Biometric Recognition

Marta Gomez-Barrero

Hochschule Darmstadt, CRISP, da/sec Security Group

Finse Winter School, May 2017

Outline

- Introduction
- Vulnerabilities of Biometric Systems
- Biometrics & Privacy

Introduction

Why biometric recognition?

- We need to identify ourselves in a daily basis
- Impossible to remember 100 different passwords

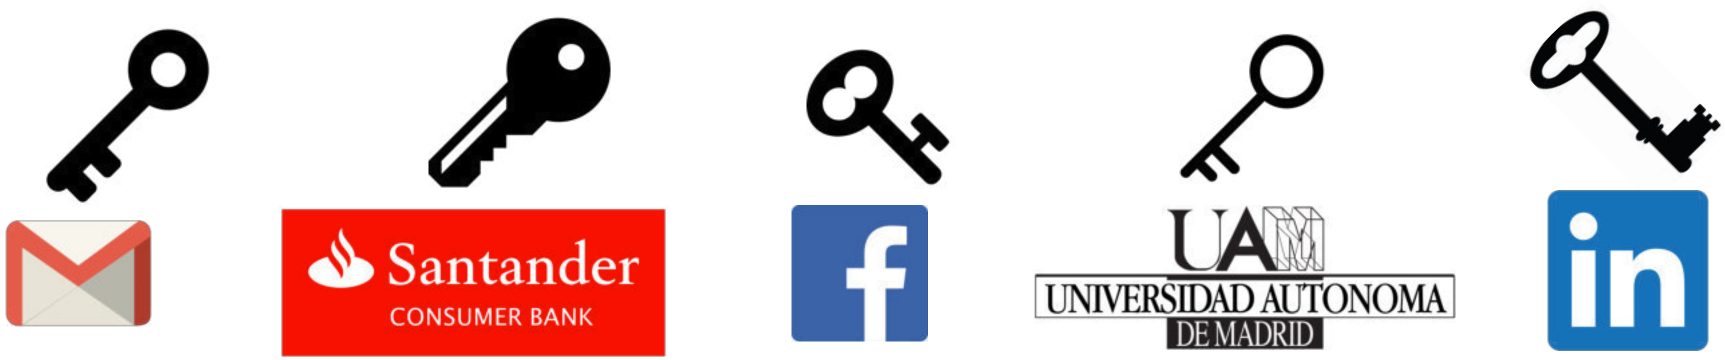

- Losing or forgetting our password / token is easy

Why not use our body features or behavioural patterns?

Biometric characteristics

- Classification:
  - Physiological
  - Behavioural

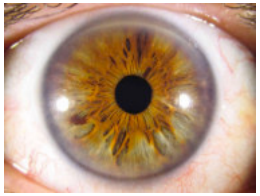

- Properties:
  - **Universality**: everybody should possess it
  - **Distinctiveness**: should have enough intervariability
  - **Permanence**: should not vary through time
  - **Collectability**: should be easy to acquire
  - **Performance**: should have good error rates
  - **Acceptability**: user should not be reluctant to use it
  - **Circumvention**: difficult to bypass

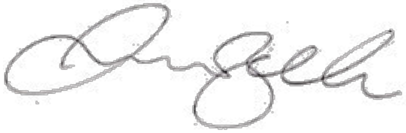

Advantages and disadvantages of biometrics

- No need to remember passwords or carry tokens
- Impersonation can be detected
- A single characteristic can be used in multiple applications, without security decrease

- Spoofing / Presentation Attacks (PA)
- Renewability
- Biometrics are no secrets
- Sensitive information

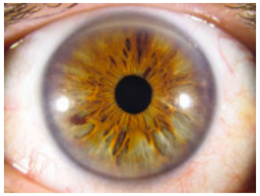

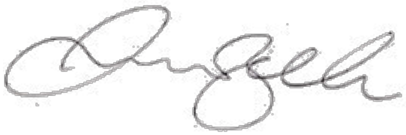

How does it work?

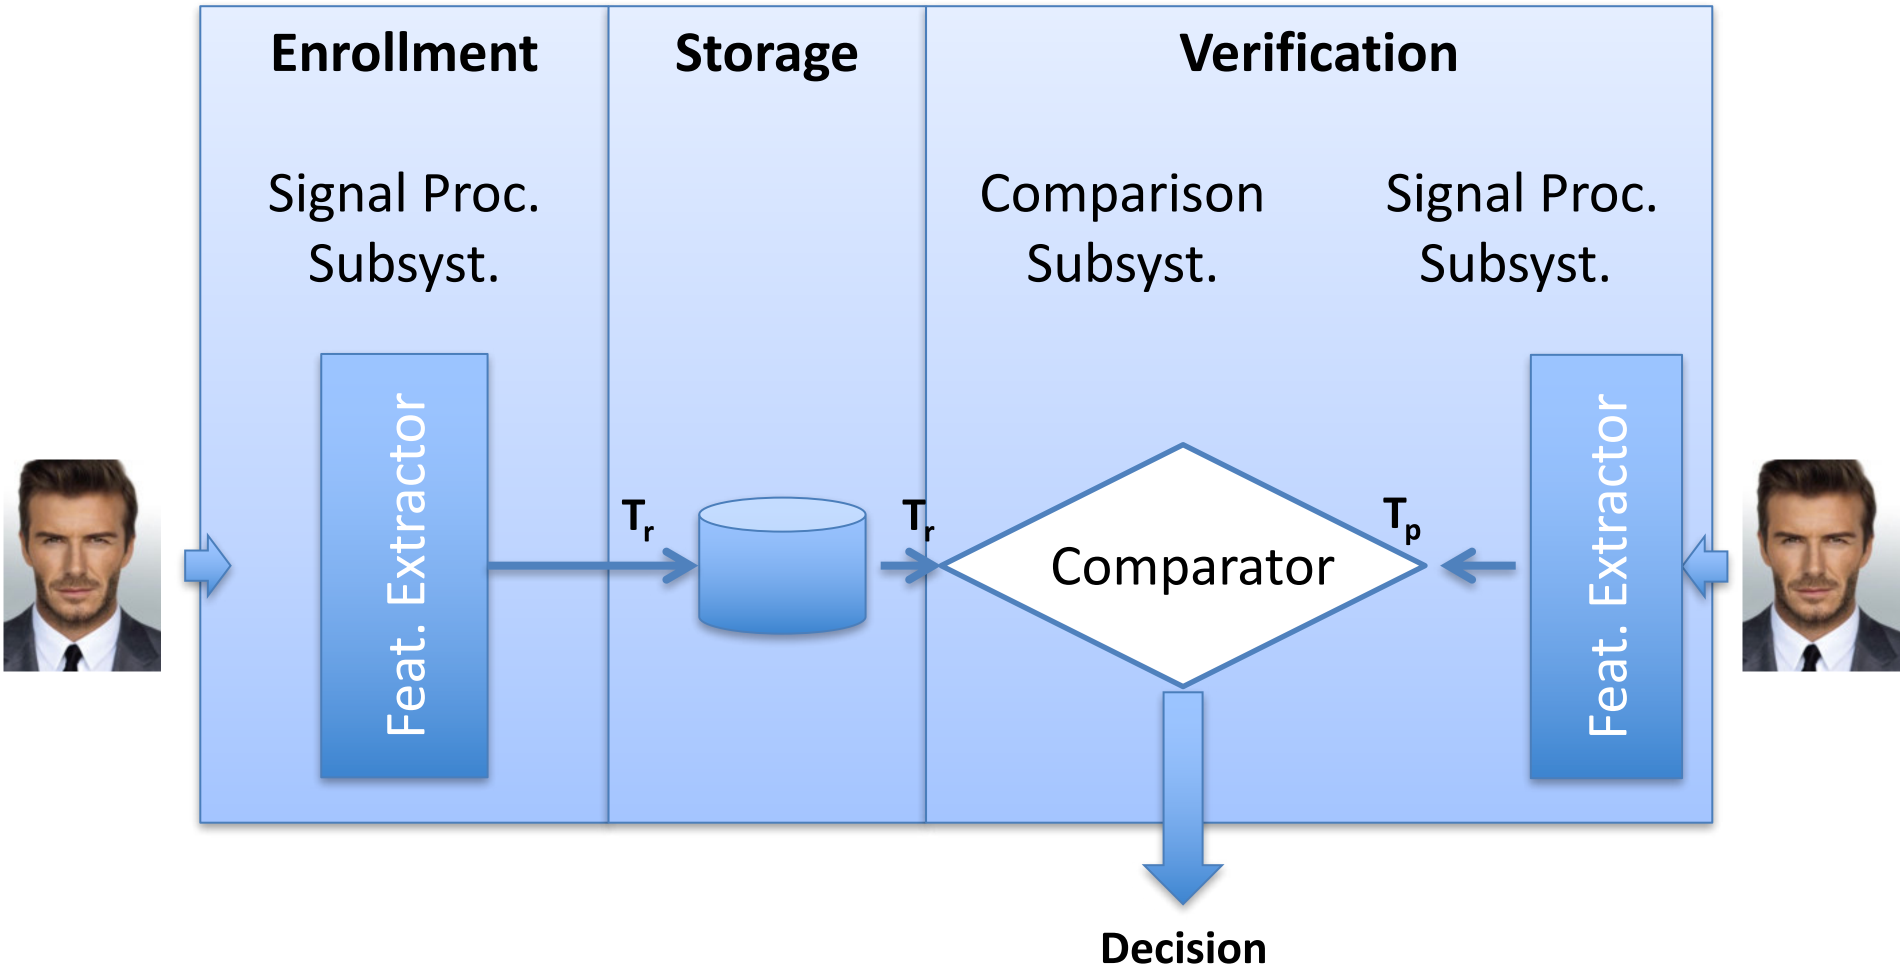

Example: iris recognition

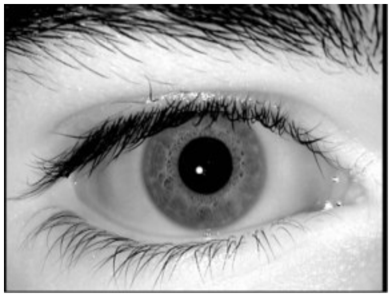

Sample

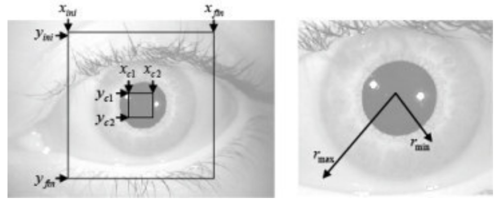

Segmentation

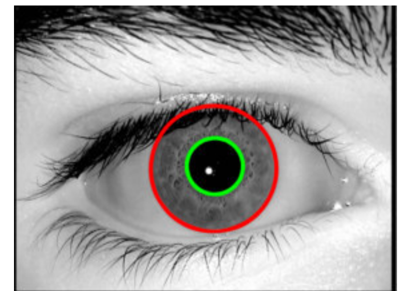

Normalization

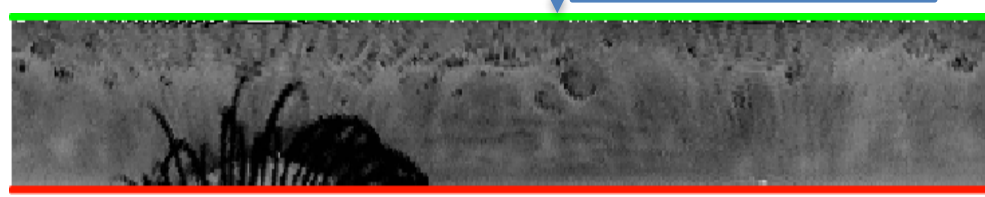

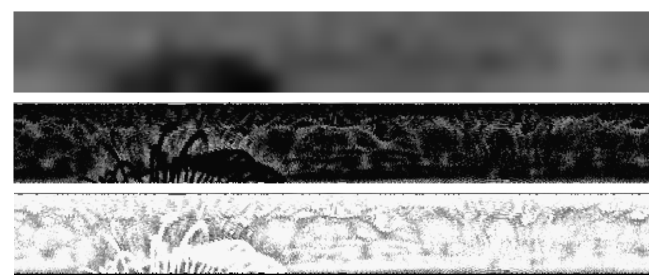

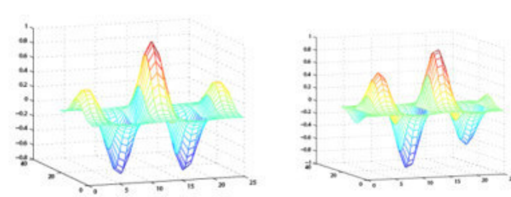

Feature Extraction

Template: T

Verification vs Identification

- Verification: I am Jon Doe (1:1)

T_r → **Yes / no**

- Identification: I am in the list (1:n)

T_1

…

T_n

→ **Identity *i* / not in the list**

Error rates

[ISO/IEC 2382-37 Harmonized Biometrics Vocabulary (HBV)]

- Two kinds of comparisons:

Mated Comparison

Non-Mated Comparison

- Two kinds of error rates:
  - **False Match Rate (FMR)** – proportion of falsely accepted non-mated comparison trials
  - **False Non-Match Rate (FNMR)** – proportion of falsely rejected mated comparison trials

Evaluating the accuracy

[ISO/IEC 19795 on Biometric performance testing and reporting]

- Plot mated and non-mated score distributions
- Establish a verification threshold: δ

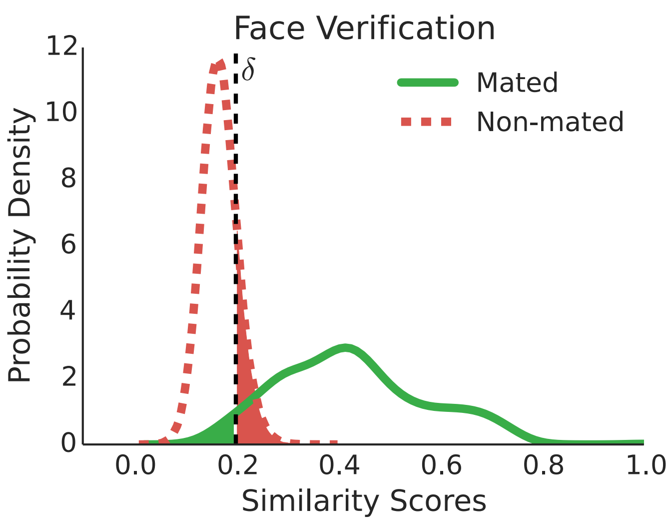

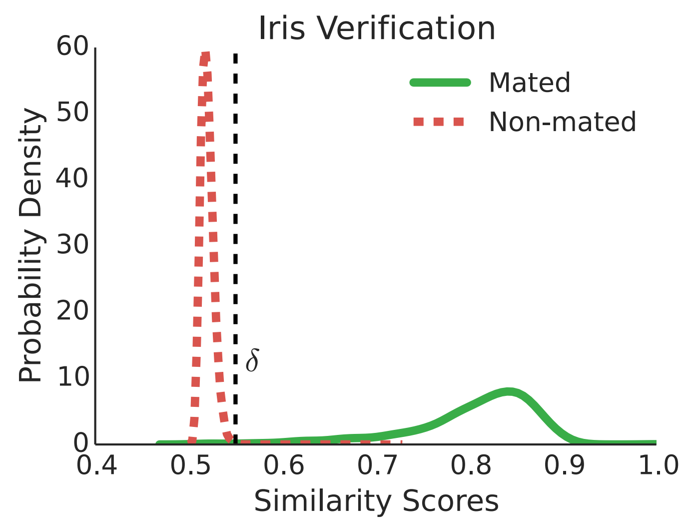

- δ determines the **FMR**
- … and the **FNMR**

Comparing systems

[ISO/IEC 19795 on Biometric performance testing and reporting]

- Compare all operating points with a **Detection Error Trade-off (DET)** curve
- The point at which FMR = FNMR is defined as **Equal Error Rate (EER)** - the lower, the better
- Report FNMR at fixed FMR – e.g., FMR = 0.1%

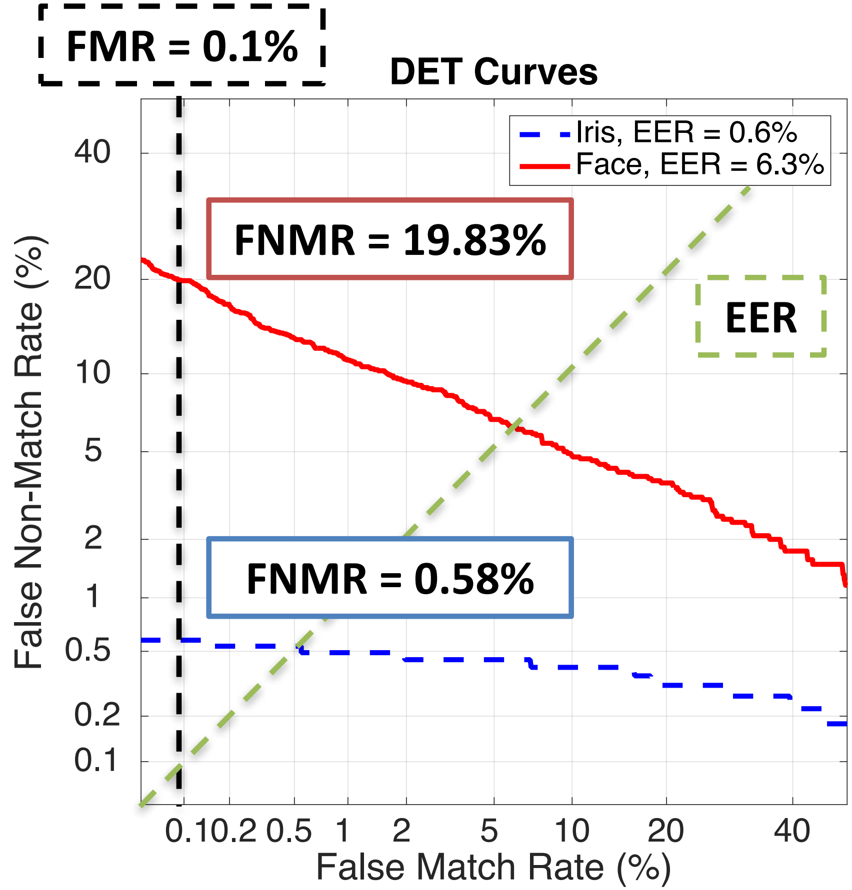

Multi-Biometric systems

[ISO/IEC TR 24722 on Multimodal and other multibiometric fusion]

- Advantages
  - Higher accuracy
  - Increased robustness to individual sensor or subsystem failures
  - Decreased number of cases where the system is not able to make a decision
  - Different levels of security
  - …

- Fusion levels:
  - Feature level
  - Score level
  - Decision level

Can be harder to achieve, but it’s preferred: reduced storage and higher security

Vulnerabilities of Biometric Systems

External Attacks

- Biometric systems are not free from external attacks.

