Lecture 13: Human Authentication

CS 181S

Fall 2020

Classes of Countermeasures

- **Authentication:** mechanisms that bind principals to actions
- **Authorization:** mechanisms that govern whether actions are permitted
- **Audit:** mechanisms that record and review actions

Classes of Principals

- **Authentication:** mechanisms that bind principals to actions

  - Authenticating Machines
  - Authenticating Programs
  - Authenticating Humans

IDENTITY

Personal identity

- Major philosophical problem
  - People are not identical to themselves over time, but their identity persists throughout changes
  - cf. Ship of Theseus
- Intrinsic identity: continuation of consciousness
- Extrinsic identity: relationship to everything else
- Control: individual's, others', no one's?

Digital identity

- Digital identity: data that describes a person and its relationship to others
  - not the person itself; not a personal identity
- A person could have many digital identities, some overlapping, some contradictory
- Data could be incorrect, outdated, incomplete

Aspects of digital identity

- Name
- NetID
- Email address
- URL
- IP address
- Citizenship
- Political party
- ...

Identity

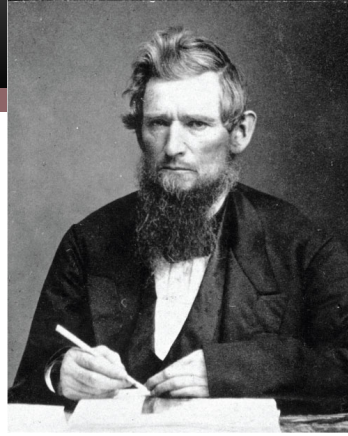

- Attribute: property of a principal
  - name is "Cecil Sagehen", birthdate is 11/29/1913
- Identity: set of attributes
  - each principal may have many identities of use in different scenarios (student, taxpayer, athlete)
- Identifier: an attribute that is unique within a population
- Verifier: an attribute that is hard to produce hence can be used as a basis for authentication

Enrollment

- Enrollment: establishing identity with a system
  - Create an account
  - Get an ID card, visa
  - Register a machine on a network
  - Get a signing key from a provider
- System might (not) verify claimed attributes during enrollment
  - Websites rarely do
  - Governments often do

HUMAN AUTHENTICATION

Authentication of humans

- Something you are
  - biometrics (e.g., fingerprints)
- Something you know
  - secret information (e.g., a password)
- Something you have
  - possession of a physical device (e.g., a particular phone)

Exercise 1: Classifying Authentication

- Come up with a list of ways you have authenticated yourself to a machine. For each, classify it as something you are, something you know, or something you have

Something you are

Something you know

Something you have

Exercise 1: Classifying Authentication

- Come up with a list of ways you have authenticated yourself to a machine. For each, classify it as something you are, something you know, or something you have

Something you are

fingerprint, retinal scan, facial scan

Something you know

password, passphrase, PIN, answers to security questions

Something you have

phone, token, physical key, ticket, {ATM, prox, credit} card

Multi-factor Authentication

- Two-factor authentication: authenticate based on two independent methods
  - ATM card plus PIN
  - password plus registered mobile phone
- Multi-factor authentication: two or more independent methods
- Best to combine separate categories, not reuse categories
  - non-example: requiring two passwords from a single human: arguably not independent
  - non-example: requiring single password from each of two humans: authenticates two humans then makes *authorization* decision

SOMETHING YOU ARE

Biometric

- Biometric: measurement of biological and behavioral attributes (something you are)
  - biological attributes can be confounded by behavior
  - biology and behavior is non-constant: variation from one measurement to the next

Example: Fingerprint

- Particular use: California social services
  - prevent applicants for welfare from defrauding state by receiving assistance under multiple identities
- Fingerprint stored as bitmap and as minutae
  - When user authenticates, computer compares minutiae
  - If they match, human additionally reviews bitmap images (about 15 out of 10000 authentications have minutiae match even though fingerprints do not)

