Fusion and Privacy in Biometrics

Arun Ross

Professor

Michigan State University

rossarun@cse.msu.edu

http://www.cse.msu.edu/~rossarun

The iPRoBe Lab

http://iprobe.cse.msu.edu

https://twitter.com/iPRoBeLab

- Integrated Pattern Recognition and Biometrics Lab
- Currently: 8 PhD Students + 1 Post-Doc + 2 UG
- Graduated: 24 MS Thesis Students + 7 PhD Students


National Science Foundation
WHERE DISCOVERIES BEGIN


IARPA
BE THE FUTURE


NIST
National Institute of
Standards and Technology
U.S. Department of Commerce


FEDERAL BUREAU OF INVESTIGATION
FBI
BIOMETRIC
Center of Excellence
Leading the World of Biometrics


ONR
OFFICE OF NAVAL RESEARCH


NIJ
NATIONAL
INSTITUTE
of JUSTICE


WV HIGHER EDUCATION
POLICY COMMISSION


CITeR
CENTER FOR
IDENTIFICATION
TECHNOLOGY
RESEARCH


BIMA
BIOMETRICS
IDENTITY
MANAGEMENT
AGENCY


U.S. DEPARTMENT OF
HOMELAND SECURITY


noblis
For the best of reasons


DRAPER
LABORATORY


NIH
National Institutes of Health
Turning Discovery Into Health

INTRODUCTION TO BIOMETRICS

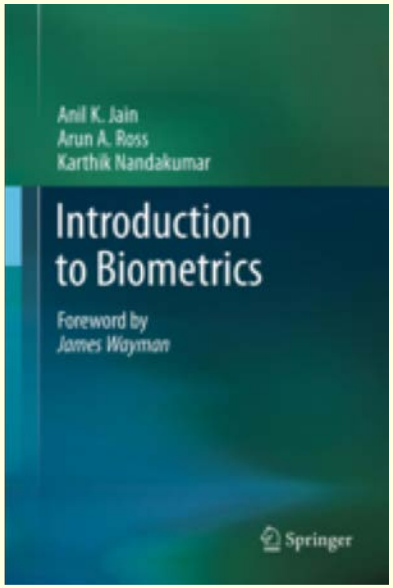

HANDBOOK OF MULTIBIOMETRICS

HANDBOOK OF BIOMETRICS

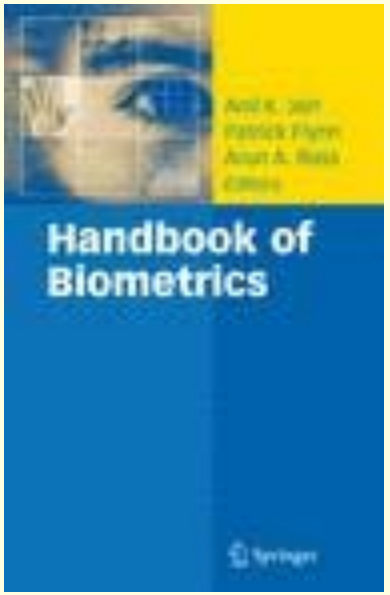

Related Papers

- M. Singh, R. Singh, A. Ross, "**A Comprehensive Overview of Biometric Fusion**," Information Fusion, 2019 (to appear)
- A. K. Jain, K. Nandakumar, A. Ross, "**50 Years of Biometric Research: Accomplishments, Challenges, and Opportunities**," Pattern Recognition Letters, Vol. 79, pp. 80 - 105, August 2016.
- A. Dantcheva, P. Elia, A. Ross, "**What Else Does Your Biometric Data Reveal? A Survey on Soft Biometrics**," IEEE Transactions on Information Forensics And Security (TIFS), Vol. 11, No. 3, pp. 441 - 467, March 2016.
- A. K. Jain and A. Ross, "**Bridging the Gap: From Biometrics to Forensics**," Philosophical Transactions of The Royal Society B, Vol. 370, Issue 1674, August 2015.
- A. K. Jain, B. Klare, A. Ross, "**Guidelines for Best Practices in Biometrics Research**," Proc. of 8th IAPR International Conference on Biometrics (ICB), (Phuket, Thailand), May 2015.

Biometric System

BIOMETRIC TRAIT

HUMAN MACHINE INTERFACE

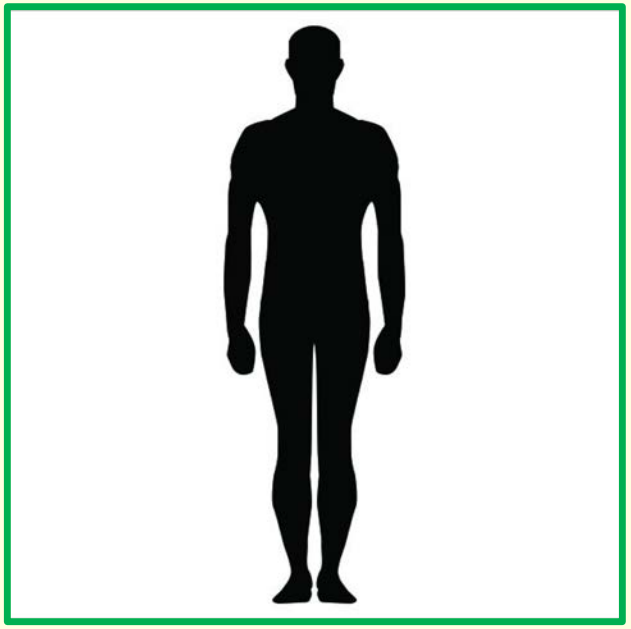

PERSON

Challenges in a Biometric System

- **Noise in sensed data**: e.g., defective sensors or unfavorable ambient/physiological conditions
- **Intra-user variations**: e.g., different types of interaction with sensor, variations in user's biometric trait, sensor characteristics are modified
- **Distinctiveness**: e.g., capacity of biometric template is limited
- **Non-universality**: e.g., all users may not be able to successfully present the trait
- **Presentation attacks**: circumvent the system by imitation or using artificial traits

Non-universality

- Some people may consistently offer poor quality fingerprint images which means they have to be identified by some other means

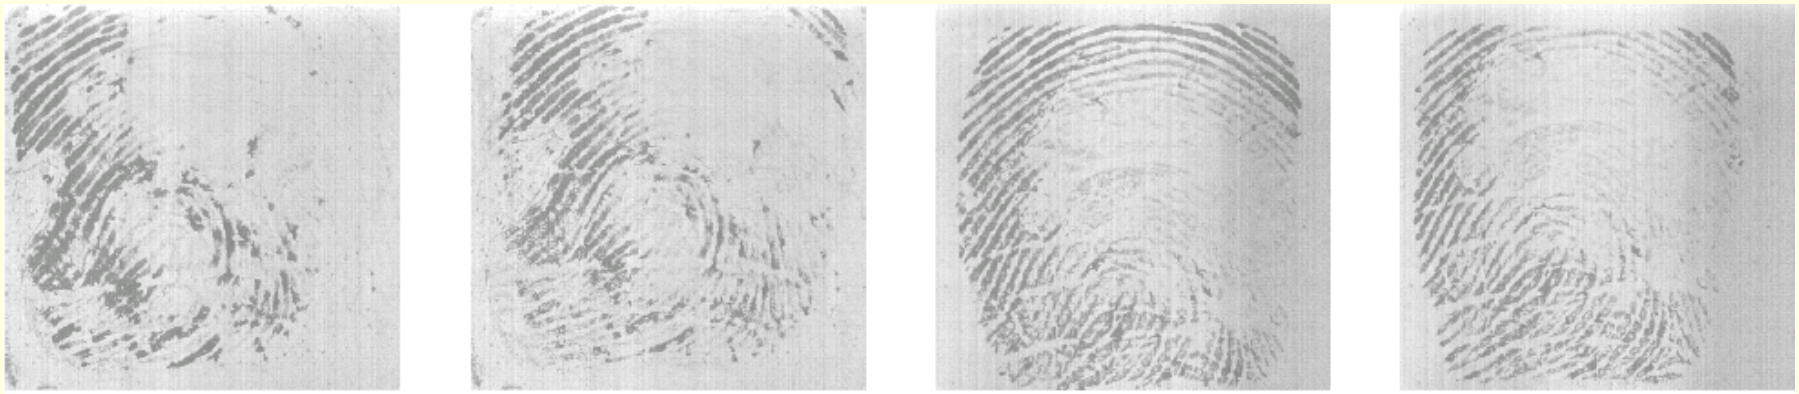

Four impressions of a user's print exhibiting incomplete ridge information

FTE: Failure-to-Enroll Problem

Jain, Prabhakar, Ross, "Fingerprint Matching: Data Acquisition and Performance Evaluation", MSU Technical Report TR99-14, 1999.

Intra-user variations

- **FNMR: False Non-Match Rate**

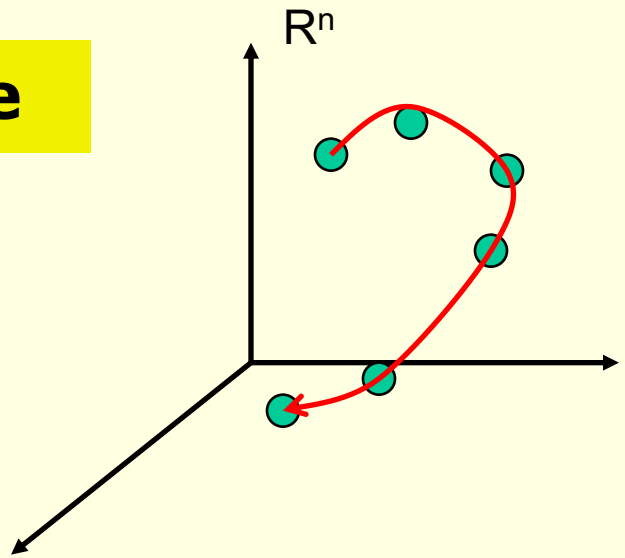

Inter-user similarity

TWIN BROTHERS
© Martin Schoeller

MOTHER DAUGHTER
© PleasantonWeekly.Com

- **FMR: False Match Rate**

Capacity of a template

- Existence of a biometric "zoo": Different categories of users impact error rates in a different manner

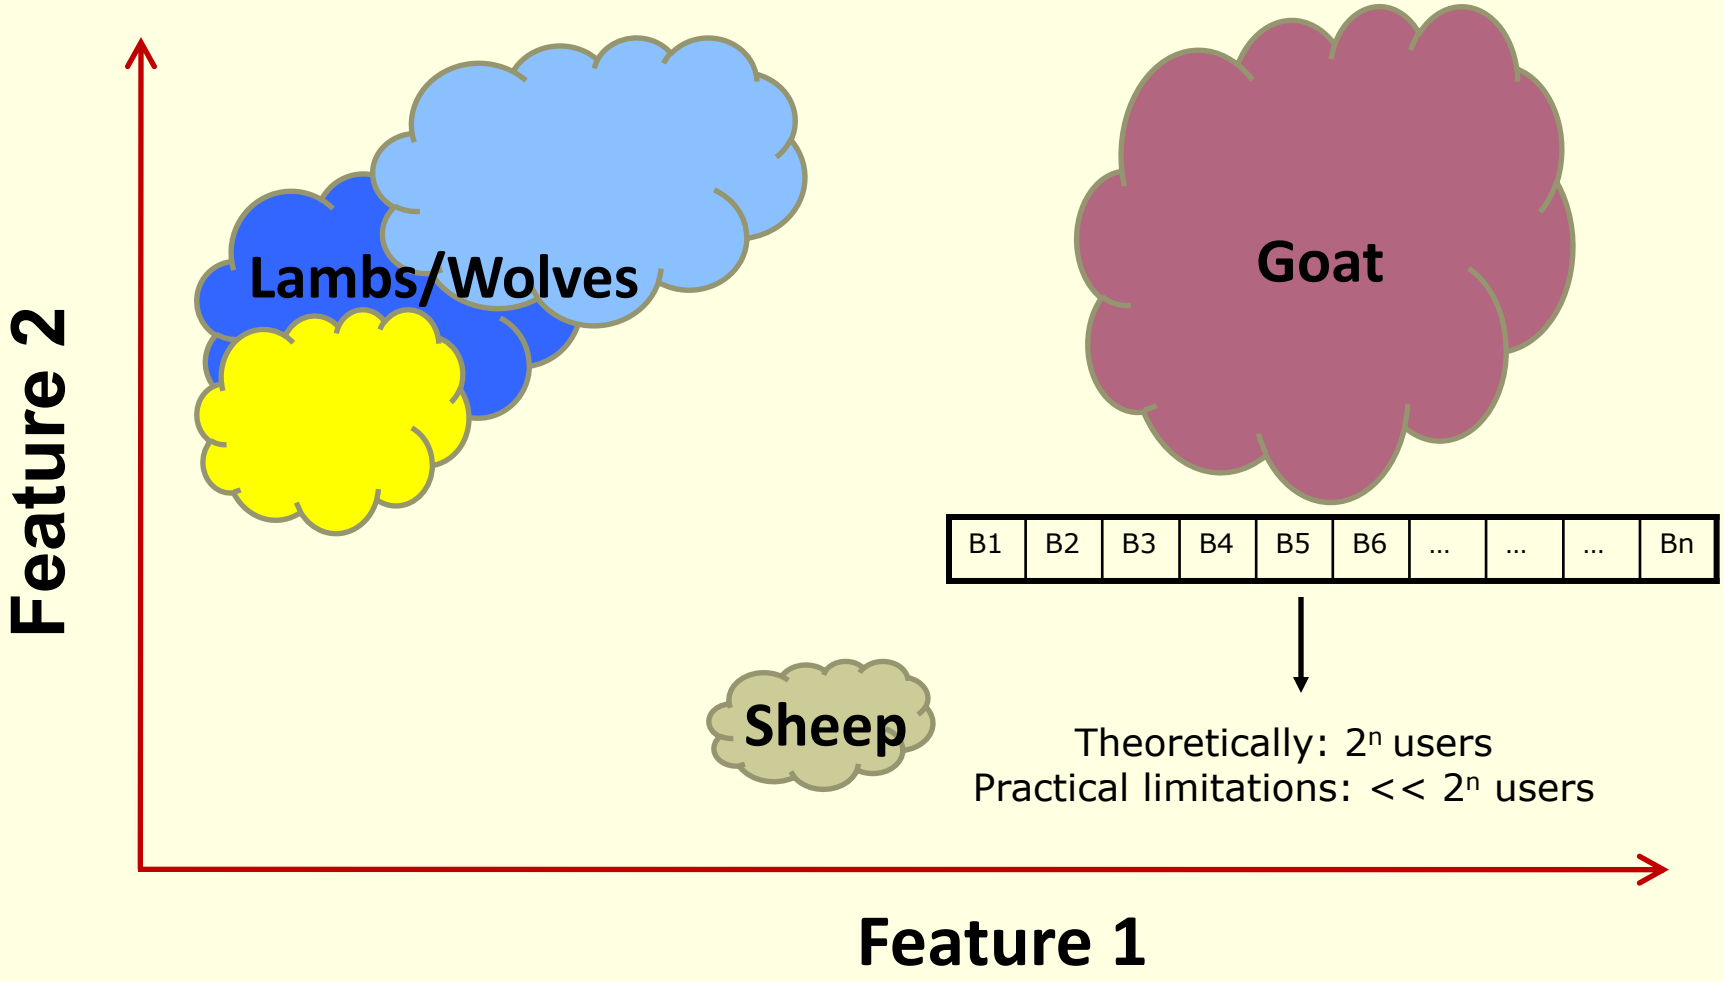

*Ross, Rattani, Tistarelli, "Exploiting the Doddington Zoo Effect in Biometric Fusion," BTAS 2009

Noisy Data

During enrolment

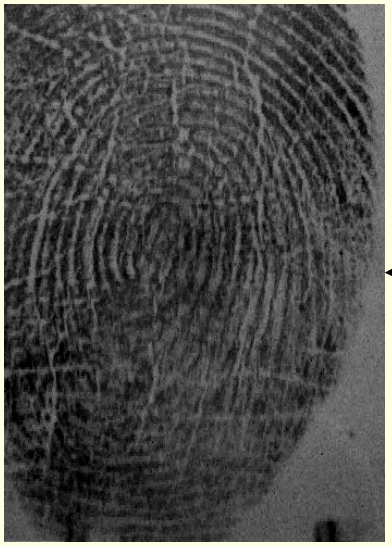

During recognition

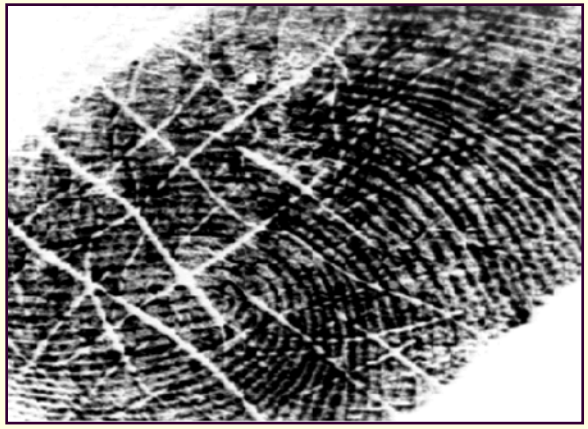

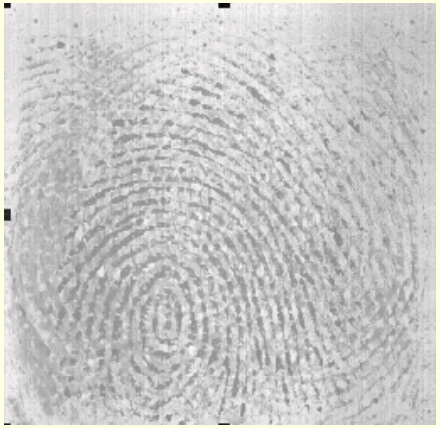

Noise due to smearing, residual deposits, cuts and folds, etc

Can impact both FMR and FNMR

Changes Due to Illumination

nachoguzman.net

Biometric Ageing

Biological age increases

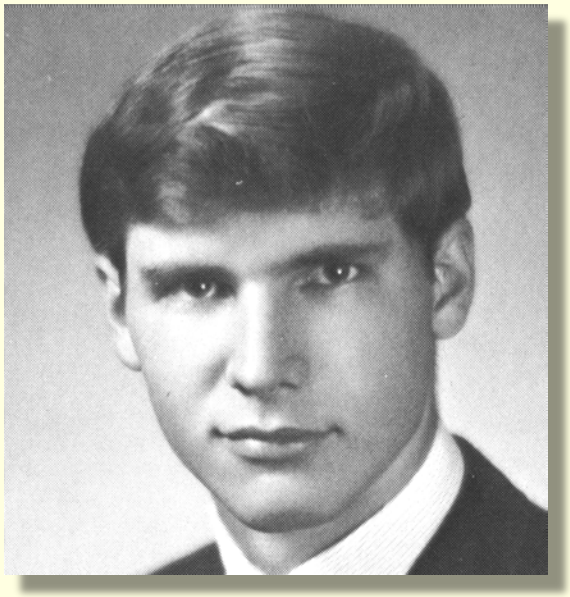

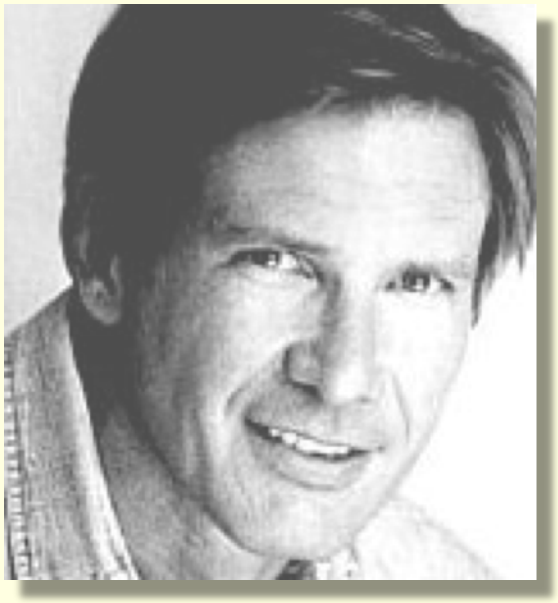

Heterogeneous Biometrics

Photo vs Sketch

Fundamental Differences in Image Formation Characteristics

Before vs After Makeup

RGB vs NIR vs THM

Young vs Old

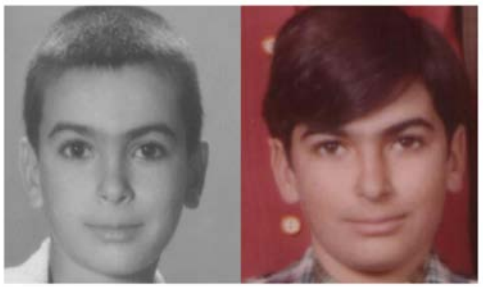

2D vs 3D

Spoofing: Presentation Attack

- Spoofing: Altering one's trait or creating a physical artifact in order to "spoof" another person's trait

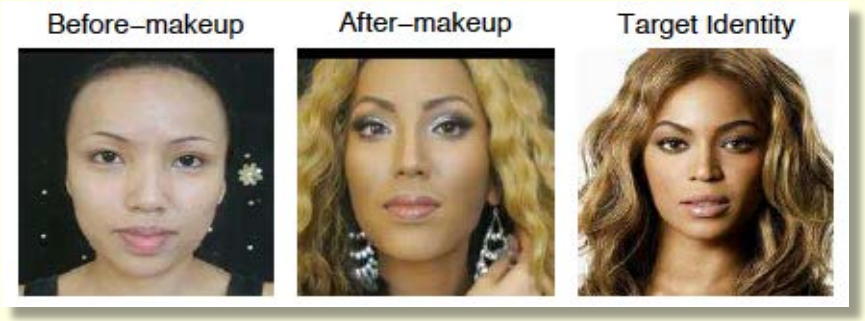

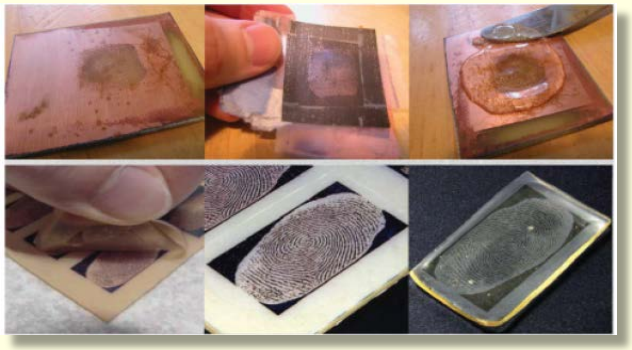

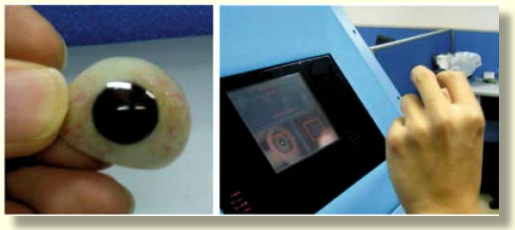

Obfuscation: Presentation Attack

- Obfuscation: Masking one's own identity by altering the trait

BEFORE

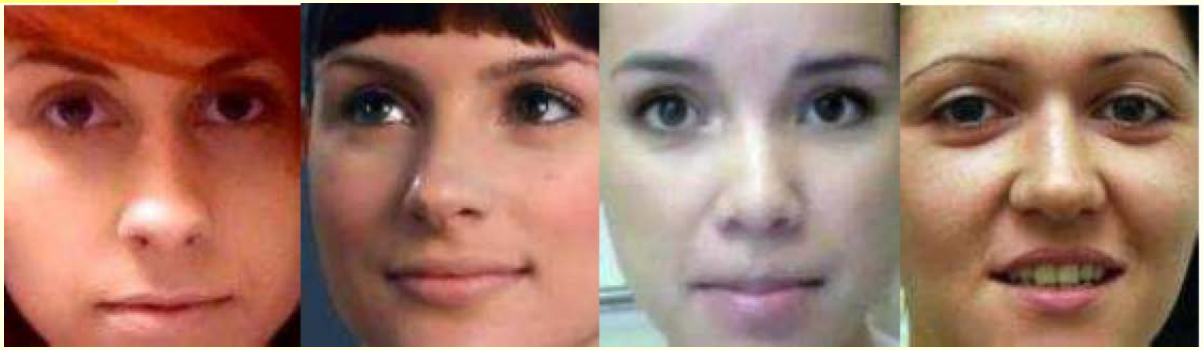

AFTER

Dantcheva et al, "Can Facial Cosmetics Affect the Matching Accuracy of Face Recognition Systems?", BTAS 2012

Fingerprint Alteration

- 1995: Alexander Guzman was arrested by Florida officials for possessing a false passport
- He was found to have mutilated fingerprints
- After a two-week search based on manually reconstructing the damaged fingerprints and searching the FBI database, the reconstructed fingerprints were linked to the fingerprints of Jose Izquiredo who was an absconding drug criminal

The "Z"-cut

- His fingerprint mutilation process consisted of three steps: making a 'Z' shaped cut on the fingertip; lifting and switching two triangles; and stitching them back.