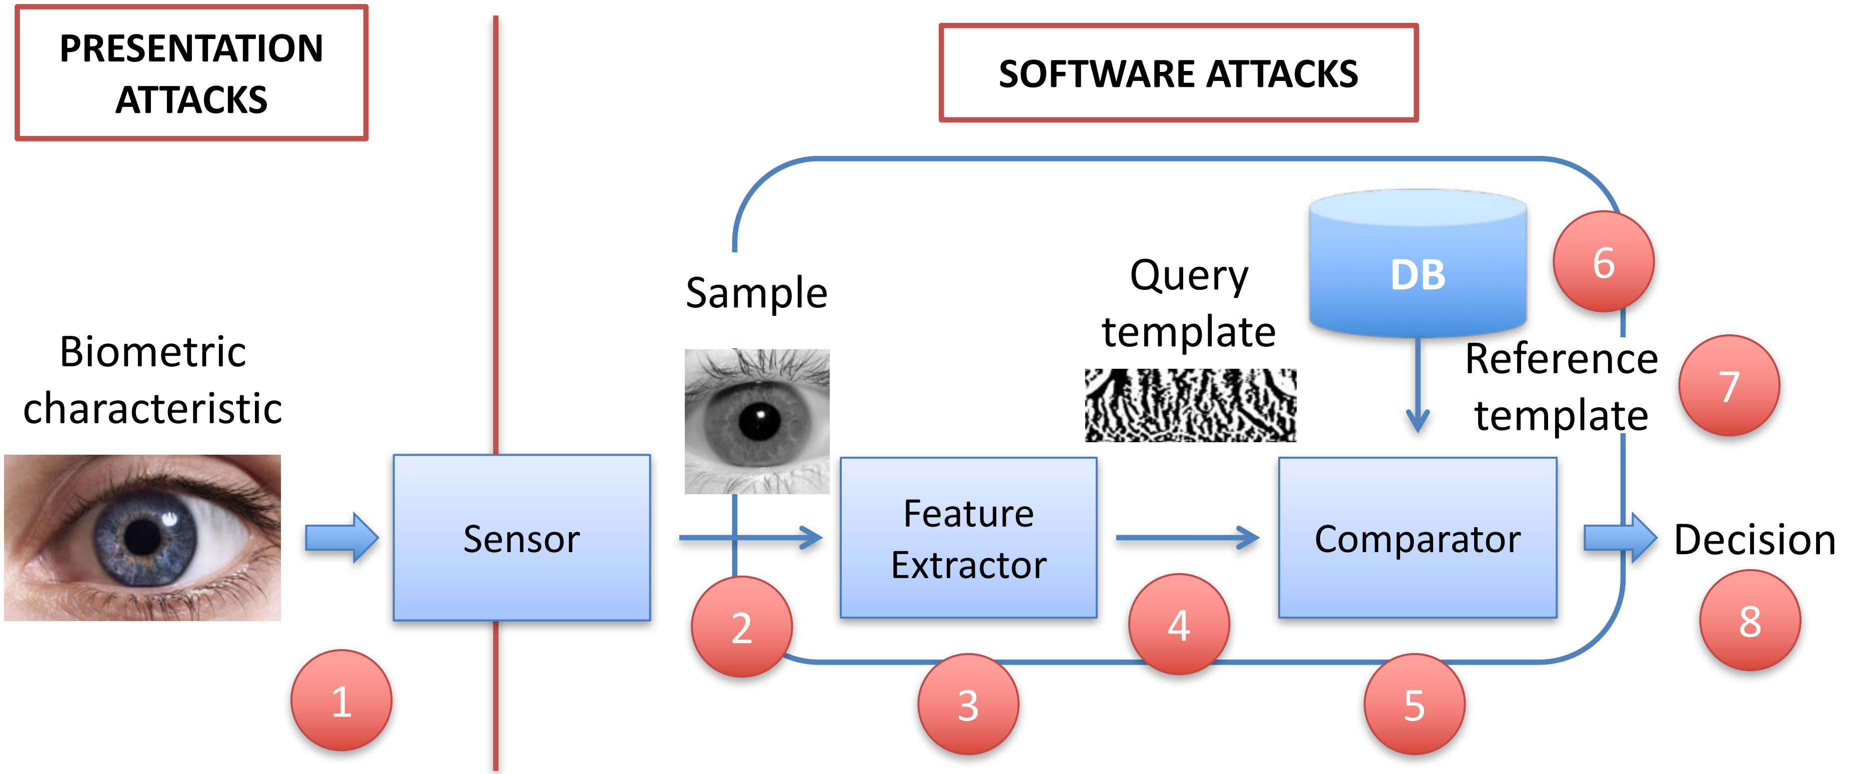

Vulnerability Analysis

- Projects
- Competitions
- Standards

Hill Climbing attacks

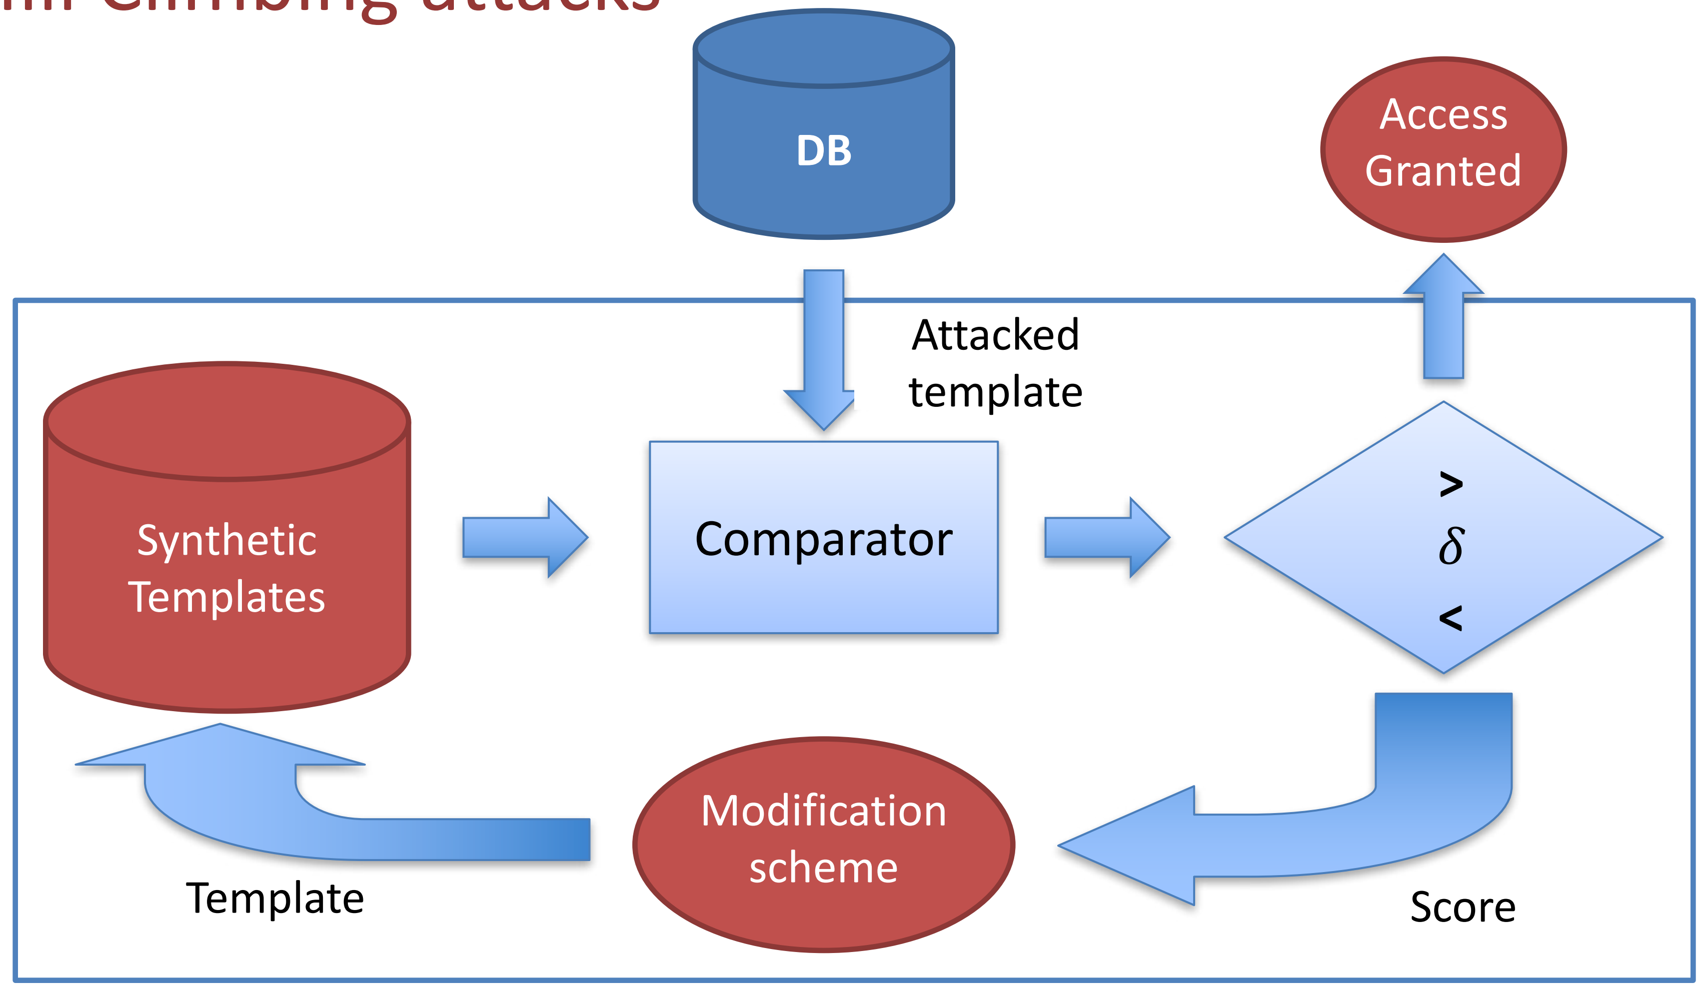

HC based on the Uphill Simplex algorithm

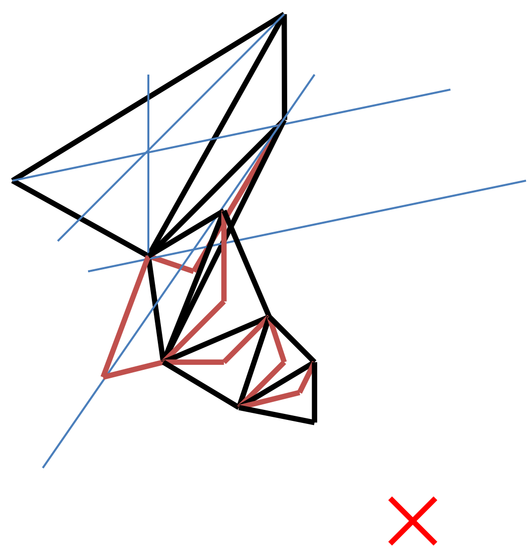

- New point:
  - Compute centroid: $\bar{\mathbf{y}} = \frac{1}{K+1} \sum_i \mathbf{y}_i$
  - Try reflection: $\mathbf{a} = (1+\alpha)\bar{\mathbf{y}} - \alpha \mathbf{y}_l$
  - Try expansion $\mathbf{b} = \gamma \mathbf{a} + (1-\gamma)\bar{\mathbf{y}}$

    or contraction: $\mathbf{b} = \beta \mathbf{y}_l + (1-\beta)\bar{\mathbf{y}}$

- Stopping criteria:
  - One of the points of the simplex is close enough => success
  - Maximum number of iterations allowed reached => failure

Vulnerabilities

Example 1: Face

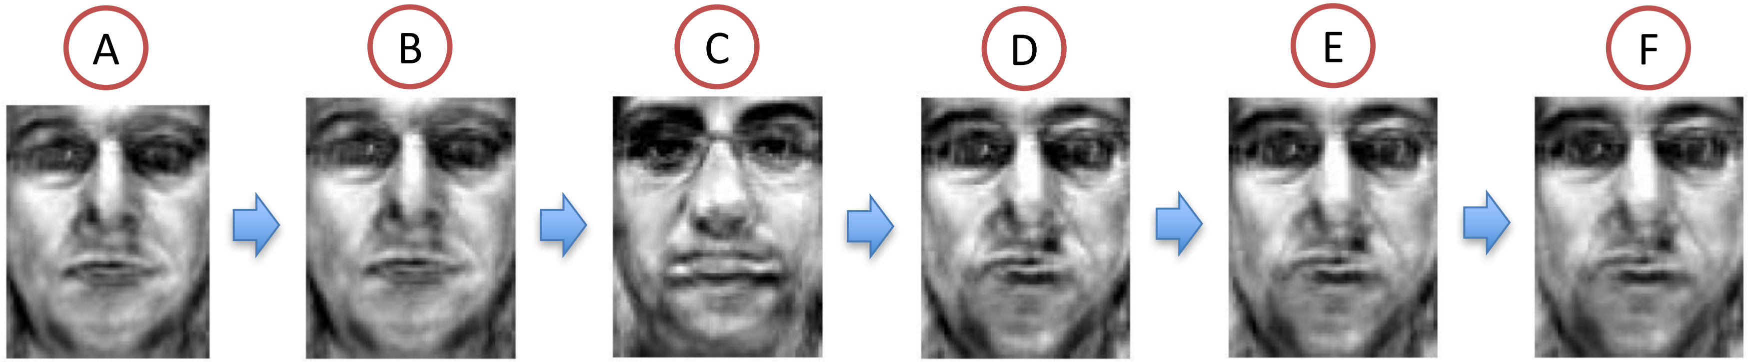

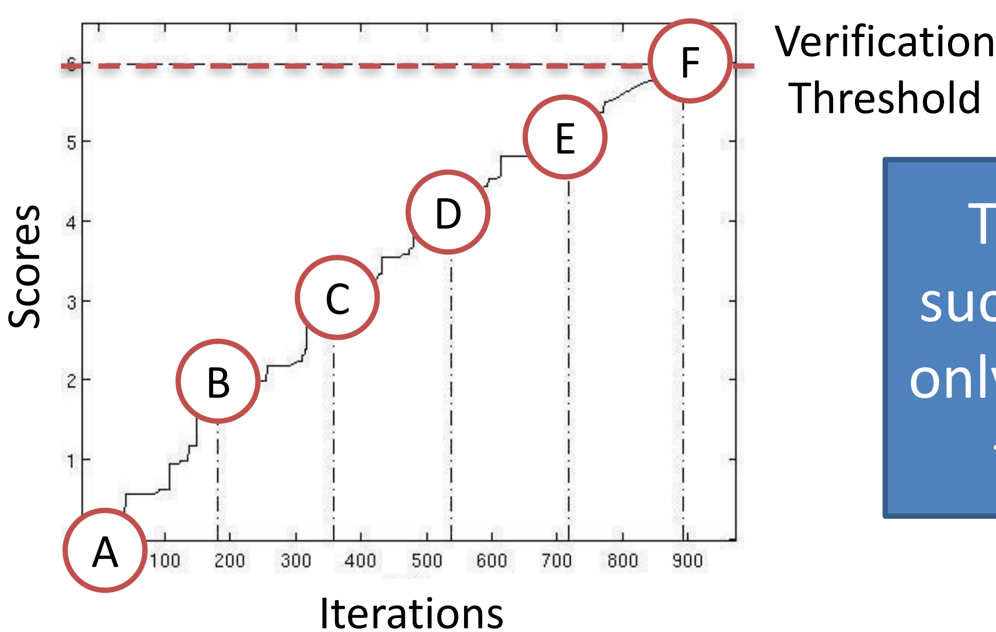

The attack was successful, and we only needed access to the scores

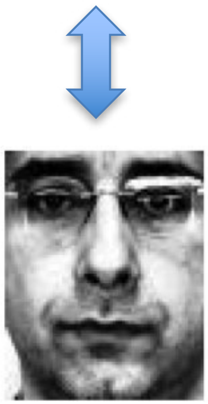

Target: Enrolled Sample

Example 2: Face and signature Success Rates (SR)

- We can evaluate how dangerous the attack is in terms of the success rate:

$$\mathrm{SR} = \frac{A_B}{A_T}$$

- At different operation points in terms of FMR

FMR (%)	Face System	Signature System
0.05%	100%	92.69%
0.01%	100%	87.84%

Hill Climbing attacks represent a real challenge to the security offered by biometric systems => Quantized Scores

HC based on genetic algorithms (I)

- We start with a random population of binary individuals
- At each iteration, we generate a new population according to four rules:
  - **Elite**: two individuals
  - **Selection**: stochastic universal sampling
  - **Crossover**: scattered crossover
  - **Mutation**: random changes
- Our fitness function is the similarity score
- Stopping criteria:
  - One of the individuals exceeds the verification threshold => success
  - Score increase in the last generations is very small => failure
  - Maximum number of iterations allowed reached => failure

HC based on genetic algorithms (II)

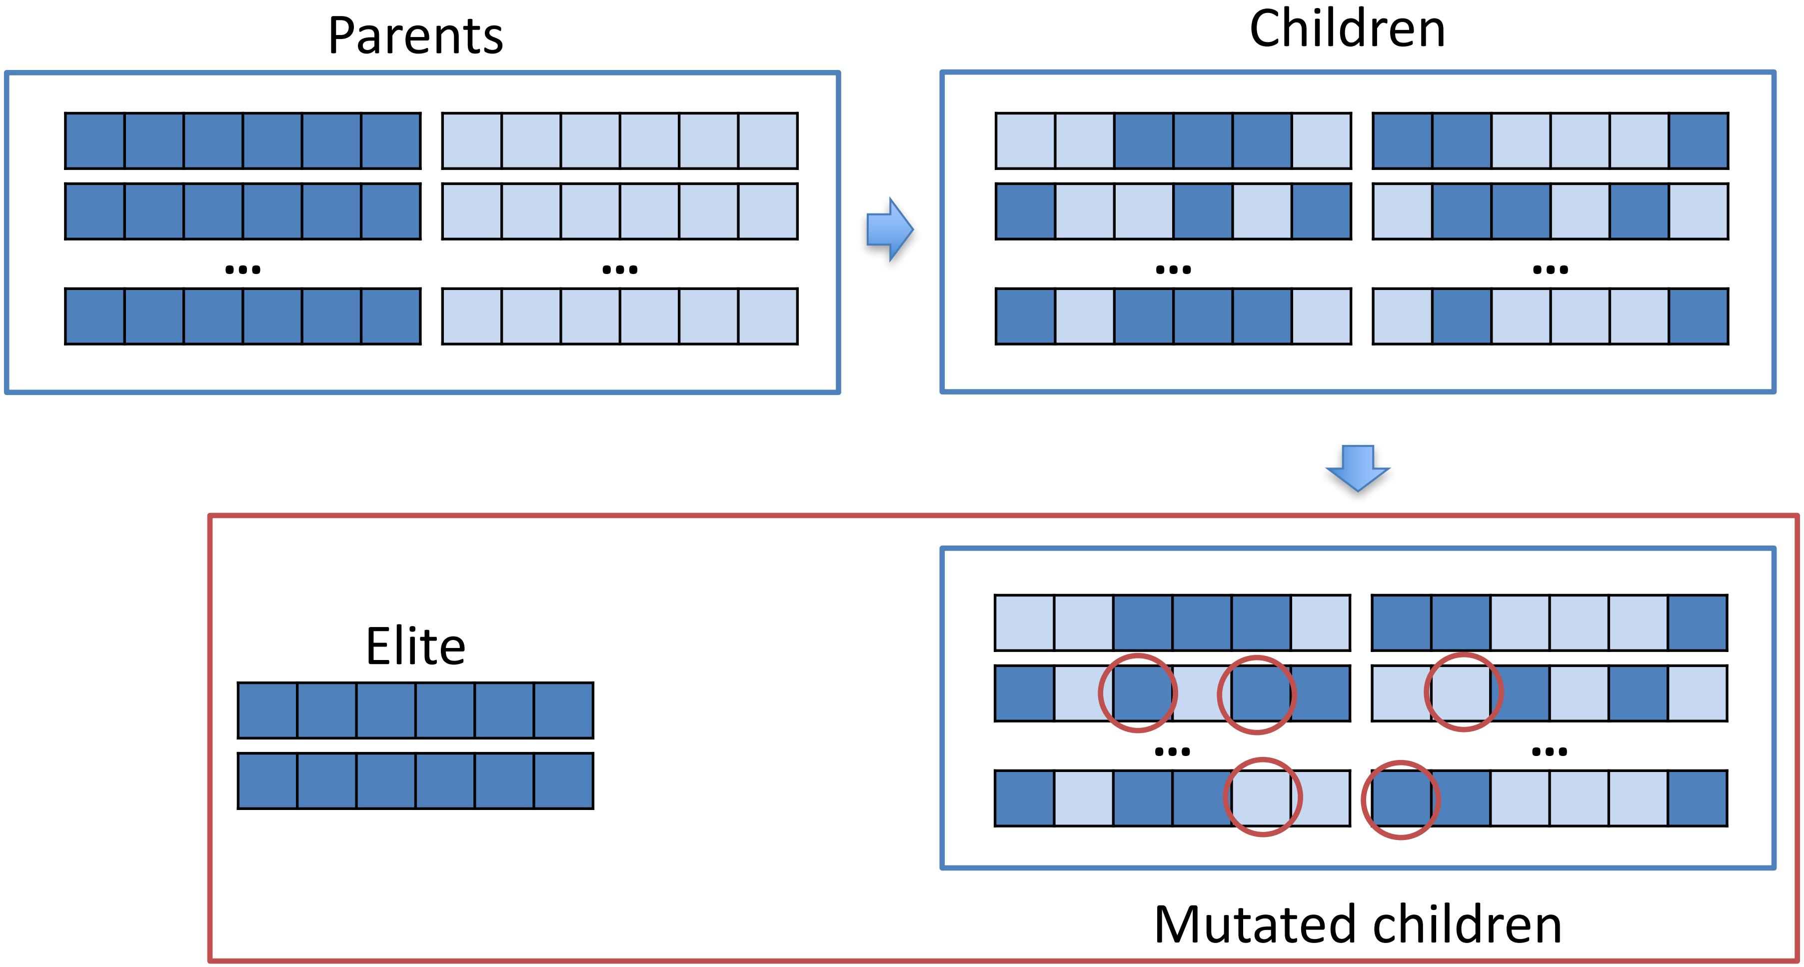

Example: Iris

FMR (%)	Iris System
0.05%	80.89%
0.01%	62.36%

Hill Climbing attacks represent a real challenge to the security offered by biometric systems => Quantized Scores

HC Attacks on multi-biometric systems

- Contrary to the belief that it is more difficult to attack a multi-biometric systems, we can combine these algorithms and succeed in our attack

FMR (%)	Face System	Iris System	Multi-Biometric
0.05%	100%	80.89%	100%
0.01%	100%	62.36%	100%

