HONG KONG BAPTIST UNIVERSITY
INSTITUTE OF AUTOMATION CHINESE ACADEMY OF SCIENCES
Sponsors: ICTS IAPR IEEE Biometrics Council

IAPR/IEEE WINTER SCHOOL ON BIOMETRICS
9-13 JANUARY 2017 | HONG KONG BAPTIST UNIVERSITY, HONG KONG

Soft Biometrics and Continuous Authentication

DR. TERENCE SIM
SCHOOL OF COMPUTING
NATIONAL UNIVERSITY OF SINGAPORE

Brief Bio

- Associate Professor & Vice Dean

NUS National University of Singapore
School of Computing

- Research: face recognition, biometrics, computational photography
- PhD from CMU, MSc from Stanfrod, SB from MIT
- Google “Terence Sim”, or tsim@comp.nus.edu.sg


12

Traditional authentication: one-time

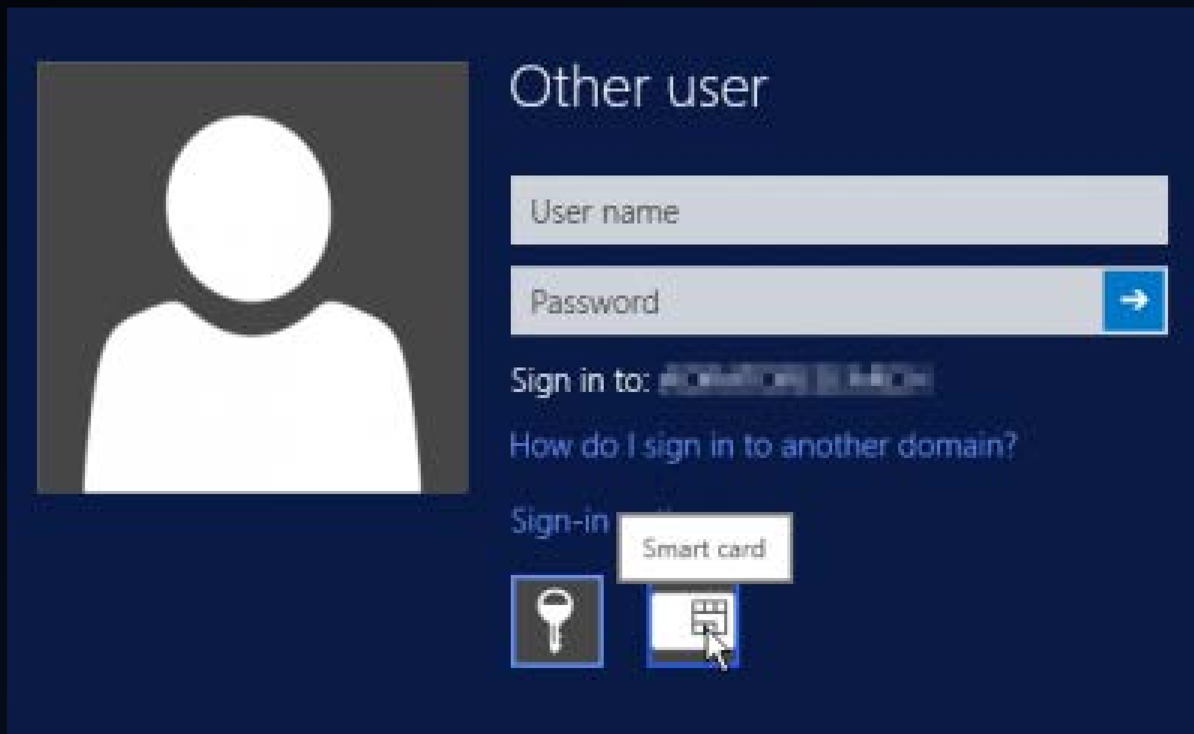

Session hijacking

System still thinks legitimate user is there!

Solution: continuous authentication

Cassandra Carrillo MSc. Thesis 2003

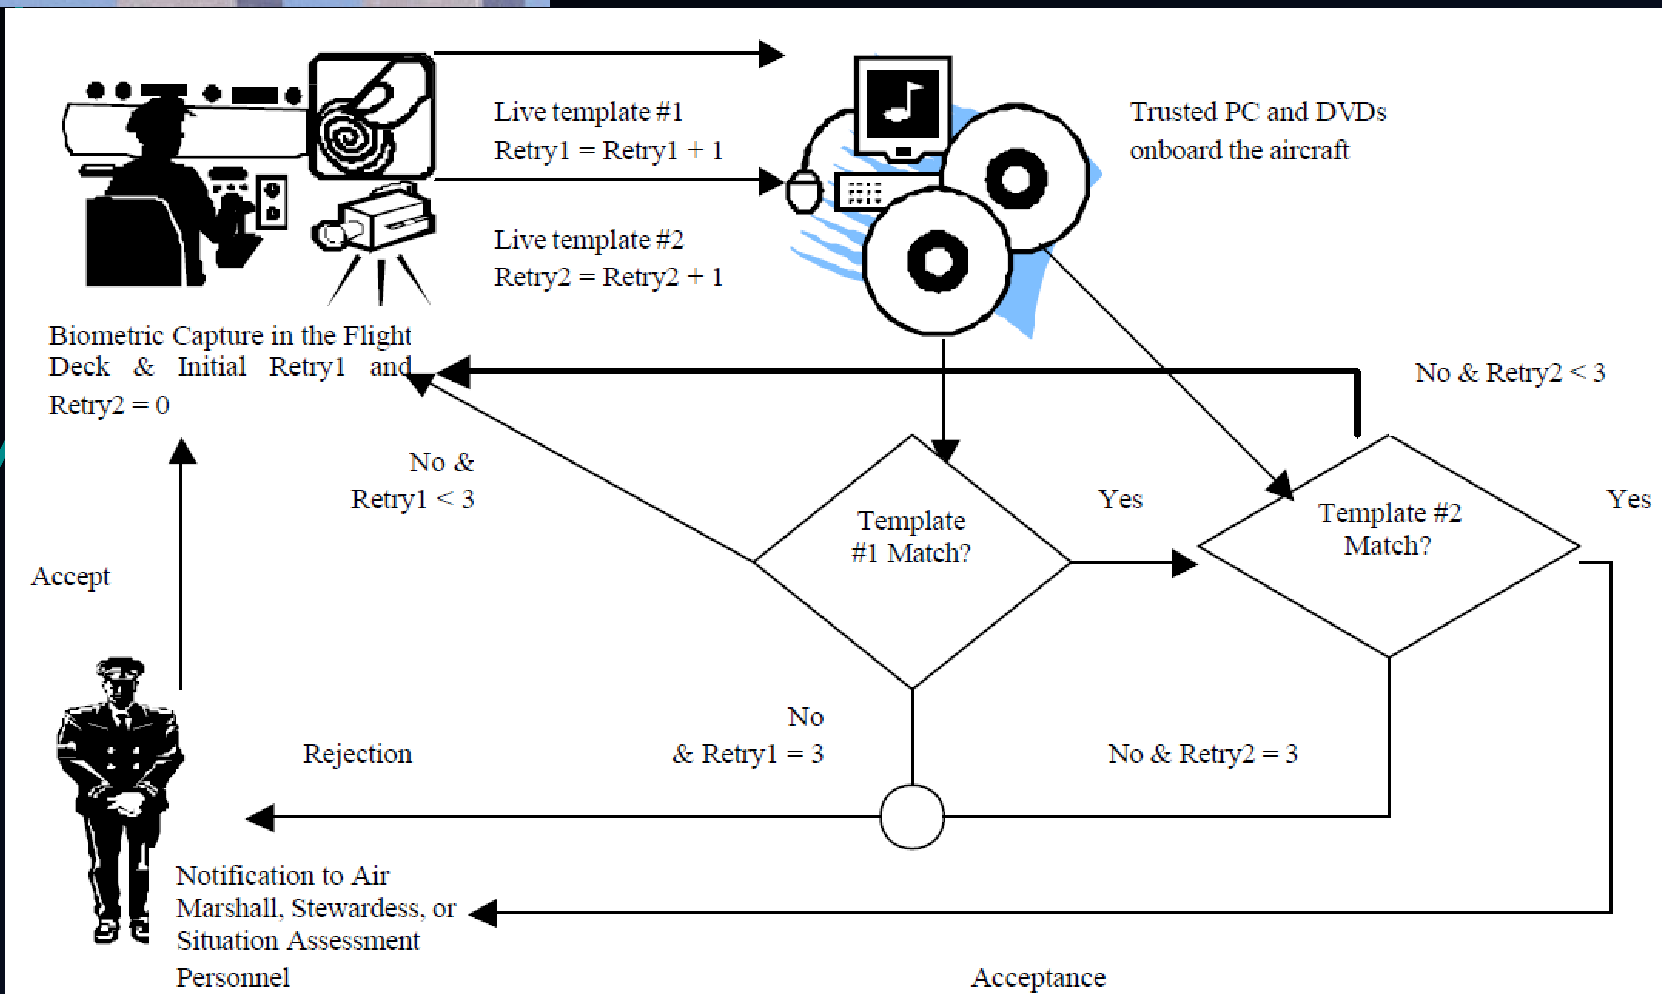

R Janakiraman, S Kumar, S Zhang, T Sim 2005

- Using Continuous Face Verification to Improve Desktop Security

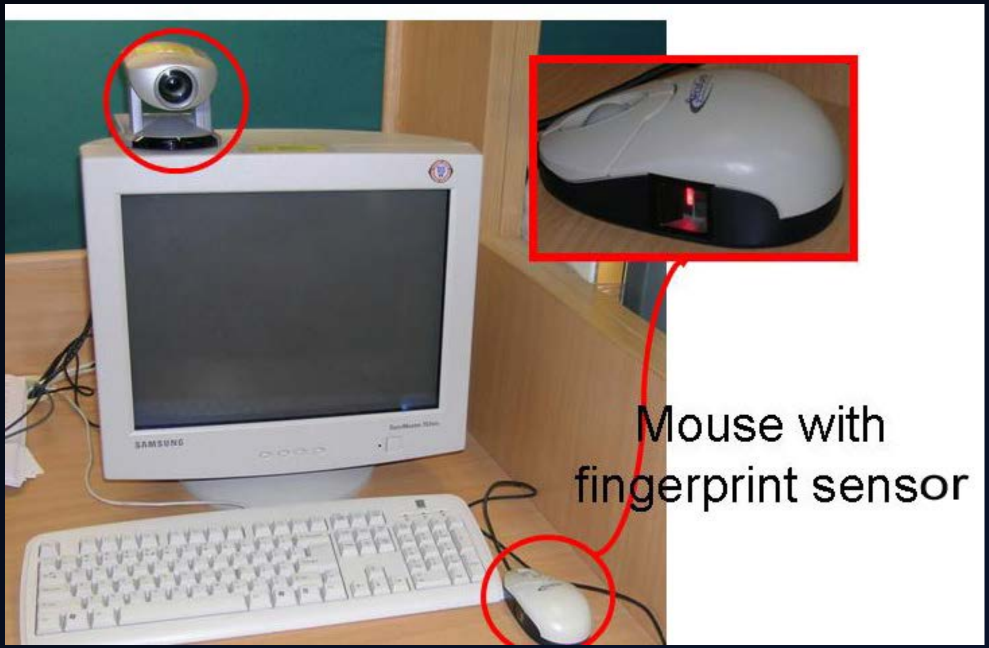

INTRODUCTION


Challenges

#1: Must be done passively

- Asking for PIN repeatedly causes frustration
- Biometrics is best suited for this

#2: Have minimal overhead

- Usability & energy issues

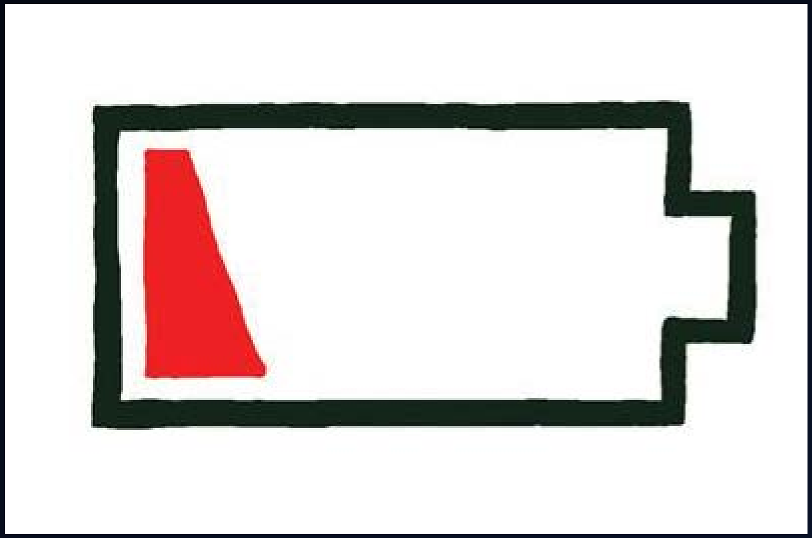

#3: Achieve low error rates

- High FAR: imposter easily takes over
- High FRR: re-login needed, user is inconvenienced
- Time must be taken into account
  - FAR & FRR not enough;
  - new performance metric needed

#4: Provide Authentication Certainty at all times

- Certainty that the legitimate user is still present
- Even when user provides no biometric signals


Challenges

CRITERIA

Observations over time

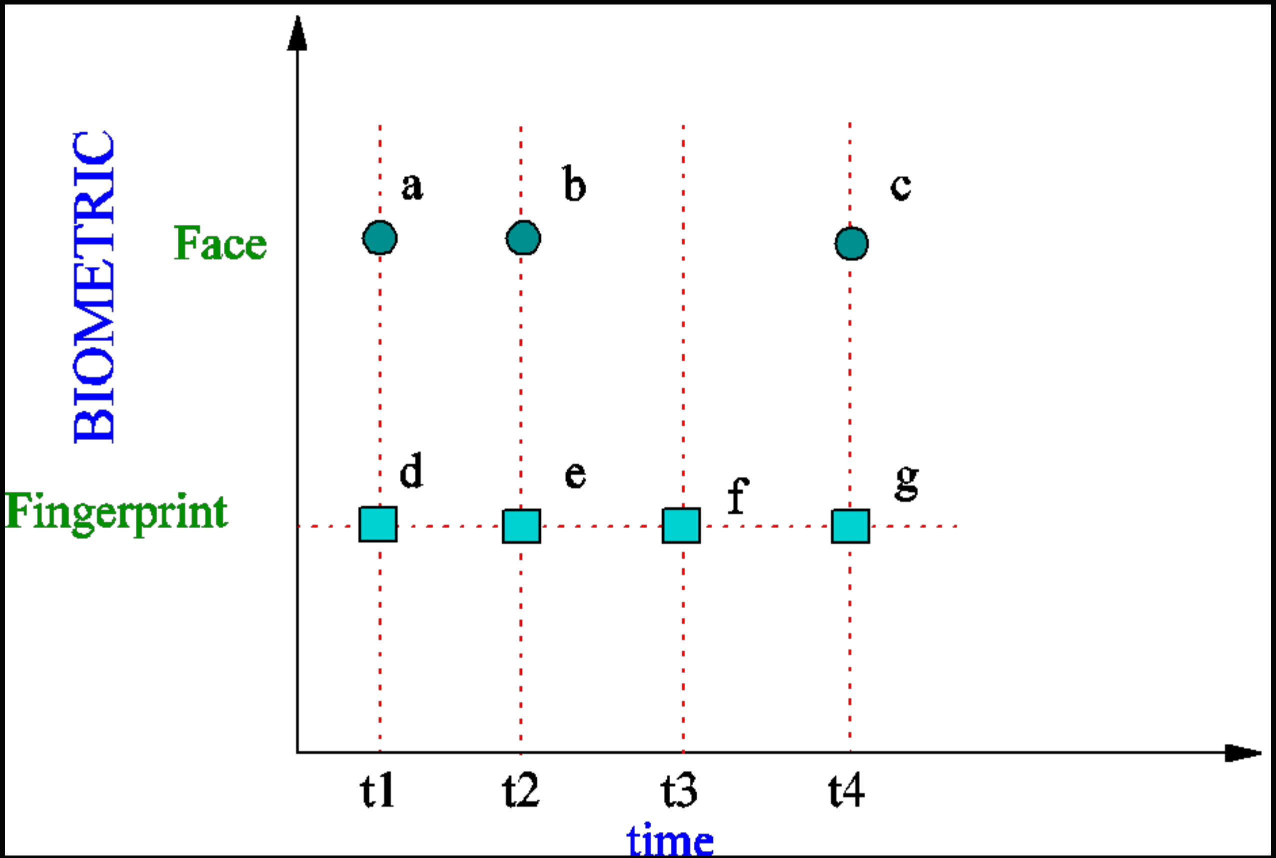

#1: Account for reliability of different modalities

- Fingerprint considered more reliable than face
- Thus must affect the authentication decision more than face

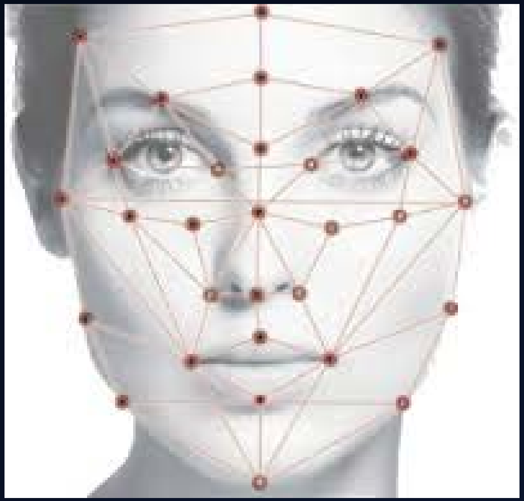

#2: Older observations must be discounted to reflect the increasing uncertainty of the continued presence of the legitimate user

- The longer the elapsed time, the more uncertain is the continued presence of the user.