Small Fingerprint Sensors

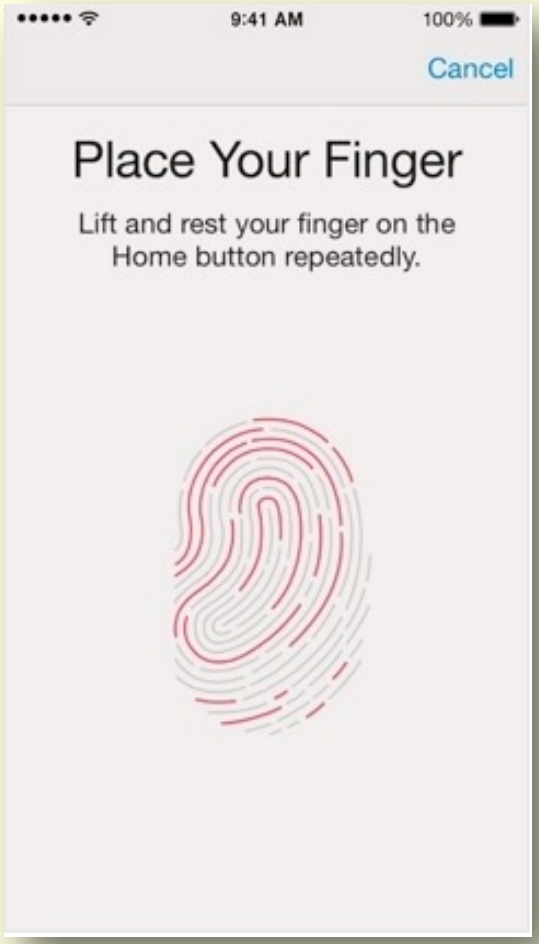

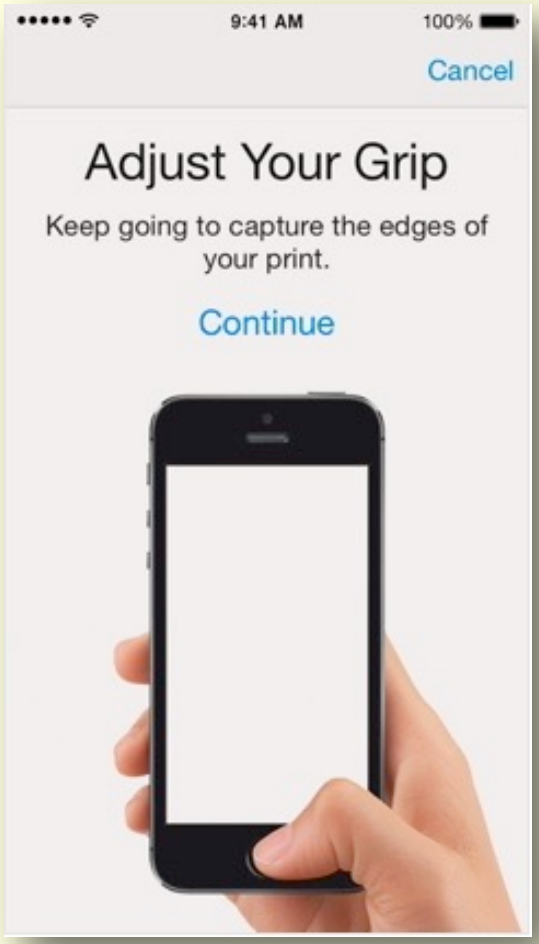

Partial Fingerprints

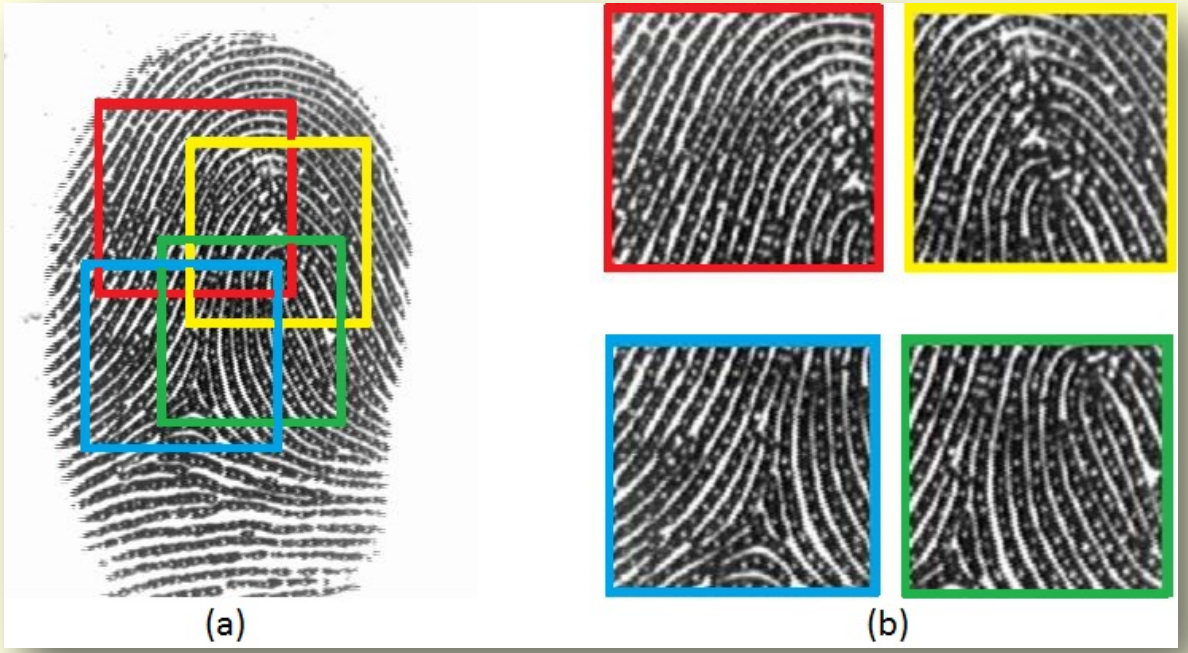

- Small sensors | Capture a limited portion of full finger
- Multiple partial fingerprints are captured | Enroll multiple fingers
- Access granted if the sensed partial fingerprint matches any one of the partial fingerprint of any enrolled finger

MasterPrints!

- Fingerprints that accidentally match with a large proportion of the fingerprint population
- Could be either full prints or partial prints

Roy, Memon, Ross, "MasterPrint: Exploring the Vulnerability of Partial Fingerprint-based Authentication Systems," TIFS 2017

Deep MasterPrints!

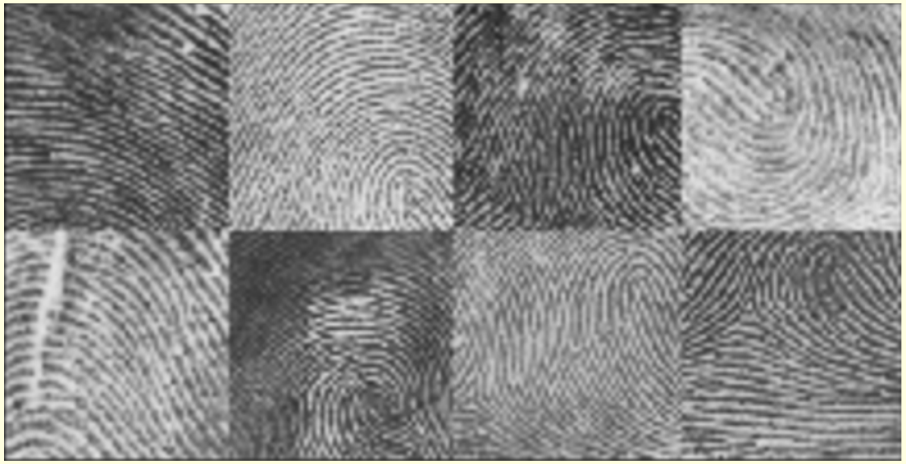

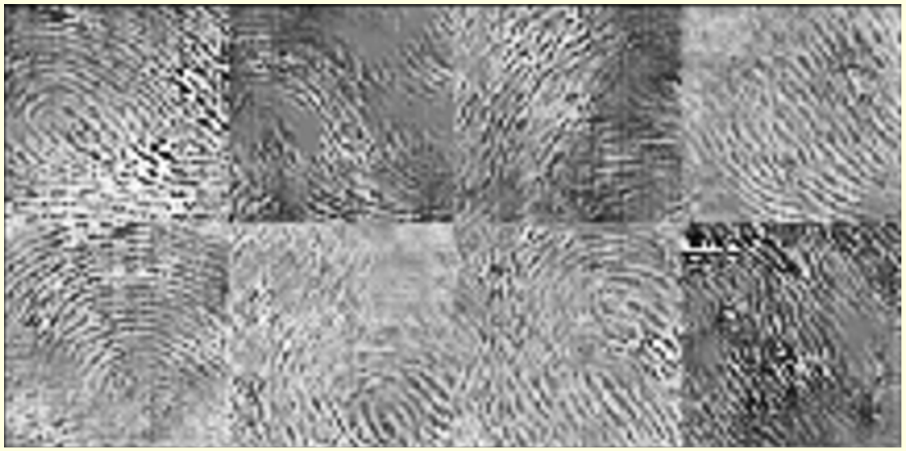

Bontrager et al, "DeepMasterPrints: Generating MasterPrints for Dictionary Attacks via Latent Variable Evolution," BTAS 2018

Observations: MasterPrints

- Dictionary of 5 MasterPrints + a maximum of 5 attempts
- It was possible to compromise:
  - 26.46% users (each having 12 impressions per finger) in the FingerPass DB7 capacitive fingerprint dataset
  - 65.20% users (each having ≈ 80 partial impressions per finger) in the FVC optical fingerprint at an FMR of 0.1%.
- **The attack accuracy varied greatly with the FMR value and the number of impressions per finger**

Roy, Memon, Ross, "MasterPrint: Exploring the Vulnerability of Partial Fingerprint-based Authentication Systems," TIFS 2017

Observations: DeepMasterPrints

	Rolled DeepMasterPrint			Capacitive DeepMasterPrint		
	0.01% FMR	0.1% FMR	1% FMR	0.01% FMR	0.1% FMR	1% FMR
	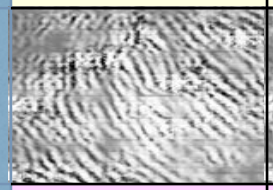	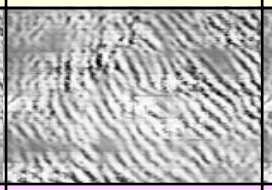	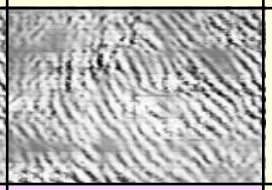	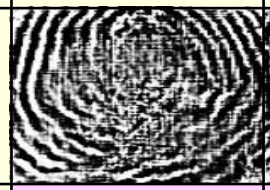	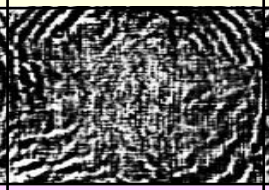	
Attack Rate	0.3%	8.6%	78.1%	1.1%	22.5%	76.7%

All evolved for the Fingerpass DB7 dataset (50% train/test split)

Bontrager et al, "DeepMasterPrints: Generating MasterPrints for Dictionary Attacks via Latent Variable Evolution," BTAS 2018

Attributes of a Biometric Trait

- **Uniqueness** (Is it distinctive across users?)
- **Permanence** (Does it change over time?)
- **Universality** (Does every user have it?)
- **Collectability** (Can it be measured quantitatively?)
- **Acceptability** (Is it acceptable to the users?)
- **Performance** (Does it meet error rate, throughput, etc.?)
- **Vulnerability** (Can it be easily spoofed or obfuscated?)
- **Integration** (Can it be embedded in the application?)

No biometric trait is "optimal", but many are "admissible"

Jain, Ross, Prabhakar. "An Introduction to Biometric Recognition," IEEE TCSVT, 2004

Evidence Accumulation and Information Fusion

It is time to stop arguing over which type of pattern classification technique is best because that depends on our context and goal. Instead we should work at a higher level of organization and discover how to build managerial systems to exploit the different virtues and evade the different limitations of each of these ways of comparing things (Minsky 1991)

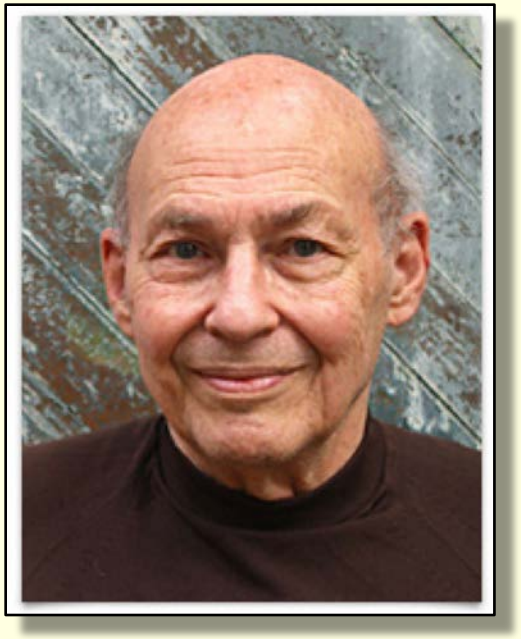

Marvin Lee Minsky
Born: August 9, 1927
Died: January 24, 2016

Biometric Fusion

- Combining multiple biometric evidence
- The identity of an individual is reinforced through multiple traits
- Especially significant in scenarios where partial biometric data is available

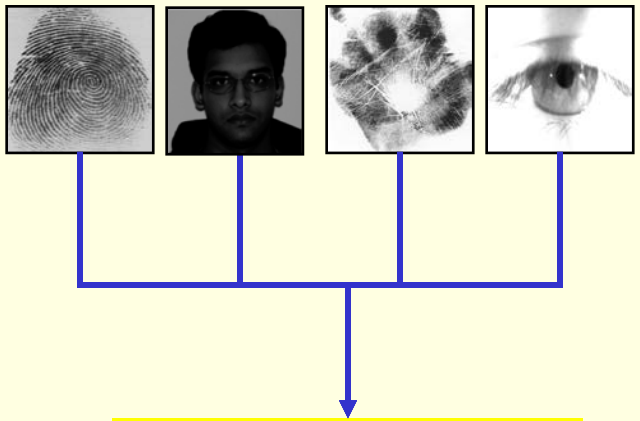

Information "Scavenging"

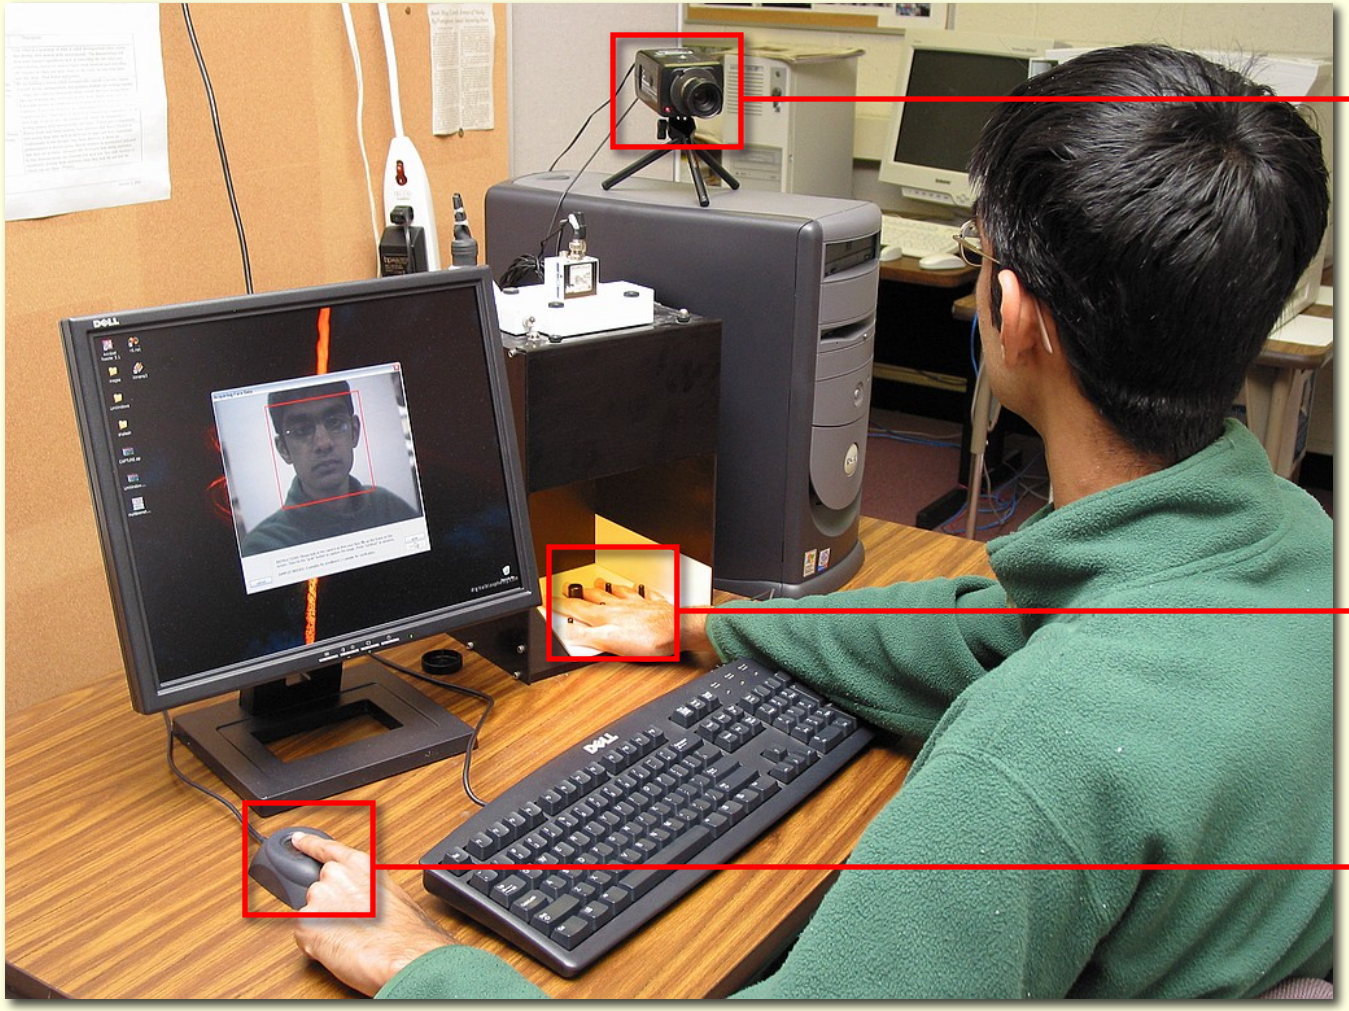

- Serial versus parallel mode of operation

Multibiometric Systems

- Multiple sources of biometric information are integrated to enhance matching performance
- Increases population coverage by reducing failure to enroll rate
- Anti-spoofing; difficult to spoof multiple traits simultaneously

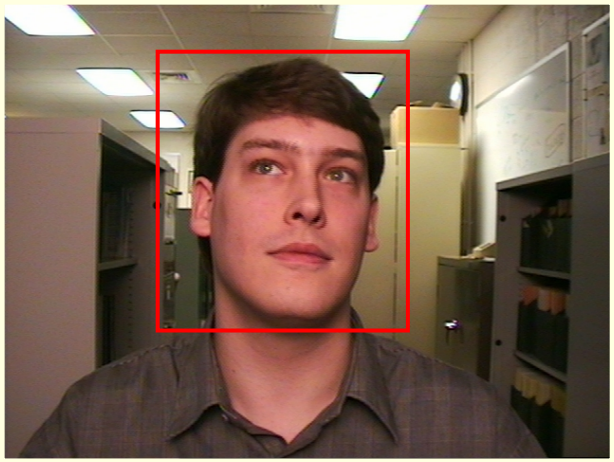

FBI and DHS

- The FBI fingerprint database has ten-print information of over 80 million individuals

- The US-VISIT (OBIM) database has information about the face and fingerprint of over 150 million individuals

Sources of Fusion

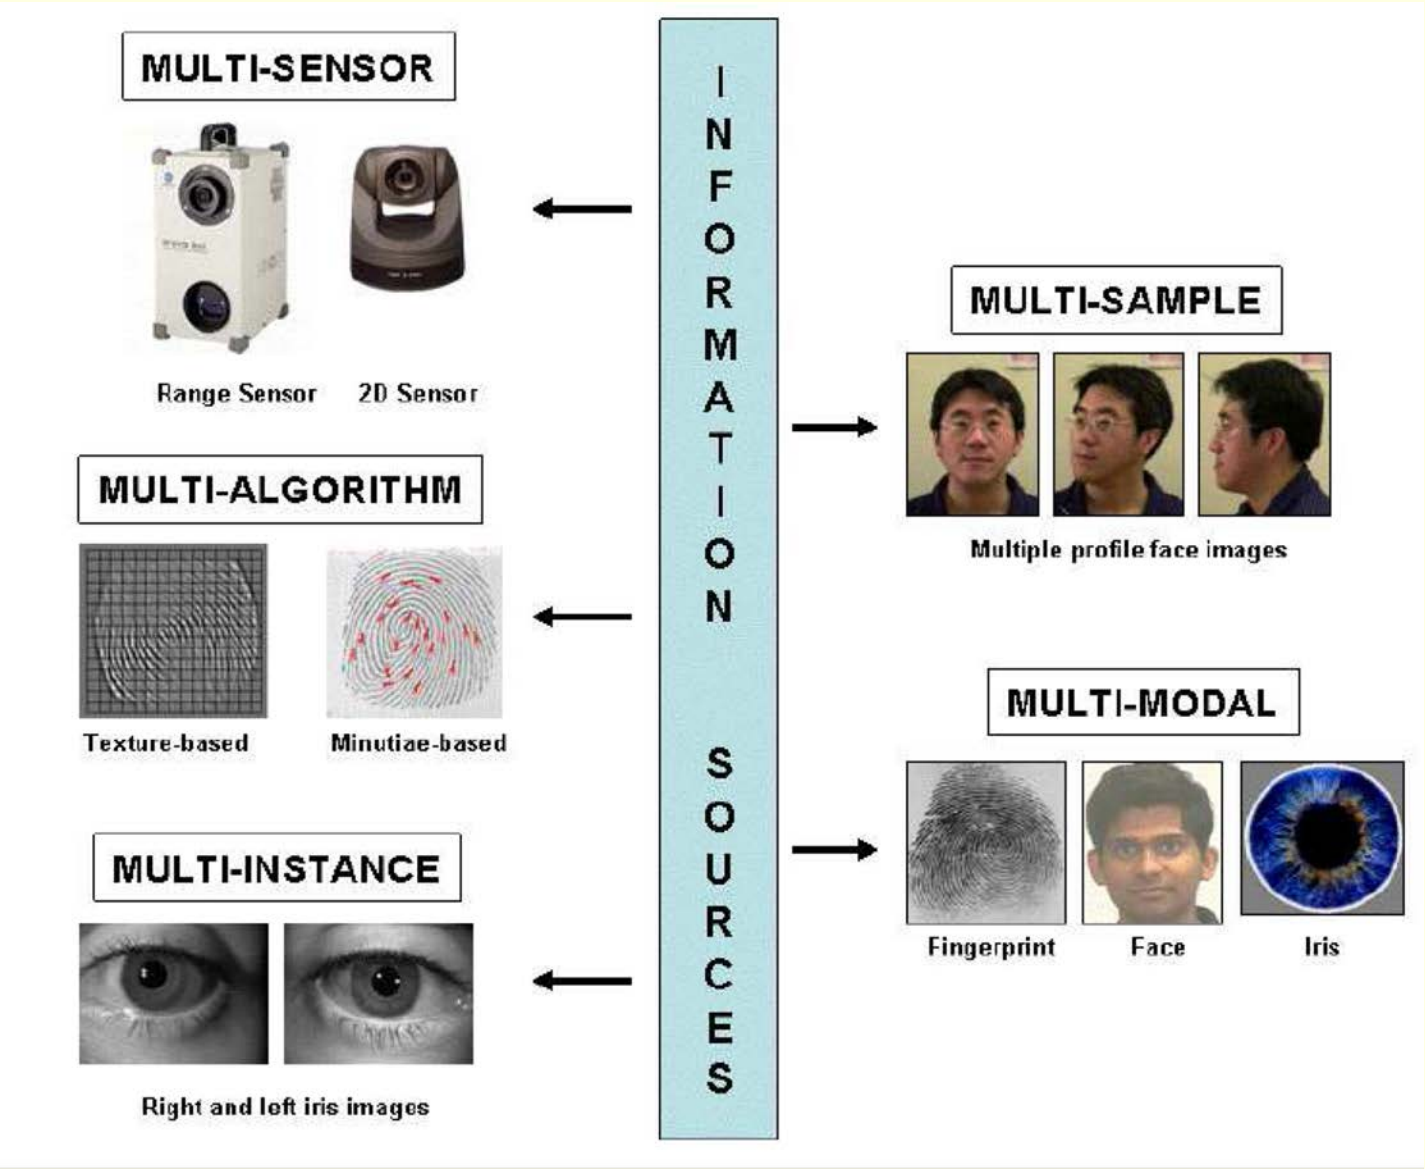

Dual Factor Authentication

Device ID + Subject ID

Who you are

+

What device you have?