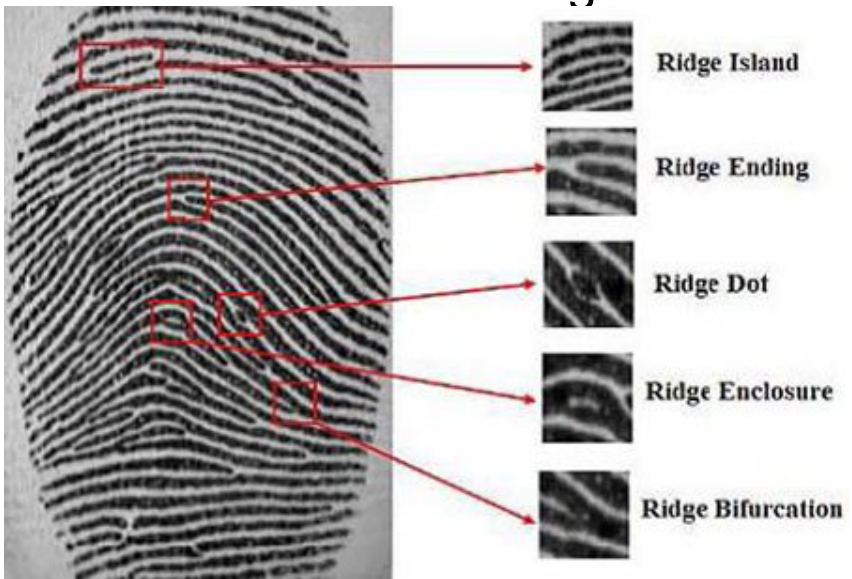

Example: Hand geometry

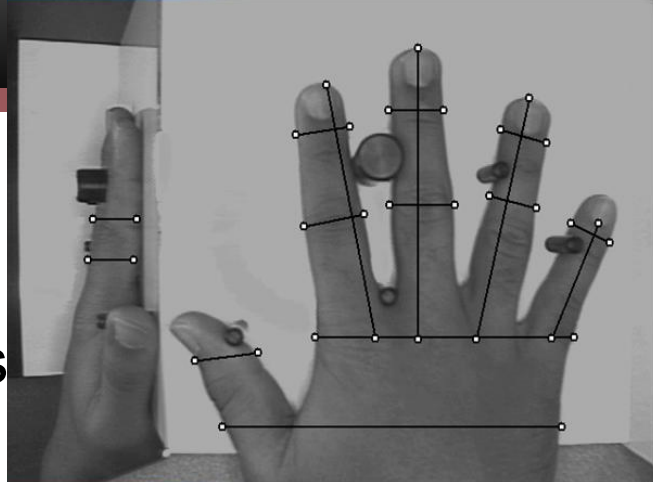

- Used in 2012 Olympic Games, Walt Dis
  nuclear facilities, data centers, ...
- Camera images palm and side of hand (no texture information)
- Images reduced to (e.g.) 31000 points then 90 measurements then 9 bytes of data
  - Final data not directly related to any source measurements
  - Data stored as a template for later comparison
- When user authenticates, another set of images taken
  - If data are close enough to stored template, user deemed authenticated
  - Can adjust threshold per-user, in case some users are difficult to authenticate
- Each time user is authenticated, template is updated to account for change over time

Example: Facial recognition

- Used in border control, Facebook, iPhone X
- Operates on 2D image or depth map
- Modern systems use ML classifiers to identify matches
  - Most systems perform poorly on profiles, low-res images
  - Most systems perform less well on women and minorities

