DIFFERENTIAL PRIVACY:

a new notion of the privacy for a new paradigm of data?

Artificial Intelligence and Agnostic Science - Doing Science in the Age of Artificial Intelligence

Edwige Cyffers

UNIVERSITÉ PARIS 1
PANTHÉON SORBONNE

A context of scandals

- Governmental surveillance
- Pricing at scale
- Sensible data leakage
- "Privacy is dead"

Rising subject in CS

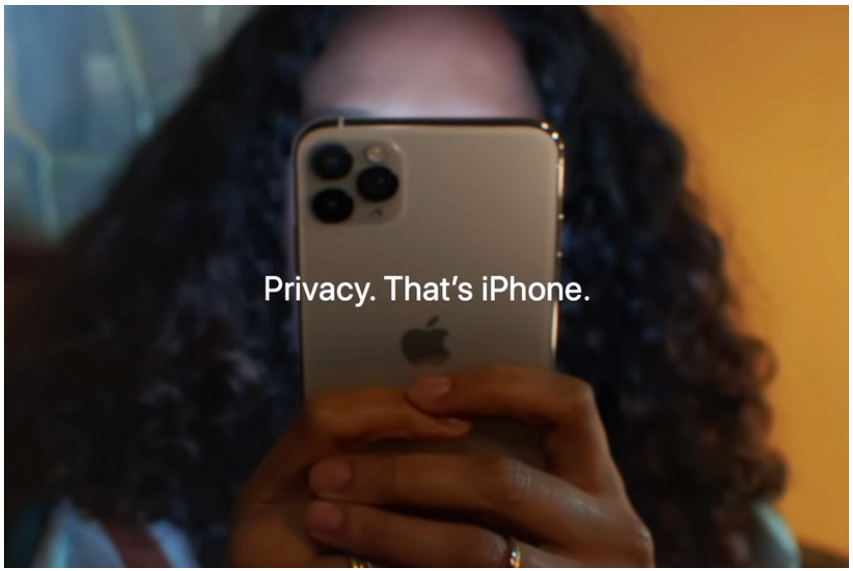

- Growing awareness that forces companies to adapt
- All the giants: Google, Apple,...
- Regulations
- Individual scale

PRIVACY

A traditional concept

Definition

the right that someone has to keep their personal life or personal information secret or known only to a small group of people

- uncountable noun
- IRL concept

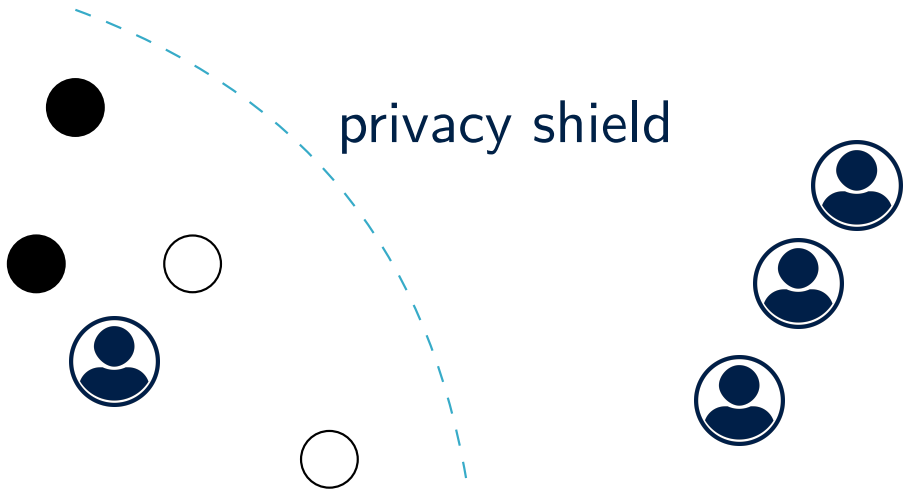

Anonymity

- A way to achieve privacy
- Free speech scenario, link with equality

set of individuals

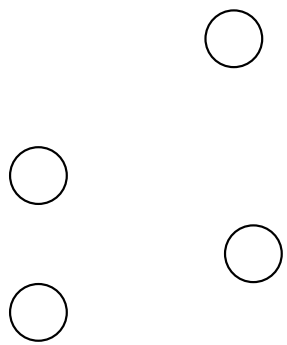

public facts

GDPR

Personal data – Art 4 GDPR

personal data means any information relating to an identified or identifiable natural person (data subject); an identifiable natural person is one who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, an identification number, location data, an online identifier or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person;

