Carnegie Mellon
DATA PRIVACY LAB

Data Privacy
Beyond Specification is Enforcement

Bradley Malin
School of Computer Science
Carnegie Mellon University
October 31, 2005

Traditional Beliefs

- We know how to protect privacy:
  - "If You Encrypt, They Must Acquit"
    - *(Cryptography, Secure Storage)*
  - "Make Strong Barriers"
    - *(Authentication, Network Security, Intrusion Detection)*
  - "Inform Collectors and Users"
    - *(Policy Specification, Auditing)*
  - "Don't Share Identity"
    - *(Federal Agencies, Data Brokers, Crediting)*

Security for Privacy?

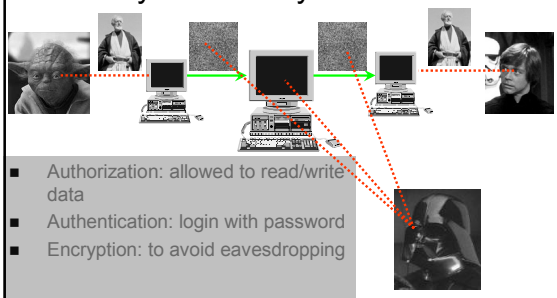

- Authorization: allowed to read/write data
- Authentication: login with password
- Encryption: to avoid eavesdropping

Security for Privacy?

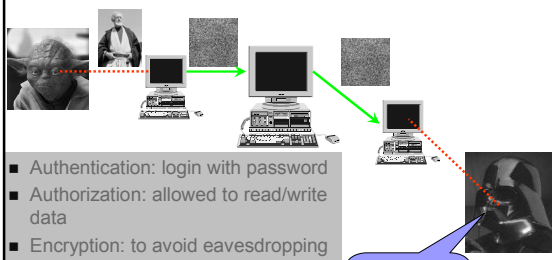

- Authentication: login with password
- Authorization: allowed to read/write data
- Encryption: to avoid eavesdropping
- But Data Can Re-identify!

Can I see some anonymous data?

Security for Privacy?

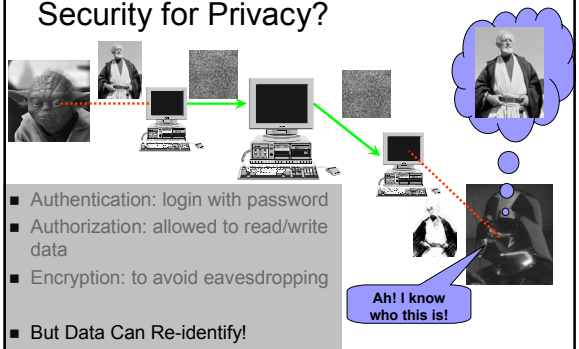

- Authentication: login with password
- Authorization: allowed to read/write data
- Encryption: to avoid eavesdropping
- But Data Can Re-identify!

Ah! I know who this is!

Ah-ha! it's the Data: Define a Privacy Policy

- Cheer for the many benefits!
  - *Procedure*
    - Specifies how data can (not) be used
  - *Logical Cognition*
    - Demands active involvement & thought regarding information
  - *Standardization*
    - equal opportunity
  - *Legal Enforcement*

L. Cranor. Web privacy with P3P. O'Reilly & Associates. Sebastopol, CA. 2002.

W. Stufflebeam, et al. Specifying privacy policies with P3P and EPAL: lessons learned. *Workshop on Privacy in the Electronic Society*. 2003.

Why? Legal Aspects

- United States
  - Federal / State level
  - Privacy Act of 1974
  - Privacy on the WWW
  - Educational (FERPA)
  - Wiretap & Surveillance Laws
  - Financial (GLB)
  - Medical (HIPAA)
  - Minors (COPPA)
  - …

Why? Legal Aspects

- United States
  - Federal / State level
  - Privacy Act of 1974
  - Privacy on the WWW
  - Educational (FERPA)
  - Financial (GLB)
  - Medical (HIPAA)
  - Minors (COPPA)
  - Wiretap and Surveillance Laws
- Europe
  - Data Directive 95/46
  - Safe Harbor and US
  - …

Let's Consider FERPA
(Buckley Amendment)

- Family Educational Right to Privacy
- Applies to: schools receiving funds from US Dept. Educ.
- If school permits the release of students' educational records w/o written consent of parents → Federal funding refusal
- Parents or eligible students have rights
  - Inspect/review student's school's education records
  - request school correct records believed to be inaccurate/misleading

FERPA

- Schools may disclose, without consent, "directory" information, such as:
  - Name
  - address
  - telephone number
  - date and place of birth
  - honors and awards
  - and dates of attendance
- Schools must alert parents/students about directory and allow request not to disclose
- Schools must notify parents and eligible students annually of their rights under FERPA

FERPA in Practice

- Many schools' privacy policies state they choose not to post any (or minimal) directory information on their students
- Example: MIT
  - "*7.2 School, department and lab web pages - Faculty, staff and students MUST EXERCISE CAUTION IN POSTING DIRECTORY and other information to a web page that is accessible to MIT and/or to the public. STUDENTS HAVE THE RIGHT TO WITHOLD DIRECTORY and other information from public distribution. FACULTY AND STAFF MUST RECEIVE PERMISSION to post personal information and identification photographs to web pages.*"

FERPA in the Face of Technology

- *RosterFinder* – software program finds online name "lists"
  - Leverages Google API
- Applied to gather undergraduate information
  - Discovered many directories of undergraduates online → not supposed to be there
  - Improper communication and privacy policy enforcement