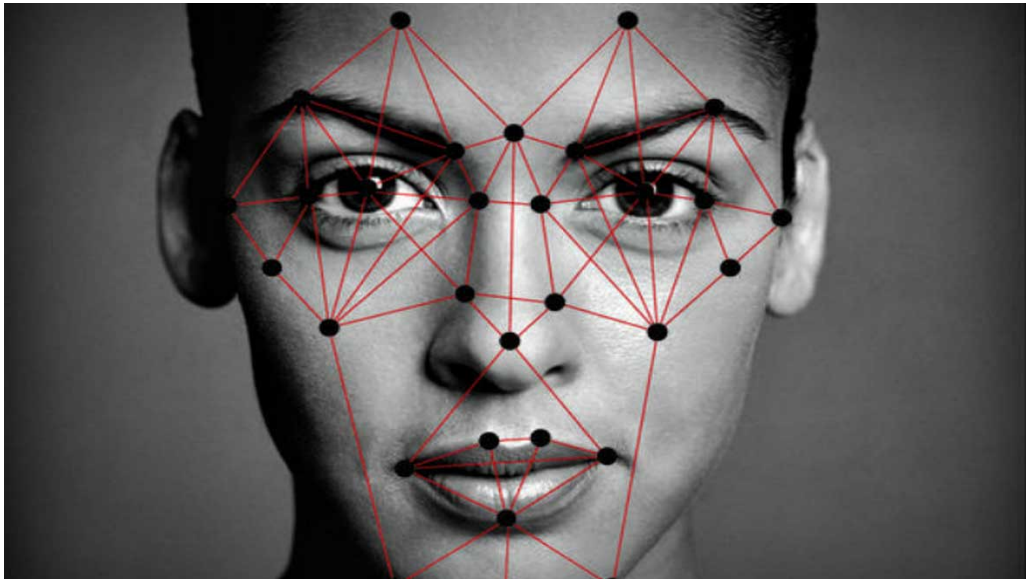

Other Biometrics

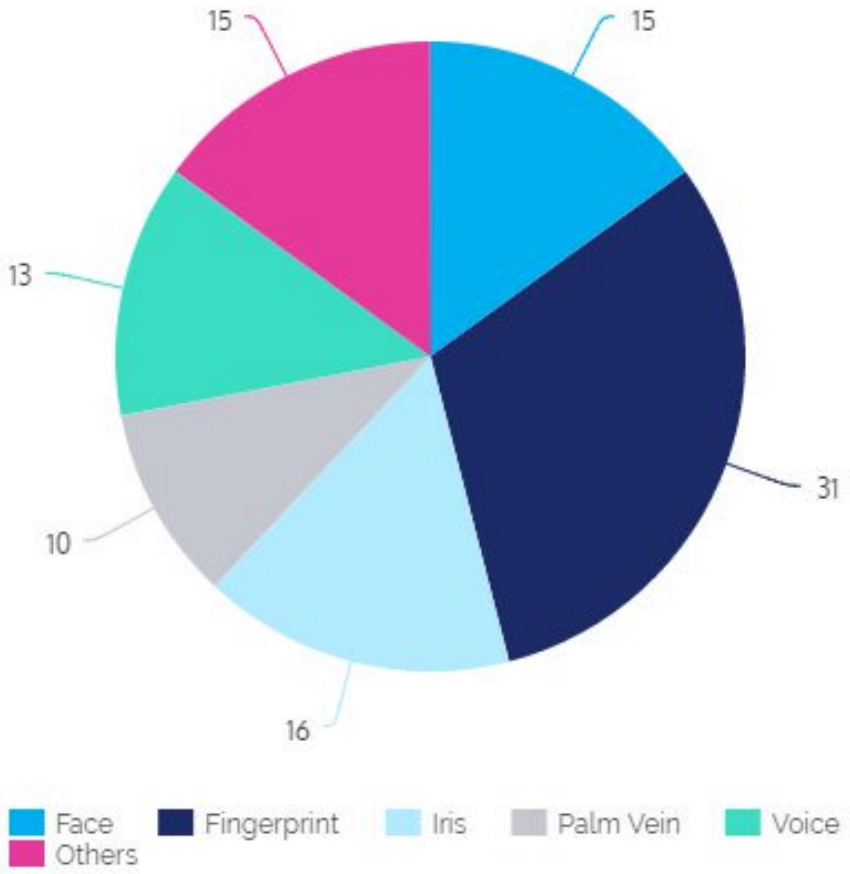

Biometric attributes as verifiers

- **Advantages:**
  - Can't lose or forget a biometric
  - Easy to use some biometrics (e.g., fingerprint scan vs. PIN on iPhone)
- **Disadvantages:**
  - Physical process with errors...
  - Updating identities after disclosure is hard (new fingerprints? new retina?)
    - So enrolling a biometric identifier places **permanent trust** in receiver, even if they go bankrupt, retroactively change privacy policies, get taken over by new administration, ...
  - Impossible to be application specific (your hand geometry is the same regardless of what system you use)
  - Fear of negative implications for privacy...

