Why do we care about privacy?

Katherine Lee

Attack GenAI, Team Lead
Google DeepMind

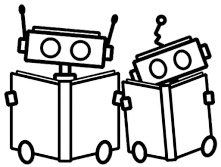

Co-founder
The GenLaw Center

What is privacy?

Comply with US Privacy Laws

The enactment of the California Consumer Privacy Act of 2018 (CCPA) on January 1, 2020 with an enforceability date of July 1, 2020, marked the first comprehensive US state pri
of privacy-related legislation at both the federal and state level followed. Although many of these bills failed to become law, several states have now managed to pass comprehe
Moreover, a federal bill known as the American Data Privacy and Protection Act (ADPPA) is making its way through Congress. The bill is significant as it marks the first federal pr
and bicameral support. If enacted, the ADPPA would preempt the majority of state and local laws, rendering any similar provisions therein invalid.

Utah	Consumer Privacy Act (UCPA)	December 31, 2023

USA State Law Tracker | Portal Advisory | DataGuidance

Definition of ε-differential privacy [edit]

Let ε be a positive real number and $\mathcal{A}$ be a randomized algorithm that takes a dataset as input (representing the actions of the trusted party holding the data).

Let $\operatorname{im} \mathcal{A}$ denote the image of $\mathcal{A}$.

The algorithm $\mathcal{A}$ is said to provide ε-differential privacy if, for all datasets D_1 and D_2 that differ on a single element (i.e., the data of one person), and all subsets S of $\operatorname{im} \mathcal{A}$:

$$\frac{\Pr[\mathcal{A}(D_1) \in S]}{\Pr[\mathcal{A}(D_2) \in S]} \leq e^{\varepsilon},$$

where the probability is taken over the randomness used by the algorithm.[11]

Differential privacy offers strong and robust guarantees that facilitate modular design and analysis of differentially private mechanisms due to its composability, robustness to post-processing, and graceful degradation in the presence of correlated data.

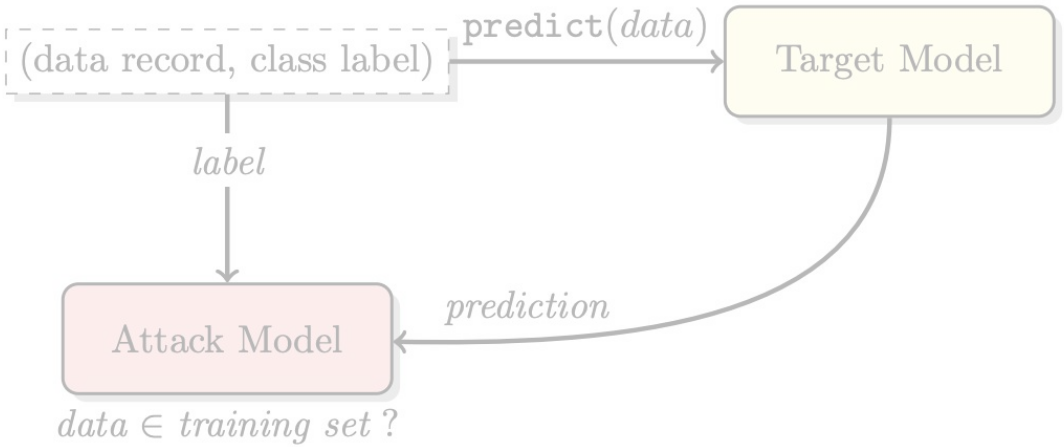

Fig. 1: Membership inference attack in the black-box setting. The attacker queries the target model with a data record and obtains the model's prediction on that record. The prediction is a vector of probabilities, one per class, that the record belongs to a certain class. This prediction vector, along with the label of the target record, is passed to the attack model, which infers whether the record was *in* or *out* of the target model's training dataset.

(z) The term "privacy-enhancing technology" means any software or hardware solution, technical process, technique, or other technological means of mitigating privacy risks arising from data processing, including by enhancing predictability, manageability, disassociability, storage, security, and confidentiality. These technological means may include secure multiparty computation, homomorphic encryption, zero-knowledge proofs, federated learning, secure enclaves, differential privacy, and synthetic-data-generation tools. This is also sometimes referred to as "privacy-preserving technology."

Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence | The White House

Why do we care?

Do we care?

Autonomy
Dignity

Human Right

Let’s take a look at how privacy shows up in the law

I am not a lawyer!!!!!

I am not a lawyer!!!!!

