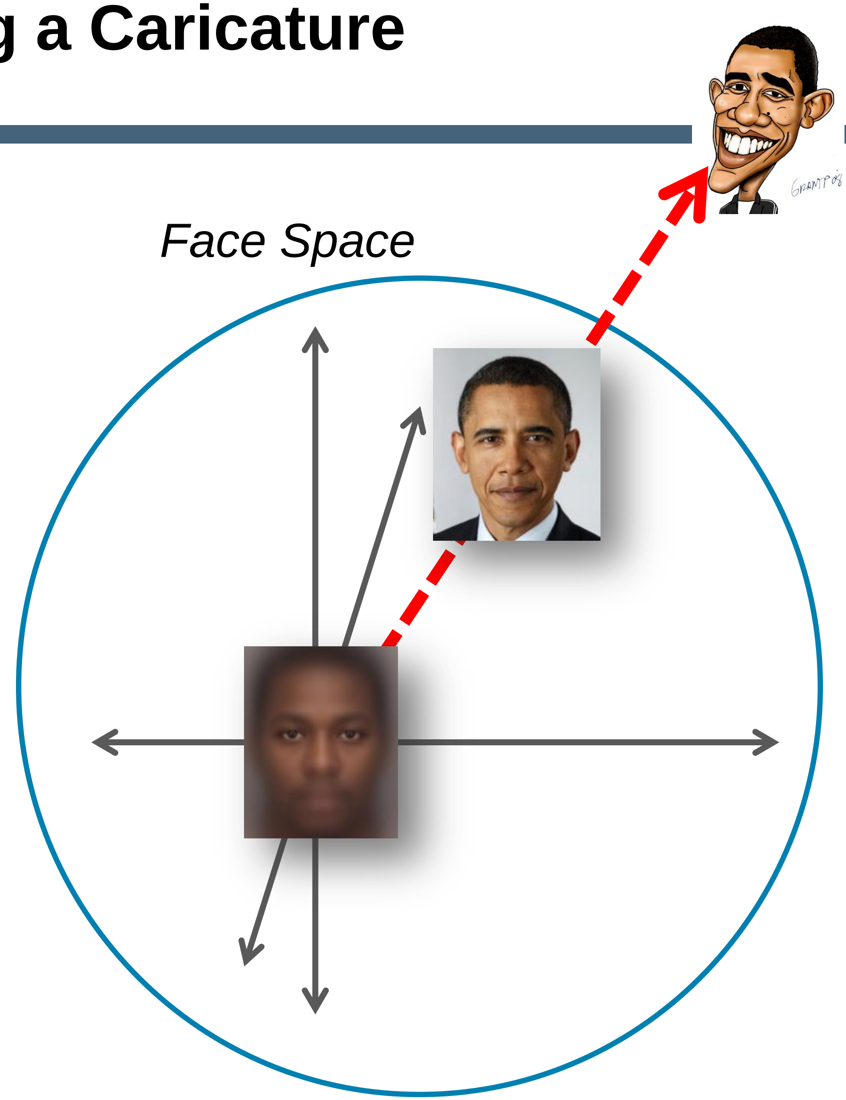

[1] R. Mauro and M. Kubovy. Caricature and face recognition.Memory & Cognition, 20(4):433–440, 1992
[2] G. Rhodes, S. Brennan, and S. Carey. Identification and ratings of caricatures: Implications for mental representations of faces. Cognitive Psychology, 19(4):473–497, 1987
[3] D. A. Leopold, A. J. O'Toole, T. Vetter, and V. Blanz.Prototype-referenced shape encoding revealed by high-levelaftereffects. Nature Neuroscience, 4:89–94, 2001.

Caricature Matching

- Encoded caricatures using qualitative features
- Collected dataset of 196 caricature/photo pairs from internet and fellow artists
- Dataset randomly split into 2/3 training (134 pairs) and 1/3 testing (62 pairs) – (averaged over 10 random splits)
- Plan to use features for face indexing and unconstrained scenarios

Method	TAR @ FAR=10.0%	TAR @ FAR=1.0%	Rank-1
Qualitative Features (no learning):			
NN_{L_2}	39.2 ± 5.4	9.4 ± 2.7	12.1 ± 5.2
Qualitative Features (learning):			
Logistic Regression	50.3 ± 2.4	11.3 ± 2.9	17.7 ± 4.2
MKL	39.5 ± 3.2	7.4 ± 3.9	11.0 ± 3.9
NN_{MKL}	46.6 ± 3.9	10.3 ± 3.6	14.4 ± 2.9
SVM	52.6 ± 5.0	12.1 ± 2.8	20.8 ± 5.6
Logistic Regression+NN_{MKL}+SVM	56.9 ± 3.0	15.5 ± 4.6	23.7 ± 3.5
Image Descriptors (learning):			
LBP with LDA	33.4 ± 3.9	11.5 ± 2.5	15.5 ± 4.6
Qualitative Features + Image Descriptors:			
Logistic Regression+NN_{MKL}+ SVM+LBP with LDA	61.9 ± 4.5	22.7 ± 3.5	32.3 ± 5.1