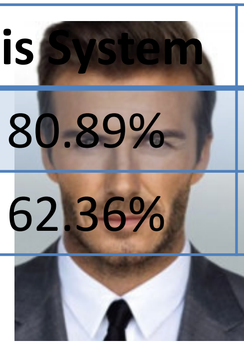

Sub-Algorithm 1: Uphill Simplex

Face Template

Score

Full Template

The multi-biometric system is as vulnerable as the most vulnerable characteristic

Biometrics & Privacy

Biometrics: sensitive data

- Wide deployment of biometrics:
  - Large scale national and international projects
  - Banking apps, ATMs
  - Smartphone unlocking
- Biometrics are classified as sensitive data

  [EU 2016/679 Data Protection Regulation]
  [EU 2016/680 Data Protection Directive]

- And we cannot prevent databases leakage

Inverse biometrics attacks

- It was a common belief that the stored templates revealed no information about the biometric characteristics:

- However, biometric samples can be recovered from the stored unprotected templates

Inverse biometrics attacks: Hill-Climbing

- Based on the HC algorithms presented before, we can reconstruct biometric samples:

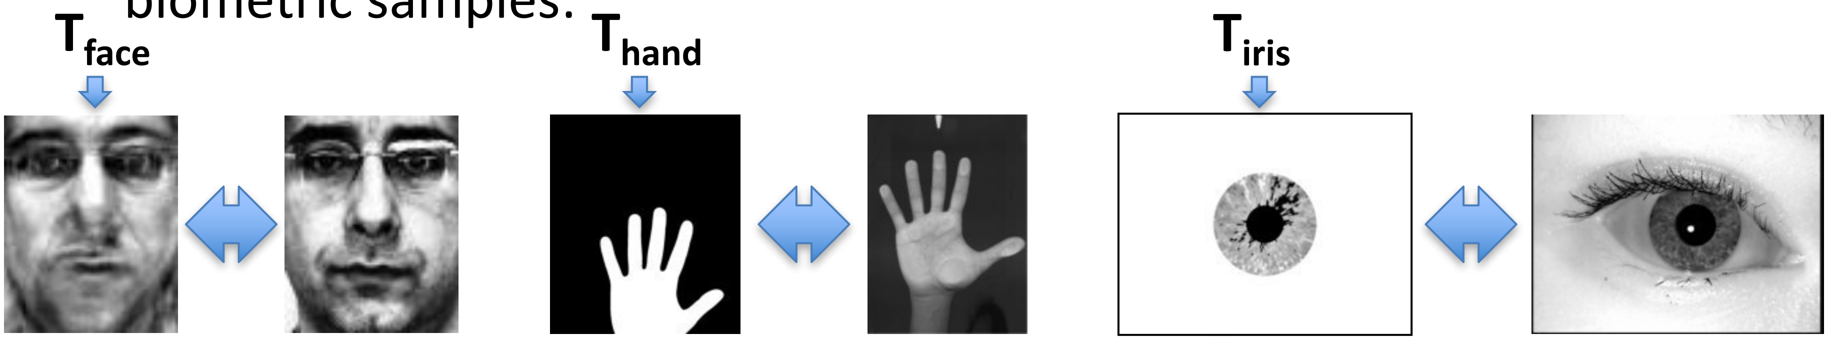

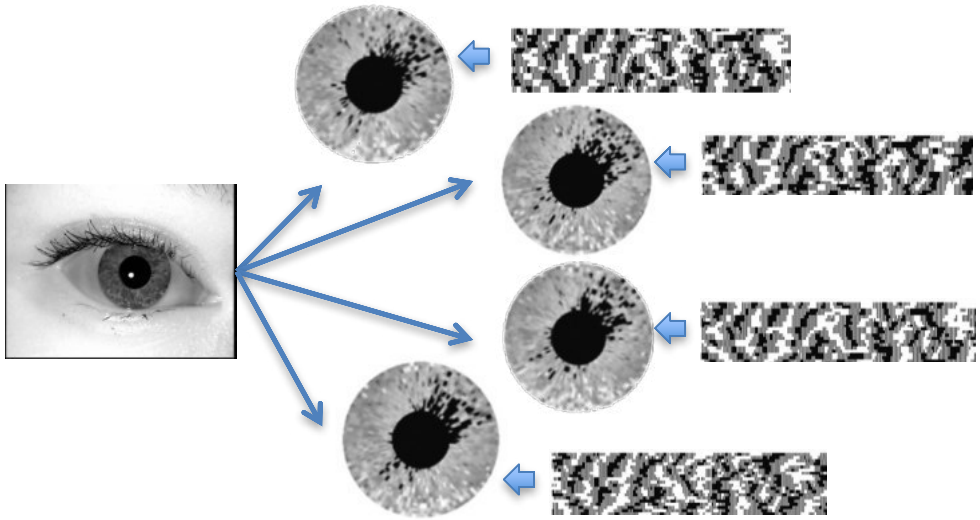

[M. Gomez-Barrero *et al.*, *Int. Conf. on Biometrics*, 2012]
[M. Gomez-Barrero *et al.*, *Information Sciences*, 2014]
[J. Galbally, *et al.*, *Computer Vision & Image Understanding*, 2013]

Inv

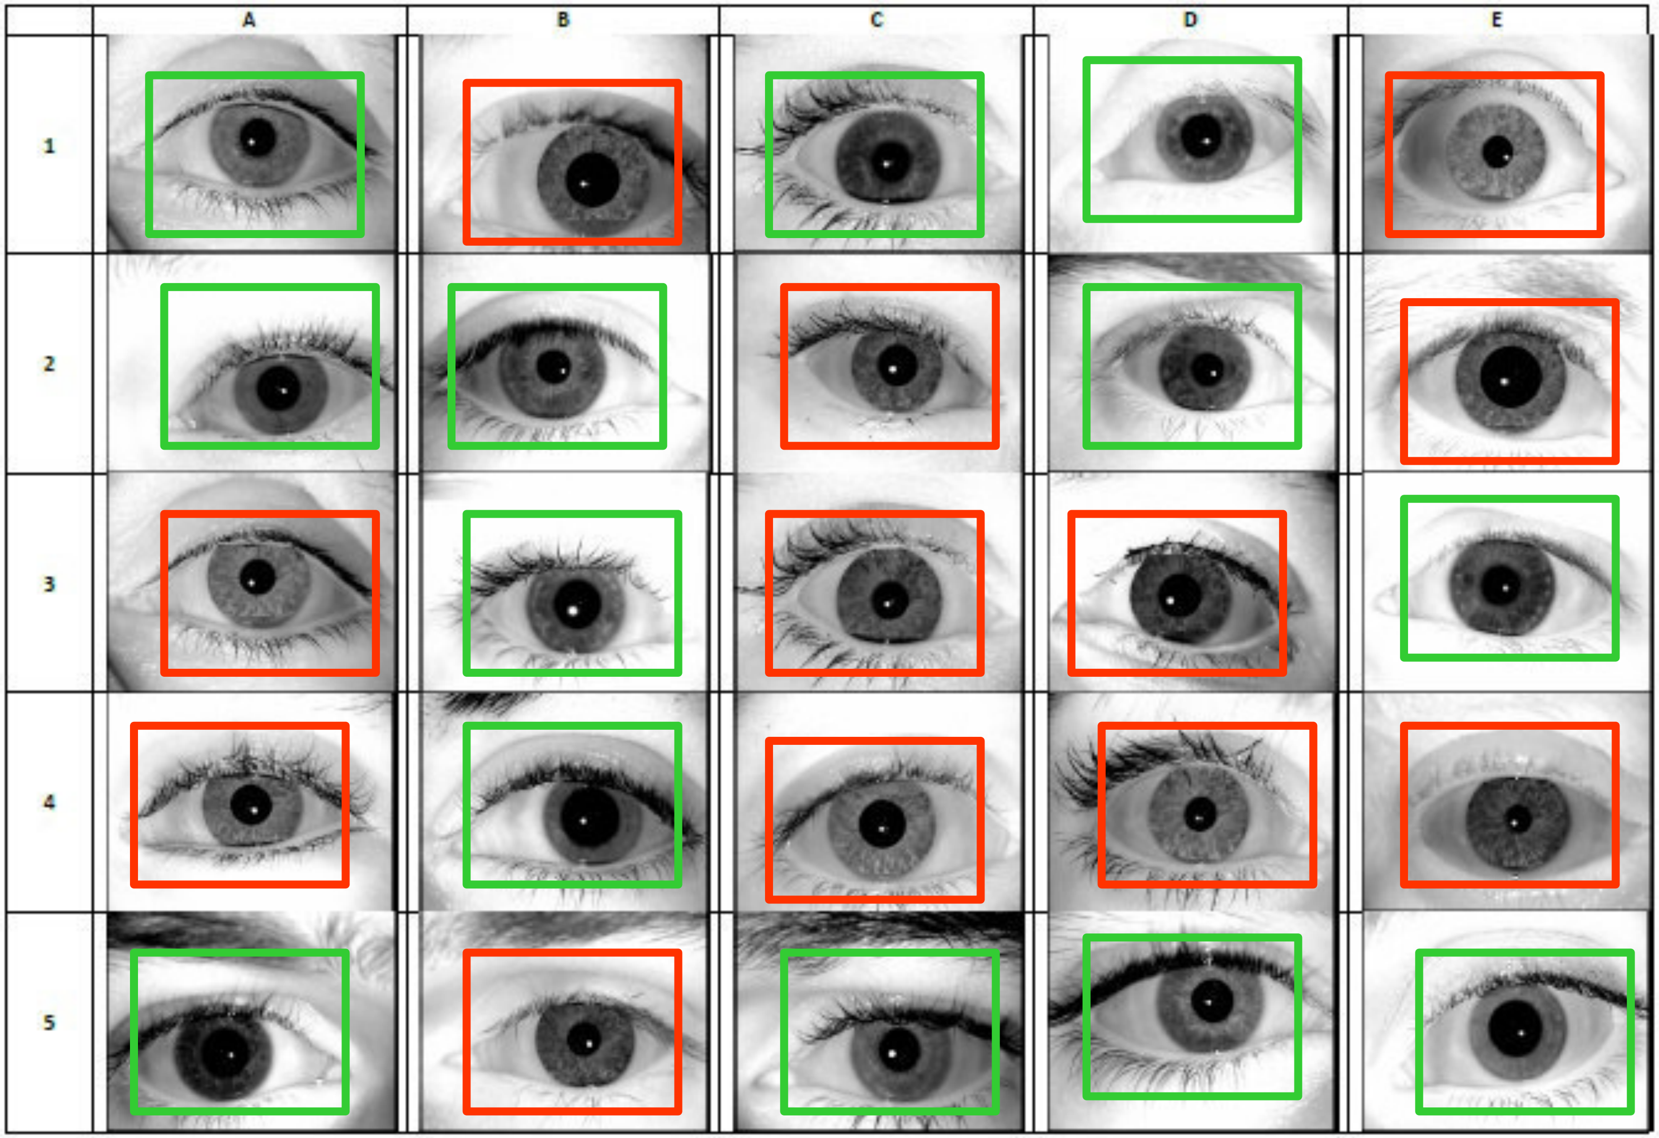

Inverse

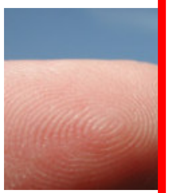

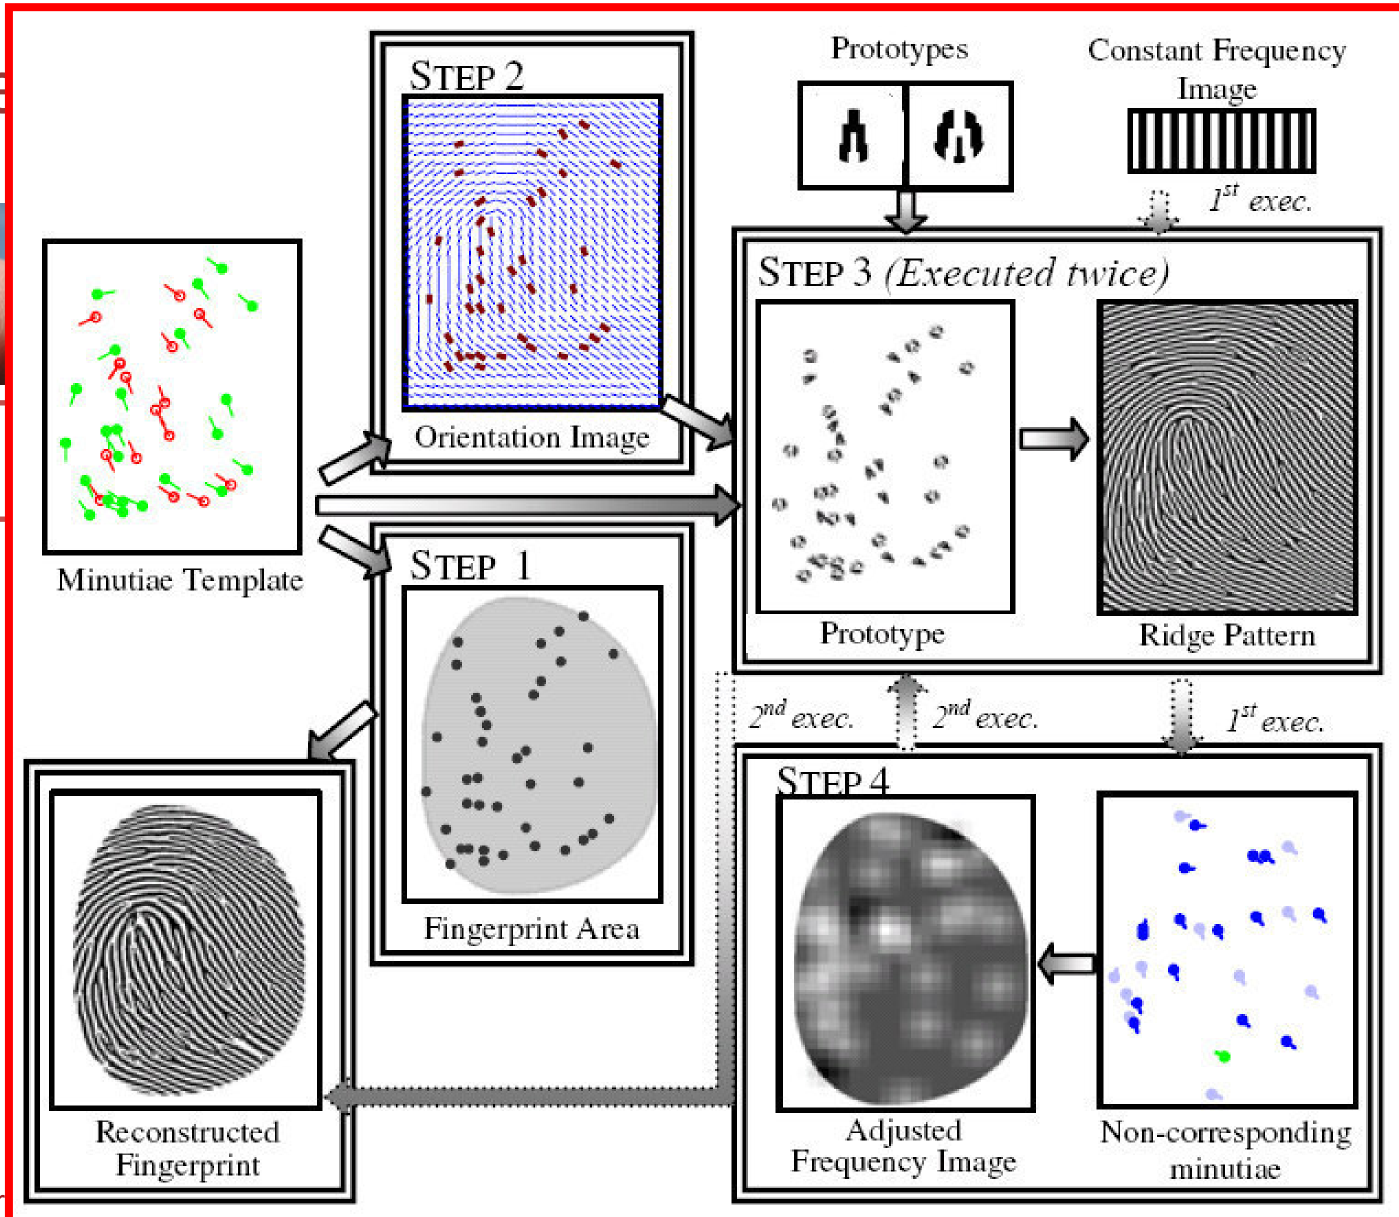

[Galbally *et*
2009]

2007]

Inverse biometrics attacks: Success Rates

FMR (%)	Iris	Fingerprint (indirect)	Fingerprint (PA)
0.05%	85.1%	98%	78%
0.01%	83.6%	92%	68%

Over 85% of the attacks are successful => Real challenge!

Lower success chances, but more difficult to detect

Templates need to be protected, so that we cannot recover the biometric sample

In addition, Presentation Attacks need to be detected

Cross-matching attacks

- We can enroll with a single characteristic in different applications

Templates need to be protected, so that no one can find out on which applications we are enrolled

Summary

- Do the stored templates reveal any information about the original biometric samples? **IRREVERSIBILITY**

- Are my enrolled templates in different recognition systems somehow related to each other? **UNLINKABILITY**

- What if someone steals a template extracted from my face? Has it been permanently compromised? **RENEWABILITY**

[ISO/IEC IS 24745 on Biometric Information Protection]

Marta Gomez-Barrero
(marta.gomez-barrero@h-da.de)

From inverse biometrics attack to PA

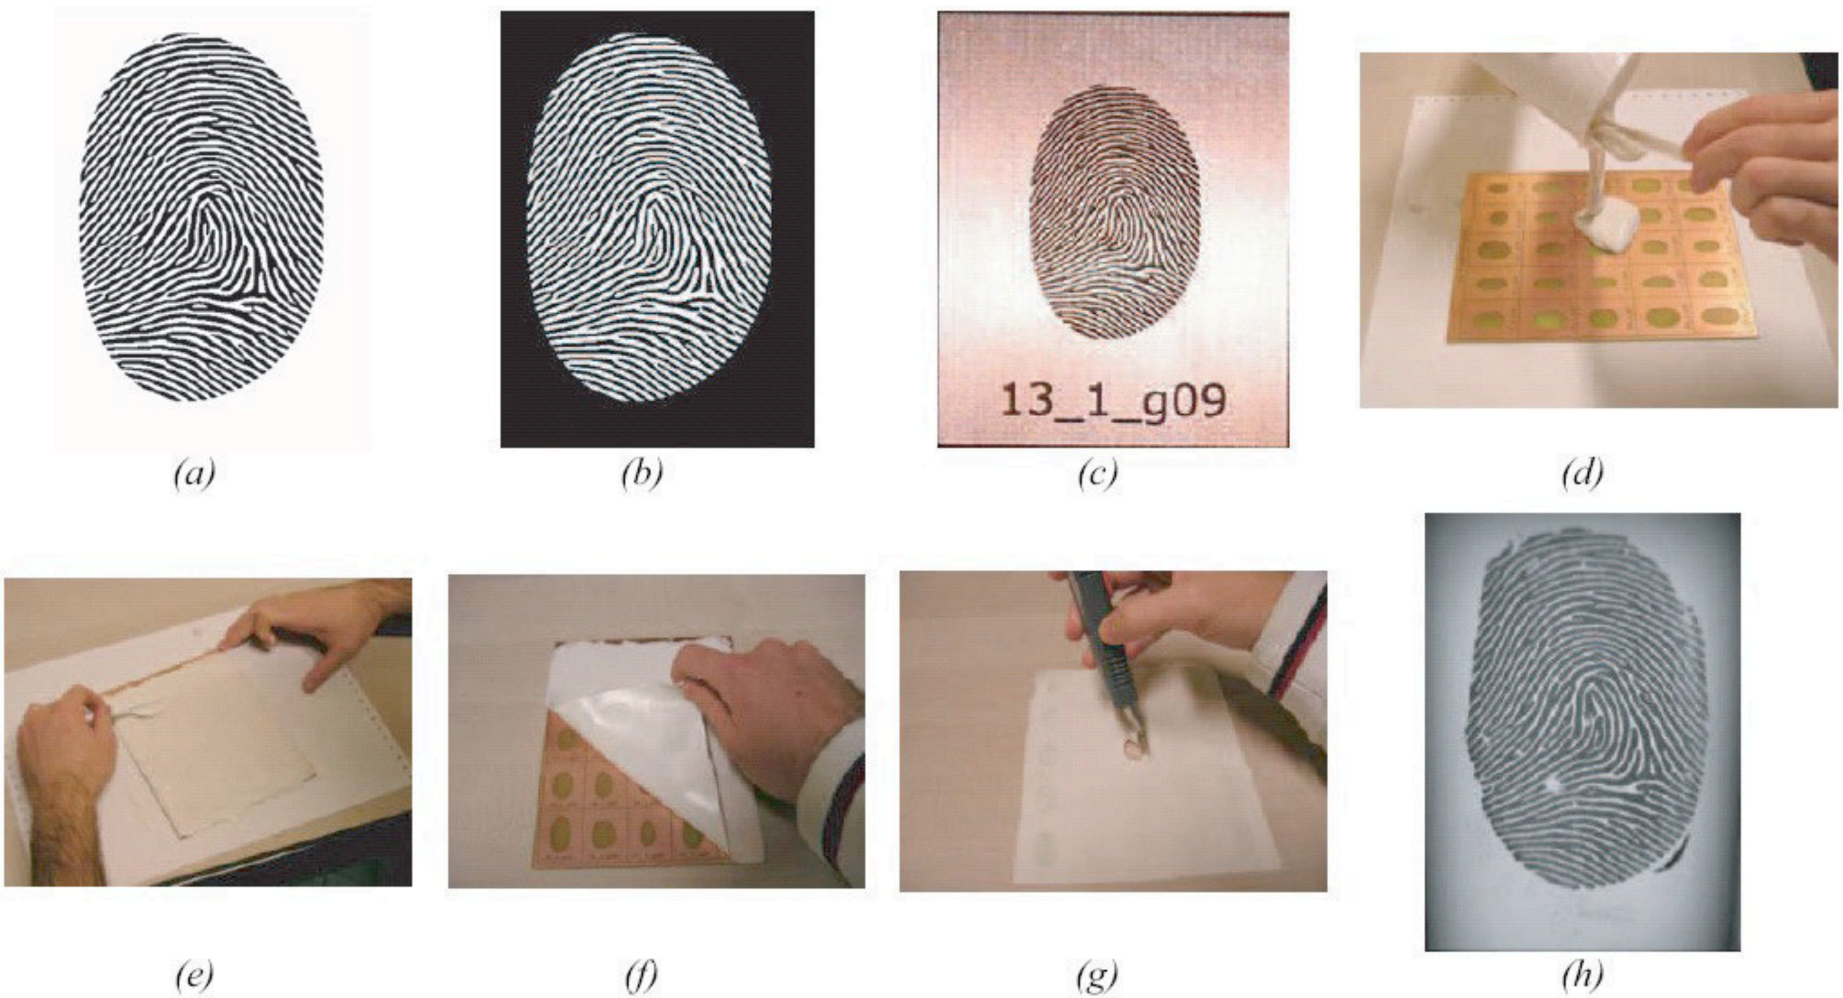

Biometric Template Protection

Marta Gomez-Barrero

Hochschule Darmstadt, CRISP, da/sec Security Group

Finse Winter School, May 2017

Outline

- Introduction
- Security and Privacy Evaluation
- Cancelable Biometrics Based on Bloom Filters
- BTP Based on Homomorphic Encryption
- Summary

