

SoK: Differential Privacies

A taxonomy of differential privacy variants and extensions

Damien Desfontaines¹ and Balázs Pejó²

¹Tumult Labs
damien@desfontain.es
²CrySyS Lab
pejo@crysys.hu

Abstract

Shortly after it was first introduced in 2006, *differential privacy* became the flagship data privacy definition. Since then, numerous variants and extensions were proposed to adapt it to different scenarios and attacker models. In this work, we propose a systematic taxonomy of these variants and extensions. We list all data privacy definitions based on differential privacy, and partition them into seven categories, depending on which aspect of the original definition is modified.

These categories act like dimensions: variants from the same category cannot be combined, but variants from different categories can be combined to form new definitions. We also establish a partial ordering of relative strength between these notions by summarizing existing results. Furthermore, we list which of these definitions satisfy some desirable properties, like composition, post-processing, and convexity by either providing a novel proof or collecting existing ones.

Contents

1	Introduction	3
2	Differential Privacy	5
2.1	The original version	6
2.2	Dimensions	6
2.3	Properties	8
2.4	Relations between definitions	9
3	Quantification of privacy loss (Q)	10
3.1	Allowing a small probability of error	10
3.2	Averaging the privacy loss	11
3.3	Controlling the tail distribution of the privacy loss	12
3.4	Extension	14
3.5	Multidimensional definitions	14

4	Neighborhood definition (N)	14
4.1	Changing the sensitive property	15
4.2	Limiting the scope of the definition	17
4.3	Applying the definition to other types of input	17
4.4	Extensions	19
4.5	Multidimensional definitions	20
5	Variation of privacy loss (V)	20
5.1	Varying the privacy level across inputs	20
5.2	Randomizing the variation of privacy levels	21
5.3	Multidimensional definitions	22
6	Background knowledge (B)	24
6.1	Conditioning the output on the sensitive property	24
6.2	Removing the effect of correlations in the data	25
6.3	Multidimensional definitions	26
7	Change in formalism (F)	28
7.1	Hypothesis testing	28
7.2	Changing the shape of the prior-posterior bounds	29
7.3	Comparing two posteriors	30
7.4	Multidimensional definitions	31
8	Relativization of the knowledge gain (R)	32
8.1	Taking into account auxiliary leakage function	32
8.2	Borrowing concepts from zero-knowledge proofs	33
8.3	Multidimensional definitions	33
9	Computational power (C)	34
9.1	Multidimensional definitions	35
10	Summarizing table	36
10.1	Proofs of properties	40
10.1.1	Axioms	40
10.1.2	Composition	41
11	Scope and related work	42
11.1	Methodology	42
11.2	Out of scope definitions	43
11.3	Local model and other contexts	45
11.4	Related work	46
12	Conclusion	47

1 Introduction

What does it mean for data to be anonymized? Samarati and Sweeney discovered that removing explicit identifiers from dataset records was not enough to prevent information from being re-identified [Sam01, Swe02], and they proposed the first definition of anonymization. This notion, called *k-anonymity*, is a property of a dataset: each combination of re-identifying fields must be present at least k times. In the following decade, further research showed that sensitive information about individuals could still be leaked when releasing k -anonymous datasets, and many variants and definitions were proposed, such as *l-diversity* [MGKV06], *t-closeness* [LLV07], and *n-confusion* [ST12].

A common shortcoming of these approaches is that they defined anonymity as a property of the *dataset*: without knowing how the dataset is generated, arbitrary information can be leaked. This approach was changed with the introduction of *differential privacy* [DM05, Dwo06] (DP): rather than being a property of the sanitized dataset, anonymity was instead defined as a property of the *process*. It was inspired by Dalenius’ privacy goal that “Anything about an individual that can be learned from the dataset can also be learned without access to the dataset” [Dal77], a goal similar to one already used in probabilistic encryption [SM84].

Thanks to its useful properties, DP quickly became the flagship of data privacy definitions. Many algorithms and statistical processes were changed to satisfy DP and were adopted by organizations like Apple [Tea17], Facebook [MDH⁺20] [HDS⁺20], Google [ABC⁺20] [BDD⁺20] [BBD⁺21] [WZL⁺20], LinkedIn [RCM⁺20] [RSP⁺20], Microsoft [DKY17], OhmConnect [MP20], and the US Census Bureau [AAG⁺10] [GAP18] [FMM19].

Since the original introduction of differential privacy, many variants and extensions have been proposed to adapt it to different contexts or assumptions. These new definitions enable practitioners to get privacy guarantees, even in cases that the original DP definition does not cover well. This happens in a variety of scenarios: the noise mandated by DP can be too large and force the data custodian to consider a weaker alternative, the risk model might be inappropriate for certain use cases, or the context might require the data owner to make stronger statements on what information the privacy mechanism can reveal.

Figure 1 shows the prevalence of this phenomenon: approximately 225 different notions¹, inspired by DP, were defined in the last years. As we show in Figure 1, this phenomenon does not seem to slow down over time. These definitions can be *extensions* or *variants* of DP. An extension encompasses the original DP notion as a special case, while a variant changes some aspect, typically to weaken or strengthen the original definition.

With so many definitions, it is difficult for new practitioners to get an overview of this research area. Many definitions have similar goals, so it is also challenging to understand which are appropriate to use in which context. These difficulties also affect experts: a number of definitions listed in this work have been defined independently multiple times (often with identical meaning but different names, or identical names but different meanings). Finally, variants are often introduced without a comparison to related notions.

This systematization of knowledge attempts to solve these problems. It is a taxonomy of variants and extensions of DP, providing short explanations of the intuition, use cases and basic properties of each. By categorizing these definitions, we attempt to simplify the understanding of existing variants and extensions, and of the relations between them. We hope to make it easier for new practitioners to understand whether their use case needs an alternative definition, and if so, which

¹We count all the definitions which are presented as “new” in the papers introducing them.




Figure 1: Accumulated number of papers which are introducing new DP notions (line) and the exact number of these definitions (bar) for the years from 2008 till 2021.

existing notions are the most appropriate, and what their basic properties are.

Contributions and organization

We systematize the scientific literature on variants and extensions of differential privacy, and propose a unified and comprehensive taxonomy of these definitions. We define seven *dimensions*: these are ways in which the original definition of DP can be modified or extended. We list variants and extensions that belong to each dimension, and we highlight representative definitions for each. Whenever possible, we compare these definitions and establish a partial ordering between the strengths of different notions. Furthermore, for each definition, we specify whether it satisfies Kifer et al.’s *privacy axioms* [KL10, KL12], (post-processing and convexity), and whether they are composable.

Our survey is organized as follows:

- In Section 2, we recall the original definition of DP and introduce our dimensions along which DP can be modified. Moreover, we present the basic properties of DP, and define how definitions can related to each other.
- In the following 7 sections (Sections 3 to 9), we introduce our dimensions, and list and compare the corresponding definitions.
- In Section 10, we summarize the results from the previous sections into a table, showing the corresponding properties with proofs, and list the known relations.
- In Section 11 we detail the methodology and scope of this work and review the related literature.
- Finally, in 12 we conclude this work.

2 Differential Privacy

In this chapter we recap the original DP definition with its basic properties, define how definitions can related to each other and introduce our dimensions along which DP can be modified.

Let \mathcal{T} denote an arbitrary set of possible *records*. We typically use t to denote the records themselves. A *dataset* is a finite indexed family of records. We denote by \mathcal{D} space of possible datasets, individuals datasets are typically called D , D' , D_1 or D_2 . The indices of a dataset are typically called i and j , with $D(i)$ referring to the i -th record of a dataset D . We denote by D_{-i} the dataset D whose i -th record has been removed.

Let \mathcal{O} denote an arbitrary set of possible *outputs*; outputs are typically called O , and sets of outputs called S . A *mechanism* is a randomized function which takes a dataset as input and returns an output. Mechanisms are typically called \mathcal{M} while $\mathcal{M}(D)$ is usually a random variable.

Probability distributions on \mathcal{T} are called π , probability distribution on \mathcal{D} are called θ , and family of probability distributions on \mathcal{D} are called Θ . Given some property ϕ , let $\mathcal{M}(D)_{|D \sim \theta, \phi}$ denote the random variable corresponding to the output of $\mathcal{M}(D)$, when D is drawn from a distribution θ conditioned on ϕ .

Table 1 summarizes the notations used throughout the paper.

Notation	Description
\mathcal{T}	Set of possible records
$t \in \mathcal{T}$	A possible record
$\mathcal{D} = \mathcal{T}^*$	Set of possible datasets (sequences of records)
$D \in \mathcal{D}$	Dataset (we also use D' , D_1 , D_2 , \dots)
$D(i)$	i -th record of the dataset ($i \leq D $)
D_{-i}	Dataset D , with its i -th record removed
\mathcal{O}	Set of possible outputs of privacy mechanisms
$S \subseteq \mathcal{O}$	Subset of possible outputs
$O \in \mathcal{O}$	Output of the privacy mechanism
$\mathcal{M} : \mathcal{D} \rightarrow \mathcal{O}$	Privacy mechanism (probabilistic)
$\mathcal{M}(D)$	The distribution (or an instance thereof) of the outputs of \mathcal{M} given input D
$d_{\mathcal{D}} : \mathcal{D} \times \mathcal{D} \rightarrow \mathbb{R}_0^+$	Distance function between datasets
$\phi \subseteq \mathcal{D}$	Predicate on datasets
Φ	Family of sensitive predicates on datasets
π	Probability distribution on \mathcal{T}
\mathcal{B}	Set of possible background knowledges
$B \in \mathcal{B}$	Background knowledge (we also use \hat{B})
Θ	Family of probability distributions on \mathcal{D} , on $\mathcal{D} \times \mathcal{B}$
$\theta \in \Theta$	Probability distribution on \mathcal{D} , on $\mathcal{D} \times \mathcal{B}$
$\mathcal{M}(D)_{ D \sim \theta, \phi}$	Distribution of outputs of \mathcal{M} given an input drawn from θ , conditioned on ϕ
Ω	Probabilistic polynomial-time Turing machine, called <i>distinguisher</i>

Table 1: Notations used in this paper.

2.1 The original version

The first DP mechanism, randomized response, was proposed in 1965 [War65], and data privacy definitions that are a property of a mechanism and not of the output dataset were already proposed in as early as 2003 [EGS03]. However, DP and the related notion of ε -indistinguishability were first formally defined in 2006 [Dwo06, DMNS06, DM05].

Definition 1 (ε -indistinguishability [DMNS06]). *Two random variables A and B are ε -indistinguishable, denoted $A \approx_\varepsilon B$, if for all measurable sets X of possible events:*

$$\mathbb{P}[A \in X] \leq e^\varepsilon \cdot \mathbb{P}[B \in X] \text{ and } \mathbb{P}[B \in X] \leq e^\varepsilon \cdot \mathbb{P}[A \in X].$$

Informally, A and B are ε -indistinguishable if their distributions are “close”. This notion originates from the cryptographic notion of indistinguishability [GM84]. A similar notion, $(1, \varepsilon)$ -privacy, was defined in [CM06], where $(1 + \varepsilon)$ used in place of e^ε , and it was also called *log-ratio distance* in [HMSS19].

The notion of ε -indistinguishability is then used to define differential privacy.

Definition 2 (ε -differential privacy [Dwo06]). *A privacy mechanism \mathcal{M} is ε -differential private (or ε -DP) if for all datasets D_1 and D_2 that differ only in one record, $\mathcal{M}(D_1) \approx_\varepsilon \mathcal{M}(D_2)$.*

Mechanisms

Besides the random response mechanism (which returns the true value with probability p and returns a random value otherwise) in general there are three places where noise can be injected to guarantee DP [MTV⁺20]: it can be added to the input, to the output, and directly to the mechanism. For instance, in machine learning context input (e.g., [PAE⁺16]) and output (e.g., [PTB19]) perturbation are equivalent with sanitizing the dataset before and the predictions after the training respectively. Concerning mechanism perturbation, there are various techniques, such as loss function perturbation (e.g., [CMS11]), and gradient perturbation (e.g., [ACG⁺16]), which insert noise to the model objective and update respectively.

The most widely used distributions the noise is sampled from are Laplace, Gauss, and Exponential. Using the first and last makes any underlying mechanism to satisfy ε -DP for continuous and discrete cases, however, they could result in large added noise. On the contrary, the middle distribution decrease the probability of such events significantly, but the obtained differential privacy guarantee is weaker (i.e. (ε, δ) -DP, introduced in Section 3).

2.2 Dimensions

Variants and extensions of differential privacy modify the original definition in various ways. To establish a comprehensive taxonomy, a natural approach is to partition them into *categories*, depending on which aspect of the definition they change. Unfortunately, this approach fails for privacy definitions, many of which modify several aspects at once, so it is impossible to have a categorization such that every definition falls neatly into only one category.

The approach we take is to define *dimensions* along which the original definition can be modified. Each variant or extension of DP can be seen as a point in a multidimensional space, where each coordinate corresponds to one possible way of changing the definition along a particular dimension. To make this representation possible, our dimensions need to satisfy two properties:

- **Mutual compatibility:** definitions that vary along different dimensions can be combined to form a new, meaningful definition.
- **Inner exclusivity:** definitions in the same dimension cannot be combined to form a new, meaningful definition (but they can be pairwise comparable).

In addition, each dimension should be *motivatable*: there should be an intuitive explanation of what it means to modify DP along each dimension. Moreover, each possible choice within a dimension should be similarly understandable, to allow new practitioners to determine quickly which kind of definition they should use or study, depending on their use case.

We introduce our dimensions by reformulating the guarantee offered by DP, highlighting aspects that have been modified by its variants or extensions. Each dimension is attributed a letter, and we note the dimension letter corresponding to each highlight. This formulation considers the point of view of an attacker, trying to find out some sensitive information about some input data using the output of a mechanism.

An attacker with **perfect background knowledge (B)** and **unbounded computation power (C)** is **unable (R)** to **distinguish (F)** anything about an individual (N), uniformly across users (V) even in the **worst-case scenario (Q)**.

This informal definition of DP with the seven highlighted aspects give us seven distinct dimensions. We denote each one by a letter and summarize them in Table 2. Each is introduced in its corresponding section.

Dimension	Description	Typical motivations
Q uantification of Privacy Loss	How is the privacy loss quantified across outputs?	Averaging risk, obtaining better composition properties
N eighborhood Definition	Which properties are protected from the attacker?	Protecting specific values or multiple individuals
V ariation of Privacy Loss	Can the privacy loss vary across inputs?	Modeling users with different privacy requirements
B ackground Knowledge	How much prior knowledge does the attacker have?	Using mechanisms that add less or no noise to data
F ormalism change	Which formalism is used to describe the attacker’s knowledge gain?	Exploring other intuitive notions of privacy
R elativization of Knowledge Gain	What is the knowledge gain relative to?	Guaranteeing privacy for correlated data
C omputational Power	How much computational power can the attacker use?	Combining cryptography techniques with DP

Table 2: The seven dimensions and their typical motivation.

Note that the interpretation of DP is subject to some debate. In [TSD20], authors summarize this debate, and show that DP can be interpreted under two possible lenses: it can be seen as an

associative property, or as a *causal* property. The difference between the two interpretations is particularly clear when one supposes that the input dataset is modeled as being generated by a probability distribution.

- In the associative view, this probability distribution is *conditioned* upon the value of one record. If the distribution has correlations, this change can affect other records as well.
- In the causal view, the dataset is first generated, and the value of one record is then *changed* before computing the result of the mechanism.

While the causal view does not require any additional assumption to capture the intuition behind DP, the associative view requires that either all records are independent in the original probability distribution (the *independence assumption*), or the adversary must know all data points except one (the *strong adversary assumption*, which we picked in the reformulation above).

These considerations can have a significant impact on DP variants and extensions, either leading to distinct variants that attempt to capture the same intuition, or to the same variant being interpreted in different ways.

2.3 Properties

In this section, we introduce three main properties of differential privacy, that we then check against variants and extensions of DP listed in this work.

Privacy Axioms

Two important properties of data privacy notions are called *privacy axioms*, proposed in [KL10, KL12]. These are not axioms in a sense that they assumed to be true; rather, they are consistency checks: properties that, if not satisfied by a data privacy definition, indicate a flaw in the definition².

Definition 3 (Privacy axioms [KL10, KL12]).

1. **Post-processing**³ (or transformation invariance): A privacy definition Def satisfies the post-processing axiom if, for any mechanism \mathcal{M} satisfying Def and any probabilistic function f , the mechanism $D \rightarrow f(\mathcal{M}(D))$ also satisfies Def.
2. **Convexity** (or privacy axiom of choice): A privacy definition Def satisfies the convexity axiom if, for any two mechanisms \mathcal{M}_1 and \mathcal{M}_2 satisfying Def, the mechanism \mathcal{M} defined by $\mathcal{M}(D) = \mathcal{M}_1(D)$ with probability p and $\mathcal{M}(D) = \mathcal{M}_2(D)$ with probability $1 - p$ also satisfies Def.

Most differential privacy variants and extensions, including the original definition of DP, satisfy these axioms, although some do not. We highlighted these in Table 3 in Section 10.

²The necessity of these were questioned in [HMQN13], where the authors showed a natural notions of anonymity that contradict them.

³This definition must be slightly adapted for some variants, see for example Proposition 6 in Section 10.

Composition

A third important property is one of differential privacy’s main strengths: *composability*. It guarantees that the output of two mechanisms satisfying a privacy definition still satisfies the definition, typically with a change in parameters. There are several types of composition: *parallel composition*, *sequential composition*, and *adaptive composition*. We introduce the first two below.

Theorem (Parallel composition [Dwo06]). *Let \mathcal{M}_1 be a ε_1 -differentially private mechanism, and \mathcal{M}_2 a ε_2 -differentially private mechanism. For any dataset D , let D_1 and D_2 be the result of an operation that separates records in two disjoint datasets. Then the mechanism \mathcal{M} defined by $\mathcal{M}(D) = (\mathcal{M}_1(D_1), \mathcal{M}_2(D_2))$ is $\max(\varepsilon_1, \varepsilon_2)$ -differentially private.*

This property allows us to build *locally differentially private* mechanisms, in which a central server can compute global statistics without accessing the raw data from each user. In this work, we focus on sequential composition, which we simply call *composition*.

Theorem (Sequential composition [Dwo06]). *Let \mathcal{M}_1 be a ε_1 -differentially private mechanism, and \mathcal{M}_2 a ε_2 -differentially private mechanism. Then the mechanism \mathcal{M} defined by $\mathcal{M}(D) = (\mathcal{M}_1(D), \mathcal{M}_2(D))$ is $(\varepsilon_1 + \varepsilon_2)$ -differentially private.*

This theorem stays true if \mathcal{M}_2 depends on the value of $\mathcal{M}_1(D)$: this variant is called *adaptive composition*. This latter property allows to quantify the gain of information over time of an attacker interacting with a differentially private query engine.

In this work, we only consider sequential composition, in the more abstract form formalized below.

Definition 4 (Composability). *A privacy definition Def with parameter α is composable if for any two mechanisms \mathcal{M}_1 and \mathcal{M}_2 satisfying respectively α_1 -Def and α_2 -Def, the mechanism $\mathcal{M}(D) = (\mathcal{M}_1(D), \mathcal{M}_2(D))$ satisfies α -Def for some (non-trivial) α .*

2.4 Relations between definitions

When learning about a new data privacy notion, it is often useful to know what are the known relations between this notion and other definitions. However, definitions have parameters that often have different meanings, and whose value is not directly comparable. To capture extensions, when a definition can be seen as a special case of another, we introduce the following definition.

Definition 5 (Extensions). *Let α -Def₁ and β -Def₂ be data privacy definitions. We say that Def₁ is extended by Def₂, and denote it as Def₁ \subset Def₂, if for all α , there is a value of β such that α -Def₁ is identical to β -Def₂.*

Concerning variants, to claim that a definition is stronger than another, we adopt the concept of ordering established in [CY16] using α and β as tuples, encoding multiple parameters. Note that we slightly changed the original definition as that only required the second condition to hold, which would classify any extension as a stronger variant.

Definition 6 (Relative strength of privacy definitions). *Let α -Def₁ and β -Def₂ be data privacy definitions. We say that Def₁ is stronger than Def₂, and denote it Def₁ \succ Def₂, if:*

1. for all α , there is a β such that α -Def₁ \implies β -Def₂;

2. for all β , there is an α such that $\alpha\text{-Def}_1 \implies \beta\text{-Def}_2$.

If Def_1 is both stronger than and weaker than Def_2 , we say that the two definitions are equivalent, and denote it $\text{Def}_1 \sim \text{Def}_2$.

Relative strength implies a partial ordering on the space of possible definitions. On the other hand, if two definitions are equivalent, this does not mean that they are equal: they could be only equal up to a change in parameters. Both relations are reflexive and transitive; and we define the symmetric counterpart of these relations as well (i.e., \prec and \supset). Moreover, for brevity, we combine these two concepts in a single notation: if $\text{Def}_1 \subset \text{Def}_2$ and $\text{Def}_1 \succ \text{Def}_2$, we say that Def_2 is a weaker extension of Def_1 , and denote it $\text{Def}_1 \subset^> \text{Def}_2$.

A summarizing table is presented at the end of this work, where for each definition, we also highlight its dimensions and its relation to other notions. In Table 3, we also specify whether these notions satisfy the privacy axioms and the composability property (\checkmark : yes, \times : no, $?$: currently unknown); in Section 10 we either provide a reference or a novel proof for each of these claims.

3 Quantification of privacy loss (Q)

The risk model associated to differential privacy is a *worst-case* property: it quantifies not only over all possible neighboring datasets but also over all possible outputs. However, in many real-life risk assessments, events with vanishingly small probability are ignored, or their risk weighted according to their probability. It is natural to consider analogous relaxations, especially since these relaxations often have better composition properties, and enable natural mechanisms like the Gaussian mechanism to be considered private [DR⁺14].

Most of the definitions within this section can be expressed using the *privacy loss random variable*, first defined in [DN03] as the *adversary's confidence gain*, so we first introduce this concept. Roughly speaking, it measures how much information is revealed by the output of a mechanism.

Definition 7 (Privacy loss random variable [DN03]). *Let \mathcal{M} be a mechanism, and D_1 and D_2 two datasets. The privacy loss random variable between $\mathcal{M}(D_1)$ and $\mathcal{M}(D_2)$ is defined as:*

$$\mathcal{L}_{\mathcal{M}(D_1)/\mathcal{M}(D_2)}(O) = \ln \left(\frac{\mathbb{P}[\mathcal{M}(D_1) = O]}{\mathbb{P}[\mathcal{M}(D_2) = O]} \right).$$

if neither $\mathbb{P}[\mathcal{M}(D_1) = O]$ nor $\mathbb{P}[\mathcal{M}(D_2) = O]$ is 0; in case only $\mathbb{P}[\mathcal{M}(D_2) = O]$ is zero then $\mathcal{L}_{\mathcal{M}(D_1)/\mathcal{M}(D_2)}(O) = \infty$, otherwise $\mathcal{L}_{\mathcal{M}(D_1)/\mathcal{M}(D_2)}(O) = -\infty$. When the mechanism is clear from context, we simply write \mathcal{L}_{D_1/D_2} .

Differential privacy bounds the *maximum value* of \mathcal{L}_{D_1/D_2} . Instead of considering the maximum value, which corresponds to the worst possible output, relaxations of this section will allow a small probability of error, consider the average of the privacy loss random variable, or describe its behavior in finer ways.

3.1 Allowing a small probability of error

The first option, whose introduction is commonly attributed to [DKM⁺06], relaxes the definition of ϵ -indistinguishability by allowing an additional small density of probability on which the upper ϵ bound does not hold. This small density, denoted δ , can be used to compensate for outputs

for which the privacy loss is larger than e^ϵ . This led to the definition of *approximate differential privacy*, often simply called (ϵ, δ) -DP. This is, by far, the most commonly used relaxation in the scientific literature. Several of the enlisted modifications from the other dimensions were introduced with this additive factor δ , but for clarity we omit these details where not crucial.