Caricature recognition framework:

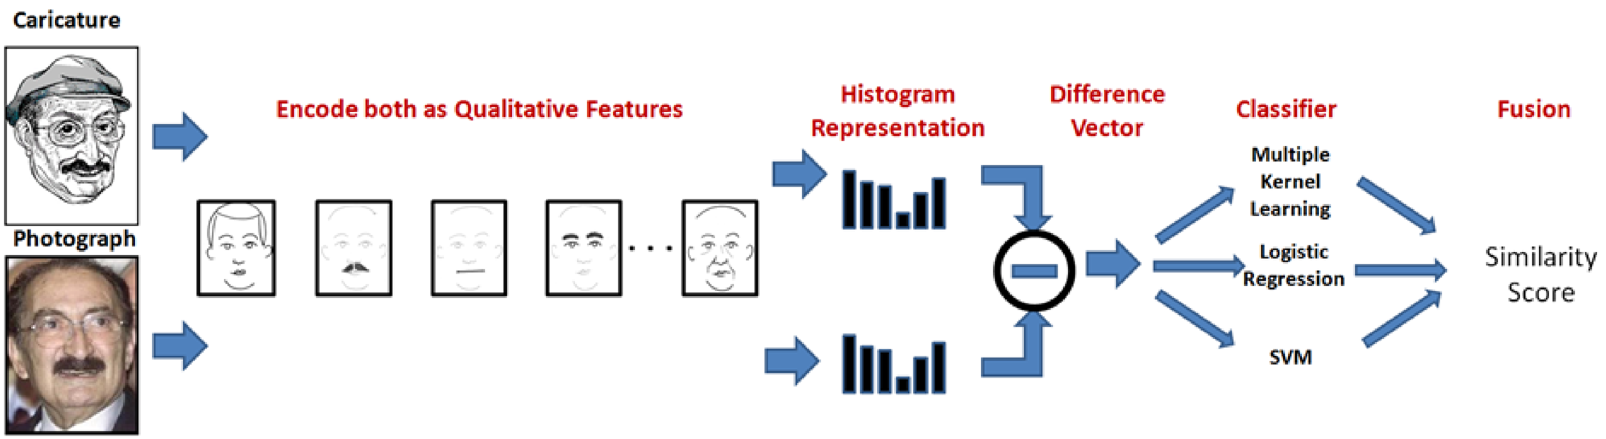

Face Recognition Challenges: Identical Twins

Pairs of identical twins:

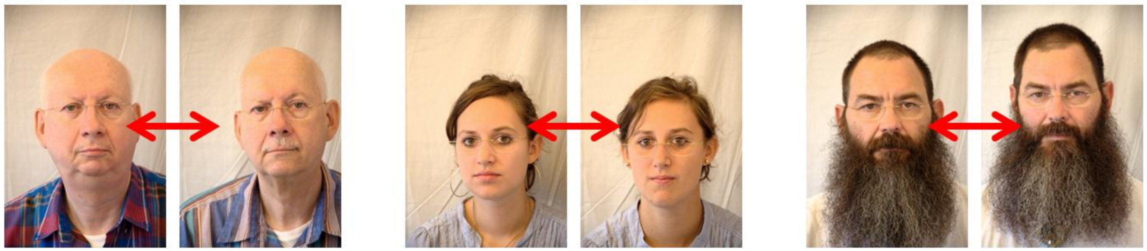

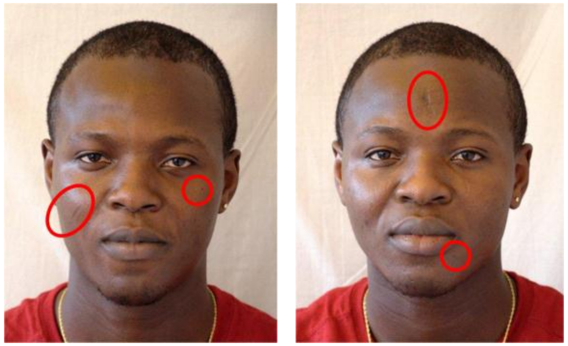

Facial marks (e.g. moles) are unique between identical twin pairs

The New York Times

Identical Twins, One Charged in a Fatal Shooting, Create Confusion for the Police

