The Data Minimization Principle in Machine Learning

Prakhar Ganesh[1,2], Cuong Tran[3], Reza Shokri[5], Ferdinando Fioretto[4]

[1]Mila, [2]McGill University, [3]Dyania Health, [4]University of Virginia, [5]National University of Singapore

The Data Minimization Principle

The Data Minimization Principle

GDPR:

Article 5(1)(b): Personal data shall be collected for specified, explicit and legitimate purposes and not further processed in a manner that is incompatible with those purposes

Article 5(1)(c): Personal data shall be adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed.

The Data Minimization Principle

GDPR:

Article 5(1)(b): Personal data shall be collected for **specified, explicit and legitimate purposes** and **not further processed in a manner that is incompatible with those purposes**

Article 5(1)(c): Personal data shall be adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed.

The Data Minimization Principle

GDPR:

Purpose Limitation

Article 5(1)(b): Personal data shall be collected for **specified, explicit and legitimate purposes** and **not further processed in a manner that is incompatible with those purposes**

Article 5(1)(c): Personal data shall be adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed.

The Data Minimization Principle

GDPR:

Article 5(1)(b): Personal data shall be collected for specified, explicit and legitimate purposes and not further processed in a manner that is incompatible with those purposes

Article 5(1)(c): Personal data shall be **adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed**

The Data Minimization Principle

GDPR:

Article 5(1)(b): Personal data shall be collected for specified, explicit and legitimate purposes and not further processed in a manner that is incompatible with those purposes

Data Relevance

Article 5(1)(c): Personal data shall be **adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed**

The Data Minimization Principle

LGPD:

Article 6: Activities of processing of personal data shall be subject to the following principles,

I: processing done for **legitimate, specific and explicit purposes** of which the data subject is informed, with **no possibility of subsequent processing that is incompatible with these purposes**;

III: limitation of the processing to the **minimum necessary to achieve its purposes**, covering data that are **relevant, proportional and non-excessive in relation to purposes of the data processing**;

The Data Minimization Principle

PIPA:

Article 3(1): The personal information controller shall **specify explicitly the purposes for which personal information is processed**; and shall collect personal information lawfully and fairly to the **minimum extent necessary for such purposes**.

The Data Minimization Principle

Use only the data necessary to achieve the explicitly defined utility.

The Data Minimization Principle

Use only the data necessary to achieve the explicitly defined utility.

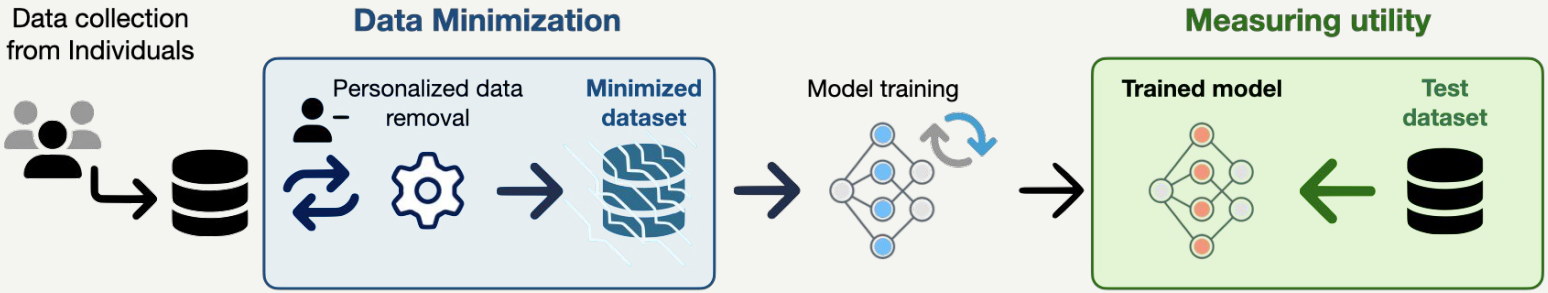

The Data Minimization Principle

Use only the data necessary to achieve the explicitly defined utility.