Definition 8 ((ϵ, δ) -differential privacy [DKM⁺06]). *A privacy mechanism \mathcal{M} is (ϵ, δ) -DP if for any datasets D_1 and D_2 that differ only on one record, and for all $S \subseteq \mathcal{O}$:*

$$\mathbb{P}[\mathcal{M}(D_1) \in S] \leq e^\epsilon \cdot \mathbb{P}[\mathcal{M}(D_2) \in S] + \delta.$$

This definition is equivalent with Max-KL stability [BNS⁺16], a special case of algorithmic stability, which requires that one change in an algorithm’s inputs does not change its output “too much”.

The δ in (ϵ, δ) -DP is sometimes explained as the probability that the privacy loss of the output is larger than e^ϵ (or, equivalently, that the ϵ -indistinguishability formula is satisfied). In fact, this intuition corresponds to a different definition, first introduced in [MKA⁺08] as *probabilistic DP* (ProDP), also called (ϵ, δ) -DP in distribution in [CO15]. A detailed explanation of the distinction between the two definitions can be found in [Mei18].

Definition 9 ((ϵ, δ) -probabilistic differential privacy [Mei18]). *A privacy mechanism \mathcal{M} is (ϵ, δ) -probabilistically DP (ProDP) if for any datasets D_1 and D_2 that differ only on one record there is a set $S_1 \subseteq \mathcal{O}$ where $\mathbb{P}[\mathcal{M}(D_1) \in S_1] \leq \delta$, such that for all measurable sets $S \subseteq \mathcal{O}$:*

$$\mathbb{P}[\mathcal{M}(D_1) \in S \setminus S_1] \leq e^\epsilon \cdot \mathbb{P}[\mathcal{M}(D_2) \in S \setminus S_1].$$

It is straightforward to show that (ϵ, δ) -ProDP is stronger than (ϵ, δ) -DP (with no change in parameters); a proof of the reverse result (with parameter change) is given in [ZWB⁺19]. Both definitions can be reformulated using the privacy loss random variable.

Theorem. *A mechanism \mathcal{M} is:*

- ϵ -DP $\Leftrightarrow \mathbb{P}_{O \sim \mathcal{M}(D_1)}[\mathcal{L}_{D_1/D_2}(O) > \epsilon] = 0$ for all neighboring D_1 and D_2 .
- (ϵ, δ) -DP $\Leftrightarrow \mathbb{E}_{O \sim \mathcal{M}(D_1)}\left[\max\left(0, 1 - e^{\epsilon - \mathcal{L}_{D_1/D_2}(O)}\right)\right] \leq \delta$ for all neighboring D_1 and D_2 .
- (ϵ, δ) -ProDP $\Leftrightarrow \mathbb{P}_{O \sim \mathcal{M}(D_1)}[\mathcal{L}_{D_1/D_2}(O) > \epsilon] \leq \delta$ for all neighboring D_1 and D_2 .

Approximate and probabilistic differential privacy can be combined to form $(\epsilon, \delta_a, \delta_p)$ -relaxed DP (RelDP) [ZQZ⁺15], which requires (ϵ, δ_a) -DP with probability at least $1 - \delta_p$.

3.2 Averaging the privacy loss

As ϵ -DP corresponds to a *worst-case* risk model, it is natural to consider relaxations to allow for larger privacy loss for some outputs. It is also natural to consider *average-case* risk models: allowing larger privacy loss values only if lower values compensate it in other cases. One such relaxation is called *Kullback-Leibler privacy* [BD14, CY16]: it considers the *arithmetic* mean of the privacy loss random variable, which measures how much information is revealed when the output of a private algorithm is observed.

Definition 10 (ε -Kullback-Leibler privacy [BD14, CY16]). *A privacy mechanism \mathcal{M} is ε -Kullback-Leibler private (KLPr) if for all D_1, D_2 differing in one record:*

$$\mathbb{E}_{O \sim \mathcal{M}(D_1)} [\mathcal{L}_{D_1/D_2}(O)] \leq \varepsilon. \quad (1)$$

Note that this formula can be expressed as $D_{KL}(\mathcal{M}(D_1) | \mathcal{M}(D_2)) \leq \varepsilon$ where D_{KL} is the Kullback-Leibler-divergence.

ε -KL privacy considers the *arithmetic* mean of the privacy loss random variable or, equivalently, the *geometric* mean of $e^{\mathcal{L}_{D_1/D_2}}$. This choice of averaging function does not attribute a lot of weight to worst-case events, where \mathcal{L}_{D_1/D_2} takes high values. *Rényi DP* extends this idea by adding a parameter $\alpha \geq 1$, which allows controlling the choice of averaging function by bounding the α th momentum of the privacy loss random variable.

Definition 11 ((α, ε) -Rényi differential privacy [Mir17]). *Given $\alpha > 1$, a privacy mechanism \mathcal{M} is (α, ε) -Rényi DP (RényiDP) if for all pairs of neighboring datasets D_1 and D_2 :*

$$\mathbb{E}_{O \sim \mathcal{M}(D_1)} \left[e^{(\alpha-1)\mathcal{L}_{D_1/D_2}(O)} \right] \leq e^{(\alpha-1)\varepsilon}.$$

Note that this formula can be expressed as $D_\alpha(\mathcal{M}(D_1) | \mathcal{M}(D_2)) \leq \varepsilon$ where D_α is the Rényi-divergence of order α .

This definition can be naturally extended by continuity to $\alpha = 1$ (where it is equivalent to ε -KL privacy) and $\alpha = \infty$ (where it is equivalent to ε -DP). Larger values of α lead to more weight being assigned to worst-case events: (α, ε) -Rényi DP \succ (α', ε) -Rényi DP iff $\alpha > \alpha'$. Besides $\alpha = 1$ and $\alpha = \infty$, Rényi DP has a simple interpretation for some values of α : $\alpha = 2$ imposes a bound on the arithmetic mean of $e^{\mathcal{L}_{D_1/D_2}}$, $\alpha = 3$ imposes it on the quadratic mean, $\alpha = 4$ on the cubic mean, etc. A related technique is the moments accountant [ACG⁺16] which keeps track of a bound on the moments of the privacy loss random variable during composition.

It is possible to use other divergence functions to obtain other relaxations. For example, in [WBK19], the authors introduce two technical definitions, *binary- $|\chi|^\alpha$ DP* (b- $|\chi|^\alpha$ DP) and *tenary- $|\chi|^\alpha$ DP* (t- $|\chi|^\alpha$ DP), as part of a proof on amplification by sampling. Other examples of divergences can lead to other variants, like ε -total variation privacy [BD14] (ε -TVPr, using the total variance) and *quantum DP* [Col16] (QDP, using the quantum divergence).

3.3 Controlling the tail distribution of the privacy loss

Some definitions go further than simply considering a worst-case bound on the privacy loss, or averaging it across the distribution. They try to obtain the benefits of (ε, δ) -DP with a smaller ε which holds in most cases, but control the behavior of the bad cases better than (ε, δ) -DP, which allows for catastrophic privacy loss in rare cases.

The first attempt to formalize this idea was proposed in [DR16], where the authors introduce *concentrated DP* (later renamed to *mean-concentrated DP* (mCoDP) in [BS16]). In this definition, a parameter controls the privacy loss variable globally, and another parameter allows for some outputs to have a greater privacy loss; while still requiring that the difference is smaller than a Gaussian distribution. In [BS16], the authors show that this definition does not satisfy the post-processing axiom, and propose another formalization of the same idea called *zero-concentrated DP* (zCoDP) [BS16], which requires that the privacy loss random variable is concentrated around zero.

Definition 12 ((ξ, ρ) -zero-concentrated differential privacy [BS16]). A mechanism \mathcal{M} is (ξ, ρ) -zero-concentrated DP if for all pairs of neighboring datasets D_1 and D_2 and all $\alpha > 1$:

$$\mathbb{E}_{O \sim \mathcal{M}(D_1)} \left[e^{(\alpha-1)\mathcal{L}_{D_1/D_2}(O)} \right] \leq e^{(\alpha-1)(\xi+\rho\alpha)}.$$

Four more variants of concentrated DP exist:

- (ξ, ρ, δ) -approximate zero-concentrated DP [BS16] (AzCoDP), which relaxes (ξ, ρ) -zCoDP by only taking the Rényi divergence on events with probability higher than $1 - \delta$ instead of on the full distribution.
- (ξ, ρ, ω) -bounded CoDP [BS16] (bCoDP) relaxes (ξ, ρ) -zCoDP by requiring the inequality to hold only for $\alpha \leq \omega$.
- (ρ, ω) -truncated CoDP [BDRS18] (tCoDP^[BDRS18]) relaxes $(0, \rho)$ -zCoDP in the same way.
- (ξ, τ) -truncated CoDP [Col16] (tCoDP^[Col16]) requires the Rényi divergence to be smaller than $\min(\xi, \alpha\tau)$ for all $\alpha \geq 1$.

The relations between these definitions and other notions in this section is well-understood. Besides the special cases (e.g., (ρ, ∞) -tCoDP^[BDRS18] is the same as $(0, \rho)$ -zCoDP) and the relations that are a direct consequence of the definitions (e.g., (ξ, ρ) -zCoDP is the same as the condition “ $(\xi + \rho\alpha)$ -RénDP for all $\alpha > 0$ ”), we list known relations below.

Theorem. For all $\varepsilon > 0$, $\delta > 0$, $\mu > 0$, $\tau > 0$, $\xi \geq 0$ and $\omega > 1$:

- ε -DP $\implies \left(\frac{\varepsilon(e^\varepsilon - 1)}{2}, \varepsilon \right)$ -mCoDP (Theorem 3.5 in [DR16])
- ε -DP $\implies \left(0, \frac{\varepsilon^2}{2} \right)$ -zCoDP (Lemma 8.3 in [BS16])
- ε -DP $\iff (\varepsilon, 0)$ -zCoDP (Lemma 3.2 in [BS16])
- (μ, τ) -mCoDP $\implies \left(\mu - \frac{\tau^2}{2}, \frac{\tau^2}{2} \right)$ -zCoDP (Lemma 4.2 in [BS16])
- (ξ, ρ) -zCoDP $\implies (\xi + \rho, O(\sqrt{\xi + 2\rho}))$ -mCoDP (Lemma 4.3 in [BS16])
- (ξ, ρ) -zCoDP $\implies \left(\xi + \rho + \sqrt{4\rho \log \left(\frac{\min(1, \sqrt{\pi\rho})}{\delta} \right)}, \delta \right)$ -DP (Lemma 3.5 and 3.6 in [BS16])
- $\left(\xi + \sqrt{\rho \log \frac{1}{\delta}} \right)$ -DP $\implies \left(\xi - \frac{\rho}{4} + 5\sqrt[4]{\rho}, \frac{\rho}{4} \right)$ -zCoDP (Lemma 3.7 in [BS16])
- (ρ, ω) -tCoDP^[BDRS18] $\implies (\hat{\varepsilon}, \delta)$ -DP, where $\hat{\varepsilon} = \rho + 2\sqrt{\rho \log \frac{1}{\rho}}$ if $\log \frac{1}{\delta} \leq (\omega - 1)^2 \rho$, and $\hat{\varepsilon} = \rho\omega + \frac{\log \frac{1}{\delta}}{\omega - 1}$ otherwise (Lemma 6 in [BDRS18])

3.4 Extension

Most definitions of this section can be seen as bounding the divergence between $\mathcal{M}(D_1)$ and $\mathcal{M}(D_2)$, for different possible divergence functions. In [BD14], the authors use this fact to generalize them and define (f, ε) -divergence DP (DivDP), which takes the particular divergence used as a parameter f .

Definition 13 ((f, ε) -divergence differential privacy [BD14]). *Let f be a convex function such as $f(1) = 0$. A privacy mechanism \mathcal{M} is (f, ε) -divergence DP if for all pairs of neighboring datasets D_1, D_2 :*

$$\mathbb{E}_{O \sim \mathcal{M}(D_1)} [f(e^{\mathcal{L}_{D_1/D_2}})] \leq \varepsilon.$$

An instance of this definition was presented in [DR18] as (f_k, ε) -divergence DP; which requires that $\mathbb{E}_{O \sim \mathcal{M}(D_1)} [|e^{\mathcal{L}_{D_1/D_2}} - 1|^k] \leq \varepsilon^k$. This definition is mainly used to prove technical results on privacy/utility tradeoffs in the local model. For any $k \leq 1$, ε -DP implies $(f_k, e^\varepsilon - 1)$ -DivDP, and when $k = 2$, it is equivalent to $(2, \log(1 + \varepsilon^2))$ -RényiDP (Section 2 in [DR18]).

Moreover, *capacity bounded differential privacy* (CBDP) was introduced in [CIM19], which uses H -restricted f -divergence: $D_f^H(P|Q) = \sup_{h \in H} [\mathbb{E}_{x \sim P} [h(x)] - \mathbb{E}_{x \sim Q} [f^*(h(x))]]$ where f is a divergence, H is a family of functions, and f^* is the Fenchel conjugate⁴. In other words, it requires the supremum condition to hold only for a selected set of functions (queries) instead of all possible ones.

Finally, most definitions in this section taking two real-valued parameters can be extended to use a *family* of parameters rather than a single pair of parameters. As shown in [SMM19] (Theorem 2) for approximate DP, probabilistic DP, and Rényi DP, finding the tightest possible family of parameters (for either definition) for a given mechanism is equivalent to specifying the behavior of its privacy loss random variable entirely.

3.5 Multidimensional definitions

Allowing a small probability of error δ by using the same concept as in (ε, δ) -DP is very common; many new DP definitions were proposed in the literature with such a parameter. Unless it creates a particularly notable effect, we do not mention it explicitly and present the definitions without this parameter.

Definitions in this section can be used as standalone concepts: (ε, δ) -DP is omnipresent in the literature, and the principle of averaging risk is natural enough for Rényi privacy to be used in practical settings, like posterior sampling [GSC17] or resistance to adversarial inputs in machine learning [PYGPA19]. Most variants in this section, however, are only used as technical tools to get better results on composition or privacy amplification [DR⁺14, WBK19, FMTT18, LK18].

4 Neighborhood definition (N)

The original definition of differential privacy considers datasets differing in one record. Thus, the datasets can differ in two possible ways: either they have the same size and differ only on one record, or one is a copy of the other with one extra record. These two options do not protect the

⁴The Fenchel conjugate for a function f with a domain R is $f^*(x) = \sup_{y \in R} [xy - f(y)]$.

same thing: the former protects the *value* of the records while the latter also protects their *presence* in the data: together, they protect any property about a single individual.

In many scenarios, it makes sense to protect a different property about their dataset, e.g., the value of a specific sensitive field, or entire groups of individuals. It is straightforward to adapt DP to protect different sensitive properties: all one has to do is change the definition of neighborhood in the original definition.

4.1 Changing the sensitive property

The original definition states that the ε -indistinguishability property should hold for “any datasets D_1 and D_2 that differ only on the data of one individual”. Modifying the set of pairs (D_1, D_2) such that $\mathcal{M}(D_1) \approx_\varepsilon \mathcal{M}(D_2)$ is equivalent to changing the protected sensitive property.

Weaker relaxations

In DP, the difference between D_1 and D_2 is sometimes interpreted as “one record value is different”, or “one record has been added or removed”. In [KM11], the authors formalize these two options as *bounded DP* and *unbounded DP*. They also introduced *attribute DP* and *bit DP*, for smaller changes within the differing record.

Definition 14 ([KM11]). *If a privacy mechanism \mathcal{M} satisfies $\mathcal{M}(D_1) \approx_\varepsilon \mathcal{M}(D_2)$ for any pair D_1, D_2 , where D_1 can be obtained from D_2 by...*

- ... adding or removing one record, then \mathcal{M} is ε -unbounded DP (*uBoDP*).
- ... changing exactly one record, then \mathcal{M} is ε -bounded DP (*BoDP*).
- ... changing one attribute in a record, then \mathcal{M} is ε -attribute DP (*AttDP*).
- ... changing one bit of an attribute in a record, then \mathcal{M} is ε -bit DP (*BitDP*).

In [KM11], authors show that ε -unbounded DP implies 2ε -bounded DP, as changing a record can be seen as deleting it and adding a new one in its place. The original definition of ε -DP is the conjunction of ε -unbounded DP and ε -bounded DP. However, bounded DP is frequently used in the literature, especially when using local differential privacy, and often simply called differential privacy. It is also sometimes renamed, like in [FMTT18], where the authors call it *per-person DP*.

Another way to relax the neighborhood definition in DP is to consider that only certain types of information are sensitive. For example, if the attacker learns that their target has cancer, this is more problematic than if they learn that their target does *not* have cancer. This idea is captured in *one-sided DP* (OSDP) [KDH⁺20]: the neighbors of a dataset D are obtained by replacing a single sensitive record with any other record (sensitive or not). The idea of sensitivity is formalized by a *policy* P , which specifies which records are sensitive. This idea cannot be captured simply by ε -indistinguishability, since one-sided DP is asymmetric.

Definition 15 ((P, ε) -one-sided differential privacy [KDH⁺20]). *Given a policy $P \subseteq \mathcal{T}$, a privacy mechanism \mathcal{M} is (P, ε) -one-sided DP iff for all datasets D_1 and D_2 , where D_2 has been obtained by replacing a record $t \in D_1 \cap P$ by any other record and for all $S \subseteq \mathcal{O}$:*

$$\mathbb{P}[\mathcal{M}(D_1) \in S] \leq e^\varepsilon \cdot \mathbb{P}[\mathcal{M}(D_2) \in S].$$

When $P = \mathcal{T}$, this is equivalent to bounded DP. Similar ideas were proposed in multiple papers:

- In [TCY21], the authors propose *asymetric DP*, which is the unbounded version of OSDP.
- In [APV19], the authors propose *sensitive privacy*, which determines which records are sensitive based on the data itself and a *normality property* N and a graph-based definition of k -neighborhood, instead of using a data-independent determination.
- In [BSW18], the authors introduce *anomaly-restricted DP*, which assumes that there is only one outlier in the dataset, and that this outlier should not be protected.

Stronger notions

More restrictive definitions are also possible. First, some definitions make the definition of neighborhood more explicit when a single person can contribute multiple times to a dataset; this is the case for *client/participant DP*, defined in [MRTZ18]. Variants of differential privacy that do not protect individuals, but single contributions (in the case where the same person can contribute multiple times to a dataset), are also often used in practice, especially for machine learning applications [MRTZ18]. Some recent works also argue that in-between definitions are appropriate: rather than protecting a single contribution or entire users contributions, authors in [ADJ19] suggest that protecting *elements* that reveal information about users, after deduplicating or clustering contributions. For example, rather than protecting all website visits by a single user, or each visit individually, one might choose to protect the fact that a user ever visited a website (but not whether the user visited the same website once or many times). They call the corresponding definition *element-level DP* (ELDP).

In [Dwo08], the authors implicitly define (c, ε) -group privacy considers datasets that do not differ in one record, but possibly several, to protect multiple individuals. This can also be interpreted as taking correlations into account when using DP: *DP under correlation* [CFYD14] uses an extra parameter to describe the maximum number of records that the change of one individual can influence.

These two definitions are formally equivalent; but the implicit interpretation of DP behind them is different. (c, ε) -group privacy is compatible with the associative view under the strong adversary assumption (the adversary knows all records except c) or the causal view (c records are changed after the data is generated). Meanwhile, DP under correlation implicitly considers the associative view with the independence assumption; and tries to relax that assumption. This last approach was further developed via *dependent DP* [LCM16], which uses “dependence relationships” to describe how much the variation in one record can influence the other records.

Definition 16 ((R, c, ε) -dependent differential privacy [LCM16]). *A privacy mechanism \mathcal{M} provides (R, c, ε) -dependent DP (DepDP) where R is the probabilistic dependence relationship and c is the dependence size, if for any pair of datasets D_1 and D_2 , where D_2 has been obtained from D_1 by changing one record and the corresponding at most $c - 1$ other records according to R , $\mathcal{M}(D_1) \approx_\varepsilon \mathcal{M}(D_2)$.*

Note that when R is the empty relation, or when $c = 1$, this definition is equivalent to bounded DP: under the associative view of DP, this represents independence between records. Similar definitions appear in [WW21a, WW21b] as *correlated indistinguishability* and *correlated tuple DP* respectively, in which correlations are defined by a correlation matrix. Furthermore, in [WDN17, WWK⁺17] as *correlated DP* (CorDP), in which correlations are defined by an observation on other

datasets, and in in [YSN15] as *bayesian DP*⁵ (BayDP^[YSN15]), where the neighborhood relation is defined by an adversary having some knowledge about correlations in the data. An extension is proposed in [LRY19] as *prior DP* (PriDP) which considers a family of adversaries instead of a single adversary.

The strongest possible variant is considered in [KM11], where the authors define *free lunch privacy* (FLPr), in which the attacker must be unable to distinguish between any two datasets, even if they are completely different. This guarantee is a reformulation of Dalenius’ privacy goal [Dal77]; as such, all mechanisms that satisfy free lunch privacy have a near-total lack of utility.

Definition 17 (ϵ -free lunch privacy [KM11]). *A privacy mechanism \mathcal{M} satisfies ϵ -free lunch privacy if $\mathcal{M}(D_1) \approx_\epsilon \mathcal{M}(D_2)$ for any pair of datasets D_1, D_2 .*

4.2 Limiting the scope of the definition

Redefining the neighborhood property can also be used to reduce the scope of the definitions. In [SCDFSM17], the authors note that DP requires ϵ -indistinguishability of results between any pair of neighboring data sets, but in practice, the data custodian has only *one* data set D they want to protect. Thus, they only require ϵ -indistinguishability between this data set D and all its neighbors, calling the resulting definition *individual DP* (InDP). An equivalent definition was proposed in [CH16] as *conditioned DP*.

Definition 18 ((D, ϵ) -individual differential privacy [SCDFSM17]). *Given a dataset $D \in \mathcal{D}$, a privacy mechanism \mathcal{M} satisfies (D, ϵ) -individual DP if for any data set D' that differs in at most one record from D , $\mathcal{M}(D) \approx_\epsilon \mathcal{M}(D')$.*

This definition was further restricted in [RW21] where besides fixing a dataset D , a record t is also fixed.

4.3 Applying the definition to other types of input

Many adaptations of DP are simply changing the neighborhood definition to protect different types of input data than datasets. A few examples follow.

- Location

In [EG16] the authors defined *location privacy*, in which neighbors are datasets which differ in at most one record, and the two differing records are at a physical distance smaller than a given threshold. This definition also appears in [CBY⁺18] as *DP on r -location set*⁶. Several more location-related DP variants were defined in [NV20]: *untrackability* (which adopts differential privacy for set of locations by protecting whether they originated from a single user or by two users), *undetectability* and *multi user untrackability*, (which extend this idea further by not assuming both sets originated from the same private data and to multiple users respectively).

- Graph

In [HLMJ09] the authors adopt DP to graph-structured data and present multiple alternative

⁵There are two other notions with the same name: introduced in [TF20, LL12], we mention them in Section 5 and 6 respectively.