Introduction

Protecting the subject's privacy

[ISO/IEC IS 24745 on Biometric Information Protection]

- Requirements of Biometric Template Protection:

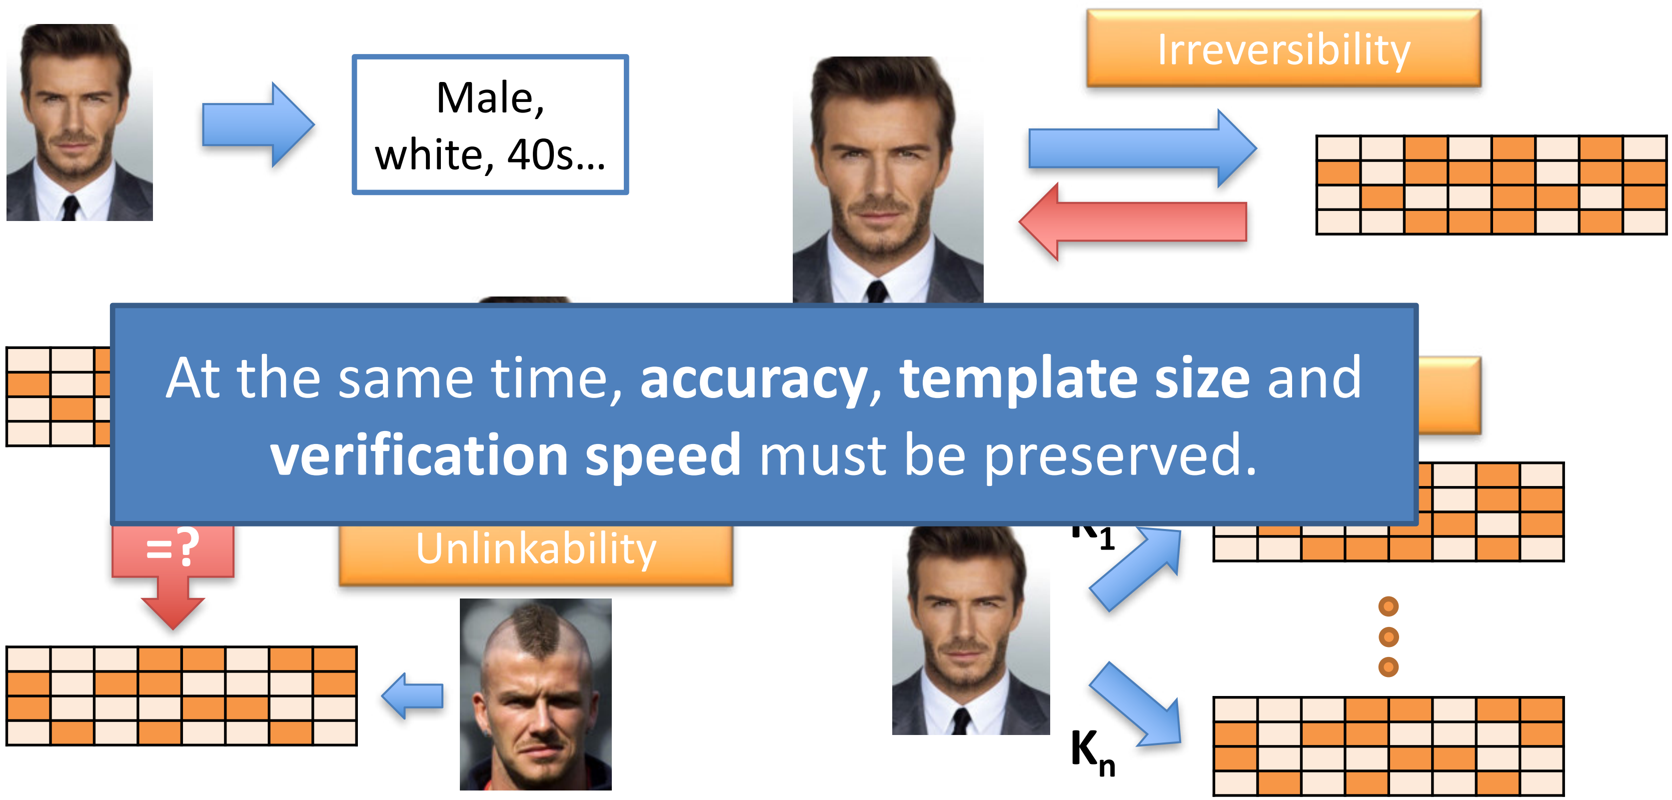

At the same time, **accuracy**, **template size** and **verification speed** must be preserved.

Biometric Template Protection Architecture

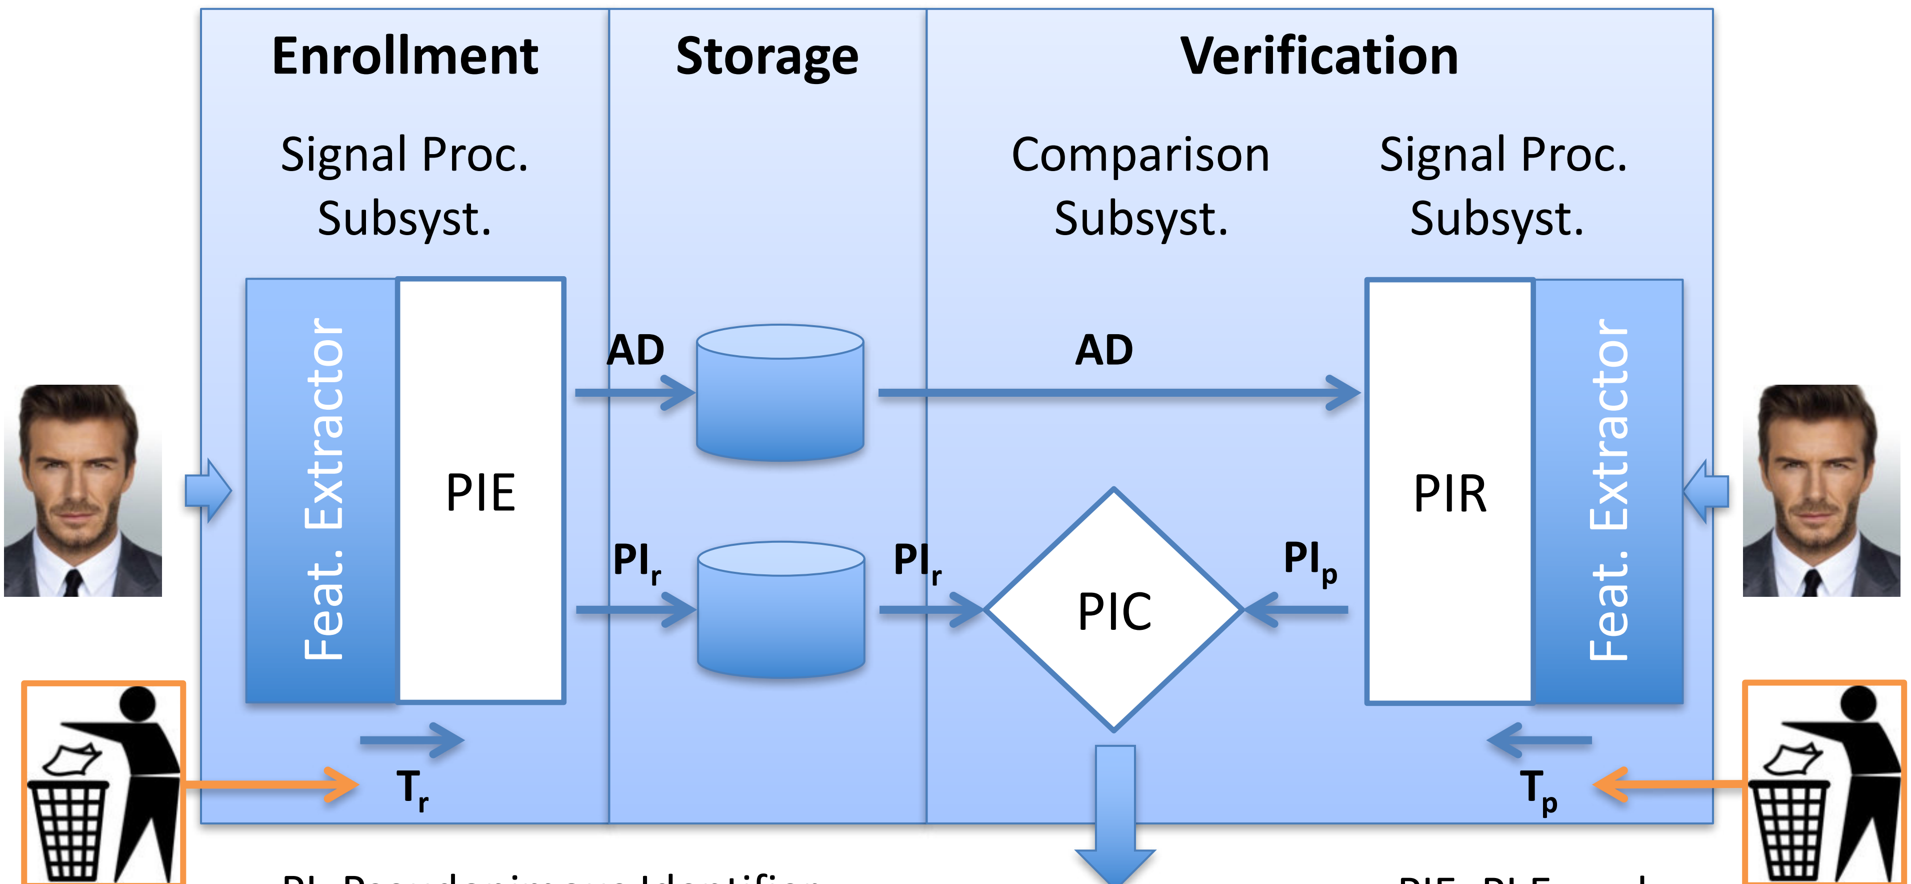

PI: Pseudonimous Identifier
AD: Auxiliary Data

PIE: PI Encoder
PIR: PI Recorder
PIC: PI Comparator

BTP Approaches

Cancelable Biometrics

- **Accuracy** drops
- Permanent **irreversibility**
- **Unlinkability** not analysed
- **Computational Complexity** Preserved

Template Protection based on Bloom filters

Cryptobiometrics

- **Accuracy** drops
- Attacks on AD (**irreversibility** compromised)
- **Unlinkability** not analysed
- **Computational Complexity** Preserved

[Campisi, *Springer* 2013]

Biometrics in the Encrypted Domain

- **Accuracy** preserved
- Permanent **irreversibility**
- **Unlinkability** granted
- **Computational Complexity** increased

Template Protection based on Homomorphic Encryption

Multi-Biometrics and BTP

- Multi-Biometrics:
  - Higher accuracy
  - Different levels of security
  - Three fusion levels: feature, score, decision [ISO/IEC TR 24722]

- Multi-Biometric Template Protection [Rathgeb and Busch, *InTech*, 2012]:
  - Alignment issues
  - Different BTP approaches for different characteristics

Security and Privacy Evaluation

Evaluation

Reproducible Research

- Public Baseline Systems
- Public DBs
- Knowledge Attacker
- Evaluation Protocol

ISO Requirements Evaluation

- Analysis 1: Accuracy
- Analysis 2: Irreversibility
- Analysis 3: Unlinkability
- Analysis 4a: Robustness to Cross-Matching Attacks
- Analysis 4b: Computational Load Increase

Cross-Matching Attacks

- We can enroll with a single characteristic in different applications

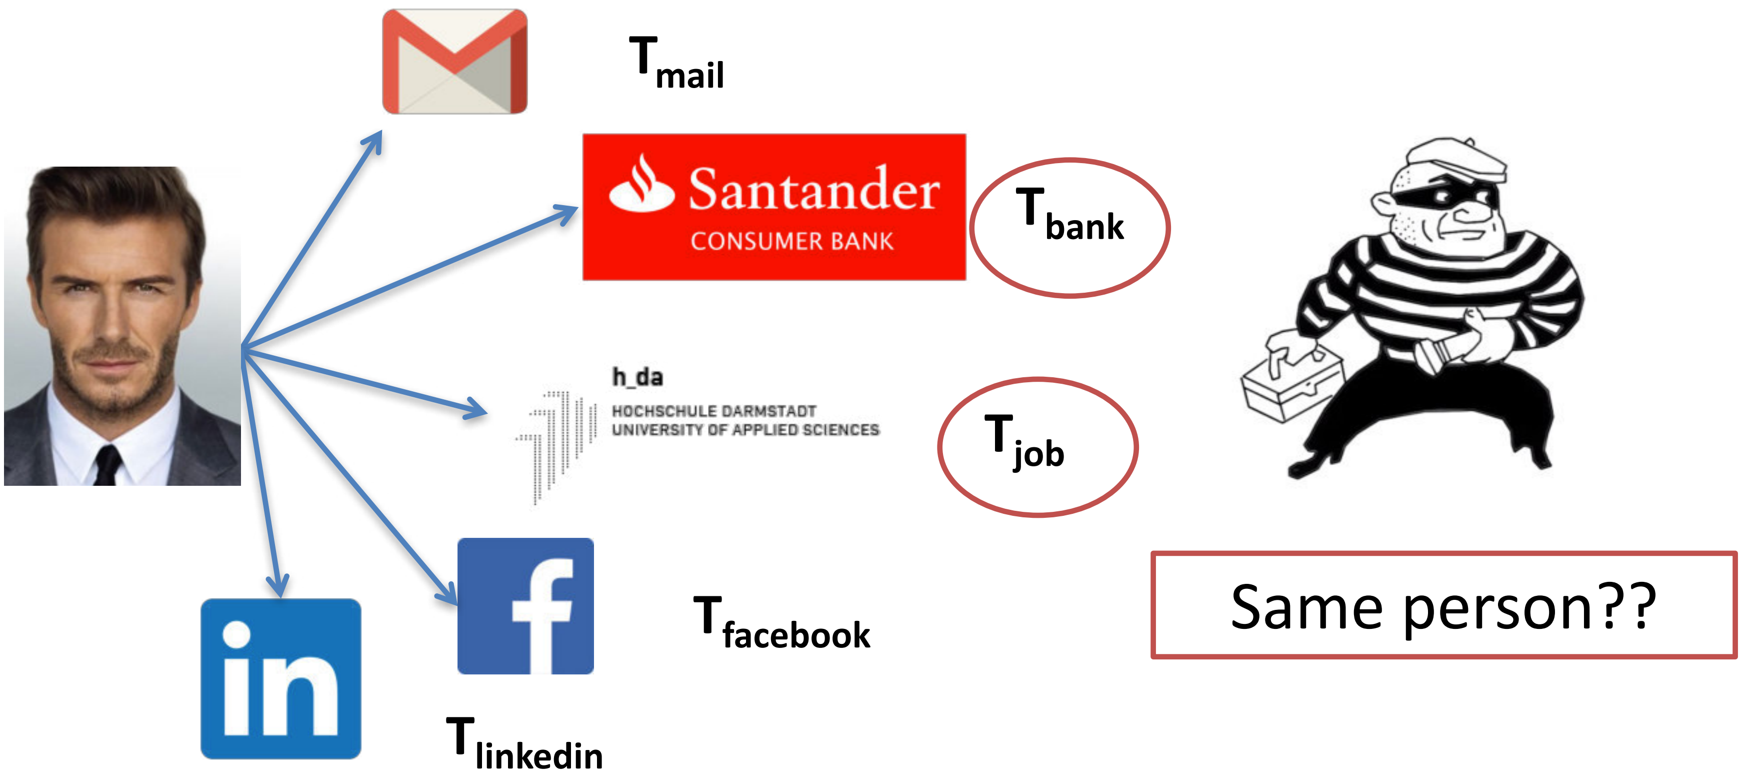

Cross-Matching Attacks: How to?

$\mathbf{T}_{\mathbf{job}}$ $\mathbf{T}_{\mathbf{bank}}$

$$s = DS\,(\mathbf{T}_{\mathbf{job}}, \mathbf{T}_{\mathbf{bank}})$$

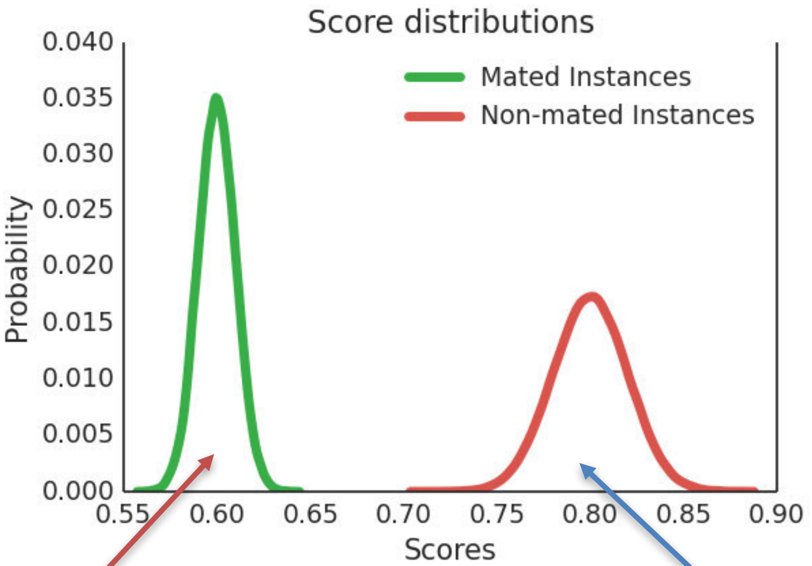

s here ➔ success!! ☺

s here ➔ try again!! ☹

s can be the dissimilarity score of the system or any other dissimilarity score, such as values extracted from partial decoding in fuzzy schemes

Unlinkability Analysis: Current Status (I)

- Advantage of the attacker over a random guessing in the indistinguishability game
  - Problem 1: assumes uniformity of data – not valid in biometrics
  - Problem 2: only analysed for fuzzy schemes – not straightforward to apply to cancelable biometrics, since calculations rely on ECC properties

[Simoens09] K. Simoens, P. Tuyls, B. Preneel, "Privacy Weaknesses in Biometric Sketches", *IEEE Symp. On Security and Privacy*, 2009.

[Buhan09] I. Buhan, J. Breebaart, M. Guajardo *et al.*, "A Quantitative Analysis of indistinguishability for a continuous Domain Biometric Cryptosystem", *Int. Workshop on Data Privacy and Management*, 2009.