Arjona et al., "Securing Minutia Cylinder Codes for Fingerprints through Physically Unclonable Functions: An Exploratory Study," *ICB 2018*

PUFs from SRAMs

- Reading the **start-up values** when the memory is switched on and no data are written
- Standard SRAM cells are composed of **two cross-coupled** inverters
- The **variability of the manufacturing process** can make inverters behave differently and the cell has a trend towards a particular start-up value

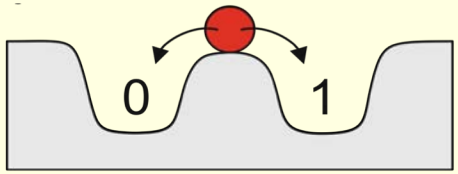

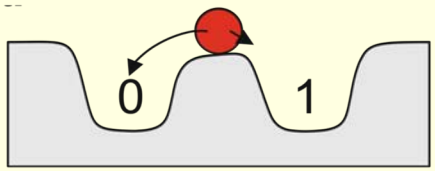

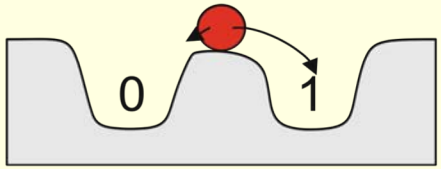

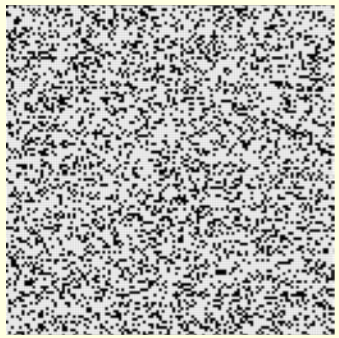

Arjona et al., "Securing Minutia Cylinder Codes for Fingerprints through Physically Unclonable Functions: An Exploratory Study," *ICB 2018*

Levels of Fusion

Modality 1

Raw Data

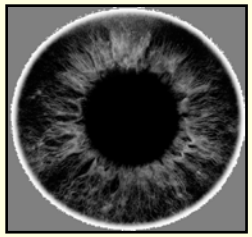

Feature vector

X1	X2	X3	X4	X5	X6	...	Xn

Match Score

S1 = 50

Rank

Rank 1: Alice
Rank 2: Bob
Rank 3: Dan

Binary Decision

Genuine

Modality 2

Y1	Y2	Y3	Y4	Y5	Y6	...	Ym

S2 = 75

Rank 1: Alice
Rank 2: Ed
Rank 3: Bob

Impostor

Data Level Fusion

- Mosaicing constructs a composite fingerprint image (or template) using multiple impressions of the same finger resulting in more information (e.g., minutiae points)

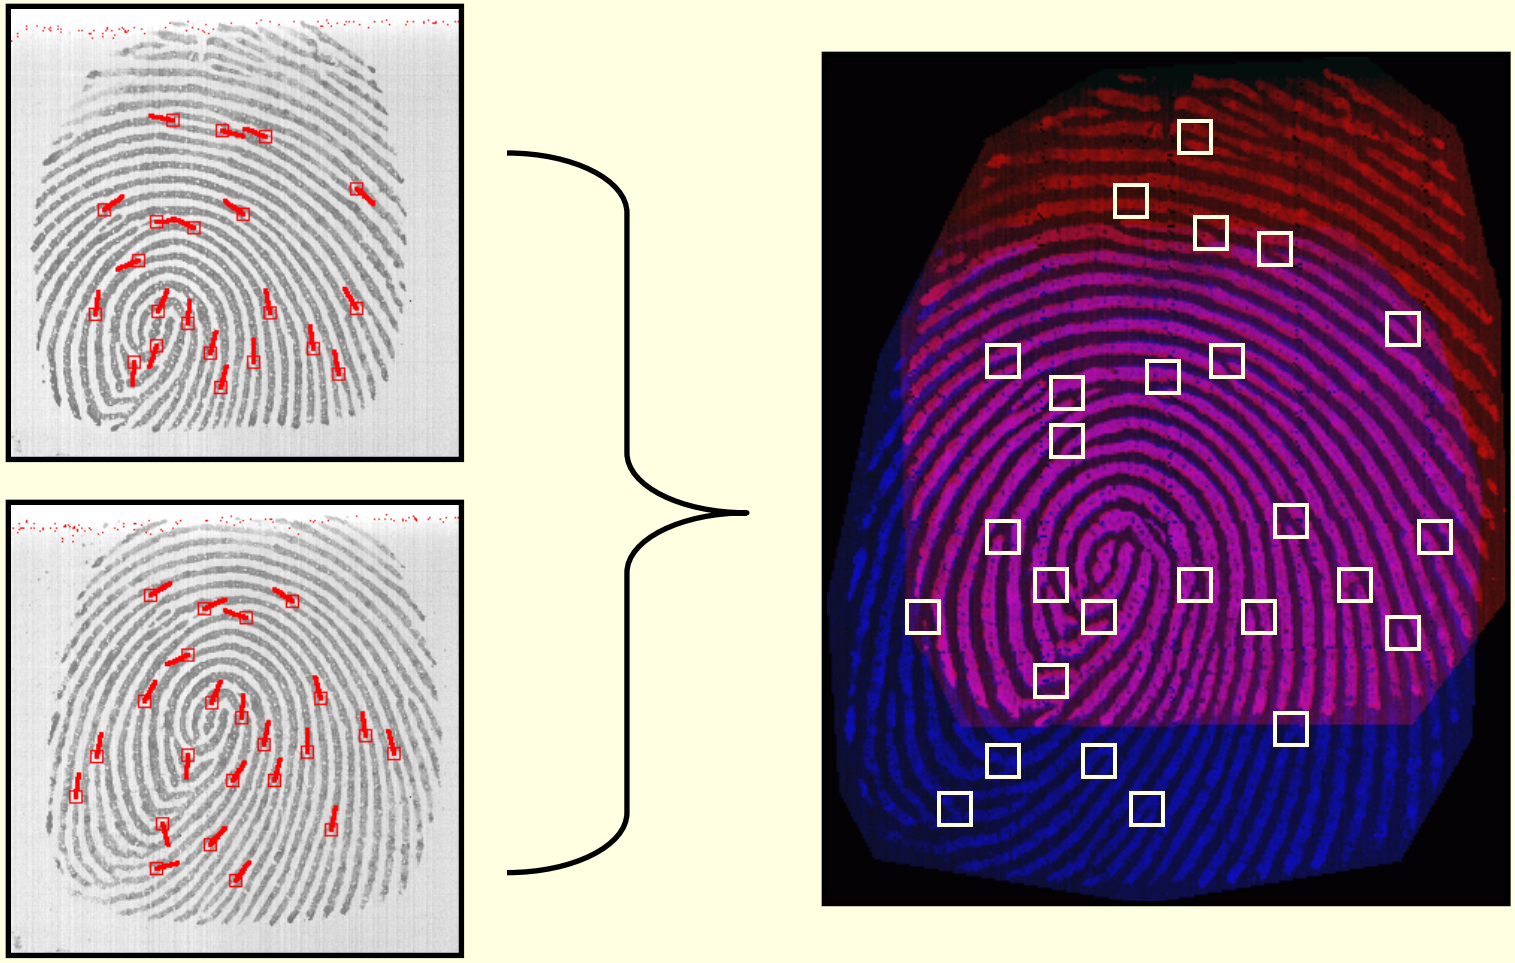

Ross et al , "Image Versus Feature Mosaicing: A Case Study in Fingerprints", SPIE, April 2006.

Data Level Fusion

- The raw data pertaining to multiple sensors are combined
  - e.g., the 2D face texture may be mapped to a 3D range image; matching performed in 3D space

2.5D range data

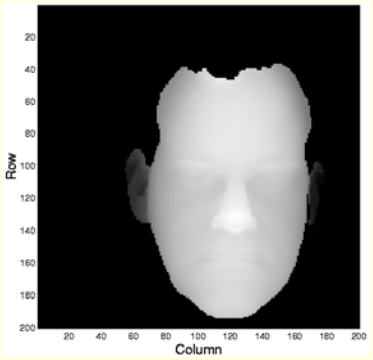

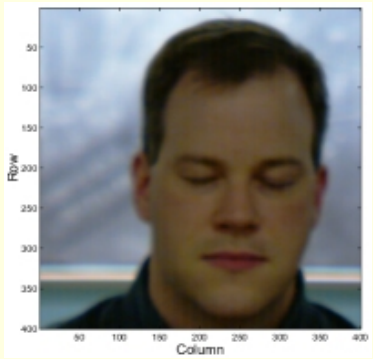

2D color texture

Texture-mapped appearance

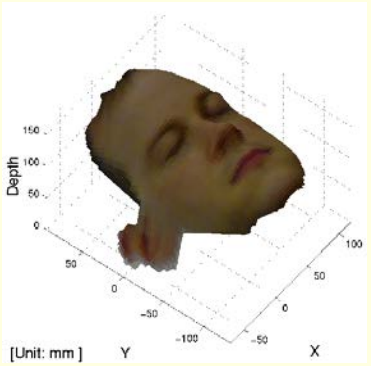

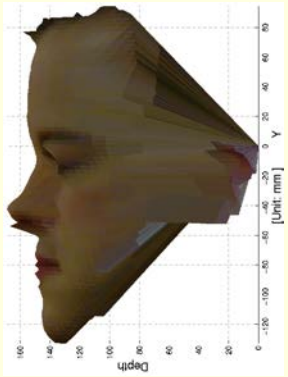

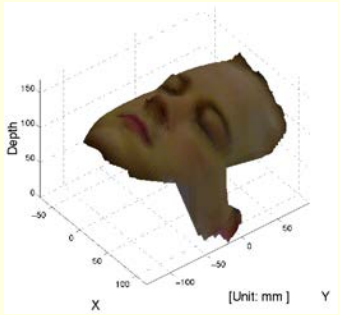

R.-L. Hsu, ``Face Detection and Modeling for Recognition'', Ph.D. Thesis, 2002

Data Level Fusion

- Goal: To de-identify fingerprint and iris images by generating a new, possibly unique, and de-identified biometric
- IrisPrint can be used directly in the feature extraction and matching stages of an existing matcher without revealing the original images

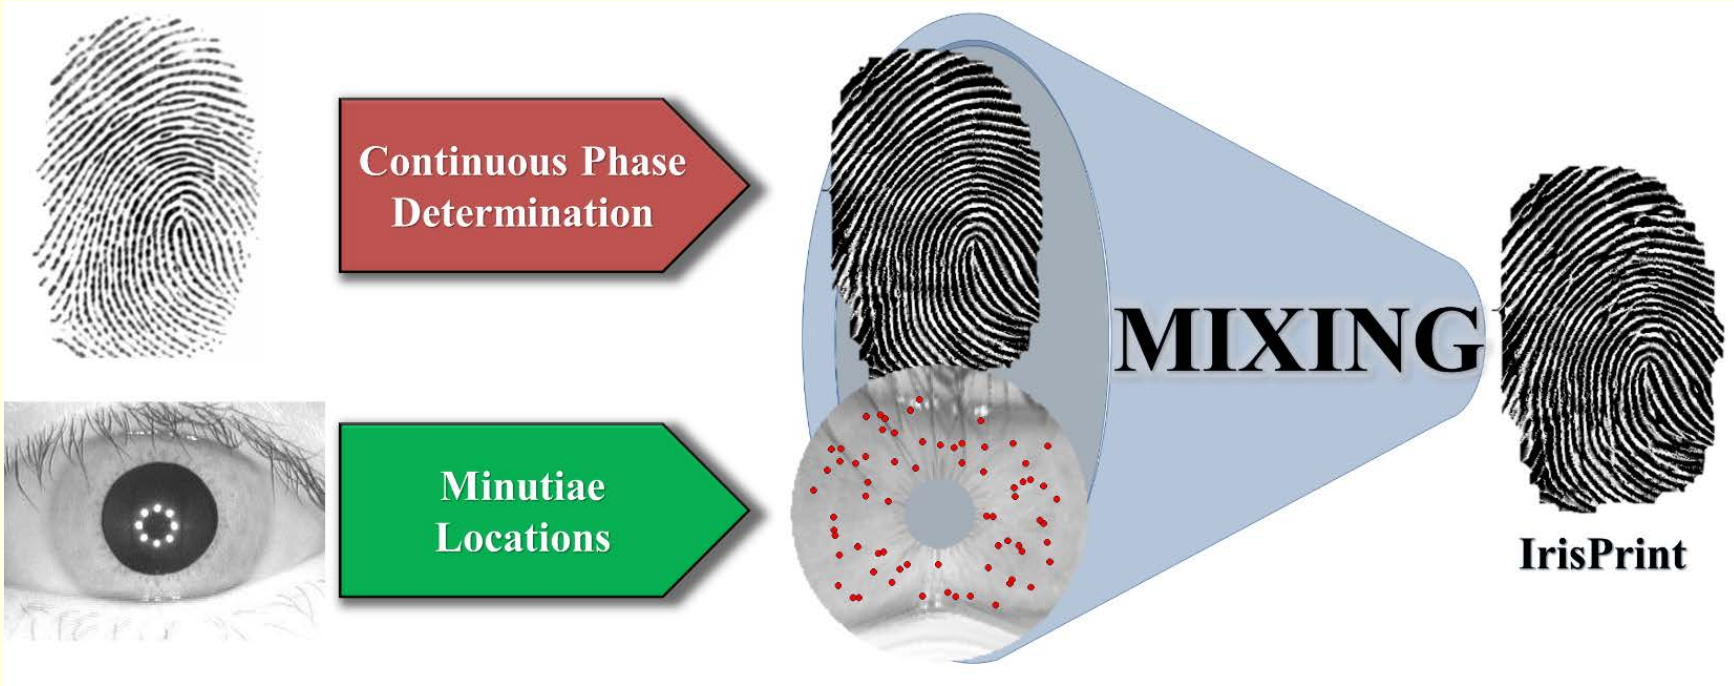

Othman and Ross, "Fingerprint + Iris = IrisPrint," SPIE 2014

Feature Level Fusion

- The feature space of two modalities are combined
  - e.g., the feature vector of face combined with that of hand geometry

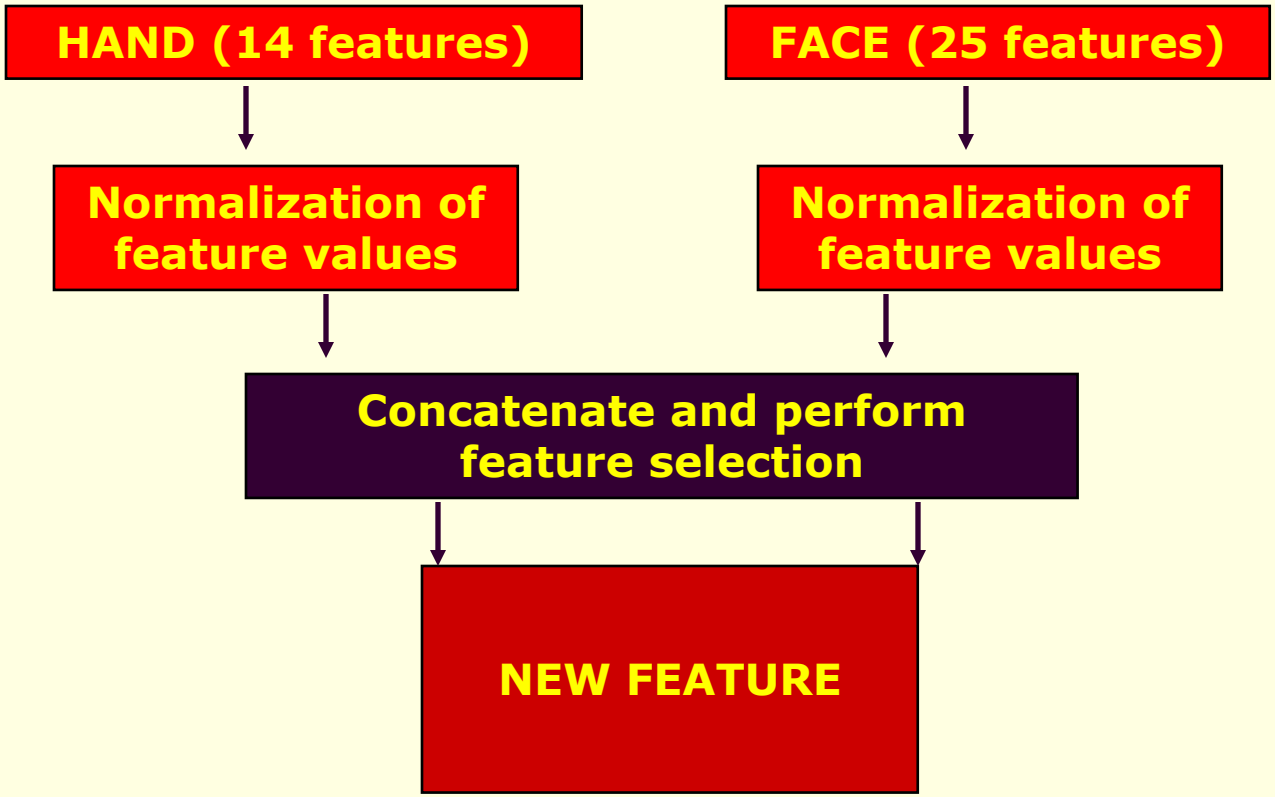

Ross, Govindarajan, "Feature Level Fusion of Hand and Face Biometrics", SPIE 2005.

Feature Level Fusion

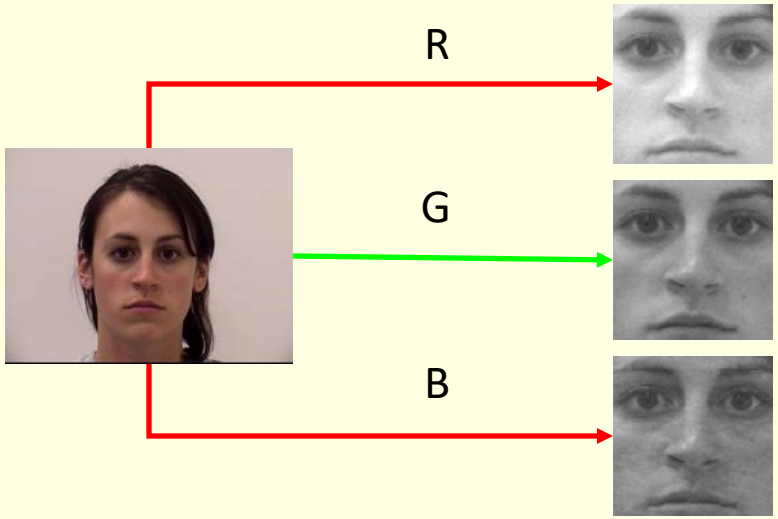

- Feature sets:
  - LDA-R : 18 features
  - LDA-G : 32 features
  - LDA-B : 40 features
- Feature-fused vector: 43 features

Ross, Govindarajan, "Feature Level Fusion of Hand and Face Biometrics", SPIE 2005.

Density-based Fusion

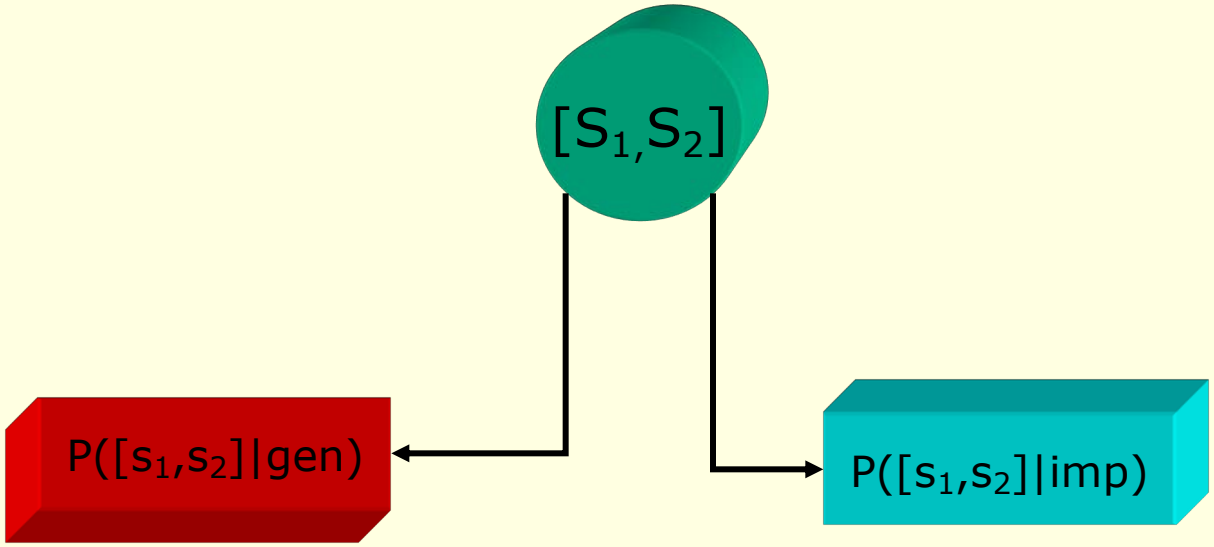

$$\frac{P([s1, s2] \mid gen)}{P([s1, s2] \mid imp)} > \text{Threshold, then Genuine}$$