⁶Distinct from *DP on δ -location set* [XX15], which we mention in Chapter 5.

definitions, which protect different parts of the graph: the strongest is *node-DP*, which protects a node and all its edges; the weakest is *edge-DP*, only protects one edge; and an intermediate definition is *k-edge-DP*, which protects a total of k edges and nodes. Similarly to one-sided DP, in [TC12, KRWY16], the authors introduce *out-link privacy*, and *protected DP* which protects all outgoing edges from a given node and guarantees that no observer can learn much about the set of edges corresponding to any protected node respectively. In addition, in [Reu18], the author introduces *QL-edged-labeled DP* which only protecting a predetermined subset of outgoing edges. In [Pin18], the author introduces *l_1 -weighted DP*, in which graph edges are weighted, and graphs are neighbors when the total weight of their differing edges is smaller than 1; this notion was also defined implicitly in [Sea16]. In [ZLBR20] the authors defined *feasible-node DP* for control-flow graphs where an addition or removal of a node results in a feasible run-time behaviors. In [SXX⁺19], the authors define *decentralized DP* which extends the graph neighborhood to two jumps. Finally, in [DWWG13] the authors introduce *seamless privacy*, which rather than protecting characteristics of a specific input graph, it ensures that certain pairs of queries on this graph return similar answers.

- Stream

Several authors adapt DP to a streaming context, where the attacker can access the mechanism’s internal states. In [DNPR10, DNP⁺10, Dwo10], authors define *pan-privacy*, which comes in two variants: *event-level* pan-privacy (called *strong DP* in [WSMD20]) protects individual events, and *user-level* pan-privacy protects all events associated to a single user. In [KPXP14], the authors extend the previous idea and propose *w-event privacy*, which protects any event sequence occurring within a window of at most w timestamps. In [Far20, NV20] this was further extended to an infinite horizon via *discounted DP* (which keep assigning smaller-and-smaller weights to further-and-further events) and *everlasting privacy* (which limit the leakage of information users suffer, no matter how many executions a mechanism had), respectively. Finally, *series-indistinguishability* [WX17] captured data correlations in the streaming context and the authors in [LNP13, PHB21] adopted DP for Kalman-filters and to time-independent power spectral densities respectively.

- RAM and PIR

In [WCM18] the authors adopt DP for Random Access Memory and Private Information Retrieval. For RAM the neighborhood is defined over the sequence of logical memory requests over time; the same notion appears in [CCMS19] as *differential obliviousness* and in [ADK⁺19] as *oblivious DP*. The adaptation of neighborhood is similar in case of PIR; a similar notion appears in [TDG16] as *ϵ -private PIR* and in [PPY19] as *ϵ -DPIR*. Additionally, in [KKNO17], the authors use a similar idea to define DP for outsourced database systems.

- Text and Images

In [JLH19], the authors adapt DP for symbolic control systems, and introduce *word-DP* and *substitution-word-DP*, protecting respectively pairs of words whose Levenshtein distance is lower than a given parameter, or whose Hamming distance is lower than a given parameter. In [ZSZ⁺18], the authors adapt DP for text vectors, and propose *text indistinguishability*, in which the neighborhood relationship between two word vectors depends on their Euclidean distance. In [YWW13, Fan18, LDX⁺21], the authors define *refinement DP*, *DP Image*, and *Pixel DP* respectively, which adopts the definition for images with neighbors given by some transformation or metric.

- Miscellaneous

Beside the already mentioned fields DP was adopted to numerous other use-cases, such as for set operations in [YLL⁺17], for gossip protocols in [HD19], for functions in [Noz19], for genomic data in [SSB16], for recommendation systems in [GKPT15], for machine learning in [LBG17], and for bandit algorithms in [TD16, BDT19].

4.4 Extensions

It is natural to generalize the variants of this section to arbitrary neighboring relationships. One example is mentioned in [KM11], under the name *generic DP*⁷ ($\text{GcDP}^{\text{[KM11]}}$), where the neighboring relation is entirely captured by a *relation* \mathcal{R} between datasets.

Definition 19 ($(\mathcal{R}, \varepsilon)$ -generic differential privacy [KM11]). *Given a relation $\mathcal{R} \subseteq \mathcal{D}^2$, a privacy mechanism \mathcal{M} satisfies $(\mathcal{R}, \varepsilon)$ -generic DP if for all $(D_1, D_2) \in \mathcal{R}$, $\mathcal{M}(D_1) \approx_\varepsilon \mathcal{M}(D_2)$.*

This definition is symmetric, but it can easily be modified to accommodate asymmetric definitions like one-sided DP.

Other definitions use different formalizations to also generalize the concept of changing the neighborhood relationship. Some (like pufferfish privacy, mentioned in Section 6) use *pairs of predicates* (ϕ_1, ϕ_2) that D_1 and D_2 must respectively satisfy to be neighbors. Others (like coupled-worlds privacy, mentioned in Section 8) use *private functions*, denoted priv , and define neighbors to be datasets D_1 and D_2 such as $\text{priv}(D_1) \neq \text{priv}(D_2)$ [BGKS13]. Others use a *distance function* d between datasets, and neighbors are defined as datasets a distance lower than a given threshold Δ ; this is the case for *DP under a neighborhood* (DPUN), introduced in [FC14], *adjacent DP* (AdjDP), introduced in [KM20]⁸, *constrained DP* (ConsDP), introduced in [ZLW09] (where the distance d captures a utility-related constraint), and *distributional privacy*⁹ ($\text{DIPr}^{\text{[ZLW09]}}$), also introduced in [ZLW09] (with additional constraints on the neighborhood definition: neighboring datasets must be part of a fixed set $S_{\mathcal{D}}$ and have elements in common). This distance can also be defined as the sensitivity of the mechanism, like in *sensitivity-induced DP* [RA17] (SIDP), or implicitly defined by a set of constraints, like what is done implicitly in [KM11] via *induced neighbors DP* (INDP).

One notable instantiation of generic DP is *blowfish privacy* (BFPr) [HMD14]. Its major building blocks are a *policy graph* G , that specifies which pairs of domain values in \mathcal{T} should not be distinguished between by an adversary; and a set of *constraints* Q that specifies the set \mathcal{I}_Q of possible datasets that the definition protects. It was inspired by the Pufferfish framework [KM12] (see Section 6), but the attacker is not assumed to have uncertainty over the data: instead, it models an attacker whose knowledge is a set of deterministic constraints on the data.

Definition 20 ($(G, \mathcal{I}_Q, \varepsilon)$ -blowfish privacy [HMD14]). *Given a policy graph $G \in \mathcal{T}^2$ and a set of datasets \mathcal{I}_Q , a privacy mechanism \mathcal{M} satisfies (G, ε) -blowfish privacy if for all datasets D_1 and D_2 in \mathcal{I}_Q that differ in only one element i such that $(D_1(i), D_2(i)) \in G$, $\mathcal{M}(D_1) \approx_\varepsilon \mathcal{M}(D_2)$.*

As noted in [HMD14], a mechanism satisfies ε -bounded DP if and only if satisfies $(K, \mathcal{I}_n, \varepsilon)$ -blowfish privacy, where K is the complete graph, and \mathcal{I}_n is any datasets of size n . A particular instantiation of this idea is explored in [KM11] as *induced DP* (IndDP), where the definition of neighbors is induced by a set of constraints.

⁷Another definition with the same name is introduced in [KL10, KL12], we mention it in Section 5.

⁸Originally simply called “differential privacy” by its authors.

⁹Another definition with the same name is introduced in [Rot10, BLR13], we mention it in Section 5.

4.5 Multidimensional definitions

Modifying the protected property is orthogonal to modifying the risk model implied by the quantification of privacy loss: it is straightforward to combine these two dimensions. Indeed, many definitions mentioned in this section were actually introduced with a δ parameter allowing for a small probability of error. One particularly general and specific example is *adjacency relation divergence DP* [KM19a], and *node-Rényi DP* [DMS⁺21], which combines an arbitrary neighborhood definition (like in generic DP) with an arbitrary divergence function (like in divergence DP) and adopts Rényi DP to graphs respectively.

As the examples in Section 4.3 show, it is very common to change the definition of neighborhood in practical contexts to adapt what aspect of the data is protected. Further, local DP mechanisms like RAPPOR [EPK14] implicitly use bounded DP: the participation of one individual is not secret, only the value of their record is protected. Variants that limit the scope of the definition to one particular dataset or user, however, provide few formal guarantees and do not seem to be used in practice.

5 Variation of privacy loss (V)

In DP, the privacy parameter ε is *uniform*: the level of protection is the same for all protected users or attributes, or equivalently, only the level of risk for the most at-risk user is considered. In practice, some users might require a higher level of protection than others or a data custodian might want to consider the level of risk across all users, rather than only considering the worst case. Some definitions take this into account by allowing the privacy loss to vary across inputs, either explicitly (by associating each user to an acceptable level of risk), or implicitly (by allowing some users to be at risk, or averaging the risk across users).

5.1 Varying the privacy level across inputs

In Section 4, we saw how changing the definition of the neighborhood can be used to adapt the definition of privacy and protect different aspects of the input data. However, the privacy protection in those variants is binary: either a given property is protected, or it was not. A possible option to generalize this idea further is to allow the privacy level to vary across possible inputs.

One natural example is to consider that some users might have higher privacy requirements than others, and make the ε vary according to which user differs between the two datasets. This is done in *personalized DP* (PerDP), a notion first defined informally in [NAD14], then independently in [JYC15, ESS15, GR15, LWS15]. An equivalent notion is also defined in [AGK16] as *heterogeneous DP*, while a location-based definition is presented in [DA18] as *personalized location DP*. Note that in these definitions, the privacy level associated with each user is not considered sensitive, and cannot depend on the data itself.

Definition 21 (Ψ -personalized differential privacy [JYC15]). *A privacy mechanism \mathcal{M} provides Ψ -personalized DP if for every pair of neighboring datasets (D, D_{-i}) and for all sets of outputs $S \subseteq \mathcal{O}$:*

$$\mathbb{P}[\mathcal{M}(D_{-i}) \in S] \leq e^{\Psi(i)} \mathbb{P}[\mathcal{M}(D) \in S]$$

where Ψ is a privacy specification, which maps each index i to the privacy level of the record $D(i)$.

As shown in [JYC15, CD20], ϵ -DP implies Ψ -PerDP, where $\Psi(i) = \epsilon$ for all indices i ; and Ψ -PerDP implies ϵ -DP where $\epsilon = \max_i \Psi(i)$.

This definition can be seen as a refinement of the intuition behind one-sided DP, which separated records into sensitive and non-sensitive ones. The idea of making the privacy level vary across inputs can be generalized further, by also making the privacy level depend on the entire dataset and the data of the differing record. This is done in [LP15], where the authors define *tailored DP* (TaiDP).

Definition 22 (Ψ -tailored differential privacy [LP15]). *A mechanism \mathcal{M} satisfies Ψ -tailored differential privacy for $\Psi : \mathcal{T} \times \mathcal{D} \rightarrow \mathbb{R}_0^\infty$ if for any dataset D , $\mathcal{M}(D) \approx_{\epsilon(D(i), D)} \mathcal{M}(D_{-i})$.*

A similar notion is *input-discriminative DP* [GLXC20] where Ψ takes the two ϵ values corresponding to the two dataset. The authors also defined *minID-DP* where Ψ is the minimum function.

This concept can be applied to strengthen or weaken the privacy requirement for a record depending on whether they are an outlier in the dataset. In [LP15], the authors formalize this idea and introduce *outlier privacy*, which tailors an individual’s protection level to their “outlierness”. Other refinements are also introduced in [LP15]: *simple outlier privacy* (SOPr), *simple outlier DP* (SODP), and *staircase outlier privacy* (SCODP). A similar idea was explored in [KLL19], which introduced *pareto DP* (ParDP): it utilizes a pareto distribution of parameters (p, r) to separate a large number of low-frequency individuals from a small number of high-frequency, and the sensitivity is calculated based on only the low-frequency individuals.

Finally, varying the privacy level across inputs also makes sense in *continuous* scenarios, where the neighborhood relationship between two datasets is not binary, but quantified. This is, for example, the case for ϵ -*geo-indistinguishability* [ABCP13], where two datasets D_1 and D_2 are considered r -*neighbors* if the only different record between D_1 and D_2 are at a distance r of each other, and the ϵ grows linearly with r .

5.2 Randomizing the variation of privacy levels

Varying the privacy level across inputs can also be done in a randomized way, by guaranteeing that some random fraction of users have a certain privacy level. One example is proposed in [HWR13] as *random DP* (RanDP): the authors note that rather than requiring DP to hold for any possible datasets, it is natural to only consider *realistic datasets*, and allow “edge-case” or very unrealistic datasets to not be protected. This is captured by generating the data randomly, and allowing a small proportion γ of cases to not satisfy the ϵ -indistinguishability property.

Definition 23 ((π, γ, ϵ) -random differential privacy [HWR13]). *Let π be a probability distribution on \mathcal{T} , D_1 a dataset generated by drawing n i.i.d. elements in π , and D_2 the same dataset as D_1 , except one element was changed to a new element drawn from π . A mechanism \mathcal{M} is (π, γ, ϵ) -random DP if $\mathcal{M}(D_1) \approx_\epsilon \mathcal{M}(D_2)$, with probability at least $1 - \gamma$ on the choice of D_1 and D_2 .*

The exact meaning of “with probability at least $1 - \gamma$ on the choice of D_1 and D_2 ” can vary slightly. In [Hal12] and in [McC15], the authors introduce *predictive DP* (PredDP) and *model-specific DP* respectively, which quantify over all possible choices of D_1 , and picks D_2 randomly in the neighborhood of D_1 . In [DJRT13], D_1 and D_2 are both taken out of a set of density larger than $1 - \gamma$, and the authors call this definition *generalized DP* (GdDP). The distribution generating the dataset is also not always assumed to be generating i.i.d. records; we denote the corresponding parameter by θ .

Random DP might look similar to probabilistic DP: in both cases, there is a small probability that the privacy loss is unbounded. On the other hand, they are very different: in random DP, this probability is computed inputs of the mechanisms (i.e., users or datasets), for probabilistic DP, it is computed across mechanism outputs. Also similarly to probabilistic DP, excluding some cases altogether creates definitional issues: random DP does not satisfy the convexity axiom (see Proposition 4 in Section 10). We postulate that using a different tool to allow some inputs to not satisfy the mechanism, similar to approximate DP or Rényi DP, could solve this problem.

Usually, data-generating distributions are used for other purposes: they typically model an adversary with partial knowledge. However, definitions in this section still compare the outputs of the mechanisms given fixed neighboring datasets: the only randomness in the indistinguishability property comes from the mechanism. By contrast, definitions of Section 6 compare the output of the mechanism on a random dataset, so the randomness comes both from the data-generating distribution and the mechanism.

5.3 Multidimensional definitions

As varying the privacy level or limiting the considered datasets are two distinct way of relaxing differential privacy, it is possible to combine them with the previously mentioned dimensions.

Combination with N

The definitions described in Section 4 (e.g., generic DP or blowfish privacy) have the same privacy constraint for all neighboring datasets. Thus, they cannot capture definitions that vary the privacy level across inputs. On the other hand, they can be combined quite easily. For instance, [DA18] and [SPB21] adopted personalized DP to location and communication graphs respectively. Varying the privacy level across inputs also makes sense in “continuous” scenarios, where the neighborhood relationship between two datasets is not binary, but quantified. This is, for example, the case for *geo-indistinguishability* [ABCP13], where two datasets D_1 and D_2 are considered “ r -neighbors” if the only different record between D_1 and D_2 are at a distance r of each other, and the ϵ grows linearly with r . Both ideas can be naturally captured together via distance functions. In [CABP13], the authors introduce $d_{\mathcal{D}}$ -privacy ($d_{\mathcal{D}}$ -Pr), in which the function $d_{\mathcal{D}}$ takes both datasets as input, and returns the corresponding maximal privacy loss (the ϵ) depending on the difference between the two datasets.

Definition 24 ($d_{\mathcal{D}}$ -privacy [CABP13]). *Let $d_{\mathcal{D}} : \mathcal{D}^2 \rightarrow \mathbb{R}_{\infty}$. A privacy mechanism \mathcal{M} satisfies $d_{\mathcal{D}}$ -privacy if for all pairs of datasets D_1, D_2 and all sets of outputs $S \subseteq \mathcal{O}$:*

$$\mathbb{P}[\mathcal{M}(D_1) \in S] \leq e^{d_{\mathcal{D}}(D_1, D_2)} \cdot \mathbb{P}[\mathcal{M}(D_2) \in S].$$

When $d_{\mathcal{D}}$ is proportional to the Hamiltonian difference between datasets, this is equivalent to ϵ -DP. In $d_{\mathcal{D}}$ -privacy, the $d_{\mathcal{D}}$ function specifies both the privacy parameter *and* the definition of neighborhood: it can simply return ∞ on non-neighboring datasets, and vary the privacy level across inputs for neighboring datasets. In the original definition, the authors impose that $d_{\mathcal{D}}$ is *symmetric*, but this condition can also be relaxed to allow $d_{\mathcal{D}}$ -privacy to extend definitions like one-sided DP.

Equivalent definitions of $d_{\mathcal{D}}$ -privacy also appeared in [EG16] as *l-privacy*, and in [KM19b] as *extended DP*. Several other definitions, such as *weighted DP* [PGM14] (WeiDP), *smooth DP* [BD14]

(SmoDP)¹⁰ and *earth mover’s privacy* [FDM19] (EMDP), can be seen as particular instantiations of $d_{\mathcal{D}}$ -privacy for specific functions d measuring the distance between datasets. This is also the case for some definitions tailored for location privacy, like *geo-graph-indistinguishability* [TCAY19, TCAY20], which specifically applies to network graphs.

Random DP can also be combined with changing the neighborhood definition: in [XX15], the authors define *DP on a δ -location set*¹¹, for which the neighborhood is defined by a set of “plausible” locations around the true location of a user. A notable definition using the same combination of dimensions is *distributional privacy*¹², introduced in [Rot10, BLR13]: it combines random DP (for a large family of distributions) and free lunch privacy.

Definition 25 ((ε, γ) -distributional privacy [Rot10, BLR13]). *An algorithm \mathcal{M} satisfies (ε, γ) -distributional privacy ($\text{DIPr}^{\text{[Rot10, BLR13]}}$) if for any distribution π over possible tuples, if $D_1, D_2 \in \mathcal{D}$ are picked by randomly drawing a fixed number n of elements from π without replacement, then with probability at least $1 - \gamma$ over the choice of D_1 and D_2 , $\mathcal{M}(D_1) \approx_{\varepsilon} \mathcal{M}(D_2)$.*

Interestingly, this definition captures an intuition similar to the variants in 8: the adversary can only learn properties of the data-generating distribution, but not about particular samples (except with probability lower than γ). The authors also prove that if \mathcal{M} is (ε, γ) -DIPr and $\gamma = o(\frac{1}{n^2})$, where n is the size of the dataset being generated, then \mathcal{M} is also ε -DP.

Combination with Q

Different risk models, like the definitions in Section 3, are also compatible with varying the privacy parameters across inputs. For example, in [Kre19], the author proposes *endogeneous DP* (EndDP), which is a combination of (ε, δ) -DP and personalized DP. Similarly, *pseudo-metric DP* (PsDP), defined in [DNM⁺13], is a combination of $d_{\mathcal{D}}$ -privacy and (ε, δ) -DP; while *extended divergence DP* (EDivDP), defined in [KM19a], is a combination of $d_{\mathcal{D}}$ -privacy and divergence DP.

Randomly limiting the scope of the definition can also be combined with ideas from the previous sections. For example, in [TF20], the authors introduce *weak Bayesian DP* (WBDP), which combines random DP and approximate DP. In [LX21] the authors introduced *Smoothed DP*¹³ which require (ε, δ) -indistinguishability to hold for all the expected dataset from a set of distributions. In [WLF16], the authors introduce *on average KL privacy*, which uses KL-divergence as quantification metric, but only requires the property to hold for an “average dataset”, like random DP; a similar notion appears in [FS18] as *average leave-one-out KL stability*. In [TF20, DBB21], the authors introduce *Bayesian DP*¹⁴ (BayDP^[TF20]) and *privacy at risk* (PAR) respectively; both definitions combine random DP with probabilistic DP, with slightly different approaches: the former quantifies over all possible datasets and changes one fixed record randomly, while the latter selects both datasets randomly, conditioned on these datasets being neighbors.

In [KL10], Kifer et al. go further and generalize the intuition from generic DP, introduced in Section 4, and generalize the indistinguishability condition entirely. The resulting definition is *also* called generic differential privacy.

¹⁰Distinct from *Smoothed DP* [LX21] introduced later in this section.

¹¹Distinct from *DP on r -location set* [CBY⁺18], which we mention in Section 4.

¹²Another definition with the same name is introduced in [ZLW09], we mention it in Section 4.

¹³Distinct from *Smooth DP* [BD14] introduced earlier in this chapter.

¹⁴There are two other notions with the same name: introduced in [YSN15, LL12], we mention them in Section 4 and 6 respectively.

Definition 26 ((\mathcal{R}, M) -generic differential privacy [KL10, KL12]). A privacy mechanism \mathcal{M} satisfies (\mathcal{R}, M) -generic DP (GcDP^[KL10]) if for all measurable sets $S \subseteq \mathcal{O}$ and for all $(D_1, D_2) \in \mathcal{R}$:

$$M_{D_1, D_2}(\mathbb{P}[\mathcal{M}(D_1) \in S]) \geq \mathbb{P}[\mathcal{M}(D_2) \in S] \quad \text{and} \quad M_{D_1, D_2}(\mathbb{P}[\mathcal{M}(D_1) \notin S]) \geq \mathbb{P}[\mathcal{M}(D_2) \notin S]$$

where $M_{D_1, D_2} : [0, 1] \rightarrow [0, 1]$ is a concave function continuous on $(0, 1)$ such as $M_{D_1, D_2}(1) = 1$.

The privacy relation \mathcal{R} is still the generalization of neighborhood and the privacy predicate is the generalization of the ε -indistinguishability to arbitrary functions. In particular, it can encompass all variants of described in Section 3 in addition to the ones in this section: for example, if $M_{D_1, D_2}(x) = \min(1, xe^\varepsilon + \delta, 1 - (1 - x - \delta)e^{-\varepsilon})$ holds for for all D_1 and D_2 , then this is equivalent to (ε, δ) -DP. This definition was an attempt at finding the most generic definition that still satisfies privacy axioms: another extension defined in the same work, *abstract DP* (AbsDP) is even more generic, but no longer satisfies the privacy axioms.

Definitions in this section are particularly used in the context of local DP¹⁵ and in particular for applications to location privacy: various metrics have been discussed to quantify how indistinguishable different places should be to provide users of a local DP mechanism with meaningful privacy protection [CEP⁺17].

6 Background knowledge (B)

In differential privacy, the attacker is implicitly assumed to have full knowledge of the dataset: their only uncertainty is whether the target belongs in the dataset or not. This implicit assumption is also present for the definitions of the previous dimensions: indeed, the attacker has to distinguish between two *fixed* datasets D_1 and D_2 . The only source of randomness in ε -indistinguishability comes from the mechanism itself. In many cases, this assumption is unrealistic, and it is natural to consider weaker adversaries, who do not have full background knowledge. One of the main motivations to do so is to use significantly less noise in the mechanism [Dua09].

The typical way to represent this uncertainty formally is to assume that the input data comes from a certain probability distribution (named *data evolution scenario* in [KM12]): the randomness of this distribution models the attacker’s uncertainty. Informally, the more random this probability distribution is, the less knowledge the attacker has. However, the definition that follows depends whether DP is considered with the associative or the causal view. In the associative view, the sensitive property changes *before* the data is generated: it conditions the data-generating probability distribution. In the causal view, however, the sensitive property is only changed *after* the data is generated. The two options lead to very distinct definitions.

6.1 Conditioning the output on the sensitive property

Using a probability distribution to generate the input data means that the ε -indistinguishability property cannot be expressed between two fixed datasets. Instead, one natural way to express it is to condition this distribution on some sensitive property. The corresponding notion, *noiseless privacy*¹⁶ (NPr) was first introduced in [Dua09] and formalized in [BBG⁺11].

¹⁵For details, see Section 11.3.

¹⁶Another definition with the same name is introduced in [Far21], we mention it in Section 11.2.

