

Privacy, Security & Regulations

Matei Zaharia
CS 320



Outline

Security goals and how to think about them

Location data case study

Micro-data and privacy

Protecting privacy

Regulations

Outline

Security goals and how to think about them

Location data case study

Micro-data and privacy

Protecting privacy

Regulations

Why Security & Privacy?

Data is valuable and can cause harm if released

- Example: medical records, purchase history, internal company documents, etc

Data releases can't usually be “undone”

Why Security & Privacy?

It's the law! many regulations & contracts about data:

HIPAA: Health Insurance Portability & Accountability Act (1996)

- Mandatory encryption, access control, training

EU GDPR: General Data Protection Regulation (2018)

- Users can ask to see & delete their data

PCI: Payment Card Industry standard (2004)

- Required in contracts with MasterCard, etc

App Tracking Transparency on iOS (2021)

Some Security Goals

Confidentiality: data is inaccessible to external parties (often via cryptography)

Integrity: data can't be modified by external parties

Privacy: limited info on “individual” users can be learned

Access Control: only the “right” users can perform actions

Auditing: system records an incorruptible logs of activity

Clarifying These Goals

Say our goal was **access control**: only Matei, Steve and James can set CS 320 student grades on Axxess

What scenarios should Axxess protect against?

1. Bobby T. (an evil student) logging into Axxess as himself and being able to change grades
2. Bobby sending hand-crafted network packets to Axxess to change grades
3. Bobby getting a job as a DB admin at Axxess
4. Bobby guessing Matei's password
5. Bobby blackmailing Matei to change his grade
6. Bobby discovering a flaw in encryption algorithms to do #2

Threat Models

To meaningfully reason about security, need a **threat model**:
what types of adversaries we want to defend against

For example, in our Axess scenario, assume:

- Adversaries only interact with Axess through its public HTTPS website
- No crypto algorithm or software bugs
- No password theft

Implementing complex security policies can be
hard even with these assumptions!

Threat Models


No useful threat model can cover everything

- Goal is to cover the most feasible scenarios for adversaries to increase the **cost** of attacks