else Impostor

Density-based Fusion

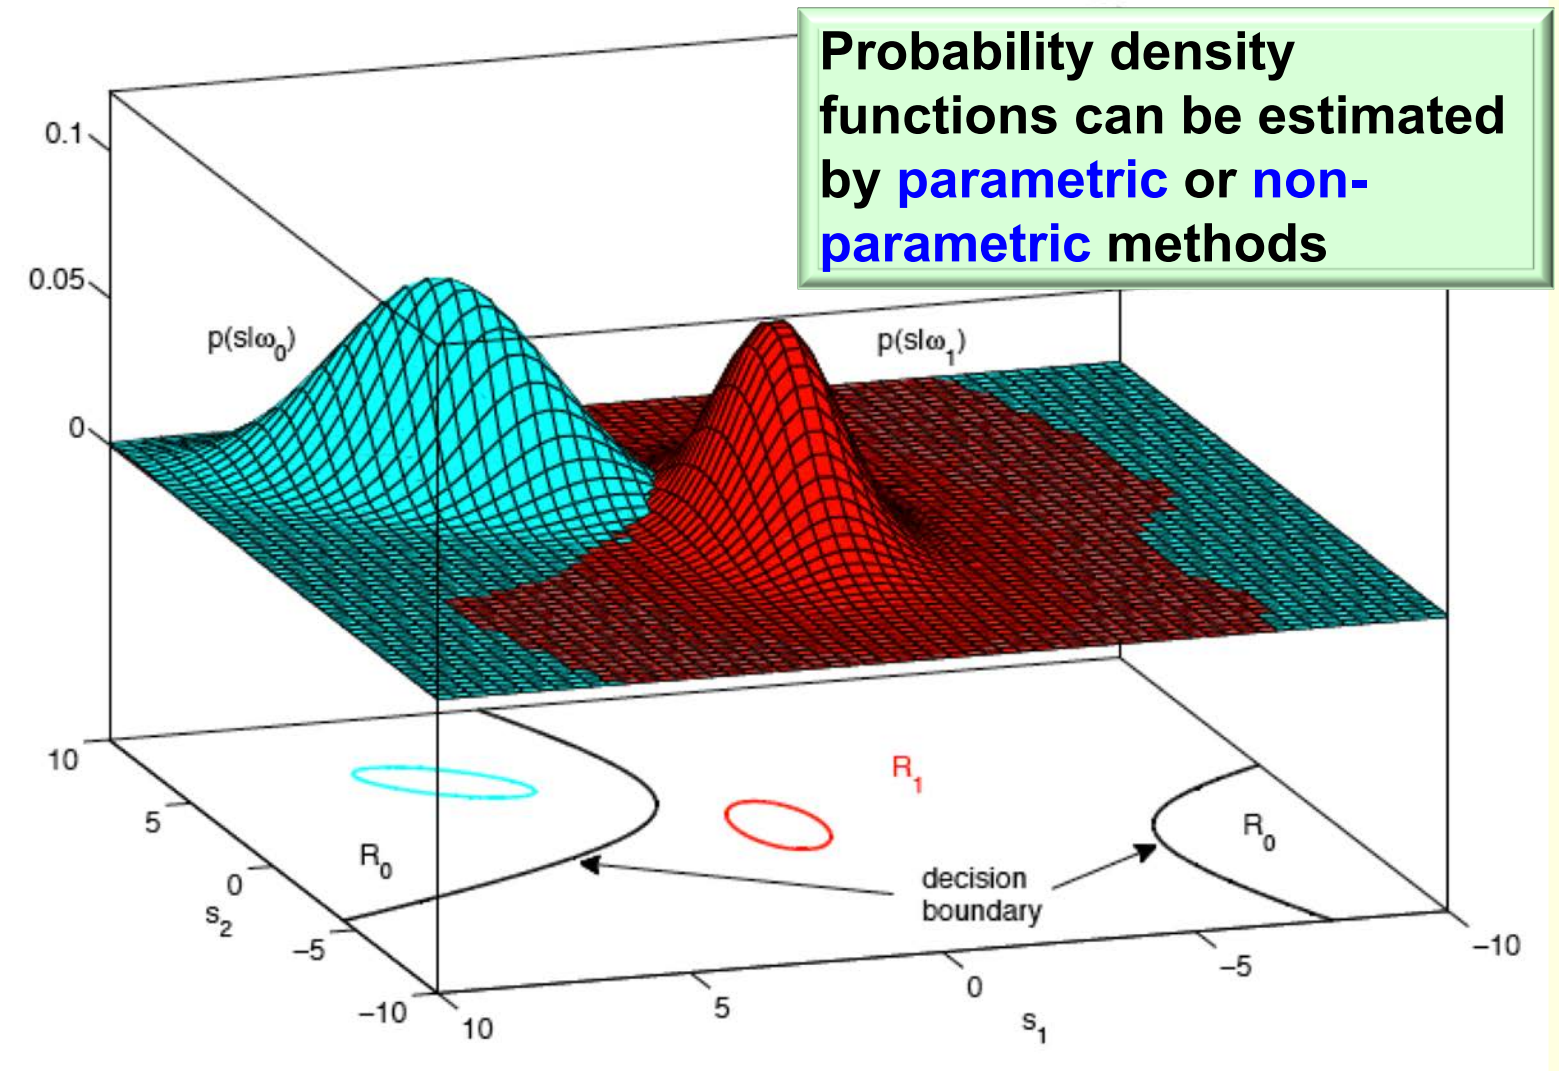

Classifier-based Fusion

- Match scores emitted by multiple sources are input to a trained classifier

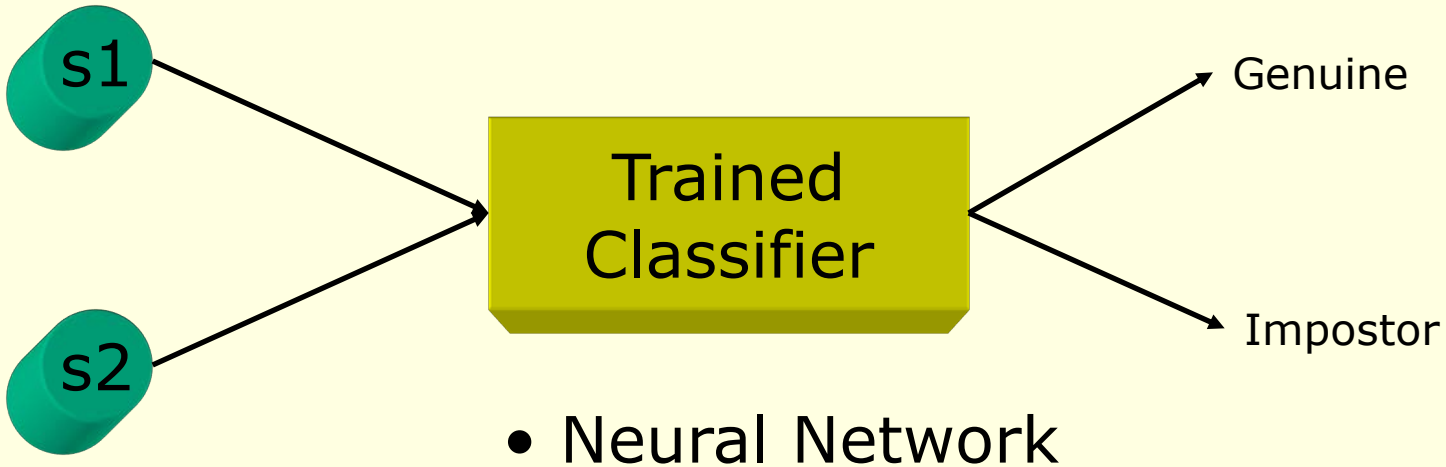

- Neural Network
- SVM
- Decision Trees
- Nearest Neighbor
- Random Forest

Transformation-based Fusion

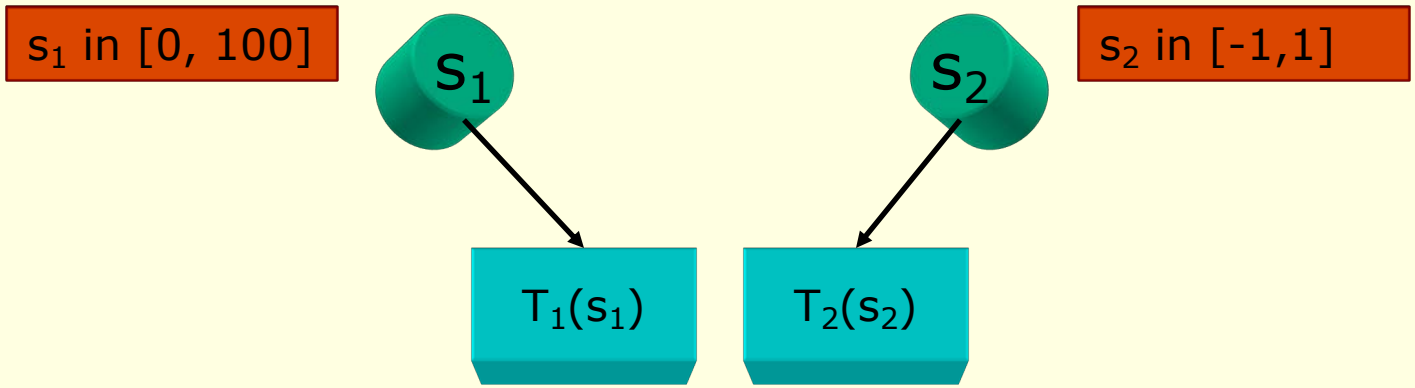

- The transformed scores can be combined using several different rules
  - min$[T_1(s_1), T_2(s_2)]$
  - max$[T_1(s_1), T_2(s_2)]$
  - sum$[T_1(s_1), T_2(s_2)]$
  - prod$[T_1(s_1), T_2(s_2)]$

T_i: Normalization Function
1. minmax
2. MAD
3. tanh

Jain, Ross, Nandakumar, "Score Normalization in Multimodal Biometric Systems," PR 2005

Simple Sum Rule

- Sum rule (weighted average of individual scores) has been shown to improve matching accuracy:

$$S = w_1 s_1 + w_2 s_2 + w_3 s_3$$

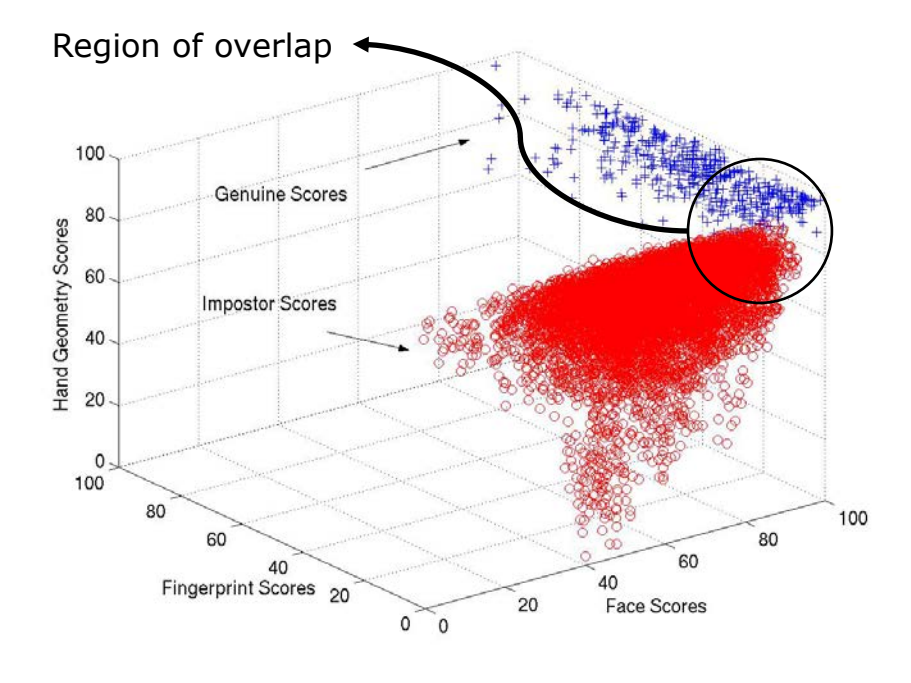

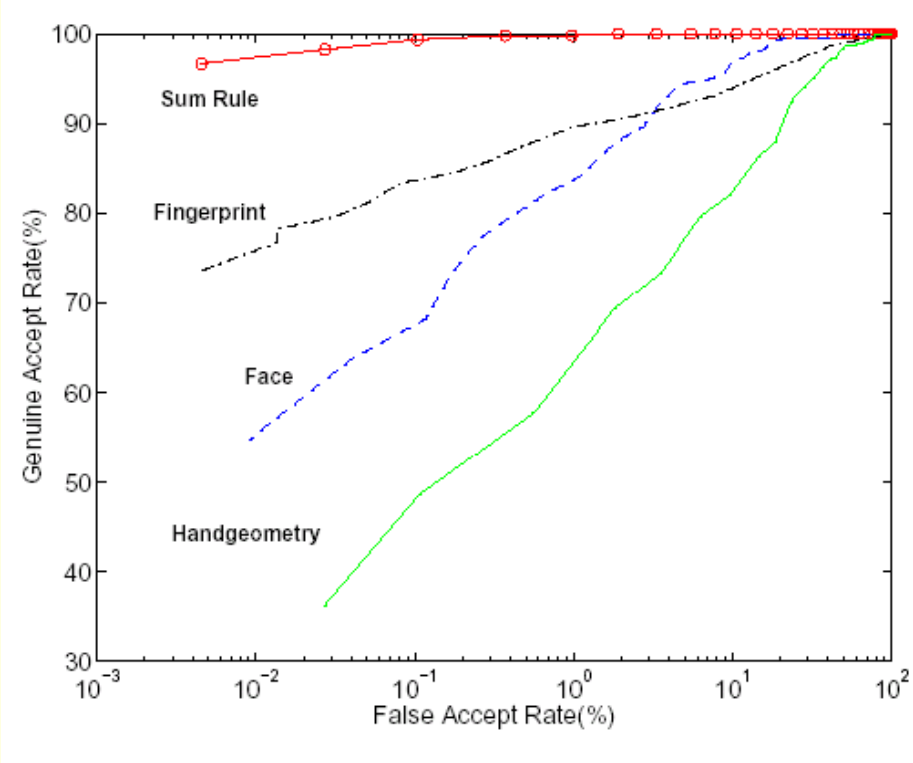

Ross and Jain, "Information Fusion in Biometrics", PRL 2003.

Score Normalization

- Scores output by individual matchers:
  - Non-homogeneous: distance or similarity
  - Ranges may be different; e.g., [0,100] or [0,1000]
  - Distributions may be different

- To facilitate fusion:
  - Modify the location and scale parameters of score distributions of individual matchers
  - Apply transformation to scores present in the genuine-impostor overlap region

- Factors to consider:
  - Robustness: Should not be affected by the outliers
  - Efficiency: Estimated parameters of the score distribution should be close to the true values

Match Score Distributions

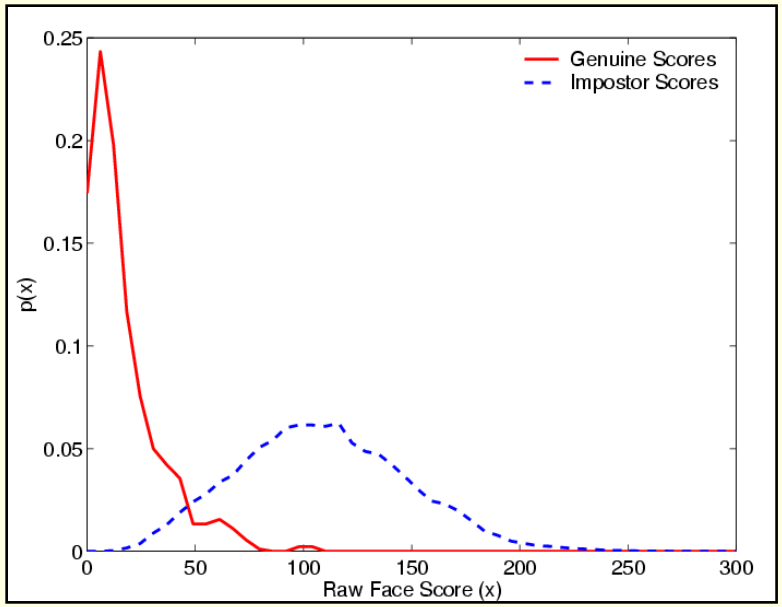

Face

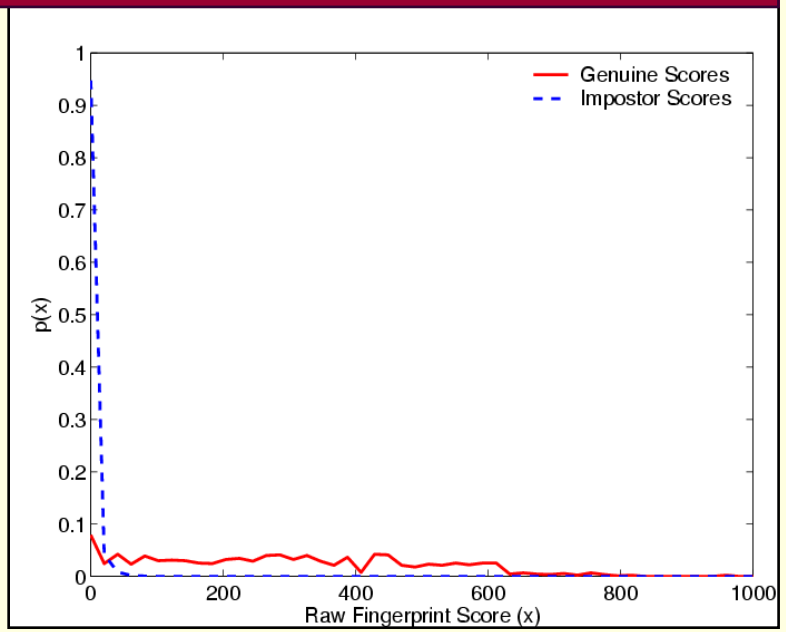

Fingerprint

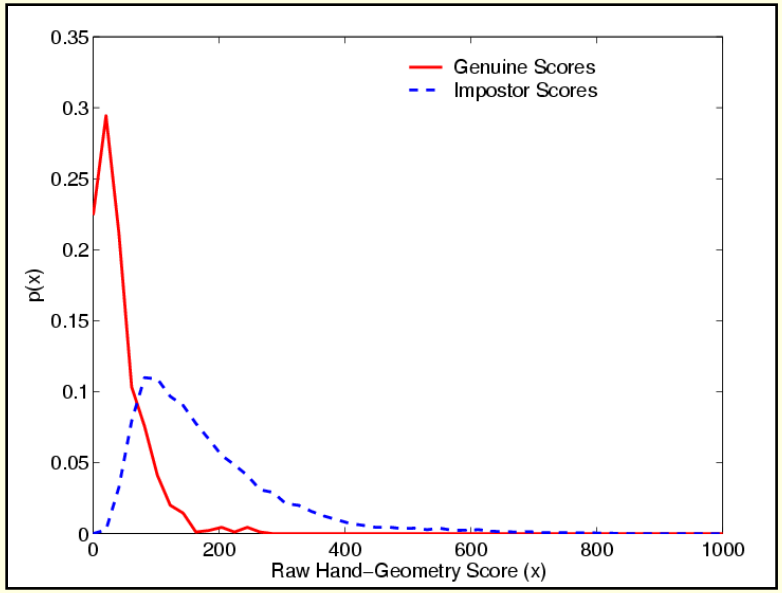

Hand-geometry

Normalization Techniques

- Min-max normalization: Given matching scores $\{s_k\}$, k=1,2,..,n the normalized scores are given by:

$$s' = \frac{s - \min\{s_k\}}{\max\{s_k\} - \min\{s_k\}}$$

- Decimal scaling: Used when scores of different matchers differ by a logarithmic factor; e.g., one matcher has scores in the range [0,1] and the other matcher has scores in the range [0, 1000]

$$s' = \frac{s}{10^n},$$

$$n = \log_{10} \max\{s_k\}$$

Normalization Techniques

- Z-score:

$$s' = \frac{s - \mu}{\sigma}$$

- Median and Median Absolute Deviation (MAD):

$$s' = \frac{(s - median)}{MAD}$$

$$MAD = median(|\{s_k\} - median|)$$

- Double Sigmoid function:

$$s' = \frac{1}{1 + \exp\left(-2\left(\frac{s - t}{r}\right)\right)}$$

r = r_1, if s < t

r = r_2, otherwise

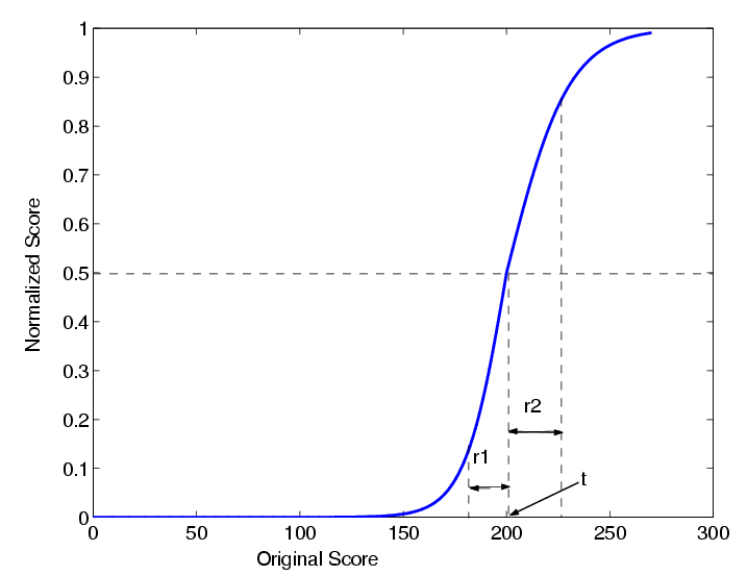

Normalization Techniques

- Tanh estimators:

$$s' = 0.5\left[\tanh\left(0.01\frac{(s-\mu_{GH})}{\sigma_{GH}}\right)+1\right],$$

where μ_{GH} and σ_{GH} are the mean and standard deviation estimates of the genuine score distribution as given by Hampel estimators*

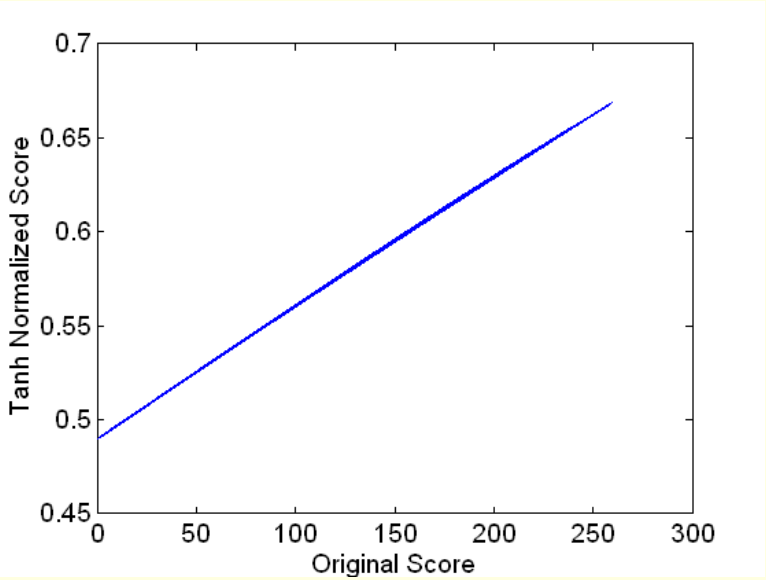

- Min-max, Z-score, and Tanh normalization schemes are efficient
- Median, Double Sigmoid, and Tanh methods are robust

**Hampel et al., Robust Statistics: The Approach Based on Influence Functions, 1986*

Overlap Region

- QLQ transformation:

$$n_{QLQ} = \begin{cases} \frac{1}{(c-\frac{w}{2})} n_{MM}^2 & n_{MM} \leq (c - \frac{w}{2}) \\ n_{MM} & (c - \frac{w}{2}) \leq n_{MM} \leq (c + \frac{w}{2}) \\ (c + \frac{w}{2}) + \sqrt{(1 - c - \frac{w}{2})(n_{MM} - c - \frac{w}{2})} & \text{otherwise} \end{cases}$$

- n_{MM} is the min-max normalized score
- c is the center of the overlap regions
- w is the width of the overlap region

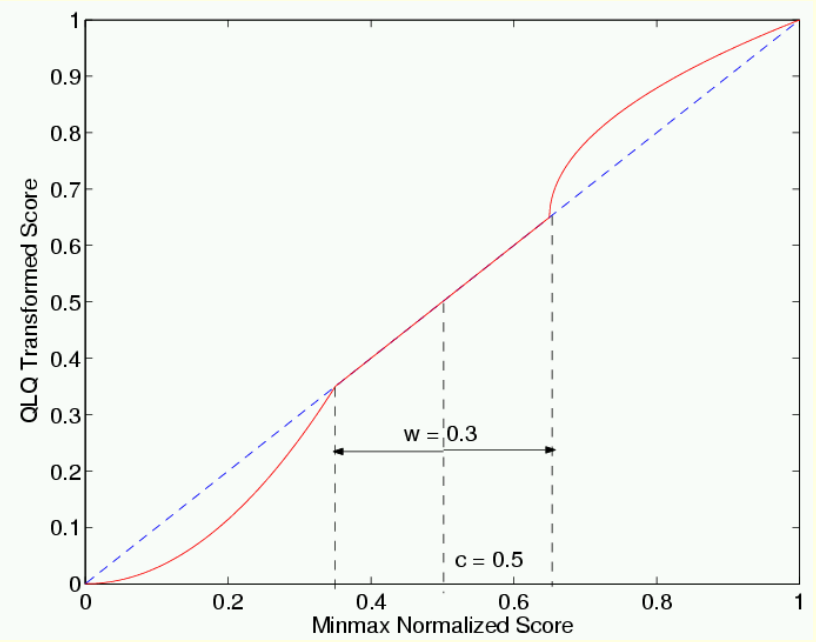

Score Level Fusion

Feature Extraction and Matching

d_f

Normalization

S_f'

Eigenfaces

Fusion Rule

Decision

Feature Extraction and Matching

s_p

Normalization

s_p'

Minutiae

Feature Extraction and Matching

d_h

Normalization

s_h'

14-D hand feature vector

Effect of Normalization

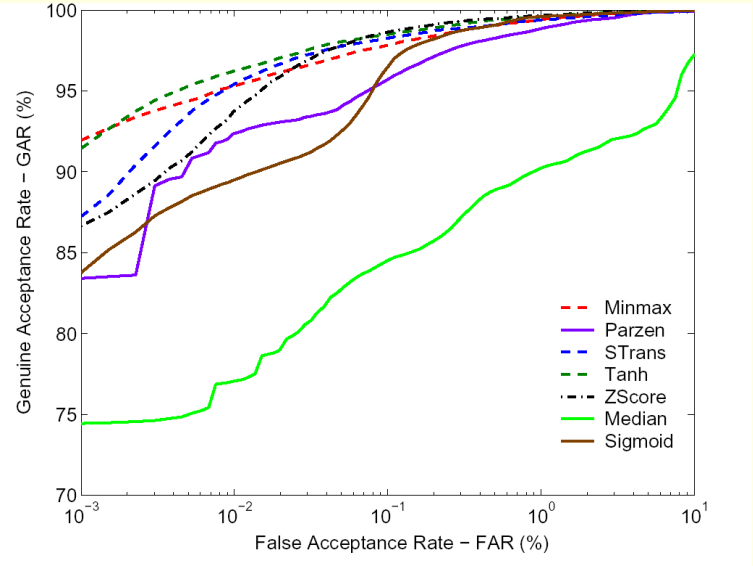

(a) Results of various schemes

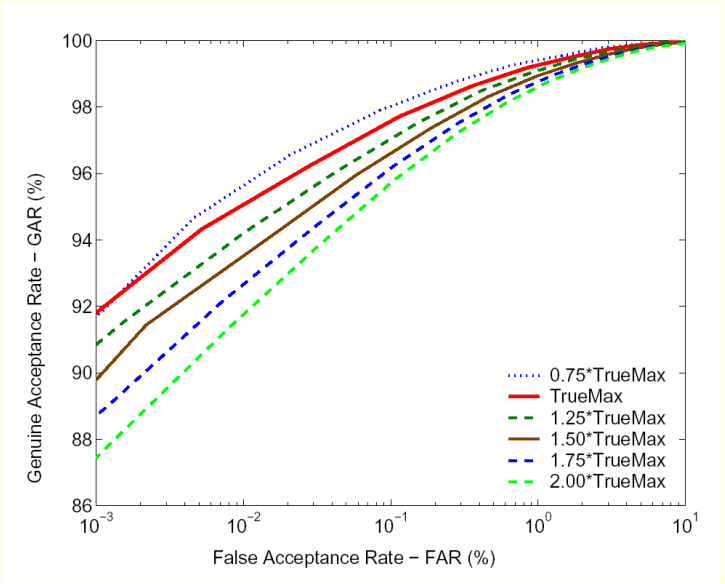

(b) Sensitivity to outliers - minmax

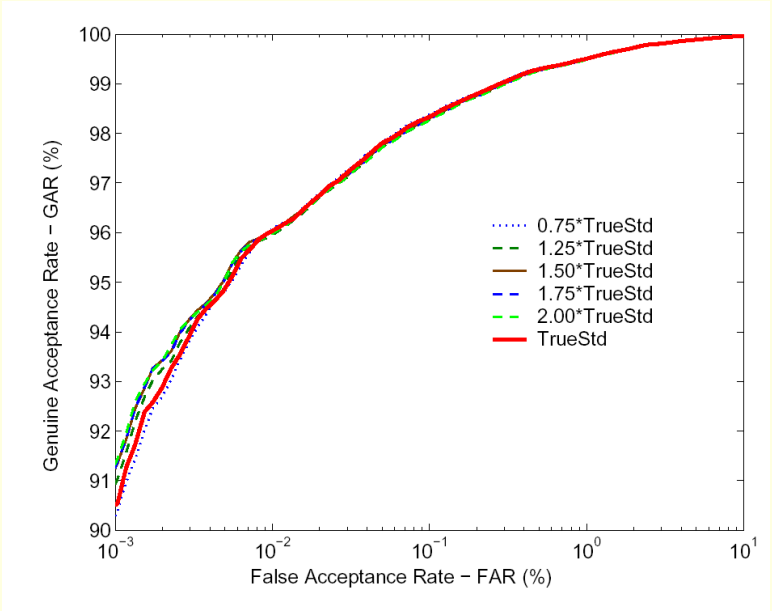

(c) Sensitivity to outliers - tanh

Jain et al, "Score Normalization in Multimodal Biometric Systems", Pattern Recognition 2005.

Is Fusion Always Beneficial?