Definition 27 ((Θ, ε) -noiseless privacy [Dua09, BBG⁺11]). *Given a family Θ of probability distribution on \mathcal{D} , a mechanism \mathcal{M} is (Θ, ε) -noiseless private if for all $\theta \in \Theta$, all i and all $t, t' \in \mathcal{T}$:*

$$\mathcal{M}(D)_{|D \sim \theta, D(i)=t} \approx_{\varepsilon} \mathcal{M}(D)_{|D \sim \theta, D(i)=t'}.$$

In the original definition, the *auxiliary knowledge* of the attacker is explicitly pointed out in an additional parameter. In the case where there is no δ of $(\varepsilon, 0)$ -DP, this syntactic add-on is not necessary [DMKB19], so we omitted it here.

This definition follows naturally from considering the associative view of DP with the strong adversary assumption, and attempting to relax this assumption. The exact way to model the adversary’s uncertainty can be changed; for example *DP under sampling* [LQS12], an instance of noiseless privacy, models it using random sampling.

This definition was not the first attempt at formalizing an adversary with restricted background knowledge. In [MGG09], authors define ε -*privacy*, which represents the background knowledge as a Dirichlet distribution (instead of an arbitrary distribution), whose parameters are interpreted as characteristics of the attacker. However, this definition imposes a condition on the output, but not on the mechanism which produced the output. As such, it does not offer strong semantic guarantees like the other definitions presented in this survey.

6.2 Removing the effect of correlations in the data

In [BGKS13], however, the authors argue that in the presence of correlations in the data, noiseless privacy can be too strong, and make it impossible to learn global properties of the data. Indeed, if one record can have an arbitrarily large influence on the rest of the data, conditioning on the value of this record can lead to very distinguishable outputs even if the mechanism only depends on global properties of the data. To fix this problem, they propose *distributional DP* (DistDP), an alternative definition that only conditions the data-generating distribution on one possible value of the target record, and quantifies the information leakage from the mechanism.¹⁷ In [DMKB19], the authors show that this creates definitional issues in the presence of background knowledge, and introduce *causal DP* (CausDP), to capture the same intuition without encountering the same problems.

Definition 28 ((Θ, ε) -causal differential privacy [DMKB19]). *Given a family Θ of probability distributions on \mathcal{D} , a mechanism \mathcal{M} satisfies (Θ, ε) -causal DP (CausDP) if for all probability distributions $\theta \in \Theta$, for all i and all $t, t' \in \mathcal{T}$:*

$$\mathcal{M}(D)_{|D \sim \theta, D(i)=t} \approx_{\varepsilon} \mathcal{M}(D_{i \rightarrow t'})_{|D \sim \theta, D(i)=t}.$$

where $D_{i \rightarrow t'}$ is the dataset obtained by changing the i -th record of D into t' .

In this definition, one record is changed *after* the dataset has been generated, so it does not affect other records through dependence relationships. These dependence relationships are the only difference between noiseless privacy and causal DP: when each record is independent of all others, this definition is equivalent to noiseless privacy [DMKB19].

¹⁷Note that the original formalization used in [BGKS13] was more abstract, and uses a simulator, similarly to variants introduced in Section 8.

6.3 Multidimensional definitions

Limiting the background knowledge of an attacker is orthogonal to the dimensions introduced previously: one can modify the risk model, introduce different neighborhood definitions, or even vary the privacy parameters across the protected properties along with limiting the attacker background knowledge.

Combination with Q

Modifying the risk model while limiting the attacker’s background knowledge has interesting consequences. In [DMKB19], the authors show that two options are possible: either consider the background knowledge as additional information given to the attacker or let the attacker *influence* the background knowledge. This distinction between an *active* and a *passive* attacker does not matter if only the worst-case scenario is considered, like in noiseless privacy. However, under different risk models, such as allowing a small probability of error, they lead to two different definitions.

Both of these definitions use an adapted version of the privacy loss random variable (PLRV) which explicitly models the attacker background knowledge: the data-generating distribution not only generates a dataset D but also some auxiliary knowledge B , with values in a set \mathcal{B} .

Definition 29 (PLRV for partial knowledge [DMKB19]). *Given a mechanism \mathcal{M} , a distribution θ with values in $\mathcal{D} \times \mathcal{B}$, an index i , and values $a, b \in \mathcal{T}$, the PLRV of an output $O \in \mathcal{O}$ given partial knowledge \hat{B} is:*

$$\mathcal{L}_{i \leftarrow a / i \leftarrow b}^{\mathcal{M}, \theta}(O, \hat{B}) = \ln \left(\frac{\mathbb{P}_{(D, B) \sim \theta} [\mathcal{M}(D) = O \mid D(i) = a, B = \hat{B}]}{\mathbb{P}_{(D, B) \sim \theta} [\mathcal{M}(D) = O \mid D(i) = b, B = \hat{B}]} \right).$$

if the three conditions are satisfied:

1. $\mathbb{P}_{(D, B) \sim \theta} [D(i) = a, B = \hat{B}] \neq 0$ and $\mathbb{P}_{D \sim \theta} [D(i) = b, B = \hat{B}] \neq 0$
2. $\mathbb{P}_{(D, B) \sim \theta} [\mathcal{M}(D) = O \mid D(i) = a, B = \hat{B}] \neq 0$
3. $\mathbb{P}_{(D, B) \sim \theta} [\mathcal{M}(D) = O \mid D(i) = b, B = \hat{B}] \neq 0$.

If condition 1 does not hold, then $\mathcal{L}_{i \leftarrow a / i \leftarrow b}^{\mathcal{M}, \theta}(O, B) = 0$. Else, if condition 2 does not hold, then $\mathcal{L}_{i \leftarrow a / i \leftarrow b}^{\mathcal{M}, \theta}(O, B) = -\infty$. Else, if condition 3 does not hold, then $\mathcal{L}_{i \leftarrow a / i \leftarrow b}^{\mathcal{M}, \theta}(O, B) = \infty$.

This formalization can then be used to adapt noiseless privacy to a risk model similar to (ε, δ) -DP, in the case of an active or a passive attacker. The active variant, *active partial knowledge differential privacy* (APKDP), quantifies over all possible values of the background knowledge. It was first introduced in [BBG⁺11, BGKS13] as *noiseless privacy* and reformulated in [DMKB19] to clarify that it implicitly assumes an active attacker.

Definition 30 ($(\Theta, \varepsilon, \delta)$ -active partial knowledge differential privacy [BBG⁺11, BGKS13, DMKB19]). *Given a family Θ of probability distribution on $\mathcal{D} \times \mathcal{B}$, a mechanism \mathcal{M} is $(\Theta, \varepsilon, \delta)$ -active partial knowledge DP (APKDP) if for all $\theta \in \Theta$, all indices i , all $t, t' \in \mathcal{T}$, and all possible values \hat{B} of the background knowledge:*

$$\mathbb{E}_{(D, B) \sim \theta_{|D(i)=t, B=\hat{B}}, O \sim \mathcal{M}(D)} \left[\max \left(0, 1 - e^{\varepsilon - \mathcal{L}_{i \leftarrow t / i \leftarrow t'}^{\mathcal{M}, \theta}(O, \hat{B})} \right) \right] \leq \delta.$$

One specialization of this definition is *DP under sampling* [LQS12] (DPuS), which mandates DP to be satisfied after a random sampling is applied to the dataset. The authors use this definition to show that applying k -anonymity to a randomly sampled dataset provides differential privacy; but this definition could also be used on its own, to model the attacker’s uncertainty using a randomly sampled distribution.

The second definition, *passive partial knowledge differential privacy* [DMKB19] (PPKDP), is strictly weaker: it models a passive attacker, who cannot choose their background knowledge, and thus cannot influence the data. In this context, δ does not only apply to the output of the mechanism, but also to the value of the background knowledge.

Definition 31 ($(\Theta, \varepsilon, \delta)$ -passive partial knowledge differential privacy [DMKB19]). *Given a family Θ of probability distribution on $\mathcal{D} \times \mathcal{B}$, a mechanism \mathcal{M} is $(\Theta, \varepsilon, \delta)$ -passive partial knowledge DP (PPKDP) if for all $\theta \in \Theta$, all indices i , and all $t, t' \in \mathcal{T}$:*

$$\mathbb{E}_{(D,B) \sim \theta|_{D(i)=t}, O \sim \mathcal{M}(D)} \left[\max \left(0, 1 - e^{\varepsilon - \mathcal{L}_{i \leftarrow t / i \leftarrow t'}^{\mathcal{M}, \theta}(O, B)} \right) \right] \leq \delta.$$

Causal DP can also be adapted to a risk model similar to (ε, δ) -DP: in [BD19], authors introduce a similar notion to causal DP as *inherit DP* (InhDP), with the small difference that the second dataset is obtained by removing one record from the first dataset, instead of replacing it; and (ε, δ) -indistinguishability is used. The authors also define *empirical DP* [BD19]¹⁸, which is identical, except the empirical distribution is used instead of the actual data distribution, in context where the latter is unknown. In both cases, the influence of δ on the attacker model is unclear.

Combination with N

Modifying the neighborhood definition is simpler: it is clearly orthogonal to the dimensions introduced in this section. In all definitions of this section so far, the two possibilities between which the adversary must distinguish are similar to bounded DP. This can easily be changed to choose other properties to protect from the attacker. This is done in *pufferfish privacy* [KM12] (PFPr), which extends the concept of neighboring datasets to neighboring *distributions* of datasets.

Definition 32 ($(\Theta, \Phi, \varepsilon)$ -pufferfish privacy [KM12]). *Given a family of probability distributions Θ on \mathcal{D} , and a family of pairs of predicates Φ on datasets, a mechanism \mathcal{M} verifies $(\Theta, \Phi, \varepsilon)$ -pufferfish privacy if for all distributions $\theta \in \Theta$ and all pairs of predicates $(\phi_1, \phi_2) \in \Phi$:*

$$\mathcal{M}(D)|_{D \sim \theta, \phi_1(D)} \approx_{\varepsilon} \mathcal{M}(D)|_{D \sim \theta, \phi_2(D)}$$

Pufferfish privacy extends the concept of neighboring datasets to neighboring *distributions* of datasets; starting with a set of data-generating distributions, then conditioning them on sensitive attributes. The result compares pairs of distributions encompasses noiseless privacy, as well as other notions. For example, it captures *bayesian DP*¹⁹ (BayDP^[LL12]), introduced in [LL12], in which neighboring records have up to k fixed elements in common and all other elements are generated randomly from a distribution π .

The same idea can be formalized by comparing pairs of distributions directly. This is done in [JB14, KM19b] via *distribution privacy* (DnPr). The two formalisms are equivalent: an arbitrary pair of distributions can be seen as a single distribution, conditioned on the value of a secret

¹⁸Another definition with the same name is introduced in [ASV13], we mention it in Section 11.2.

¹⁹There are two other notions with the same name: introduced in [YSN15, TF20], we mention them in Section 4 and 5 respectively.

parameter. Distribution privacy was instantiated in [GC19] via *profile-based DP* (PBDP), in which the attacker tries to distinguish between different probabilistic user profiles. A similar idea was proposed in [BCDP19] as *robust privacy*, which uses lossy Wasserstein distance over the corresponding outputs to define the neighbourhood of the inputs.

Further relaxations encompassing the introduced dimensions are *probabilistic distribution privacy* [KM19b] (PDnPr), a combination of distribution privacy and probabilistic DP, *extended distribution privacy* [KM19b] (EDnPr), a combination of distribution privacy and $d_{\mathcal{D}}$ -privacy, *divergence distribution privacy* [KM19a] (DDnPr), a combination of distribution privacy, and *extended divergence distribution privacy* [KM19a] (EDDnPr), which combines the latter two definitions. Finally, *divergence distribution privacy with auxiliary inputs* [KM19a] considers the setting where the attacker might not know the input probability distribution perfectly.

Definitions of this section are an active area of research; a typical question is to quantify in which conditions deterministic mechanisms can provide some level privacy. However, they are not used a lot in practice, likely because of their fragility: if the assumptions about the attacker’s limited background knowledge are wrong in practice, then the definitions do not provide any guarantee of protection.

7 Change in formalism (F)

The definition of differential privacy using ε -indistinguishability compares the distribution of outputs given two neighboring inputs. This is not the only way to capture the idea that a attacker should not be able to gain too much information on the dataset. Other formalisms have been proposed, which model the attacker more explicitly.

One such formalism reformulates DP in terms of hypothesis testing by limiting the type I and the type II error of the hypothesis that the output O of a mechanism originates from D_1 (instead of D_2). Other formalisms model the attacker explicitly, by formalizing their prior belief as a probability distribution over all possible datasets. This can be done in two distinct ways. Some variants consider a specific prior (or family of possible priors) of the attacker, implicitly assuming a limited background knowledge, like in Section 6. We show that these variants can be interpreted as changing the prior-posterior bounds of the attacker. Finally, rather than comparing prior and posterior, a third formalism compares two possible posteriors, quantifying over all possible priors.

Definitions in this section provide a deeper understanding of the guarantees given by differential privacy, and some of them lead to tighter and simpler theorems on differential privacy, like composition or amplification results.

7.1 Hypothesis testing

First, differential privacy can be interpreted in terms of *hypothesis testing* [WZ10, KOV17]. In this context, an adversary who wants to know whether the output O of a mechanism originates from D_1 (the *null hypothesis*) or D_2 (the *alternative hypothesis*). Calling S the rejection region, the probability of false alarm (type I error), when the null hypothesis is true but rejected, is $\mathbb{P}_{FA} = \mathbb{P}[\mathcal{M}(D_1) \in S]$. The probability of missed detection (type II error), when the null hypothesis is false but retained, is $\mathbb{P}_{MD} = \mathbb{P}[\mathcal{M}(D_2) \in \mathcal{O}/S]$.

It is possible to use these probabilities, to reformulate DP:

$$\begin{aligned}
\varepsilon\text{-DP} &\Leftrightarrow \begin{aligned} &\mathbb{P}_{FA} + e^\varepsilon \mathbb{P}_{MD} \geq 1 \\ &e^\varepsilon \mathbb{P}_{FA} + \mathbb{P}_{MD} \geq 1 \end{aligned} \quad \text{for all } S \subseteq \mathcal{O} \\
\varepsilon\text{-DP} &\Leftrightarrow \mathbb{P}_{FA} + \mathbb{P}_{MD} \geq \frac{2}{1+e^\varepsilon} \\
(\varepsilon, \delta)\text{-DP} &\Leftrightarrow (\mathbb{P}_{FA}, \mathbb{P}_{MD}) = \{(\alpha, \beta) \in [0, 1] \times [0, 1] : \\ &\quad (1 - \alpha \leq e^\varepsilon \beta + \delta)\}
\end{aligned}$$

This hypotheses testing interpretation was used in [DRS21] to define *f-differential privacy* (*f*-DP), which avoids difficulties associated with divergence based relaxations. Specifically, its composition theorem is lossless as it provides a computationally tractable tool for analytically approximating the privacy loss. Moreover, there is a general duality between *f*-DP and infinite collections of (ε, δ) -DP guarantees.

Definition 33 (*f*-differential privacy [DRS21]). *Let $f : [0, 1] \rightarrow [0, 1]$ be a convex, continuous, and non-increasing function such that for all $x \in [0, 1]$, $f(x) \leq 1 - x$. A privacy mechanism \mathcal{M} satisfies *f*-DP if for all neighboring D_1, D_2 and all $x \in [0, 1]$:*

$$\inf_S \{1 - \mathbb{P}[\mathcal{M}(D_2) \in S] \mid \mathbb{P}[\mathcal{M}(D_1) \in S] \leq x\} \geq f(x).$$

Here, S is the rejection region; and the infimum is the trade-off function between $\mathcal{M}(D_1)$ and $\mathcal{M}(D_2)$. The authors also introduce *Gaussian differential privacy* (GaussDP) as an instance of *f*-differential privacy, which tightly bounds from below the hardness of determining whether an individual’s data was used in a computation than telling apart two shifted Gaussian distributions. Besides, *weak* and *strong federated f-DP* [ZCLS21] are also defined that adopts the definition for federated learning [M⁺21]. They describe the privacy guarantee against an individual and a group of adversaries respectively.

7.2 Changing the shape of the prior-posterior bounds

Differential privacy can be interpreted as giving a bound on the posterior of a Bayesian attacker as a function of their prior. This is exactly the case in *indistinguishable privacy* (IndPr), an equivalent reformulation of differential privacy defined in [LXL13]: suppose that the attacker is trying to distinguish between two options $D = D_1$ and $D = D_2$, where D_1 corresponds to the option “ $t \in D$ ” and D_2 to “ $t \notin D$ ”. Initially, they associate a certain prior probability $\mathbb{P}[t \in D]$ to the first option. When they observe the output of the algorithm, this becomes the posterior probability $\mathbb{P}[t \in D \mid \mathcal{M}(D) = O]$. Combining the definition of ε -DP and the Bayes Theorem, we have²⁰:

$$\frac{\mathbb{P}[t \in D \mid \mathcal{M}(D) = O]}{\mathbb{P}[t \notin D \mid \mathcal{M}(D) = O]} \leq e^\varepsilon \cdot \frac{\mathbb{P}[t \in D]}{\mathbb{P}[t \notin D]} \Rightarrow \mathbb{P}[t \in D \mid \mathcal{M}(D) = O] \leq \frac{e^\varepsilon \cdot \mathbb{P}[t \in D]}{1 + (e^\varepsilon - 1) \mathbb{P}[t \in D]}.$$

A similar, symmetric lower bound can be obtained. Hence, differential privacy can be interpreted as bounding the posterior level of certainty of a Bayesian attacker as a function of its prior. We visualize these bounds on the first two figures of Figure 2.

Some variants of differential privacy use this idea in their formalism, rather than obtaining the posterior bound as a corollary to the classical DP definition. For example, *positive membership privacy* (PMPPr) [LQS⁺13] requires that the posterior does not increase too much compared to the prior. Like noiseless privacy, it assumes an attacker with limited background knowledge.

²⁰Note that the original formalization used in [LXL13] was more abstract, and it used polynomially bounded adversaries what we introduce in Section 9.

Definition 34 ((Θ, ε) -positive membership privacy [LQS⁺13]). *A privacy mechanism \mathcal{M} provides (Θ, ε) -positive membership privacy if for any distribution $\theta \in \Theta$, any record $t \in \mathcal{D}$ and any $S \subseteq \mathcal{O}$:*

$$\mathbb{P}_{D \sim \theta} [t \in D | \mathcal{M}(D) \in S] \leq e^\varepsilon \mathbb{P}_{D \sim \theta} [t \in D] \quad \text{and} \quad \mathbb{P}_{D \sim \theta} [t \notin D | \mathcal{M}(D) \in S] \geq e^{-\varepsilon} \mathbb{P}_{D \sim \theta} [t \notin D].$$

Note that this definition is *asymmetric*: the posterior is bounded from above, but not from below. In the same paper, the authors also define *negative membership privacy* (NMP_r), which provides the symmetric lower bound, and *membership privacy*²¹ (MP_r), which is the conjunction of positive and negative membership privacy. They show that this definition can represent differential privacy (in its bounded and unbounded variants), as well as other definitions like *differential identifiability* [LC12] and *sampling DP* [LQS12], which we mention in Section 11.2. Bounding the ratio between prior and posterior by e^ε is also done in the context of location privacy: in [DGY⁺18], authors define ε -DP location obfuscation, which formalizes the same intuition as membership privacy.

A previous attempt at formalizing the same idea was presented in [RHMS09] as *adversarial privacy* (AdvPr). This definition is similar to positive membership privacy, except only the first relation is used, and there is a small additive δ as in approximate DP. We visualize the corresponding bounds on the third figure of Figure 2.

Definition 35 ($(\Theta, \varepsilon, \delta)$ -adversarial privacy [RHMS09]). *An algorithm \mathcal{M} is $(\Theta, \varepsilon, \delta)$ -adversarial private if for all $S \subseteq \mathcal{O}$, tuples t , and distributions $\theta \in \Theta$:*

$$\mathbb{P}_{D \sim \theta} [t \in D | \mathcal{M}(D) \in S] \leq e^\varepsilon \cdot \mathbb{P}_{D \sim \theta} [t \in D] + \delta.$$

Adversarial privacy (without δ) was also redefined in [WXH17] as *information privacy*²².

Finally, *aposteriori noiseless privacy* (ANPr) is a similar variant of noiseless privacy introduced in [BBG⁺11]; the corresponding bounds can be seen on the last figure of Figure 2.

Definition 36 ((Θ, ε) -aposteriori noiseless privacy [BBG⁺11]). *A mechanism \mathcal{M} is said to be (Θ, ε) -aposteriori noiseless private (ANPr) if for all $\theta \in \Theta$, all $S \subseteq \mathcal{O}$ and all i :*

$$D(i)_{|D \sim \theta, \mathcal{M}(D) \in S} \approx_\varepsilon D(i)_{|D \sim \theta}.$$

We visualize the prior/posterior bounds for these various definitions in Figure 2.

7.3 Comparing two posteriors

In [GKS08, KS14], the authors propose an approach that captures an intuitive idea proposed by Dwork in [Dwo06]: “any conclusions drawn from the output of a private algorithm must be similar whether or not an individual’s data is present in the input or not”. They define *semantic privacy* (SemPr): instead of comparing the posterior with the prior belief like in DP, this bounds the difference between two posterior belief distributions, depending on which dataset was secretly chosen. The distance chosen to represent the idea that those two posterior belief distributions are close is the statistical distance. One important difference between the definitions in the previous subsection is that semantic privacy quantifies over all possible priors: like in DP, the attacker is assumed to have arbitrary background knowledge.

²¹Another definition with the same name is introduced in [SDS⁺19], we mention it in Section 11.1.

²²Another definition with the same name is introduced in [dPCF12], we mention it later in this section.




Figure 2: From left to right, using $\varepsilon = \ln 3$: posterior-prior bounds in differential privacy, positive membership privacy, adversarial privacy (with $\delta = 0.05$) and aposteriori noiseless privacy.

Definition 37 (ε -semantic privacy [GKS08, KS14]). *A mechanism \mathcal{M} is ε -semantically private if for any distribution over datasets θ , any index i , any $S \subseteq \mathcal{O}$, and any set of datasets $X \subseteq \mathcal{D}$:*

$$|\mathbb{P}_{D \sim \theta} [D \in X \mid \mathcal{M}(D) \in S] - \mathbb{P}_{D \sim \theta} [D \in X \mid \mathcal{M}(D_{-i}) \in S]| \leq \varepsilon.$$

A couple of other definitions also compare posteriors directly: *inferential privacy* [GK17] is a reformulation of noiseless privacy, and *range-bounded privacy* [DR19] (RBPr) requires that two different values of the PLRV are close to each other (instead of being between centered around zero like in ε -DP). It is equivalent to ε -DP up to a change in parameters, and is used as a technical tool to prove composition results.

7.4 Multidimensional definitions

This dimension could be combined with other dimensions fairly easily; several DP modifications from this section does belong to multiple dimensions, but to introduce the concept we overlooked these details.

Definitions that limit the background knowledge of the adversary explicitly formulate it as a probability distribution. As such, they are natural candidates for Bayesian reformulations. In [WXH17], the authors introduce *identity DP*, which is an equivalent Bayesian reformulation of noiseless privacy.

Another example is *inference-based causal DP* [DMKB19], similar to aposteriori noiseless DP, except it uses causal DP instead of noiseless DP.

Definition 38 ((Θ, ε) -inference-based causal differential privacy [DMKB19]). *Given a family Θ of probability distribution on \mathcal{D} , a mechanism \mathcal{M} satisfies (Θ, ε) -inference-based distributional DP if (IBCDP) for all probability distributions $\theta \in \Theta$, for all i , all $a, b \in \mathcal{T}$ and all outputs O :*

$$D(i)_{|D \sim \theta, \mathcal{M}(D)=O} \approx_{\varepsilon} D(i)_{|D \sim \theta, \mathcal{M}(D_{i \rightarrow b})=O}.$$

where $D_{i \rightarrow b}$ is the dataset obtained by changing the i -th record of D into b .