Threat models also let us divide security tasks across different components

- E.g. auth system handles passwords, 2FA

Threat Models



Outline

Security goals and how to think about them

Location data case study

Micro-data and privacy

Protecting privacy

Regulations

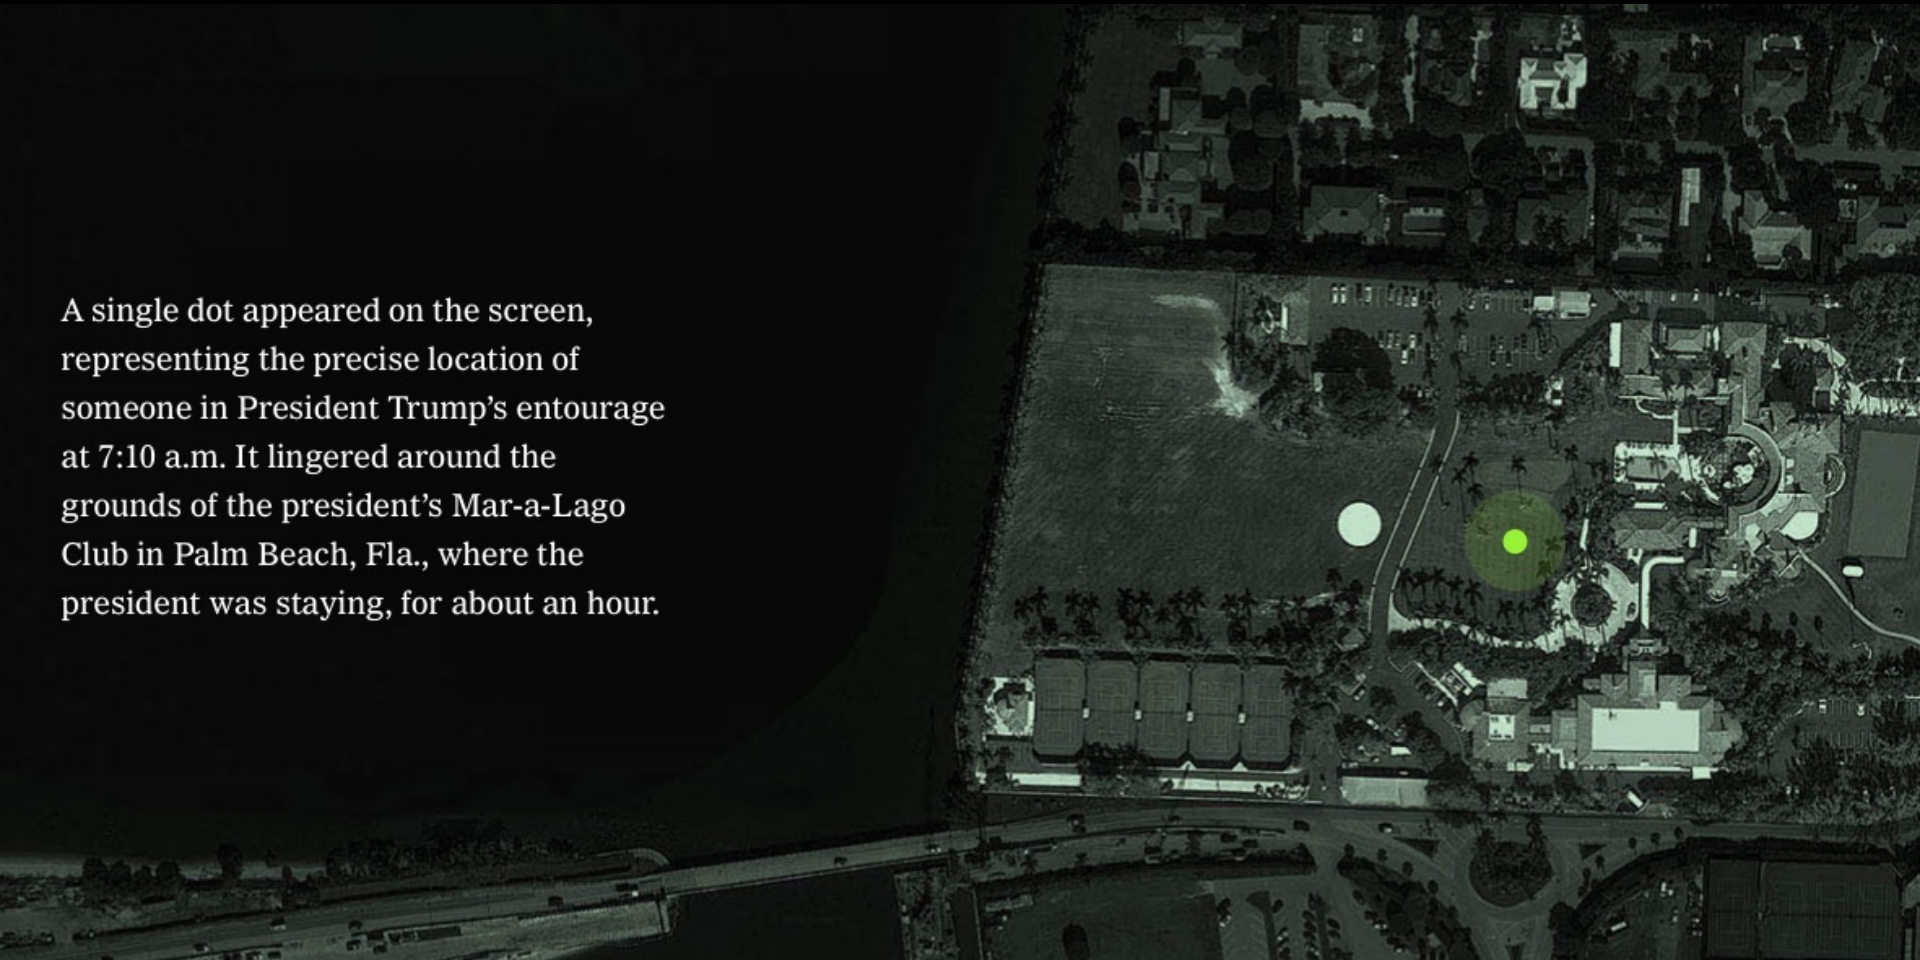


... at the White House ...

Where is This Data From?

What Threat Models are Relevant Here?