EVALUATING BIOMETRICS

Biometric attributes as verifiers

Requirements:

- Identifier
- Easy to measure
- Small variation over time and measurement
- Difficult to spoof
- Acceptable to users

Accuracy

- False accept: authenticate a principal with wrong identity (fraud)
- False reject: fail to authenticate a principal under right identity (insult)
- Hypothesis testing:
  - null hypothesis: human being authenticated has claimed identity
  - false reject = type I error
  - false accept = type II error
- Tunable trade off of sensitivity between which error is more likely
  - False acceptance rate (FAR): percentage of attempts in which imposters are authenticated (with wrong identity)
  - False reject rate (FRR): percentage of attempts in which legitimate users are denied authentication

Sensitivity

Receiver operating characteristics (ROC) curve: graph of FRR vs. FAR (or perhaps 1-FAR, perhaps nonlinear axes)

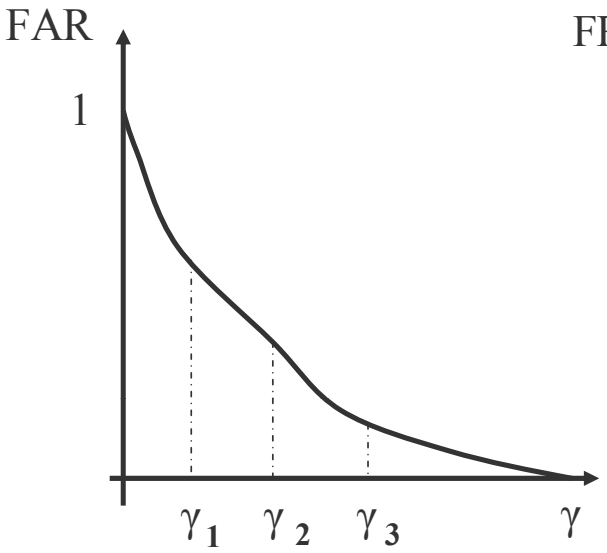