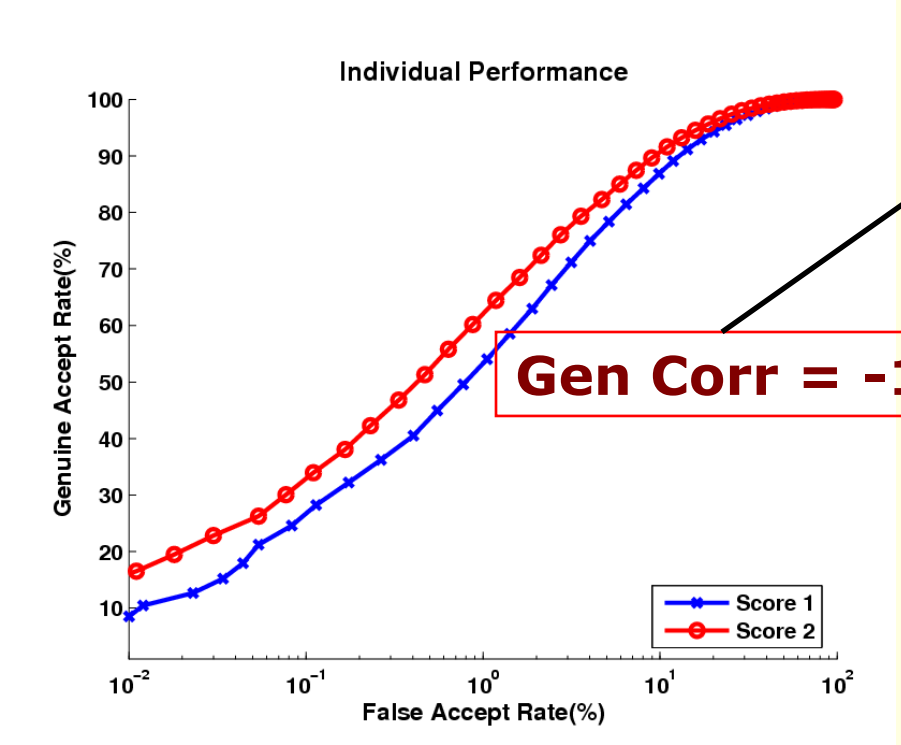

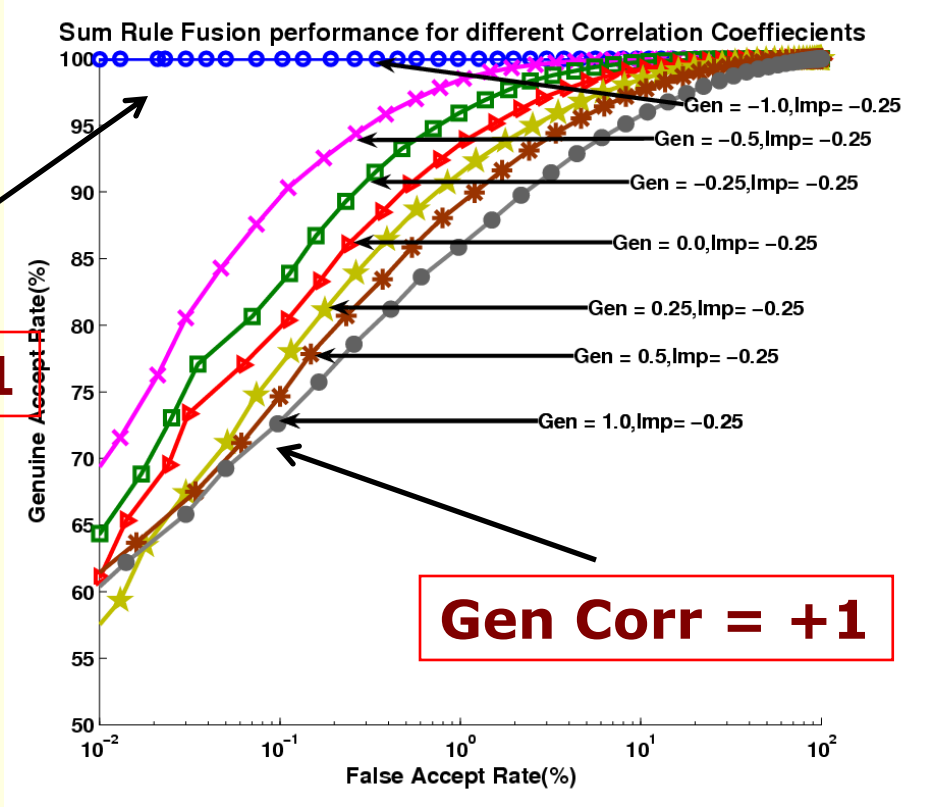

SINGLE MODALITY

SUM RULE FUSION

- Negatively correlated or uncorrelated classifiers preferable

Identification Systems

- Given an input image:
  - Compare input against the enrolled identities using the matcher
  - Generate a ranking of the enrolled identities based on their match scores
- Ranks versus Scores
  - The score-normalization problem is avoided
  - The "absolute distance" between identities is lost

Rank-level Fusion

- Every biometric matcher ranks the identities in the databases
- Rank-level fusion consolidates the ranks associated with every subject

Database

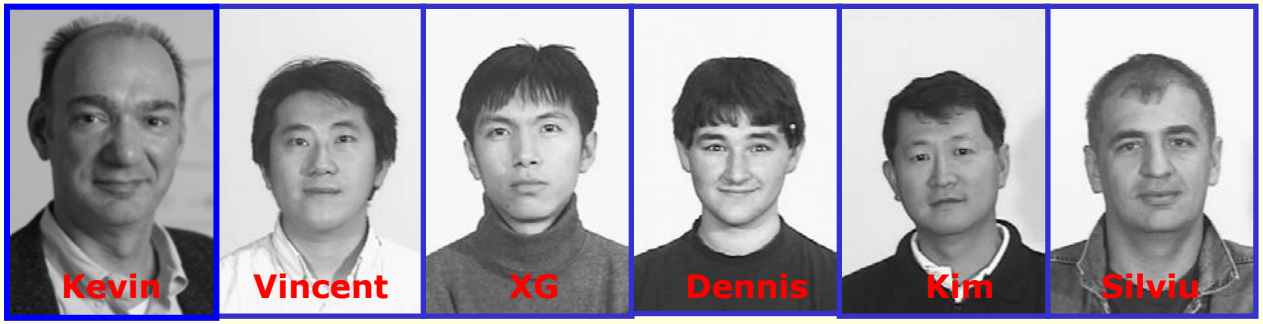

Face Matcher	1	4	5	2	6	3
Finger Matcher	1	3	2	5	6	4
Iris Matcher	2	4	6	1	5	3

Notation Used

- N: number of users enrolled in the database
- C: number of matchers
- r_{ij}: the rank assigned to user j by the i^{th} matcher
- R_j: the rank for user j after applying rank level fusion

Fusion Schemes

- **Highest Rank Fusion**: The fused rank of a user is computed as the best rank generated by different matchers

$$R_j = \min_{i=1}^{C} \left\{ r_{i,j} \right\}$$

- **Borda Count Fusion**: The fused rank of a user is computed as the sum of the ranks generated by different matchers

$$R_j = \sum_{i=1}^{C} r_{i,j}$$

T. Ho, J. Hull, and S. Srihari. Decision combination in multiple classifier systems. *IEEE Transaction on Pattern Analysis and Machine Intelligence (PAMI), 16(1):66–75, 1994.*

Decision-level Fusion

- Genuine or impostor?
  - 1 or 0?
- Fusion schemes
  - AND [Very strict]
  - OR [Very relaxed]
  - Majority Voting
  - Behavior Knowledge Space (BKS)

Ross et al, "Handbook of Multibiometrics," Springer 2006

Importance of Privacy

- "Privacy is the right to be let alone" [Samuel Warren and Louis Brandeis (1890)]
- "Privacy is the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others" [Alan Westin (1970)]
- "Privacy is the right of people to conceal information about themselves that others might use to their disadvantage" [Richard Posner (1983)]

PRIVACY IS DIFFERENT FROM SECURITY

Biometric Recognition

- Automated recognition of individuals based on their biological and behavioral characteristics
- Biological and behavioral characteristic of an individual from which distinguishing, repeatable biometric features can be extracted

H.T. F. Rhodes, Alphonse Bertillon: Father of Scientific Detection, Harrap, 1956

Identity vs Recognition

- We do not necessarily want to elicit identity
- We want to recognize a person

INPUT

Based on a single fingerprint image, we cannot say this belongs to *Jane Doe*

REFERENCE

We need a reference fingerprint image that is known to belong to *Jane Doe* in order to make this assessment

Reference Biometric Images

- Some biometric systems may store the raw images of an individual as a reference image
  - e.g., face or fingerprint or iris image

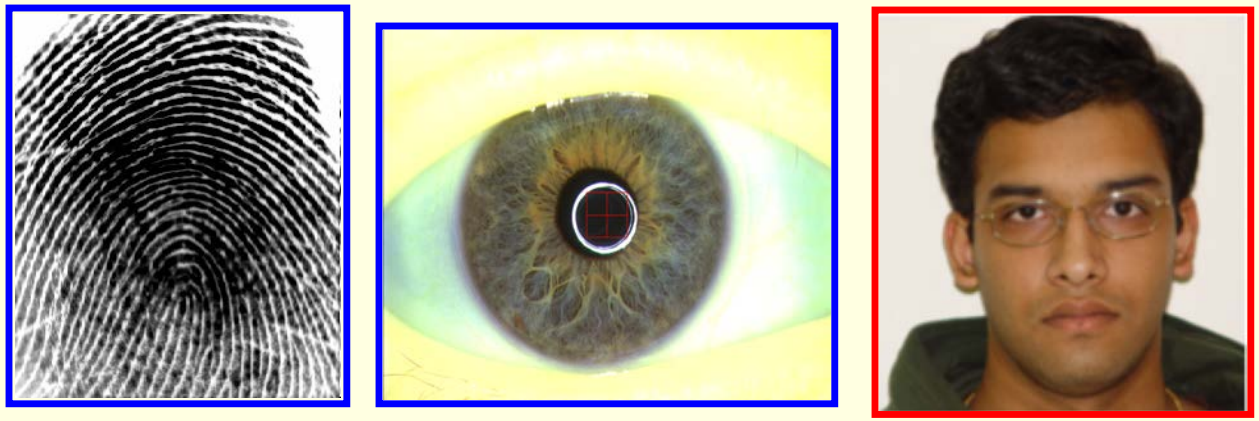

- From a visual standpoint, face images are perceived to divulge more information about a person

Privacy of Biometric Data

- Age, Gender, Ethnicity, can be **automatically derived** from the face image
- That is, a **trained classifier or a regressor** may be used to automatically deduce certain soft biometric attributes

- Gender: Male
- Age: 25
- Health: Very good
- Eye Sight: Wears glasses
- Ethnicity: Asian Indian

Dantcheva, Elias, Ross, "What Else Does Your Biometric Data Reveal? A Survey on Soft Biometrics," TIFS 2016

Iris: Levels of Information

- Biographical:

  Age, Gender, Race

- Anatomical:

  Distribution of crypts, Wolfflin nodules, pigmentation spots

- Environmental:

  Sensor, Illumination wavelength, Indoor/Outdoor

- Pathological:

  Stromal Atrophy

- Other:

  Pupil dilation level, Contact Lens

Not all information can be reliably extracted

But information can be aggregated

From Image to Sensor

IMAGE

SENSORS

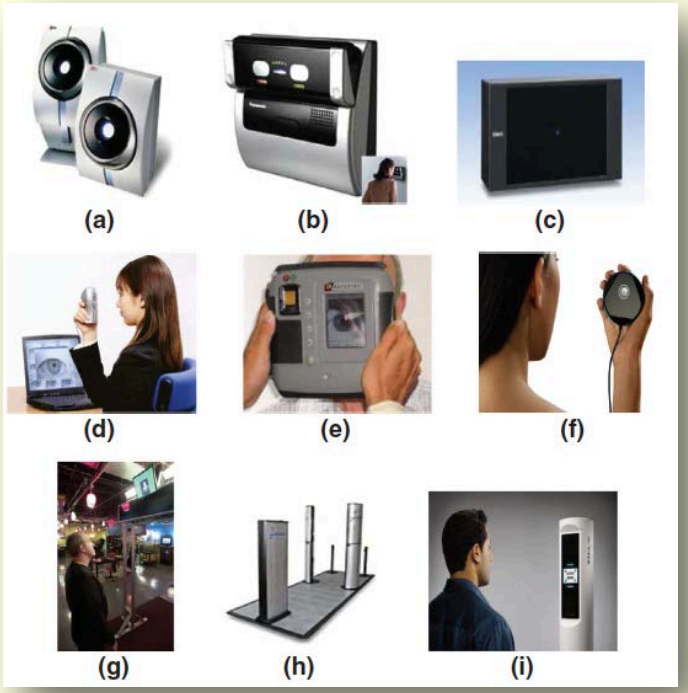

Photo Response Non Uniformity (PRNU):
Sensor Pattern Noise Present in Images

Biometrics + Forensics

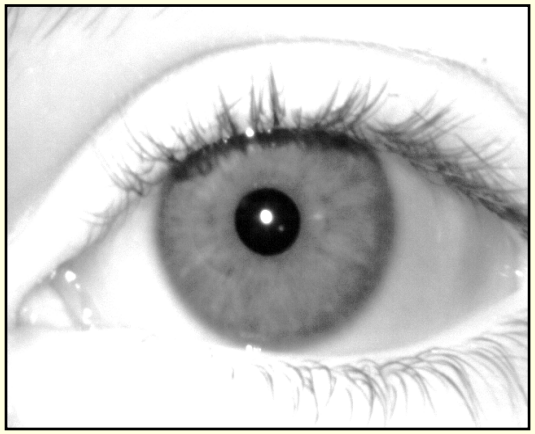

- Subject is a Male (90% Confidence), White (85% Confidence)
- Image taken using an Aoptix camera
- Iris stroma is plain textured
- Highly constricted pupil suggests strong ambient illumination

Bridges the gap between human and machine description of data

OR

Compromises privacy?

Face2Gene

MEGAN MOLTENI SCIENCE 01.09.17 01:00 PM

THANKS TO AI, COMPUTERS CAN NOW SEE YOUR HEALTH PROBLEMS

"In hindsight it was all clear to me," says Gripp, who is chief of the Division of Medical Genetics at A.I. duPont Hospital for Children in Delaware, and had been seeing the patient for years. "But it hadn't been clear to anyone before." What had taken Patient Number Two's doctors 16 years to find took Face2Gene just a few minutes.

Face2Gene is a suite of phenotyping applications that facilitate comprehensive and precise genetic evaluations.

Identifying People on the Web

- **Faces of Facebook: Privacy in the Age of Augmented Reality (Alessandro Acquisti)**
- Convergence of three technologies:
  - face recognition, cloud computing, online social networks
- Started from an anonymous face in the street
- Ended up with very sensitive information about that person → **data accretion**
- Combined face recognition with the algorithms they developed in 2009 to predict SSNs from public data

Importance of Privacy

- "Privacy is the right to be let alone" [Samuel Warren and Louis Brandeis (1890)]
- "Privacy is the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others" [Alan Westin (1970)]
- "Privacy is the right of people to conceal information about themselves that others might use to their disadvantage" [Richard Posner (1983)]

PRIVACY IS DIFFERENT FROM SECURITY

Privacy Visor

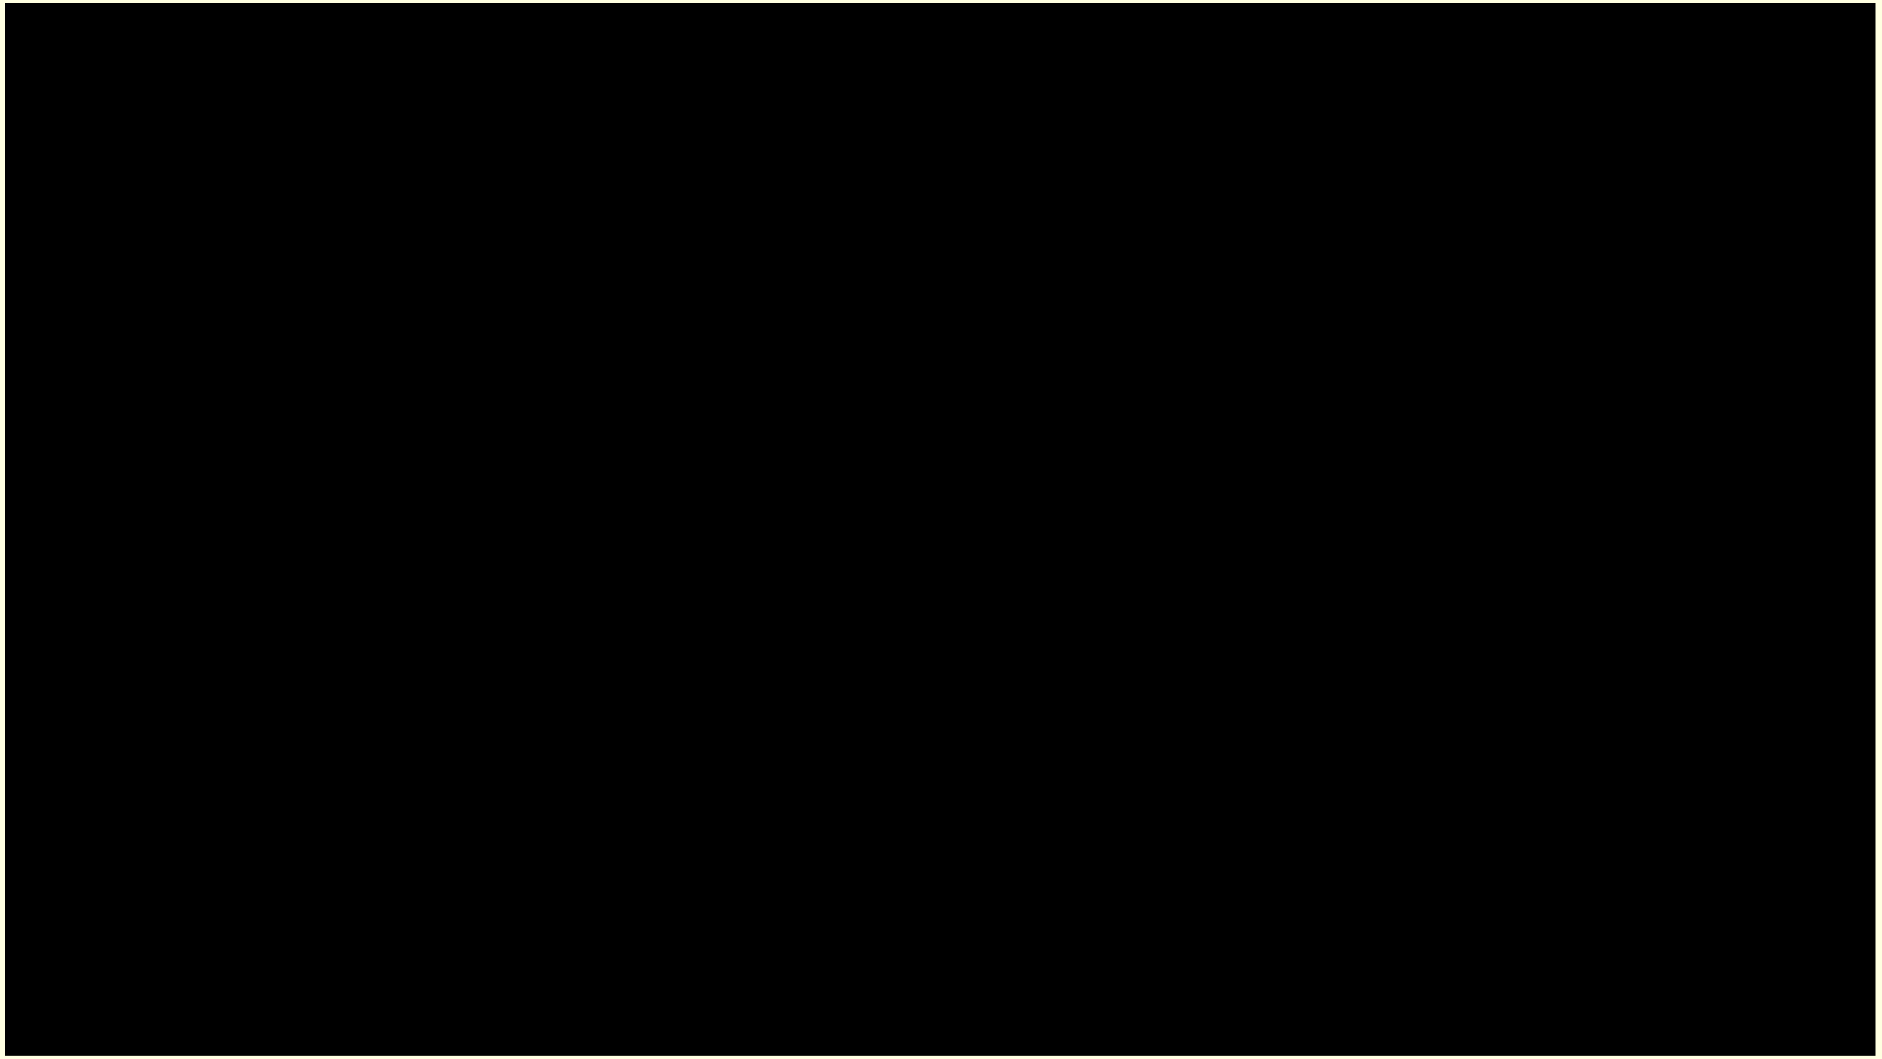

https://www.youtube.com/watch?v=LRj8whKmN1M

Anti-Face!