L. Sweeney. Finding Lists of People on the Web. *ACM Computers and Society*, 34 (1) April 2004.

FERPA in the Face of Technology

- Precision of rosters by RosterFinder results
  - M: manually
  - R: RosterFinder
  - Pos: ranked position
  - Tot: total number of docs
- Increasing ability to gather data & infringe on privacy!
- But can also automated policy enforcement

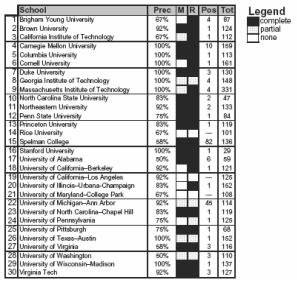

	School	Prec	M	R	Pos	Tot
1	Brigham Young University	67%			4	87
2	Brown University	92%			1	124
3	California Institute of Technology	67%			1	112
4	Carnegie Mellon University	100%			10	159
5	Columbia University	100%			1	113
6	Cornell University	100%			1	161
7	Duke University	100%			3	130
8	Georgia Institute of Technology	100%			4	148
9	Massachusetts Institute of Technology	100%			4	331
10	North Carolina State University	83%			2	87
11	Northeastern University	92%			2	133
12	Penn State University	75%			1	84
13	Princeton University	83%			1	119
14	Rice University	67%			---	101
15	Spelman College	58%			82	136
16	Stanford University	100%			1	29
17	University of Alabama	50%			6	59
18	University of California--Berkeley	92%			1	121
19	University of California--Los Angeles	92%			---	125
20	University of Illinois--Urbana-Champaign	83%			1	162
21	University of Maryland--College Park	67%			---	108
22	University of Michigan--Ann Arbor	92%			45	114
23	University of North Carolina--Chapel Hill	83%			1	119
24	University of Pennsylvania	75%			1	125
25	University of Pittsburgh	75%			1	68
26	University of Texas--Austin	100%			1	162
27	University of Virginia	58%			3	116
28	University of Washington	50%			3	110
29	University of Wisconsin--Madison	100%			1	131
30	Virginia Tech	92%			3	127

Legend: complete, partial, none

L. Sweeney. Finding Lists of People on the Web. *ACM Computers and Society*, 34 (1) April 2004.

Great, you Defined a Privacy Policy? But Wait a Minute…

- But consider some of the limitations
  - Need robust language (*P3P and EPAL are the beginning*)
  - Scope of world / interaction
  - Syntax, not semantics
  - Need enforcement
- Enter data privacy:
  - *WHERE* does data come from?
  - *WHAT* does data reveal?
  - *HOW* do we prove data does not reveal more than specified?

What is Data Privacy?

- The study of computational solutions for releasing data such that
  - data remains practically useful while
  - aspects of subjects are not revealed.
- Privacy Protection ("data protectors"):
  - release information such that entity-specific properties (e.g. identity) are controlled
  - restrict what can be learned
- Data Linkage ("data detectives")
  - combining disparate pieces of entity-specific information to learn more about an entity

"Privacy" Is Complex

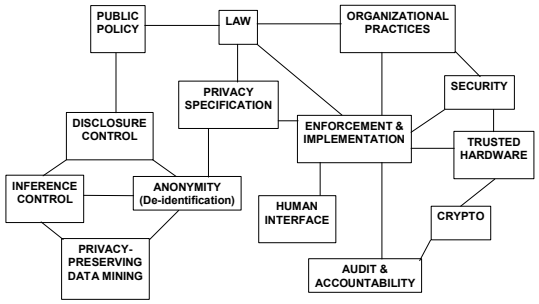

Diagram courtesy of Michael Shamos.

Data Privacy is Interdisciplinary

	anonymity	*rights mgt*	*database*	*ubiquitous*
AI	heavy	some	light	heavy
learning	some	light	light	heavy
theory	heavy	light	some	heavy
database	light	light	heavy	some
language	light	heavy	light	some
security		light	some	some
IS		light	light	light

- "AI" primarily concerns knowledge representation and semantics
- "Learning" focuses on data mining algorithms
- "Theory" includes zero-knowledge proofs and multi-party computations

Data. Data. Data.

- What kind of data? *All kinds!*
  - Field Structured Databases
  - Text Documents
  - Genomic
  - Image
  - Video
  - Network (Physical or Social)
  - Communications

Information Explosion (Sweeney '97)

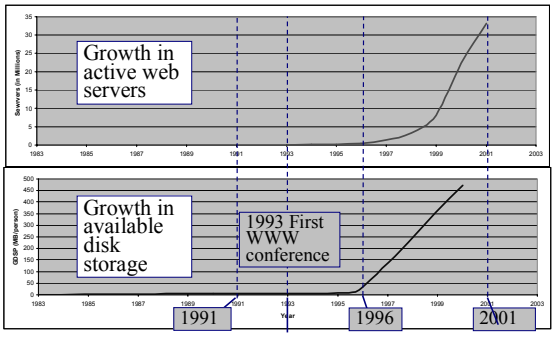

Anonymity & De-identification

- **Anonymous:** Data can not be "manipulated" or linked to identify an individual
- **De-identified:** All explicit identifiers, such as name, address, & phone number are removed, generalized, or replaced with made up values
- Does Anonymous = De-identified?