#3: It must be possible to determine authentication certainty at any point in time, even when there is no observations in one or more modalities

- At any time, the system must be able to check if the legitimate user is still present.

CRITERIA

Mouse with fingerprint sensor

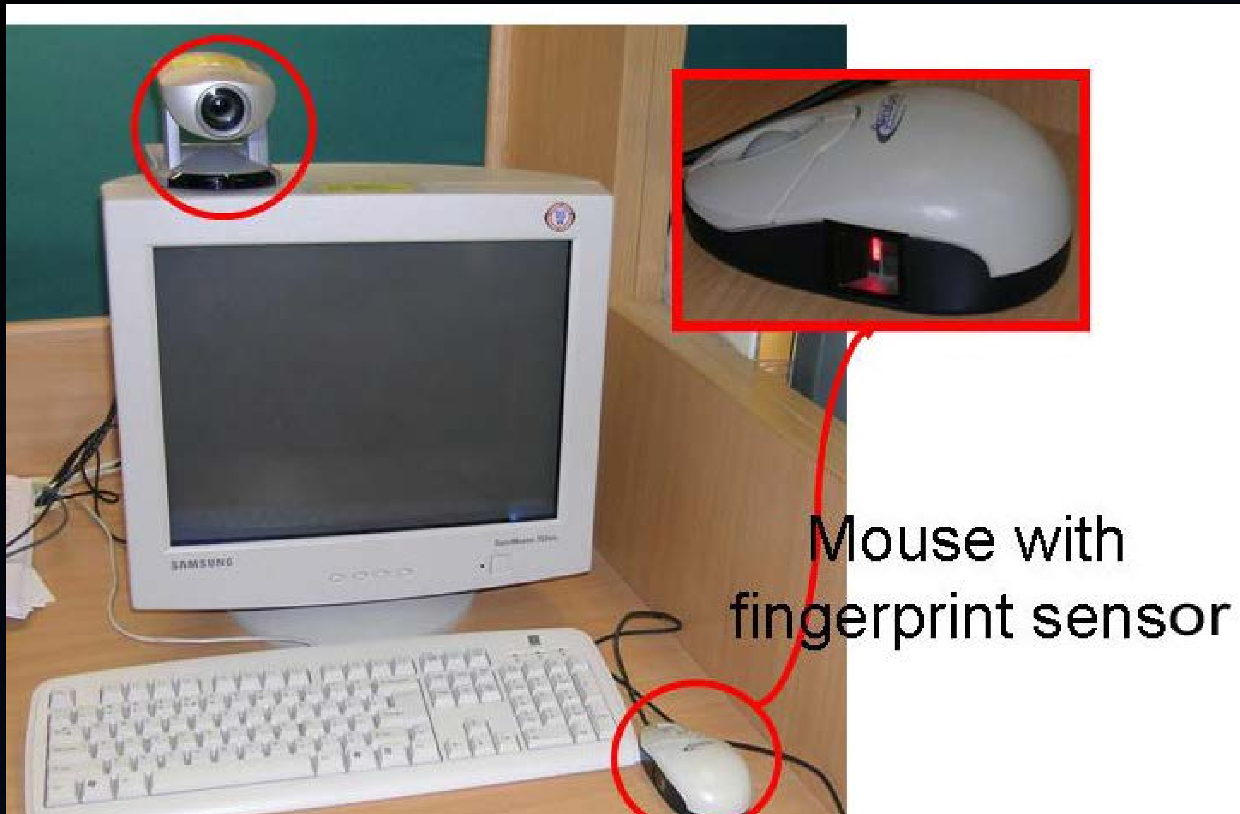
SAMSUNG
Mouse with
fingerprint sensor

System Architecture

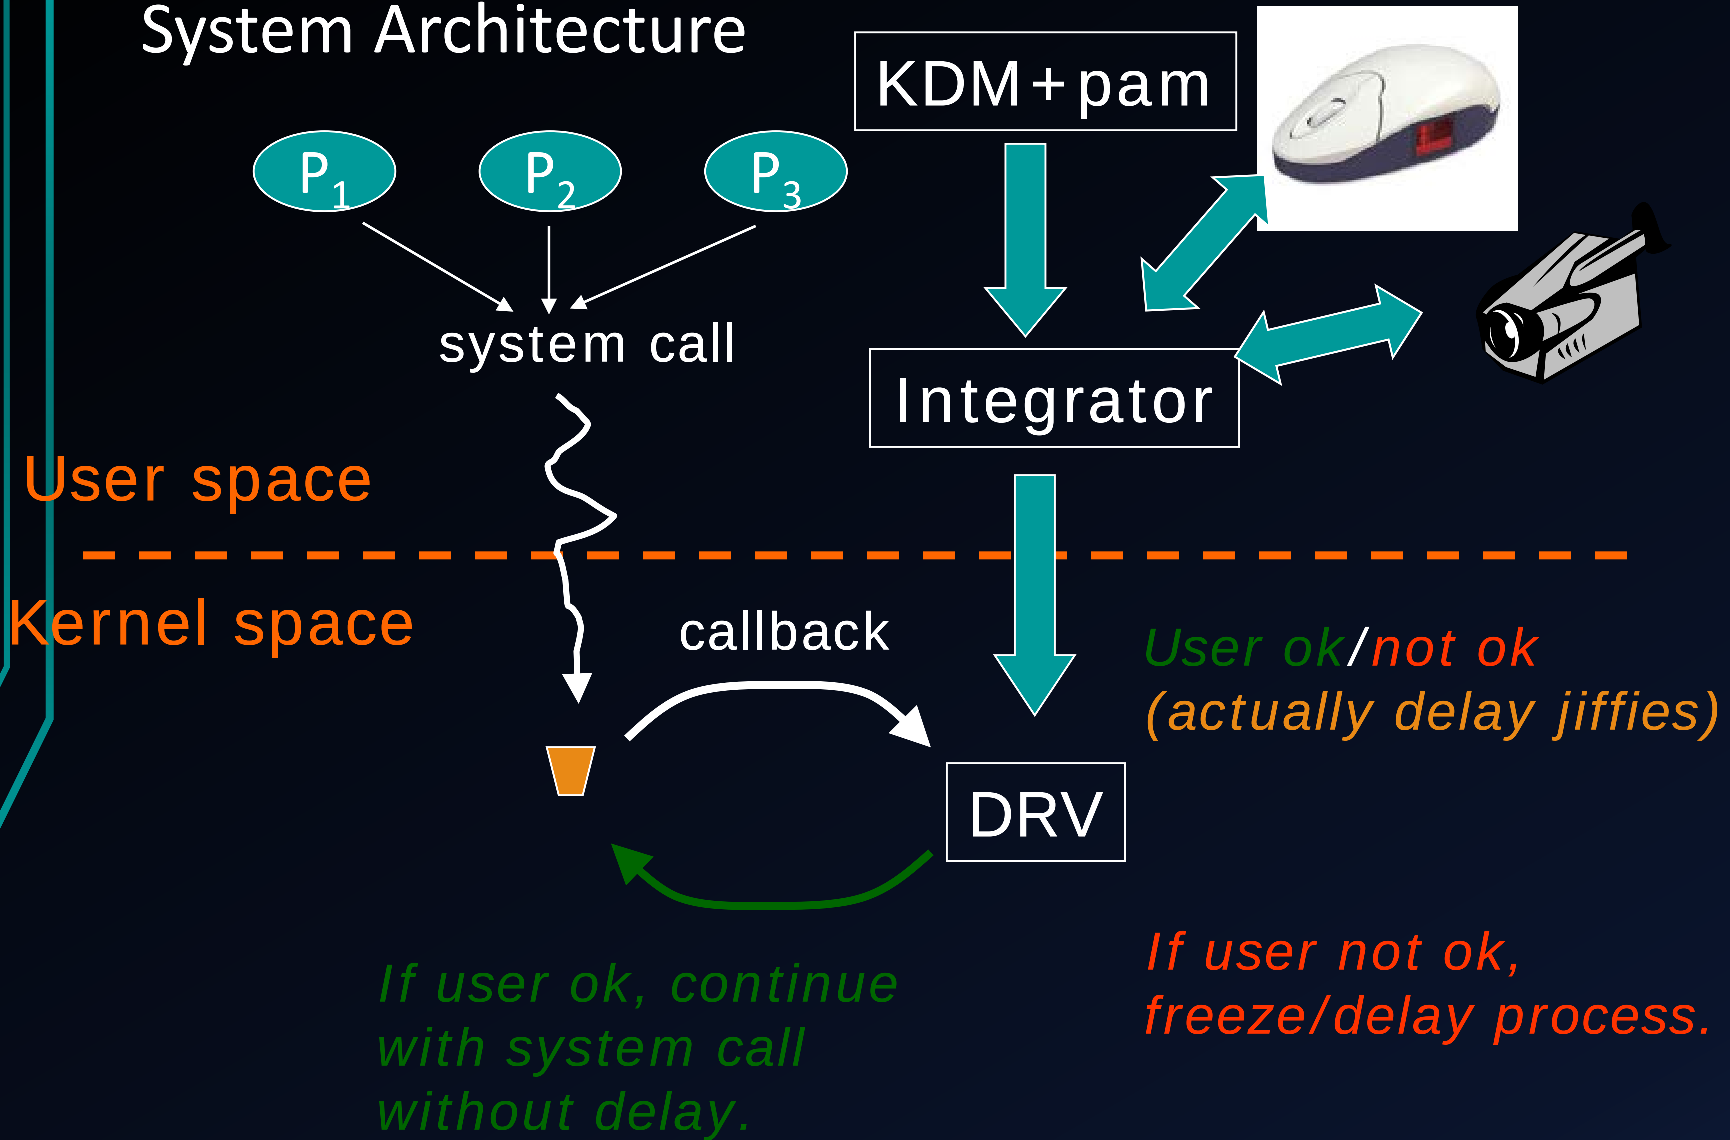

Probabilistic Approach

- The Integrator computes a probabilistic estimate of user presence, $\boldsymbol{P_{safe}}$.
- The OS is tuned with a threshold for verification, $\boldsymbol{T_{safe}}$.
  - If $\boldsymbol{P_{safe}} < \boldsymbol{T_{safe}}$, then user deemed absent.
- OS processes belonging to the user's *interactive* session are *suspended* or *delayed* as a function of

  (P_{safe} - T_{safe}, syscall)

Hidden Markov Model

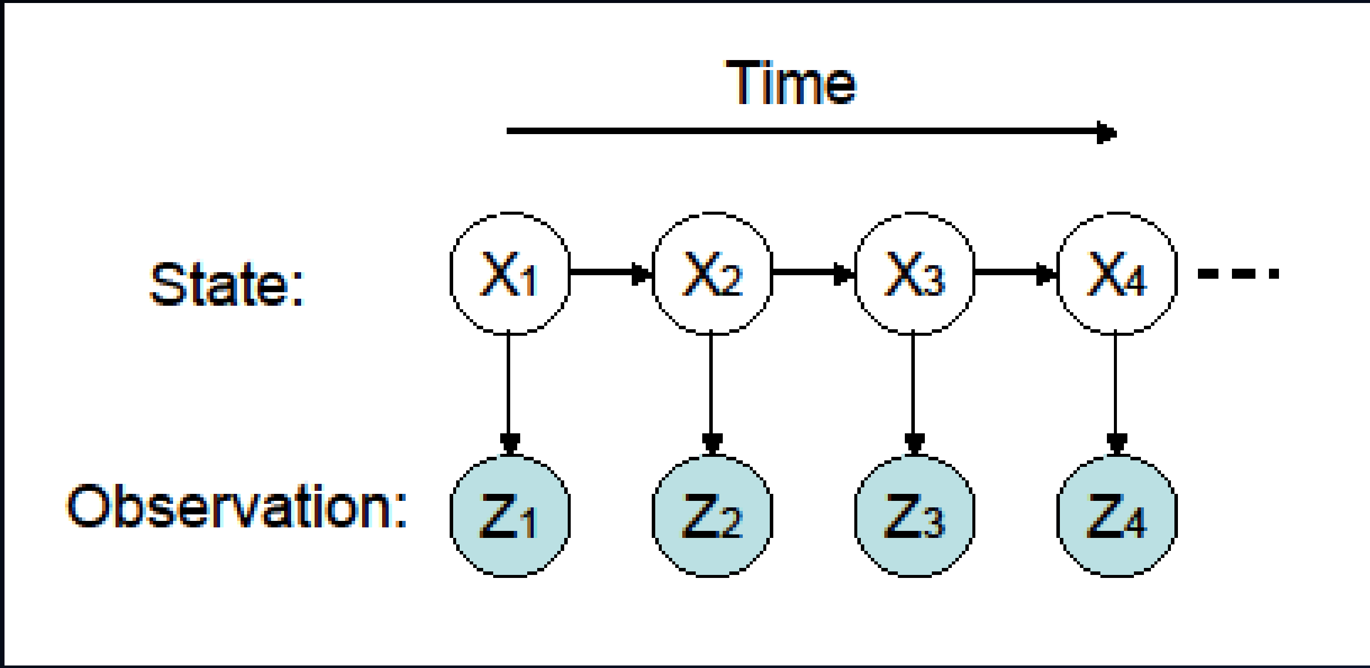

HMM States

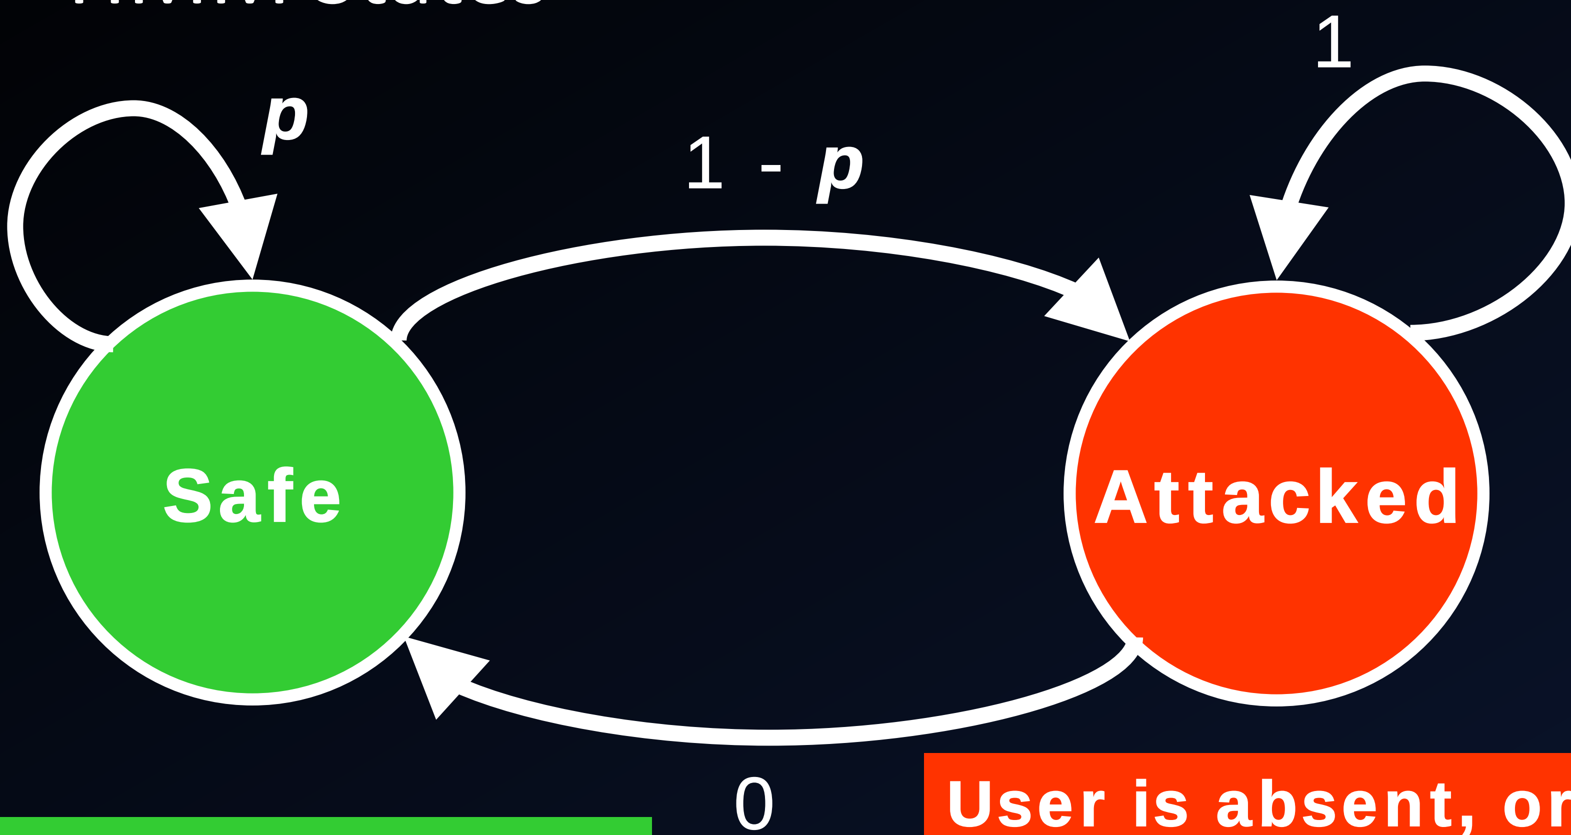

User still present at console.

