





Et si au lieu de partager
les données, on partageait
le savoir?



The elementary blocks of data analysis




The elementary blocks of data analysis



Input privacy:
Personal data is not exposed
during computation

The elementary blocks of data analysis




How to guarantee outputs are anonymous?

1/ Remove PII?

Name	Address	Email	Movie #1	Movie #2	Movie #3	Movie #4
------	---------	-------	----------	----------	----------	----------

1/ Remove PII?



The fields are likely to be unique and easy to match to individuals

=> REMOVE

This combination of fields is probably unique, but it would be bad luck if someone uses it to identify someone.

=> KEEP

1/ Remove PII



0105v2 [cs.CR] 22 Nov 2007

Robust De-anonymization of Large Datasets (How to Break Anonymity of the Netflix Prize Dataset)

Arvind Narayanan and Vitaly Shmatikov
The University of Texas at Austin
February 5, 2008

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 that subscriber's record in the dataset. Using the Internet Movie Database as the source of information, we successfully identified the Netflix records of known users, uncovering their names and other potentially sensitive information.

are increasingly becoming... Some datasets

NetFlix Cancels Recommendation Contest

wired.com/2010/03/netflix-cancels-contest/

RYAN SINGEL | BACKCHANNEL | BUSINESS | CULTURE | GEAR | IDEAS | MORE

SECURITY 03.12.2010 02:48 PM

NetFlix Cancels Recommendation Contest After Privacy Lawsuit

Netflix is canceling its second \$1 million Netflix Prize to settle a legal challenge that it breached customer privacy as part of the first contest's race for a better movie-recommendation engine. Friday's announcement came five months after Netflix had announced a successor to its algorithm-improvement contest. The company at the time said it intended to [...]

f t e

2/ Aggregate?

- ▶ As long as there are unique rows, there is re-identification risk
- ▶ Aggregation does improve things but:
 - ▷ It destroys most of the potential for machine learning
 - ▷ It is not a silver bullet and can still lead to re-identification (triangulation?)
 - ▷ It requires ad hoc decision making

=> Sharing 'anonymous' data is probably an illusion

A “mathematical” definition of anonymous information

“Anonymous information: information which does not relate to an identified or identifiable natural person.”

“Data which could be attributed to a natural person by the use of additional information should be considered to be information on an identifiable natural person”. ([GDPR, Recital 26](#))

A “mathematical” definition of anonymous information

“Anonymous information: information which does not relate to an identified or identifiable natural person.”

“Data which could be attributed to a natural person by the use of additional information should be considered to be information on an identifiable natural person”. ([GDPR, Recital 26](#))

=> I should not learn something on any given individual

=> It should not depend on what I already know

A mathematical definition of anonymous information: *Differential privacy*

An algorithm A is (ϵ, δ) -differentially private if for any two neighboring datasets D and D' and any event S :

$$\Pr[\mathcal{A}(\mathcal{D}) \in \mathcal{S}] \leq e^\epsilon \Pr[\mathcal{A}(\mathcal{D}') \in \mathcal{S}] + \delta$$


A mathematical definition of anonymous information: *Differential privacy*

An algorithm A is (ϵ, δ) -differentially private if for any two neighboring datasets D and D' and any event S :


$$\Pr[\mathcal{A}(\mathcal{D}) \in \mathcal{S}] \leq e^\epsilon \Pr[\mathcal{A}(\mathcal{D}') \in \mathcal{S}] + \delta$$

$$\Pr[\mathcal{A}(\{\text{data}, \text{Alice}\}) = s] \approx \Pr[\mathcal{A}(\{\text{data}, \text{SpongeBob}\}) = s]$$

Not quite sharing “data”

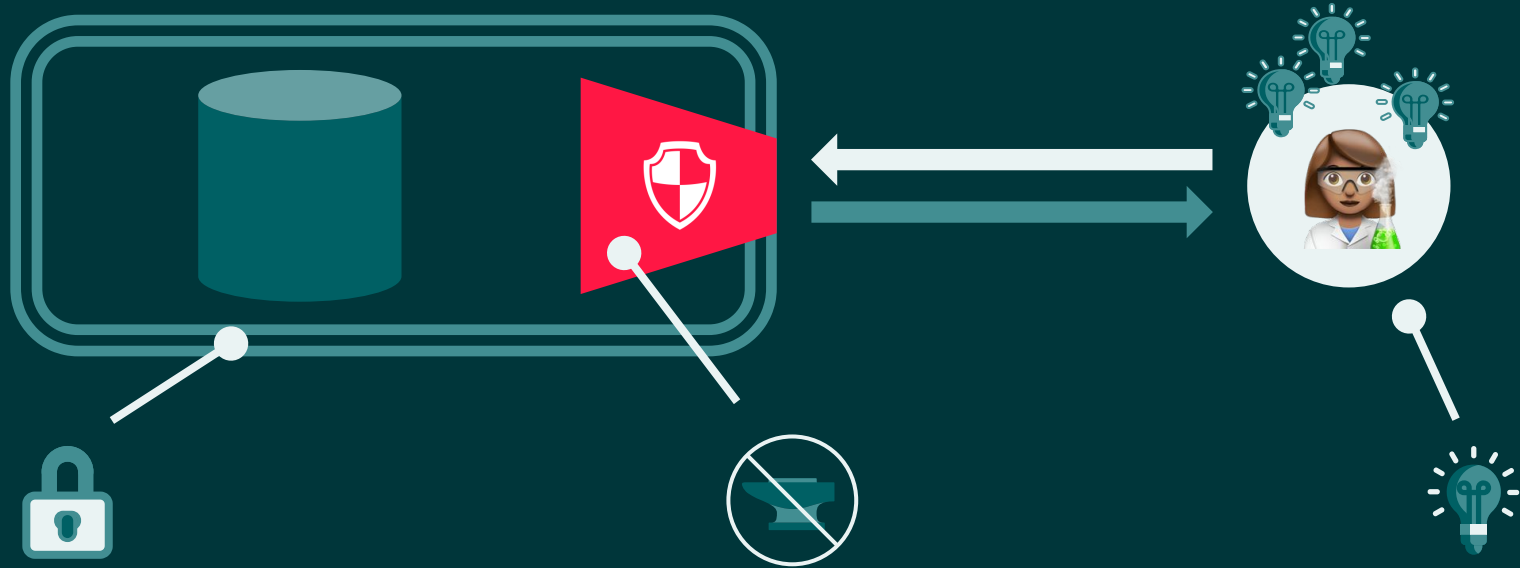


Not quite sharing “data”



- Any data, no matter how sensitive
- Any learning objective
- No assumption on what information may be used

Sarus: Learn from Data You cannot See with Privacy Guarantees




Sarus Vault: A Remote Execution Environment

Differential privacy is built into every API call

Use **Synthetic Data** and native SQL/ML SDK

Sarus: Learn from Data You cannot See with Privacy Guarantees




Sarus Vault: A Remote Execution Environment


Differential privacy is built into every API call

Use **Synthetic Data** and native SQL/ML SDK


Privacy-by-design data-centric organization



Secure collaboration with external partners



Leverage data across organization silos or geographic borders



Thank you!

Maxime Agostini
ma@sarus.tech

We're hiring! 🚀

And opening a private beta, request an
invite! 📺

invite@sarus.tech