[Buhan10] I. Buhan, E. Kelkboom, J. Guajardo, "Efficient Strategies for Playing the Indistinguishability Game for Fuzzy Sketches", *IEEE Workshop on Information Forensics and Security*, 2010.

Unlinkability Analysis: Current Status (II)

- Plot a DET curve of genuine and impostor scores, comparing templates enrolled in different system

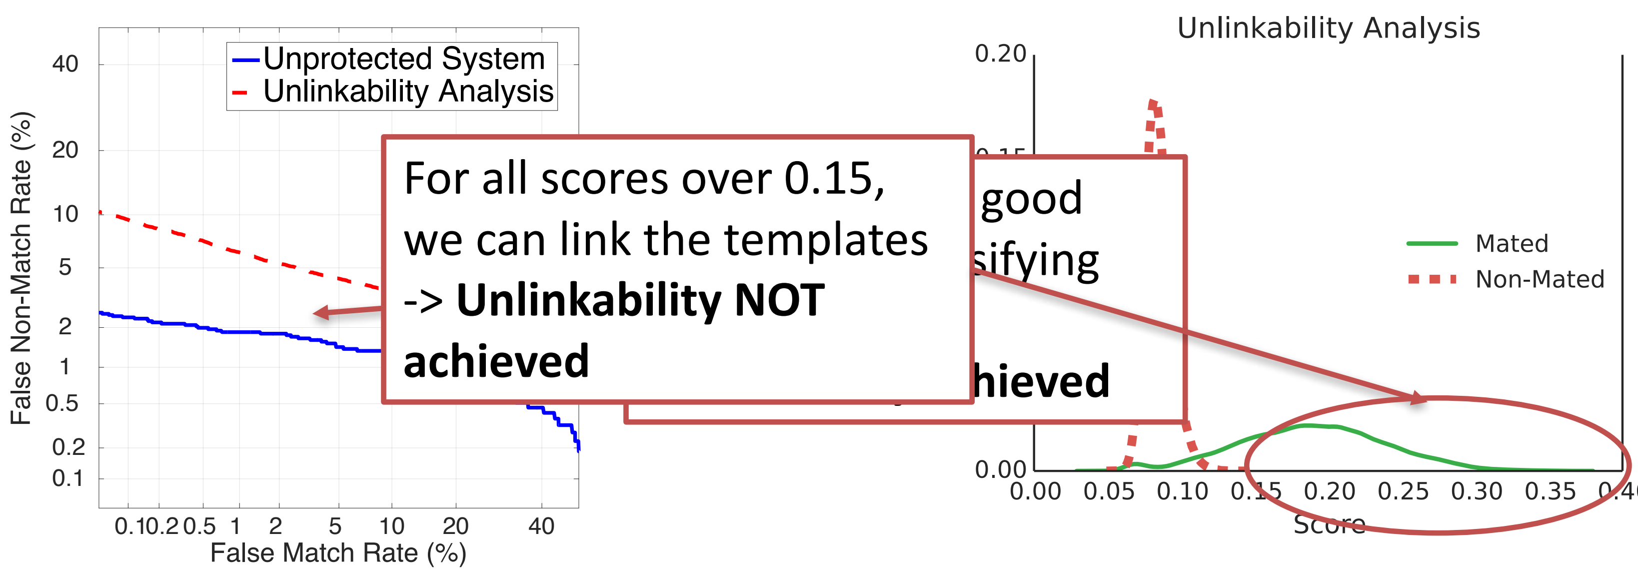

[Nagar10] A. Nagar, K. Nandakumar, A. K. Jain, "Biometric Template Protection Transformation: A Security Analysis", *SPIE, Electronic Imaging, Media Forensics and Security*, 2010.

[Kelkboom11] E. Kelkboom, J. Breebart, T. Kevenaar *et al.*, "Preventing the Decodability Attack based Cross-Matching in a Fuzzy Commitment Scheme", *IEEE TIFS*, 2011.

Unlinkability Analysis: Current Status (III)

- Plot *Mated* and *Non-mated samples* distributions, for templates protected with different keys.
- How to analyse those distributions? $\Rightarrow$ Kullback-Leibler (D_{KL}) divergence

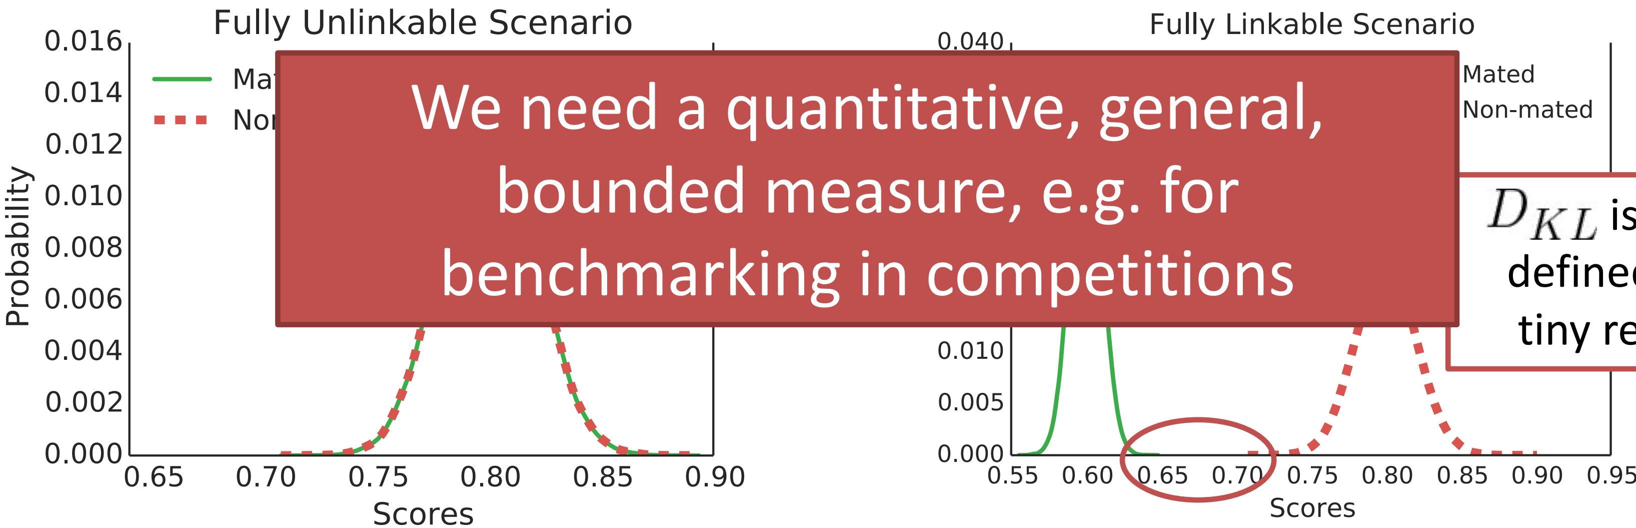

We need a quantitative, general, bounded measure, e.g. for benchmarking in competitions

D_{KL} is only defined in a tiny region

D_{KL} = 0.0

D_{KL} = 0.0005

D_{KL} is not bounded: $D_{KL} \in [0, \infty)$ $\Rightarrow$ difficult to compare systems

Unlinkability Analysis: Proposal

- Two measures:
  - Local measure $D_{\leftrightarrow}(s)$ ➔ for which scores is the system vulnerable?
  - Global measure $\mathrm{D}_{\leftrightarrow}^{sys}$ ➔ how can we compare two systems globally?
- Both bounded in [0,1], and defined for all dissimilarity scores.
- General measures, valid for all BTP schemes

Full Unlinkability

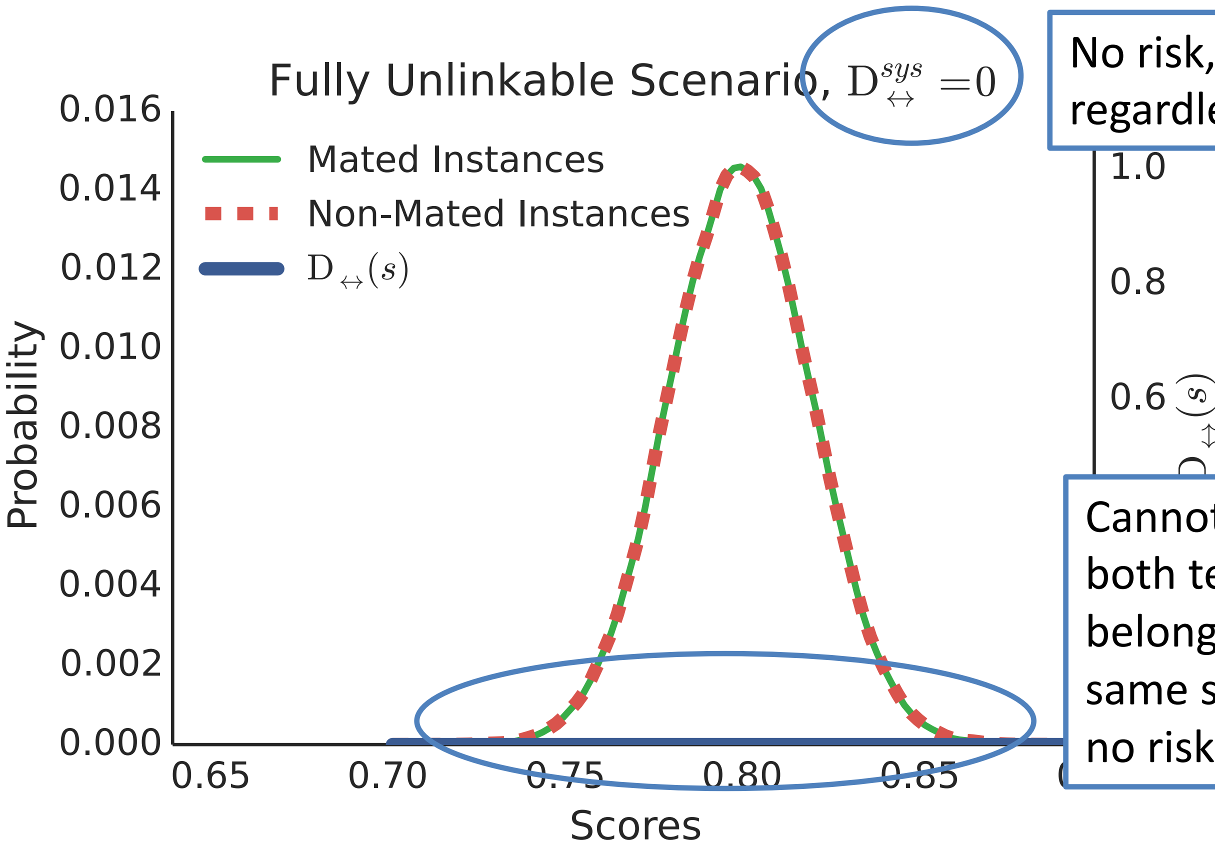

No risk, regardless of *s*

Cannot ensure both templates belong to the same subject ➔ no risk

Full Linkability

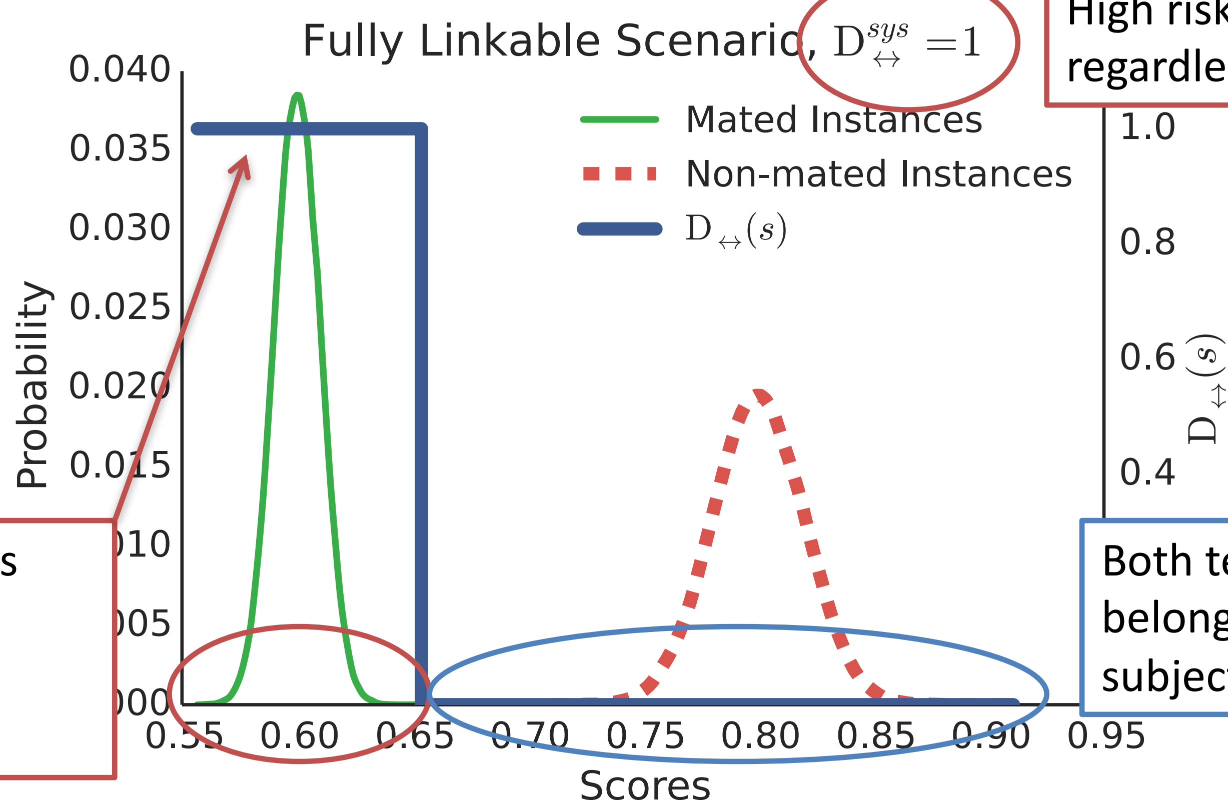

High risk, regardless of *s*

Both templates belong to the same subject ➔high risk

Both templates belong to different subjects ➔ no risk

Semi-Linkable Scenario A

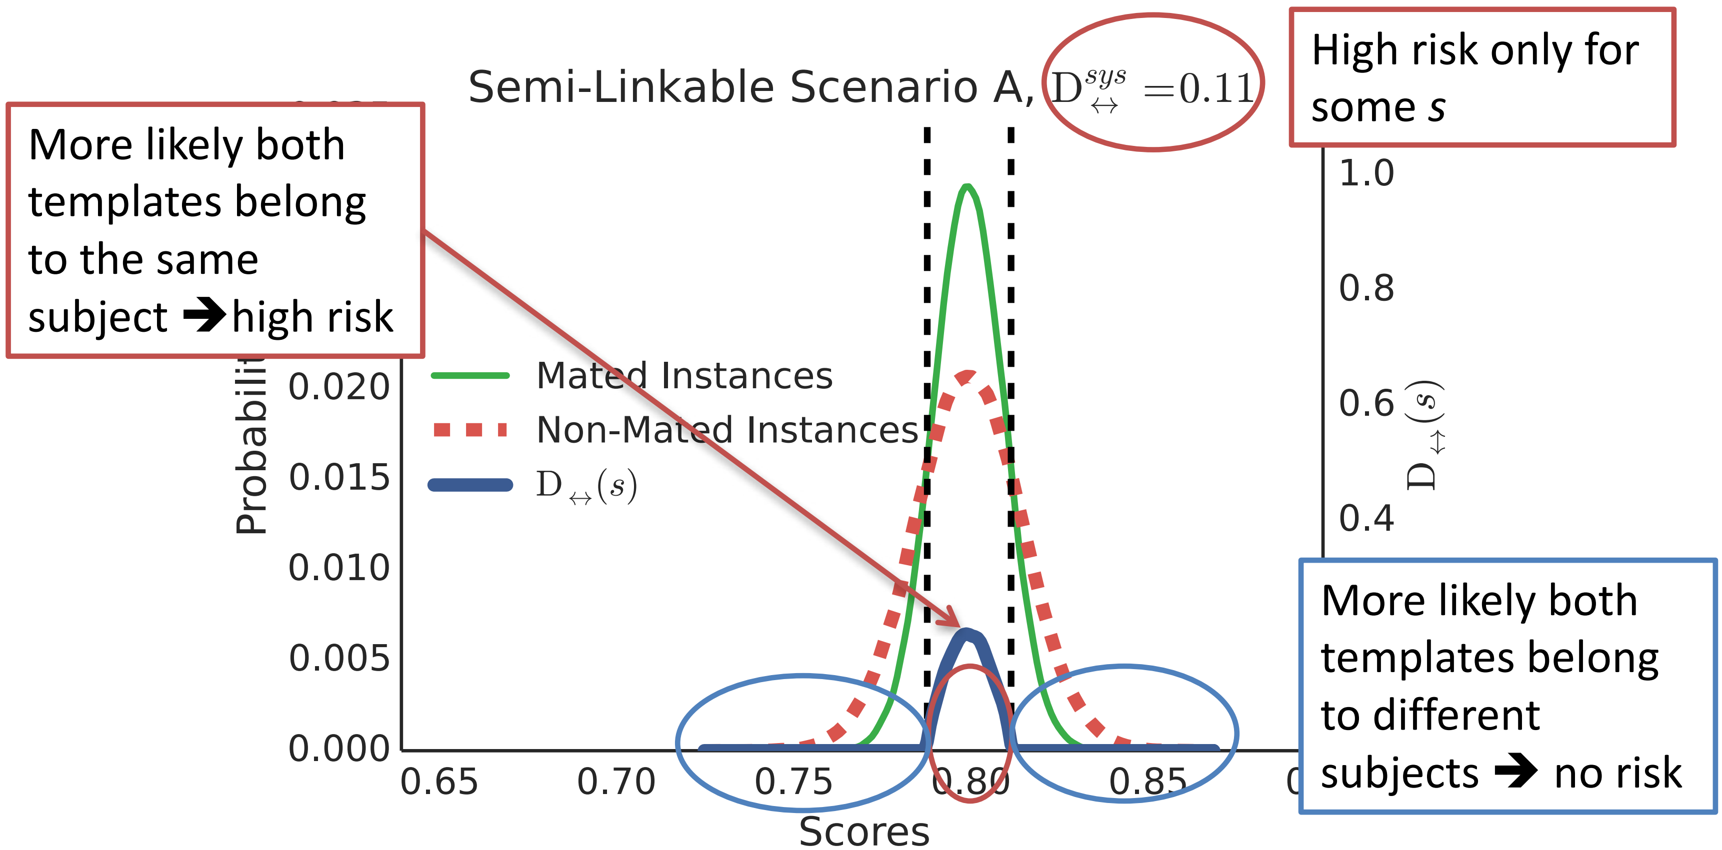

Semi-Linkable Scenario B

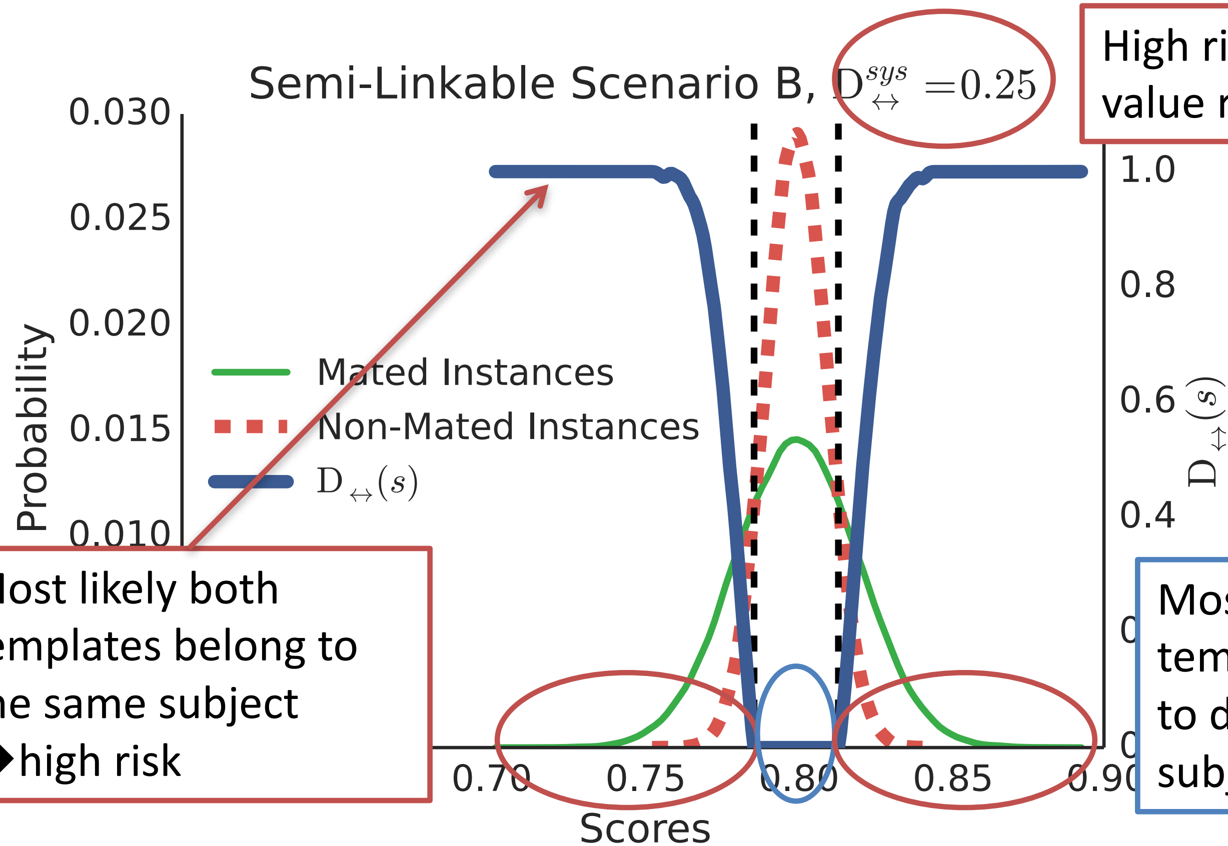

High risk, for larger value range *s*

Most likely both templates belong to the same subject ➔high risk

Most likely both templates belong to different subjects ➔ no risk

Local measure: Background

- We are interested in evaluating: $D_{\leftrightarrow}(s) = p(H_m|s) - p(H_{nm}|s)$
- But we don't know $p(H_m|s)$, $p(H_{nm}|s)$
- He can use LRs: $LR(s) = \frac{p(s|H_m)}{p(s|H_{nm})} = \frac{p(H_m|s)}{p(H_{nm}|s)} \cdot \frac{p(H_{nm})}{p(H_m)}$
- Doing some tricks, we get:

$$p(H_m|s) = \frac{LR(s) \cdot \omega}{1 + LR(s) \cdot \omega} \qquad \omega = p(H_m)/p(H_{nm})$$