User is absent, or Imposter has hijacked console.

***p*: prob. of remaining in *Safe* state at next time instant.**

Bayesian Inference

- Let $\boldsymbol{z}_t$ be a biometric observation (face or fingerprint) at time t.
- Let $\boldsymbol{x}_t$ be the state at time t.
- Given the current and past observations, what is the most likely current state?
- Bayesian inference: select the larger of

$P(\boldsymbol{x}_t\text{=Safe} \mid \boldsymbol{z}_1, \boldsymbol{z}_2, \ldots \boldsymbol{z}_t)$ and

$P(x_t\text{=Attacked} \mid z_1, z_2, \ldots z_t)$

Bayesian Inference

- $P(\boldsymbol{x}_t \mid \boldsymbol{z}_1, \ldots, \boldsymbol{z}_t)$ is efficiently computed in terms of
- $P(\boldsymbol{z}_t \mid \boldsymbol{x}_t)$: prob. of getting current observation given current state
- $P(\boldsymbol{x}_t \mid \boldsymbol{x}_{t-1})$: transition probabilities
- $P(\boldsymbol{x}_{t-1} \mid \boldsymbol{z}_1, \ldots, \boldsymbol{z}_{t-1})$: previous state given previous observations (recursion)
- Upon initial login,
  - $t=0$, and $P(\boldsymbol{x}_0=\text{Safe}) = 1$

Face Biometric

- We use a Bayesian classifier.
- From 500 training face images of legitimate user, and 1200 images of other people (imposter), we learn:

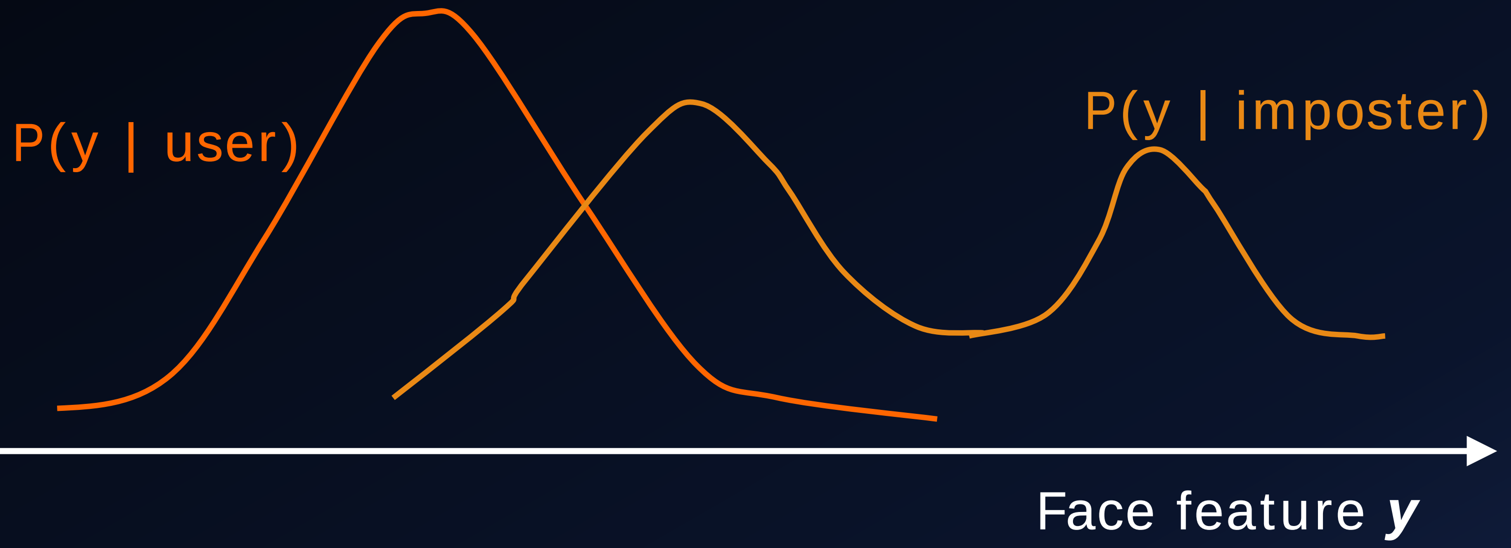

Face Biometric

- Note that
  - $P(\boldsymbol{z}_t \mid \boldsymbol{x}_t = \text{Safe})$ is just P(y | user)
  - $P(\boldsymbol{z}_t \mid \boldsymbol{x}_t = \text{Attacked})$ is just P(y | imposter)

Fingerprint Biometric

- Also Bayesian classifier.
- Vendor's proprietary algorithm matches 2 fingerprint images.
  - Outputs a matching score, *s*
- From training images, we learn:
  - P(*s* | user) and P(*s* | imposter)
- Which become
  - $P(\boldsymbol{z}_t \mid \boldsymbol{x}_t = \text{Safe})$ and $P(\boldsymbol{z}_t \mid \boldsymbol{x}_t = \text{Attacked})$ respectively

Further Comments

- $\boldsymbol{P}_{safe}$ = P($\boldsymbol{x}_t$=Safe | $\boldsymbol{z}_1$, …, $\boldsymbol{z}_t$)

- We can compute $\boldsymbol{P}_{safe}$ anytime.
  - If no observation at time *t*, then use most recent observation: $\boldsymbol{P}_{safe}$ = P($\boldsymbol{x}_t$=Safe | $\boldsymbol{z}_1$, …, $\boldsymbol{z}_{t-1}$)
  - But decay transition probability ***p*** by time lapse.

    $$\boldsymbol{p} = e^{k\Delta t}$$

  - This reflects increasing uncertainty about presence of user when no observations available.

Further Comments

- In theory, we want the larger of

  P($\boldsymbol{x}_t$=Safe | $\boldsymbol{z}_1, \ldots, \boldsymbol{z}_t$) and P($\boldsymbol{x}_t$=Attacked | $\boldsymbol{z}_1, \ldots, \boldsymbol{z}_t$)
- Equivalent to: $\boldsymbol{P}_{safe} > 0.5$
- But in practice, we use $\boldsymbol{P}_{safe} > \boldsymbol{T}_{safe}$
  - More flexible: different $\boldsymbol{T}_{safe}$ for different process actions (e.g. reads vs. writes)
  - Avoids "close call" cases when both probabilities almost equal.
- Math details in paper.