HIPAA

(Health Insurance Portability & Accountability Act)

- Rationale: Inconsistent state laws promulgating unnecessary difficulties in standardization, transfer, and sharing of health-related information
- A covered entity may not use or disclose protected health information
- Exceptions
  - To the individual that the information corresponds
  - With consent: to carry out treatment, payment, or health care operations
  - If consent is not required: same as above, but not with respect to psychotherapy notes

Data Sharing Under HIPAA

- Safe Harbor
  - Data that can be given away
  - requires removal of 18 direct and other quasi-identifiers
    Includes: name, address, zip code, phone number, birthdate, no geographic smaller than a state
- Limited Release
  - Recipient contractually agrees to not use or disclose the information for purposes other than prespecified research and will not identify or contact the individuals who are the subjects
  - May include specific geographic locations (i.e. zip code)
- Statistical or Scientific Standard (*we'll return to this*)

Healthcare Reform At Work

- Collect and disseminate hospital discharge data
- Attributes recommended by *National Association of Health Data Organizations* for disclosure

BUT… this is outside the jurisdiction of HIPAA

-9)

- Patient Gender
- Patient Racial Background
- Patient Number
- Visit Date
- Physician ID Number
- Physician Zip Code
- Total Charges

Linkage

- Use combination of attributes to determine the uniqueness of an entity in a dataset
- Second dataset with identified subjects is used to make the re-identification by drawing inferences between the two datasets on the related attributes
- The attributes do not have to be equal, but there must exist some ability for inference of between attributes.

Linking to “Re-identify” Data

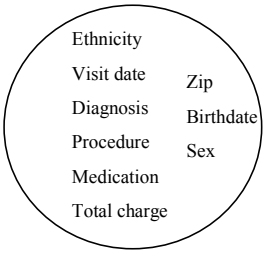

Medical Data

Linking to “Re-identify” Data

Voter List

Linking to “Re-identify” Data

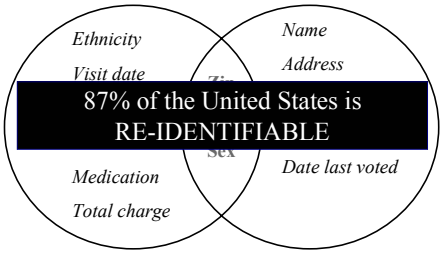

Medical Data **Voter List**

L. Sweeney. Weaving technology and policy to maintain confidentiality. *Journal of Law, Medicine, and Ethics*. 1997:

{*date of birth*, *gender*, *5-digit ZIP*} uniquely identifies 87.1% of USA

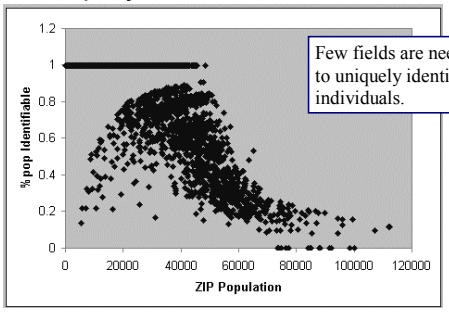

L. Sweeney. Uniqueness. Data Privacy Laboratory Technical Report. 2000.

{*date of birth*, *gender*, *5-digit ZIP*} uniquely identifies 87.1% of USA

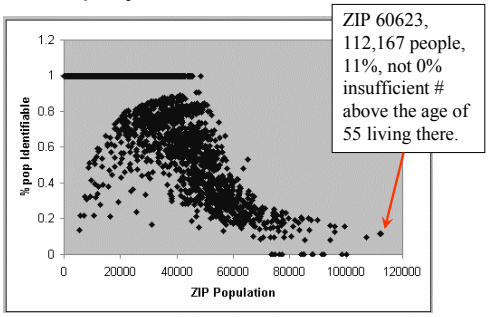

{*date of birth*, *gender*, *5-digit ZIP*} uniquely identifies 87.1% of USA

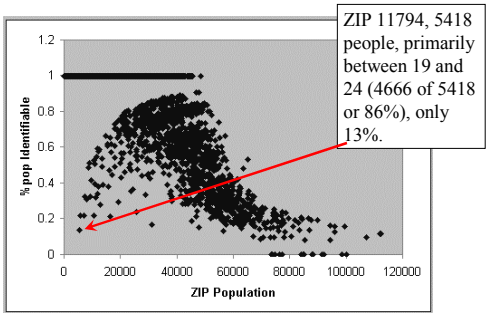

Chain of Links

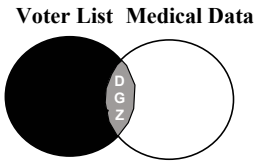

Chain of Links

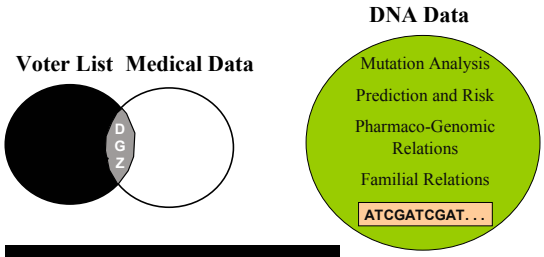

So what do you do?

DNA - Discharge Inferences Exist

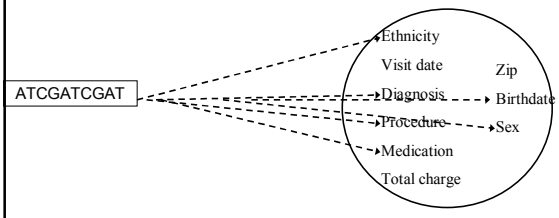

- Inferences can lead to re-identification

B. Malin and L. Sweeney. Determining the identifiability of DNA database entries. In Proceedings of the 2005 AMIA Annual Symposium. 2000: 547-551.