The authors of semantic privacy also combined it with probabilistic DP, and simply called it (ε, δ) -semantic privacy. Further, it is possible to consider different definitions of neighborhood. In [dPCF12], authors introduce *information privacy*²³ (InfPr), which can be seen as a posteriori

²³Another definition with the same name is introduced in [WXH17], we mention it earlier in this section.

noiseless privacy combined with free lunch privacy: rather than only considering the knowledge gain of the adversary on one particular user, it considers its knowledge gain about any possible group of values of the dataset.

Definition 39 ((Θ, ε) -information privacy [dPCF12]). *A mechanism \mathcal{M} satisfies (Θ, ε) -information privacy if for all probability distributions $\theta \in \Theta$, all $D \in \mathcal{D}$ and all $O \in \mathcal{O}$, $D|_{D \sim \theta} \approx_\varepsilon D|_{D \sim \theta, \mathcal{M}(D)=O}$.*

The authors further prove that if Θ contains a distribution whose support is \mathcal{D} , then (Θ, ε) -InfPr implies 2ε -DP.

Apart from the hypothesis testing reformulations, that can be used to improve composition and amplification results, the definitions in this section mostly appear in theoretical research papers, to provide a deeper understanding of guarantees offered by DP and its alternatives. They do not seem to be used in practical applications.

8 Relativization of the knowledge gain (R)

A differentially private mechanism does not reveal more than a bounded amount of probabilistic information about a user. This view does not explicitly take into account other ways information can leak, like side-channel functions or knowledge about the structure of a social network. We found two approaches that attempt to include such auxiliary functions in DP variants. One possibility is to weaken DP by allowing or disregarding a certain amount of leakage; another option is to explicitly forbid the mechanism to reveal more than another function, considered to be safe for release.

8.1 Taking into account auxiliary leakage function

In [LPR20], authors define *bounded leakage DP*, which quantifies the privacy that is maintained by a mechanism despite bounded, additional leakage of information by some leakage function. Interestingly, this leakage function P shares the randomness of the privacy mechanism: it can, for example, capture side-channel leakage from the mechanism’s execution. In the formal definition of this DP variant, the randomness is *explicit*: the privacy mechanism and the leakage takes the random bits $r \in \{0, 1\}^*$ as an additional parameter.

Definition 40 ((P, ε, δ) -bounded leakage differential privacy [LPR20]). *Let $P : \mathcal{D} \times \{0, 1\}^*$ be a leakage function. A privacy mechanism \mathcal{M} is (P, ε, δ) -bounded leakage differentially private (BLDP) if for all pairs of neighboring datasets D_1 and D_2 , all outputs O_P of P such that $\mathbb{P}[P(D_1, r) = O_P] \neq 0$ and $\mathbb{P}[P(D_2, r) = O_P] \neq 0$, and all sets of outputs $S \subseteq \mathcal{O}$:*

$$\mathbb{P}[\mathcal{M}(D_1, r) \in S \mid P(D_1, r) = O_P] \leq e^\varepsilon \cdot \mathbb{P}[\mathcal{M}(D_2, r) \in S \mid P(D_2, r) = O_P] + \delta$$

where the randomness is taken over the random bits r .

As expected, if there is no leakage (P is a constant function), this is simply (ε, δ) -DP. The authors also show that it is closed for post-processing and composable. Furthermore, if the privacy mechanism is independent from the leakage function, it is strictly weaker than differential privacy.

8.2 Borrowing concepts from zero-knowledge proofs

When using the associative interpretation with the independence assumption, it is unclear how to adapt DP to correlated datasets like social networks: data about someone’s friends might reveal sensitive information about this person. The causal interpretation of DP does not suffer from this problem, but how to adapt the associative view to such correlated contexts? Changing the definition of the neighborhood is one possibility (see Section 4.1), but it requires knowing in advance the exact impact of someone on other records. A more robust option is to impose that the information released does not contain more information than the result of some predefined algorithms on the data, without the individual in question. The method for formalizing this intuition borrows ideas from *zero-knowledge proofs* [GMR89].

Instead of imposing that the result of the mechanism is roughly *the same* on neighboring datasets D_1 and D_2 , the intuition is to impose that the result of the mechanism on D_1 can be *simulated* using only some information about D_2 . The corresponding definition, called *zero-knowledge privacy* and introduced in [GLP11], captures the idea that the mechanism does not leak more information on a given target than a certain class of aggregate metrics. This class, called *model of aggregate information* in [GLP11], is formalized by a family of (possibly randomized) family of algorithms Agg .

Definition 41 ((Agg, ε)-zero-knowledge privacy [GLP11]). *Let Agg be a family of (possibly randomized) algorithms agg . A privacy mechanism \mathcal{M} is (Agg, ε)-zero-knowledge private (ZKPr) if there exists an algorithm $\text{agg} \in \text{Agg}$ and a simulator Sim such as for all datasets D and indices i , $\mathcal{M}(D) \approx_\varepsilon \text{Sim}(\text{agg}(D_{-i}))$.*

In [GLP11], authors show that (Agg, ε)-ZKPr implies 2ε -DP for any Agg , while ε -DP implies (Agg, ε)-ZKPr if the identity function is in Agg . This is yet another way to formalize the intuition that differential privacy protects against attackers who have full background knowledge.

Another approach is to ignore some leakage instead of limiting it. This is done in *subspace DP* [GGY21], which require DP to holds for a sub-space within the projection of the range of \mathcal{M} . Furthermore, *induced subspace DP* [GGY21] is also defined which ensures that the output of \mathcal{M} does meet the invariant-linear external constraints. Essentially, they say for a mechanism to satisfy subspace DP, the function $D \rightarrow f(\mathcal{M}(D))$ must be DP for some function f .

8.3 Multidimensional definitions

Using a simulator allows making statements of the type “this mechanism does not leak more information on a given target than a certain class of aggregate metrics”. Similarly to noiseless privacy, it is possible to explicitly limit the attacker’s background knowledge using a data-generating probability distribution, as well as vary the neighborhood definitions to protect other types of information than the presence and characteristics of individuals. This is done in [BGKS13] as *coupled-worlds privacy* (CWPr), a generalization of distributional DP, where a family of functions priv represents the protected attribute.

Definition 42 (($\Theta, \Gamma, \varepsilon$)-coupled-worlds privacy [BGKS13]). *Let Γ be a family of pairs of functions (agg, priv). A mechanism \mathcal{M} satisfies ($\Theta, \Gamma, \varepsilon$)-coupled-worlds privacy if there is a simulator Sim such that for all distributions $\theta \in \Theta$, all $(\text{agg}, \text{priv}) \in \Gamma$, and all possible values p :*

$$\mathcal{M}(D)_{|D \sim \theta, \text{priv}(D)=p} \approx_\varepsilon \text{Sim}(\text{agg}(D))_{|D \sim \theta, \text{priv}(D)=p}$$

A special case of coupled-worlds privacy is also introduced in [BGKS13] as *distributional DP*, already mentioned in Section 6: each function `priv` captures the value of a single record, and the corresponding function `agg` outputs all other records.

Coupled-worlds privacy is a good example of combining variants from different dimensions: it changes several aspects of the original definition according to from \mathbf{N} , \mathbf{B} and \mathbf{R} . Moreover, \mathbf{Q} and \mathbf{F} can easily be integrated with this definition by using (ε, δ) -indistinguishability with a Bayesian reformulation. This is done explicitly in *inference-based coupled-worlds privacy* [BGKS13]; the same paper also introduces inference-based distributional differential privacy (IBDDP).

Definition 43 ($(\Theta, \Gamma, \varepsilon, \delta)$ -inference-based coupled-worlds privacy [BGKS13]). *Given a family Θ of probability distributions on $\mathcal{D} \times \mathcal{B}$, and a family Γ of pairs of functions $(\text{agg}, \text{priv})$, a mechanism \mathcal{M} satisfies $(\Theta, \Gamma, \varepsilon, \delta)$ -inference-based coupled-worlds privacy (IBCWP) if there is a simulator Sim such that for all distributions $\theta \in \Theta$, and all $(\text{agg}, \text{priv}) \in \Gamma$:*

$$\text{priv}(D)_{|(D,B) \sim \theta, \mathcal{M}(D)=O, B=\hat{B}} \approx_{(\varepsilon, \delta)} \text{priv}(D)_{|(D,B) \sim \theta, \text{Sim}(\text{agg}(D))=O, B=\hat{B}}$$

with probability at least $1 - \delta$ over the choice of O and \hat{B} .

Random DP was also combined with an idea similar to ZKPr: in [BF16], the authors introduce *typical stability* (TypSt), which combines random DP with approximate DP, except that rather using (ε, δ) -indistinguishability between two outputs of the mechanism, it uses a simulator that only knows data-generating distribution.

Definition 44 ($(\Theta, \gamma, \varepsilon, \delta)$ -typical stability [BF16]). *Given a family Θ of probability distributions on \mathcal{D} , a mechanism \mathcal{M} satisfies $(\Theta, \gamma, \varepsilon, \delta)$ -typical stability (TypSt) if for all distributions $\theta \in \Theta$, there is a simulator Sim such that with probability at least $1 - \gamma$ over the choice of $D \sim \theta$, $\mathcal{M}(D) \approx_{\varepsilon, \delta} \text{Sim}(\theta)$.*

In the same paper, the authors introduce a variant of the same definition with the same name, which compares two outputs of the mechanism; this is essentially a combination between $\text{DIPr}^{\text{[Rot10, BLR13]}}$ and approximate DP.

We did not find any evidence that the variants and extensions of this section are used outside of theoretical papers exploring the guarantees they provide.

9 Computational power (C)

The ε -indistinguishability property in DP is *information-theoretic*: the attacker is implicitly assumed to have infinite computing power. This is unrealistic in practice, so it is natural to consider definitions where the attacker only has polynomial computing power. Changing this assumption leads to weaker data privacy definitions. In [MPRV09], two approaches have been proposed to formalize this idea: either modeling the distinguisher explicitly as a polynomial Turing machine, either allow a mechanism not to be technically differentially private, as long as one cannot distinguish it from a truly differentially private one.

Using a polynomial distinguisher on the output

The attacker is not explicit in the formalization of DP based on ε -indistinguishability. It is possible change the definition to make this attacker explicit: model it as a *distinguisher*, who tries to guess whether a given output O comes from a dataset D_1 or its neighbor D_2 . In doing so, it becomes

straightforward to require this attacker to be computationally bounded: simply model it as a probabilistic polynomial-time Turing machine. In [MPRV09], the authors introduce IndCDP, short for *Indistinguishability-based Computational DP*.

Definition 45 (ε_κ -IndCDP [MPRV09]). *A family $(\mathcal{M}_\kappa)_{\kappa \in \mathbb{N}}$ of privacy mechanisms \mathcal{M}_κ provides ε_κ -IndCDP if there exists a negligible function neg such that for all non-uniform probabilistic polynomial-time Turing machines Ω (the distinguisher), all polynomials $p(\cdot)$, all sufficiently large $\kappa \in \mathbb{N}$, and all datasets $D_1, D_2 \in \mathcal{D}$ of size at most $p(\kappa)$ that differ only one record:*

$$\mathbb{P}[\Omega(\mathcal{M}(D_1)) = 1] \leq e^{\varepsilon_\kappa} \cdot \mathbb{P}[\Omega(\mathcal{M}(D_2)) = 1] + \mathit{neg}(\kappa)$$

where neg is a function that converges to zero asymptotically faster than the reciprocal of any polynomial.

This definition can also be expressed using the authors define *differential indistinguishability*, a notion defined in [BKMR14] that adapts ε -indistinguishability to a polynomially bounded attacker.

Using a polynomial distinguisher on the mechanism

Another natural option is to require that the mechanism “looks like” a truly differentially private mechanism, at least to a computationally bounded distinguisher. In [MPRV09], the authors introduce SimCDP, short for *Simulation-based Computational DP*.

Definition 46 (ε_κ -SimCDP [MPRV09]). *A family $(\mathcal{M}_\kappa)_{\kappa \in \mathbb{N}}$ of privacy mechanisms \mathcal{M}_κ provides ε_κ -SimCDP if there exists a family $(\mathcal{M}'_\kappa)_{\kappa \in \mathbb{N}}$ of ε_κ -DP and a negligible function neg such that for all non-uniform probabilistic polynomial-time Turing machines Ω , all polynomials $p(\cdot)$, all sufficiently large $\kappa \in \mathbb{N}$, and all datasets $D \in \mathcal{D}$ of size at most $p(\kappa)$:*

$$\mathbb{P}[\Omega(\mathcal{M}(D)) = 1] - \mathbb{P}[\Omega(\mathcal{M}'(D)) = 1] \leq \mathit{neg}(\kappa)$$

where neg is a function that converges to zero asymptotically faster than the reciprocal of any polynomial.

In [MPRV09], the authors show that ε_κ -SimCDP implies ε_κ -IndCDP.

9.1 Multidimensional definitions

IndCDP has been adapted to different settings, and extended to arbitrary neighborhood relationships. In *output constrained DP* (OCDP), introduced in [HMFS17], the setting is a two-party computation, each party can have its own set of privacy parameters (ε_A , ε_B , δ_A , and δ_B —the δ parameters correspond to the $\mathit{neg}(\kappa)$ term in IndCDP), and the neighborhood relationship is determined by a function f . The authors also propose *DP for Record Linkage* (DPRL), an instance of OCDP that uses for a specific function f that captures the need to protect non-matching records during the execution of a private record linkage protocol.

Some DP variants which explicitly model an adversary with a simulator can relatively easily be adapted to model a computationally bounded adversary, simply by imposing that the simulator must be polynomial. This is done explicitly in [GLP11], where the authors define *computational zero-knowledge privacy* (CZKPr), which could also be adapted to e.g., the two coupled-worlds privacy definitions as well.

Further, although we have not seen this done in practice in existing literature, the idea behind SimCDP can in principle be adapted to any other definition: rather than requiring that a given definition holds in an information-theoretic fashion, it should be possible to require that the mechanism “looks like” a mechanism which genuinely satisfies the definition.

Limiting the computational power of the attacker is a reasonable assumption, but for a large class of queries, it cannot provide significant benefits over classical DP in the typical client-server setting [GKY11]. Thus, existing work using it focuses on multi-party settings [BHE⁺18].

10 Summarizing table

In this section we summarize the results from the previous 7 sections into a table where we list the known relations and show the properties with either referring to the original proof or creating a novel one.

In Table 3 we list the differential privacy variants and extensions introduced in this work. For each, we specify their name, parameters and where they were introduced (column 1), which dimensions they belong to (column 2), which axioms they satisfy (column 3, post-processing on the left and convexity on the right), whether they are composable (column 4) and how they relate to other differential privacy notions (column 5). We do not list definitions whose only difference is that they apply DP to other types of input, like those from Section 4.3, or geolocation-specific definitions.

Name & References	Dimension	Axioms		Cp. 11	Relations
		P.P. 9	Cv. 10		
(ε, δ) -approximate DP [DKM ⁺ 06]	\mathcal{Q}	\checkmark^1	\checkmark^1	\checkmark^{13}	(ε, δ) - DP $\supset^{\prec} \varepsilon$ -DP
(ε, δ) -probabilistic DP [Mei18]	\mathcal{Q}	$\boldsymbol{\chi}^{13}$	$\boldsymbol{\chi}^2$	\checkmark^{13}	(ε, δ) -DP $\sim (\varepsilon, \delta)$ - ProDP
$(\varepsilon, \delta_a, \delta_p)$ -Relaxed DP [ZQZ ⁺ 15]	\mathcal{Q}	$\boldsymbol{\chi}^{13}$	$\boldsymbol{\chi}^2$	\checkmark^{13}	(ε, δ_p) -ProDP $\subset^{\succ} (\varepsilon, \delta_a, \delta_p)$ - RelDP (ε, δ_a) -DP $\subset^{\succ} (\varepsilon, \delta_a, \delta_p)$ - RelDP
ε -Kullback-Leiber Pr [CY16]	\mathcal{Q}	\checkmark^1	\checkmark^1	\checkmark^{13}	(ε, δ) -DP $\prec \varepsilon$ - KLPPr $\prec \varepsilon$ -DP
(α, ε) -Rényi DP [Mir17]	\mathcal{Q}	\checkmark^1	\checkmark^1	\checkmark^{13}	ε -KLPPr $\subset^{\prec} (\alpha, \varepsilon)$ - RényiDP $\supset^{\prec} \varepsilon$ -DP
(μ, τ) -mean concentrated DP [DR16]	\mathcal{Q}	$\boldsymbol{\chi}^{13}$?	\checkmark^{13}	(ε, δ) -DP $\prec (\mu, \tau)$ - mCoDP $\prec \varepsilon$ -DP
(ξ, ρ) -zero concentrated DP [BS16]	\mathcal{Q}	\checkmark^1	\checkmark^1	\checkmark^{13}	(ξ, ρ) - zCoDP $\sim (\mu, \tau)$ -mCoDP
(ξ, ρ, δ) -approximate CoDP [BS16]	\mathcal{Q}	$\boldsymbol{\chi}^1$?	\checkmark^{13}	(ε, δ) -DP $\succ (\xi, \rho, \delta)$ - ACoDP $\supset^{\prec} (\xi, \rho)$ -zCoDP
(ξ, ρ, ω) -bounded CoDP [BS16]	\mathcal{Q}	\checkmark^1	\checkmark^1	\checkmark^{13}	(ξ, ρ, ω) - BCoDP $\supset^{\prec} (\xi, \rho)$ -zCoDP
(η, τ) -truncated CoDP [Col16]	\mathcal{Q}	\checkmark^1	\checkmark^1	\checkmark^{13}	(η, τ) - TCoDP ^[Col16] $\sim \varepsilon$ -DP
(ρ, ω) -truncated CoDP [BDRS18]	\mathcal{Q}	\checkmark^1	\checkmark^1	\checkmark^{13}	(ρ, ω) - TCoDP ^[BDRS18] $\subset^{\prec} (\xi, \rho, \omega)$ -bCoDP
(f, ε) -divergence DP [BD14]	\mathcal{Q}	\checkmark^1	\checkmark^1	?	(f, ε) - DivDP \supset most def. in \mathcal{Q}
(f_k, ε) -divergence DP [DR18]	\mathcal{Q}	\checkmark^1	\checkmark^1	?	(f_k, ε) - kDivDP $\subset (f, \varepsilon)$ -DivDP

Name & References	Dimension	Axioms		Cp. 11	Relations
		P.P. 9	Cv. 10		
(H, f, ε) -capacity bounded DP [CIM19]	\mathbb{Q}	\checkmark^1	\checkmark^1	?	(H, f, ε) -CBDP $\subset^{\prec} (f, \varepsilon)$ -DivDP
ε -unbounded DP [KM11]	\mathbb{N}	\checkmark^{13}	\checkmark^{13}	\checkmark^7	ε -DP $\sim \varepsilon$ -uBoDP $\subset^{\sim} (c, \varepsilon)$ -GrDP
ε -bounded DP [KM11]	\mathbb{N}	\checkmark^{13}	\checkmark^{13}	\checkmark^7	ε -BoDP $\prec \varepsilon$ -DP
(P, ε) -one-sided DP [KDH ⁺ 20]	\mathbb{N}	\checkmark^{13}	\checkmark^{13}	\checkmark^7	(P, ε) -OnSDP $\supset^{\prec} \varepsilon$ -BoDP
(P, ε) -asymmetric DP [TCY21]	\mathbb{N}	\checkmark^{13}	\checkmark^{13}	\checkmark^7	(P, ε) -AsDP $\supset^{\prec} \varepsilon$ -uBoDP
ε -client DP [MRTZ18]	\mathbb{N}	\checkmark^{13}	\checkmark^{13}	\checkmark^7	ε -ClDP $\succ \varepsilon$ -DP
ε -element DP [ADJ19]	\mathbb{N}	\checkmark^{13}	\checkmark^{13}	\checkmark^7	ε -ElDP $\succ \varepsilon$ -EIDP $\succ \varepsilon$ -DP
(c, ε) -group DP [Dwo08]	\mathbb{N}	\checkmark^{13}	\checkmark^{13}	\checkmark^7	(c, ε) -GrDP $\supset^{\sim} \varepsilon$ -DP
(R, c, ε) -dependent DP [LCM16]	\mathbb{N}	\checkmark^{13}	\checkmark^{13}	\checkmark^7	(R, c, ε) -DepDP $\supset (c, \varepsilon)$ -GrDP
$(\mathcal{A}, \varepsilon)$ -bayesian DP [YSN15]	\mathbb{N}	\checkmark^{13}	\checkmark^{13}	\checkmark^7	$(\mathcal{A}, \varepsilon)$ -BayDP ^[YSN15] $\supset (R, c, \varepsilon)$ -DepDP
$(\mathcal{A}, \varepsilon)$ -correlated DP [WDN17, WWK ⁺ 17]	\mathbb{N}	\checkmark^{13}	\checkmark^{13}	\checkmark^7	$(\mathcal{A}, \varepsilon)$ -CorDP $\subset (\mathcal{A}, \varepsilon)$ -BayDP ^[YSN15]
$(\mathcal{A}, \varepsilon)$ -prior DP [LRY19]	\mathbb{N}	\checkmark^{13}	\checkmark^{13}	\checkmark^7	$(\mathcal{A}, \varepsilon)$ -PriDP $\supset (\mathcal{A}, \varepsilon)$ -BayDP ^[YSN15]
ε -free lunch Pr [KM11]	\mathbb{N}	\checkmark^{13}	\checkmark^{13}	\checkmark^7	ε -FLPr \succ all def. in \mathbb{N}
(D, ε) -individual DP [SCDFSM17]	\mathbb{N}	\checkmark^{13}	\checkmark^{13}	\checkmark^7	(D, ε) -IndDP $\prec \varepsilon$ -DP
(D, t, ε) -per-instance DP [RW21]	\mathbb{N}	\checkmark^{13}	\checkmark^{13}	\checkmark^7	(D, t, ε) -PIDP $\prec (D, \varepsilon)$ -IndDP
$(\mathcal{R}, \varepsilon)$ -generic DP [KM11, FC14]	\mathbb{N}	\checkmark^{13}	\checkmark^{13}	\checkmark^7	$(\mathcal{R}, \varepsilon)$ -GcDP ^[KM11] \supset most def. in \mathbb{N}
(d, Δ, ε) -constrained DP [ZLW09]	\mathbb{N}	\checkmark^{13}	\checkmark^{13}	\checkmark^7	(d, Δ, ε) -ConsDP $\sim (\mathcal{R}, \varepsilon)$ -GcDP
$(d, \Delta, S_{\mathcal{D}}, \varepsilon)$ -distributional Pr [ZLW09]	\mathbb{N}	\checkmark^{13}	\checkmark^{13}	\checkmark^7	$(d, \Delta, S_{\mathcal{D}}, \varepsilon)$ -DiPr ^[ZLW09] $\subset (\mathcal{R}, \varepsilon)$ -GcDP
(f, ε) -sensitivity-induced DP [RA17]	\mathbb{N}	\checkmark^{13}	\checkmark^{13}	\checkmark^7	ε -SIDP $\subset (\mathcal{R}, \varepsilon)$ -GcDP
$(\mathcal{I}_Q, \varepsilon)$ -induced-neighbors DP [KM11]	\mathbb{N}	\checkmark^{13}	\checkmark^{13}	\checkmark^7	$(\mathcal{I}_Q, \varepsilon)$ -INDP $\subset (\mathcal{R}, \varepsilon)$ -GcDP
$(G, \mathcal{I}_Q, \varepsilon)$ -blowfish Pr [HMD14]	\mathbb{N}	\checkmark^{13}	\checkmark^{13}	\checkmark^7	$(G, \mathcal{I}_Q, \varepsilon)$ -BFPr $\subset (\mathcal{R}, \varepsilon)$ -GcDP
Ψ -personalized DP [JYC15]	\mathbb{V}	\checkmark^3	\checkmark^3	\checkmark^7	Ψ -PerDP $\supset \varepsilon$ -DP
Ψ -tailored DP [LP15]	\mathbb{V}	\checkmark^3	\checkmark^3	\checkmark^7	Ψ -TaiDP $\supset \Psi$ -PerDP
$(\Psi, \bar{\varepsilon})$ -input-discriminative DP [GLXC20]	\mathbb{V}	\checkmark^3	\checkmark^3	\checkmark^7	Ψ -IDDP $\supset \Psi$ -PerDP
$\varepsilon(\cdot)$ -outlier Pr [LP15]	\mathbb{V}	\checkmark^3	\checkmark^3	\checkmark^7	$\varepsilon(\cdot)$ -OutPr $\subset \Psi$ -TaiDP
(ε, p, r) -Pareto DP [LP15]	\mathbb{V}	\checkmark^3	\checkmark^3	\checkmark^7	(ε, p, r) -ParDP $\subset \Psi$ -TaiDP
$(\pi, \gamma, \varepsilon)$ -random DP [HWR13]	\mathbb{V}	?	χ^4	\checkmark^{13}	$(\pi, \gamma, \varepsilon)$ -RandDP $\supset^{\prec} \varepsilon$ -DP
$d_{\mathcal{D}}$ -Pr [CABP13]	\mathbb{N}, \mathbb{V}	\checkmark^3	\checkmark^3	\checkmark^7	ε -DP $\subset d_{\mathcal{D}}$ -Pr
$(d_{\mathcal{D}}, \varepsilon)$ -smooth DP [BD14]	\mathbb{N}, \mathbb{V}	\checkmark^3	\checkmark^3	\checkmark^7	$(d_{\mathcal{D}}, \varepsilon)$ -SmDP $\sim d_{\mathcal{D}}$ -Pr