Local measure: final definition

$$D_{\leftrightarrow}(s) = \begin{cases} 0 & \text{if } LR(s)\cdot\omega \leq 1 \\ 2\frac{LR(s)\cdot\omega}{1+LR(s)\cdot\omega} - 1 & \text{if } LR(s)\cdot\omega > 1 \end{cases}$$

- If we know $p(H_m)$, $p(H_{nm})$, use them to set ω
- Otherwise, assume $p(H_m) = p(H_{nm})$ and $\omega = 1$

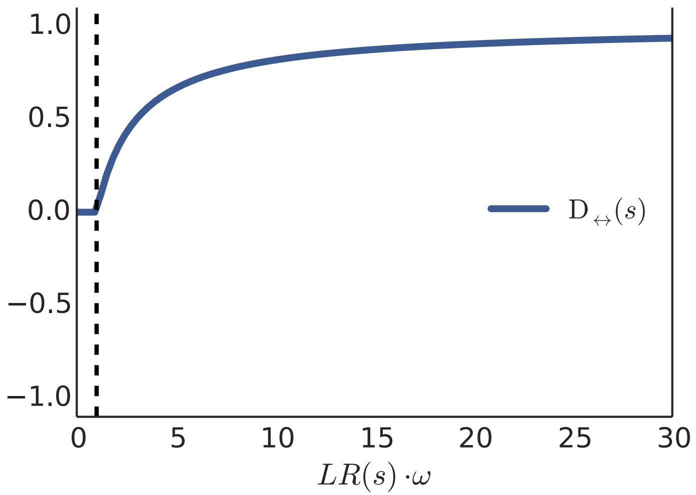

Global measure

- Global measure

$$\int_{s_{min}}^{s_{max}} p\left(H_m \cap s\right) - p\left(H_{nm} \cap s\right) \mathrm{d}s = \int_{s_{min}}^{s_{max}} p\left(s\right) \cdot \left(p\left(H_m|s\right) - p\left(H_{nm}|s\right)\right) \mathrm{d}s$$

$$= p\left(H_m\right) \int_{s_{min}}^{s_{max}} p\left(s|H_m\right) \cdot \left(p\left(H_m|s\right) - p\left(H_{nm}|s\right)\right) \mathrm{d}s +$$

$$p\left(H_{nm}\right) \int_{s_{min}}^{s_{max}} p\left(s|H_{nm}\right) \cdot \left(p\left(H_m|s\right) - p\left(H_{nm}|s\right)\right) \mathrm{d}s$$

$$p\left(H_m|s\right) > p\left(H_{nm}|s\right)$$

$$\mathrm{D}_{\leftrightarrow}^{sys} = \int_{s_{min}}^{s_{max}} p\left(s|H_m\right) \cdot \mathrm{D}_{\leftrightarrow}\left(s\right) \mathrm{d}s$$

Linkability Scenarios: Summary

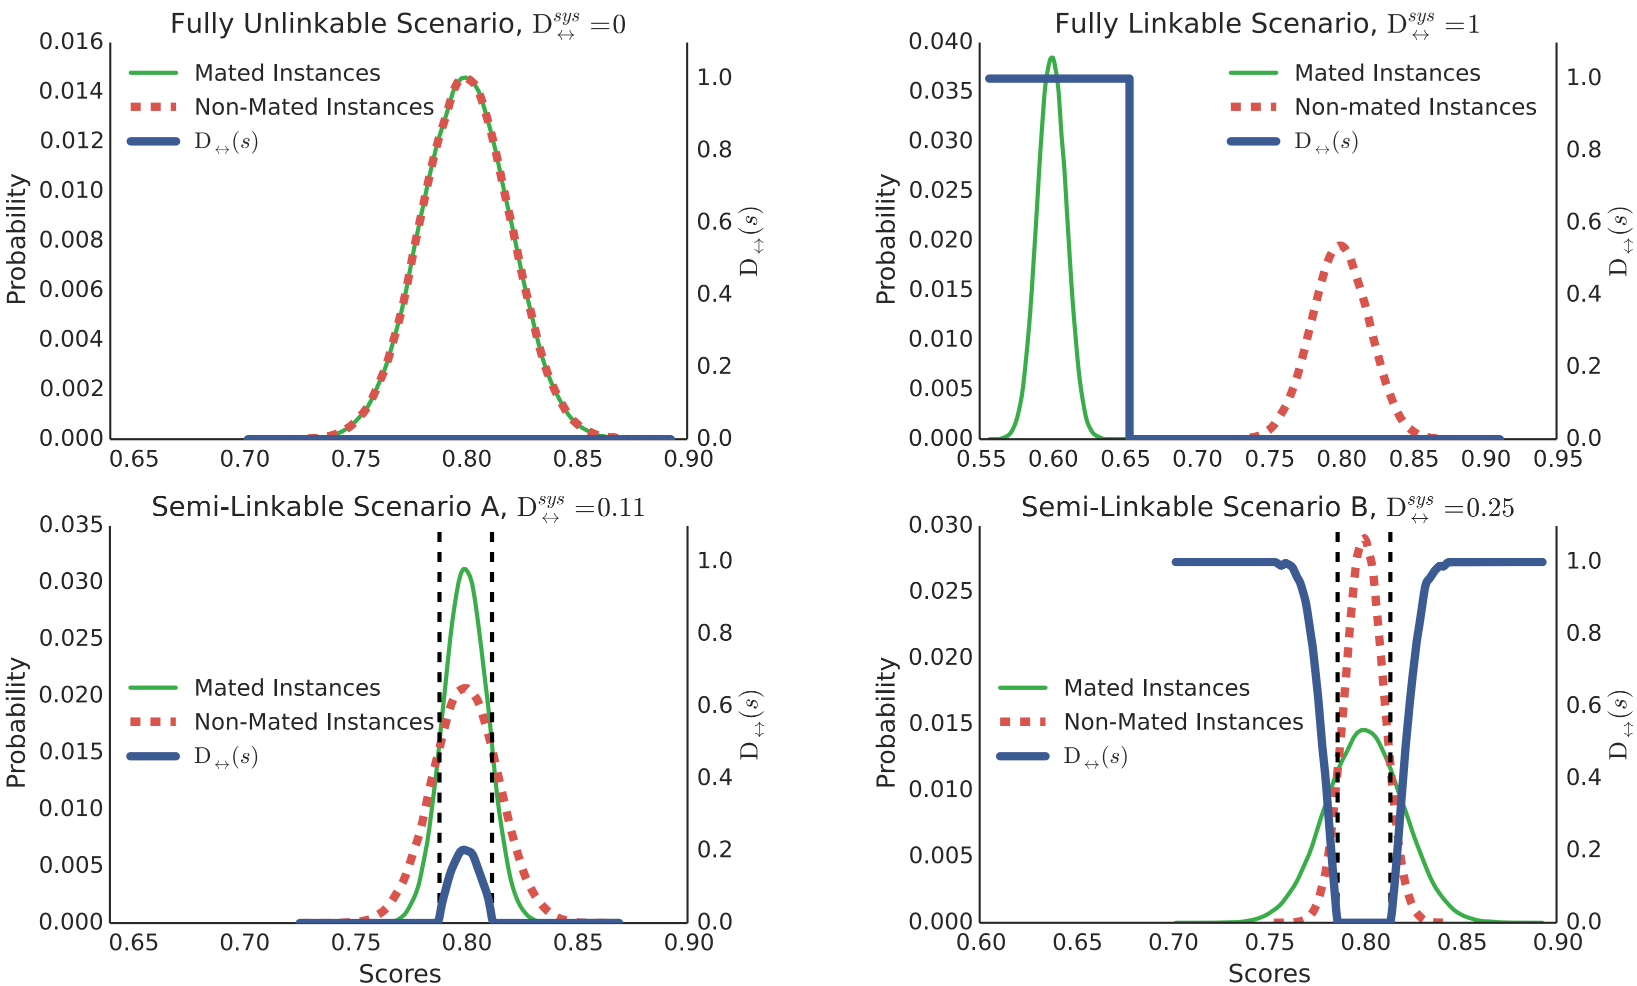

Cancelable Biometrics Based on Bloom Filters

Why Bloom filters?

[Bloom, *Comm. of the ACM* 1970]
[Broder and Mitzenmacher, *Internet Mathematics* 2004]

- Biometric Template Protection based on Bloom filters:
  - **General**: successfully applied to iris, face, fingerprint, fingervein
  - **Multimodal**: feature level fusion
  - **Irreversibility** achieved
  - **Accuracy**, depending on the configuration, preserved
  - **Template size**: similar or compressed
  - **Verification speed** similar

- But we need to add **unlinkability**
- And find a way to fuse templates of different sized (**Multi-Biometrics**)