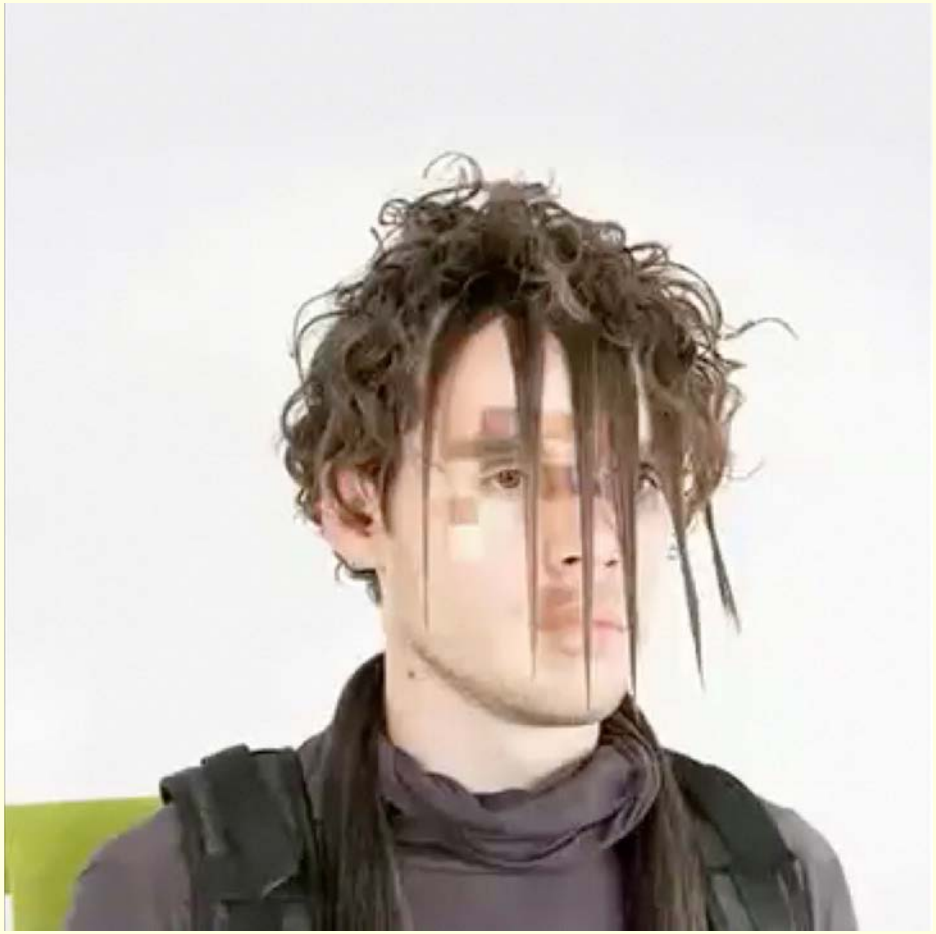

https://cvdazzle.com/

“Differential” Privacy

Differential Privacy

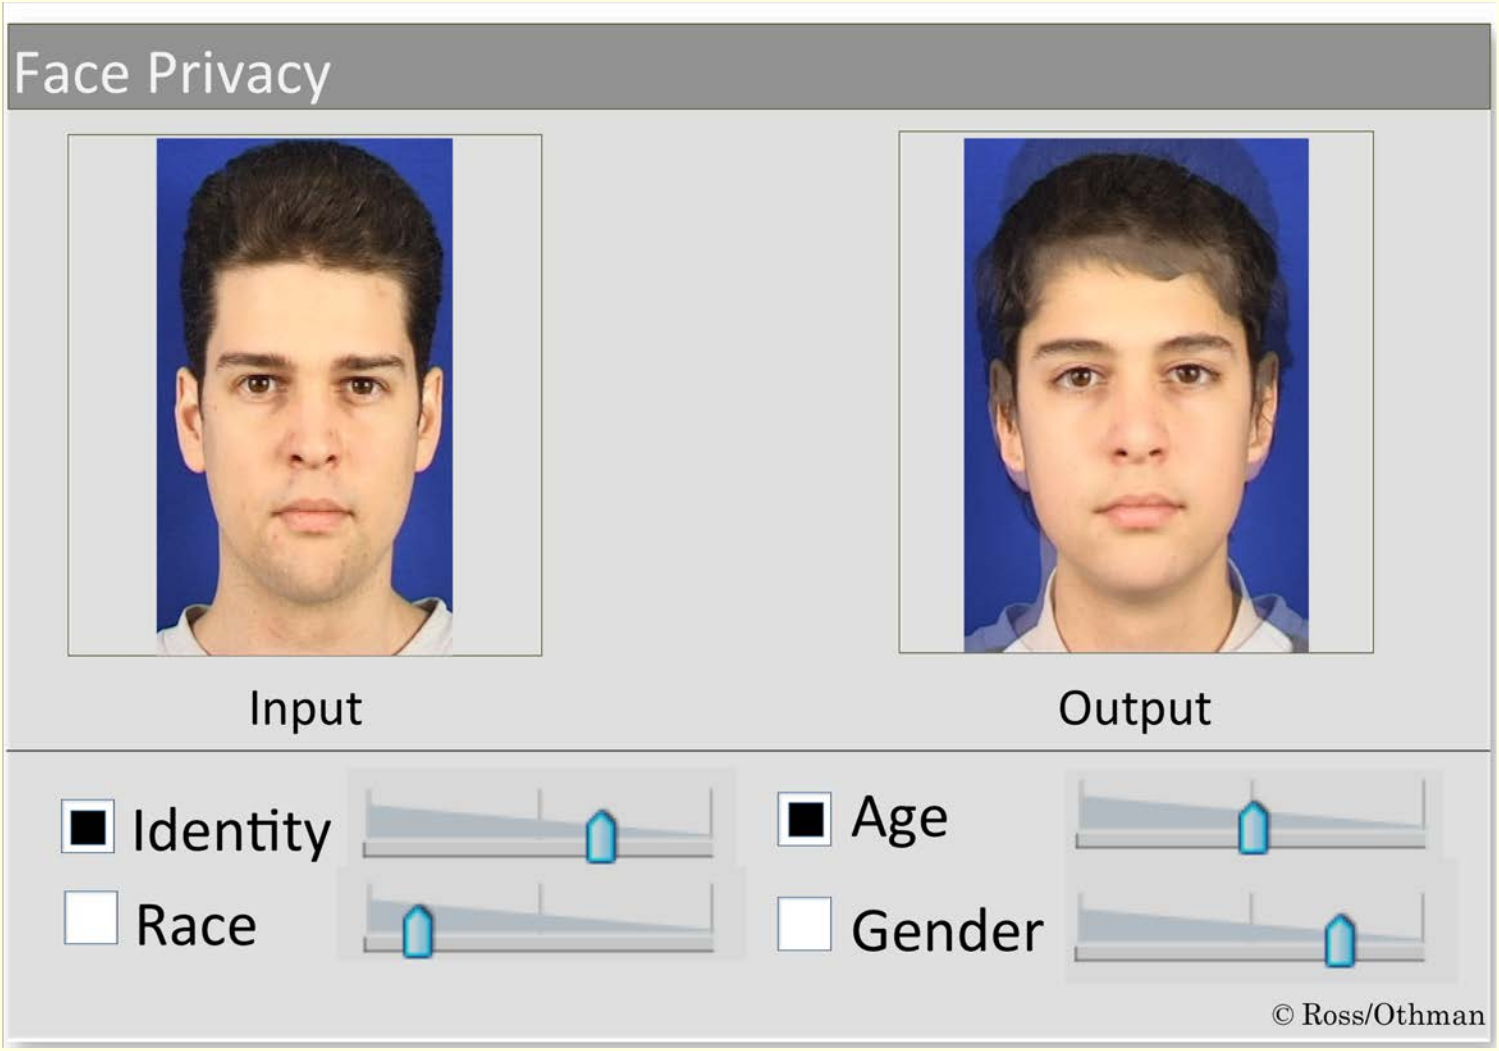

Othman and Ross, "Privacy of Facial Soft Biometrics," ECCVW 2014

Differential Privacy

- We investigate the possibility of preserving the contextual integrity of face images stored in a central biometric database
- We consider the problem of suppressing a soft biometric attribute of a face
- This modification should not drastically impact the accuracy of the automated face matcher

Soft Biometric Privacy

- Gender attribute of an input face image is progressively suppressed
- With respect to a face matcher the recognition capability is preserved

Input image Transformed images

Name	Alice	Alice	Alice	Alice
Gender	Female (confident)	Female (less confident)	Male (less confident)	Male (confident)

Othman and Ross, "Privacy of Facial Soft Biometrics: Suppressing Gender But Retaining Identity", ECCV Workshop, 2014

Semi-Adversarial Networks (SAN)

- Design a transformation model to:
  - Confound gender attribute ➔ gender classifiers will not work
  - Retain recognition capability ➔ face matchers will still work

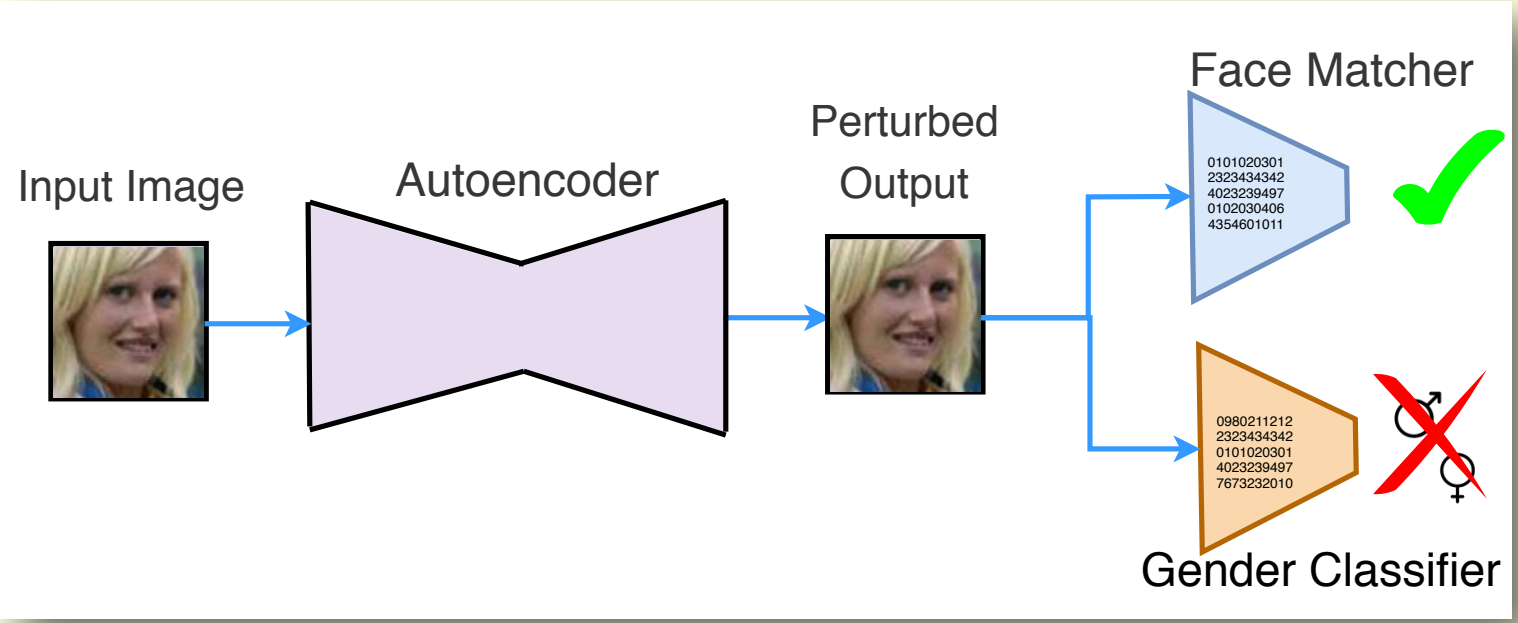

Mirjalili et al., Semi-Adversarial Networks: Convolutional Autoencoders for Imparting Privacy to Face Images, ICB 2018

General Architecture of SAN Model

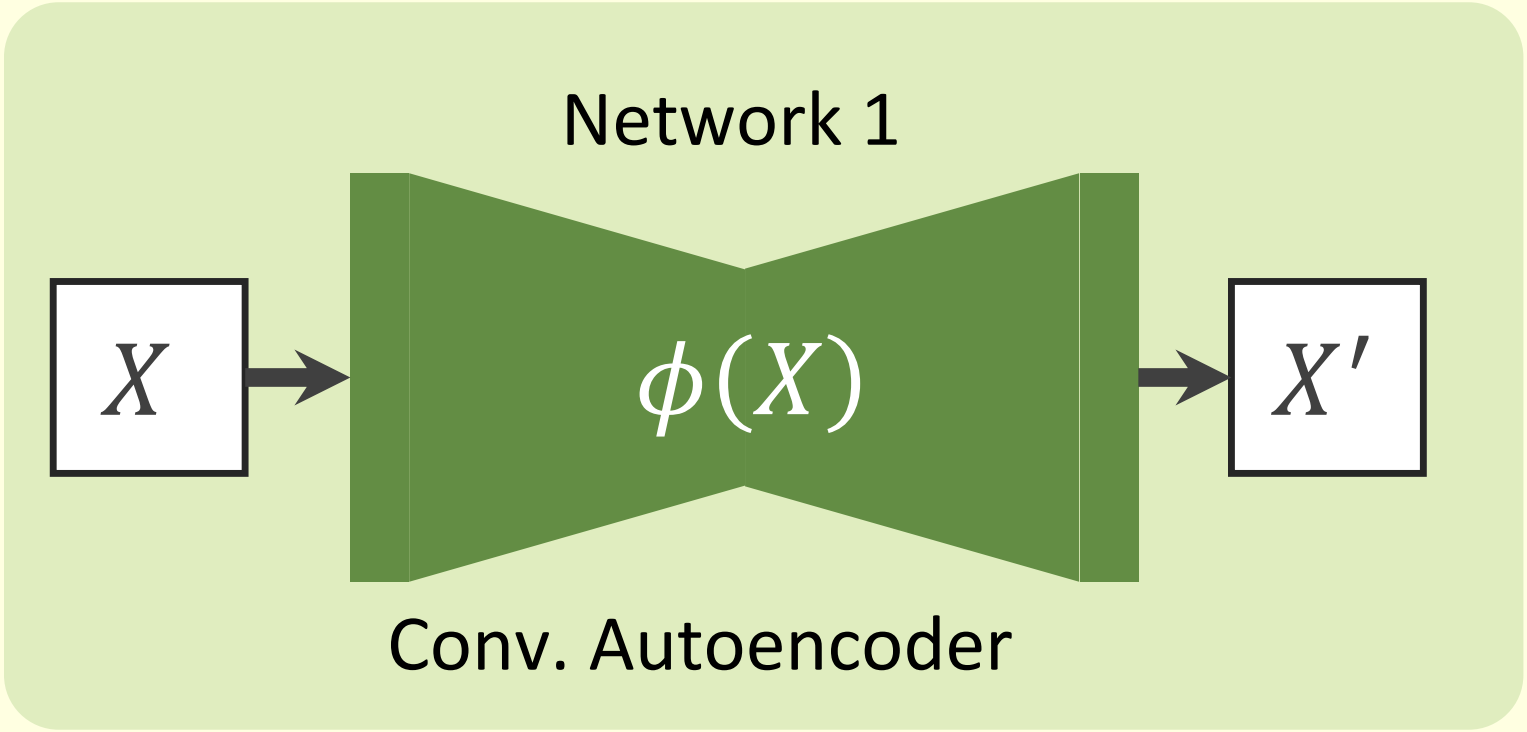

Mirjalili et al., Semi-Adversarial Networks: Convolutional Autoencoders for Imparting Privacy to Face Images, ICB 2018

General Architecture of SAN Model

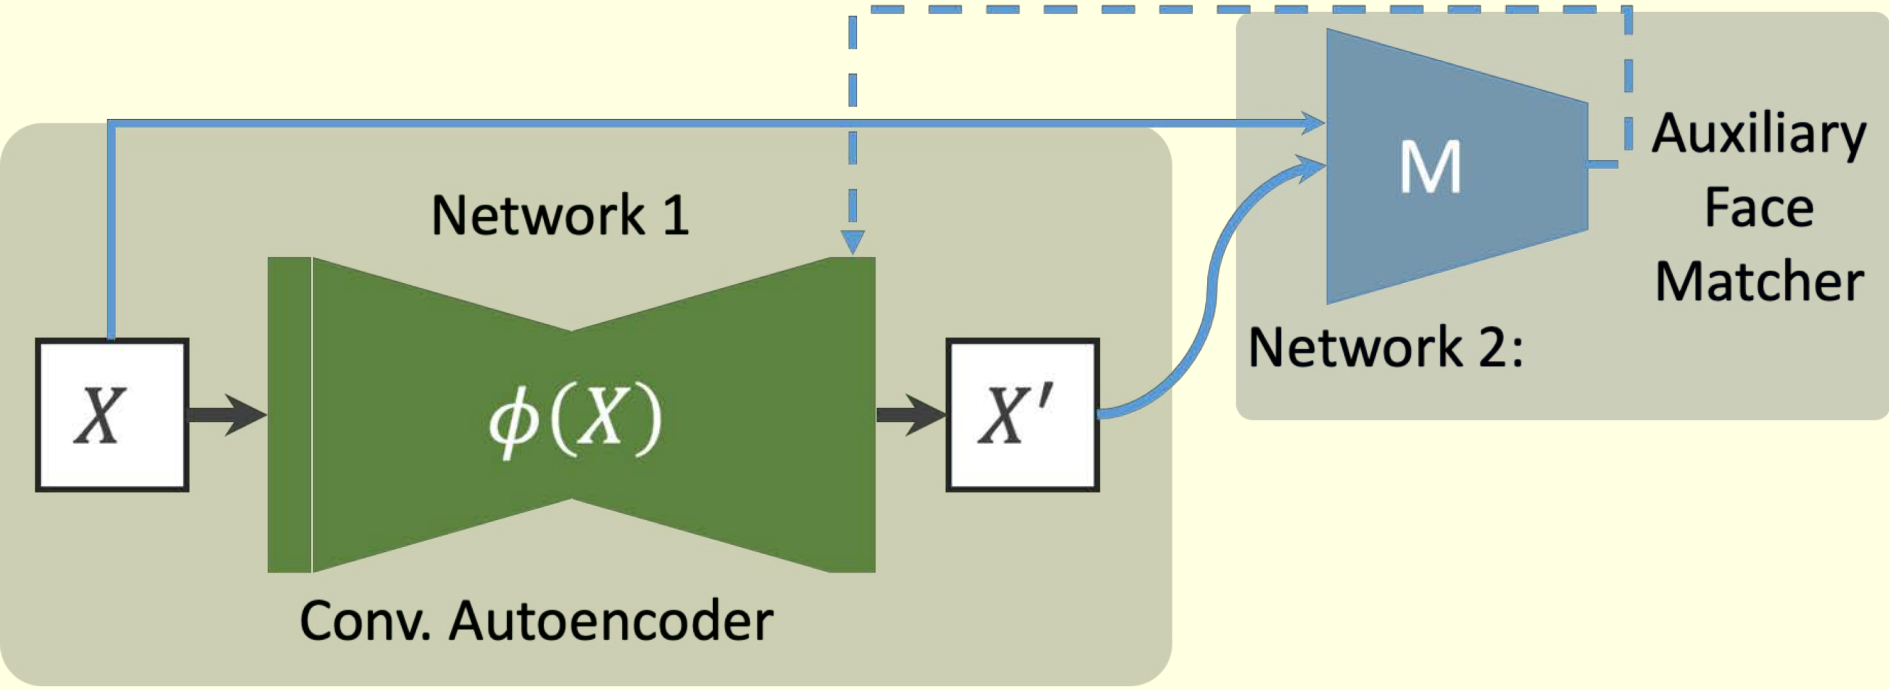

Mirjalili et al., Semi-Adversarial Networks: Convolutional Autoencoders for Imparting Privacy to Face Images, ICB 2018

General Architecture of SAN Model

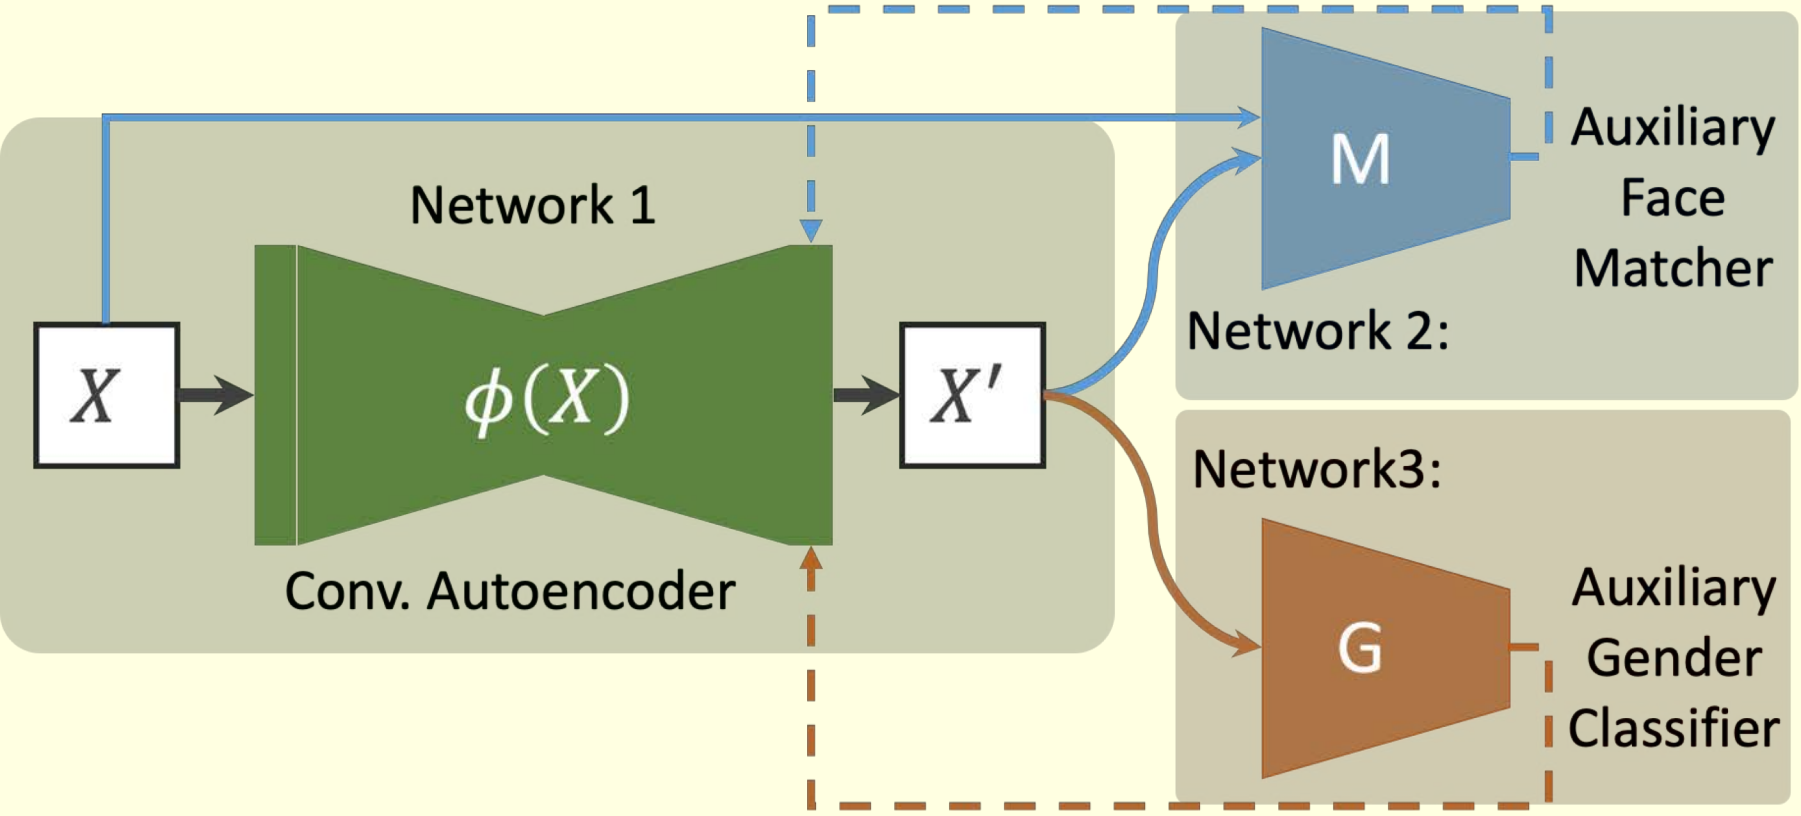

Mirjalili et al., Semi-Adversarial Networks: Convolutional Autoencoders for Imparting Privacy to Face Images, ICB 2018

Ensemble of SANs

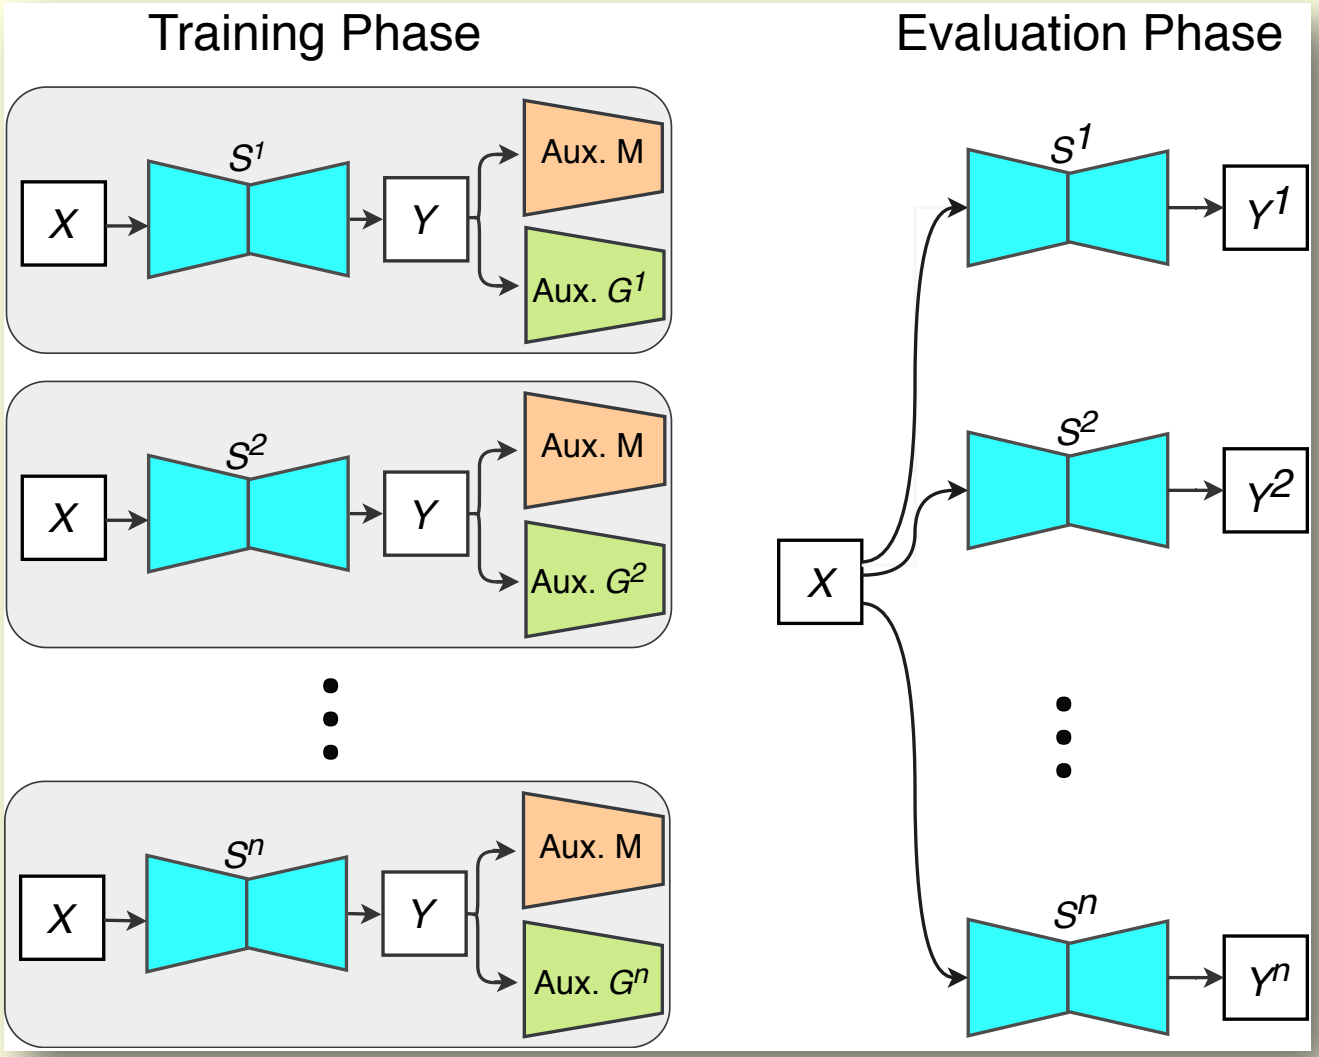

V. Mirjalili, S. Raschka, A. Ross, "Gender Privacy: An Ensemble of Semi Adversarial Networks for Confounding Arbitrary Gender Classifiers," BTAS 2018

Cost Functions for Semi-Adversarial Learning

1. Pixel-wise similarity term

$$J_D(X, X'_{SM}) = \sum_{k=1}^{N} S\left(X^{(k)}, X'^{(k)}_{SM}\right)$$

- Only used during the pre-training of Autoencoder

2. Loss term related to gender attribute

- Correctly predict gender of X'_{SM}
- Flip the gender prediction on X'_{OP}

$$J_G(X, X'_{SM}, X'_{OP}, y; f_G) = S\left(y, f_G(X'_{SM})\right) + S\left(1 - y, f_G(X'_{OP})\right)$$