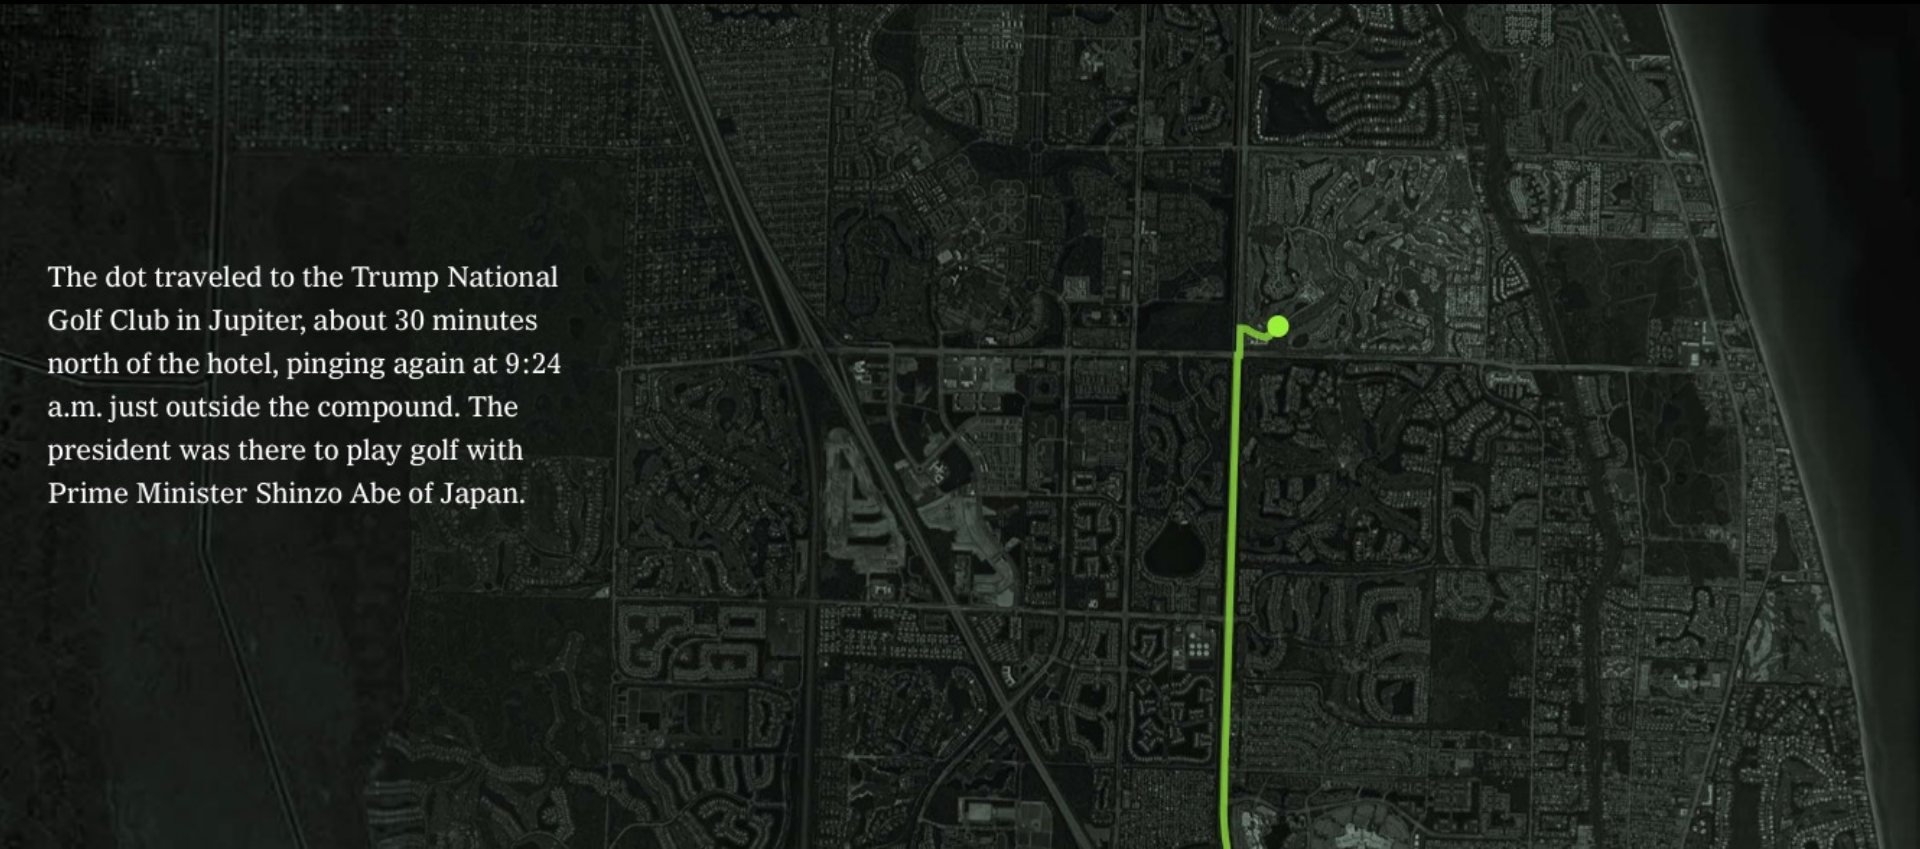
A single dot appeared on the screen, representing the precise location of someone in President Trump's entourage at 7:10 a.m. It lingered around the grounds of the president's Mar-a-Lago Club in Palm Beach, Fla., where the president was staying, for about an hour.