Genotype-Phenotype Relations

Can infer genotype-phenotype relationships out of both DNA and medical databases

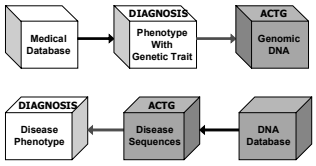

B. Malin and L. Sweeney. Determining the identifiability of DNA database entries. In Proceedings of the 2005 AMIA Annual Symposium. 2000: 547-551.

False Protection Example

Name	Address	Diagnosis	Treatment	DNA
John Doe	1 Some Way	33321	1230	accta…a
Jane Doh	2 No Way	82132	9012	agctt…c

Name	Address	Diagnosis	Treatment
John Doe	1 Some Way	33321	1230
Jane Doh	2 No Way	82132	9012

DNA
accta…a
agctt…c

DNA Re-identification

- Many deployed genomic privacy technologies leave DNA susceptible to re-identification
- DNA is re-identified by automated methods, such as:
  - Genotype – Phenotype (G-P) Inference

DNA *CLINICAL*

Sequence
accta…a
agctt…c

Name	Address	ICD9 Code
John Doe	1 Some Way	3334
Jane Doh	2 No Way	277

Linkage Prediction

Carnegie Mellon DATA PRIVACY LAB

DNA Re-identification

- Many deployed genomic privacy technologies leave DNA susceptible to re-identification
- DNA is re-identified by automated methods, such as:
  - Genotype – Phenotype (G-P) Inference

DNA	INFERRED DISEASE
accta…a	Cystic Fibrosis
agctt…c	Huntington's Disease

INFERRED DISEASE	Name	Address	ICD-9
Huntington's Disease	John Doe	1 Some Way	3334
Cystic Fibrosis	Jane Doh	2 No Way	277

Carnegie Mellon DATA PRIVACY LAB

DNA Re-identification

- Many deployed genomic privacy technologies leave DNA susceptible to re-identification
- DNA is re-identified by automated methods, such as:
  - Genotype – Phenotype (G-P) Inference

DNA	INFERRED DISEASE
accta…a	Cystic Fibrosis
agctt…c	Huntington's Disease

INFERRED DISEASE	Name	Address	ICD-9
Huntington's Disease	John Doe	1 Some Way	3334
Cystic Fibrosis	Jane Doh	2 No Way	277

Linkage Prediction – Unique Re-identification!

Carnegie Mellon DATA PRIVACY LAB

Longitudinal Genomic Learning Model

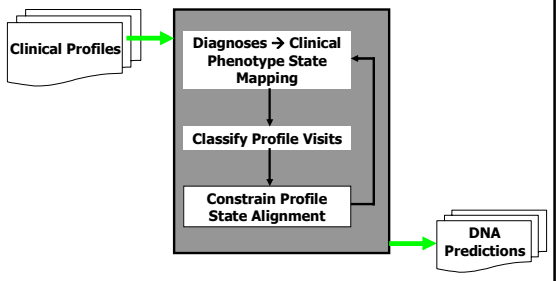

B. Malin and L. Sweeney. Inferring genotype from clinical phenotype through a knowledge-based algorithm. In Proceedings of the Pacific Symposium on Biocomputing. 2002: 41-52.

Carnegie Mellon DATA PRIVACY LAB

Learning DNA from Phenotype

- Example: Huntington's disease
  - Exists strong correlation between age of onset and DNA mutation (# of CAG repeats)
  - Given longitudinal clinical info, accurately infer age of onset in 20 of 22 cases

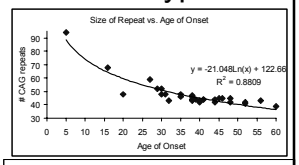

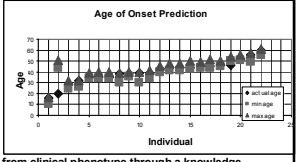

B. Malin and L. Sweeney. Inferring genotype from clinical phenotype through a knowledge-based algorithm. In Proceedings of the Pacific Symposium on Biocomputing. 2002: 41-52.

Carnegie Mellon DATA PRIVACY LAB

So What Do We Do?

- Some say, "You Can't Release Any Data"

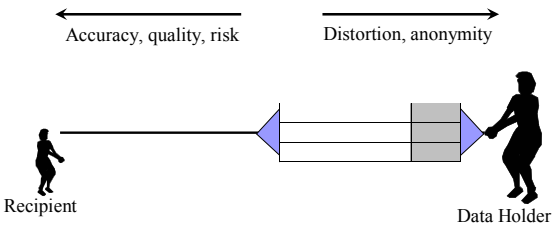

Carnegie Mellon DATA PRIVACY LAB

So What Do We Do?

- Others* say, "Privacy is Dead, Get Over It"

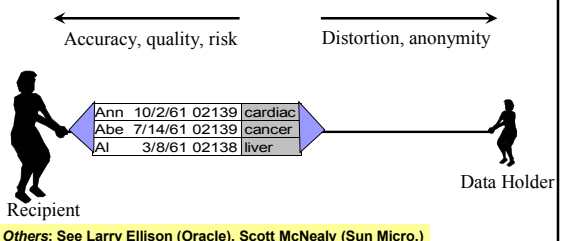

***Others*: See Larry Ellison (Oracle), Scott McNealy (Sun Micro.)**

So What Do We Do?