3. Loss term related to face identity matching

$$J_M\left(X, X'_{SM}; R_{vgg}\right) = \left|\left|R_{vgg}(X'_{SM}) - R_{vgg}(X)\right|\right|_2^2$$

Training Protocol

- **Auxiliary subnetworks**
  - Auxiliary gender predictor is trained on CelebA dataset, and its parameters are frozen during training of Conv. Autoencoder
  - Publicly available parameters for VGG are used for the auxiliary face matcher
- **Training the Autoencoder**

Step1: pre-training the Conv. Autoencoder with two loss terms: pixel-wise similarity + gender term

Step2: replace the pixel-wise similarity term with the matching term based on VGG subnetwork (trained for 20 epochs)

Examples of Inputs and Outputs

Male:
99%

Female:
98%

Male:
97%

Male:
100%

Female:
69%

Male:
99%

Male:
71%

Female:
58%

Mirjalili et al., Semi-Adversarial Networks: Convolutional Autoencoders for Imparting Privacy to Face Images, ICB 2018

Examples of Inputs and Outputs

Male:
98%

Male:
99%

Female:
100%

Female:
99%

Female:
79%

Female:
53%

Male:
63%

Male:
67%

Mirjalili et al., Semi-Adversarial Networks: Convolutional Autoencoders for Imparting Privacy to Face Images, ICB 2018

Examples of Inputs and Outputs

Male:
100%

Male:
85%

Female:
100%

Female:
99%

Female:
95%

Female:
51%

Male:
75%

Male:
78%

Mirjalili et al., Semi-Adversarial Networks: Convolutional Autoencoders for Imparting Privacy to Face Images, ICB 2018

Examples of Inputs and Outputs

Male:
99%

Male:
88%

Male:
99%

Male:
94%

Male:
52%

Female:
91%

Female:
56%

Female:
93%

Mirjalili et al., Semi-Adversarial Networks: Convolutional Autoencoders for Imparting Privacy to Face Images, ICB 2018

Examples of Inputs and Outputs

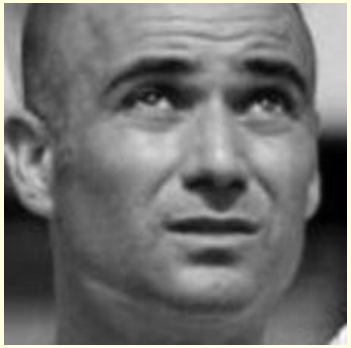

Male:
98%

Female:
72%

Female:
94%

Female:
99%

Male:
85%

Female:
80%

Male:
95%

Male:
52%

Mirjalili et al., Semi-Adversarial Networks: Convolutional Autoencoders for Imparting Privacy to Face Images, ICB 2018

Datasets Statistics

Dataset	# Samples	# Subjects	# Male Images	# Female Images
CelebA-train	157,350	--	65,160	92,190
CelebA-test	39,411	--	16,318	23,093
MUCT	3,754	276	1,844	1,910
LFW	12,988	5,658	10,083	2,905
AR-face	3,286	136	1,821	1,465

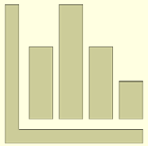

- CelebA dataset was split into train and test
- CelebA-train was used for training the autoencoder as well as the auxiliary gender predictor

Experimental Design

- Six unseen gender Classifiers
  - G-COTS [Commercial]
  - IntraFace [De la Torre et al., 2015]
  - AFFACT [Günther et al., 2017]
  - 3 CNN models [in-house]
- Four unseen face Matchers
  - M-COTS [Commercial]
  - DR-GAN [Tran et al., 2017]
  - FaceNet [Schroff et al., 2015]
  - OpenFace [Amos et al., 2016]

Unseen:
the classifier or face matcher is not used during training of the SAN models

Performance Assessment on MUCT dataset: Confound gender classifiers

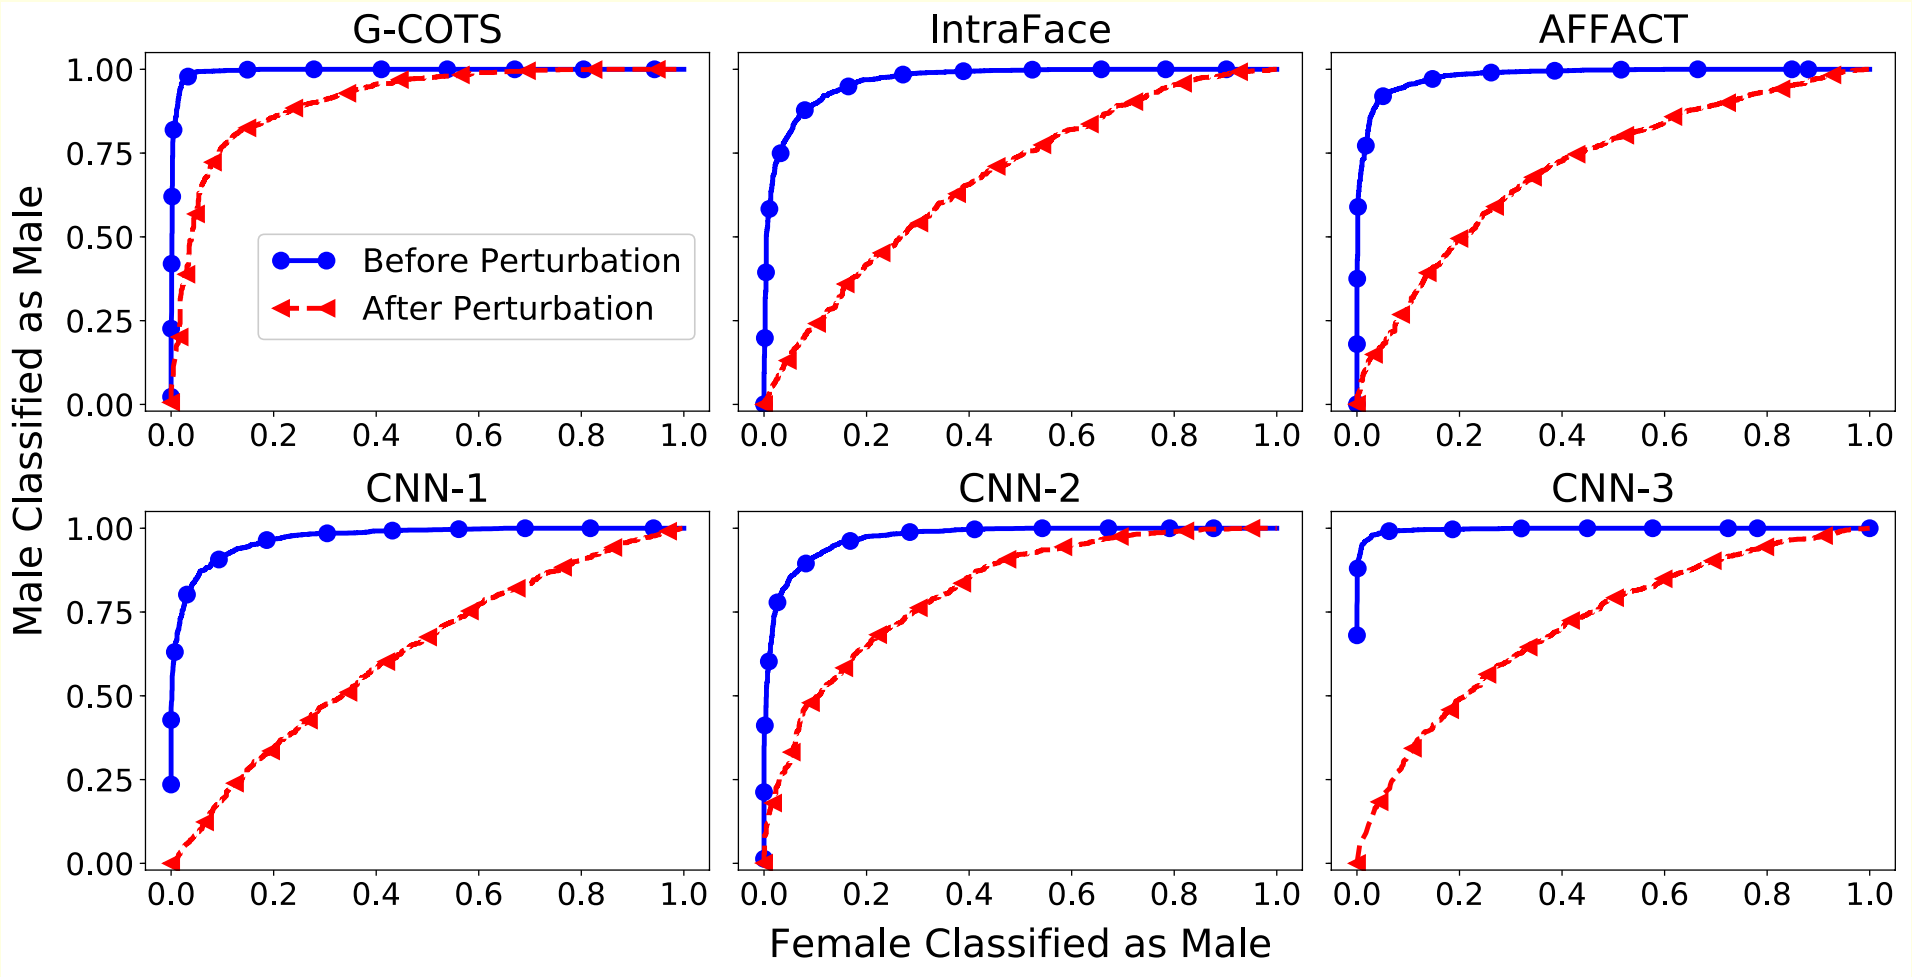

Performance Assessment on MUCT dataset: Retain Matching Capability

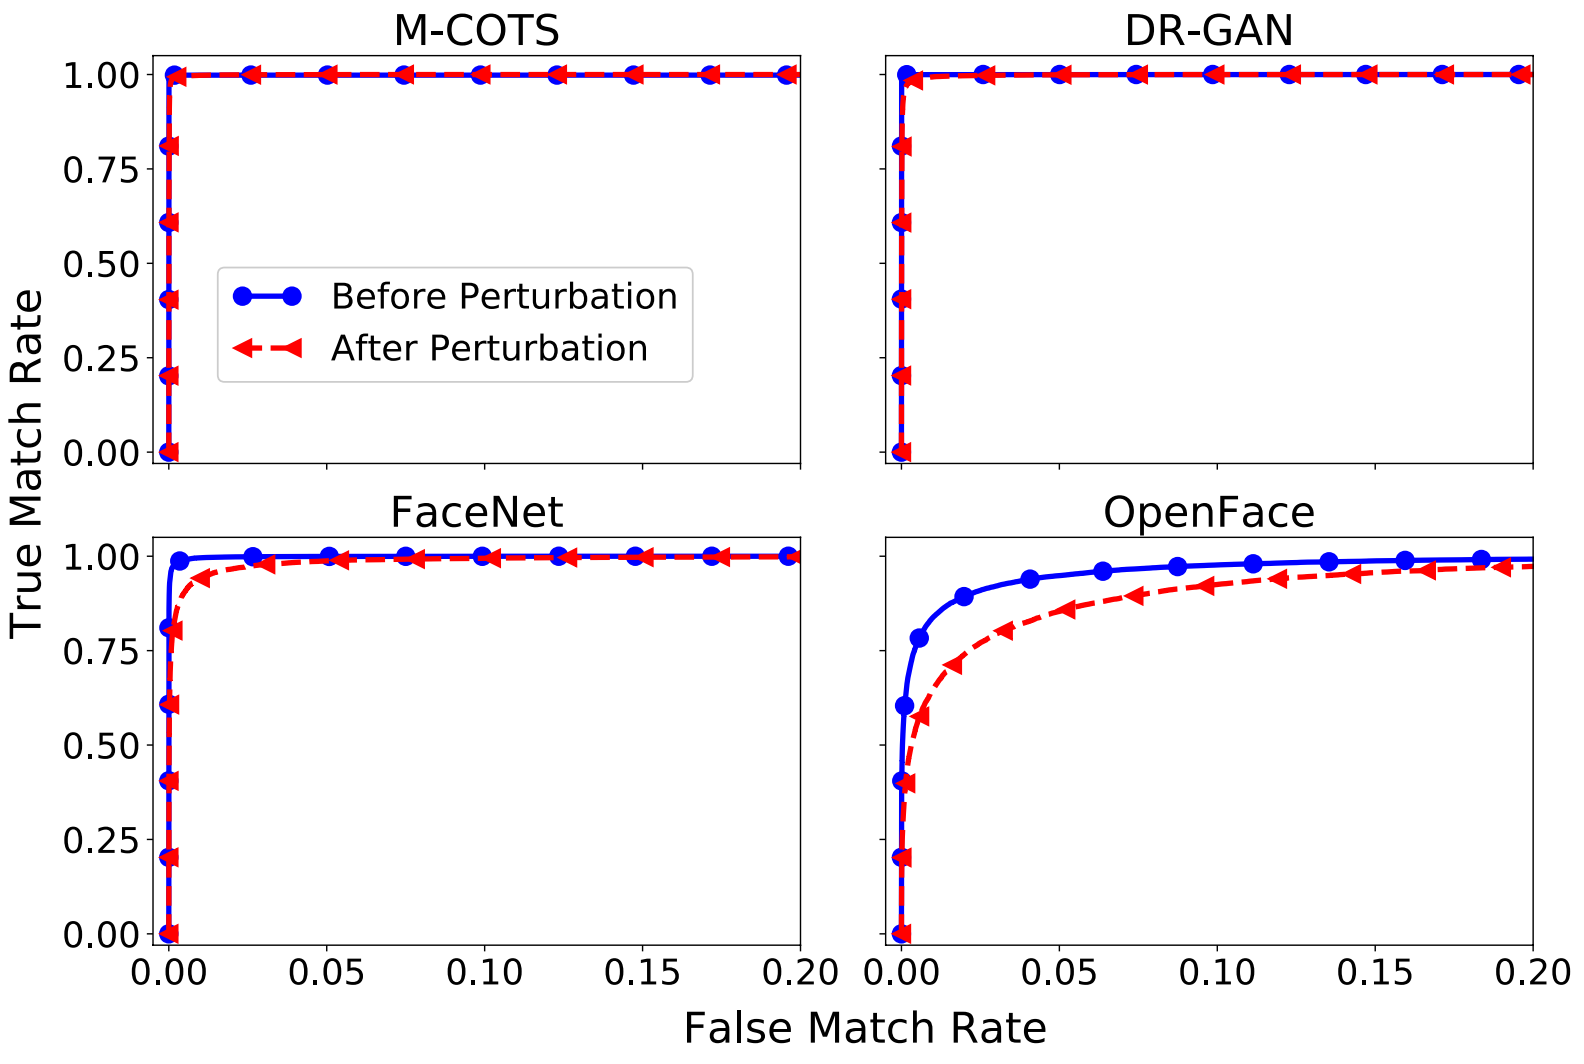

Recent Publications

- V. Mirjalili. S. Raschka, A. Ross, "**Gender Privacy: An Ensemble of Semi Adversarial Networks for Confounding Arbitrary Gender Classifiers**," BTAS 2018.

- V. Mirjalili, S. Raschka, A. Namboodiri, A. Ross, "**Semi-Adversarial Networks: Convolutional Autoencoders for Imparting Privacy to Face Images**," ICB 2018.

- V. Mirjalili and A. Ross, "**Soft Biometric Privacy: Retaining Biometric Utility of Face Images while Perturbing Gender**," IJCB 2017.

De-identification via Collaboration

Decomposing Face Images

- The input face image is decomposed and stored in two separate servers: either server will be unable to deduce original face image by themselves

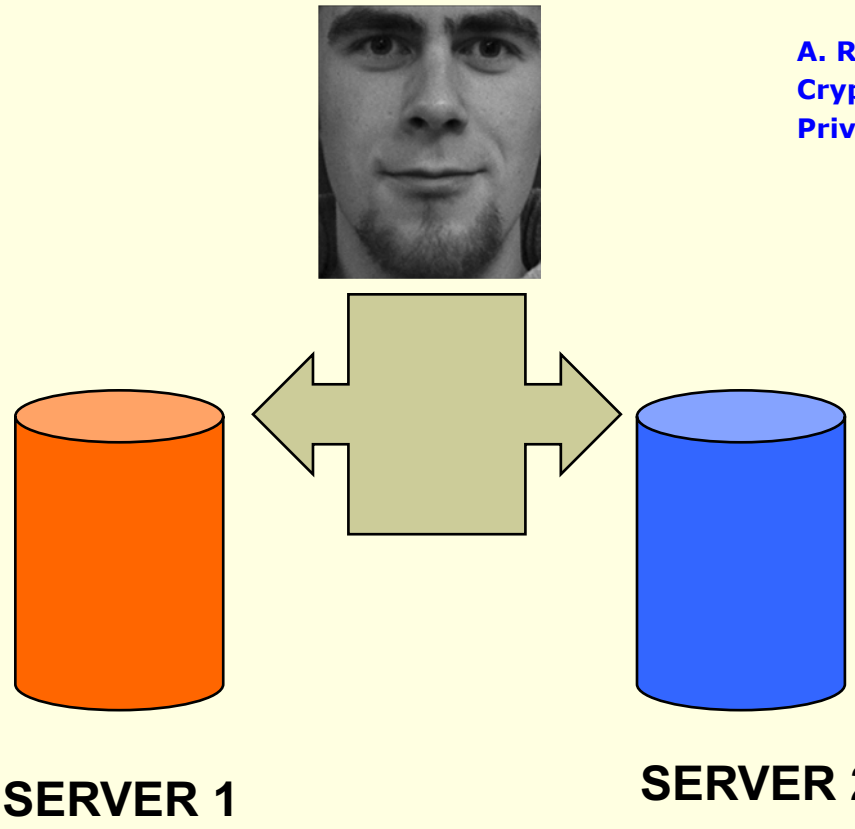

A. Ross and A. Othman, "Visual Cryptography for Biometric Privacy," TIFS 2011

Visual Cryptography*

- Given an original binary image *T*, it is encrypted in *n* images, such that:

$$T = S_{h_1} \oplus S_{h_2} \oplus S_{h_3} \oplus \ldots \oplus S_{h_k}$$

where $\oplus$ is a Boolean operation , S_{hi} is an image which appears as noise, $k \leq n$, and n is the number of noisy images

- This is referred to as *k-out-of-n* VCS

* M. Naor and A. Shamir, "Visual cryptography," in EUROCRYPT, pp. 1–12, 1994.

2-out-of-2 VCS

Pixel	Probability	Shares #1	Shares #2	Superposition of the two shares	
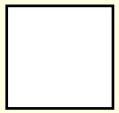	$p = 0.5$				White Pixels
	$p = 0.5$				
	$p = 0.5$				Black Pixels
	$p = 0.5$				

Decomposing a Binary Image

- Decomposing a fingerprint into two random images using Visual Cryptography

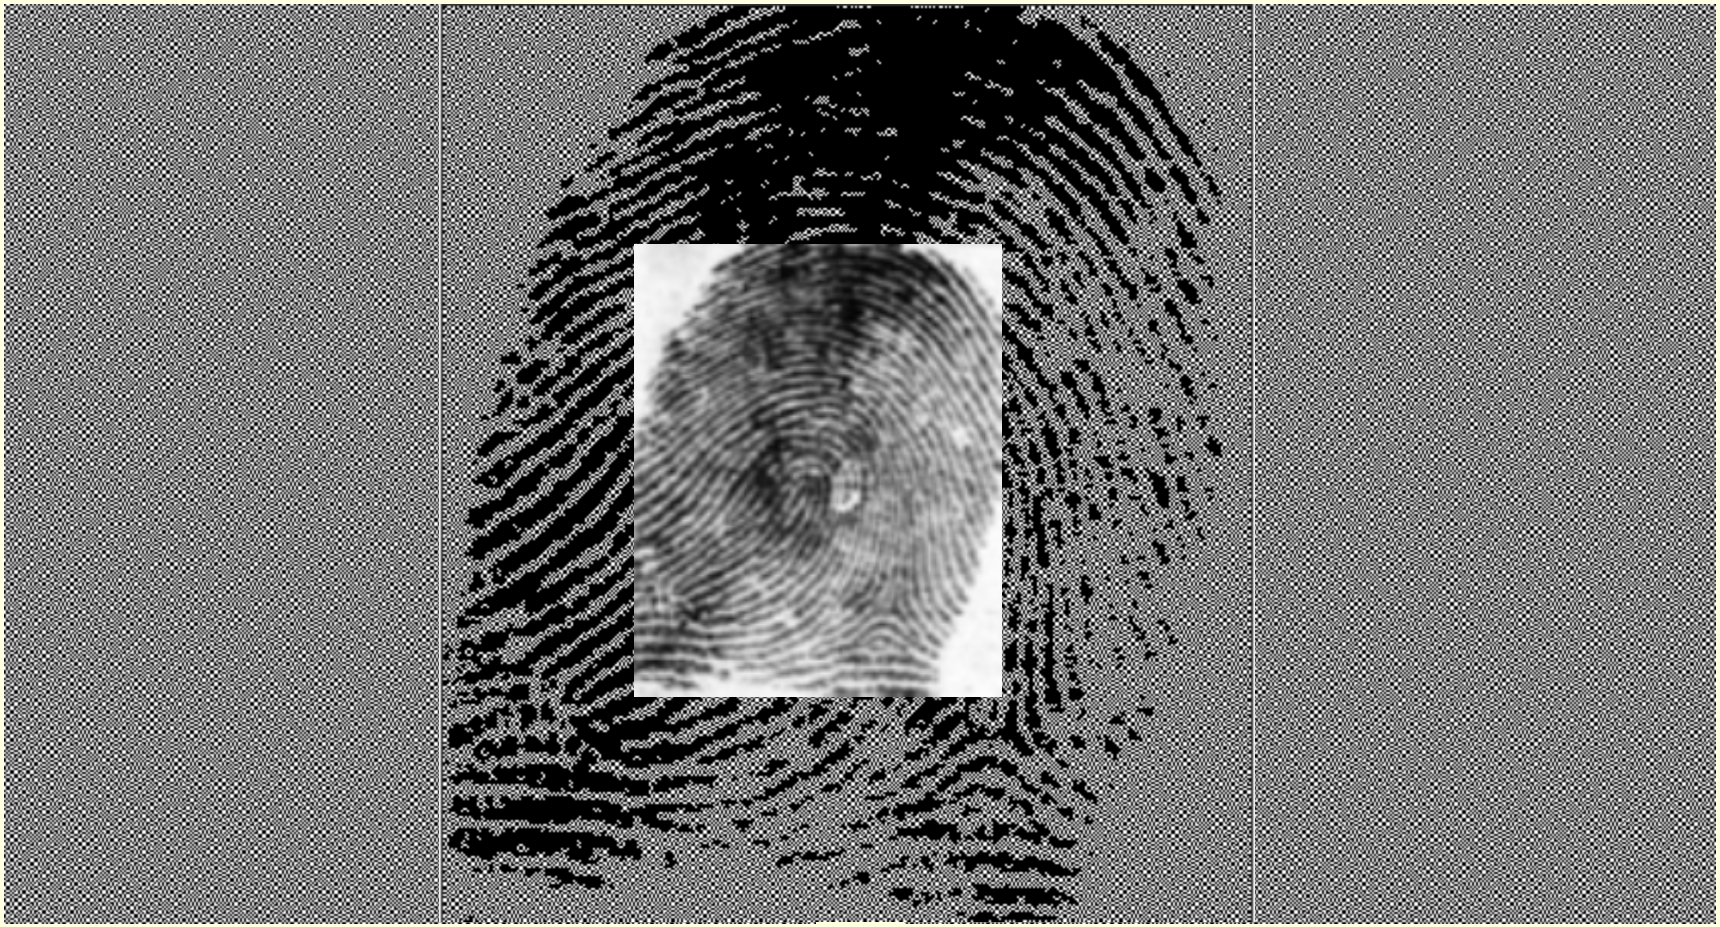

Decomposing a Face Image

- Decomposing a face into two random images?
  Problematic!

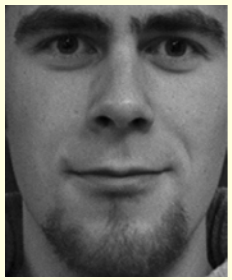

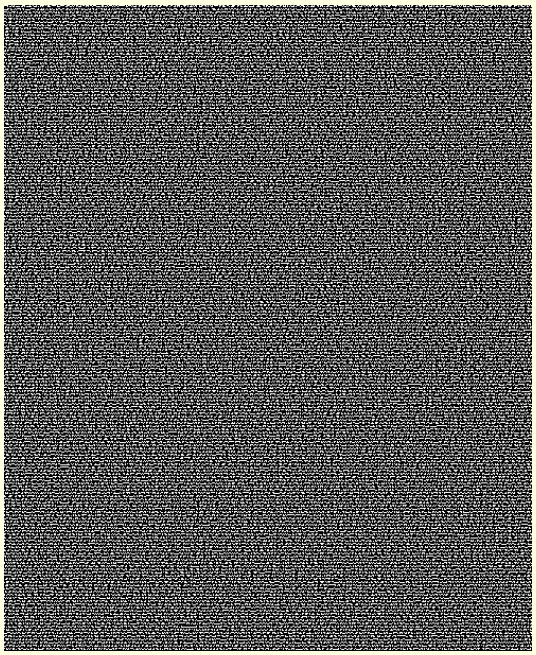

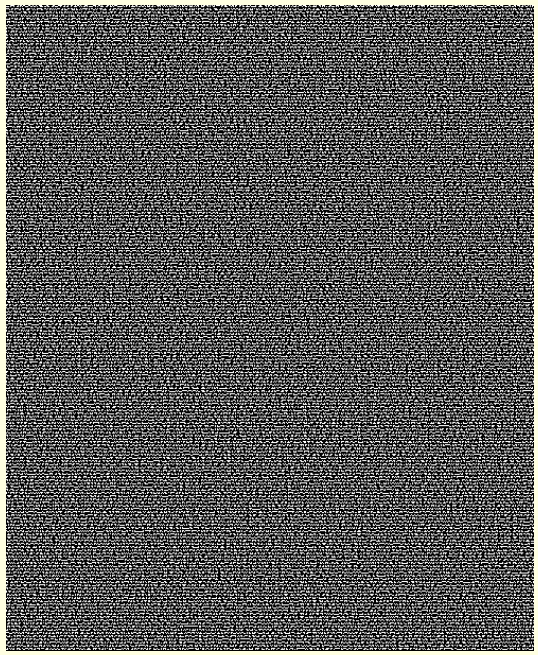

Gray-level Extended Visual Cryptography Scheme (GEVCS)

- VCS allows us to encode a secret image into n sheet images
- These sheets appear as a random set of pixels
- The sheets could be reformulated as natural images – known as host images

M. Nakajima and Y. Yamaguchi, "Extended visual cryptography for natural images," Journal of WSCG 10(2), pp. 303–310, 2002.