Then it was on the move.



The dot traveled to the Trump National Golf Club in Jupiter, about 30 minutes north of the hotel, pinging again at 9:24 a.m. just outside the compound. The president was there to play golf with Prime Minister Shinzo Abe of Japan.



Outline

Security goals and how to think about them

Location data case study

Micro-data and privacy

Protecting privacy

Regulations

Location Isn't the Only Sensitive Data

The key problem with location data was that very few people have a specific **combination** of locations over time

- Would making the data more coarse-grained help?

Any dataset with records about individuals (**micro-data**) is at risk of these types of attacks

Robust De-anonymization of Large Sparse Datasets

Arvind Narayanan and Vitaly Shmatikov

The University of Texas at Austin

Abstract

We present a new class of statistical de-anonymization attacks against high-dimensional micro-data, such as individual preferences, recommendations, transaction records and so on. Our techniques are robust to perturbation in the data and tolerate some mistakes in the adversary's background knowledge.

We apply our de-anonymization methodology to the Netflix Prize dataset, which contains anonymous movie ratings of 500,000 subscribers of Netflix, the world's largest online movie rental service. We demonstrate that an adversary who knows only a little bit about an individual subscriber can easily identify this subscriber's record in the dataset. Using the Internet Movie Database as the source of background knowledge, we successfully identified the Netflix records of known users, uncovering their apparent political preferences and other potentially sensitive information.

and sparsity. Each record contains many attributes (*i.e.*, columns in a database schema), which can be viewed as dimensions. Sparsity means that for the average record, there are no “similar” records in the multi-dimensional space defined by the attributes. This sparsity is empirically well-established [7, 4, 19] and related to the “fat tail” phenomenon: individual transaction and preference records tend to include statistically rare attributes.

Our contributions. Our first contribution is a formal model for privacy breaches in anonymized micro-data (section 3). We present two definitions, one based on the probability of successful de-anonymization, the other on the amount of information recovered about the target. Unlike previous work [25], we do not assume *a priori* that the adversary's knowledge is limited to a fixed set of “quasi-identifier” attributes. Our model thus encompasses a much broader class of de-anonymization attacks than simple cross-database correlation.

Our second contribution is a very general class of de-anonymization algorithms, demonstrating the fundamental limits of privacy in public micro-data (section 4).

The similarity measure **Sim** is a function that maps a pair of attributes (or more generally, a pair of records) to the interval $[0, 1]$. It captures the intuitive notion of two values being “similar.” Typically, **Sim** on attributes will behave like an indicator function. For example, in our analysis of the Netflix Prize dataset, **Sim** outputs 1 on a pair of movies rated by different subscribers if and only if both the ratings and the dates are within a certain threshold of each other; it outputs 0 otherwise.