- We say, "Share Data While Providing Guarantees of Anonymity"

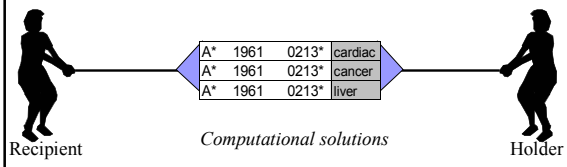

Computational solutions

Example: Camera-Happy World

- Over 3 million cameras in the US
  - Manhattan has over 2500 cameras
  - Average American caught on camera 8-10 times / day
- Over 4 million cameras in the UK
  - Average Londoner is caught > 300 times a day

Some Camera Watch Images

CMU Camera Watch Project

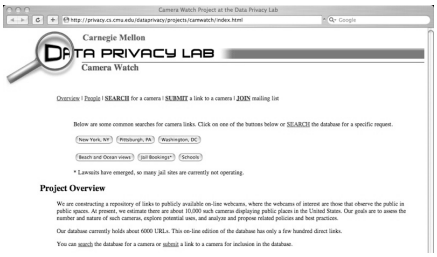

http://privacy.cs.cmu.edu/dataprivacy/projects/camwatch/index.html

Video Data Privacy Goal

- Modify video images so that
  - Privacy: automated attempts to recognize faces fail
  - Utility: knowledge learned from data is useful
- Solution to problem
  - Enables sharing of data for specified purposes
  - Protects rights as specified in policy
    - e.g. your identity won't be revealed unless you have done something illegal

The "Good" Side of Surveillance

- Homeland security monitoring
  - Monitor number of faces over time

- Early bioterrorism detection
  - Monitor for respiratory distress

L. Sweeney and R. Gross. Mining images from publicly-available cameras for homeland security. In Proceedings of the AAAI Spring Symposium. 2005.

Protection Post / During Capture

- Can we study video and image information for surveillance purposes with identity protection?
- Example: can we track people, but withhold identity?

A Solution: The "Dot" Approach

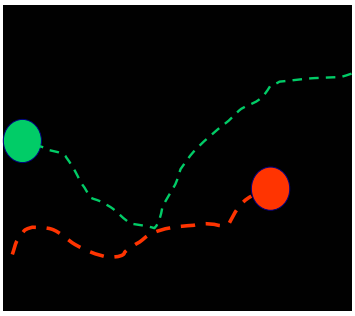

De-identifying People

- More detailed
- Silhouettes and coloring for tracking

(Andrew Senior --- IBM)

A. Senior, et. al. Enabling video privacy through computer vision. *IEEE Security and Privacy Magazine.* May-June 2005; 3(3): 50-57.

Alternative De-identification

- Masking and environmental suppression

(More from Andrew Senior - IBM)

Original | People Removed | Back Removed People Silhouette

Andrew Senior. http://www.research.ibm.com/people/a/aws/

Can we make Video Privacy More Formal?

- De-identification for some uses can be achieved by
  - replacing people with dots or
  - replacing faces with blobs.
- In each case, de-identification is achieved but not necessarily anonymity
- What if we need to see what a face is expressing?

De-identifying Video Surveillance Data

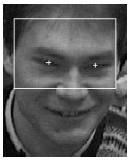

Example use.
Tracking coughs (bio-surveillance) or suspicious behavior in public spaces.

De-identification, not anonymity

Separating machines from humans

Example: De-identification of Faces

Captured images are below. Here is a known image of Bob. Which person is Bob?

Example: De-identification of Faces

Captured images are below. Here is a known image of Bob. Which person is Bob?

Face Recognition: The Big Idea

Identity Unknown

Name of Person

PCA-Based Face Recognition Systems

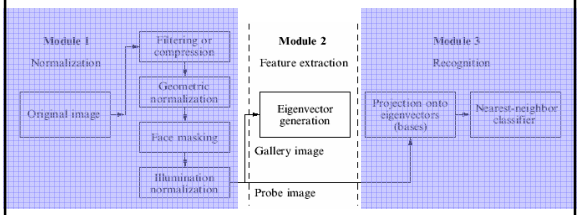

Goal: Limit success of Module 2

Module 2: Eigenfaces / PCA

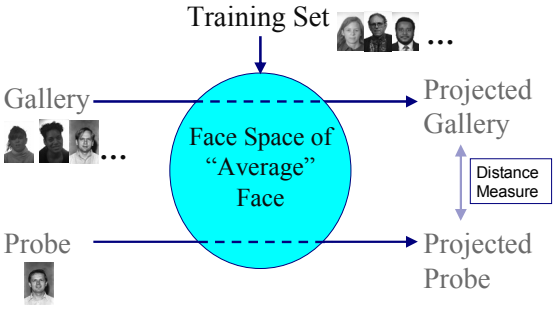

Face Recognition Software: ~70%

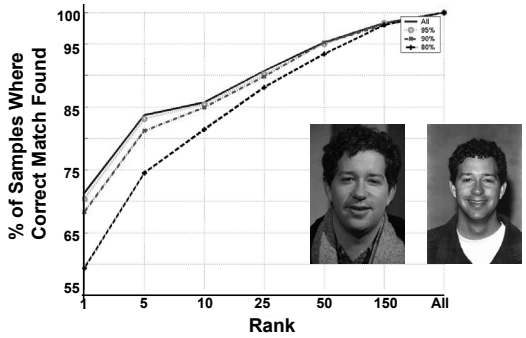

Eigenvectors (i.e. “Concepts”)

The characteristic function:

$$|(\mathbf{A}-\lambda\mathbf{I})| = 0$$

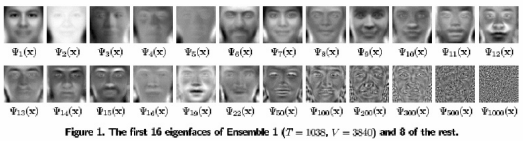

Figure 1. The first 16 eigenfaces of Ensemble 1 ($T = 1038$, $V = 3840$) and 8 of the rest.

where A is the covariance matrix

De-identification: T-mask

Example continued...