- Basis of regulation: "data subject"-based
- National regulations

MACHINE LEARNING

Processing Data

Processing – Art 4 GDPR

processing means any operation or set of operations which is performed on personal data or on sets of personal data, whether or not by automated means, such as collection, recording, organisation, structuring, storage, adaptation or alteration, retrieval, consultation, use, disclosure by transmission, dissemination or otherwise making available, alignment or combination, restriction, erasure or destruction;

- Profiling
- Statistics
- Health
- and every other subject

Inversion of data role

- From a data processed by an algorithm to an algorithm shaped by the data
- Algorithm do have prior on the data, but the exact data dependence is not known before the training

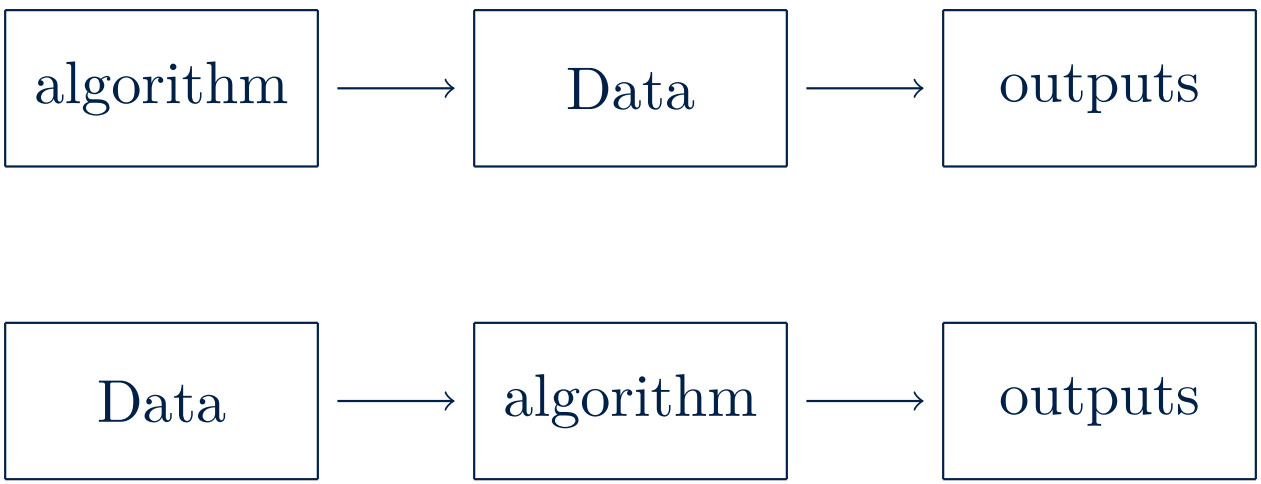

Generalization

Data uncovers a model that will fit any other points

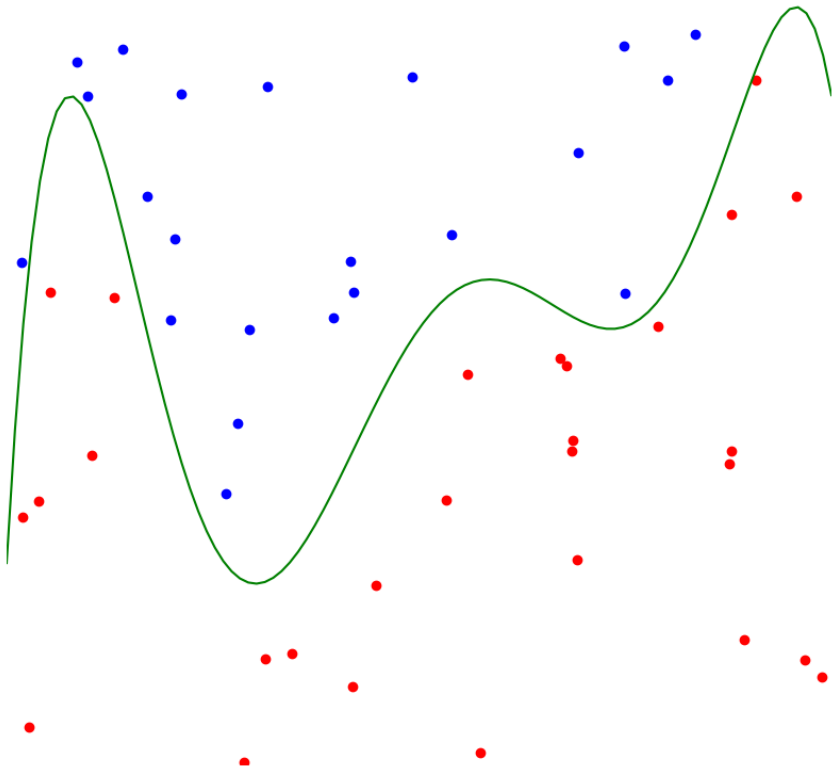

Generalization

Data uncovers a model that will fit any other points