General architecture

- Adding unlinkability:
  - Small complexity
  - Small impact on accuracy

Random shuffling of bits ⇒ ↑EER > 40%

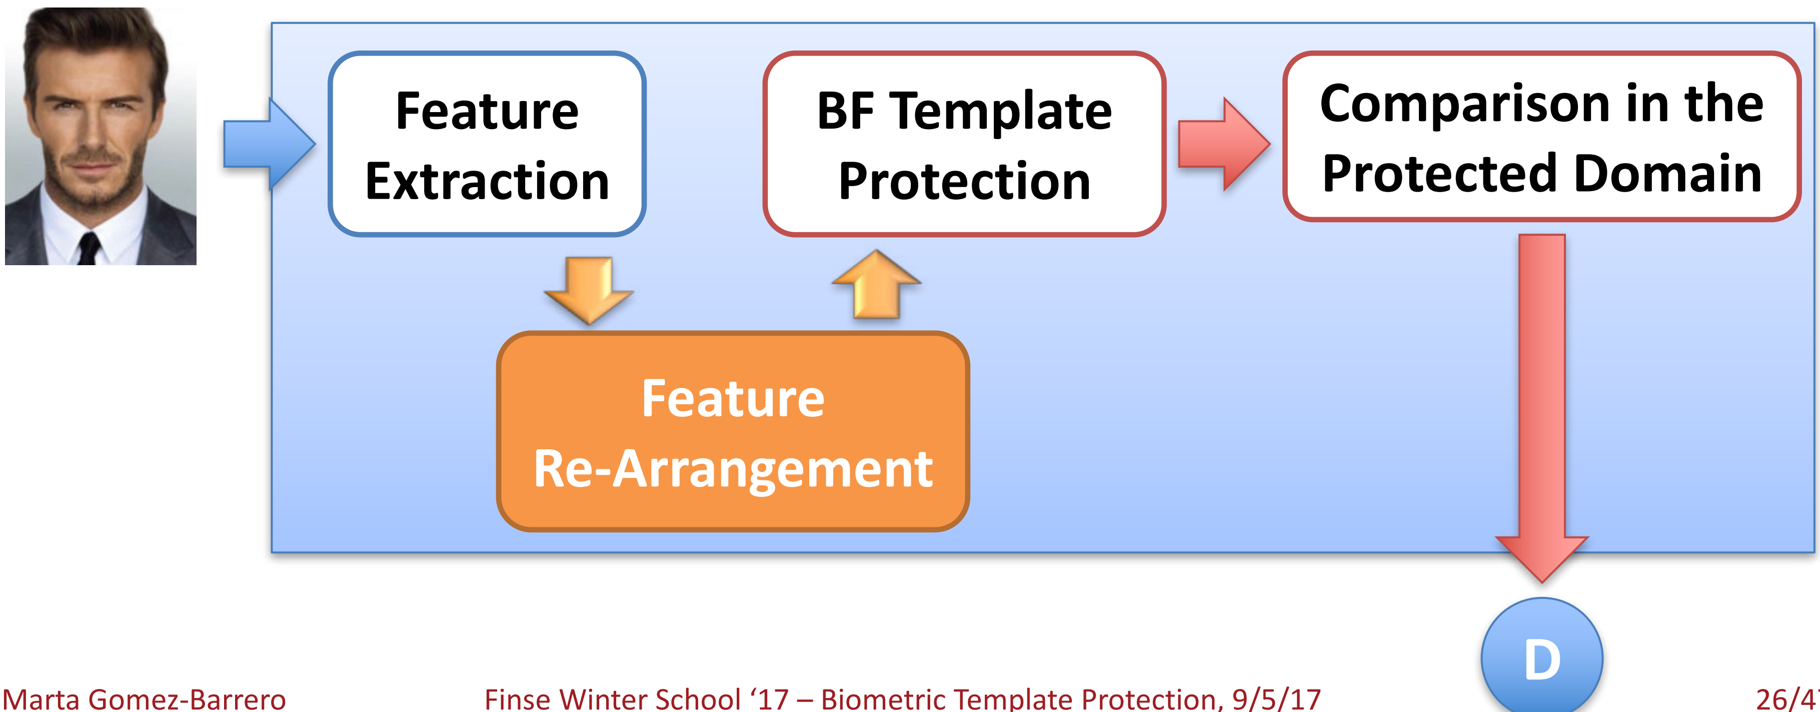

Bloom filters

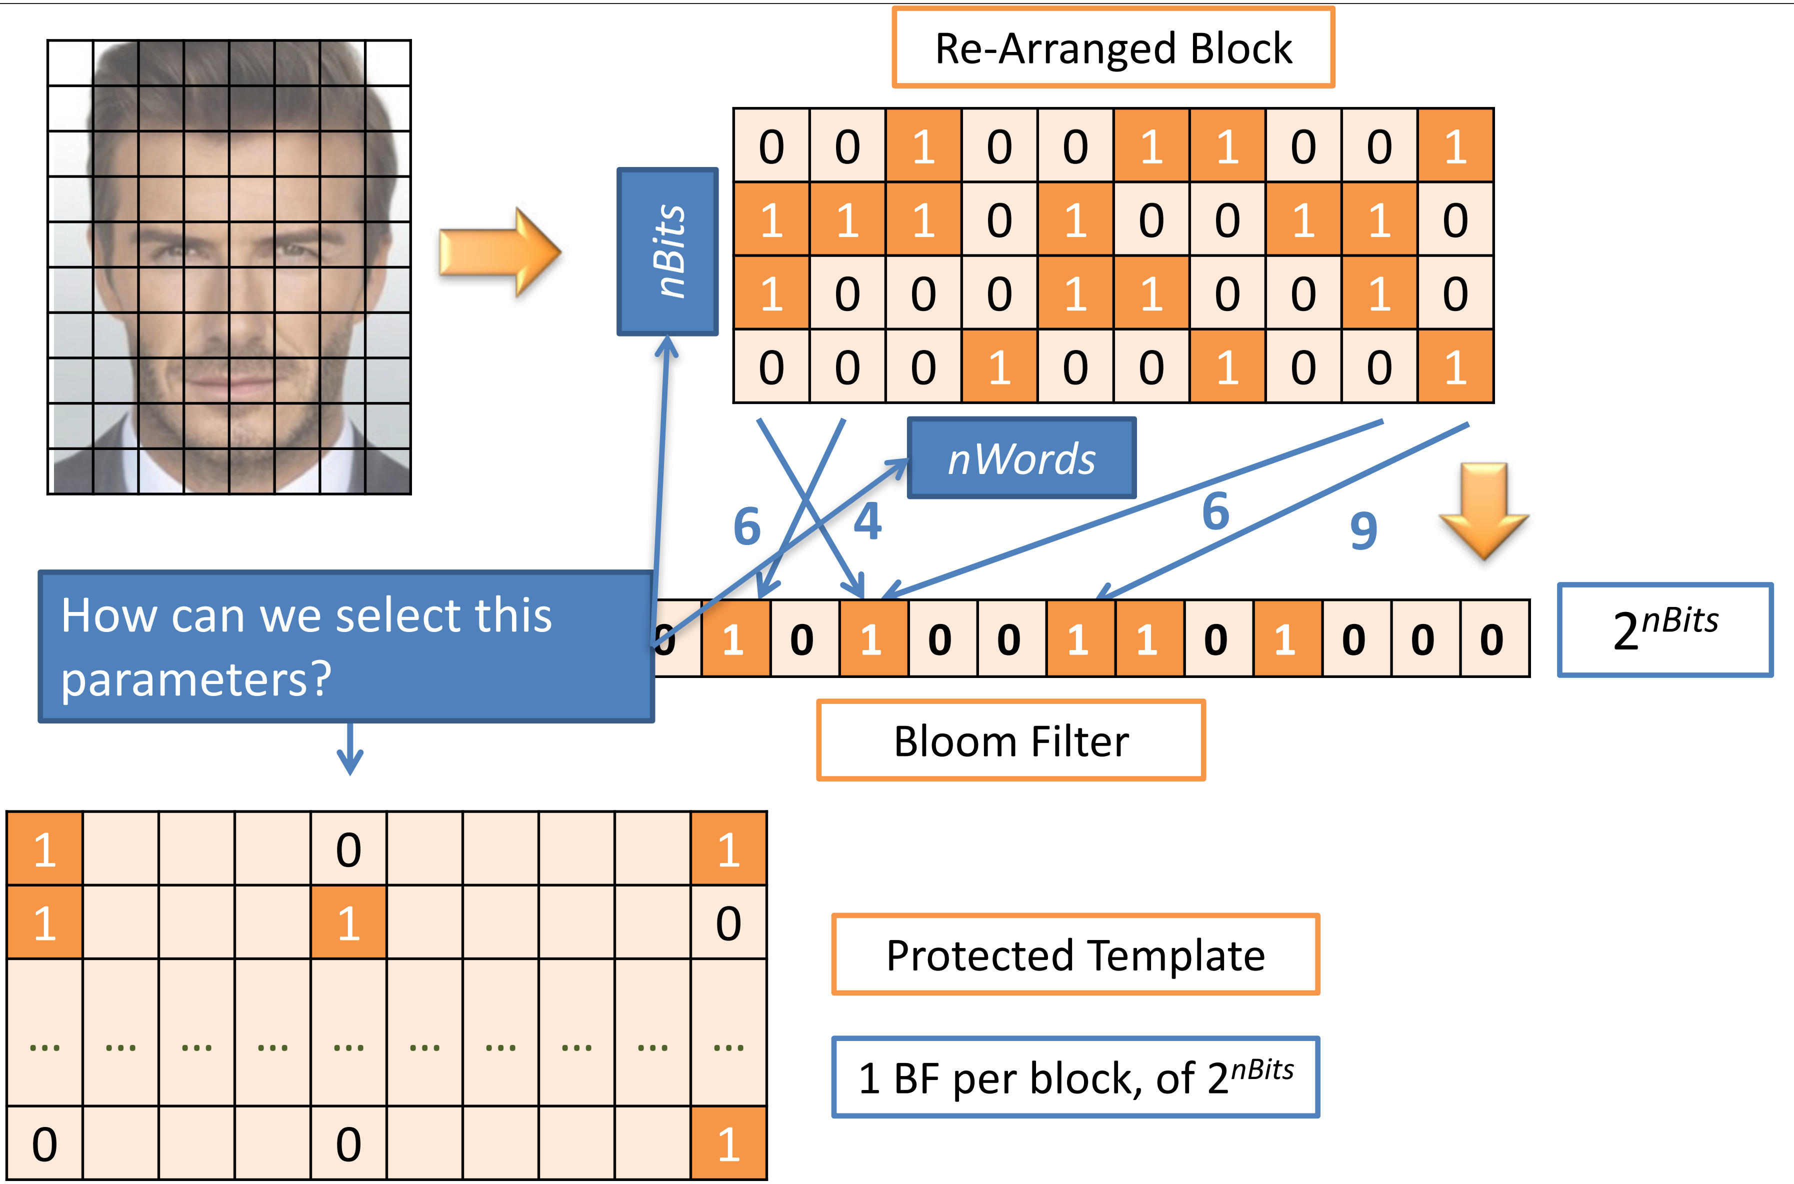

Bloom filters

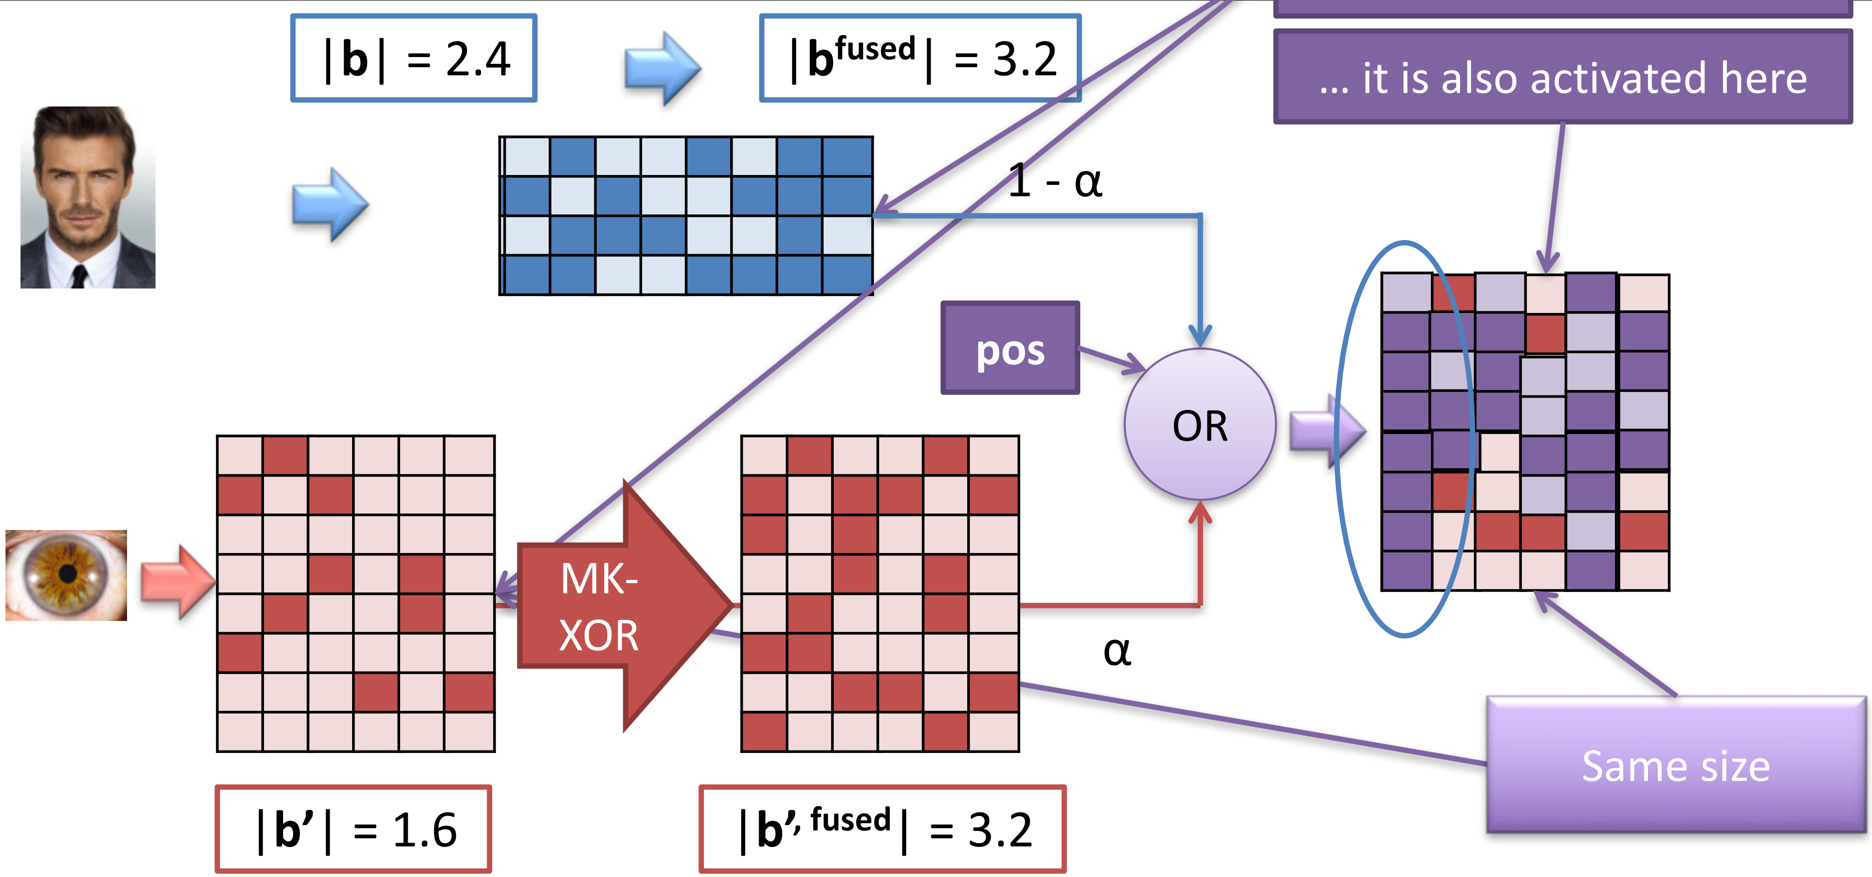

To achieve a a fusion weight α:

Different number of keys
=> different α

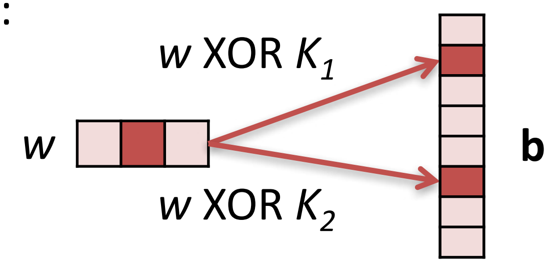

Set number of keys in terms of:
$|\mathbf{b}^{fused}| \,/\, |\mathbf{b}'|$

Accuracy Analysis

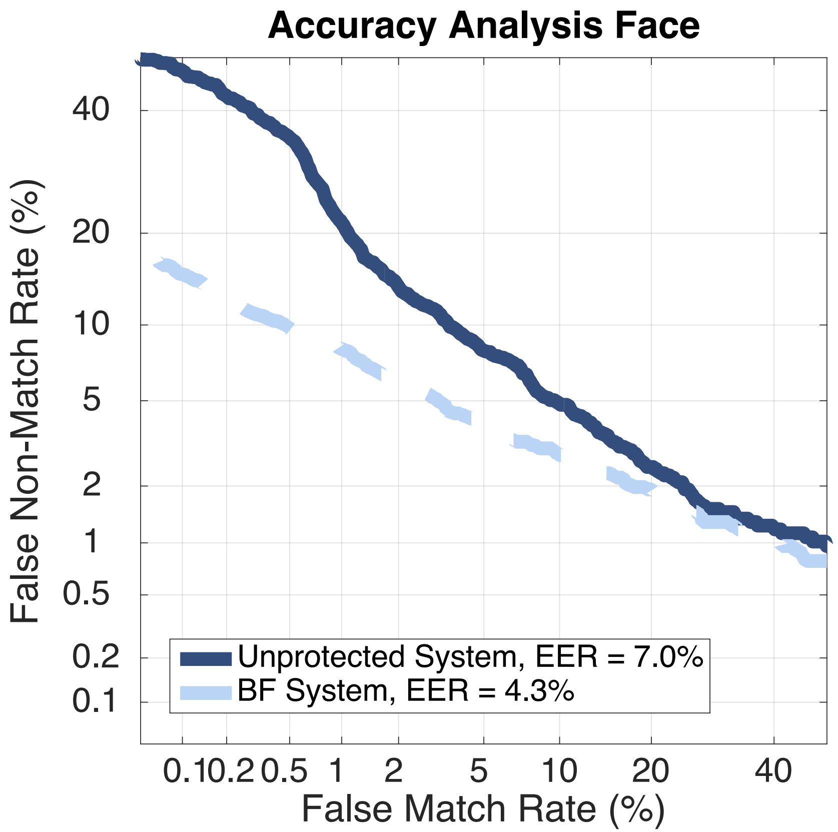

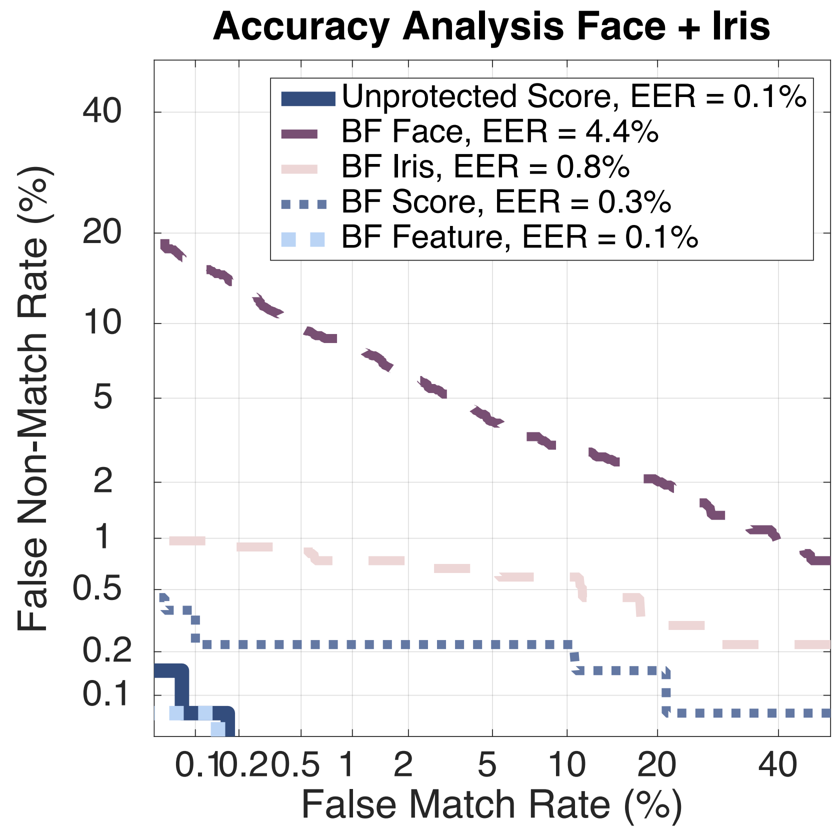

Accuracy is preserved at all operating points

For the fusion, best accuracy for protected feature level

Irreversibility analysis

➢ Are the reconstructed unprotected templates similar to the original ones?

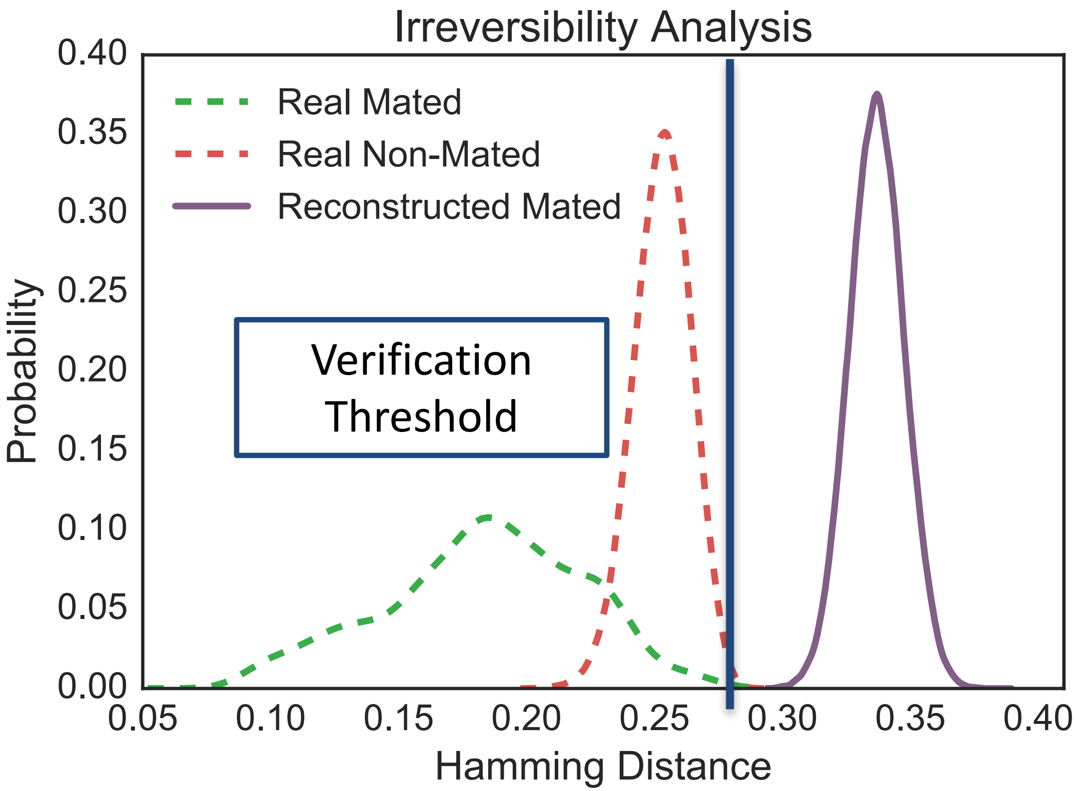

Irreversible: HD bigger than impostor comparisons

[Bringer *et al.*, *ICB* 2015]

Unlinkability analysis (I)

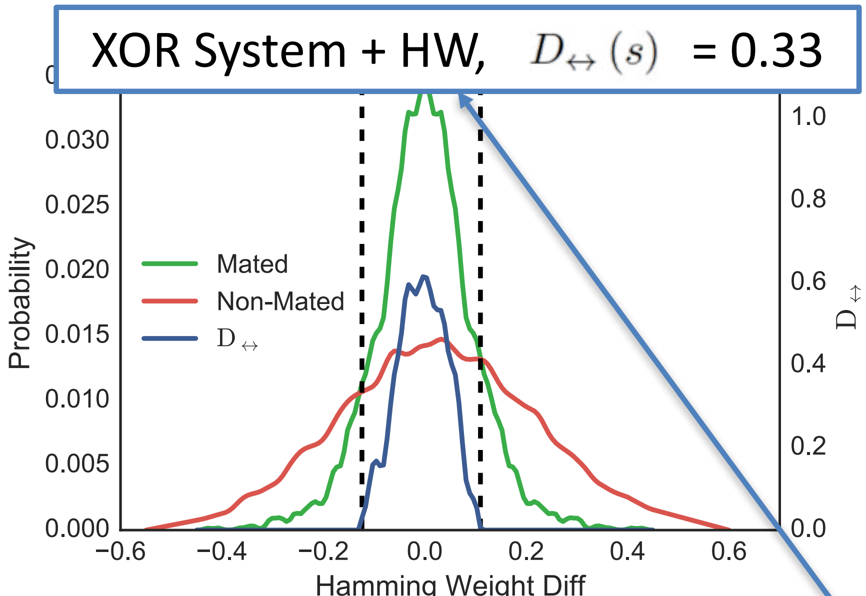

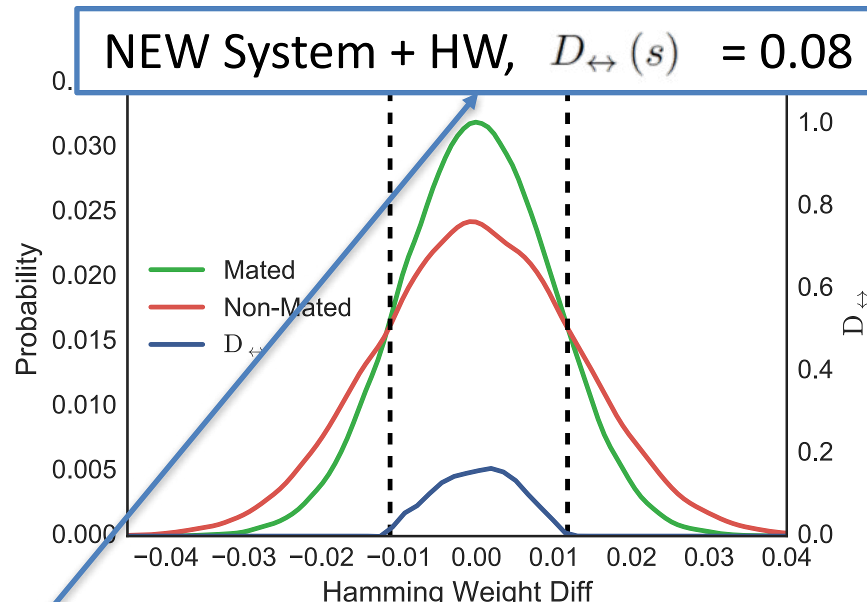

Linkability has decreased! ☺

Unlinkability analysis (II)

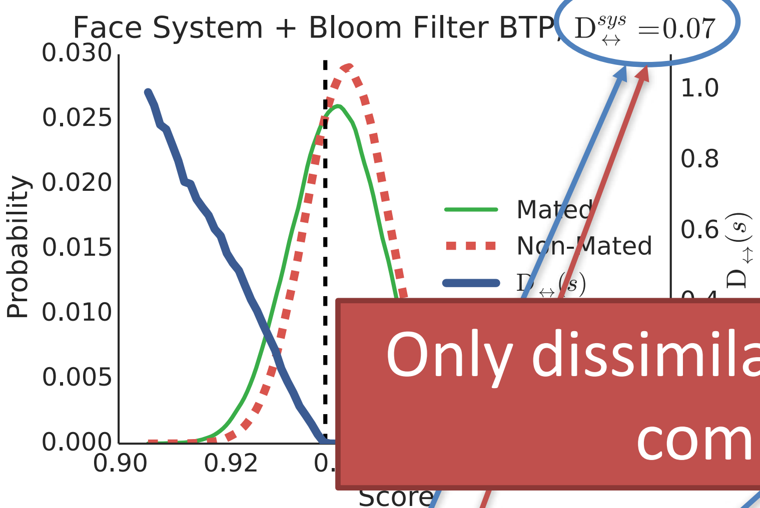

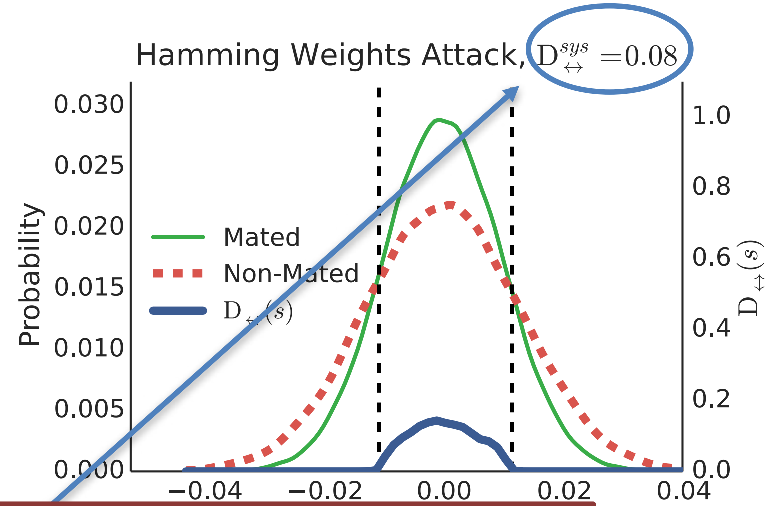

Only dissimilarity scores are needed to compute the metrics

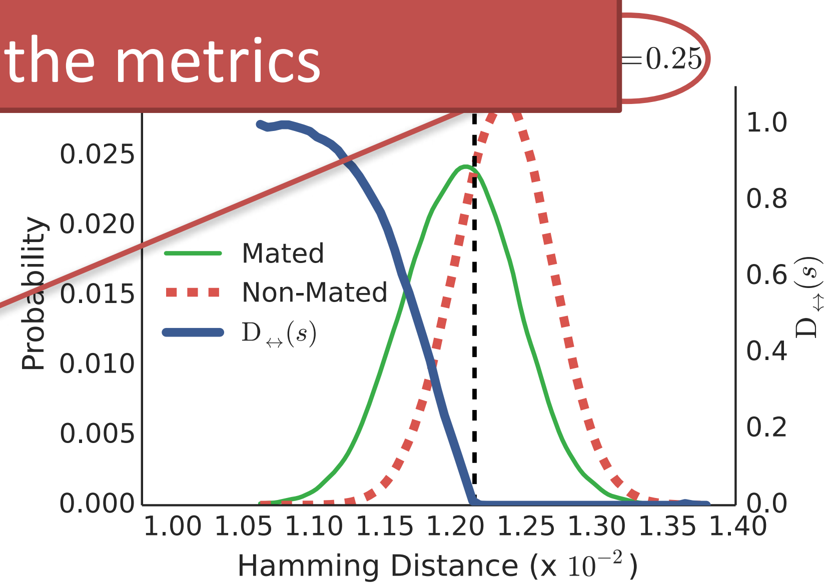

Linkability has barely increased ☺

Still room for improvement

BTP Based on Homomorphic Encryption

Why Homomorphic Encryption?