To define **Sim** over two records r_1, r_2 , we “generalize” the cosine similarity measure:

$$\text{Sim}(r_1, r_2) = \frac{\sum \text{Sim}(r_{1i}, r_{2i})}{|\text{supp}(r_1) \cup \text{supp}(r_2)|}$$

Definition 1 (Sparsity) *A database D is (ϵ, δ) -sparse w.r.t. the similarity measure **Sim** if*

$$\Pr_r[\text{Sim}(r, r') > \epsilon \forall r' \neq r] \leq \delta$$

As a real-world example, in fig. 1 we show that the Netflix Prize dataset is overwhelmingly sparse. For the vast majority of records, there isn’t a *single* record with similarity score over 0.5 in the entire 500,000-record dataset, even if we consider only the sets of movies rated without taking into account numerical ratings or dates.

Threat Models for Micro-Data

1. Using the micro-data alone (adversary has no additional information about individual(s) of interest)
 - What attacks might someone do here?
2. Using additional, external data about an individual (e.g. Donald Trump met Shinzo Abe at 9 AM)

Protecting Micro-Data

Many forms of perturbation, etc don't work well because the dataset remains sparse (according to previous definition)

Conclusion: To a first approximation, any personally identifiable micro-data is a significant security risk



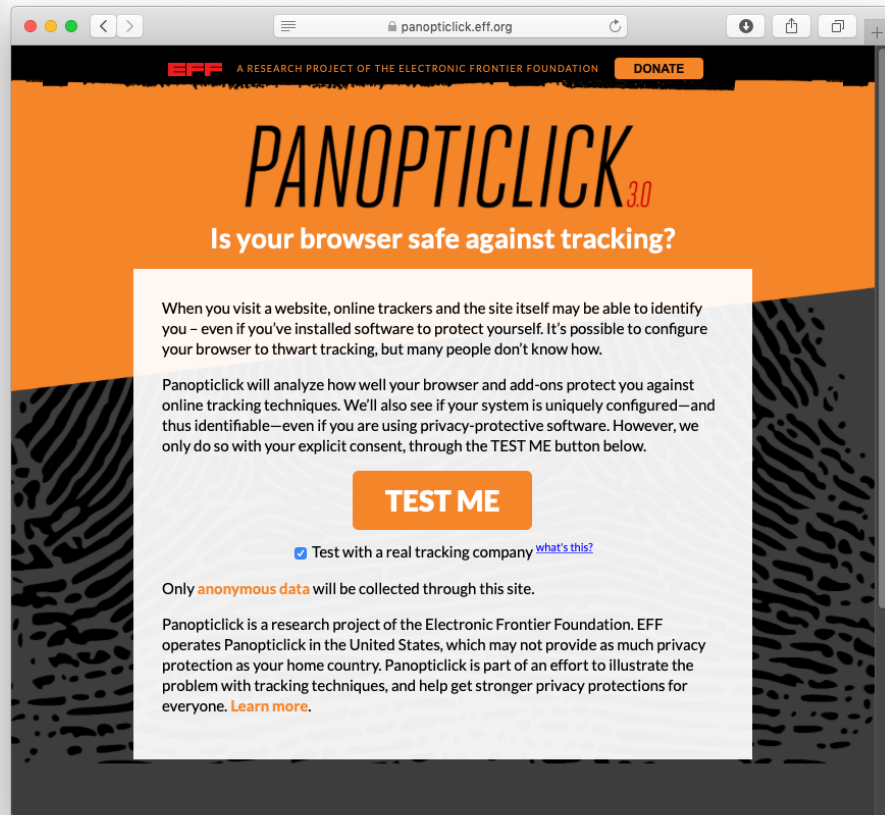
How Unique Is Your Web Browser?

Peter Eckersley*

Electronic Frontier Foundation,
pde@eff.org

Abstract. We investigate the degree to which modern web browsers are subject to “device fingerprinting” via the version and configuration information that they will transmit to websites upon request. We implemented one possible fingerprinting algorithm, and collected these fingerprints from a large sample of browsers that visited our test side, panopticlick.eff.org. We observe that the distribution of our fingerprint contains at least 18.1 bits of entropy, meaning that if we pick a browser at random, at best we expect that only one in 286,777 other browsers will share its fingerprint. Among browsers that support Flash or Java, the situation is worse, with the average browser carrying at least 18.8 bits of identifying information. 94.2% of browsers with Flash or Java were unique in our sample.