Captured images are de-identified below. Here is a known image of Bob. Which person is Bob?

De-identification: T-mask

Example continued...

Captured images are de-identified below. Here is a known image of Bob. Which person is Bob?

Automated Recognition Fails!

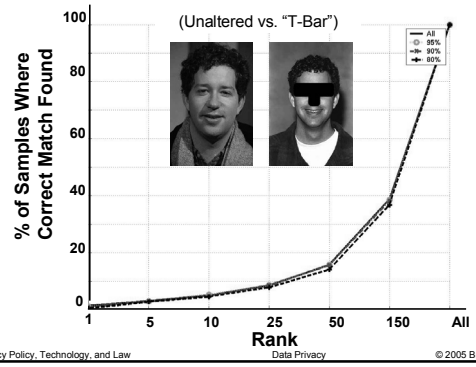

De-identification: pixel reduction

Example continued...

Captured images are de-identified below. Here is a known image of Bob. Which person is Bob?

De-identification: pixel reduction

Example continued...

Captured images are de-identified below. Here is a known image of Bob. Which person is Bob?

De-identification: pixel reduction

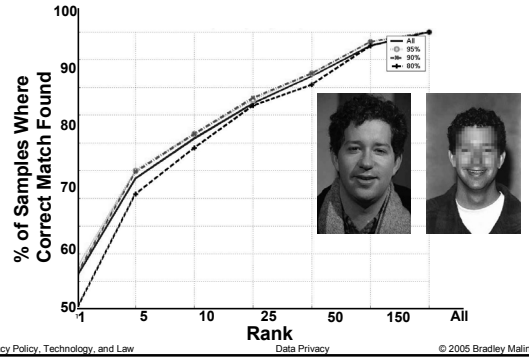

Pixelation: Automated Recognition Easier!

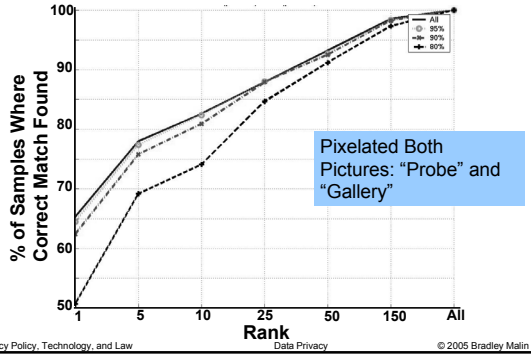

Why Try These Crazy Things?

- Many people and organizations claim they work

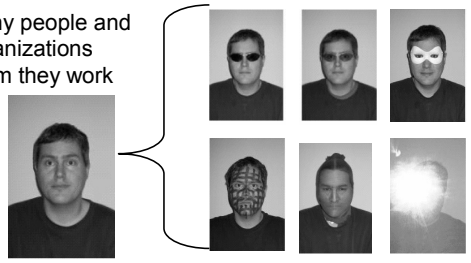

J. Alexander and J. Smith. Engineering privacy in public: confounding face recognition. Third Privacy Enhancing Technologies Workshop. 2003

Why Try These Crazy Things?

- Many people and organizations claim they work

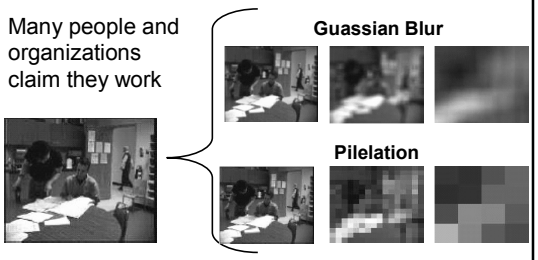

M. Boyle, C. Edwards, and S. Greenberg. The effects of filtered video on awareness and privacy. ACM Conference on Computer Supported Cooperative Work. 2000.

But Why Should We Care? Policy Sidebar

- European Data Directive
  - Collected video and images can not be released unless they have been sufficiently protected
- Contends "pixelation" is sufficient criteria for identity protection

More De-identification Ideas!