$$\underset{\boldsymbol{B}\in\{\perp,1\}^{n\times p}}{\text{Minimize}} \quad \|\boldsymbol{B}\|_1 \qquad s.t.: \quad J(\hat{\theta};\mathbf{X},\mathbf{Y}) - J(\overset{\star}{\theta};\mathbf{X},\mathbf{Y}) \leq \alpha \quad \text{(2a)}$$

$$\text{where} \quad \hat{\theta} = \arg\min_{\theta} \frac{1}{n}\sum_{i=1}^{n} \ell\left(f_\theta(\boldsymbol{x}_i \odot \boldsymbol{B}_i), y_i\right) \quad \text{(2b)}$$

Binary Minimization Assumption

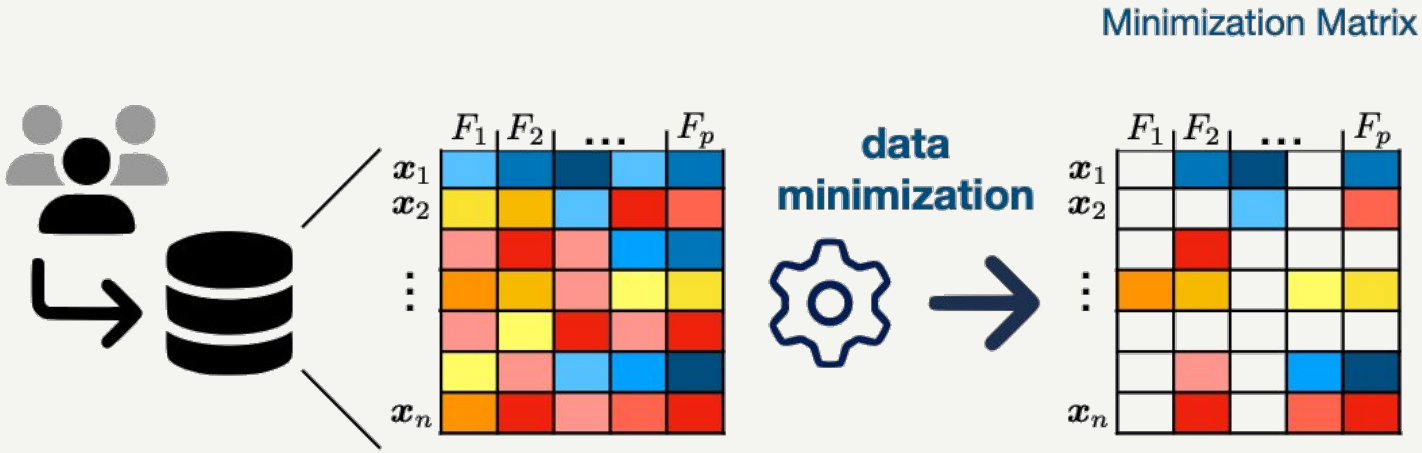

Narrative around DM and Privacy

Narrative around DM and Privacy

GDPR:

Article 5(1)(b): **Personal data** shall be collected for specified, explicit and legitimate purposes and not further processed in a manner that is incompatible with those purposes

Article 5(1)(c): **Personal data** shall be adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed.

Narrative around DM and Privacy

GDPR:
Article 4(1): “personal data” means any information relating to an **identified or identifiable natural person** (“data subject”);

Narrative around DM and Privacy

GDPR:
Article 4(1): “personal data” means any information relating to an **identified or identifiable natural person** (“data subject”);

PIPA:
Article 2(1): The term “personal information” means any of the following information relating to a living individual:
(a) Information that **identifies a particular individual** [...];
(b) Information which, even if it by itself does not identify a particular individual, **may be easily combined with other information to identify a particular individual** [...];

Narrative around DM and Privacy

Data minimization refers to the principle of limiting data collection and retention to the bare minimum necessary to accomplish a given purpose. It's a key principle embedded in privacy laws and regulations, such as the European General Data Protection Regulation (GDPR). Data minimization not only reduces the risk of data breaches, but it also mandates good data governance and enhances consumer trust. In this respect, its importance cannot be overstated.

https://www.kiteworks.com/risk-compliance-glossary/data-minimization/

Strong data minimization provisions in bills in Maine and Vermont are also progressing — Vermont's House just unanimously passed a relatively tough bill — giving privacy advocates hope that after a string of 14 weak state comprehensive data privacy laws with toothless data minimization standards, the tide is finally turning. California is seen as the only state with tough data privacy regulations on the books, but its language is not as strong on the minimization issue as Maryland's will likely be.

https://therecord.media/lawmakers-set-sights-on-data-minimization-with-new-bills

Why Data Minimization is a Key Principle of Data Privacy

https://www.k2view.com/blog/data-minimization/

Narrative around DM and Privacy

EU AI Act (Recital 69):

Updated on 31 July 2024 based on the version published in the Official Journal of the EU dated 12 July 2024 and entered into force on 1 August 2024.