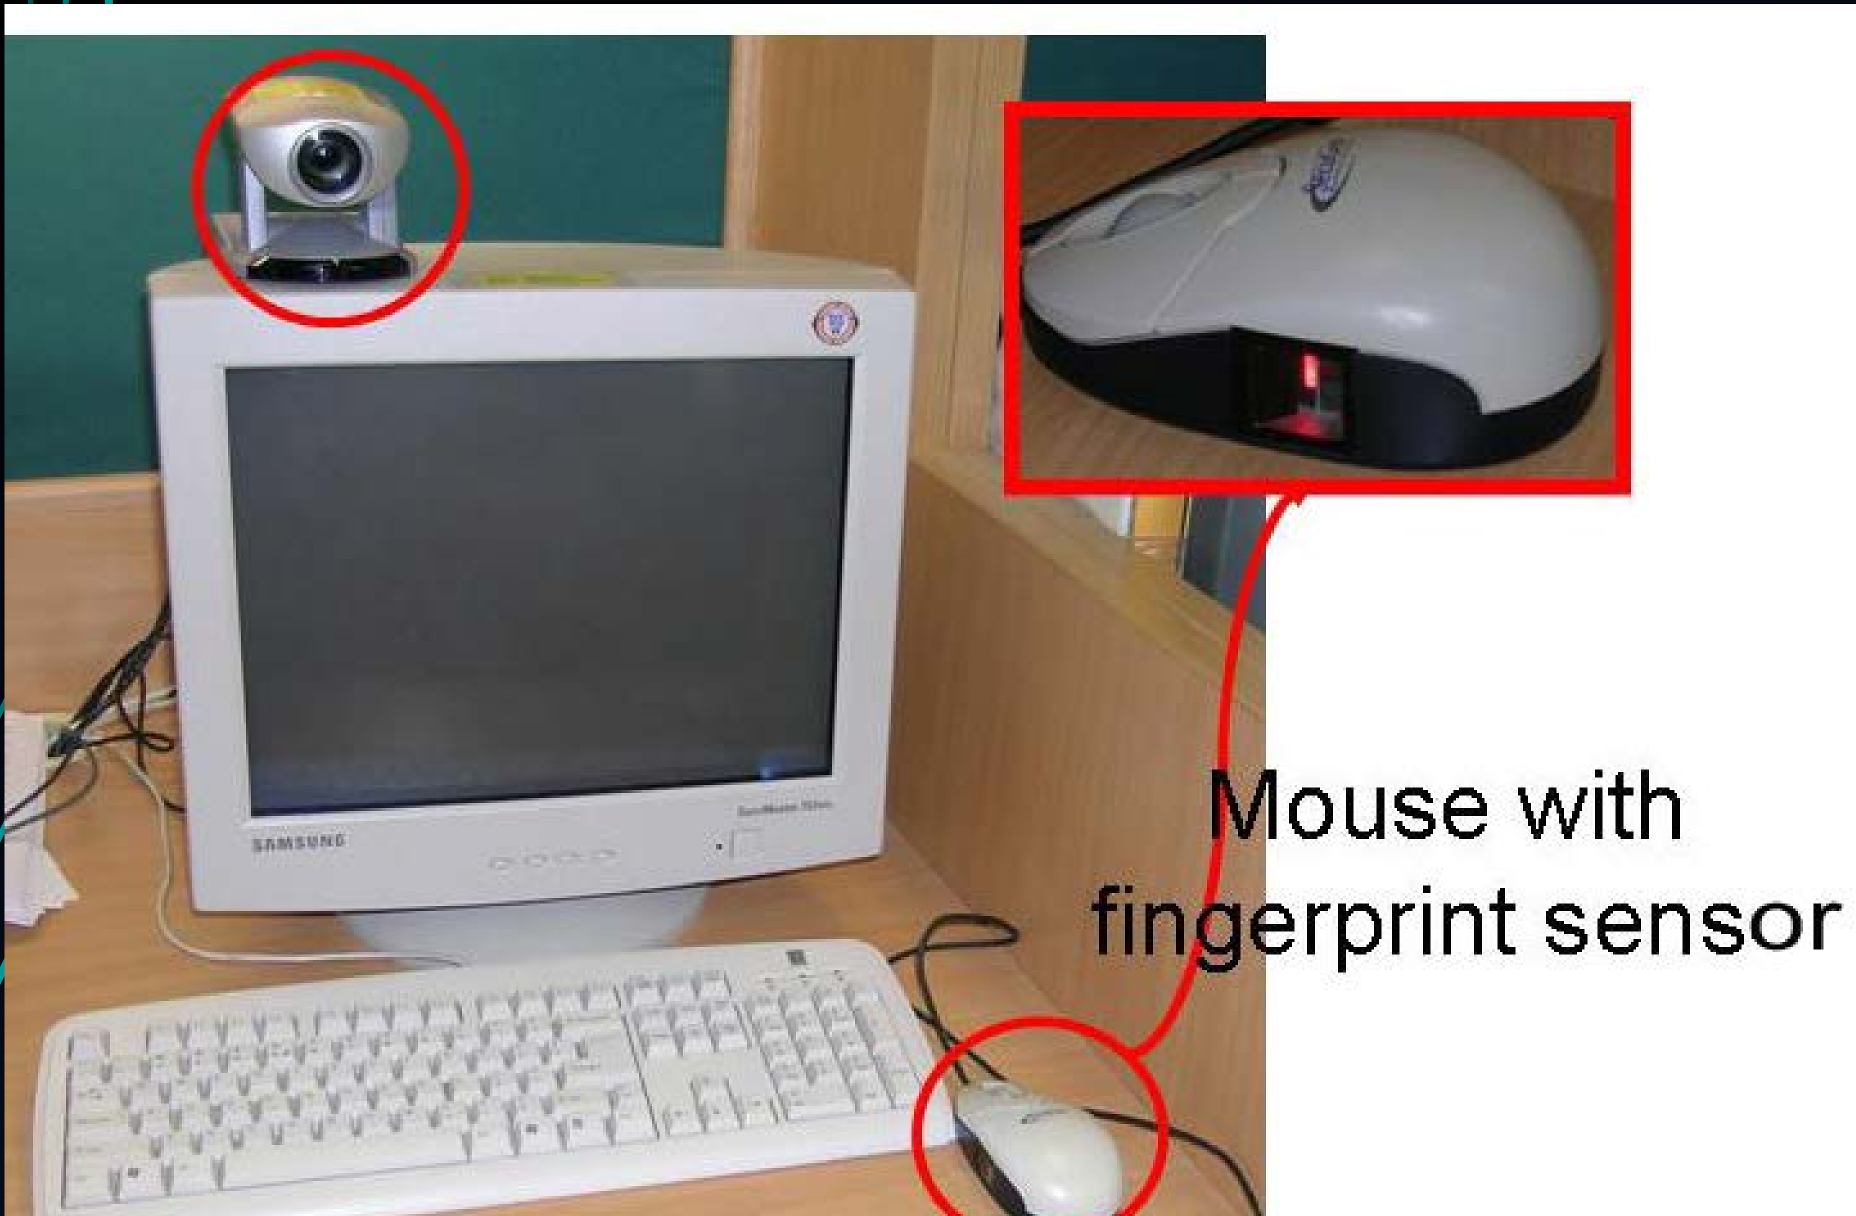
Mouse with
fingerprint sensor
SAMSUNG

Evaluation

Other Fusion Methods

Temporal-first

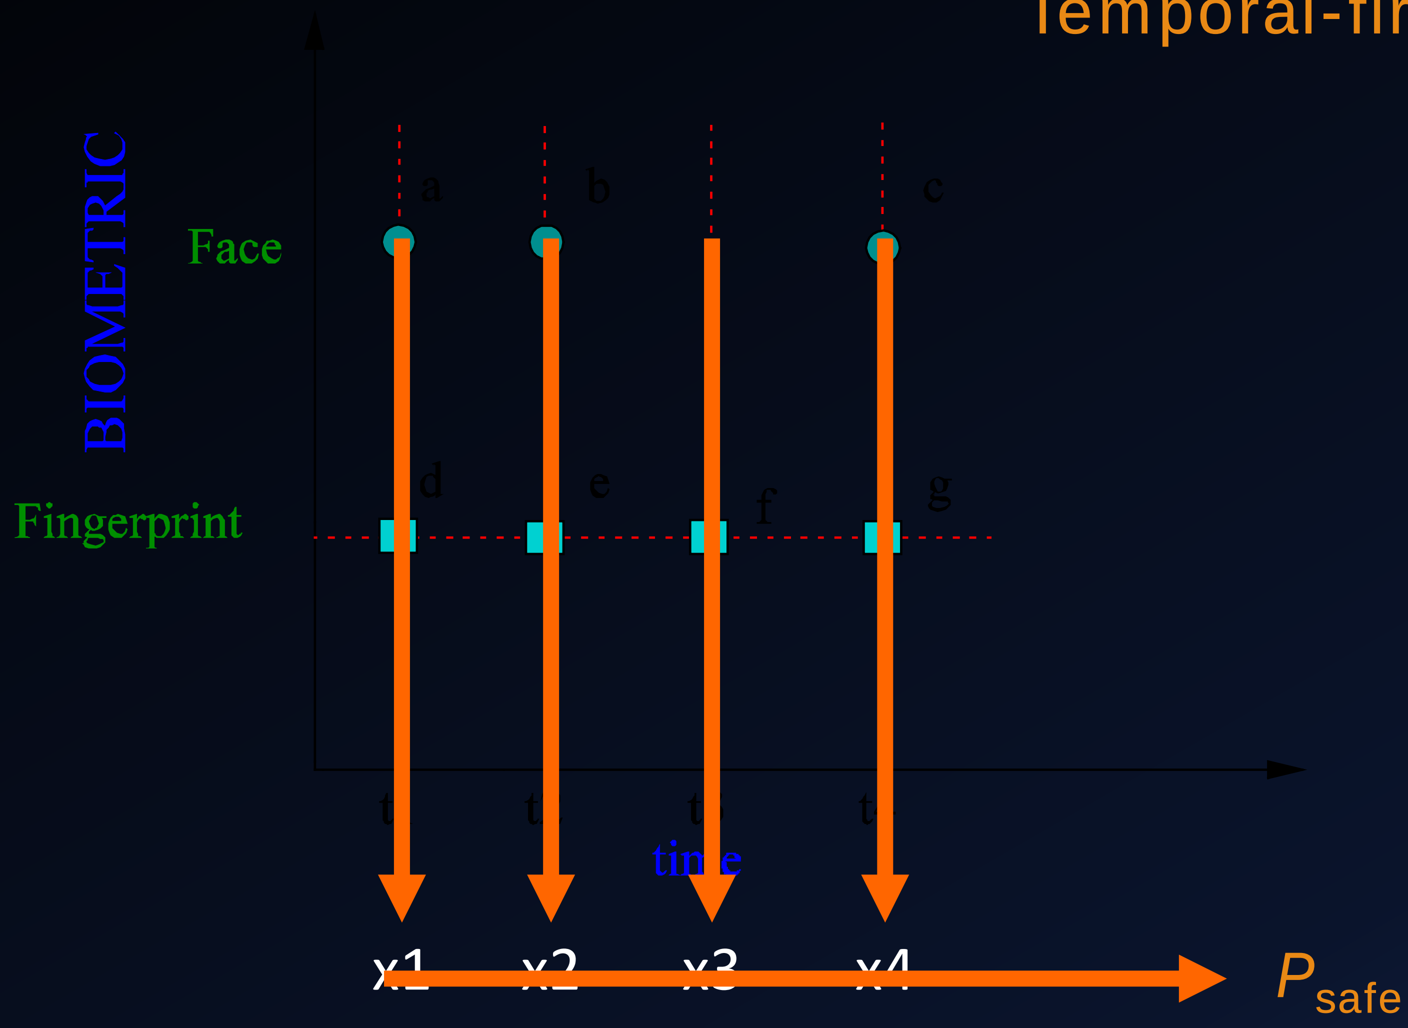

Other Fusion Methods

Modality-first

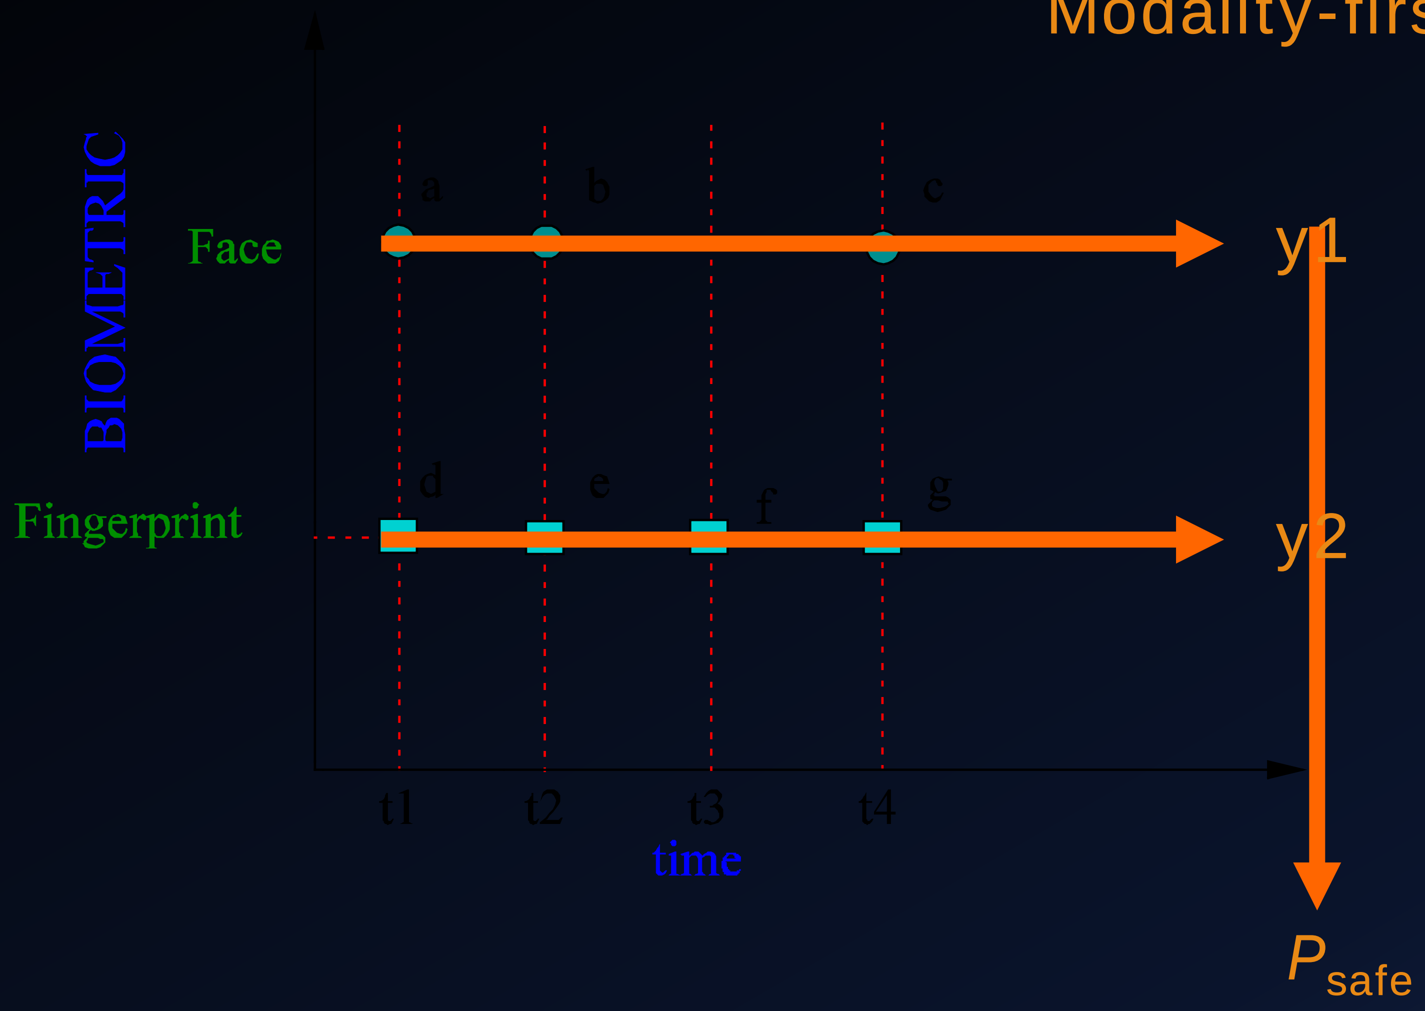

Naïve Integration

- Idea: use the most reliable modality available at any time instant.
- Since fingerprint more reliable than face, use it whenever available.
- Else use face.
- If no modality available, use the previous one, but decay it appropriately.

Reliability

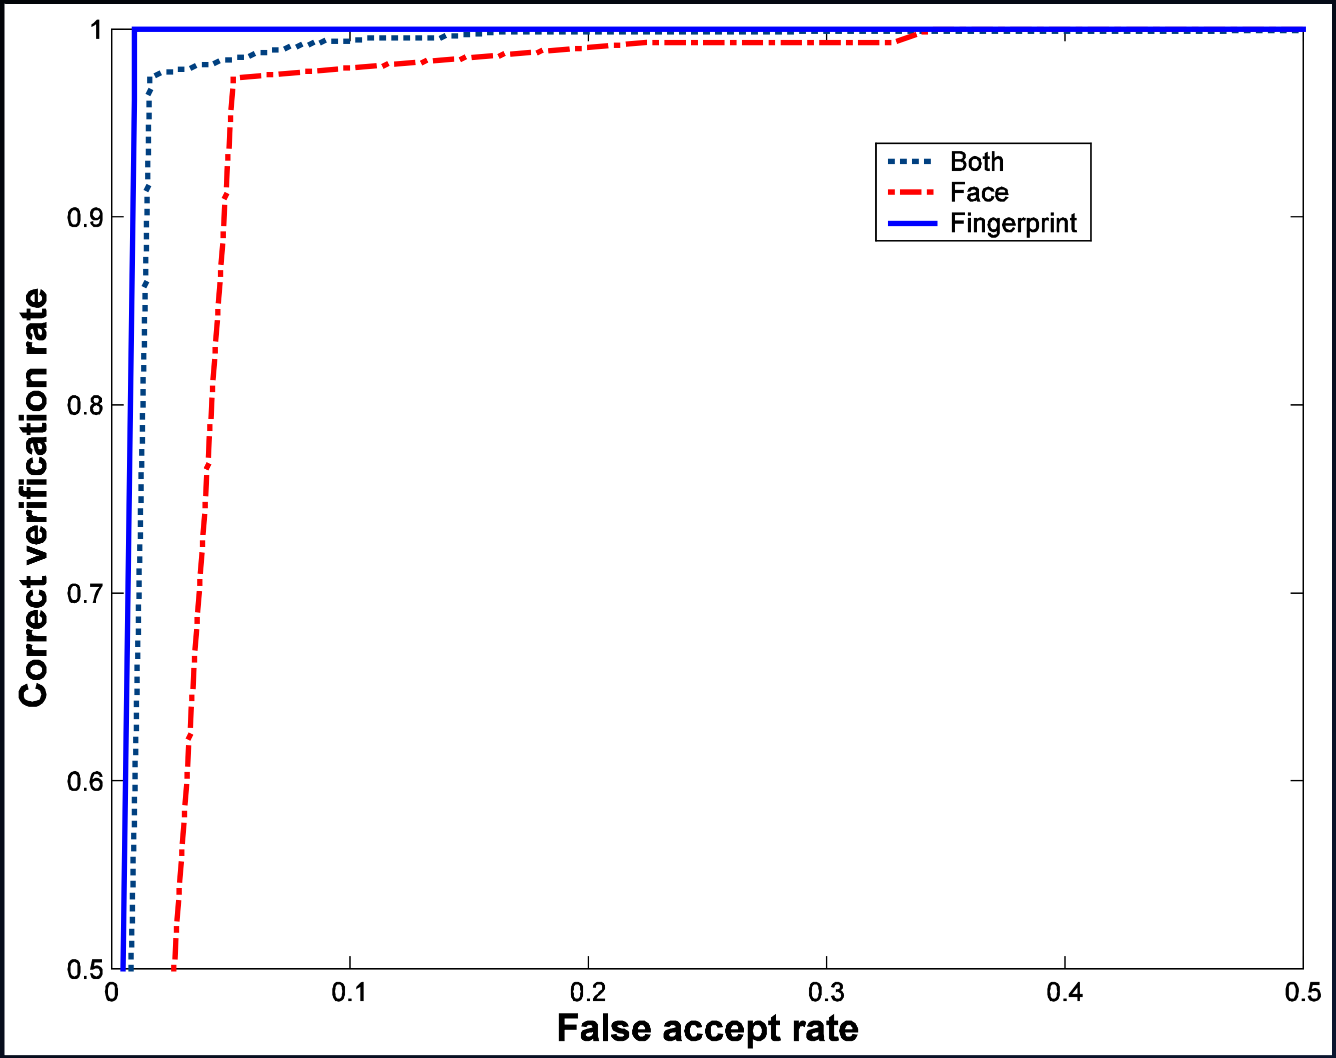

Experiment: Legitimate User

- Indiv. Probabilities sporadic ➔ significant FAR/FRR for any threshold T_{safe}
- FAR = security breach!
- FRR = inconvenience
- Holistic Fusion closest to ideal.
- Abrupt drop in Temporal-first, Modality-first curves.

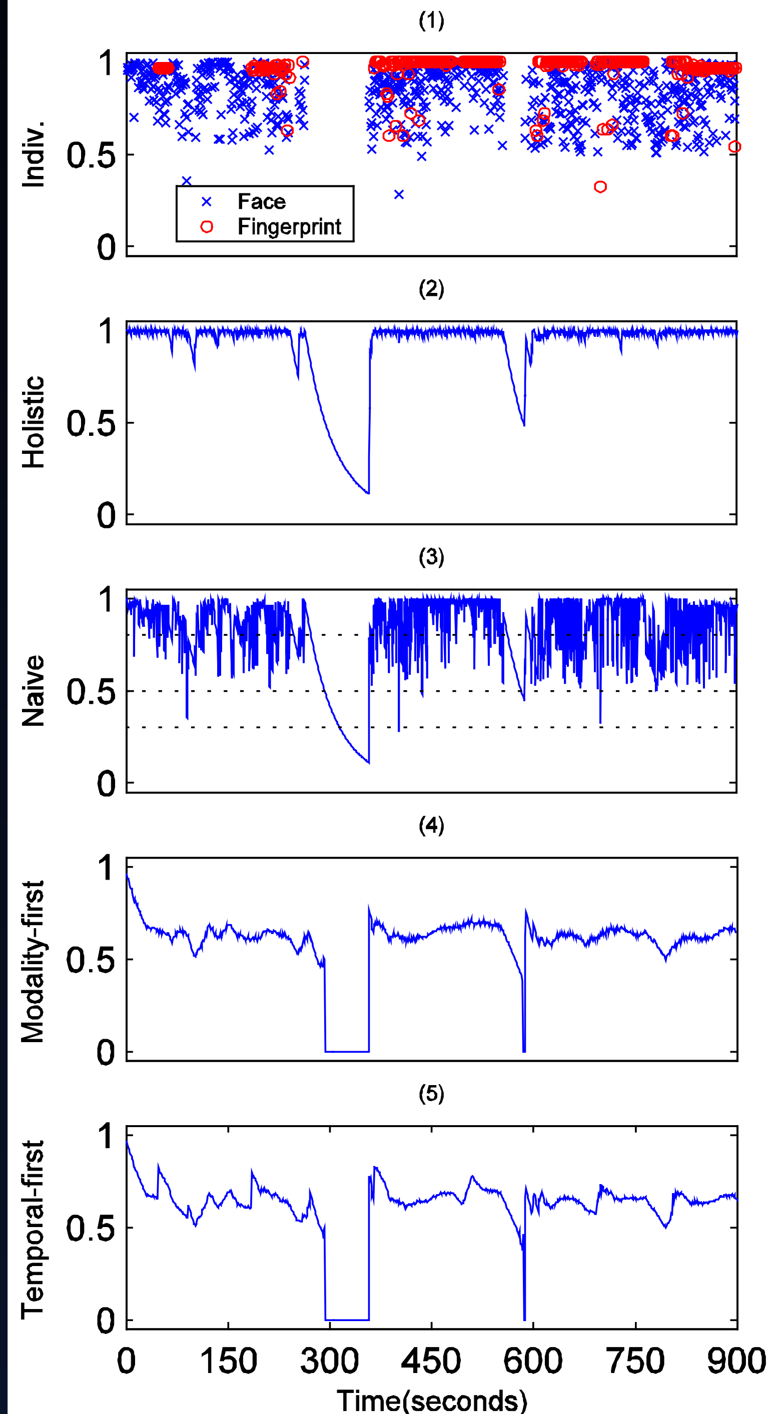

Experiment: Imposter

- Imposter hijacks session at time = 38s
- Detect by change in slope.
- Holistic Fusion and Naïve Integration detects hijacking sooner than others (time = 43s).

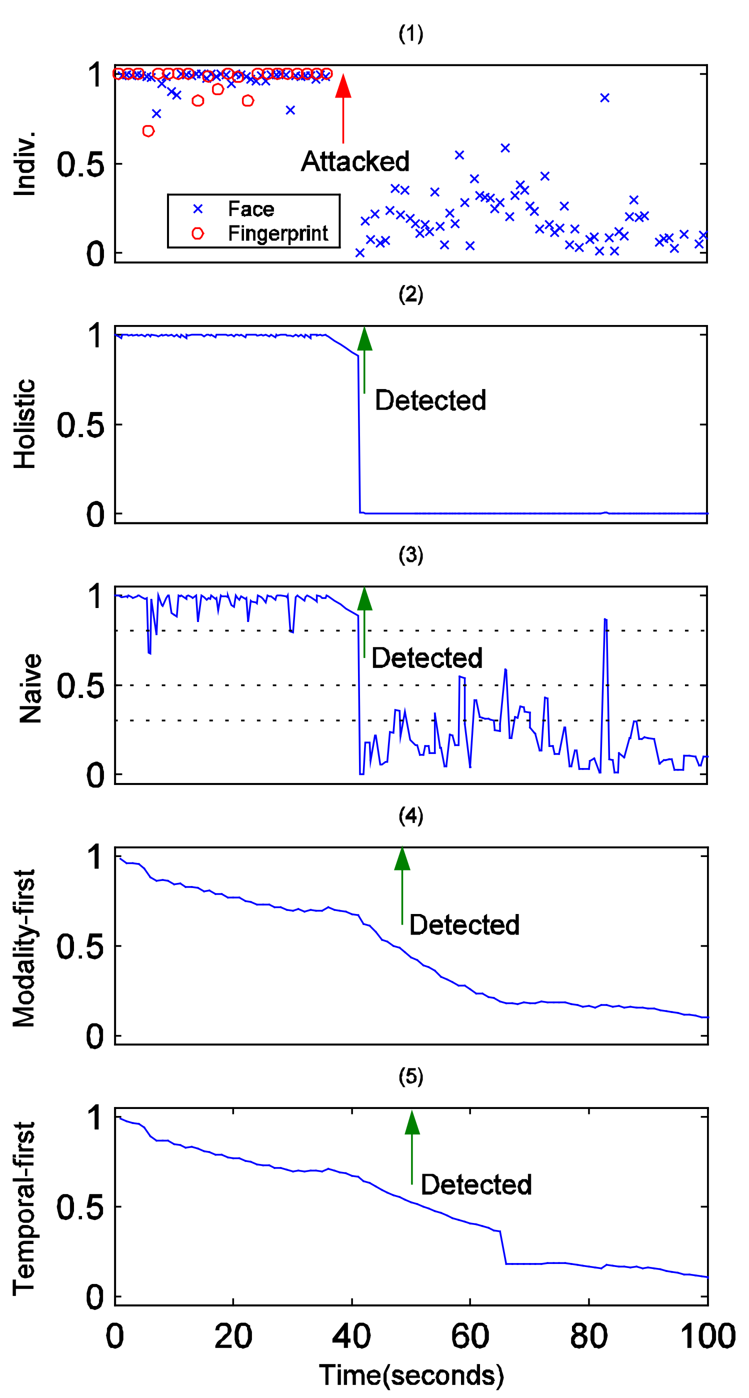

Experiment: Partial Impersonation

- Successfully faked fingerprint, but not face.
- This is easily detected by Holistic and Naïve, but not by others.