By MARC LACEY
Published: August 8, 2011

CHANDLER, Ariz. — At first, the murder case against Orlando Nembhard seemed solid, as witness after witness came forward to say they saw someone that looked just like him brandish a pistol in February outside a

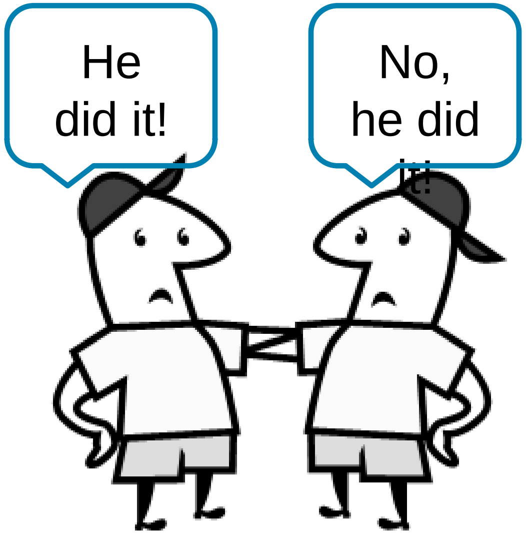

B. F. Klare, A. Paulino, and A. K. Jain, "Analysis of Facial Features in Identical Twins", Proceedings of the International Joint Conference on Biometrics, 2011.

noblis

Quality-based Score Normalization

- Uniqueness-based nonmatch estimates (UNE) framework demonstrates the ability to improve face recognition performance of any face matcher [1]
- Uses novel metric for measuring the uniqueness of a given individual, called the impostor-based uniqueness measure (IUM)
- UNE maps face match scores into non-match probability estimates conditionally dependent IUM
- Framework demonstrates: (i) improved matching accuracy, (ii) improved human interoperability (iii) the predictive ability of IUM towards face recognition accuracy
- Study conducted on an operational dataset with 16,000 subjects using three different face matchers

UNE framework:

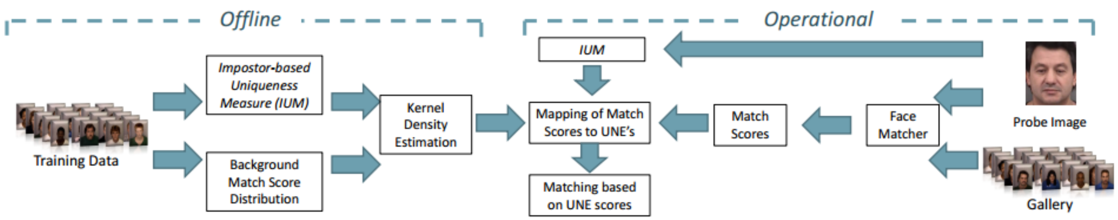

IUM-based match estimates:

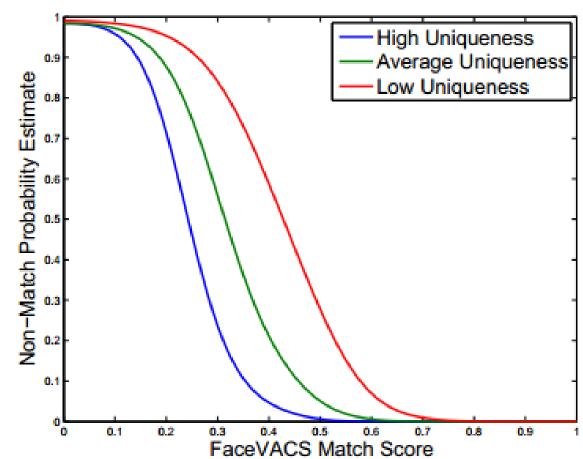

Matching Results:

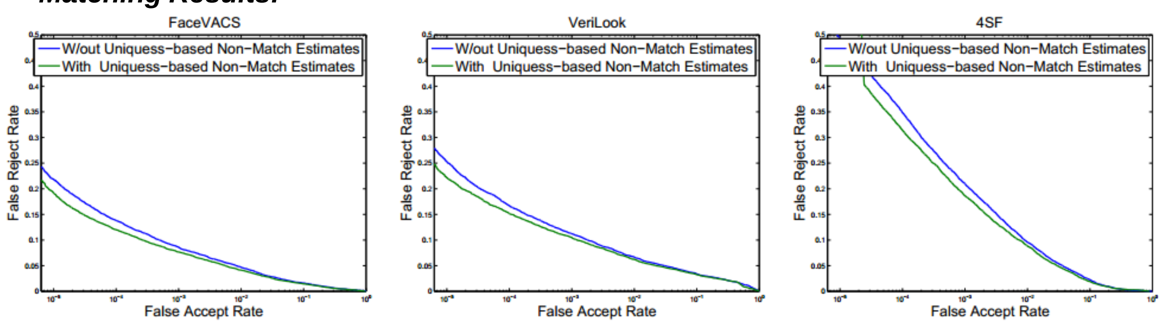

[1] B. Klare and A. K. Jain, "Face Recognition: Impostor-based Measures of Uniqueness and Quality", Proc. of IEEE Conference on Biometrics: Theory, Applications and Systems (BTAS) 2012.

Component-based Face Recognition

- Proposed component-based face alignment and representation framework [1]
- Aligns faces per component, extracts LBP features, and learns RSLDA subspaces
- Motivated by recent evidence from the cognitive science community demonstrating the efficacy of component-based facial representations [2]
- Proposed component-based representations: (i) are more robust to changes in facial pose, and (ii) improve recognition accuracy on occluded face images in forensic scenarios
- Demonstrates need for accurate landmark detection

Component-based approach:

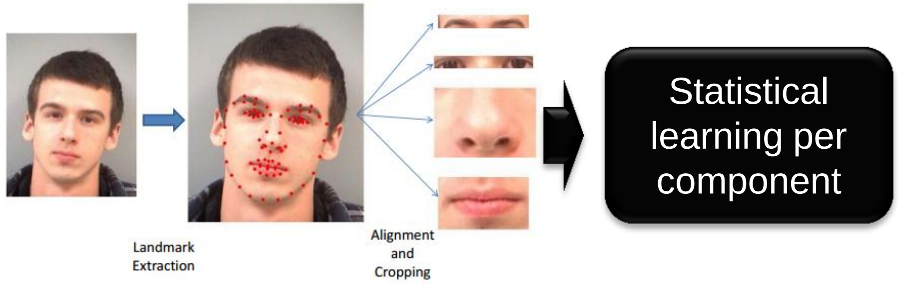

Recognition results on FERET database:

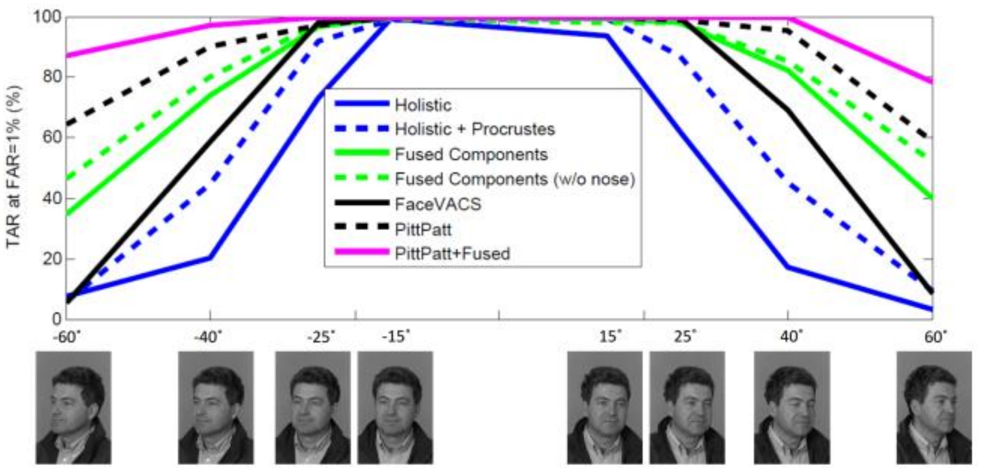

[1] K. Bonnen, B. Klare, and A. K. Jain, "Component-Based Representation in Automated Face Recognition", IEEE Transactions on Information Forensics and Security, Vol. 8, No. 1, pp. 239-253, January 2013.
[2] J. Gold, P. Mundy, and B. Tjan, “The perception of a face is no more than the sum of its parts,” Psychological Sciences, March 2012.

Summary

- Discussed how face recognition systems work, and the need to improve all key stages
- Heterogeneous face recognition algorithms presented to handle sketch and infrared recognition
- Discussed how other forms of heterogeneity (age, demographics, pose) can effect the face recognition process
- This is all moving towards unconstrained face recognition algorithms

Next generation FR algorithms are expected to handle unconstrained face images:

noblis