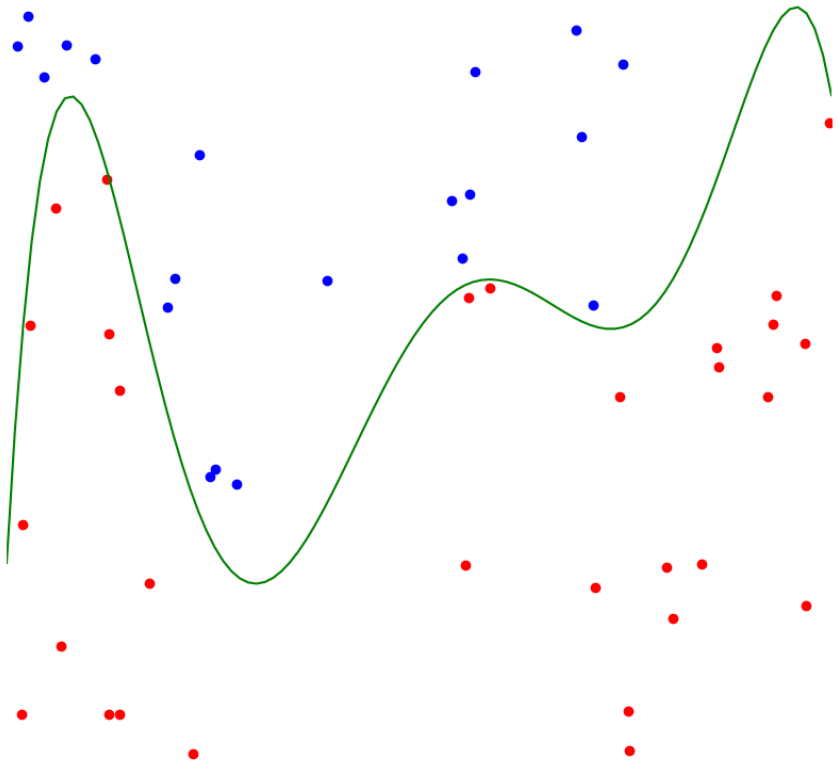

END OF ANONYMIZATION

Removal of Personally Identifiable Information

- Intuition of anonymization: the best of the two world, we keep the value of the data and cut the link with the owner

Id	Name	X	Y	Z
19385923	P. Williams	b	.13	TRUE
19385924	A. Smith	c	.64	TRUE
19385925	J. Brown	c	.92	FALSE
19385926	C. Doe	a	.15	FALSE
19385927	E. Garcia	b	.53	TRUE

Removal of Personally Identifiable Information

- Intuition of anonymization: the best of the two world, we keep the value of the data and cut the link with the owner

Id	Name	X	Y	Z
19385923	P. Williams	b	.13	TRUE
19385924	A. Smith	c	.64	TRUE
19385925	J. Brown	c	.92	FALSE
19385926	C. Doe	a	.15	FALSE
19385927	E. Garcia	b	.53	TRUE

Failure of Pseudoanonymization

Robust De-anonymization of Large Datasets

We demonstrate that an adversary who knows only a little bit about an individual subscriber can easily identify this subscribers 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