And I’m definitely not simultaneously a US + Canadian + UK + EU + Chinese + Japanese + …. lawyer

I am not a lawyer!!!!!

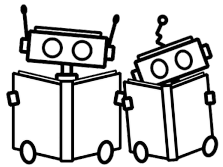

http://genlaw.org/

General Data Protection Regulation

GDPR

Welcome to gdpr-info.eu. Here you can find the official PDF of the Regulation (EU) 2016/679 (General Data Protection Regulation) in the current version of the OJ L 119, 04.05.2016; cor. OJ L 127, 23.5.2018 as a neatly arranged website. All Articles of the GDPR are linked with suitable recitals. The European Data Protection Regulation is applicable as of May 25th, 2018 in all member states to harmonize data privacy laws across Europe. If you find the page useful, feel free to support us by sharing the project.

Art. 4 GDPR

Definitions

For the purposes of this Regulation:

(1) '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;

Art. 4 GDPR – Definitions - General Data Protection Regulation (GDPR)

General Data Protection Regulation

!= privacy

DATA PRIVACY

YOU SHOULD BE PROTECTED FROM ABUSIVE DATA PRACTICES VIA BUILT-IN PROTECTIONS AND YOU SHOULD HAVE AGENCY OVER HOW DATA ABOUT YOU IS USED