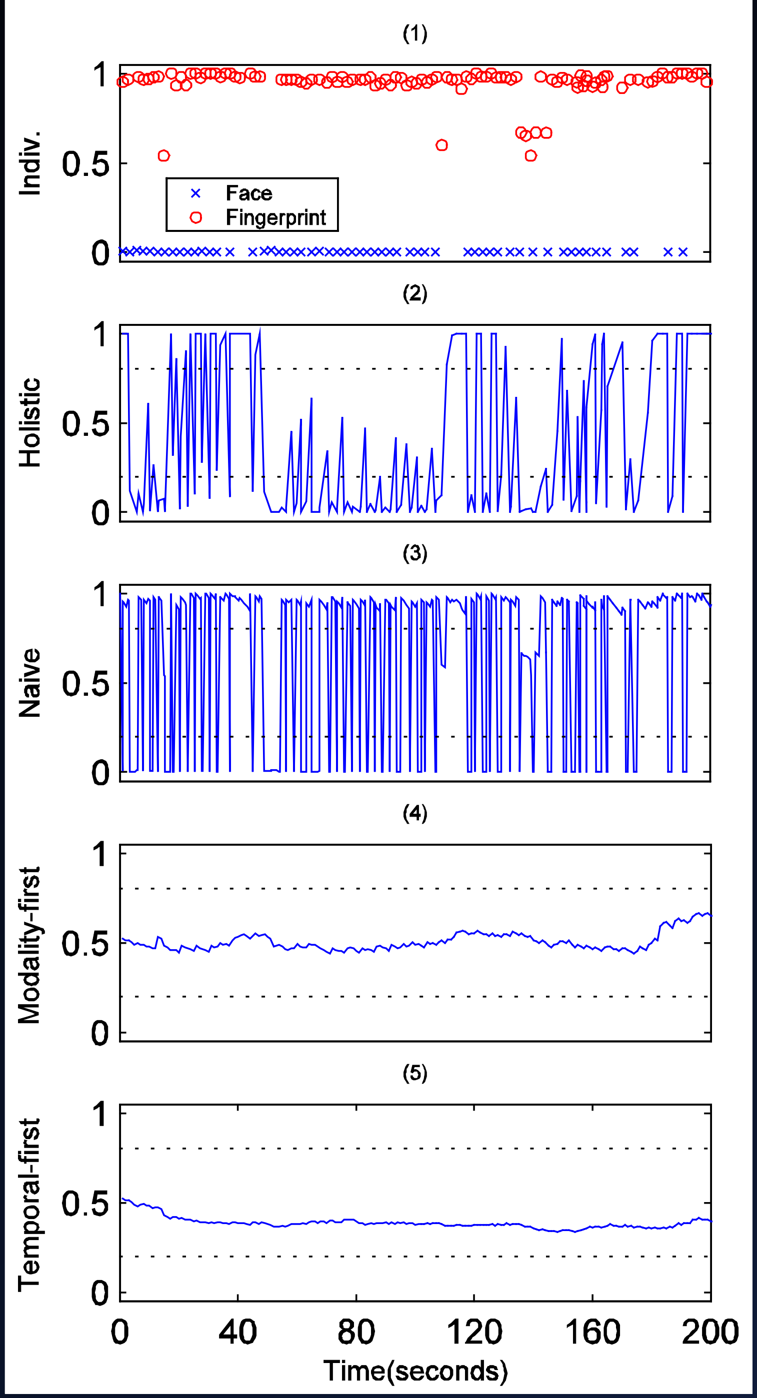

P_{safe} for different tasks

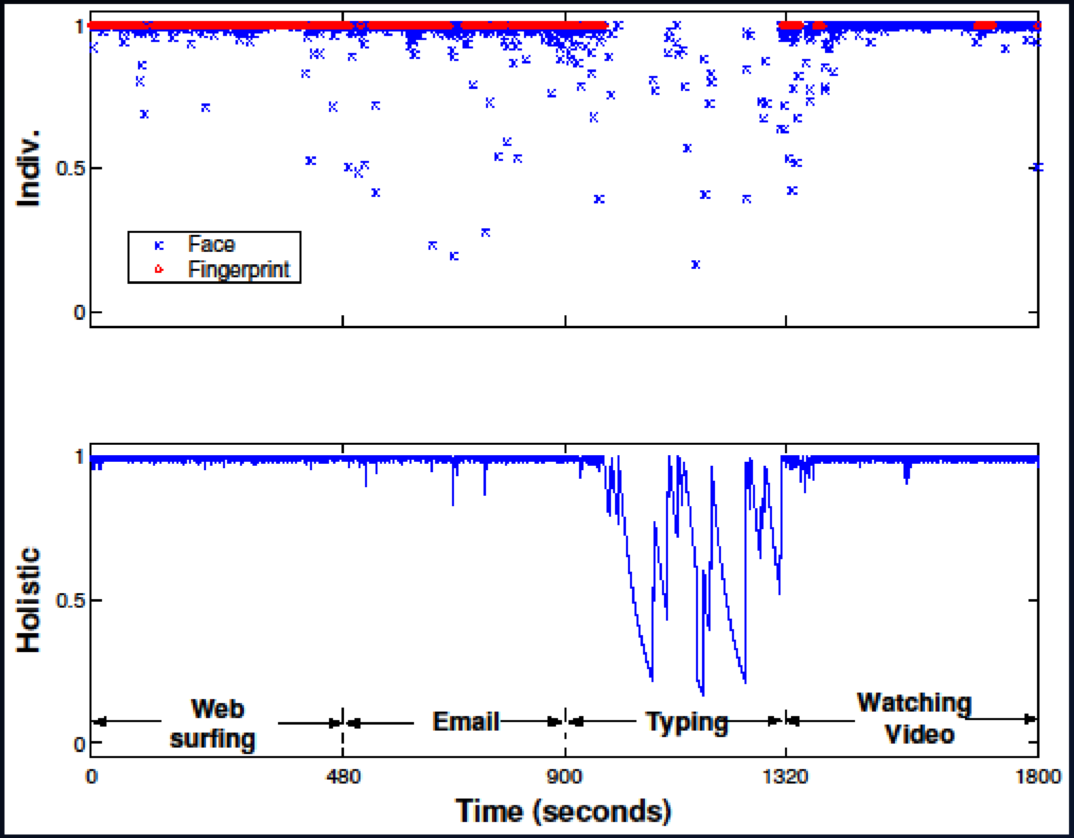

Usability test

- 58 people to perform different tasks

Usability test

- CBAS verifies users at a **low** FRR, and **low** FAR.
- Surprising result: (a) **no statistical evidence** to show that CBAS overhead affects task efficiency; (b) system performance degradation was imperceptible by users.
- Many users felt uncomfortable being "watched" by webcam. Discreet placement may solve this.
- A biometric solution for **continuous authentication** is **practical** and **usable**.
- Multi-core processors will further reduce the overhead.

New Performance Metric

- Time to Correct Reject (TCR)
- The interval between the start of the first action taken by the imposter to the time instant that the system decides to (correctly) reject him.
- Ideally, TCR = 0.
  - Practically, TCR < W (minimum time for the imposter to damage the system, eg. To type “rm –rf *”)
  - As long as TCR < W, system integrity is assured

New Performance Metric

- Probability of Time to Correct Reject (PTCR)
- The probability that TCR is less than W
- Ideally, PTCR = 1.
  - Practically, PTCR < 1 may be tolerable
  - This means that sometimes, the system can take longer than W seconds to correctly reject an imposter.
  - If system always fails to correctly reject, then PTCR = 0 for all W
  - PTCR is analogous to FAR

New Performance Metric

- Usability
- the fraction of the total time that the user is granted access to the protected resource
  - eg. User logs in for a total duration of T, but system sometimes rejects user
  - Let t be the total time user is accepted
  - Then Usability = t / T
- Ideally, Usability = 1.
  - Usability is analogous to FRR

New Performance Metric

- Usability-Security Characteristic Curve (USC)
- Plot of Usability vs PTCR
- Analogous to ROC curve

USC curve for our system

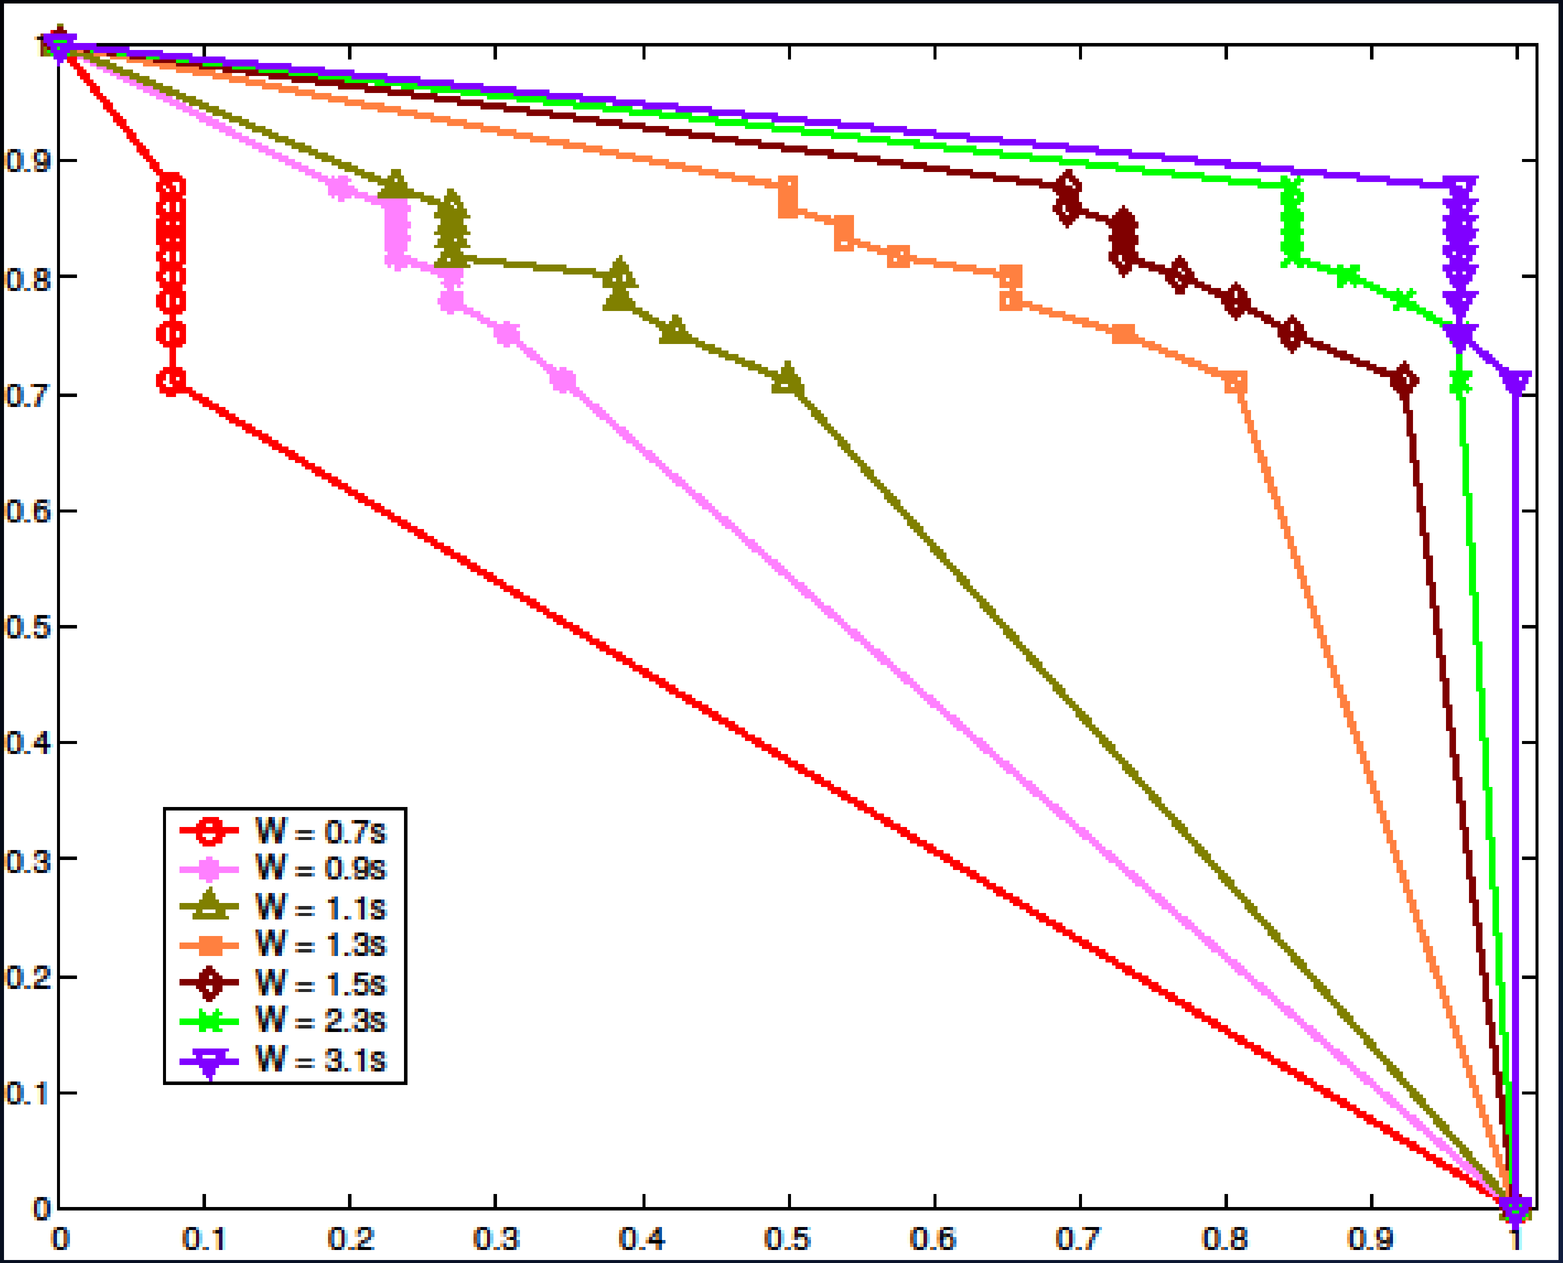

Evaluation


GROWERS
SWEET 100
VINE TOMATO

Soft Biometric Traits for Continuous User Authentication

Koichiro Niinuma, Unsang Park, *Member, IEEE*, and Anil K. Jain, *Fellow, IEEE*

Soft biometrics: Definition

- those characteristics that provide some information about the individual, but lack the distinctiveness and permanence to sufficiently differentiate any two individuals under normal circumstance
  - e.g. gender, clothes color

System

- Hard biometric: face recognition (eigenface)
- Soft biometric: face color histogram, clothes color histogram