Name & References	Dimension	Axioms		Cp. 11	Relations
		P.P. 9	Cv. 10		
(ε, γ) -distributional Pr [ZLW09, Rot10]	N, V	?	?	?	ε -FLPr \subset (ε, γ) - DIPr ^[Rot10, BLR13]
$(\pi, \varepsilon, \delta)$ -weak Bayesian DP [TF20]	Q, V	\checkmark^{13}	?	\checkmark^7	(ε, δ) -DP \succ $(\pi, \varepsilon, \delta)$ - WBDP \prec $(\pi, \gamma, \varepsilon)$ -RanDP
(Θ, ε) -on average KL Pr [WLF16]	Q, V	\checkmark^{13}	?	? ¹⁴	ε -KLPr \succ (Θ, ε) - avgKLPr \prec $(\pi, \gamma, \varepsilon)$ -RanDP
$(\pi, \varepsilon, \delta)$ -Bayesian DP [TF20]	Q, V	\checkmark^{13}	χ^4	\checkmark^7	(ε, δ) -ProDP \succ $(\pi, \varepsilon, \delta)$ - BayDP ^[TF20] $(\pi, \gamma, \varepsilon)$ -RanDP \succ $(\pi, \varepsilon, \delta)$ - BayDP ^[TF20]
$(d_{\mathcal{D}}, \varepsilon, \delta)$ -pseudo-metric DP [DNM ⁺ 13]	Q, N, V	?	?	\checkmark^{13}	(ε, δ) -DP \subset $(d_{\mathcal{D}}, \varepsilon, \delta)$ - PsDP \supset^{\prec} $d_{\mathcal{D}}$ -Pr
(f, d, ε) -extended divergence DP [KM19a]	Q, N, V	\checkmark^3	\checkmark^3	?	$d_{\mathcal{D}}$ -Pr \subset (f, d, ε) - EDivDP \supset (f, ε) -Div DP
(\mathcal{R}, M) -generic DP [KL10]	Q, N, V	\checkmark^{13}	\checkmark^{13}	?	(\mathcal{R}, M) - GcDP ^[KL10] \supset (ε, δ) -DP
(\mathcal{R}, q) -abstract DP [KL10]	Q, N, V	χ^{13}	χ^{13}	?	(\mathcal{R}, q) - AbsDP \supset (\mathcal{R}, M) -GcDP ^[KL10]
(Θ, ε) -noiseless Pr [BBG ⁺ 11]	B	\checkmark^{13}	\checkmark^{13}	χ^8	ε -DP \supset^{\prec} (Θ, ε) - NPr
(Θ, ε) -causal DP [DMKB19]	B	\checkmark^{13}	\checkmark^{13}	χ^8	(Θ, ε) - CausDP \supset^{\prec} ε -DP
(β, ε) -DP under sampling [LQS12]	Q, B	\checkmark^{13}	\checkmark^{13}	χ^8	(Θ, ε) -NPr \supset (β, ε) - SamDP \supset^{\prec} ε -DP
$(\Theta, \varepsilon, \delta)$ -active PK DP [DMKB19]	Q, B	\checkmark^{13}	\checkmark^{13}	χ^8	$(\Theta, \varepsilon, \delta)$ - APKDP \supset^{\prec} (Θ, ε) -NPr
$(\Theta, \varepsilon, \delta)$ -passive PK DP [DMKB19]	Q, B	\checkmark^{13}	\checkmark^{13}	χ^8	$(\Theta, \varepsilon, \delta)$ -APKDP \succ $(\Theta, \varepsilon, \delta)$ - PPKDP (Θ, ε) -NPr \subset^{\succ} $(\Theta, \varepsilon, \delta)$ - PPKDP
$(\Theta, \Phi, \varepsilon)$ -pufferfish Pr [KM12]	N, B	\checkmark^{13}	\checkmark^{13}	χ^8	(Θ, ε) -NPr \subset $(\Theta, \Phi, \varepsilon)$ - PFPr \supset $(\mathcal{R}, \varepsilon)$ -GcDP
(π, k, ε) -Bayesian DP [LL12]	N, B	\checkmark^{13}	\checkmark^{13}	χ^8	BayDP ^[LL12] \subset $(\Theta, \Phi, \varepsilon)$ -PFPr
$(\Theta, \varepsilon, \delta)$ -distribution Pr [JB14, KM19b]	Q, N, B	\checkmark^3	\checkmark^3	χ^8	$(\Theta, \varepsilon, \delta)$ - DnPr \supset $(\Theta, \varepsilon, \delta)$ -APKDP
$(\Theta, \varepsilon, \delta)$ -profile-based DP [GC19]	Q, N, B	\checkmark^3	\checkmark^3	χ^8	$(\Theta, \varepsilon, \delta)$ - PBDP \subset $(\Theta, \varepsilon, \delta)$ -DnPr
$(\Theta, \varepsilon, \delta)$ -probabilistic DnPr [KM19b]	Q, N, B	χ^{13}	χ^2	χ^8	ε -ProDP \subset $(\Theta, \varepsilon, \delta)$ - PDnPr \supset (Θ, ε) -DnPr
(f, Θ, ε) -divergence DnPr [KM19a]	Q, N, B	\checkmark^3	\checkmark^3	χ^8	(f, ε) -DP \subset (f, Θ, ε) - DDnPr \supset (Θ, ε) -DnPr
(d, Θ, ε) -extended DnPr [KM19b]	N, V, B	\checkmark^3	\checkmark^3	χ^8	$d_{\mathcal{D}}$ -Pr \subset (d, Θ, ε) - EDnPr \supset (Θ, ε) -DnPr
$(d, f, \Theta, \varepsilon)$ -ext. div. DnPr [KM19a]	Q, N, V, B	\checkmark^3	\checkmark^3	χ^8	(f, Θ, ε) -DDPr \subset $(d, f, \Theta, \varepsilon)$ - EDDnPr (d, Θ, ε) -EDnPr \subset $(d, f, \Theta, \varepsilon)$ - EDDnPr
ε -indistinguishable Pr [LXL13]	F	\checkmark^{12}	\checkmark^{12}	\checkmark^{12}	ε - IndPr \sim ε -DP
ε -semantic Pr [GKS08, KS14]	F	?	?	?	ε - SemPr \sim ε -DP
ε -range-bounded Pr [DR19]	F	?	?	?	ε - RBPr \sim ε -DP
f -DP [DRS21]	Q, F	\checkmark^{13}	?	\checkmark^{13}	f - DP \supset (ε, δ) -DP
Gaussian DP [DRS21]	Q, F	\checkmark^{13}	?	\checkmark^{13}	GaussDP \subset f -DP
weak-Federated f -DP [ZCLS21]	Q, F	\checkmark^{13}	?	\checkmark^{13}	f - wFDP \subset GaussDP

Name & References	Dimension	Axioms		Cp. 11	Relations
		P.P. 9	Cv. 10		
strong-Federated f -DP [ZCLS21]	\mathcal{Q}, \mathcal{F}	\checkmark^{13}	?	\checkmark^{13}	$f\text{-wFDP} \prec f\text{-sFDP} \subset \text{GaussDP}$
(Θ, ε) -aposteriori noiseless Pr [BBG ⁺ 11]	\mathcal{B}, \mathcal{F}	\checkmark^5	\checkmark^5	?	$(\Theta, \varepsilon)\text{-ANPr} \sim (\Theta, \varepsilon)\text{-NPr}$
(Θ, ε) -positive membership Pr [LQS ⁺ 13]	\mathcal{B}, \mathcal{F}	\checkmark^5	\checkmark^5	\times^8	$(\Theta, \varepsilon)\text{-PMPPr} \supset \varepsilon\text{-BoDP}$
(Θ, ε) -negative membership Pr [LQS ⁺ 13]	\mathcal{B}, \mathcal{F}	\checkmark^5	\checkmark^5	\times^8	$(\Theta, \varepsilon)\text{-NMPPr} \supset \varepsilon\text{-BoDP}$
(Θ, ε) -membership Pr [LQS ⁺ 13]	\mathcal{B}, \mathcal{F}	\checkmark^5	\checkmark^5	\times^8	$(\Theta, \varepsilon)\text{-PMPPr} \prec (\Theta, \varepsilon)\text{-MPPr} \succ (\Theta, \varepsilon)\text{-NMPPr}$
$(\Theta, \varepsilon, \delta)$ -adversarial Pr [RHMS09]	$\mathcal{Q}, \mathcal{B}, \mathcal{F}$	\checkmark^5	\checkmark^5	\times^8	$(\varepsilon, \delta)\text{-DP} \subset (\Theta, \varepsilon, \delta)\text{-AdvPr} \prec (\Theta, \varepsilon)\text{-PMPPr}$
(Θ, ε) -information Pr [dPCF12]	$\mathcal{N}, \mathcal{B}, \mathcal{F}$?	?	?	$(\Theta, \varepsilon)\text{-InfPr} \succ \varepsilon\text{-DP}$
$(\text{Agg}, \varepsilon)$ -zero-knowledge Pr [GLP11]	\mathcal{R}	\checkmark^{13}	\checkmark^{13}	? ¹⁴	$(\text{Agg}, \varepsilon)\text{-ZKPr} \succ \varepsilon\text{-DP}$
$(\mathcal{P}_\nu, \varepsilon, \delta)$ -subspace DP [GGY21]	\mathcal{Q}, \mathcal{R}	? ¹⁴	?	\checkmark^{13}	$(\mathcal{P}_\nu, \varepsilon, \delta)\text{-SubDP} \prec (\varepsilon, \delta)\text{-DP}$
(P, ε) -bounded leakage DP [LPR20]	\mathcal{Q}, \mathcal{R}	\checkmark^{13}	\checkmark^{13}	\checkmark^{13}	$(P, \varepsilon)\text{-BLDP} \supset (\varepsilon, \delta)\text{-DP}$
$(\Theta, \Gamma, \varepsilon)$ -coupled-worlds Pr [BGKS13]	$\mathcal{N}, \mathcal{B}, \mathcal{R}$	\checkmark^{13}	\checkmark^{13}	\times^3	$(\Theta, \Gamma, \varepsilon)\text{-CWPPr} \supset \varepsilon\text{-DP}$
$(\Theta, \Gamma, \varepsilon, \delta)$ -inference CW Pr [BGKS13]	$\mathcal{Q}, \mathcal{N}, \mathcal{B}, \mathcal{F}, \mathcal{R}$?	?	\times^3	$(\Theta, \Gamma, \varepsilon, \delta)\text{-IBCWPPr} \succ (\Theta, \Gamma, \varepsilon)\text{-CWPPr}$
$(\Theta, \Gamma, \varepsilon, \delta)$ -inference DistDP [BGKS13]	$\mathcal{Q}, \mathcal{N}, \mathcal{B}, \mathcal{F}, \mathcal{R}$?	?	\times^3	$(\Theta, \varepsilon)\text{-DDP} \succ (\Theta, \Gamma, \varepsilon, \delta)\text{-IBDDP} \subset (\Theta, \Gamma, \varepsilon, \delta)\text{-IBCWPPr}$
Typical Stability [BF16]	$\mathcal{Q}, \mathcal{V}, \mathcal{R}$?	?	\times^3	$(\Theta, \gamma, \varepsilon, \delta)\text{-TySt}$
ε_κ -SIM-computational DP [MPRV09]	\mathcal{C}	\checkmark^6	\checkmark^6	\checkmark^{15}	$\varepsilon_\kappa\text{-SimCDP} \prec \varepsilon\text{-DP}$
ε_κ -IND-computational DP [MPRV09]	\mathcal{C}	\checkmark^6	\checkmark^6	\checkmark^{15}	$\varepsilon_\kappa\text{-IndCDP} \prec \varepsilon_\kappa\text{-SimCDP}$
$(\text{Agg}, \varepsilon)$ -computational ZK Pr [GLP11]	\mathcal{R}, \mathcal{C}	\checkmark^6	\checkmark^6	?	$(\text{Agg}, \varepsilon)\text{-CZKPr} \supset^{\prec} (\text{Agg}, \varepsilon)\text{-ZKPr}$
(ε, δ, f) -output constrained DP [HMFS17]	$\mathcal{N}, \mathcal{V}, \mathcal{C}$	\checkmark^6	\checkmark^6	\checkmark^{15}	$(\varepsilon, \delta, f)\text{-OCDP} \supset \varepsilon_\kappa\text{-IndCDP}$
(ε, δ) -DP for Record Linkage [HMFS17]	$\mathcal{N}, \mathcal{V}, \mathcal{C}$	\checkmark^6	\checkmark^6	\checkmark^{15}	$(\varepsilon, \delta)\text{-RLDP} \subset (\varepsilon, \delta, f)\text{-OCDP}$

Table 3: Summary of variants/extensions of DP representing the main options in each combination of dimensions.

1. See Proposition 1.	9. Post-processing	appears in the paper introducing the definition.	definition
2. See Proposition 2.	10. Convexity		• V : Variation of privacy loss
3. See Proposition 3.	11. Composition	15. This claim appears in [Mir17], its proof is in the unpublished full version.	• B : Background knowledge
4. See Proposition 4.	12. Follows from its equivalence with ε -DP.	16. Abbreviations used for dimensions:	• F : Formalism of privacy loss
5. See Proposition 5.	13. A proof appears in the paper introducing the definition.	• Q : Quantification of privacy loss	• R : Relativization of knowledge gain
6. See Proposition 6.	14. A proof for a special case	• N : Neighborhood	• C : Computational power
7. See Proposition 7.			
8. See Proposition 8.			

10.1 Proofs of properties

10.1.1 Axioms

Proposition 1. *All instantiations of DivDP satisfy both privacy axioms. In particular, approximate DP, MIDP, KLPr, RenDP, and zCoDP satisfy both axioms.*

Proof. The post-processing axiom follows directly from the monotonicity property of the f -divergence. The convexity axiom follows directly from the joint convexity property of the f -divergence. \square

Proposition 2. *ProDP and ACoDP do not satisfy the convexity axiom.*

Proof. Consider the following mechanisms \mathcal{M}_1 and \mathcal{M}_2 , with input and output in $\{0, 1\}$.

- $\mathcal{M}_1(0) = 0$, $\mathcal{M}_1(1) = 1$ with probability δ , and $\mathcal{M}_1(1) = 0$ with probability $1 - \delta$.
- $\mathcal{M}_2(0) = \mathcal{M}_2(1) = 1$.

Both mechanisms are $\left(\frac{1}{1-\delta}, \delta\right)$ -ProDP. Now, consider the mechanism \mathcal{M} which applies \mathcal{M}_1 with probability $1 - 2\delta$ and \mathcal{M}_2 with probability 2δ . \mathcal{M} is a convex combination of \mathcal{M}_1 and \mathcal{M}_2 , but the reader can verify that it is not $\left(\frac{1}{1-\delta}, \delta\right)$ -ProDP. The result for (ξ, ρ, δ) -ACoDP is a direct corollary, since it is equivalent to (ξ, δ) -ProDP when $\rho = 0$. \square

Proposition 3. *$d_{\mathcal{D}}$ -Pr satisfies both privacy axioms. Further, EDivDP also satisfies both privacy axioms.*

Proof. The proof corresponding to PFPr in [KM12] is a proof by case analysis on every possible protected property. The fact that ε is the same for every protected property has no influence on the proof, so we can directly adapt the proof to $d_{\mathcal{D}}$ -Pr, and its combination with PFPr. Similarly, the proof can be extended to arbitrary divergence functions, like in Proposition 1. \square

Proposition 4. *RanDP does not satisfy the convexity axiom.*

Proof. Let π be the uniform distribution on $\{0, 1\}$, let D_1 be generated by picking 10 records according to π , and D_2 by flipping one record at random. Let \mathcal{M}_0 return 0 if all records are 0, and \perp otherwise. Let \mathcal{M}_1 return 1 if all records are 1, and \perp otherwise.

Note that both mechanisms are $(\pi, 2^{-9}, 0)$ -RanDP. Indeed, \mathcal{M}_0 will only return 0 for D_1 with probability 2^{-10} , and for D_2 with probability 2^{-10} (if D_1 only has one 1, which happens with probability $10 \cdot 2^{-10}$, and this record is flipped, which happens with probability 0.1). In both cases, \mathcal{M}_0 will return \perp for the other database; which will be a distinguishing event. Otherwise, \mathcal{M}_0 will return τ for both databases, so $\mathcal{M}(D_1) \approx_0 \mathcal{M}(D_2)$. The reasoning is the same for \mathcal{M}_1 .

However, the mechanism $\mathcal{M}_{0.5}$ obtained by applying either \mathcal{M}_0 or \mathcal{M}_1 uniformly randomly doesn't satisfy $(\pi, 2^{-9}, 0)$ -RanDP: the indistinguishability property does not hold if D_1 or D_2 have all their records set to *either* 0 or 1, which happens twice as often as either option alone. \square

Proposition 5. *All variants of MPr, AdvPr, and ANPr satisfy both axioms. As a direct corollary, InfPr also satisfies both axioms.*

Proof. We prove it for AdvPr. A mechanism \mathcal{M} satisfies $(\Theta, \varepsilon, \delta)$ -AdvPr if for all $t \in \mathcal{T}$, $\theta \in \Theta$, and $S \subseteq \mathcal{O}$, $\mathbb{P}_{D \sim \theta}[t \in D \mid \mathcal{M}(D) \in S] \leq e^\varepsilon \cdot \mathbb{P}_{D \sim \theta}[t \in D] + \delta$. We first prove that it satisfies the convexity axiom. Suppose \mathcal{M} is a convex combination of \mathcal{M}_1 and \mathcal{M}_2 . Simplifying $\mathbb{P}_{D \sim \theta}[\dots]$ into $\mathbb{P}[\dots]$, we have:

$$\begin{aligned} \mathbb{P}[t \in D \mid \mathcal{M}(D) \in S] &= \frac{\mathbb{P}[t \in D \text{ and } \mathcal{M}(D) \in S \text{ and } \mathcal{M} = \mathcal{M}_1]}{\mathbb{P}[\mathcal{M}(D) \in S]} \\ &\quad + \frac{\mathbb{P}[t \in D \text{ and } \mathcal{M}(D) \in S \text{ and } \mathcal{M} = \mathcal{M}_2]}{\mathbb{P}[\mathcal{M}(D) \in S]} \end{aligned}$$

Denoting $X_i = \mathbb{P}[\mathcal{M}(D) \in S \text{ and } \mathcal{M} = \mathcal{M}_i]$ for $i \in \{1, 2\}$, this gives:

$$\begin{aligned} \mathbb{P}[t \in D \mid \mathcal{M}(D) \in S] &= \\ &\frac{X_1 \cdot \mathbb{P}[t \in D \mid \mathcal{M}_1(D) \in S]}{X_1 + X_2} + \frac{X_2 \cdot \mathbb{P}[t \in D \mid \mathcal{M}_2(D) \in S]}{X_1 + X_2} \\ &\leq \frac{X_1 (e^\varepsilon \cdot \mathbb{P}[t \in D]) + \delta}{X_1 + X_2} + \frac{X_2 (e^\varepsilon \cdot \mathbb{P}[t \in D]) + \delta}{X_1 + X_2} \\ &\leq e^\varepsilon \cdot \mathbb{P}[t \in D] + \delta \end{aligned}$$

The proof for the post-processing axiom is similar, summing over all possible outputs $\mathcal{M}(D)$. It is straightforward to adapt the proof to all other definitions which change the shape of the prior-posterior bounds. \square

Proposition 6. *Both versions of CDP satisfy both privacy axioms; where the post-processing axiom is modified to only allow post-processing with functions computable on a probabilistic polynomial time Turing machine. CZKPr also satisfies both privacy axioms.*

Proof. For Ind-CDP and the post-processing axiom, the proof is straightforward: if post-processing the output could break the ε -indistinguishability property, the attacker could do this on the original output and break the ε -indistinguishability property of the original definition.

For Ind-CDP and the convexity axiom, without loss of generality, we can assume that the sets of possible outputs of both mechanisms are disjoint (otherwise, this give strictly less information to the attacker). The proof is then the same as for the post-processing axiom.

For SimCDP, applying the same post-processing function to the “true” differentially private mechanism immediately leads to the result, since DP satisfies post-processing. The same reasoning holds for convexity.

The proof that CWPr satisfies both privacy axioms can be found in [BGKS13]; as an immediate corollary, CZKPr also satisfies both axioms. \square

10.1.2 Composition

In this section, if \mathcal{M}_1 and \mathcal{M}_2 are two mechanisms, we denote \mathcal{M}_{1+2} the mechanism defined by $\mathcal{M}_{1+2}(D) = (\mathcal{M}_1(D), \mathcal{M}_2(D))$.