The right to privacy and to protection of personal data must be guaranteed throughout the entire lifecycle of the AI system. In this regard, the principles of data minimisation and data protection by design and by default, as set out in Union data protection law, are applicable when personal data are processed. Measures taken by providers to ensure compliance with those principles may include not only anonymisation and encryption, but also the use of technology that permits algorithms to be brought to the data and allows training of AI systems without the transmission between parties or copying of the raw or structured data themselves, without prejudice to the requirements on data governance provided for in this Regulation.

Narrative around DM and Privacy

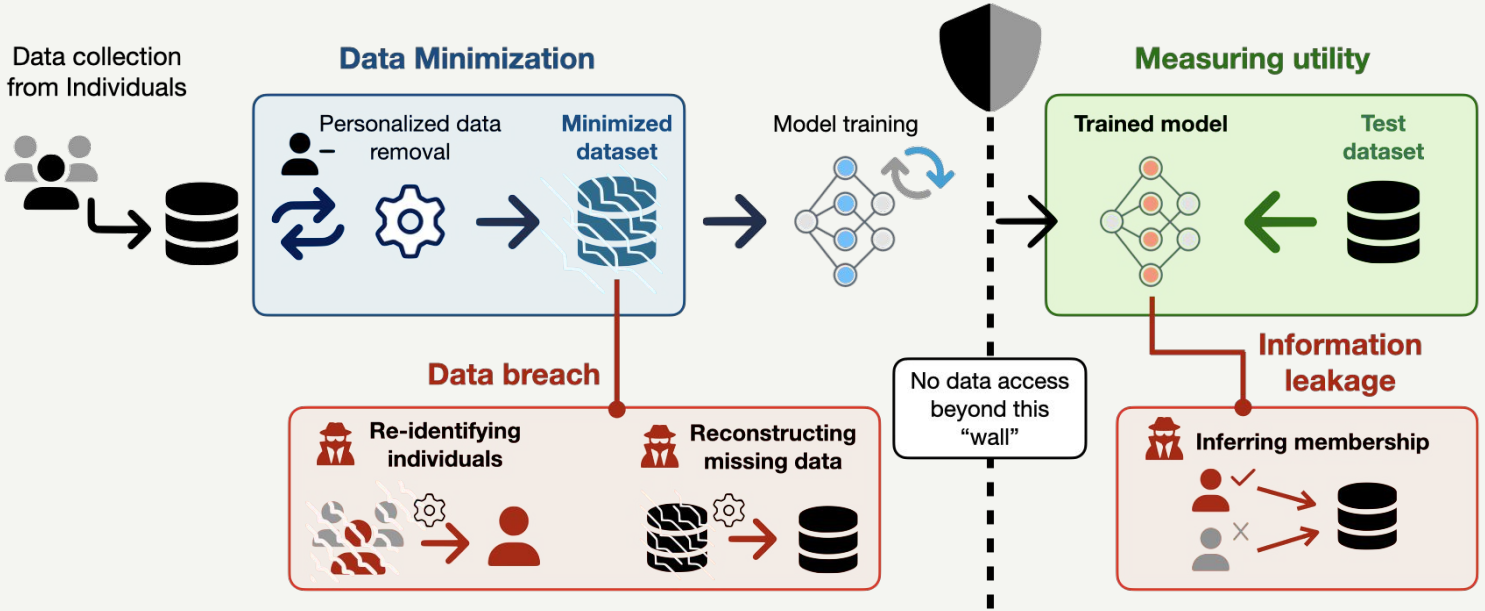

Does DM necessarily protect Privacy?

Does DM necessarily protect Privacy?

Unfortunately, no.

Does DM necessarily protect Privacy?

Name	Zipcode	Age	Gender	Genetic Marker for Cancer
Alice	117068	27	F	Yes
Bob	167056	64	M	No
Charlie	118567	32	M	No
David	191504	81	M	No

Does DM necessarily protect Privacy?