4 modes

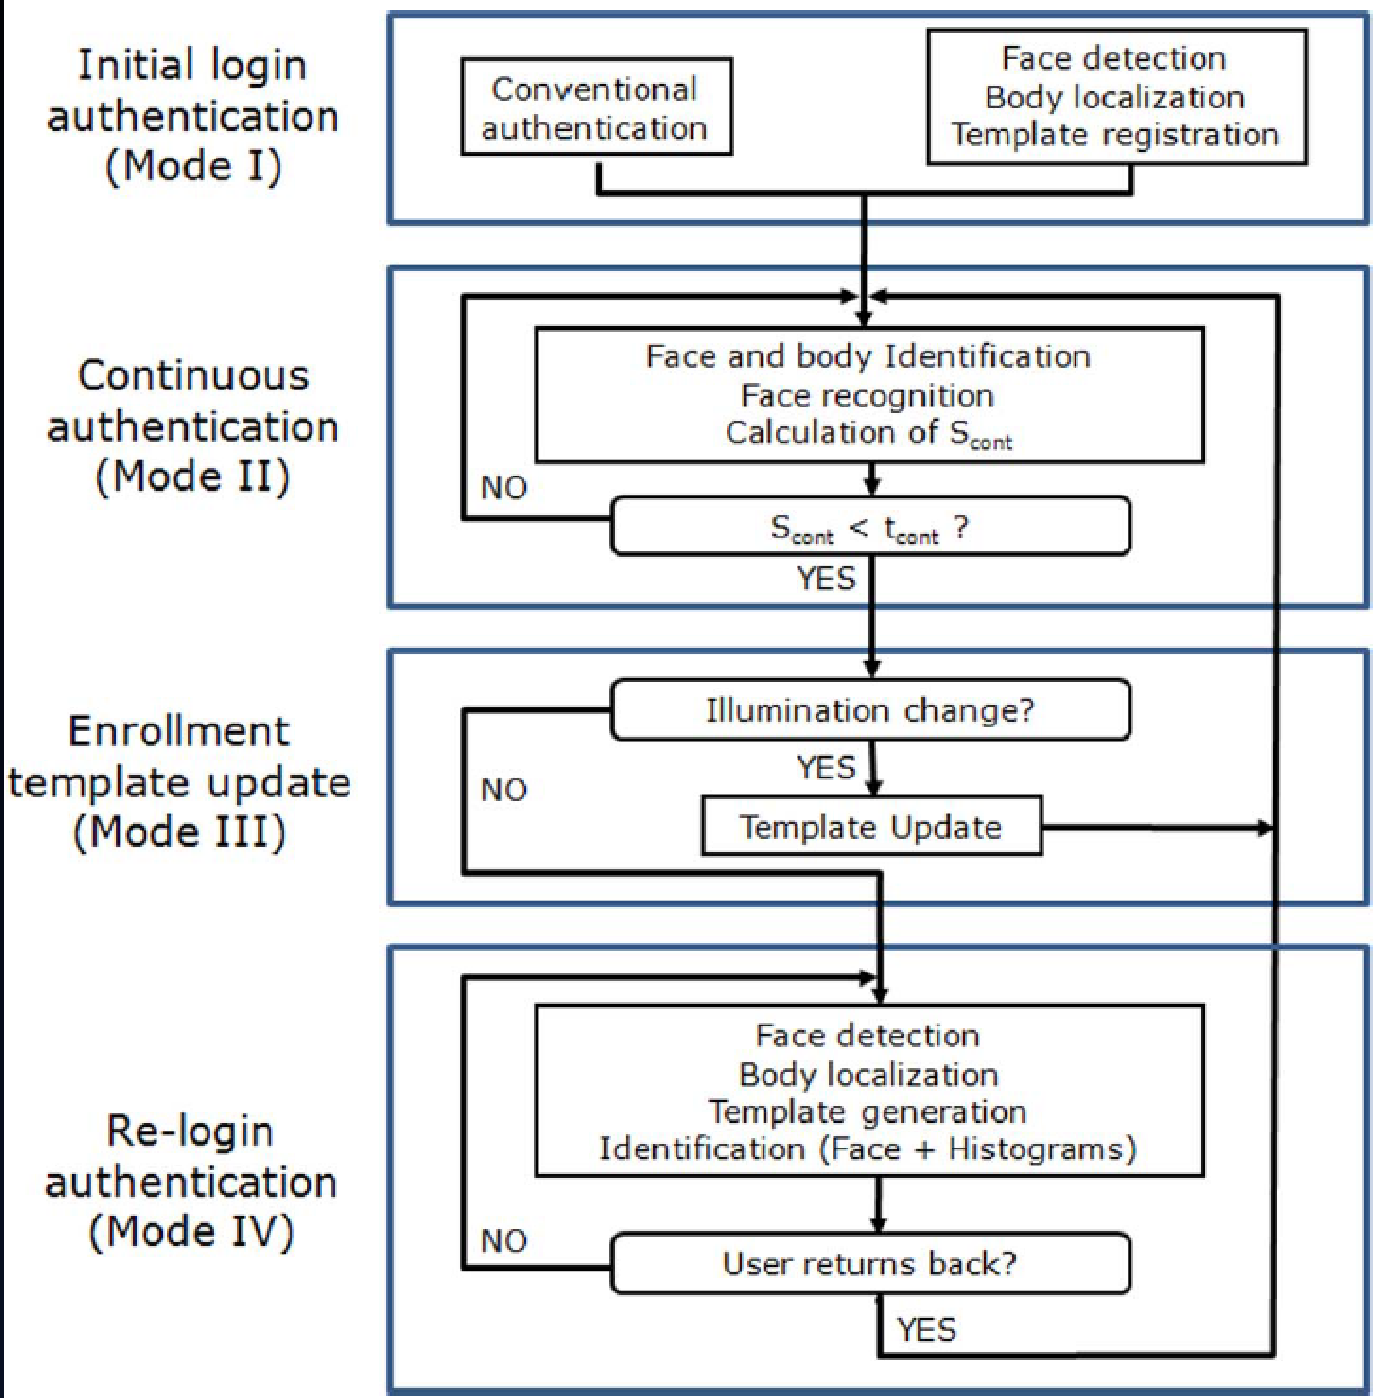

Initial login authentication (Mode I)
Conventional authentication
Face detection
Body localization
Template registration
Continuous authentication (Mode II)
Face and body Identification
Face recognition
Calculation of S_{cont}
$S_{cont} < t_{cont}$?
NO
YES
Enrollment template update (Mode III)
Illumination change?
YES
NO
Template Update
Re-login authentication (Mode IV)
Face detection
Body localization
Template generation
Identification (Face + Histograms)
User returns back?
NO
YES

Hard vs Soft biometrics

	Hard biometrics	Soft biometrics
Confidence of decision with each observation	High to medium	Medium to low
Frequency of observation	Medium to low	High
Pre-registration	Required	Not required
Ω_{intra} and Ω_{inter}	Available	Not available

Hard vs Soft biometrics

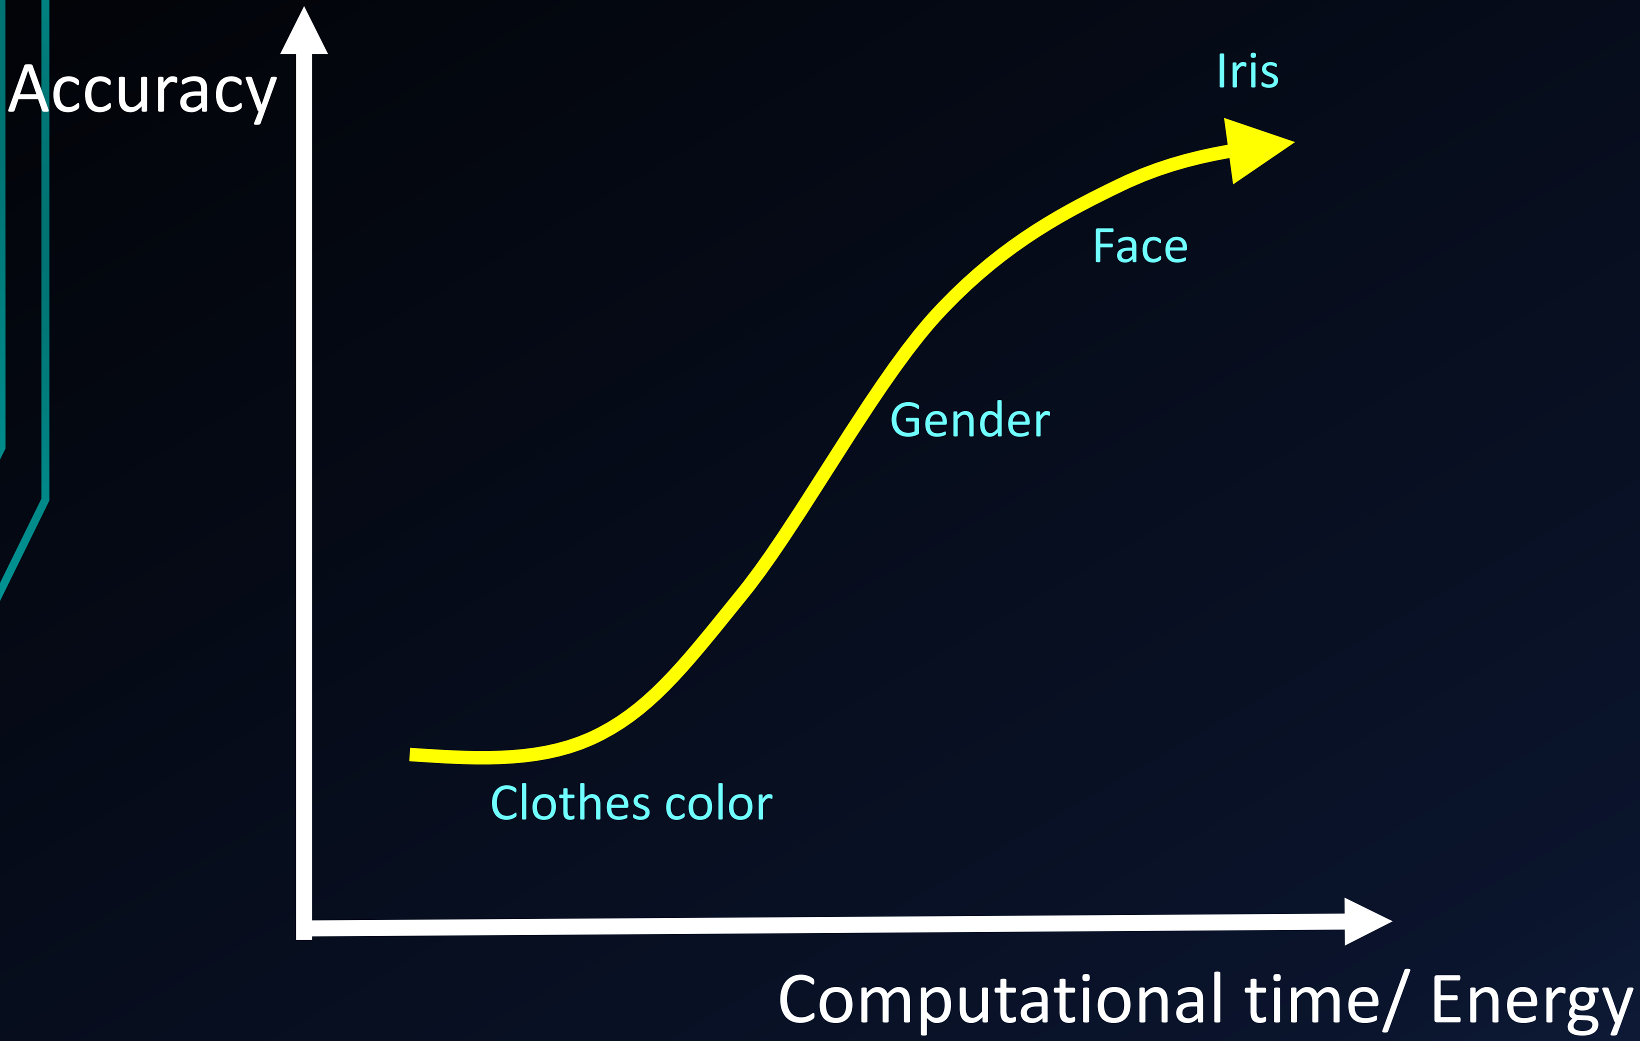

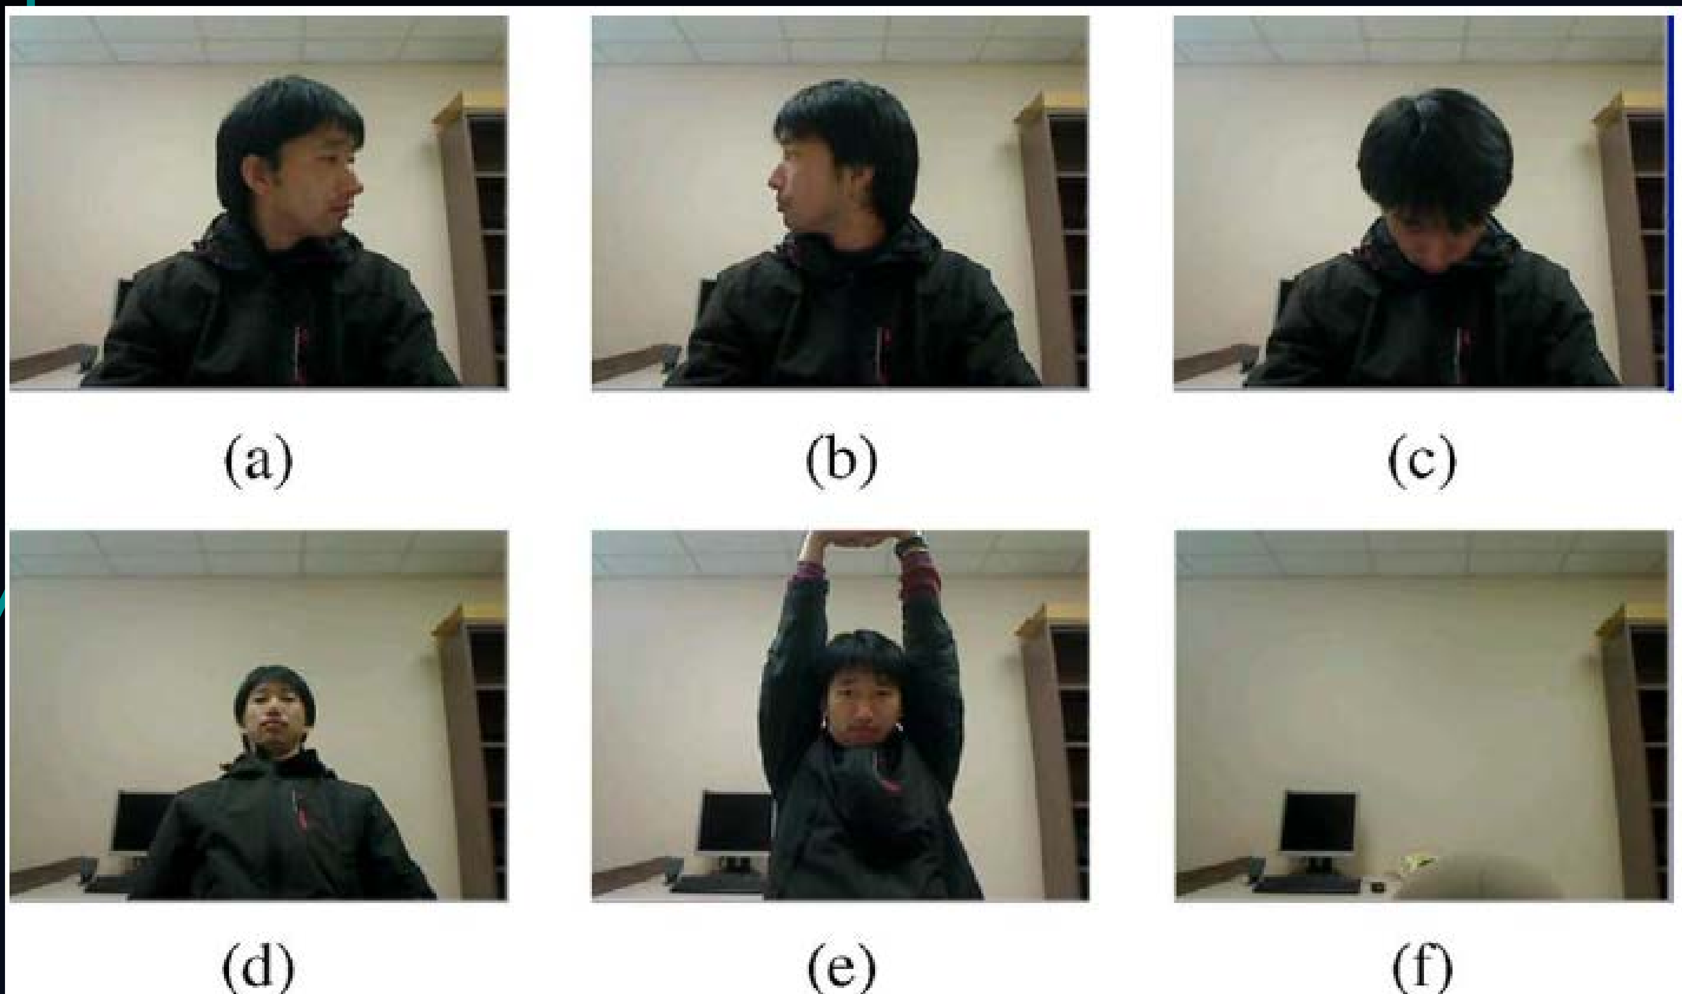
(a)
(b)
(c)
(d)
(e)
(f)