Proposition 7. *If \mathcal{M}_1 is $d_{\mathcal{D}}^1$ -private and \mathcal{M}_2 is $d_{\mathcal{D}}^2$ -private, then \mathcal{M}_{1+2} is $d_{\mathcal{D}}^{1+2}$ -private, where $d_{\mathcal{D}}^{1+2}(D_1, D_2) = d_{\mathcal{D}}^1(D_1, D_2) + d_{\mathcal{D}}^2(D_1, D_2)$.*

Proof. The proof is essentially the same as for ϵ -DP. \mathcal{M}_1 's randomness is independent from \mathcal{M}_2 's, so:

$$\begin{aligned} \mathbb{P} \left[\begin{array}{l} \mathcal{M}_1(D_1) = O_1 \ \& \\ \mathcal{M}_2(D_1) = O_2 \end{array} \right] &= \mathbb{P}[\mathcal{M}_1(D_1) = O_1] \cdot \mathbb{P}[\mathcal{M}_2(D_1) = O_2] \\ &\leq e^{d_{\mathcal{D}}^1(D_1, D_2)} \cdot \mathbb{P}[\mathcal{M}_2(D_2) = O_1] \cdot e^{d_{\mathcal{D}}^2(D_1, D_2)} \cdot \mathbb{P}[\mathcal{M}_2(D_2) = O_2] \\ &\leq e^{d_{\mathcal{D}}^{1+2}(D_1, D_2)} \cdot \mathbb{P}[\mathcal{M}_1(D_2) = O_1 \text{ and } \mathcal{M}_2(D_2) = O_2] \end{aligned}$$

Most definition can also be combined with $d_{\mathcal{D}}$ -privacy, and the composition proofs can be similarly adapted. \square

Proposition 8. *In general, definitions which assume limited background knowledge from the adversary do not compose.*

Proof. The proof of Proposition 7 cannot be adapted to a context in which the attacker has limited background knowledge: as the randomness partially comes from the data-generating distribution, the two probabilities are no longer independent. A typical example considers two mechanisms which answer e.g., queries “how many records satisfy property P ” and “how many records satisfy property P and have an ID different from 4217”: the randomness in the data might make each query private, but the combination of two queries trivially reveals something about a particular user. Variants of this proof can easily be obtained for all definitions with limited background knowledge. \square

11 Scope and related work

In this section, we detail our criteria for excluding particular data privacy definitions from our work, we list some relevant definitions that were excluded by this criteria, and we list related works and existing surveys in the field of data privacy.

11.1 Methodology

Whether a data privacy definition fits our description is not always obvious, so we use the following criterion: the attacker’s capabilities must be clearly defined, and the definition must prevent this attacker from learning about a protected property. Consequently, we do not consider:

- definitions which are a property of the output data and not of the mechanism;
- variants of technical notions that are not data privacy properties, like the different types of sensitivity;
- definitions whose only difference with DP is in the context and not in the formal property, like the distinction between local and global models.

In Section 11.2, we give a list of notions that we found during our survey, and considered to be out of scope for our book. To find a comprehensive list of DP notions, besides the definitions we were aware of or were suggested to us by experts, we conducted a wide literature review using

two research datasets: BASE²⁴ and Google Scholar²⁵. The exact queries were run several times: November 1st in 2018, August 1st in 2019, June 1st in 2020, and 1st of September in 2021. The final result count are summarized in Table 4.

Query (BASE)	Hits
“differential privacy” relax	194
“differential privacy” variant -relax	225
Query (Google Scholar)	Hits
“differential privacy” “new notion”	332
“differential privacy” “new definition” -“new notion”	248

Table 4: Queries for the literature review.

First, we manually reviewed each paper and filtered them out until we had only papers which either contained a new definition or were applying DP in a new setting. All papers which defined a variant or extension of DP are cited in this work.

11.2 Out of scope definitions

As detailed in the previous section, we considered certain data privacy definitions to be out of scope for our work, even when they seem to be related to differential privacy. This section elaborates on such definitions.

Lack of semantic guarantees

Some definitions do not provide clear semantic privacy guarantees, or are only used as a tool in order to prove links between existing definitions. As such, we did not include them in our survey.

- ϵ -privacy²⁶, introduced in [MGG09], was a first attempt at formalizing an adversary with restricted background knowledge by using Dirichlet distribution. This definition imposes a condition on the output, but not on the mechanism, consequently it does not offer strong semantic guarantees like noiseless privacy [Dua09, BBG⁺11] (introduced in Section 6).
- *Relaxed indistinguishability*, introduced in [RHMS09] is a relaxation of adversarial privacy that provides a plausible deniability by requiring for each tuple, that at least l tuples must exist with ϵ -indistinguishability. However, it does not provide any guarantee against Bayesian adversaries.
- *Mutual-information DP*, introduced in [CY16] is an alternative way to average the privacy loss, similar to Section 3. It formalize the intuition that any individual record should not “give out too much information” on the output of the mechanism (or vice-versa).

²⁴<https://www.base-search.net/>

²⁵<https://scholar.google.com/>

²⁶Another definition with the same name is introduced later in this section.

- *Individual privacy*, introduced in [LGHC21] is also a mutual information based extension of DP which can encapsulate the (deterministic) secure multiparty computations [GM84], hence, we exclude it due to the lack of clear semantic privacy guarantees.
- *Differential identifiability*, introduced in [LC12], bounds the probability that a given individual’s information is included in the input datasets but does not measure the “change” in probabilities between the two alternatives. As such, it does not provide any guarantee against Bayesian adversaries²⁷.
- *Crowd-blending privacy*, introduced in [GHL12], combines differential privacy with k -anonymity. As it is strictly weaker than any mechanism which always returns a k -anonymous dataset, the guarantees it provides against a Bayesian adversary are unclear. It is mainly used to show that combining crowd-blending privacy with pre-sampling implies zero-knowledge privacy [GHL12, LP15].
- (k, ε) -*anonymity*, introduced in [HABMA17], first performs k -anonymisation on a subset of the quasi identifiers and then ε -DP on the remaining quasi-identifiers with different settings for each equivalence class of the k -anonymous dataset. The semantic guarantees of this definition are not made explicit.
- *Integral privacy*, introduced in [TNA16] looks at the inverse of the mechanism \mathcal{M} , and require all outputs O to be generated by a large and diverse set of databases D . This could capture definitions like k -anonymity, but offer no semantic guarantee in the Bayesian sense.
- *Membership privacy*²⁸, introduced in [SDS⁺19] is tailored to membership inference attacks on machine learning models; the guarantees it provides are not clear.
- *Posteriori DP*, introduced in [WYZ16], compares two posteriors in a way similar to inferential privacy in Section 7, but does not make the prior (and thus, the attacker model) explicit.
- ε -*privacy*²⁹, introduced in [PZJL20], adopts DP to location data and protect only the secret locations, but it only offers absolute bounds irrespectively of the prior.
- *Noiseless privacy*³⁰, and *measure of privacy*, introduced in [Far21] and [Far19] limits the change in the number of possible outputs when one record in the dataset changes and determines privacy using non-stochastic information theory respectively. Consequently, they do not bound the change in “probabilities” of the mechanism, so do not seem to offer clear guarantees against a Bayesian adversary.
- *Weak DP*, introduced in [WSMD20], adapts DP for streams, but it only provides a DP guarantee for the “average” of all possible mechanism outputs³¹, rather than for the mechanism itself. Thus, its semantics guarantees are also unclear.
- *Data-privacy* and *multi-dimensional data-privacy*, introduced in [HCG18] and [SZH⁺20] ensures that an attacker’s inference accuracy and disclosure probability is below some thresholds.

²⁷Although it was reformulated in [LQS⁺13] as an instance of membership privacy introduced in Section 7.

²⁸Another definition with the same name is introduced in [LQS⁺13], we mention it in Section 7

²⁹Another definition with the same name is introduced earlier in this section.

³⁰Another definition with the same name is introduced in [Dua09, BBG⁺11], we mention it in Section 6.

³¹It also assumes that some uncertainty comes from the data itself, similarly to definitions in Section 7

- *Error Preserving Privacy*, introduced in [DN GRV18], states that the “variance” of the adversary’s error when trying to guess a given user’s record does not change significantly after. The exact adversary model is not specified.

Variants of sensitivity

A important technical tool used when designing differentially private mechanisms is the “sensitivity” of the function that we try to compute. There are many variants to the initial concept of global sensitivity [DMNS06], including local sensitivity [NRS07], smooth sensitivity [NRS07], restricted sensitivity [BBDS13], empirical sensitivity [CZ13], empirical differential privacy³²³³ [ASV13], recommendation-aware sensitivity [ZLR⁺13], record and correlated sensitivity [ZXLZ15], dependence sensitivity [LCM16], per-instance sensitivity [RW21], individual sensitivity [CD20], elastic sensitivity [JNS18], distortion sensitivity [BCDP19], and derivative & partial sensitivity [LPM18, MZU⁺21] respectively. We did not consider these notions as these do not modify the actual definition of differential privacy.

11.3 Local model and other contexts

In this book we focused on DP modifications typically used in the “global model”, in which a central entity has access to the whole dataset. It is also possible to use DP in other contexts, without formally changing the definition. The main alternative is the “local model”, where each individual randomizes their own data before sending it to an aggregator. This model is surveyed in [YLZ⁺20], and formally introduced in [DJW13].

Many definitions we listed were initially presented in the local model, such as $d_{\mathcal{D}}$ -privacy [CABP13], geo-indistinguishability [ABCP13], earth mover’s Pr [FDM19], location Pr [EG16], profile-based DP [GC19], input-discriminative DP [GLXC20], divergence DP and smooth DP from [BD14], and extended DP, distribution Pr, and extended distribution Pr from [KM19b]. Below, we list the definitions that are the same as previously listed definitions, but used in a different attacker setting; the list also includes alternatives to the local and global models.

- In [SCR⁺11], the authors introduce *distributed DP*, which corresponds to local DP, with the additional assumption that only a portion of participants are honest.
- In [KPRU14], the authors define *joint DP*, to model a game in which each player cannot learn the data from any other player, but are still allowed to observe the influence of their data on the mechanism output.
- In [WHWX16], authors define a slightly different version of joined DP, called *multiparty DP*, in which the view of each “subgroup” of players is differentially private in respect to other players inputs.
- In [BEM⁺17], the authors define *DP in the shuffled model*, which falls in-between the global and the local model: the local model is augmented by an anonymous channel that randomly permutes a set of user-supplied messages, and differential privacy is only required of the output of the shuffler.

³²Even though it is introduced as a variant of DP, it was later shown to be a measure of sensitivity [CH16].

³³Another definition with the same name is introduced in [BD19], we mention it in Section 6.

- In [BDT19], the authors defined *local DP for bandits*, a local version of *instantaneous DP* (mentioned in Section 4).
- In [LKCT19], the authors define *task-global DP* and *task-local DP*, which are equivalents of element-level DP (mentioned in Section 4) in a meta-learning context.
- In [MK19], the authors define *utility-optimized local DP*, a local version of one-sided differential privacy (mentioned in Section 4) which additionally guarantees that if the data is considered sensitive, then a certain set of outputs is forbidden.
- In [DPZ⁺18, NYH⁺18, ABK⁺20, SXY21], the authors define *personalized local DP*, a local version of personalized DP (mentioned in Section 5).
- In [ACPP18], the authors define $d_{\mathcal{D}}$ -*local DP*, a local version of $d_{\mathcal{D}}$ -DP (mentioned in Section 5); this was defined as *condensed local DP* in [GTT⁺19].
- In [JLT18], the authors define *localized information privacy*, a local version of information privacy (mentioned in Section 7).

11.4 Related work

Some of the earliest surveys focusing on DP summarize algorithms achieving DP and applications [Dwo08, Dwo09]. The more detailed “privacy book” [DR⁺14] presents an in-depth discussion about the fundamentals of DP, techniques for achieving it, and applications to query-release mechanisms, distributed computations or data streams. Other recent surveys [NR19, RP21] focuses on the release of histograms and synthetic data with DP and describing existing privacy metrics and patterns while providing an overall view of different mathematical privacy preserving framework.

In [HZNF15], the authors classify different privacy enhancing technologies (PETs) into 7 complementary dimensions. Indistinguishability falls into the “Aim” dimension, but within this category, only k -anonymity and oblivious transfer are considered; differential privacy is not mentioned. In [AGM18], the authors survey privacy concerns, measurements and privacy-preserving techniques used in online social networks and recommender systems. They classify privacy into 5 categories; DP falls into “Privacy-preserving models” along with e.g., k -anonymity. In [WE18] the authors classified 80+ privacy metrics into 8 categories based on the output of the privacy mechanism. One of their classes is “Indistinguishability”, which contains DP as well as several variants. Some variants are classified into other categories; for example Rényi DP is classified into “Uncertainty” and mutual-information DP into *Information gain/loss*. The authors list 8 differential privacy variants; our taxonomy can be seen as an extension of the contents of their work (and in particular of the “Indistinguishability” category).

In [WYZ16], authors establish connections between differential privacy (seen as the additional disclosure of an individual’s information due to the release of the data), “identifiability” (seen as the posteriors of recovering the original data from the released data), and “mutual-information privacy” (which measures the average amount of information about the original dataset contained in the released data).

The relation between the main syntactic models of anonymity and DP was studied in [CT13], in which the authors claim that the former is designed for privacy-preserving data publishing (PPDP), while DP is more suitable for privacy preserving data mining (PPDM). The survey [ZZLZ20] investigate the issue of privacy loss due to data correlation under DP models and classify existing

literature into three categories: using parameters to describe data correlation, using models to describe data correlation, and describing data correlation based on the framework.

Other surveys focus on location privacy. In [MH18], the authors highlight privacy concerns in this context and list mechanisms with formal provable privacy guarantees; they describe several variants of differential privacy for streaming (e.g., pan-privacy) and location data (e.g., ge-indistinguishability) along with extensions such as pufferfish and blowfish privacy. In [CEP⁺17], the authors analyze different kinds of privacy breaches and compare metrics that have been proposed to protect location data.

Finally, the appropriate selection of the privacy parameters for DP was also exhaustively studied. This problem is not trivial, and many factors can be considered: in [HGH⁺14], the authors used economic incentives, in [LC11, Kre19, PTB19], the authors looked at individual preferences, and in [LHC⁺19, LP19], the authors took into account an adversary’s capability in terms of hypothesis testing and guessing advantage respectively.

12 Conclusion

We proposed a classification of DP variants and extensions using the concept of dimensions. When possible, we compared definitions from the same dimension, and we showed that definitions from the different dimensions can be combined to form new, meaningful definitions. In theory, it means that even if there were only three possible ways to change a dimension (e.g., making it weaker or stronger), this would result in $3^7 = 2187$ possible definitions: the ≈ 225 existing definitions shown in Figure 1 are only scratching the surface of the space of possible notions. Using these dimensions, we unified and simplified the different notions proposed in the literature. We highlighted their properties such as composability and whether they satisfy the privacy axioms by either collecting the existing results or creating new proofs, and whenever possible, we showed their relative relations to one another. We hope that this work will make the field of data privacy more organized and easier to navigate, especially for new practitioners.

Acknowledgments

The authors would like to thank James Bailie, David Basin, Alex Kulesza, Esfandiar Mohammadi, Jeremy Seeman, and the anonymous reviewers for their helpful comments. This work was partially funded by Google, and started when Balázs Pejó was at University of Luxembourg.

References

- [AAG⁺10] John M Abowd, Fredrik Andersson, Matthew Graham, Lars Vilhuber, and Jeremy Wu. Formal privacy guarantees and analytical validity of onthemap public-use data. <https://ecommons.cornell.edu/handle/1813/47672>, 2010.
- [ABC⁺20] Ahmet Aktay, Shailesh Bavadekar, Gwen Cossoul, John Davis, Damien Desfontaines, Alex Fabrikant, Evgeniy Gabrilovich, Krishna Gadepalli, Bryant Gipson, Miguel Guevara, Chaitanya Kamath, Mansi Kansal, Ali Lange, Chinmoy Mandayam, Andrew Oplinger, Christopher Pluntke, Thomas Roessler, Arran Schlosberg, Tomer Shekel, Swapnil Vispute, Mia Vu, Gregory Wellenius, Brian Williams, and Royce J Wilson.

Google covid-19 community mobility reports: Anonymization process description (version 1.1), 2020.

- [ABCP13] Miguel E Andrés, Nicolás E Bordenabe, Konstantinos Chatzikokolakis, and Catuscia Palamidessi. Geo-indistinguishability: Differential privacy for location-based systems. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*. ACM, 2013.
- [ABK⁺20] Jayadev Acharya, Kallista Bonawitz, Peter Kairouz, Daniel Ramage, and Ziteng Sun. Context aware local differential privacy. In *International Conference on Machine Learning*, pages 52–62. PMLR, 2020.
- [ACG⁺16] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pages 308–318, 2016.
- [ACPP18] Mário Alvim, Konstantinos Chatzikokolakis, Catuscia Palamidessi, and Anna Pazii. Local differential privacy on metric spaces: optimizing the trade-off with utility. In *2018 IEEE 31st Computer Security Foundations Symposium (CSF)*. IEEE, 2018.
- [ADJ19] Hilal Asi, John Duchi, and Omid Javidsbakht. Element level differential privacy: The right granularity of privacy. *arXiv preprint arXiv:1912.04042*, 2019.
- [ADK⁺19] Joshua Allen, Bolin Ding, Janardhan Kulkarni, Harsha Nori, Olga Ohrimenko, and Sergey Yekhanin. An algorithmic framework for differentially private data analysis on trusted processors. In *Advances in Neural Information Processing Systems*, 2019.
- [AGK16] Mohammad Alaggan, Sébastien Gambs, and Anne-Marie Kermarrec. Heterogeneous differential privacy. *Journal of Privacy and Confidentiality*, 7(2):127–158, 2016.
- [AGM18] Erfan Aghasian, Saurabh Garg, and James Montgomery. User’s privacy in recommendation systems applying online social network data, a survey and taxonomy. *arXiv preprint arXiv:1806.07629*, 2018.
- [APV19] Hafiz Asif, Periklis A Papakonstantinou, and Jaideep Vaidya. How to accurately and privately identify anomalies. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2019.
- [ASV13] John M Abowd, Matthew J Schneider, and Lars Vilhuber. Differential privacy applications to bayesian and linear mixed model estimation. *Journal of Privacy and Confidentiality*, 2013.
- [BBD⁺21] Shailesh Bavadekar, Adam Boulanger, John Davis, Damien Desfontaines, Evgeniy Gabrilovich, Krishna Gadepalli, Badih Ghazi, Tague Griffith, Jai Gupta, Chaitanya Kamath, Dennis Kraft, Ravi Kumar, Akim Kumok, Yael Mayer, Pasin Manurangsi, Arti Patankar, Irippuge Milinda Perera, Chris Scott, Tomer Shekel, Benjamin Miller, Karen Smith, Charlotte Stanton, Mimi Sun, Mark Young, and Gregory Wellenius. Google covid-19 vaccination search insights: Anonymization process description, 2021.

- [BBDS13] Jeremiah Blocki, Avrim Blum, Anupam Datta, and Or Sheffet. Differentially private data analysis of social networks via restricted sensitivity. In *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*. ACM, 2013.
- [BBG⁺11] Raghav Bhaskar, Abhishek Bhowmick, Vipul Goyal, Srivatsan Laxman, and Abhradeep Thakurta. Noiseless database privacy. In *International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 2011.
- [BCDP19] Aman Bansal, Rahul Chunduru, Deepesh Data, and Manoj Prabhakaran. Extending the foundations of differential privacy: Robustness and flexibility. 2019.
- [BD14] Rina Foygel Barber and John C Duchi. Privacy and statistical risk: Formalisms and minimax bounds. *arXiv preprint arXiv:1412.4451*, 2014.
- [BD19] Paul Burchard and Anthony Daoud. Empirical differential privacy. *arXiv preprint arXiv:1910.12820*, 2019.
- [BDD⁺20] Shailesh Bavadekar, Andrew Dai, John Davis, Damien Desfontaines, Ilya Eckstein, Katie Everett, Alex Fabrikant, Gerardo Flores, Evgeniy Gabrilovich, Krishna Gade-palli, Shane Glass, Rayman Huang, Chaitanya Kamath, Dennis Kraft, Akim Kumok, Hinali Marfatia, Yael Mayer, Benjamin Miller, Adam Pearce, Irippuge Milinda Perera, Venky Ramachandran, Karthik Raman, Thomas Roessler, Izhak Shafran, Tomer Shekel, Charlotte Stanton, Jacob Stimes, Mimi Sun, Gregory Wellenius, and Masrour Zoghi. Google covid-19 search trends symptoms dataset: Anonymization process description (version 1.0), 2020.
- [BDRS18] Mark Bun, Cynthia Dwork, Guy N Rothblum, and Thomas Steinke. Composable and versatile privacy via truncated cdp. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*. ACM, 2018.
- [BDT19] Debabrota Basu, Christos Dimitrakakis, and Aristide Tossou. Differential privacy for multi-armed bandits: What is it and what is its cost? *arXiv preprint arXiv:1905.12298*, 2019.
- [BEM⁺17] Andrea Bittau, Úlfar Erlingsson, Petros Maniatis, Ilya Mironov, Ananth Raghunathan, David Lie, Mitch Rudominer, Ushasree Kode, Julien Tinnes, and Bernhard Seefeld. Prochlo: Strong privacy for analytics in the crowd. In *Proceedings of the 26th Symposium on Operating Systems Principles*. ACM, 2017.
- [BF16] Raef Bassily and Yoav Freund. Typical stability. *arXiv preprint arXiv:1604.03336*, 2016.
- [BGKS13] Raef Bassily, Adam Groce, Jonathan Katz, and Adam Smith. Coupled-worlds privacy: Exploiting adversarial uncertainty in statistical data privacy. In *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*. IEEE, 2013.
- [BHE⁺18] Johes Bater, Xi He, William Ehrich, Ashwin Machanavajjhala, and Jennie Rogers. Shrinkwrap: efficient sql query processing in differentially private data federations. *Proceedings of the VLDB Endowment*, 2018.

- [BKMR14] Michael Backes, Aniket Kate, Sebastian Meiser, and Tim Ruffing. Differential indistinguishability for cryptography with (bounded) weak sources. *Grande Region Security and Reliability Day (GRSRD)*, 2014.
- [BLR13] Avrim Blum, Katrina Ligett, and Aaron Roth. A learning theory approach to non-interactive database privacy. *Journal of the ACM (JACM)*, 2013.
- [BNS⁺16] Raef Bassily, Kobbi Nissim, Adam Smith, Thomas Steinke, Uri Stemmer, and Jonathan Ullman. Algorithmic stability for adaptive data analysis. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*. ACM, 2016.
- [BS16] Mark Bun and Thomas Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of Cryptography Conference*. Springer, 2016.
- [BSW18] Daniel M Bittner, Anand D Sarwate, and Rebecca N Wright. Using noisy binary search for differentially private anomaly detection. In *International Symposium on Cyber Security Cryptography and Machine Learning*. Springer, 2018.
- [CABP13] Konstantinos Chatzikokolakis, Miguel E Andrés, Nicolás Emilio Bordenabe, and Catuscia Palamidessi. Broadening the scope of differential privacy using metrics. In *International Symposium on Privacy Enhancing Technologies Symposium*. Springer, 2013.
- [CBY⁺18] Zhili Chen, Xianyue Bao, Zuobin Ying, Ximeng Liu, and Hong Zhong. Differentially private location protection with continuous time stamps for vanets. In *International Conference on Algorithms and Architectures for Parallel Processing*. Springer, 2018.
- [CCMS19] TH Hubert Chan, Kai-Min Chung, Bruce M Maggs, and Elaine Shi. Foundations of differentially oblivious algorithms. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2019.
- [CD20] Rachel Cummings and David Durfee. Individual sensitivity preprocessing for data privacy. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2020.
- [CEP⁺17] Kostantinos Chatzikokolakis, Ehab ElSalamouny, Catuscia Palamidessi, Pazii Anna, et al. Methods for location privacy: A comparative overview. *Foundations and Trends® in Privacy and Security*, 2017.
- [CFYD14] Rui Chen, Benjamin C Fung, Philip S Yu, and Bipin C Desai. Correlated network data publication via differential privacy. *The VLDB Journal—The International Journal on Very Large Data Bases*, 2014.
- [CH16] Anne-Sophie Charest and Yiwei Hou. On the meaning and limits of empirical differential privacy. *Journal of Privacy and Confidentiality*, 2016.
- [CIM19] Kamalika Chaudhuri, Jacob Imola, and Ashwin Machanavajjhala. Capacity bounded differential privacy. In *Advances in Neural Information Processing Systems*, 2019.
- [CM06] Kamalika Chaudhuri and Nina Mishra. When random sampling preserves privacy. In *Annual International Cryptology Conference*. Springer, 2006.