By observing returning visitors, we estimate how rapidly browser fingerprints might change over time. In our sample, fingerprints changed quite rapidly, but even a simple heuristic was usually able to guess when a fingerprint was an “upgraded” version of a previously observed browser’s fingerprint, with 99.1% of guesses correct and a false positive rate of only 0.86%.



Outline

Security goals and how to think about them

Location data case study

Micro-data and privacy

Protecting privacy

Regulations

Approach 1: Don't Hold Too Much Data

Only collect data that the business actually needs

Only hold data for a limited time

Automatically redact sensitive data

Let users ask for data to be deleted

Form of Payment:

VISA

Last Four Digits 1913


GDPR, PCI, etc require some of these measures!

Access Misconfiguration for Customer Support Database

[MSRC](#) / [By MSRC Team](#) / [January 22, 2020](#) / [Misconfiguration](#), [Privacy](#)

Today, we concluded an investigation into a misconfiguration of an internal customer support database used for Microsoft support case analytics. While the investigation found no malicious use, and although most customers did not have personally identifiable information exposed, we want to be transparent about this incident with all customers and reassure them that we are taking it very seriously and holding ourselves accountable.

Our investigation has determined that a change made to the database's [network security group](#) on December 5, 2019 contained misconfigured [security rules](#) that enabled exposure of the data. Upon notification of the issue, engineers remediated the configuration on December 31, 2019 to restrict the database and prevent unauthorized access. This issue was specific to an internal database used for support case analytics and does not represent an exposure of our commercial cloud services.



As part of Microsoft's standard operating procedures, data stored in the support case analytics database is redacted using automated tools to remove personal information. Our investigation confirmed that the vast majority of records were cleared of personal information in accordance with our standard practices. In some scenarios, the data may have remained unredacted if it met specific conditions. An example of this occurs if the information is in a non-standard format, such as an email address separated with spaces instead of written in a standard format (for example, "XYZ @contoso com" vs "XYZ@contoso.com"). We have begun notifications to customers whose data was present in this redacted database.

Approach 2: Security Controls

Encrypt data at rest and on networks (HIPAA, PCA, ...)

Limit physical access to facilities

Limit which employees can access data

Audit all accesses

Implement external audits, penetration tests, etc

Approach 3: Privacy-Preserving Analysis

Strong security controls only get us so far: some analysts still need to access the data, and might learn stuff about users

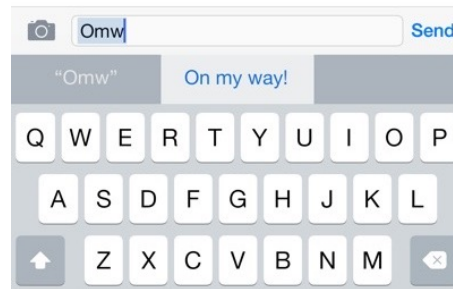
- Also, infeasible to share data with other organizations

Can we design schemes that allow **aggregate analysis** over a dataset without revealing individual information?

Examples


External app: iOS auto-complete

- Want to learn across many users, but make sure model leaks no personal info



Internal app: limit access of users analyzing health data, census data, ...

Data sharing: publish anonymized data



First Question: How to Define Privacy?

k-anonymity: the data record released for each individual cannot be distinguished from at least $k-1$ others


Differential privacy: results of a query are similar whether or not a particular individual is in the database

Many other definitions...

First Question: How to Define Privacy?


Probably need to look at threat models:

- Does the adversary know anything auxiliary about the individuals?
- Does the adversary know all data about all individuals except the one they are trying to find information about?



Stronger model; we'll explore this one in the context of differential privacy

Example Threat Model



- Database software is working correctly
- Adversaries have limited # of user accounts
- Adversaries can know all data except one field of user in question

How to Define Privacy?

This is conceptually very tricky! How to distinguish between

`SELECT TOP(disease) FROM patients WHERE state="California"`
and

`SELECT TOP(disease) FROM patients WHERE name="Matei"`

How to Define Privacy?

Also want to defend against adversaries who have some side-information; for instance:

```
SELECT TOP(disease) FROM patients WHERE  
birth_year="19XX" AND gender="M" AND born_in="Romania"  
AND ...
```

 Side information about Matei


Also consider adversaries who do multiple queries (e.g. subtract 2 results)

Differential Privacy

Privacy definition that tackles these concerns and others by looking at **possible databases**

- Idea: results that an adversary saw should be “nearly as likely” for a database without Matei

A randomized algorithm M is ϵ -differentially private if for all $S \subseteq \text{Range}(M)$ and all sets A, B that differ in 1 element,

$$\Pr[M(A) \in S] \leq \Pr[M(B) \in S] e^\epsilon$$


$\approx 1+\epsilon$


What Does It Mean?