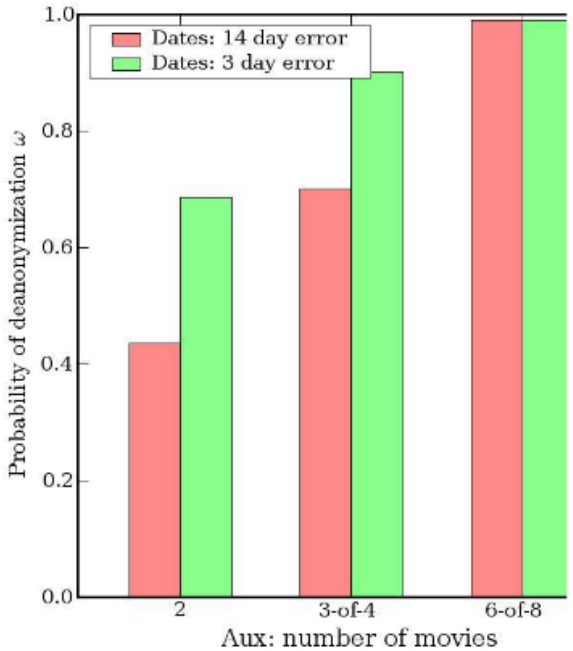

DIFFERENTIAL PRIVACY

Framework

High level idea Let's see how much a specific input is revealed by the algorithm

- An Algorithm $\mathcal{M}$
- Database X
- A neighboring relation $\sim$

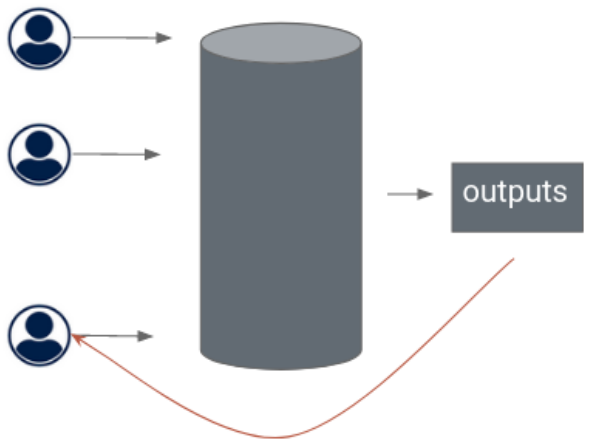

Definition

Differential Privacy

Let $\varepsilon \leq 0$. A mechanism $\mathcal{M} : \mathcal{X} \to \mathbb{P}(Z)$ is ε-differentially private with respect to the neighboring relation $\sim$ if for every pair of databases $X \sim X'$ and every measurable subset $S \subset Z$:

$$\mathbb{P}(\mathcal{M}(x) \in \mathcal{S}) \leq \exp(\varepsilon)\mathbb{P}(\mathcal{M}(y) \in \mathcal{S})$$

- "Distance" between distribution
- Rényi differential privacy $\mathcal{D}_\alpha(\mathcal{M}(X)||\mathcal{M}(X')) \leq \varepsilon$

Key Properties

- Composition
- Robust to post processing
- Participation based

Participation based

- Economic interpretation
- Example of Randomized response
- Similar to the "fair" deal Data versus Service

Algorithm 1: Randomized response with coin tossing.

```
Toss a coin;
if tails then
    Answer truthfully
else
    Toss a coin;
    if tails then
        Answer Yes;
    else
        Answer No;
```