- BTP based on Homomorphic Encryption:
  - **General**
  - **Accuracy fully preserved**
  - **Permanent protection:** all computations in the encrypted domain
  - **Irreversibility** and **unlinkability** achieved
  - **Renewability** with no re-acquisition

- Limitation on the number of operations in the encrypted domain
- Secret key + protected template = unprotected template compromised

[Fontaine *et al.*, *EURASIP J. Inf. Sec.* 2007]
[Lagendijk *et al.*, *IEEE SP Mag.* 2013]

Homomorphic Encryption

- Practical implementation: Paillier Cryptosystem [P. Paillier, EUROCRYPT, 1999]
- HE- Paillier: based on the DECISIONAL COMPOSITE RESIDUOSITY ASSUMPTION

DCRA: given a composite n and and integer z, it is (very) hard to decide whether there exists y such that:

$$z = y^n \pmod{n^2}$$

Additive Homomorphic Encryption

$$D_{sk}(m_1^* \cdot m_2^* \mod n^2) = m_1 + m_2 \mod n$$

Product of ciphertexts

Sum of plain texts

$$D_{sk}\left((m_1^*)^l \mod n^2\right) = m_1 \cdot l \mod n$$

Exponentiation of ciphertext and plain text

Product of plain texts

General architecture

Additionally, only integer values can be handled

Problem 1: what do we store in the database?

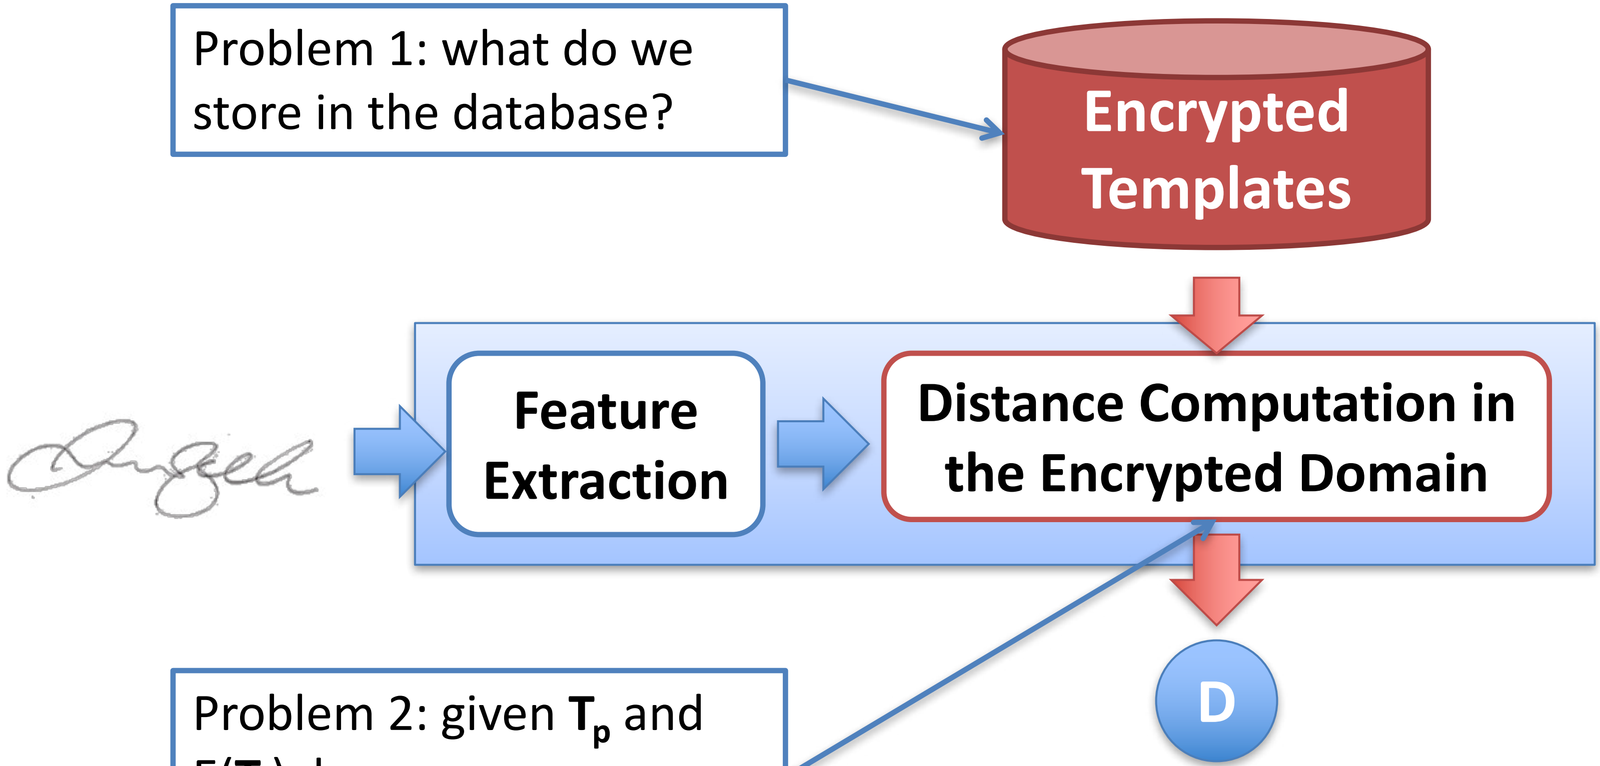

Problem 2: given $\mathbf{T}_p$ and $E(\mathbf{T}_r)$, how can we compute $E(d(\mathbf{T}_p, \mathbf{T}_r))$?

Multi-Biometrics

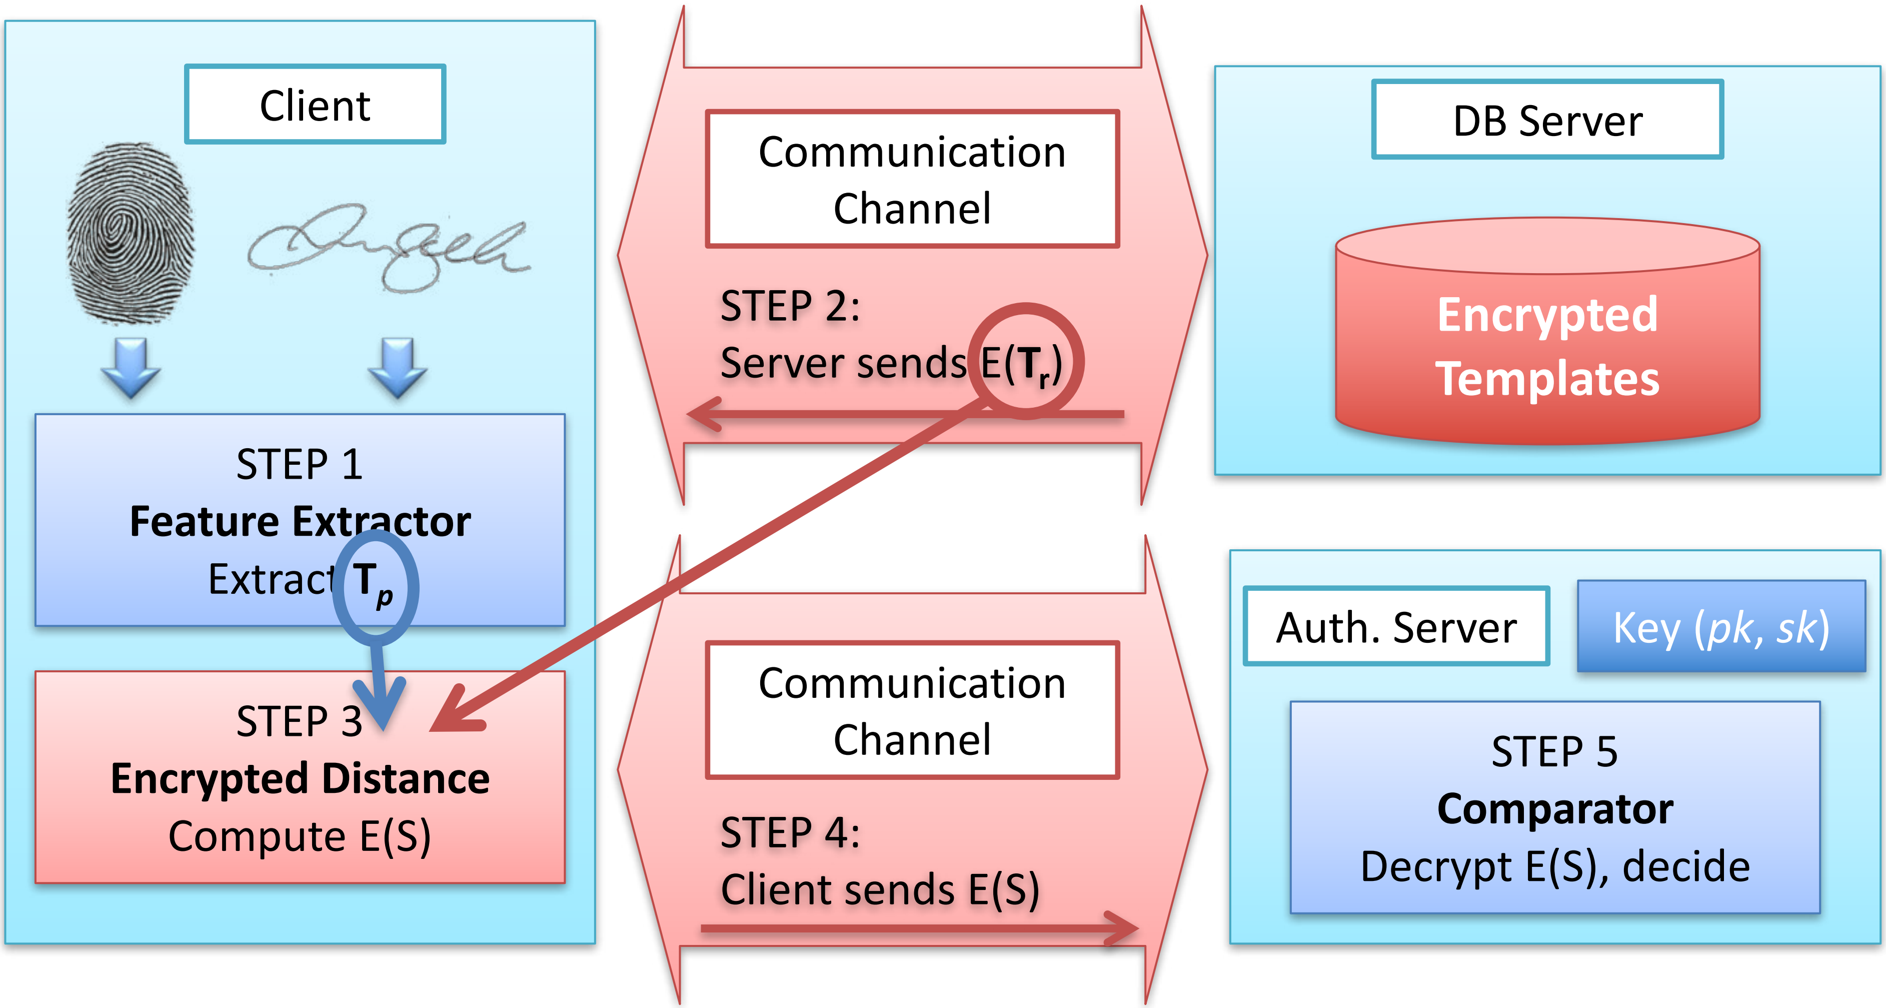

Encrypted distance computation

Euclidean distance: Given two vectors $\mathbf{T}_p$ and E($\mathbf{T}_r$), of length *F*

$$S_{euc} = \sum_{f=1}^{F} p_f^2 + r_f^2 - 2 p_f r_f$$

Encrypted Euclidean distance: Given two vectors $\mathbf{T}_p$ and E($\mathbf{T}_r$), of length *F*

$$E\left(S_{euc}\right) = \prod_{f=1}^{F} E\left(1\right)^{p_f^2} \cdot E\left(r_f^2\right) \cdot E\left(r_f\right)^{-2p_f}$$

Encrypted reference template stored in DB

Probe template

BTP & HE

Cosine similarity: Given two vectors $\mathbf{T}_p$ and $\mathbf{T}_r$, of length F

$$d_{cos}\left(\mathbf{T}_p, \mathbf{T}_r\right) = \frac{\mathbf{T}_p \cdot \mathbf{T}_r}{\|\mathbf{T}_p\| \cdot \|\mathbf{T}_r\|} = \sum_{f=1}^{F} \frac{p_f \cdot r_f}{\|\mathbf{T}_p\| \cdot \|\mathbf{T}_r\|}$$

$$d_{cos}\left(\mathbf{T}_p, \mathbf{T}_r\right) \in [0, 1] \quad \Rightarrow \quad S_{cos} = 10^{12} d_{cos}\left(\mathbf{T}_p, \mathbf{T}_r\right)$$

Encrypted Cosine similarity: Given two vectors $\mathbf{T}_p$ and $E(\mathbf{T}_r)$, of length F

$$E\left(S_{cos}\right) = \prod_{f=1}^{F} E\left(\frac{10^6 r_f}{\|\mathbf{T}_r\|}\right)^{10^6 p_f / \|\mathbf{T}_p\|}$$

Encrypted reference template stored in DB

Probe template

Accuracy Eva

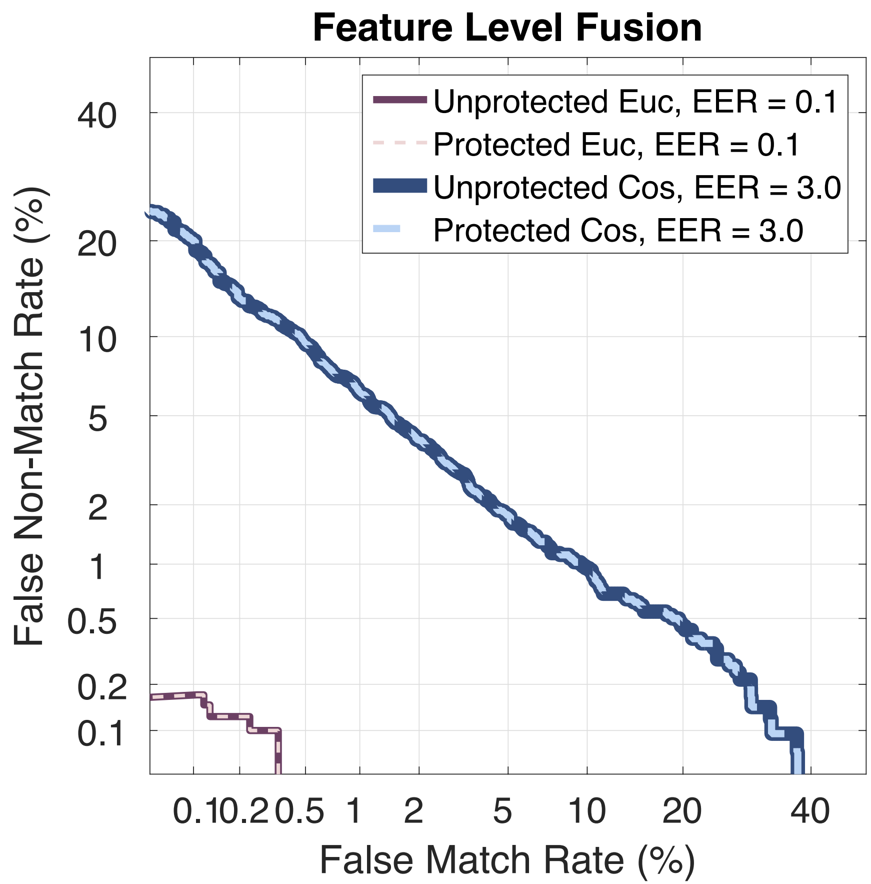

BioSecurID DB [Fierrez *et al.*, *PAA* 2009]

Accuracy is fully preserved at all operating points

Unlinkability Analysis

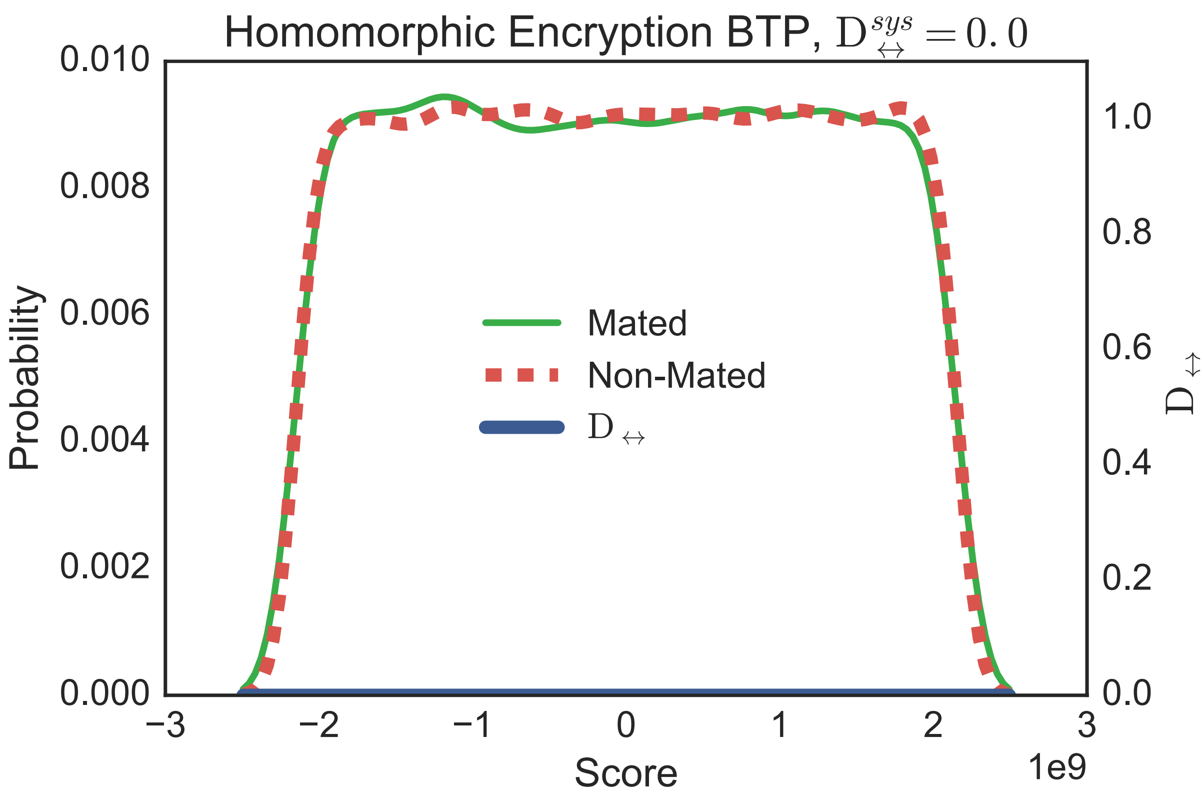

Full unlinkability, as long as the secret key is not compromised

Computational Overhead

- 1 real value (16 bits) → 2,048 bits encrypted → x 128 increase factor
- Depending on distance, more values need to be stored

Unprotected template:
F real values → 0.27 KB

Euclidean distance template:
$2F + 1$ encrypted values → 70.25 KB

Cosine distance template:
F encrypted values → 35 KB

Storage requirements and communication bandwidth multiplied by 128 - 256

However, templates are still small enough for real time apps

Summary

Summary

- Methodology for a standardized security and privacy evaluation of BTP schemes

- BTP schemes based on Bloom filters or Homomorphic Encryption comply with ISO/IEC IS 24745, providing irreversibility, unlinkability, renewability and accuracy preservation

- MBTP schemes based on Bloom filters or Homomorphic Encryption achieve higher accuracy and privacy protection

Summary

- Bloom filters advantages:
  - **Compressed** templates
  - **Irreversibility** even if **key is compromised**
  - **Low** computational **load**

- HE advantages:
  - **Full accuracy** preservation
  - Revocability with **no re-acquisition**
  - **Higher** degree of **unlinkability**

- Bloom filters limitations:
  - **Some accuracy degradation** depending on feature extractors
  - **Some accuracy degradation** at low FMRs

- HE limitations:
  - **Key compromised → reversible**
  - **Storage** requirements **x 128**

Marta Gomez-Barrero
(marta.gomez-barrero@h-da.de)