Name	Zipcode	Age	Gender	Genetic Marker for Cancer
Alice	117068	27	F	Yes
Bob	167056	64	M	No
Charlie	118567	32	M	No
David	191504	81	M	No

Who has genetic markers for cancer?
Alice has genetic markers for cancer.

Does DM necessarily protect Privacy?

Name	Zipcode	Age	Gender	Genetic Marker for Cancer
	117068	27	F	Yes
	167056	64	M	No
	118567	32	M	No
	191504	81	M	No

Who has genetic markers for cancer?
Female, age 27, from zipcode 117068 has genetic markers for cancer.

Does DM necessarily protect Privacy?

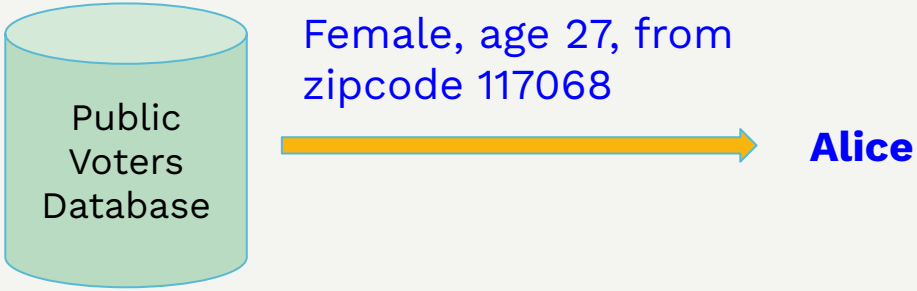

Who has genetic markers for cancer?
Female, age 27, from zipcode 117068 has genetic markers for cancer.

Does DM necessarily protect Privacy?

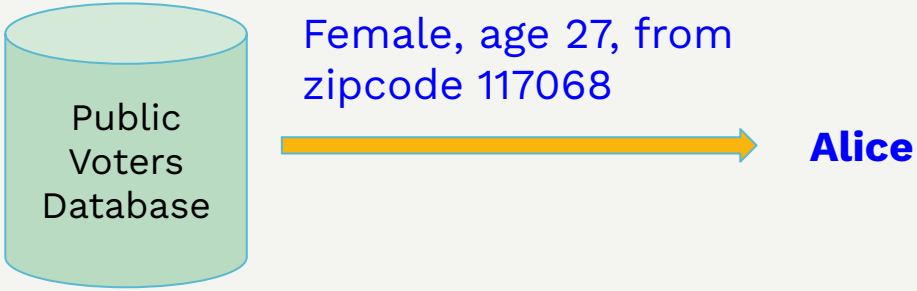

Who has genetic markers for cancer?
Alice has genetic markers for cancer.

Empirical Results: Some Takeaways

Empirical Results: Some Takeaways

- **We formalize the optimization problem for data minimization in machine learning, and show how to minimize significant amounts of data (upto 60-70% data in many cases) without losing utility (less than 1-2% drop in accuracy).**

Empirical Results: Some Takeaways

- **We formalize the optimization problem for data minimization in machine learning, and show how to minimize significant amounts of data (upto 60-70% data in many cases) without losing utility (less than 1-2% drop in accuracy).**

- **We find multiplicity in minimization, i.e., there can be many different 'minimized datasets' of same size that achieve the same utility.**

Empirical Results: Some Takeaways

- **Using real-world adversarial attacks, we highlight the misalignment between the expectations of data minimization and actual benefits of privacy.**

Empirical Results: Some Takeaways

- **Using real-world adversarial attacks, we highlight the misalignment between the expectations of data minimization and actual benefits of privacy.**

- **By explicitly incorporating privacy constraints into the optimization problem, we are able to reduce privacy risks while maintaining the same utility.**

The Data Minimization Principle in Machine Learning

There is a disconnect between the expectation of privacy from data minimization and the actual privacy benefits.
We need to explicitly incorporate privacy into the objectives of data minimization, rather than treating it as an afterthought.

Prakhar Ganesh
Mila, McGill University

Cuong Tran
Dyania Health

Reza Shokri
National University of Singapore

Ferdinando Fioretto
University of Virginia

Paper