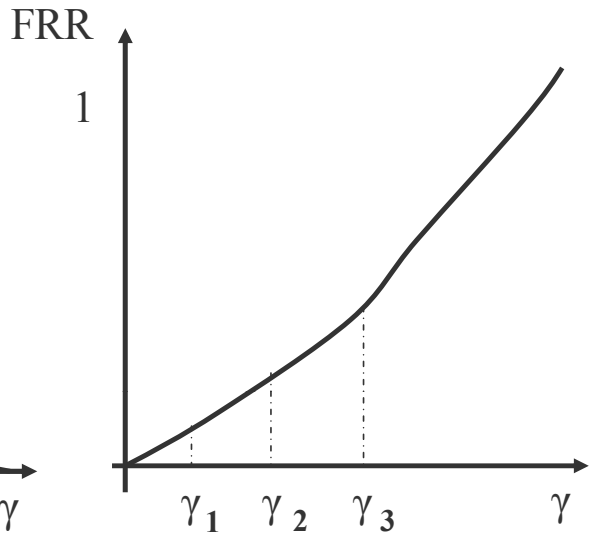

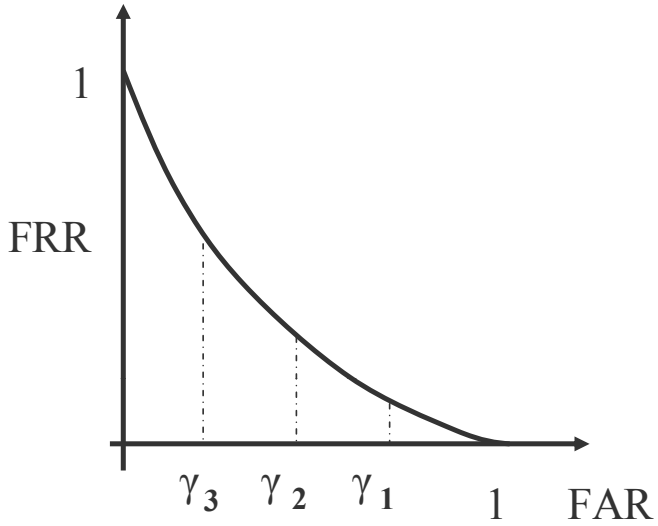

γ = sensitivity

ROC comparison

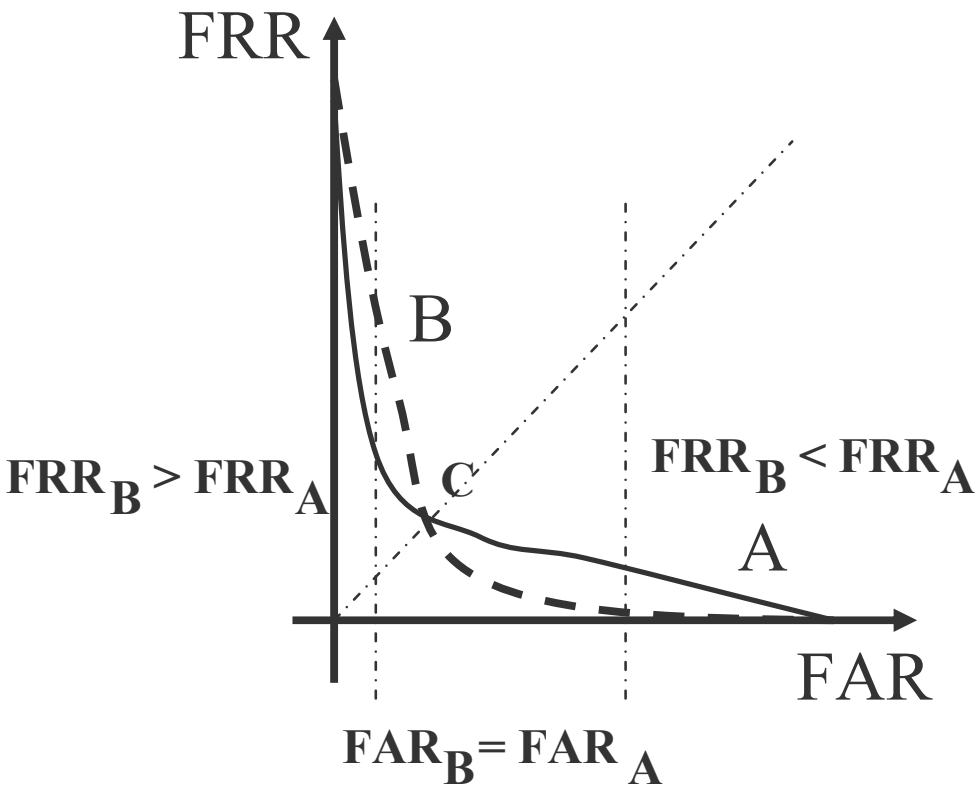

- Two matchers (A=solid; B=dashed)
- At point C, matchers have same FAR and FRR
- To the left of C, matcher A has lower FRR for same FAR
- To the right, matcher B has lower FRR for same FAR

ROC comparison

- Crossover error rate (CER): value on ROC at which FAR=FRR (aka *equal error rate, ERR)*
- Many other statistics for comparison possible
  - Anytime a graph is reduced to a single number, we lose information

- *What matters most for biometrics is the use case/threat model*

Use cases

- **Entry to military facility:**
  - letting imposters in might be worse than (temporarily) delaying entry of personnel
  - so prefer low false accept rate
- **Entry to hotel lobby:**
  - letting non-guests in might be better than (temporarily) delaying entry of guests
  - so prefer low false reject rate