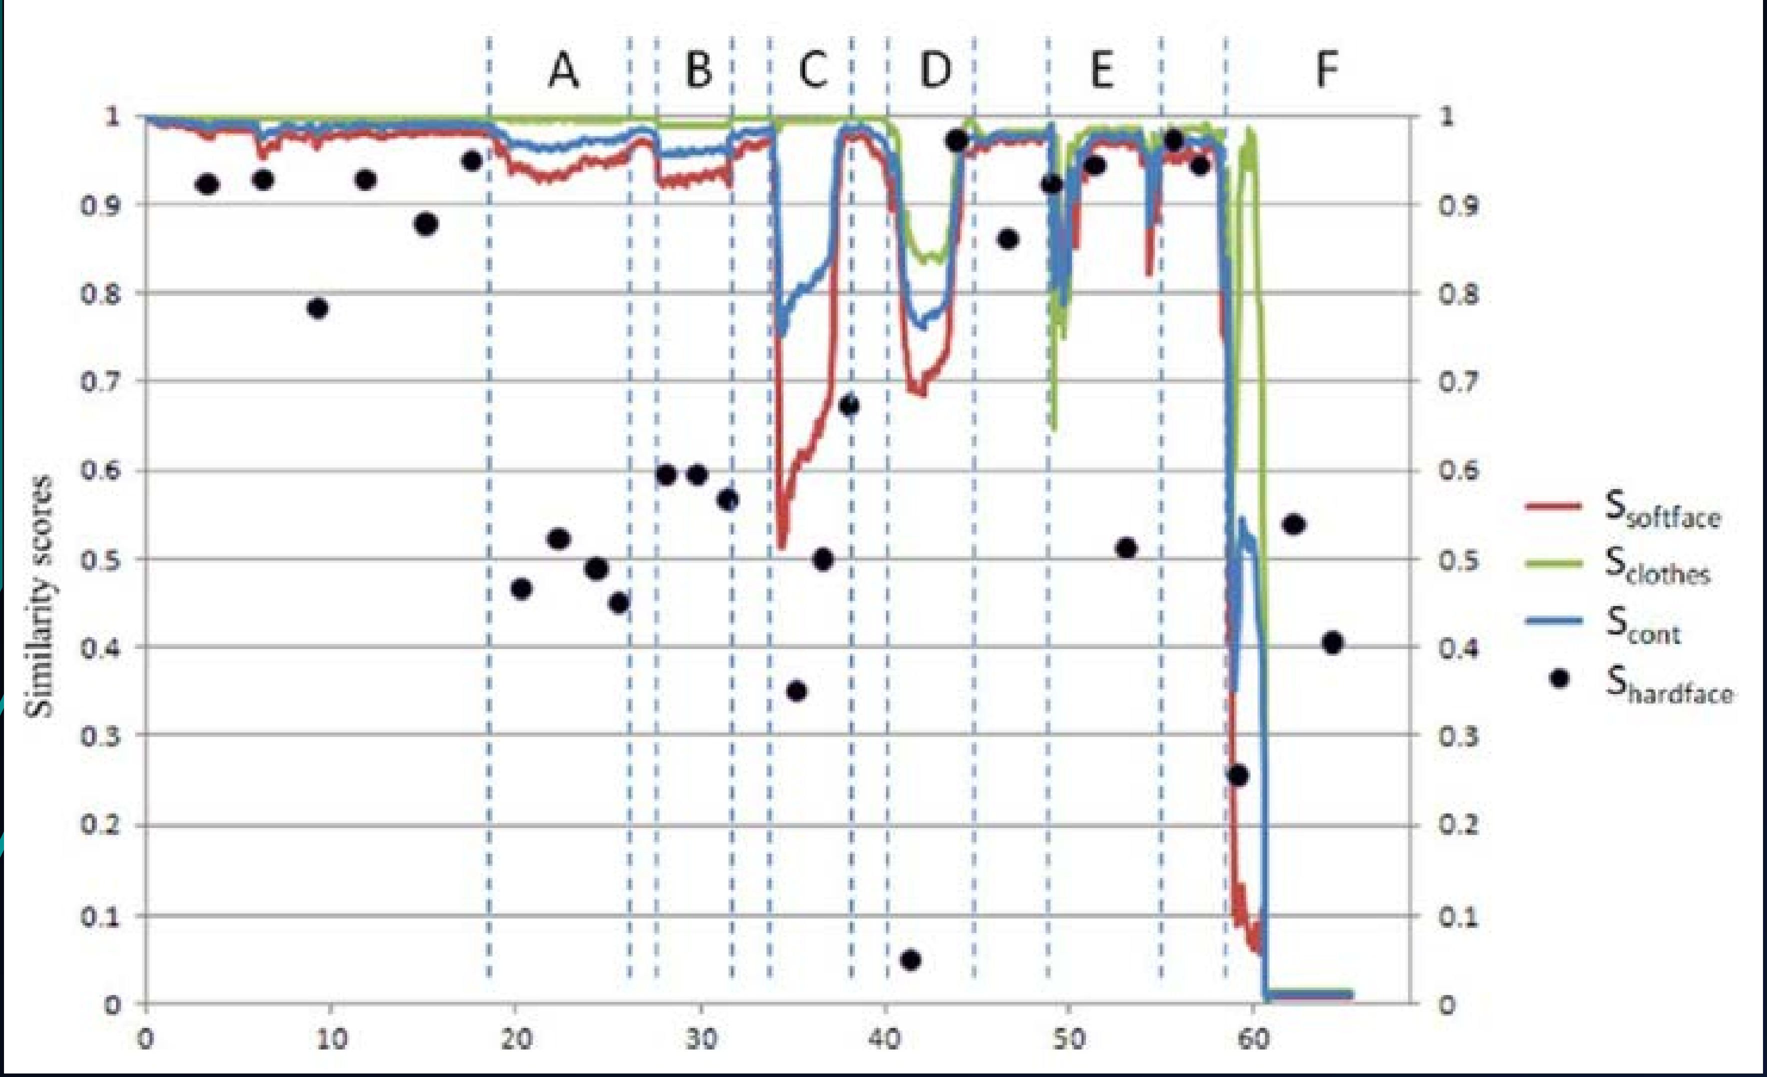

A
B
C
D
E
F
Similarity scores
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
0
10
20
30
40
50
60
$S_{softface}$
$S_{clothes}$
S_{cont}
$S_{hardface}$

Coping with illum change

(a) (b)

Coping with illum change

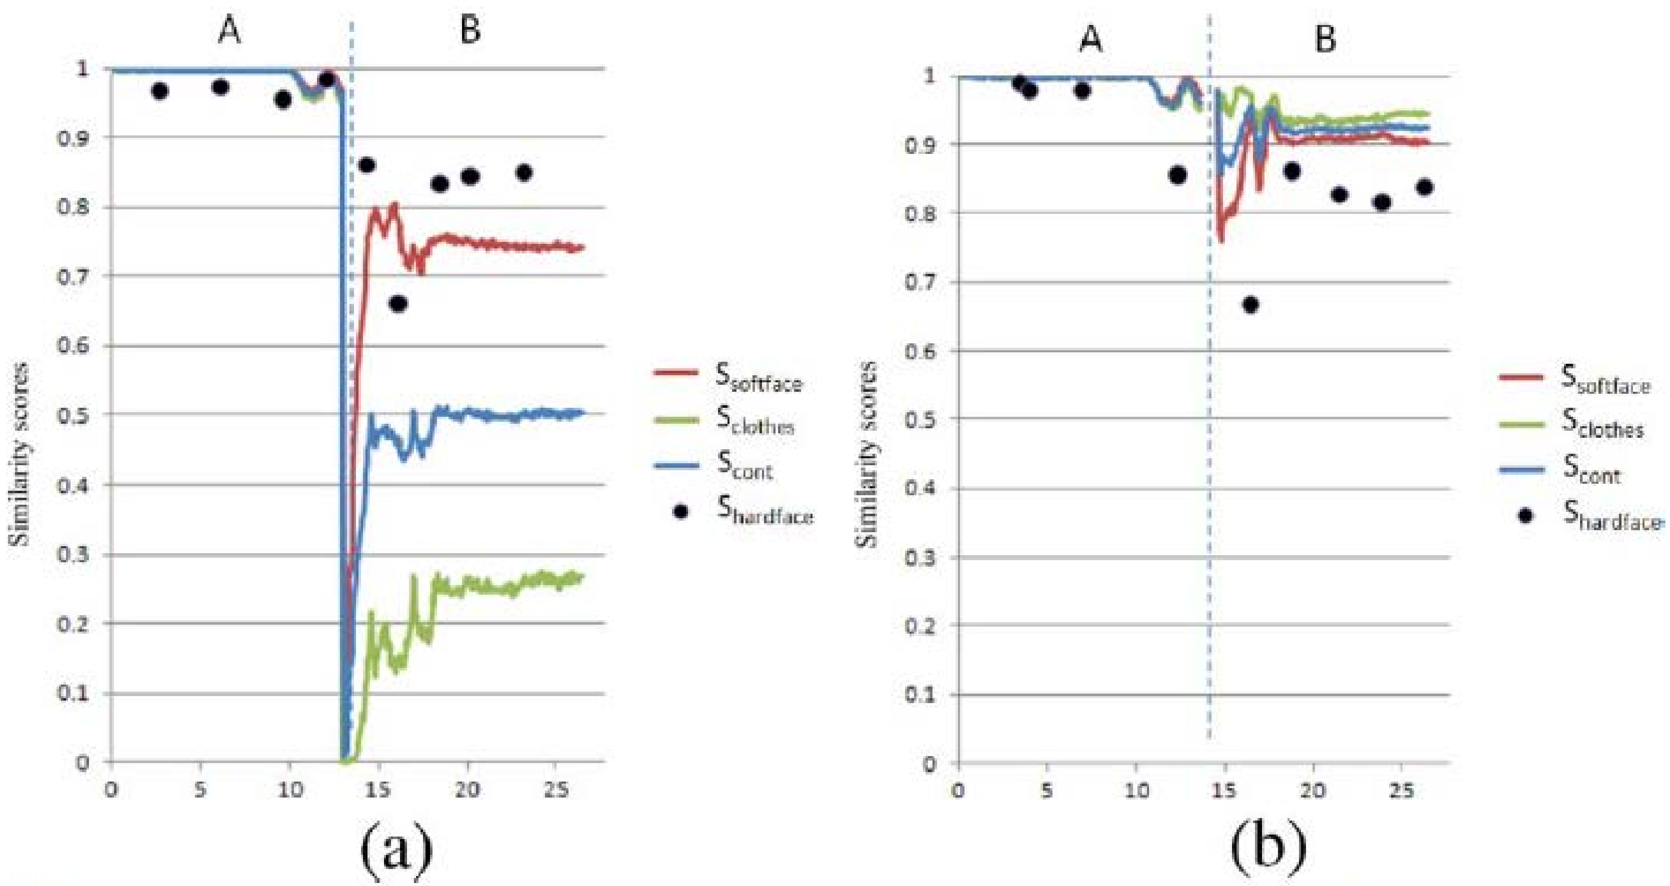

Fig. 18. Example 2 of similarity score versus time graphs with and without enrollment update. (a) Without enrollment update. (b) With enrollment update.

Evaluation

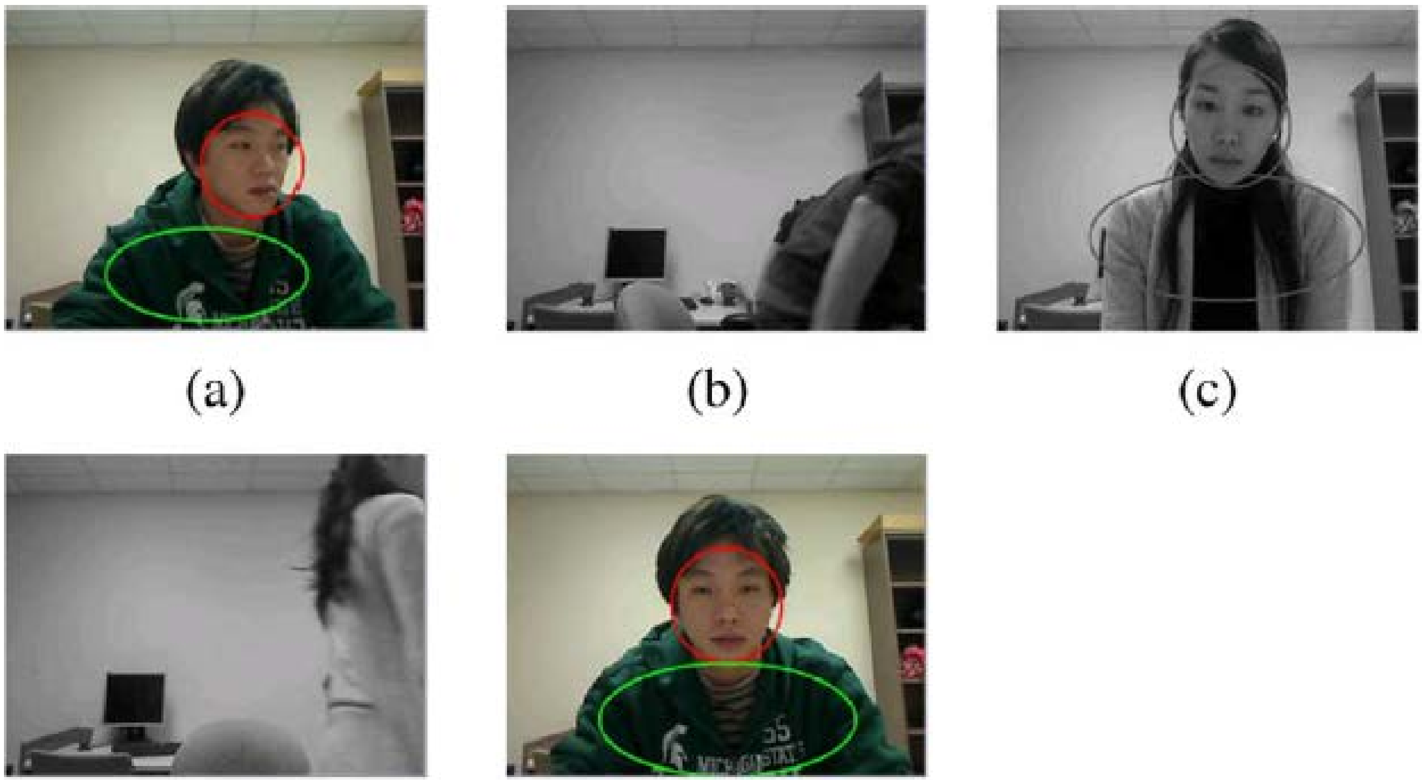

(a) (b) (c)

(d) (e)

Fig. 20. Example results of relogin authentication experiments. (a) Authentic user; (b) authentic user walks away; (c) imposter user; (d) imposter user walks away; and (e) authentic user returns.