Ad Hoc Methods = Poor Protection

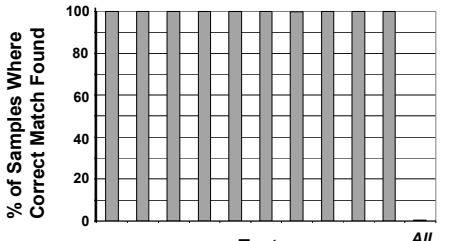

Not Looking Good…

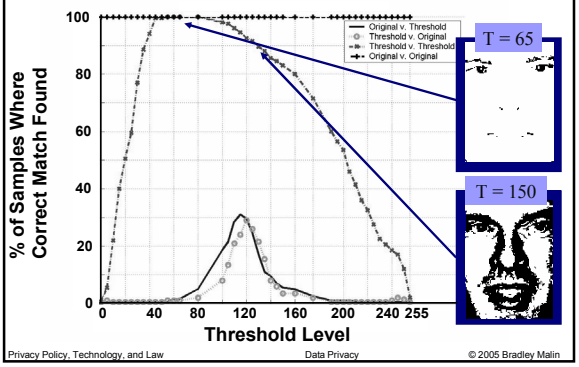

Random Changes to Grayscale Images

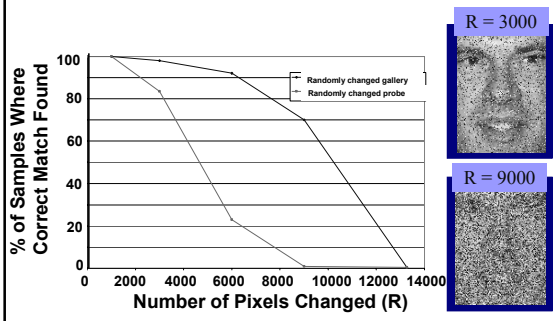

Don't be Naïve

- Again, de-identified ≠ anonymous
- Masks can be removed and trained against
- Some cases naïve de-identification even harms privacy!
  - pixelation and "blur" may improve performance
- Time to get logical…

k-Protection Models

- k-Anonymity: For every record, there are at least k individuals to whom it refers (realized upon release).
- k-Same: For every face, there are at least k people to whom that face refers. No face actually refers to a single real person.

E. Newton, L. Sweeney, and B. Malin. Preserving privacy by de-identifying facial images. IEEE Transactions on Knowledge and Data Engineering. 2005; 17(2): 232-243.

Formal Models of Anonymity

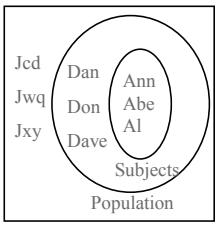

Jcd	cardiac
Jwq	cancer
Jxy	liver

Null-Map

Al	3/8/1961	02138 cardiac
Ann	10/2/1961	02139 cancer
Abe	7/14/1961	02139 liver

Wrong-Map

Ann	10/2/61	02139 cardiac
Abe	7/14/61	02139 cancer
Al	3/8/61	02138 liver

Private Information

A*	1961	0213*	cardiac
A*	1961	0213*	cancer
A*	1961	0213*	liver

k-Anonymity

Model Examples

k-map: For each tuple record in the release,

record must refer to at least *k* entities in the population

A*	1963	0213*	cardiac
A*	1961	0213*	cancer
A*	1964	0213*	liver

***k*-anonymity:** "*k* in the release"

A*	1961	0213*	cardiac
A*	1961	0213*	cancer
A*	1961	0213*	liver

Subexample: Population Registers

Register

There are three colors with frequencies: 1 red, 3 green and 2 blue.

There are 2 types of figures, with 2 of one type and 4 of the other.

The combination of color and figure labeled as Hal and Len are each unique.

Formal Protection Example

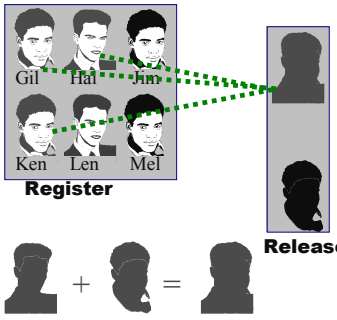

Register

Release

To achieve *k*-map where *k*=2, agents for *Gil*, *Hal* and *Ken* merge their info

Information released about any of them results in the same merged image.

Ranking of Faces

How does everyone rank against each other? Who is closest? Who is farthest?

k-Anonymity: Face Style!

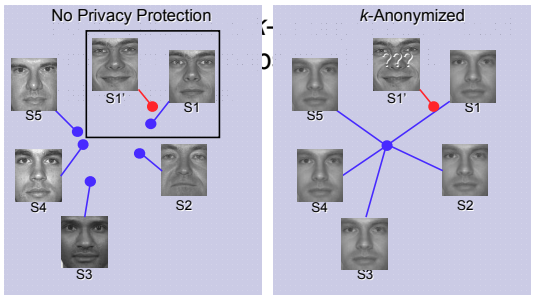

k-Same Algorithm

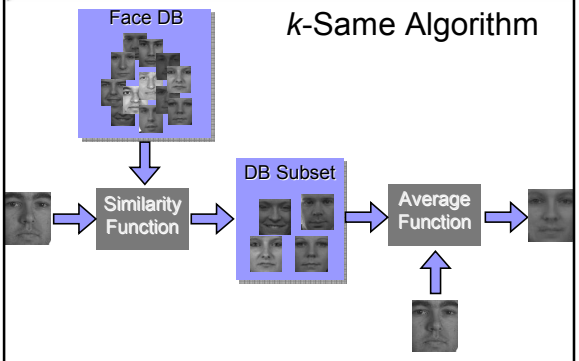

Example of 2-Same

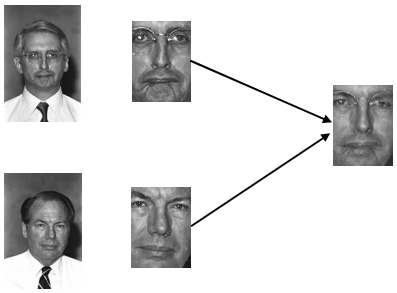

k-Same Example (More Depth)