Say an adversary runs some query and observes a result X

Adversary had some set of results, S , that lets them infer something about Mate_i if $X \in S$

Then:

$$\Pr[X \in S \mid \text{Mate}_i \in \text{DB}] \leq e^\epsilon \Pr[X \in S \mid \text{Mate}_i \notin \text{DB}]$$

 $\approx 1+\epsilon$

and

$$\Pr[X \notin S \mid \text{Mate}_i \in \text{DB}] \leq e^\epsilon \Pr[X \notin S \mid \text{Mate}_i \notin \text{DB}]$$

Similar outcomes whether or not Mate_i in DB

What Does It Mean?

Example (assume $\epsilon=0.1$):

```
SELECT TOP(diagnosis) FROM patients WHERE age<35  
AND city="Palo Alto" → flu
```

```
SELECT TOP(diagnosis) FROM patients WHERE age<35  
AND city="Palo Alto" AND born="Romania" → drug overdose
```

Does this mean Matei specifically takes drugs?

- Result would have been nearly as likely (within 10%) even if Matei were not in the database
- Could be that we just got a low-probability result
- Could be that *most* Romanians do drugs (no info on Matei)

Nice Property of Differential Privacy

Composition: can compute the privacy effect of multiple (even dependent) queries

Let queries M_i each provide ϵ_i -differential privacy; then the sequence of queries $\{M_i\}$ provides $(\sum_i \epsilon_i)$ -differential privacy

Disadvantages of Differential Privacy

Each user can only make a limited number of queries (more precisely, limited total ϵ) in the model we described

- There are schemes that bound total ϵ but limit possible queries

How to set ϵ in practice?

- Hard to tell what various values mean, though there is a nice Bayesian interpretation
- Apple set $\epsilon=6$ and researchers said it's too high

Can't query using arbitrary code (must know ϵ)

Another Use of Differential Privacy

“Randomized response”: clients add noise to data they send instead of relying on provider




Example: statistics collection about iOS features

Other Private Computing Tools

Secure multi-party computation

Computing on encrypted data
(homomorphic encryption)

Hardware enclaves



Outline

Security goals and how to think about them

Location data case study

Micro-data and privacy

Protecting privacy

Regulations

Data Protection Regulations

EU GDPR, California Consumer Protection Act (CCPA), HIPAA, and many industry standards

Major concerns:

- Geographic location of data
- Declaring and minimizing uses
- Enforcing security best practices
- Letting users see & delete own data

GDPR: Who Are the Participants?

Data subjects

Data controllers

Data processors

GDPR: What is Personal Data?

GDPR: Key Provisions

State the purposes of data collection

Limit data processing to a specific set of lawful purposes

Have privacy controls and use highest privacy levels by default

Store EU data inside the EU

Report data breaches

Rights for data subjects: access, information, rectification, erasure, limiting automated processing

Assign a Data Protection Officer (DPO)

Exercise: Purposes of Data Collection



Deutsche Bank

Lawful Bases for Data Processing

Consent

Contract

Public task

Vital interest

Legitimate interest (unless conflicting with data subject)