Comparing Biometrics

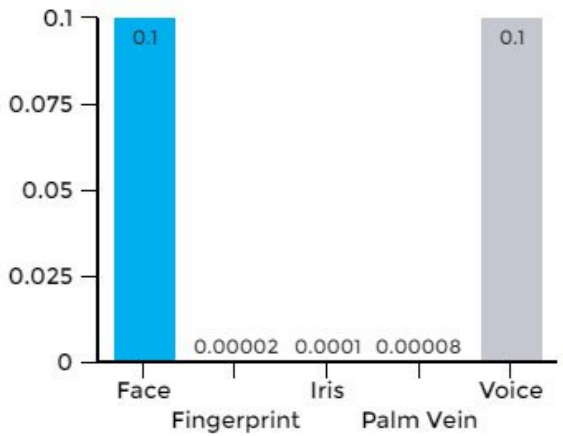

False Acceptance Rate

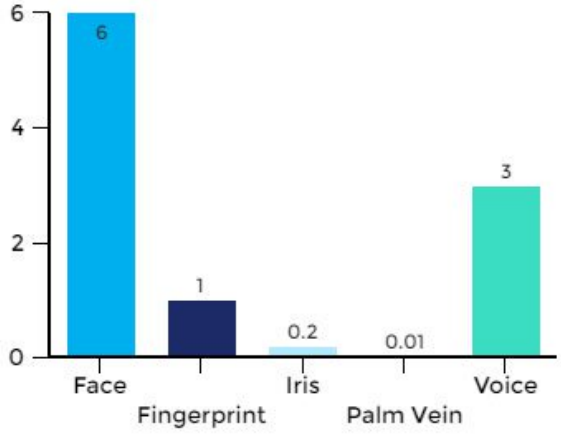

False Rejection Rate

Biometric Technology	Accuracy	Cost	Devices required	Social acceptability
ADN	High	High	Test equipment	Low
Iris recognition	High	High	Camera	Medium-low
Retinal Scan	High	High	Camera	Low
Facial recognition	Medium-low	Medium	Camera	High
Voice recognition	Medium	Medium	Microphone, telephone	High
Hand Geometry	Medium-low	Low	Scanner	High
Fingerprint	High	Medium	Scanner	Medium
Signature recognition	Low	Medium	Optic pen, touch panel	High

Spoofing

- Active adversary fools sensor with artificial object
- Solution:
  - better sensors
  - better biometrics
  - multi-factor authentication

Gummy Bear Attack

Face ID Attack

Exercise 2: Evaluating Biometrics

Consider the use of voice authentication as a biometric. With voice authentication, the human is asked to say a specific passphrase and their response compared to a recorded voice print by a machine learning system.

1. What are potential advantages of this biometric?
2. What are potential disadvantages of this biometric?

Privacy concerns

- Humans might have concerns about measurements (have photo taken, parts of body scanned)
- Humans might not want to disclose attributes during enrollment (SSN, political party)
- Humans might not want action bound to their identity (buying medication)
- Humans might not want their actions linked to other actions, exposing them to inference about what they thought were unrelated activities.

Privacy and biometrics

- Biometrics can violate intrinsic privacy by requiring submission to bodily contact or measurement
  - Fear of germs
  - Religious prohibitions
- Biometrics can violate informational privacy
  - Biometric identifiers might effectively become a standard, universal identifier, enabling linking

Principles for privacy

- **Seek consent:** get permission to authenticate and store identity
- **Select minimal identity:** use the smallest possible set of attributes
- **Limit storage:** don't save information about identity or authentication without need, and delete when no longer needed
- **Avoid linking:** don't reuse identifiers across systems

Exercise 3: Feedback

1. Rate how well you think this recorded lecture worked
    1. Better than an in-person class
    2. About as well as an in-person class
    3. Less well than an in-person class, but you still learned something
    4. Total waste of time, you didn't learn anything

2. How much time did you spend on this video lecture (including time spent on exercises)?

3. Do you have particular questions you would like me to address class?

4. Do you have any other comments or feedback?