-Pixel

-Eigen

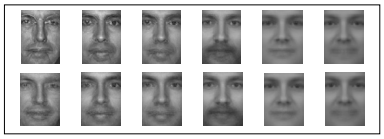

k = 2 3 5 10 50 100

k-Same Algorithm

- Guarantee
  - Image sets de-identified using *k*-Same are *k*-anonymized

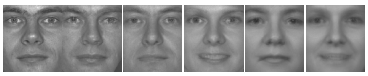

Performance of *k*-Same Algorithms

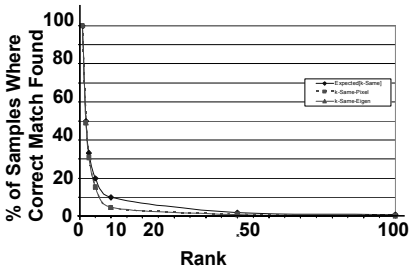

Upper-bound on Recognition Performance =

$$\frac{1}{k}$$

Can Guarantee this bound for ANY Recognition System

Some Intuition: Blurring

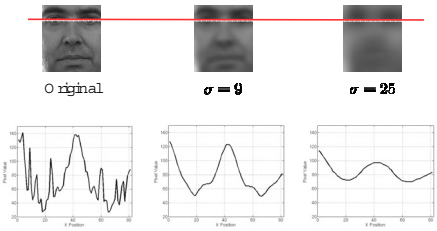

Some Intuition: Pixelation

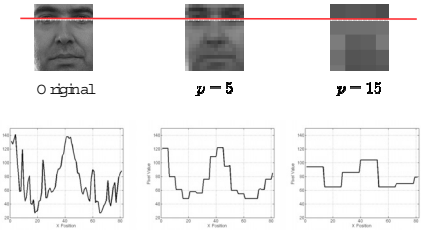

Some Intuition: *k*-same

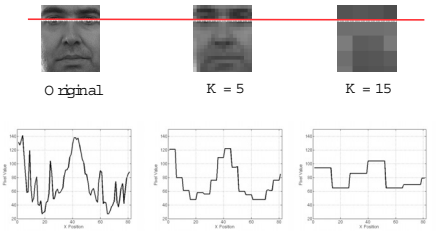

k-Same Algorithm Concerns

- Guarantee
  - Image sets de-identified using *k*-Same are *k*-anonymized

- But
  - Changes in "face expression"
  - Changes in "gender"
  - Noticeable blurring…

Extending *k*-Same to "*k*-Same-Select"

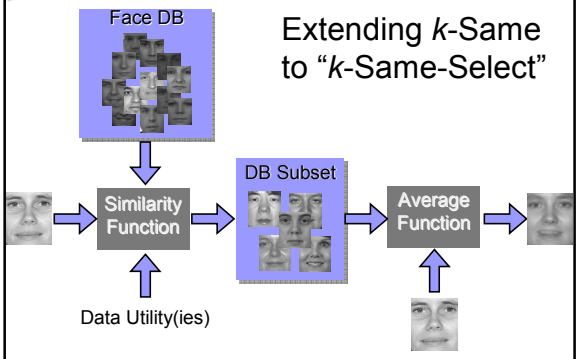

k-Same-Select Results

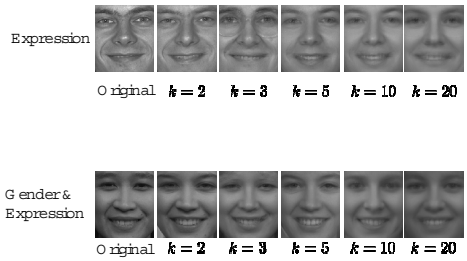

Gender Classification: Ad Hoc

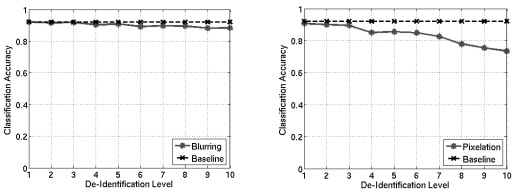

Small performance decrease for blurring
Noticeable decrease for pixelation

Expression Classification

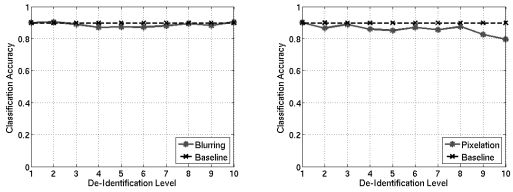

Similar results similar to gender classification

Expression Classification

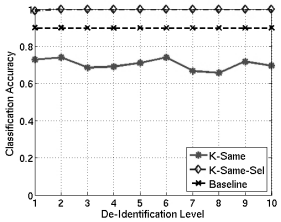

- k-Same *decreases* data utility
- k-Same-Select *increases* data utility

Demonstration Time!

- K-Same Demo
  (http://privacy.cs.cmu.edu/dataprivacy/projects/video/datainfo.html)

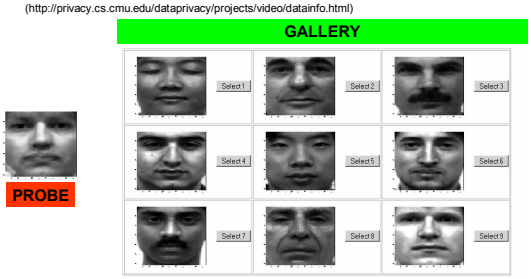

Some Parting Thoughts

- Security + Policy does not guarantee Privacy
- Privacy is not dead, but it requires “intelligence”
- Interdisciplinary approach is necessary
  - Understand policy & law
  - Understand the technology
  - Understand the goals of data use

Thanks!
malin@cs.cmu.edu

- Some slides adapted from:
  - Ralph Gross
  - Elaine Newton
  - Michael Shamos
  - Latanya Sweeney
- More information:
  - http://privacy.cs.cmu.edu
  - http://www.cs.cmu.edu/~malin