DIFFERENTIAL PRIVACY IN CONTEXT

Move beyond individual based framework

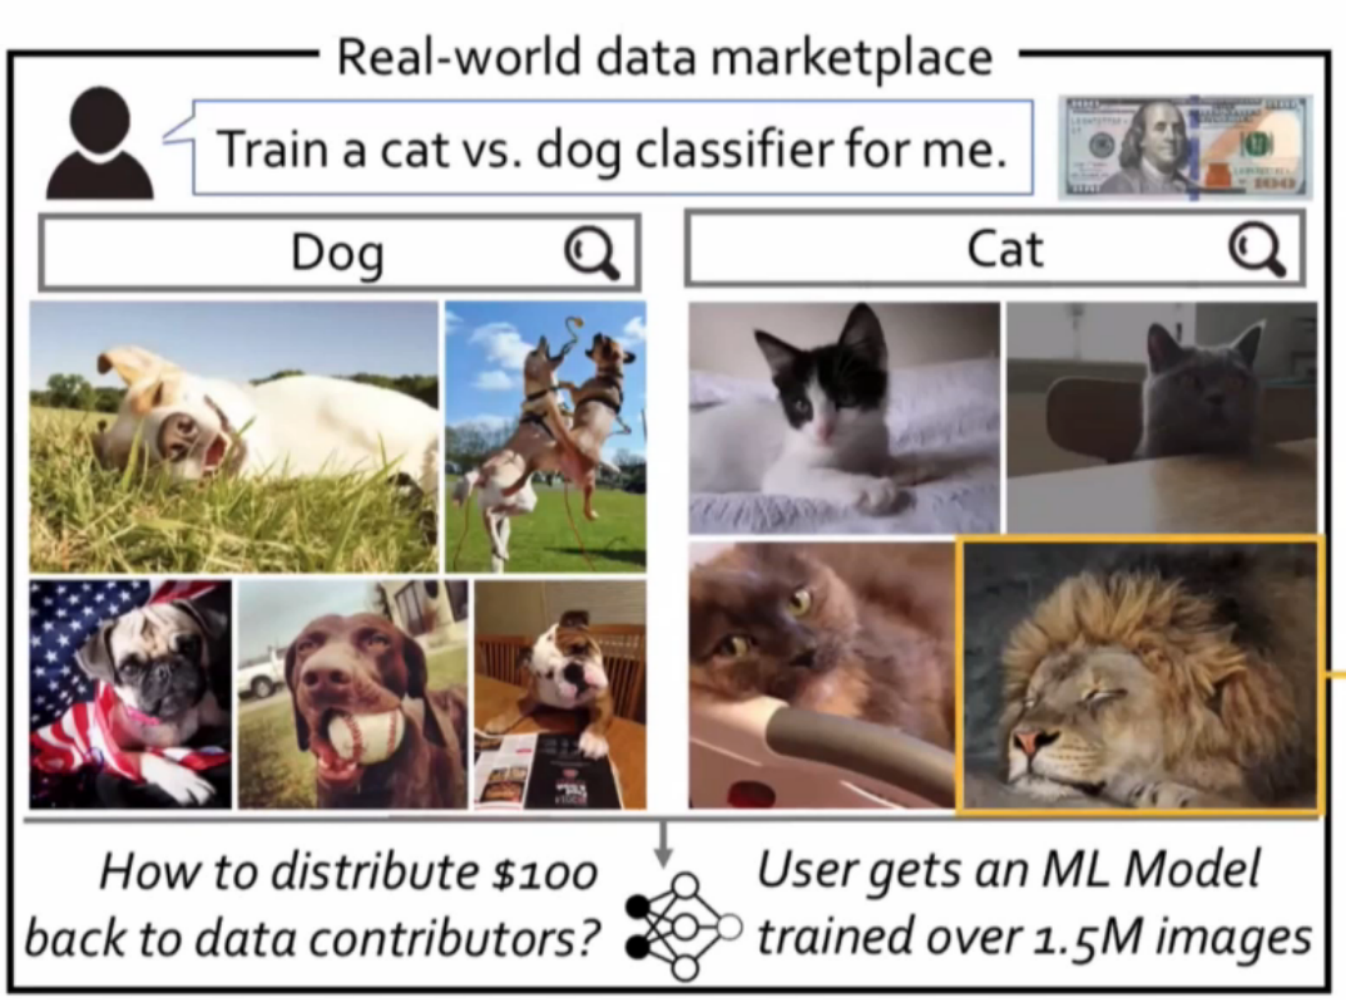

Keynote Talk of Dawn Song, NeurIPS Workshop SpicyFL

Individual based protection ?

False sense of security as features don't match who we think we are

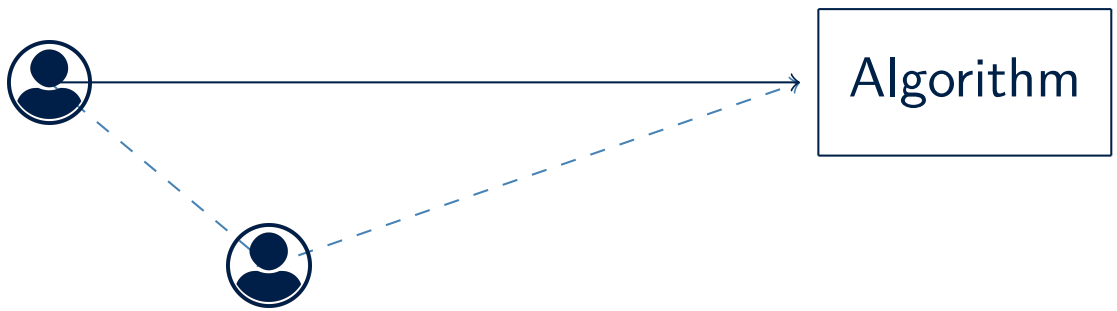

Real applications

- Google, Apple and Microsoft
- US Census

The U.S. Census Bureau Adopts Differential Privacy

John M. Abowd
Chief Scientist and Associate Director for Research and Methodology
U.S. Census Bureau
24th ACM SIGKDD Conference on Knowledge Discovery and Data Mining
London, United Kingdom
August 23, 2018

United States Census Bureau | U.S. Department of Commerce
Economics and Statistics Administration
U.S. CENSUS BUREAU
census.gov

Summary

- Change of scale in term of quantity and possibility of automated processing
- Privacy in mathematical terms
- Privacy guarantees towards a individual that must be defined

Open questions

- Which criterion to define the correct neighboring relation ?
- Why is there a focus on differential privacy as a method rather than as a metric ?
- Which part of privacy are missing ? Which mathematical approach for them ?