 ▸ OSTP ▸ BLUEPRINT FOR AN AI BILL OF RIGHTS

+ Data Privacy | OSTP | The White House

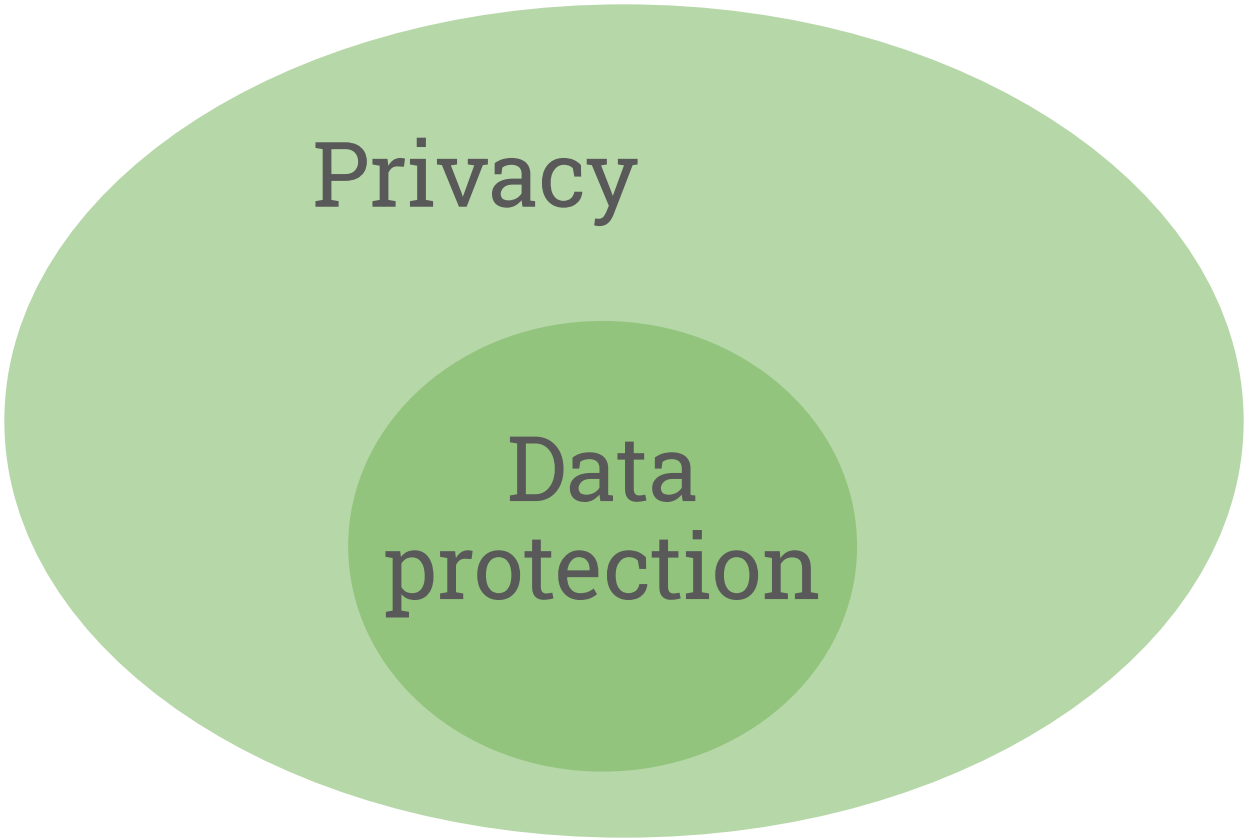
Privacy
Data protection

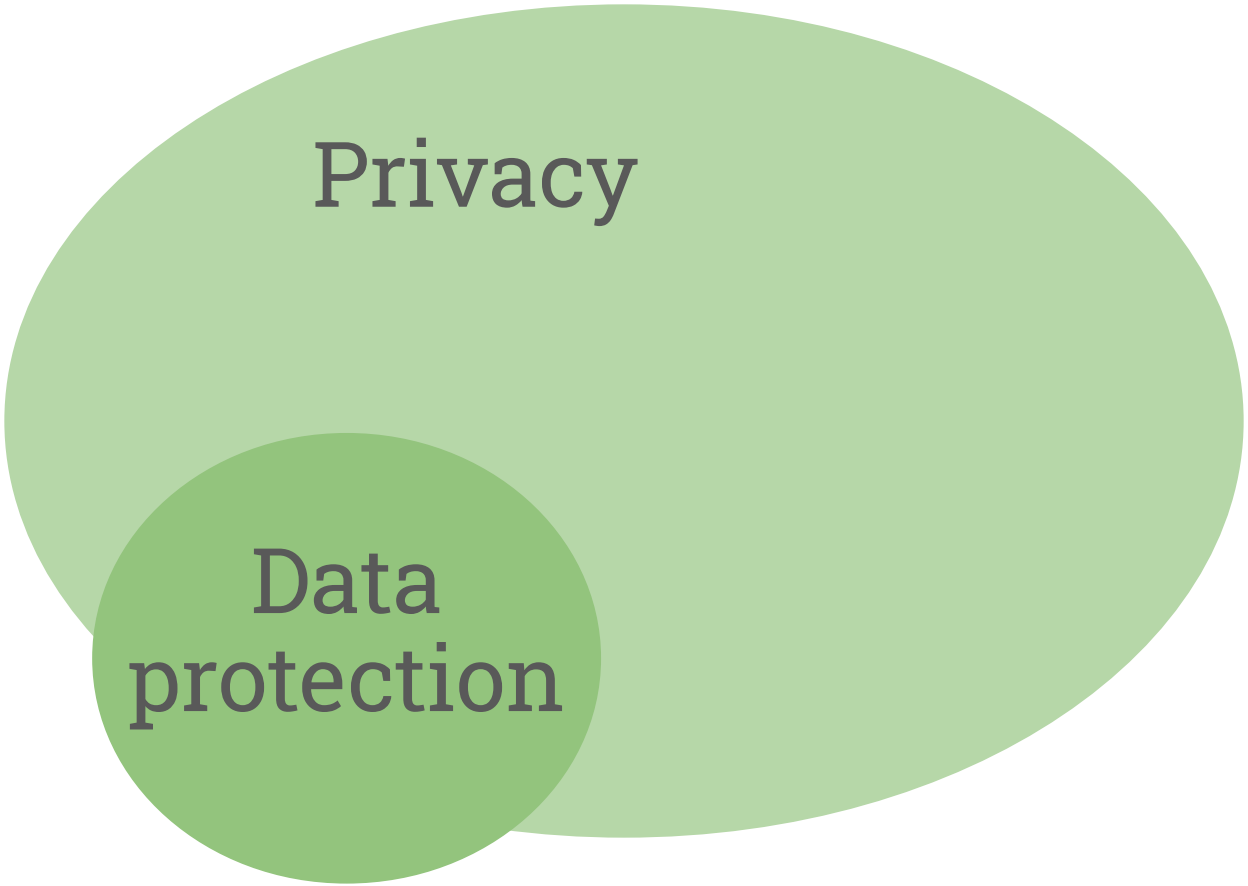
Privacy
Data protection

WHAT IS THE RIGHT TO PRIVACY?