Legal requirement

Rights for Data Subjects

Access / portability

Information

Rectification

Erasure

Automated processing

GDPR Lawsuits

Date ↕	Organisation ↕	Amount ▼	Issued by ↕	Reason(s)
2021-06-16	Amazon Europe Core Srl	€746,000,000	Luxembourg (CNPD)	The largest fine for violating GDPR to date. ^{[71][72]}
2021-09-02	WhatsApp Ireland Ltd	€225 M	Ireland	^[73]
2019-07-08	British Airways	£183,000,000	UK (ICO)	Use of poor security arrangements that resulted in a 2018 web skimming attack affecting 500,000 consumers. ^{[24][25][26]} Was later reduced to £20 million
2020-12-10	Google LLC	€60,000,000	France (CNIL)	Deposit of cookies without obtaining consent, lack of information provided to users and defective "opposition" mechanism ^[48]
2019-01-21	Google LLC	€50,000,000	France (CNIL)	Insufficient transparency, control, and consent over the processing of personal data for the purposes of behavioural advertising . ^{[4][5]}
2020-12-10	Google Ireland Limited	€40,000,000	France (CNIL)	Deposit of cookies without obtaining consent, lack of information provided to users and defective "opposition" mechanism ^[48]
2020-10-01	H&M	€35,300,000	Germany (HmbBfDI)	Illegal surveillance of several hundred employees ^[46]
2020-12-10	Amazon Europe Core Srl	€35,000,000	France (CNIL)	Deposit of cookies without obtaining consent and lack of information provided to users ^[47]
2020-01-15	TIM S.p.A.	€27,800,000	Italy (GPDP)	Unlawful processing for marketing purposes ^[42]
2020-10-30	Marriott International	£18,400,000	UK (ICO)	Failure to keep millions of customers' personal data secure ^[27]
2019-12-09	1&1 Ionos	€9,550,000	Germany (BfDI)	Insufficient protection of personal data, failing to put "sufficient technical and organizational measures" in place to protect customer data in its call centers. Violation of article 32 of GDPR ^[40]

The decision, [revealed by Bloomberg](#), suffers from no ambiguity: **the targeted ad system that Amazon forces onto us is not based on free consent**, which is a violation of the GDPR. As such, the corporation is fined to the tune of 746 million euros. This is a new European record for breaching GDPR rules (the previous high-mark was the 50 million euros fine the CNIL, the French DPA, levied against Google, again as a result of our collective legal action).

The Data Protection Commission (DPC) has today announced a conclusion to a GDPR investigation it conducted into WhatsApp Ireland Ltd. The DPC's investigation commenced on 10 December 2018 and it examined whether WhatsApp has discharged its GDPR transparency obligations with regard to the provision of information and the transparency of that information to both users and non-users of WhatsApp's service. This includes information provided to data subjects about the processing of information between WhatsApp and other Facebook companies.

GDPR Lawsuits

2019-06-18	Sergic (real estate services)	€400,000	France (CNIL)	Failure to implement appropriate security measures; failure to define appropriate data retention periods for the personal data of unsuccessful rental candidates. ^[18]
2019-12-17	Doorstep Dispensaree	£275,000	UK (ICO)	"cavalier attitude to data protection", having left 500,000 patient records in an unsecured location ^[41]
2019-06	La Liga	€250,000	Spain (AEPD)	Poorly disclosing purpose for requesting GPS and microphone permissions within the football league's mobile app . When the app was open, it transmitted the user's location if it detected an acoustic fingerprint embedded within game telecasts. This was used to help pinpoint the locations of venues that may be screening the games from unauthorized feeds . ^{[14][15]}

2019-05-28	Unnamed Belgian mayor	€2,000	Belgium (GBA/APD)	Misuse of personal data collected for local administrative purposes for election campaign purposes. ^[13]
2019-03-07	Unnamed bank	€1,560	Hungary (NAIH)	Failure to erase and correct data at the request of the data subject. ^[6]
2019-03-07	Unnamed debt collector	€1,560	Hungary (NAIH)	Breaching the principles of transparency and data minimisation. ^[7]
2019-06-18	Unnamed police officer	€1,400	Germany (LfDI)	Autonomously processing personal data for non-legal purposes. ^[17]

Downsides of GDPR

For businesses?

For individuals?

Summary

Security and privacy are essential concerns for businesses and products based on data

Threat models are a systematic way to measure security and reason about designs

Micro-data is generally not private at all; need to actively protect this data if we want to protect users

Big Data for Security


Data collection is cheap, so log and audit everything

- Flows sent on network, remote procedure calls, authentication calls, logins, etc
- Hard for adversaries to hide

“Security Information Management” (SIM) systems are products for this purpose

Attacks on Machine Learning

Adversarial inputs to ML models



Extracting training data from a trained model




Poisoning models with bad data

Figure 1: An image recovered using a new model inversion attack (left) and a training set image of the victim (right). The attacker is given only the person's name and access to a facial recognition system that returns a class confidence score.

New Security Tools

Enclaves and confidential computing



Computing on encrypted data
(homomorphic encryption)

Secure multi-party computation