Gray-level Extended Visual Cryptography Scheme (GEVCS)

PRIVATE IMAGE

HOSTS (PUBLIC IMAGES)

PRIVATE IMAGE
AFTER DECRYPTION

HOSTS AFTER ENCRYPTION

Ross and Othman, "Visual Cryptography for Biometrics Privacy", TIFS 2011

Automated Host Image Selection

- The original image is encrypted into two dynamically selected host images

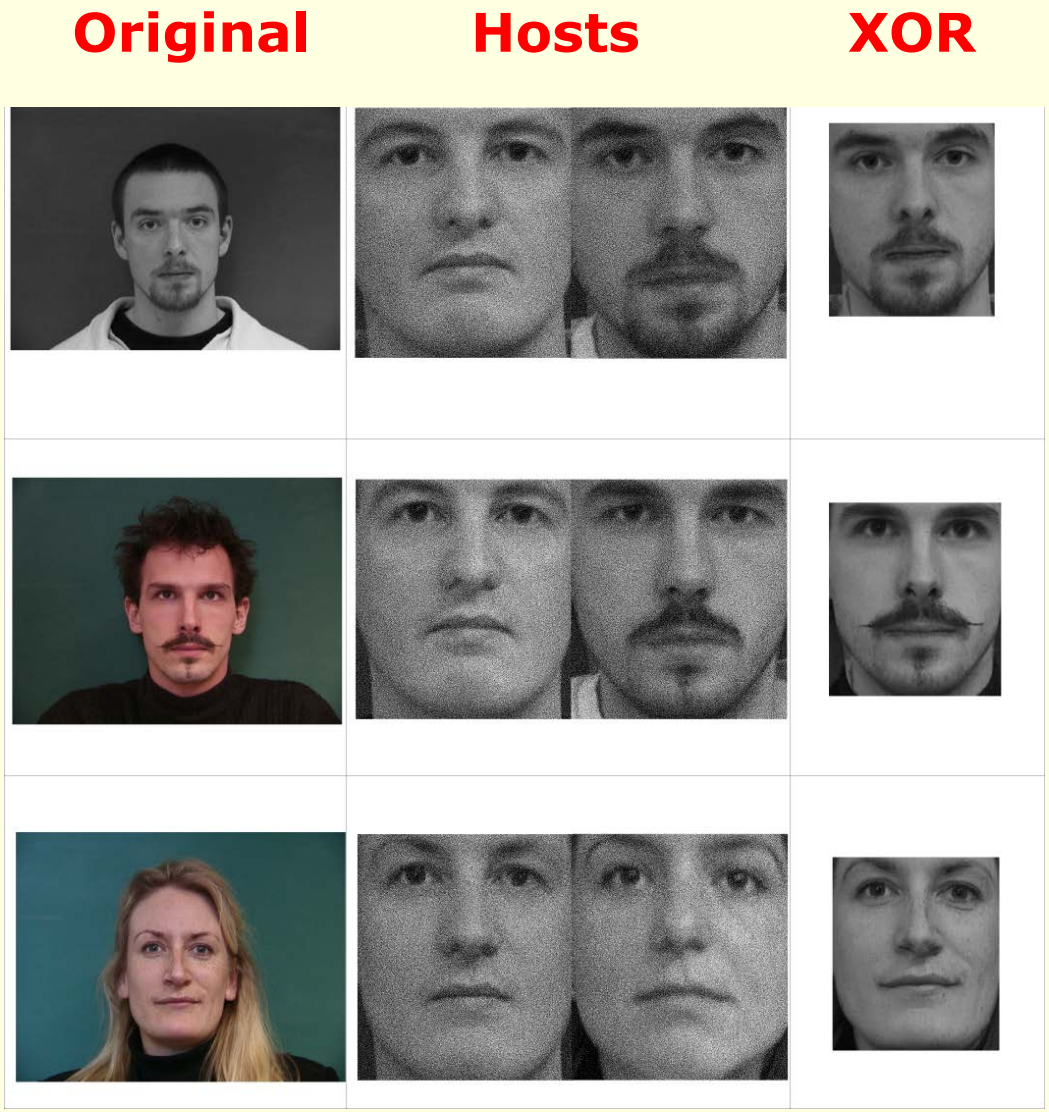

Face Visual Cryptography

Actual Face

=

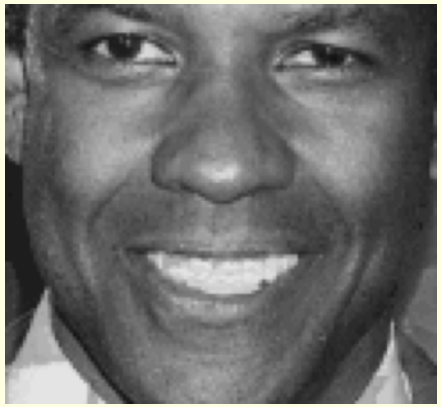

HOST IMAGE IN SERVER 1

+ ***Simple XOR operator***

HOST IMAGE IN SERVER 2

Ross and Othman, "Visual Cryptography for Biometrics Privacy", TIFS 2011

Face De-identification: Results

- Method to protect privacy of face images by decomposing it into two independent host (public) face images
- Original face image can be reconstructed only when both host images are available
- Either host image does not expose the identity of the original face image

De-identification via Mixing

Mixing Fingerprints

- An input fingerprint image is mixed with another fingerprint (e.g., from a different finger)
  - produces a new mixed fingerprint image that obscures the identity of the original fingerprint
- We consider the problem of mixing two fingerprint images in order to generate a new cancelable fingerprint image

Othman and Ross, "On Mixing Fingerprints", TIFS 2013

Applications

- To obscure the information present in an individual's fingerprint image prior to storing it in a central database
- To generate a cancelable template, i.e., the template can be reset if the mixed fingerprint is compromised
- To generate virtual identities by mixing fingerprint images pertaining to an individual

Othman and Ross, "On Mixing Fingerprints", TIFS 2013

Mixing Fingerprints

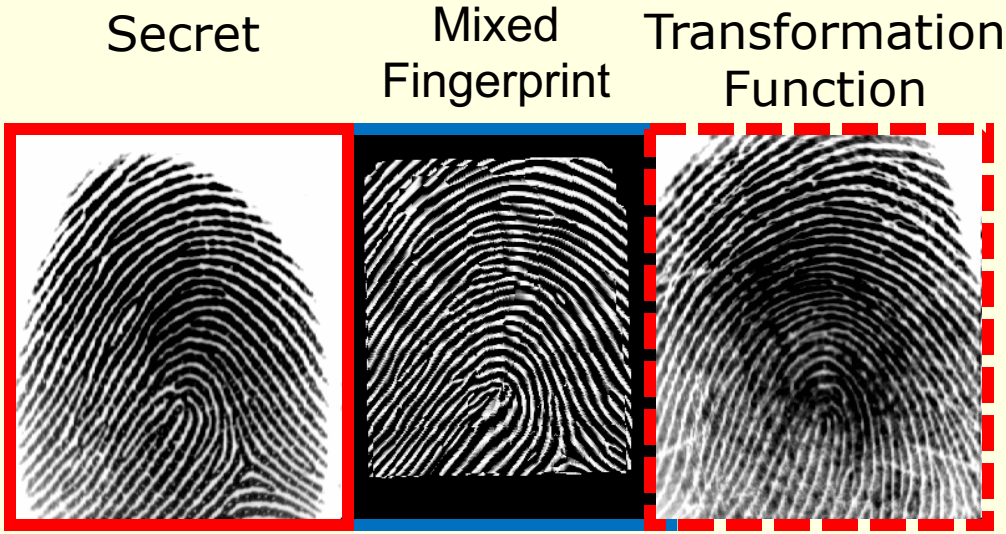

- Mixing fingerprints creates a new entity that looks like a plausible fingerprint:
  - It can be processed by conventional fingerprint algorithms
  - An eavesdropper may not be able to determine if a given fingerprint is mixed or not

Hologram Model

- The ridge flow of a fingerprint can be represented as a 2D Amplitude and Frequency Modulated (AM-FM) signal:

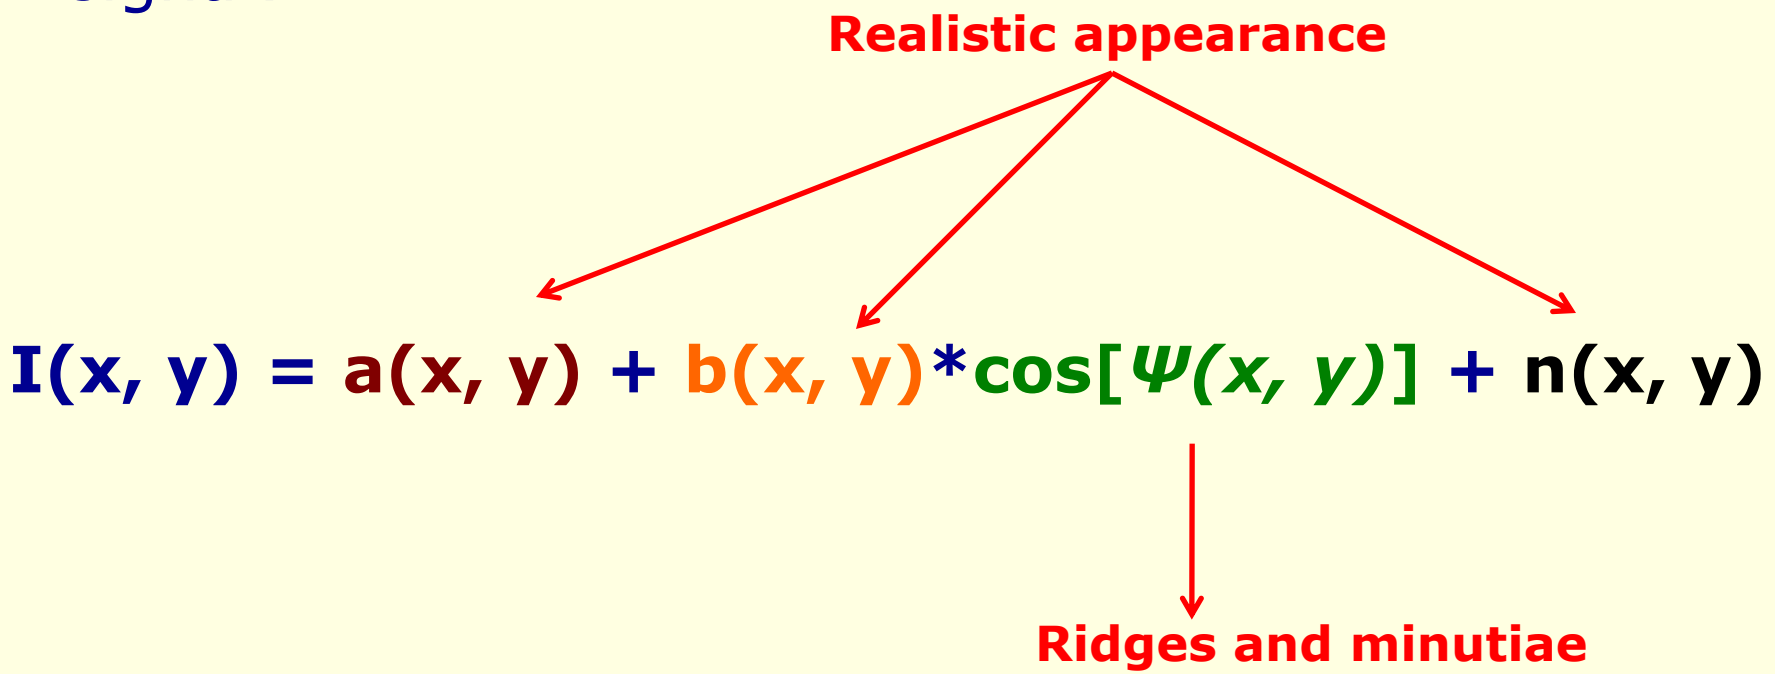

K. G. Larkin and P. A. Fletcher. A coherent framework for fingerprint analysis: are fingerprints holograms? Opt. Express, 15(14):8667–8677, 2007.

Helmholtz Decomposition

- Based on the Helmholtz Decomposition theorem, the phase **Ψ(x, y)** can be uniquely decomposed into two components:

$$\Psi(x, y) = \Psi c(x, y) + \Psi s(x, y)$$

- The continuous component, **Ψc(x, y)**, defines the local ridge orientation
- The spiral component, **Ψs(x, y)**, characterizes the minutiae locations

Decomposition: Left Loop

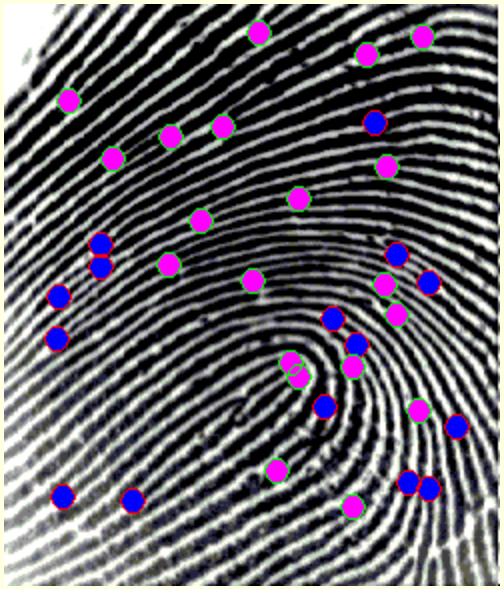

Original

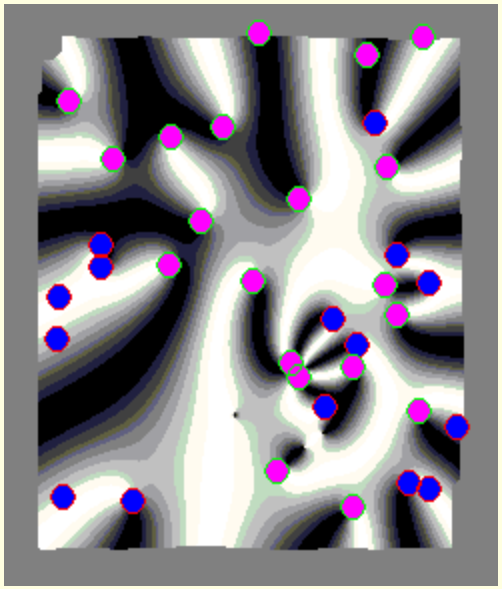

Spiral Phase

Continuous Phase

Othman and Ross, "On Mixing Fingerprints", TIFS 2013

Decomposition: Right Loop

Original

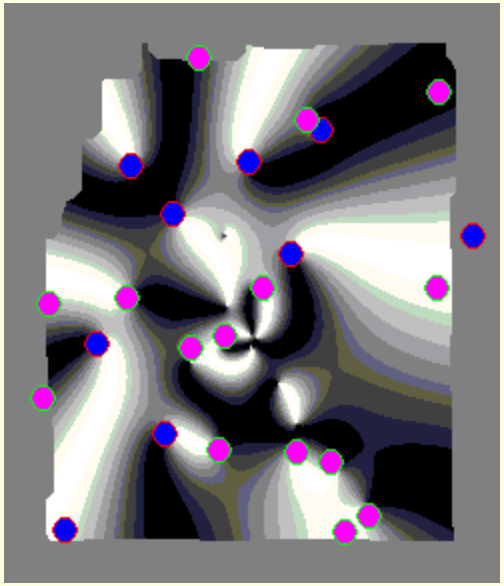

Spiral Phase

Continuous Phase

Othman and Ross, "On Mixing Fingerprints", TIFS 2013

Decomposition: Whorl

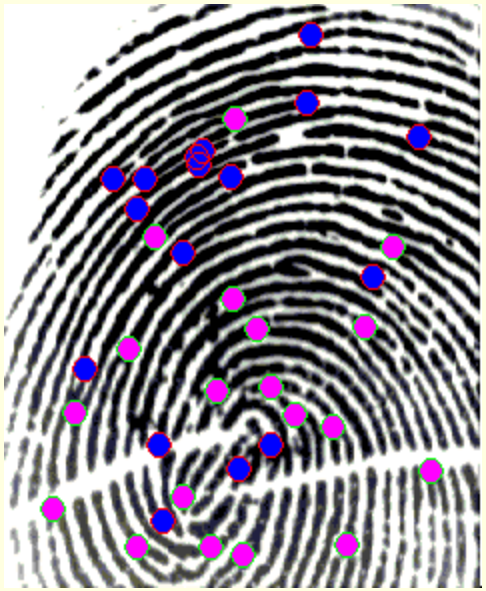

Original

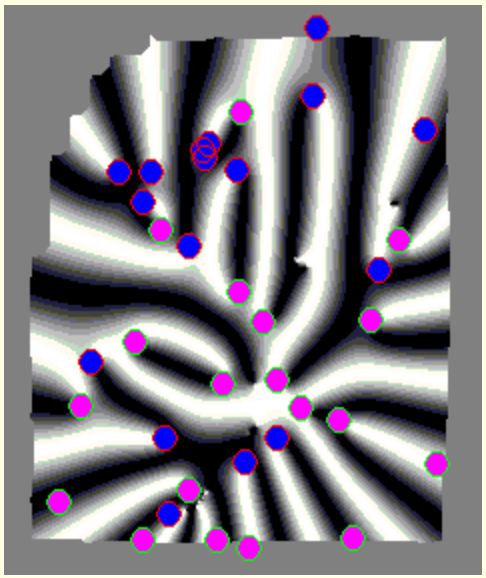

Spiral Phase

Continuous Phase

Othman and Ross, "On Mixing Fingerprints", TIFS 2013

Decomposition: Arch

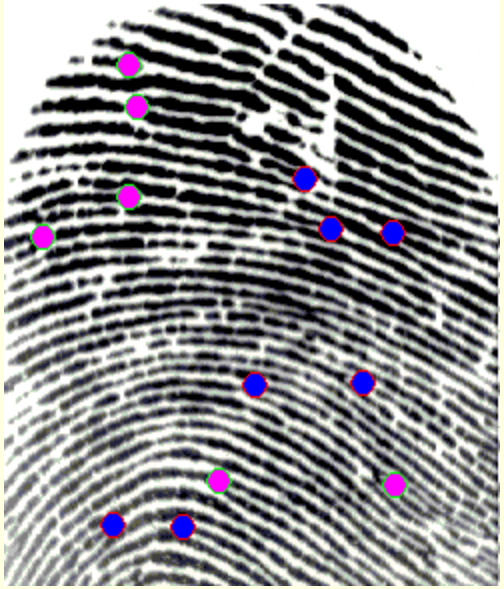

Original

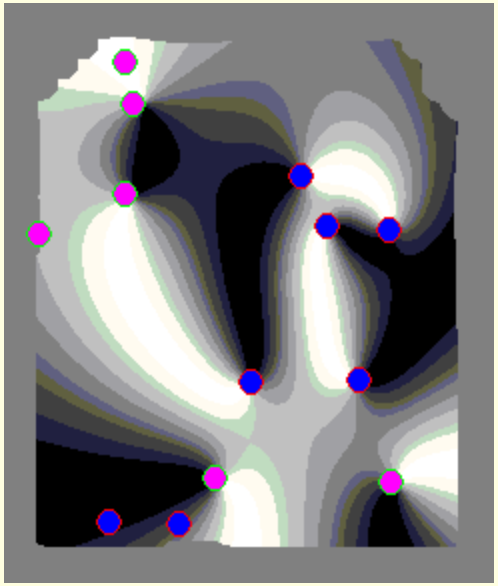

Spiral Phase

Continuous Phase

Othman and Ross, "On Mixing Fingerprints", TIFS 2013

Mixing Fingerprints

- Let F_1 and F_2 be two different fingerprint images from different fingers, and let $\boldsymbol{\Psi c_i(x, y)}$ and $\boldsymbol{\Psi s_i(x, y)}$ be the pre-aligned continuous and spiral phases, i = 1,2.

$$MF_1 = \cos[\Psi c_2(x, y) + \Psi s_1(x, y)]$$

$$MF_2 = \cos[\Psi c_1(x, y) + \Psi s_2(x, y)]$$

- The continuous phase of F_2 is combined with the spiral phase of F_1 which generates a new fused fingerprint image $\boldsymbol{MF_1}$

Othman and Ross, "On Mixing Fingerprints", TIFS 2013

Mixed Fingerprint Images

Othman and Ross, "On Mixing Fingerprints", TIFS 2013

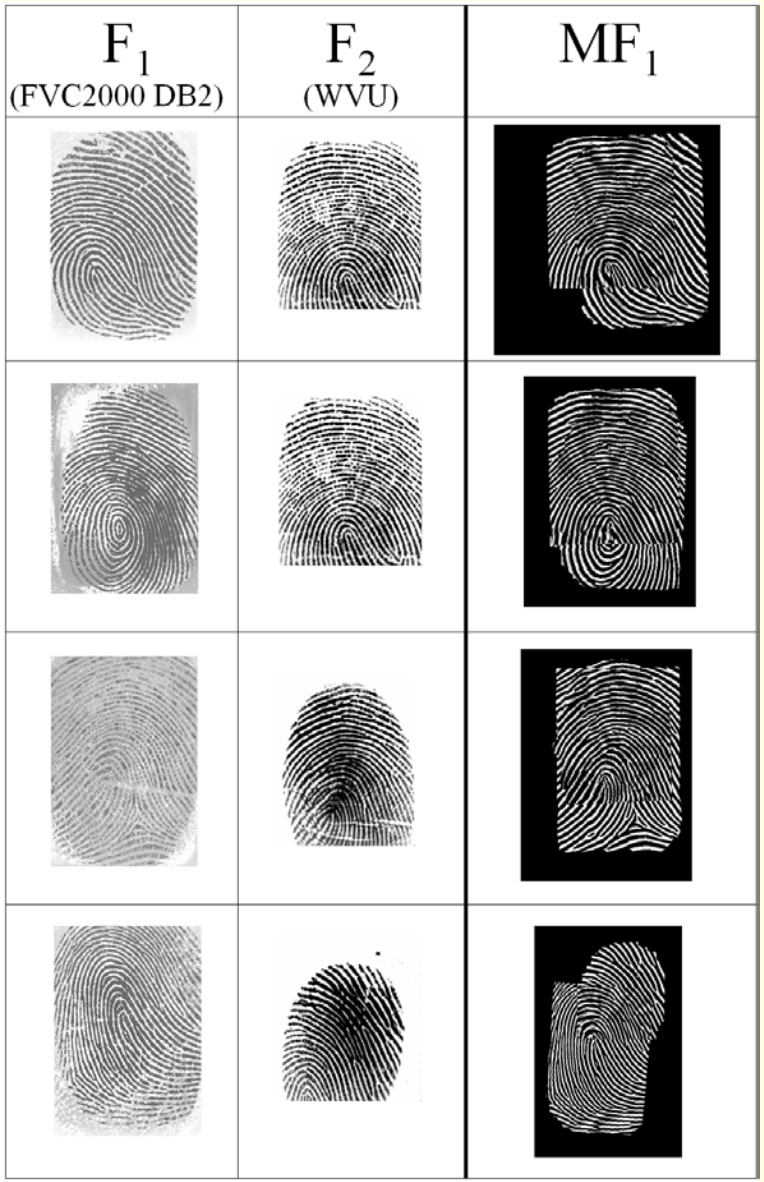

Mixing Fingerprints: Results

- Can the mixed fingerprint be used as a new biometric identity? (Yes)
- Are the original fingerprint and the mixed fingerprint correlated? (No)
- Does mixing result in cancelable templates? (Yes)
- If two different fingerprints are mixed with a common fingerprint, are the mixed fingerprints similar? (No)

Othman and Ross, "On Mixing Fingerprints", TIFS 2013

Privacy Enhancing Technology

- Preserving the privacy of a user's stored biometric data
  - Regulate cross-linking across applications
  - Regulate gleaning additional information from biometric data (e.g., medical condition)

Need to
- Define Privacy and Privacy Metrics
- Guarantee Privacy
- Develop Differential Privacy Schemes

Summary

- Perturbing soft biometric information in face images by perturbing/transforming face images
- Visual Cryptography for decomposing a face image and storing it in two separate servers
  - Individual servers cannot identify the face
- Mixing fingerprints by combining the spiral and continuous phase components of two fingerprints
  - Cancellable fingerprints
  - Joint identity/Group Authentication