The right to privacy is not mentioned in the Constitution, but the Supreme Court has said that several of the amendments create this right. One of the amendments is the Fourth Amendment, which stops the police and other government agents from searching us or our property without "probable cause" to believe that we have committed a crime. Other amendments protect **our freedom to make certain decisions about our bodies and our private lives without interference from the government - which includes the public schools.**

ACLU, Students: Your Right to Privacy

The Fourth Amendment

The Fourth Amendment reads:

> The right of the people to be secure in their persons, houses, papers, and effects, against unreasonable searches and seizures, shall not be violated, and no warrants shall issue, but upon probable cause, supported by oath or affirmation, and particularly describing the place to be searched, and the persons or things to be seized.

At the outset, two threshold issues are particularly important. First, only searches by the government implicate the Fourth Amendment; it does not apply to "a

What is a search?

What is a “reasonable expectation of privacy”?

The Fifth Amendment

The Fifth Amendment provides:

> *No person shall* be held to answer for a capital, or otherwise infamous crime, unless on a presentment or indictment of a Grand Jury, except in cases arising in the land or naval forces, or in the Militia, when in actual service in time of War or public danger; nor shall any person be subject for the same offence to be twice put in jeopardy of life or limb; nor shall *be compelled in any criminal case to be a witness against himself,* nor be deprived of life, liberty, or property, without due process of law; nor shall private property be taken for public use, without just compensation. (emphasis added)

Grimmelmann, "Internet Law: Cases & Problems"

Laws from different jurisdictions can interfere

The USA *Patriot Act* and Canada's *Privacy Act*

The USA *Patriot Act* was passed following the attacks on September 11, 2001. The Act provides US law enforcement with measures to expand their surveillance capability while minimizing procedural "legal" obstacles. The *Patriot Act* permits law enforcement officials such as the FBI, to seek a court order allowing them to access the personal records of any individual that is under the control of an American company in the United States or an American affiliate operating in Canada for the purpose of an anti-terrorism investigation, without informing individuals or agencies that such disclosure has occurred. In theory, US officials could access information about Canadians through US firms and/or their affiliates, even if the data is located in Canada. There are no provisions in the *Patriot Act* for challenging a US order and refusing to comply with the order may constitute contempt.

Canada

Privacy: a fundamental right in Canada

Comply with US Privacy Laws

The enactment of the California Consumer Privacy Act of 2018 (CCPA) on January 1, 2020 with an enforceability date of July 1, 2020, marked the first comprehensive US state pri
of privacy-related legislation at both the federal and state level followed. Although many of these bills failed to become law, several states have now managed to pass comprehe
Moreover, a federal bill known as the American Data Privacy and Protection Act (ADPPA) is making its way through Congress. The bill is significant as it marks the first federal pr
and bicameral support. If enacted, the ADPPA would preempt the majority of state and local laws, rendering any similar provisions therein invalid.

Utah	Consumer Privacy Act (UCPA)	December 31, 2023

USA State Law Tracker | Portal Advisory | DataGuidance

Whose data is protected by the GDPR vs. U.S. data protection laws? What types of data are protected?

	EU	California		Virginia	Colorado
	GDPR	**CCPA**	**CPRA**	**VCDPA**	**CPA**
Whose data is protected?					
Statutory term	Data subject	Consumer	Consumer	Consumer	Consumer
Defined as	Natural person in the EU	Natural person who is a CA resident	Natural person who is a CA resident	Natural person who is a VA resident	Individual who is a CO resident
What types of data are protected?					
Statutory term	Personal data	Personal information	Personal information	Personal data	Personal data
Defined as	Any information relating to an identified or identifiable natural person	Information that identifies, relates to, describes, is reasonably capable of being associated with, or could reasonably be linked, directly or indirectly, with a particular consumer or household	Information that identifies, relates to, describes, is reasonably capable of being associated with, or could reasonably be linked, directly or indirectly, with a particular consumer or household	Any information that is linked or reasonably linkable to an identified or identifiable natural person	Information that is linked or reasonably linkable to an identified or identifiable individual

Bloomberg Law: Comparing US State Data Privacy Laws vs. the EU's GDPR

There are also domain specific privacy legislation

United States [edit]

Under the U.S. Health Insurance Portability and Accountability Act (HIPAA), PHI that is linked based on the following list of 18 identifiers must be treated with special care:

1. Names
2. All geographical identifiers smaller than a state, except for the initial three digits of a zip code if, according to the current publicly available data from the U.S. Bureau of the Census: the geographic unit formed by combining all zip codes with the same three initial digits contains more than 20,000 people; the initial three digits of a zip code for all such geographic units containing 20,000 or fewer people is changed to 000
3. Dates (other than year) directly related to an individual
4. Phone Numbers
5. Fax numbers
6. Email addresses
7. Social Security numbers
8. Medical record numbers
9. Health insurance beneficiary numbers
10. Account numbers
11. Certificate/license numbers
12. Vehicle identifiers and serial numbers, including license plate numbers;
13. Device identifiers and serial numbers;
14. Web Uniform Resource Locators (URLs)
15. Internet Protocol (IP) address numbers
16. Biometric identifiers, including finger, retinal and voice prints
17. Full face photographic images and any comparable images
18. Any other unique identifying number, characteristic, or code except the unique code assigned by the investigator to code the data

Protected health information - Wikipedia

In the US, we care about privacy harms

In the US, we care about privacy harms

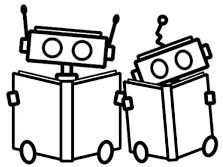

http://genlaw.org/

PRIVACY HARMS

DANIELLE KEATS CITRON* & DANIEL J. SOLOVE**

ABSTRACT

The requirement of harm has significantly impeded the enforcement of privacy law. In most tort and contract cases, plaintiffs must establish that they have suffered harm. Even when legislation does not require it, courts have taken it upon themselves to add a harm element. Harm is also a requirement to establish standing in federal court. In Spokeo, Inc. v. Robins *and* TransUnion LLC v. Ramirez, *the Supreme Court ruled that courts can override congressional judgment about cognizable harm and dismiss privacy claims.*

Case law is an inconsistent, incoherent jumble with no guiding principles. Countless privacy violations are not remedied or addressed on the grounds that there has been no cognizable harm.

What’s harm????

There’s a whole world of literature on harm

Privacy Violations

VOL. 154 JANUARY 2006 NO. 3

ARTICLES

A TAXONOMY OF PRIVACY

DANIEL J. SOLOVE†

Privacy is a concept in disarray. Nobody can articulate what it means. As one commentator has observed, privacy suffers from "an embarrassment of

The Boundaries of Privacy Harm

M. RYAN CALO*

Subjective:

“internal to the person harmed.”
“perception of unwanted observation.”
“acute or ongoing”

“range in severity from mild discomfort at the presence of a security camera to "mental pain and distress[] far greater than could be inflicted by mere bodily injury.""

Objective:

“external to the person harmed.”

“This set of harms involves the forced or unanticipated use of information about a person against that person.”

Cases

What type of harm does DP protect?

What type of harm does DP protect?

Never

Data collection

Data processing

Decisional inference

What type of harm does DP protect?

Never	**Sometimes**
Data collection	Subjective
Data processing	Objective
Decisional inference	Disclosure
	Exposure
	Breach of confidentiality
	Identification
	Surveillance
	Aggregation

What type of harm does DP protect?

Never	Sometimes	Always
Data collection	Subjective	
Data processing	Objective	
Decisional inference	Disclosure	
	Exposure	
	Breach of confidentiality	
	Identification	
	Surveillance	
	Aggregation	

What type of harm does text sanitization protect?

Never	Sometimes	Always
Data collection	Subjective	
Data processing	Objective	
Decisional inference	Disclosure	
	Exposure	
	Breach of confidentiality	
	Identification	
	Surveillance	
	Aggregation	

What type of harm does text sanitization protect?

Never	Sometimes	Always
Data collection	Subjective	
Data processing	Objective	
Decisional inference	Disclosure	
	Exposure	
	Breach of confidentiality	
	Identification	
	Surveillance	
	Aggregation	

But also, it’s a little too expansive to solve *all privacy harms*

What goes into a general privacy evaluation?

- Text extraction
- Personal information detection
- Sensitive text identification
- Contextualized tests

We can’t get there in one step

We can’t get there in one step

But we have to ***iterate***

We can’t get there in one step

But we have to ***iterate***

And get ***clear*** about what problem we’re solving

Just a few more thoughts

Privacy ==
Security ==
Safety??

Privacy vs. Security vs. Safety??

It’s all very
blended right now

Who is your dang audience?

Thank you


Katherine Lee


Nicholas Carlini


Daphne Ippolito


Milad Nasr

Get clear on what problem are you solving
    And check with the relevant parties

“Harm” is a legal concept and important for privacy

Data protection != Privacy

We have to start somewhere
    But be *clear about the limitations*


Matthew Jagielski


Chris Choquette


Chiyuan Zhang


Florian Tramèr


James Grimmelmann


A. Feder Cooper


Jon Hayase


Erics Wallace