Evaluation

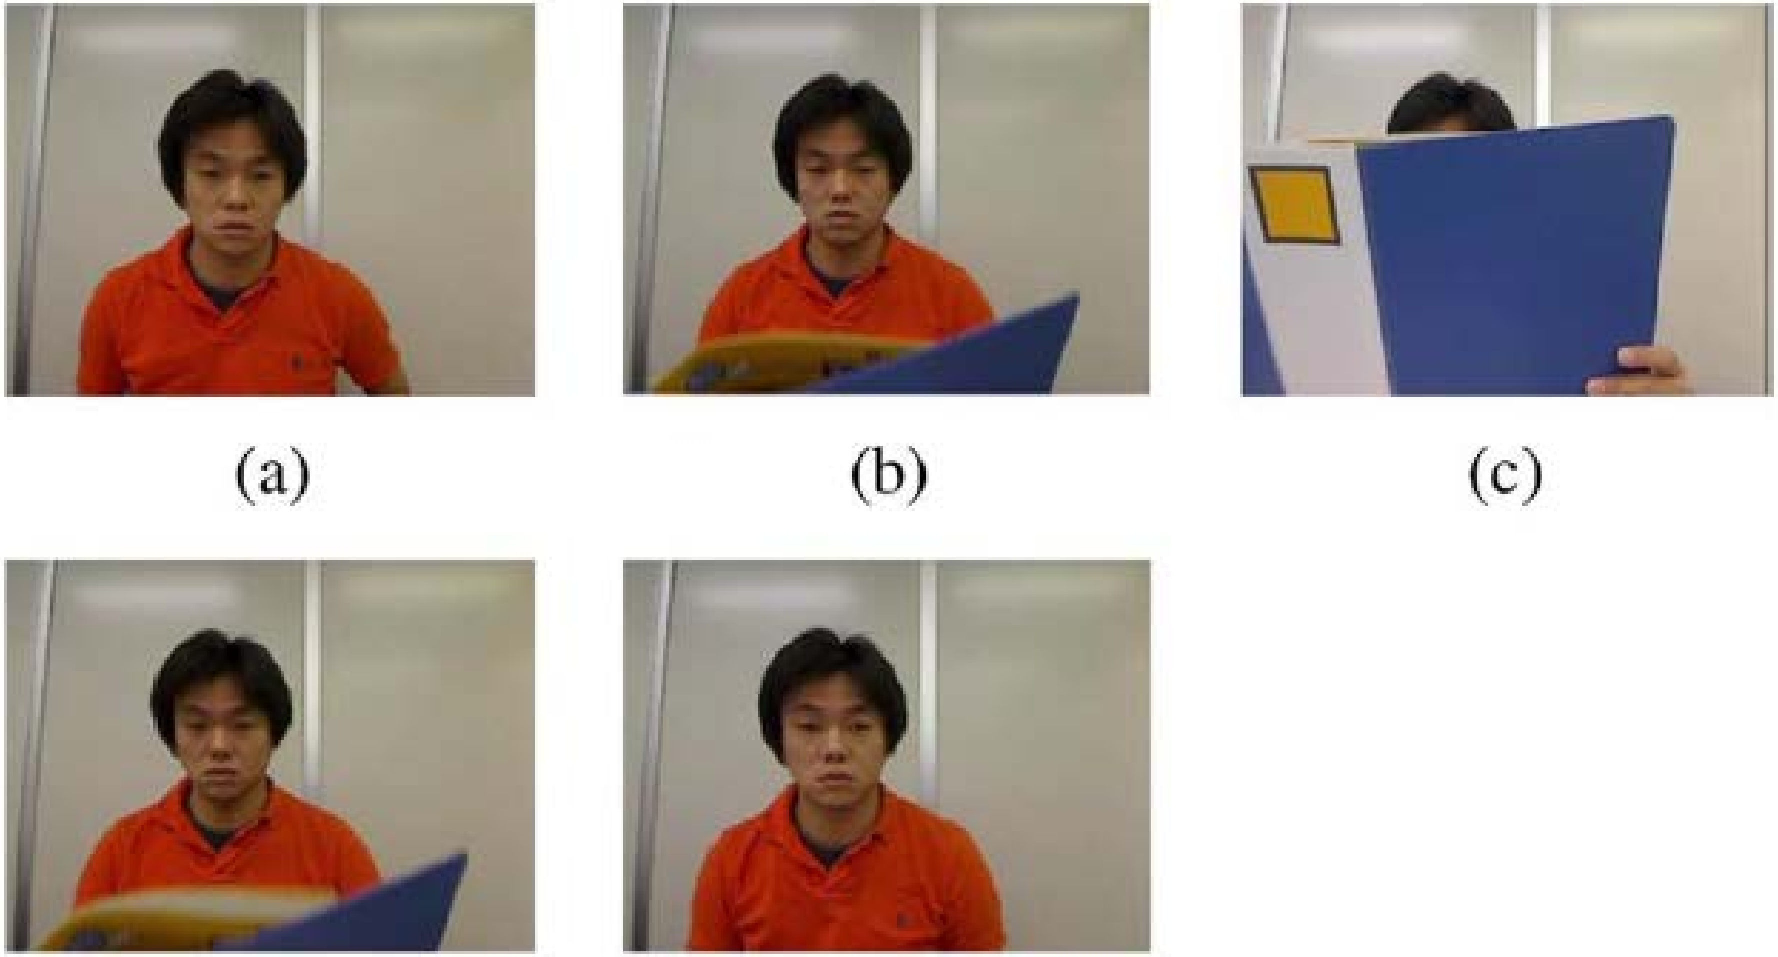

(a) (b) (c)

(d) (e)

Evaluation

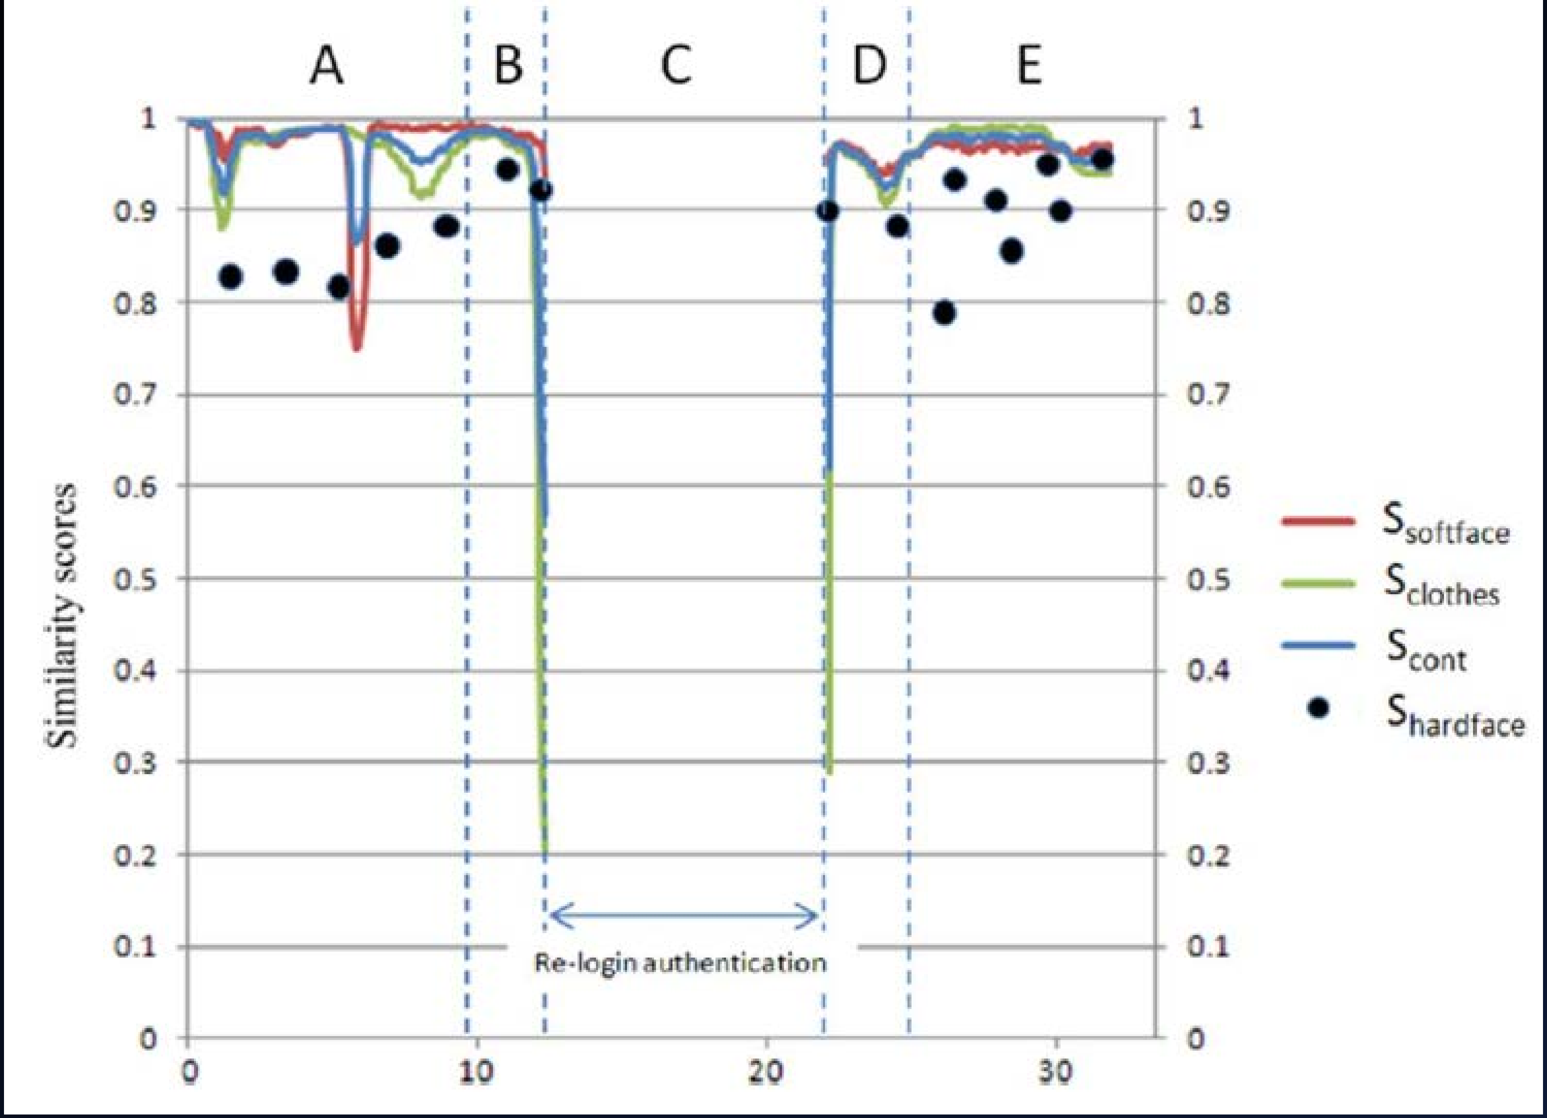

A
B
C
D
E
Similarity scores
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0
10
20
30
Re-login authentication
$S_{softface}$
$S_{clothes}$
S_{cont}
$S_{hardface}$


GROWERS
VINE TOMATO


Google
38%
23:15
Gismeteo
Камера
Диспетчер фай
Navitel
ivi.ru
Kingsoft Office
Спорт-Экспрес
Wexler Play
PLAY

Smartphones

- New opportunity for Continuous Authentication
- Rich sensors:

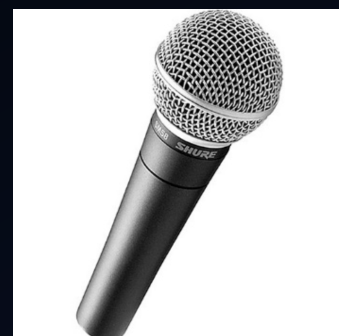

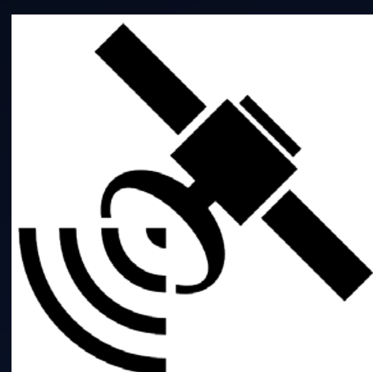

Possible biometrics

- Face: gender, identity, age, race, expression
- Iris?
- Voice
- Gait
- Keystroke dynamics (touch)
- Fingerprint
- Location
- Wifi signature
- Cellular signature

Energy usage is critical!

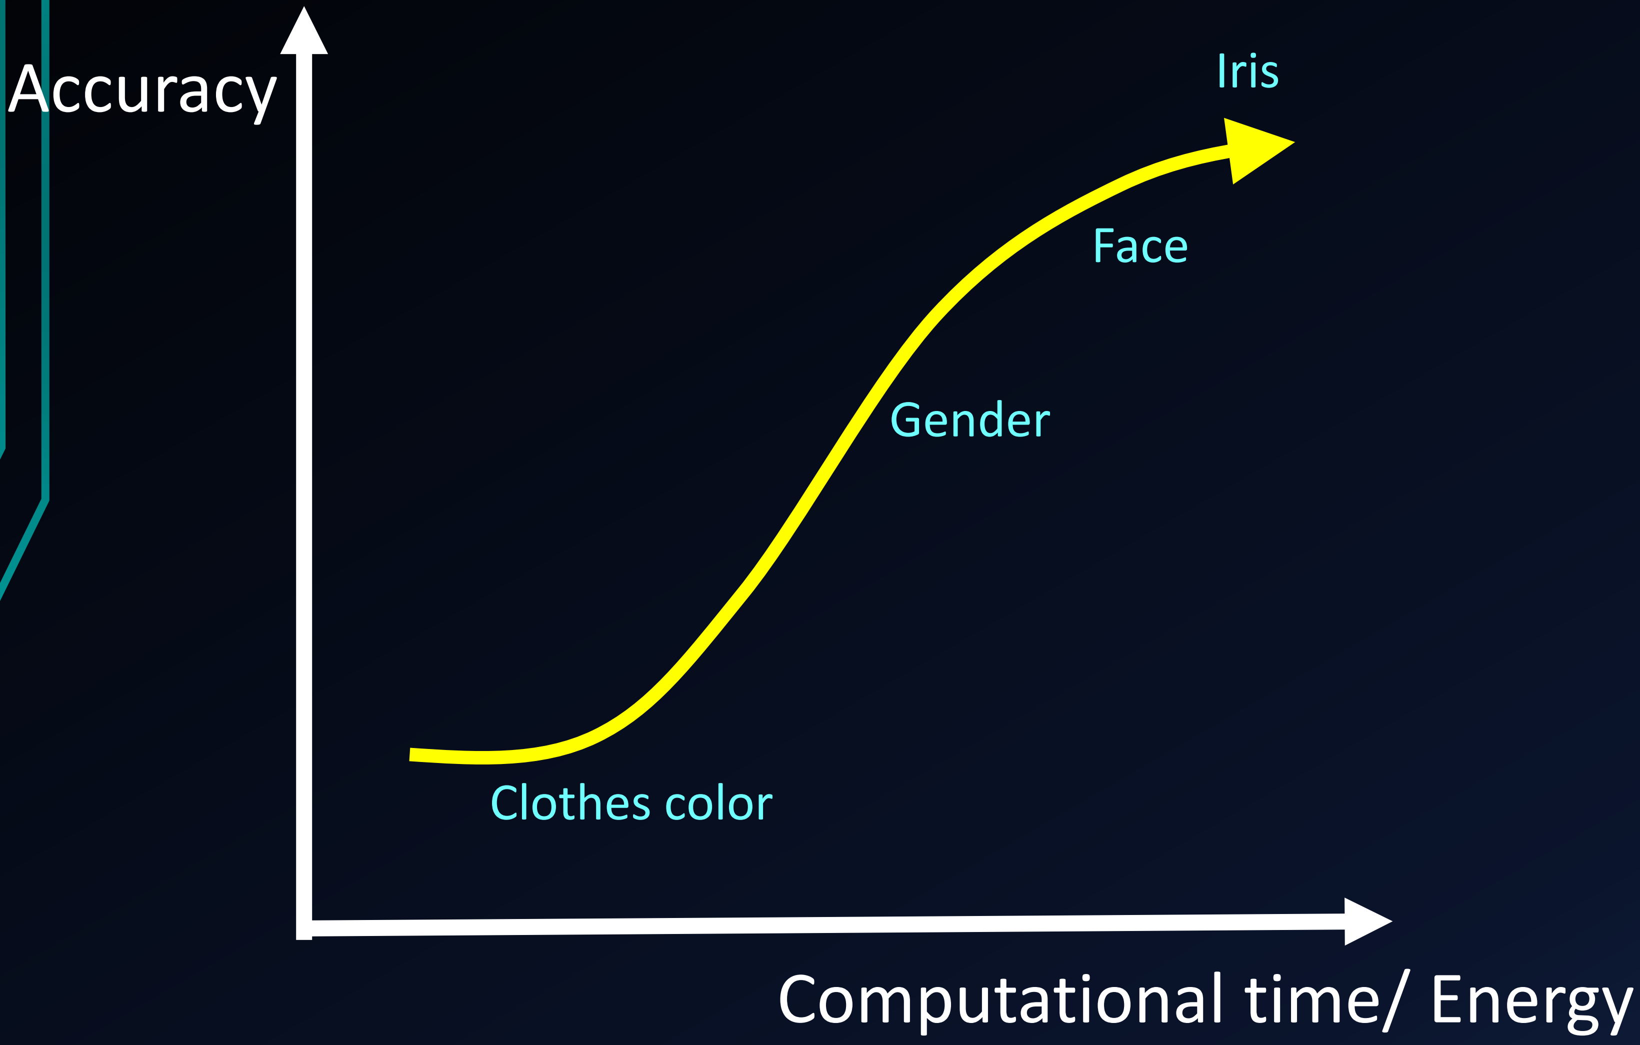

- Most research use touch dynamics
- Multimodal biometrics will be more useful
- Computational efficiency not yet considered
- Possibility for forensics use


Google
38%
23:15
Gismeteo
Камера
Диспетчер фай
Navitel
ivi.ru
Kingsoft Office
Спорт-Экспрес
PLAY
Wexler Play

References

- Sim, Terence, Sheng Zhang, Rajkumar Janakiraman, and Sandeep Kumar. "Continuous verification using multimodal biometrics." IEEE transactions on pattern analysis and machine intelligence 29, no. 4 (2007): 687-700.
- Kwang, Geraldine, Roland HC Yap, Terence Sim, and Rajiv Ramnath. "An usability study of continuous biometrics authentication." In International Conference on Biometrics, pp. 828-837. Springer Berlin Heidelberg, 2009.
- Niinuma, Koichiro, Unsang Park, and Anil K. Jain. "Soft biometric traits for continuous user authentication." IEEE Transactions on information forensics and security 5, no. 4 (2010): 771-780.
- Janakiraman, Rajkumar, and Terence Sim. "Keystroke dynamics in a general setting." In International Conference on Biometrics, pp. 584-593. Springer Berlin Heidelberg, 2007.
- Traore, Issa, ed. Continuous Authentication Using Biometrics: Data, Models, and Metrics: Data, Models, and Metrics. IGI Global, 2011.