- [CMS11] Kamalika Chaudhuri, Claire Monteleoni, and Anand D Sarwate. Differentially private empirical risk minimization. *Journal of Machine Learning Research*, 12(3), 2011.
- [CO15] Sébastien Canard and Baptiste Olivier. Differential privacy in distribution and instance-based noise mechanisms. *IACR Cryptology ePrint Archive*, 2015.
- [Col16] Léo Colisson. L3 internship report: Quantum analog of differential privacy in term of rényi divergence. 2016.
- [CT13] Chris Clifton and Tamir Tassa. On syntactic anonymity and differential privacy. In *2013 IEEE 29th International Conference on Data Engineering Workshops (ICDEW)*. IEEE, 2013.
- [CY16] Paul Cuff and Lanqing Yu. Differential privacy as a mutual information constraint. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2016.
- [CZ13] Shixi Chen and Shuigeng Zhou. Recursive mechanism: towards node differential privacy and unrestricted joins. In *Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data*. ACM, 2013.
- [DA18] Fatemeh Deldar and Mahdi Abadi. Pldp-td: Personalized-location differentially private data analysis on trajectory databases. *Pervasive and Mobile Computing*, 2018.
- [Dal77] Tore Dalenius. Towards a methodology for statistical disclosure control. *statistik Tidskrift*, 1977.
- [DBB21] Ashish Dandekar, Debabrota Basu, and Stéphane Bressan. Differential privacy at risk: Bridging randomness and privacy budget. *Proceedings on Privacy Enhancing Technologies*, 1:64–84, 2021.
- [DGY⁺18] Kai Dong, Taolin Guo, Haibo Ye, Xuansong Li, and Zhen Ling. On the limitations of existing notions of location privacy. *Future Generation Computer Systems*, 2018.
- [DJRT13] Kashyap Dixit, Madhav Jha, Sofya Raskhodnikova, and Abhradeep Thakurta. Testing lipschitz property over product distribution and its applications to statistical data privacy. *Theory of Cryptography - Lecture Notes in Computer Science*, 2013.
- [DJW13] John C Duchi, Michael I Jordan, and Martin J Wainwright. Local privacy and statistical minimax rates. In *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*. IEEE, 2013.
- [DKM⁺06] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In *Eurocrypt*. Springer, 2006.
- [DKY17] Bolin Ding, Janardhan Kulkarni, and Sergey Yekhanin. Collecting telemetry data privately. In *Advances in Neural Information Processing Systems*, 2017.
- [DM05] Cynthia Dwork and Frank McSherry. Differential data privacy, 2005.

- [DMKB19] Damien Desfontaines, Esfandiar Mohammadi, Elisabeth Kraemer, and David Basin. Differential privacy with partial knowledge. *arXiv preprint arXiv:1905.00650*, 2019.
- [DMNS06] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography Conference*. Springer, 2006.
- [DMS⁺21] Ameya Daigavane, Gagan Madan, Aditya Sinha, Abhradeep hakurta Guha, Gaurav Aggarwal, and Prateek Jain. Node-level differential private graph neural networks. *arXiv preprint arXiv:2111.15521*, 2021.
- [DN03] Irit Dinur and Kobbi Nissim. Revealing information while preserving privacy. In *Proceedings of the twenty-second ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*. ACM, 2003.
- [DNGRV18] Thanh Dai Nguyen, Sunil Gupta, Santu Rana, and Svetha Venkatesh. A privacy preserving bayesian optimization with high efficiency. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*. Springer, 2018.
- [DNM⁺13] Christos Dimitrakakis, Blaine Nelson, Aikaterini Mitrokotsa, Benjamin Rubinstein, et al. Bayesian differential privacy through posterior sampling. *arXiv preprint arXiv:1306.1066*, 2013.
- [DNP⁺10] Cynthia Dwork, Moni Naor, Toniann Pitassi, Guy N Rothblum, and Sergey Yekhanin. Pan-private streaming algorithms. In *ICS*, 2010.
- [DNPR10] Cynthia Dwork, Moni Naor, Toniann Pitassi, and Guy N Rothblum. Differential privacy under continual observation. In *Proceedings of the forty-second ACM symposium on Theory of computing*. ACM, 2010.
- [dPCF12] Flávio du Pin Calmon and Nadia Fawaz. Privacy against statistical inference. In *Communication, Control, and Computing (Allerton), 2012 50th Annual Allerton Conference on*. IEEE, 2012.
- [DPZ⁺18] Roel Dobbe, Ye Pu, Jingge Zhu, Kannan Ramchandran, and Claire Tomlin. Customized local differential privacy for multi-agent distributed optimization. *arXiv preprint arXiv:1806.06035*, 2018.
- [DR⁺14] Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 2014.
- [DR16] Cynthia Dwork and Guy N Rothblum. Concentrated differential privacy. *arXiv preprint arXiv:1603.01887*, 2016.
- [DR18] John C Duchi and Feng Ruan. The right complexity measure in locally private estimation: It is not the fisher information. *arXiv preprint arXiv:1806.05756*, 2018.
- [DR19] David Durfee and Ryan M Rogers. Practical differentially private top-k selection with pay-what-you-get composition. In *Advances in Neural Information Processing Systems*, 2019.

- [DRS21] Jinshuo Dong, Aaron Roth, and Weijie Su. Gaussian differential privacy. *Journal of the Royal Statistical Society*, 2021.
- [Dua09] Yitao Duan. Privacy without noise. In *Proceedings of the 18th ACM conference on Information and knowledge management*. ACM, 2009.
- [Dwo06] Cynthia Dwork. Differential privacy. In *Proceedings of the 33rd international conference on Automata, Languages and Programming*. ACM, 2006.
- [Dwo08] Cynthia Dwork. Differential privacy: A survey of results. In *International Conference on Theory and Applications of Models of Computation*. Springer, 2008.
- [Dwo09] Cynthia Dwork. The differential privacy frontier. In *Theory of Cryptography Conference*. Springer, 2009.
- [Dwo10] Cynthia Dwork. Differential privacy in new settings. In *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms*. SIAM, 2010.
- [DWWG13] Xuan Ding, Wei Wang, Meng Wan, and Ming Gu. Seamless privacy: Privacy-preserving subgraph counting in interactive social network analysis. In *Cyber-Enabled Distributed Computing and Knowledge Discovery (CyberC), 2013 International Conference on*. IEEE, 2013.
- [EG16] Ehab ElSalamouny and Sébastien Gambs. Differential privacy models for location-based services. *Transactions on Data Privacy*, 2016.
- [EGS03] Alexandre Evfimievski, Johannes Gehrke, and Ramakrishnan Srikant. Limiting privacy breaches in privacy preserving data mining. In *Proceedings of the twenty-second ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*. ACM, 2003.
- [EPK14] Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. RAPPOR: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security*. ACM, 2014.
- [ESS15] Hamid Ebadi, David Sands, and Gerardo Schneider. Differential privacy: Now it's getting personal. In *Acm Sigplan Notices*. ACM, 2015.
- [Fan18] Liyue Fan. Image pixelization with differential privacy. In *IFIP Annual Conference on Data and Applications Security and Privacy*, pages 148–162. Springer, 2018.
- [Far19] Farhad Farokhi. Development and analysis of deterministic privacy-preserving policies using non-stochastic information theory. *IEEE Transactions on Information Forensics and Security*, 14(10):2567–2576, 2019.
- [Far20] Farhad Farokhi. Discounted differential privacy: privacy of evolving datasets over an infinite horizon. In *ACM/IEEE Int. Conf. on Cyber-Physical Systems (ICCPS), Sydney, Australia*, 2020.
- [Far21] Farhad Farokhi. Noiseless privacy: Definition, guarantees, and applications. *IEEE Transactions on Big Data*, 2021.

- [FC14] Chengfang Fang and Ee-Chien Chang. Differential privacy with delta-neighbourhood for spatial and dynamic datasets. In *Proceedings of the 9th ACM symposium on Information, computer and communications security*. ACM, 2014.
- [FDM19] Natasha Fernandes, Mark Dras, and Annabelle McIver. Generalised differential privacy for text document processing. In *International Conference on Principles of Security and Trust*. Springer, 2019.
- [FMM19] Andrew David Foote, Ashwin Machanavajjhala, and Kevin McKinney. Releasing earnings distributions using differential privacy: Disclosure avoidance system for post-secondary employment outcomes (pseo). *Journal of Privacy and Confidentiality*, 9(2), Oct. 2019.
- [FMTT18] Vitaly Feldman, Ilya Mironov, Kunal Talwar, and Abhradeep Thakurta. Privacy amplification by iteration. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 2018.
- [FS18] Vitaly Feldman and Thomas Steinke. Calibrating noise to variance in adaptive data analysis. *Proceedings of Machine Learning Research*, 2018.
- [GAP18] Simson L Garfinkel, John M Abowd, and Sarah Powazek. Issues encountered deploying differential privacy. In *Proceedings of the 2018 Workshop on Privacy in the Electronic Society*. ACM, 2018.
- [GC19] Joseph Geumlek and Kamalika Chaudhuri. Profile-based privacy for locally private computations. In *2019 IEEE International Symposium on Information Theory (ISIT)*, pages 537–541. IEEE, 2019.
- [GGY21] Jie Gao, Ruobin Gong, and Fang-Yi Yu. Subspace differential privacy. *arXiv preprint arXiv:2108.11527*, 2021.
- [GHL12] Johannes Gehrke, Michael Hay, Edward Lui, and Rafael Pass. Crowd-blending privacy. In *Advances in Cryptology—CRYPTO 2012*. Springer, 2012.
- [GK17] Arpita Ghosh and Robert Kleinberg. Inferential privacy guarantees for differentially private mechanisms. In *8th Innovations in Theoretical Computer Science Conference (ITCS 2017)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017.
- [GKPT15] Rachid Guerraoui, Anne-Marie Kermarrec, Rhicheek Patra, and Mahsa Taziki. D 2 p: distance-based differential privacy in recommenders. *Proceedings of the VLDB Endowment*, 2015.
- [GKS08] Srivatsava Ranjit Ganta, Shiva Prasad Kasiviswanathan, and Adam Smith. Composition attacks and auxiliary information in data privacy. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2008.
- [GKY11] Adam Groce, Jonathan Katz, and Arkady Yerukhimovich. Limits of computational differential privacy in the client/server setting. In *Theory of Cryptography Conference*. Springer, 2011.

- [GLP11] Johannes Gehrke, Edward Lui, and Rafael Pass. Towards privacy for social networks: A zero-knowledge based definition of privacy. In *Theory of Cryptography Conference*. Springer, 2011.
- [GLXC20] Xiaolan Gu, Ming Li, Li Xiong, and Yang Cao. Providing input-discriminative protection for local differential privacy. In *2020 IEEE 36th International Conference on Data Engineering (ICDE)*, pages 505–516. IEEE, 2020.
- [GM84] Shafi Goldwasser and Silvio Micali. Probabilistic encryption. *Journal of computer and system sciences*, 1984.
- [GMR89] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on computing*, 1989.
- [GR15] Arpita Ghosh and Aaron Roth. Selling privacy at auction. *Games and Economic Behavior*, 2015.
- [GSC17] Joseph Geumlek, Shuang Song, and Kamalika Chaudhuri. Renyi differential privacy mechanisms for posterior sampling. *Advances in Neural Information Processing Systems*, 2017.
- [GTT⁺19] Mehmet Emre Gursoy, Acar Tamersoy, Stacey Truex, Wenqi Wei, and Ling Liu. Secure and utility-aware data collection with condensed local differential privacy. *IEEE Transactions on Dependable and Secure Computing*, 2019.
- [HABMA17] Naoise Holohan, Spiros Antonatos, Stefano Braghin, and Pól Mac Aonghusa. (k,e)-anonymity: k-anonymity with e-differential privacy. *arXiv preprint arXiv:1710.01615*, 2017.
- [Hal12] Rob Hall. *New Statistical Applications for Differential Privacy*. PhD thesis, PhD thesis, Carnegie Mellon, 2012.
- [HCG18] Jianping He, Lin Cai, and Xinping Guan. Preserving data-privacy with added noises: Optimal estimation and privacy analysis. *IEEE Transactions on Information Theory*, 64(8):5677–5690, 2018.
- [HD19] Yufan Huang and Huaiyu Dai. Quantifying differential privacy of gossip protocols in general networks. *arXiv preprint arXiv:1905.07598*, 2019.
- [HDS⁺20] Amaç Herdağdelen, Alex Dow, Bogdan State, Payman Mohassel, and Alex Pompe. Protecting privacy in facebook mobility data during the covid-19 response. <https://research.fb.com/blog/2020/06/protecting-privacy-in-facebook-mobility-data-during-the-covid-19-response/>, 2020.
- [HGH⁺14] Justin Hsu, Marco Gaboardi, Andreas Haeberlen, Sanjeev Khanna, Arjun Narayan, Benjamin C Pierce, and Aaron Roth. Differential privacy: An economic method for choosing epsilon. In *2014 IEEE 27th Computer Security Foundations Symposium*. IEEE, 2014.
- [HLMJ09] Michael Hay, Chao Li, Gerome Miklau, and David Jensen. Accurate estimation of the degree distribution of private networks. In *Data Mining, 2009. ICDM'09. Ninth IEEE International Conference on*. IEEE, 2009.

- [HMD14] Xi He, Ashwin Machanavajjhala, and Bolin Ding. Blowfish privacy: Tuning privacy-utility trade-offs using policies. In *Proceedings of the 2014 ACM SIGMOD international conference on Management of data*. ACM, 2014.
- [HMFS17] Xi He, Ashwin Machanavajjhala, Cheryl Flynn, and Divesh Srivastava. Composing differential privacy and secure computation: A case study on scaling private record linkage. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2017.
- [HMQN13] Matthias Huber, Jörn Müller-Quade, and Tobias Nilges. Defining privacy based on distributions of privacy breaches. In *Number Theory and Cryptography*. Springer, 2013.
- [HMSS19] Iftach Haitner, Noam Mazon, Ronen Shaltiel, and Jad Silbak. Channels of small log-ratio leakage and characterization of two-party differentially private computation. In *Theory of Cryptography Conference*. Springer, 2019.
- [HWR13] Robert Hall, Larry Wasserman, and Alessandro Rinaldo. Random differential privacy. *Journal of Privacy and Confidentiality*, 4(2), 2013.
- [HZNF15] Johannes Heurix, Peter Zimmermann, Thomas Neubauer, and Stefan Fenz. A taxonomy for privacy enhancing technologies. *Computers & Security*, 2015.
- [JB14] Márk Jelasity and Kenneth P Birman. Distributional differential privacy for large-scale smart metering. In *Proceedings of the 2nd ACM workshop on Information hiding and multimedia security*. ACM, 2014.
- [JLH19] Austin Jones, Kevin Leahy, and Matthew Hale. Towards differential privacy for symbolic systems. In *2019 American Control Conference (ACC)*. IEEE, 2019.
- [JLT18] Bo Jiang, Ming Li, and Ravi Tandon. Context-aware data aggregation with localized information privacy. In *2018 IEEE Conference on Communications and Network Security (CNS)*, pages 1–9. IEEE, 2018.
- [JNS18] Noah Johnson, Joseph P Near, and Dawn Song. Towards practical differential privacy for sql queries. *Proceedings of the VLDB Endowment*, 2018.
- [JYC15] Zach Jorgensen, Ting Yu, and Graham Cormode. Conservative or liberal? personalized differential privacy. In *Data Engineering (ICDE), 2015 IEEE 31st International Conference on*. IEEE, 2015.
- [KDH⁺20] Ios Kotsogiannis, Stelios Doudalis, Sam Haney, Ashwin Machanavajjhala, and Sharad Mehrotra. One-sided differential privacy. In *2020 IEEE 36th International Conference on Data Engineering (ICDE)*, pages 493–504. IEEE, 2020.
- [KKNO17] Georgios Kellaris, George Kollios, Kobbi Nissim, and Adam O’Neill. Accessing data while preserving privacy. *arXiv preprint arXiv:1706.01552*, 2017.
- [KL10] Daniel Kifer and Bing-Rong Lin. Towards an axiomatization of statistical privacy and utility. In *Proceedings of the twenty-ninth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*. ACM, 2010.

- [KL12] Daniel Kifer and Bing-Rong Lin. An axiomatic view of statistical privacy and utility. *Journal of Privacy and Confidentiality*, 2012.
- [KLL19] Hasan B Kartal, Xiaoping Liu, and Xiao-Bai Li. Differential privacy for the vast majority. *ACM Transactions on Management Information Systems (TMIS)*, 2019.
- [KM11] Daniel Kifer and Ashwin Machanavajjhala. No free lunch in data privacy. In *Proceedings of the 2011 ACM SIGMOD International Conference on Management of data*. ACM, 2011.
- [KM12] Daniel Kifer and Ashwin Machanavajjhala. A rigorous and customizable framework for privacy. In *Proceedings of the 31st ACM SIGMOD-SIGACT-SIGAI symposium on Principles of Database Systems*. ACM, 2012.
- [KM19a] Yusuke Kawamoto and Takao Murakami. Local distribution obfuscation via probability coupling. In *2019 57th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, 2019.
- [KM19b] Yusuke Kawamoto and Takao Murakami. Local obfuscation mechanisms for hiding probability distributions. In *European Symposium on Research in Computer Security*. Springer, 2019.
- [KM20] Vishaal Krishnan and Sonia Martínez. A probabilistic framework for moving-horizon estimation: Stability and privacy guarantees. *IEEE Transactions on Automatic Control*, 66(4):1817–1824, 2020.
- [KOV17] Peter Kairouz, Sewoong Oh, and Pramod Viswanath. The composition theorem for differential privacy. *IEEE Transactions on Information Theory*, 2017.
- [KPRU14] Michael Kearns, Mallesh Pai, Aaron Roth, and Jonathan Ullman. Mechanism design in large games: Incentives and privacy. In *Proceedings of the 5th conference on Innovations in theoretical computer science*. ACM, 2014.
- [KPXP14] Georgios Kellaris, Stavros Papadopoulos, Xiaokui Xiao, and Dimitris Papadias. Differentially private event sequences over infinite streams. *Proceedings of the VLDB Endowment*, 2014.
- [Kre19] Sara Krehbiel. Choosing epsilon for privacy as a service. *Proceedings on Privacy Enhancing Technologies*, 2019.
- [KRWY16] Michael Kearns, Aaron Roth, Zhiwei Steven Wu, and Grigory Yaroslavtsev. Private algorithms for the protected in social network search. *Proceedings of the National Academy of Sciences*, 2016.
- [KS14] Shiva P Kasiviswanathan and Adam Smith. On the ‘semantics’ of differential privacy: A bayesian formulation. *Journal of Privacy and Confidentiality*, 2014.
- [LBG17] Yunhui Long, Vincent Bindschaedler, and Carl A Gunter. Towards measuring membership privacy. *Unknown Journal*, 2017.
- [LC11] Jaewoo Lee and Chris Clifton. How much is enough? choosing ϵ for differential privacy. In *International Conference on Information Security*. Springer, 2011.

- [LC12] Jaewoo Lee and Chris Clifton. Differential identifiability. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2012.
- [LCM16] Changchang Liu, Supriyo Chakraborty, and Prateek Mittal. Dependence makes you vulnerable: Differential privacy under dependent tuples. In *NDSS*, 2016.
- [LDX⁺21] Bo Liu, Ming Ding, Hanyu Xue, Tianqing Zhu, Dayong Ye, Li Song, and Wanlei Zhou. Dp-image: Differential privacy for image data in feature space. *arXiv preprint arXiv:2103.07073*, 2021.
- [LGHC21] Qiongxiu Li, Jaron Skovsted Gundersen, Richard Heusdens, and Mads Græsbøll Christensen. Privacy-preserving distributed processing: Metrics, bounds and algorithms. *IEEE Transactions on Information Forensics and Security*, 16:2090–2103, 2021.
- [LHC⁺19] Changchang Liu, Xi He, Thee Chanyaswad, Shiqiang Wang, and Prateek Mittal. Investigating statistical privacy frameworks from the perspective of hypothesis testing. *Proceedings on Privacy Enhancing Technologies*, 2019.
- [LK18] Jaewoo Lee and Daniel Kifer. Concentrated differentially private gradient descent with adaptive per-iteration privacy budget. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2018.
- [LKCT19] Jeffrey Li, Mikhail Khodak, Sebastian Caldas, and Ameet Talwalkar. Differentially private meta-learning. In *International Conference on Learning Representations*, 2019.
- [LL12] Samantha Leung and Edward Lui. Bayesian mechanism design with efficiency, privacy, and approximate truthfulness. In *International Workshop on Internet and Network Economics*. Springer, 2012.
- [LLV07] Ninghui Li, Tiancheng Li, and Suresh Venkatasubramanian. t-closeness: Privacy beyond k-anonymity and l-diversity. In *Data Engineering, 2007. ICDE 2007. IEEE 23rd International Conference on*. IEEE, 2007.
- [LNP13] Jerome Le Ny and George J Pappas. Differentially private filtering. *IEEE Transactions on Automatic Control*, 59(2):341–354, 2013.
- [LP15] Edward Lui and Rafael Pass. Outlier privacy. In *Theory of Cryptography Conference*. Springer, 2015.
- [LP19] Peeter Laud and Alisa Pankova. Interpreting epsilon of differential privacy in terms of advantage in guessing or approximating sensitive attributes. *arXiv preprint arXiv:1911.12777*, 2019.
- [LPM18] Peeter Laud, Alisa Pankova, and Pettai Martin. Achieving differential privacy using methods from calculus. *arXiv preprint arXiv:1811.06343*, 2018.
- [LPR20] Katrina Ligett, Charlotte Peale, and Omer Reingold. Bounded-leakage differential privacy. In *1st Symposium on Foundations of Responsible Computing (FORC 2020)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020.

- [LQS12] Ninghui Li, Wahbeh Qardaji, and Dong Su. On sampling, anonymization, and differential privacy or, k-anonymization meets differential privacy. 2012.
- [LQS⁺13] Ninghui Li, Wahbeh Qardaji, Dong Su, Yi Wu, and Weining Yang. Membership privacy: a unifying framework for privacy definitions. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*. ACM, 2013.
- [LRY19] Yanan Li, Xuebin Ren, Shusen Yang, and Xinyu Yang. Impact of prior knowledge and data correlation on privacy leakage: A unified analysis. *IEEE Transactions on Information Forensics and Security*, 2019.
- [LWS15] Ziqi Liu, Yu-Xiang Wang, and Alexander Smola. Fast differentially private matrix factorization. In *Proceedings of the 9th ACM Conference on Recommender Systems*. ACM, 2015.
- [LX21] Ao Liu and Lirong Xia. Smoothed differential privacy. *arXiv preprint arXiv:2107.01559*, 2021.
- [LXL13] Jinfei Liu, Li Xiong, and Jun Luo. Semantic security: Privacy definitions revisited. *Trans. Data Privacy*, 2013.
- [M⁺21] H Brendan McMahan et al. Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, 14(1), 2021.
- [McC15] David R McClure. *Relaxations of differential privacy and risk/utility evaluations of synthetic data and fidelity measures*. PhD thesis, Duke University, 2015.
- [MDH⁺20] Solomon Messing, Christina DeGregorio, Bennett Hillenbrand, Gary King, Saurav Mahanti, Zagreb Mukerjee, Chaya Nayak, Nate Persily, Bogdan State, and Arjun Wilkins. Facebook Privacy-Protected Full URLs Data Set, 2020.
- [Mei18] Sebastian Meiser. Approximate and probabilistic differential privacy definitions. *Cryptology ePrint Archive, Report 2018/277*, 2018.
- [MGG09] Ashwin Machanavajjhala, Johannes Gehrke, and Michaela Götz. Data publishing against realistic adversaries. *Proceedings of the VLDB Endowment*, 2009.
- [MGKV06] Ashwin Machanavajjhala, Johannes Gehrke, Daniel Kifer, and Muthuramakrishnan Venkatasubramanian. l-diversity: Privacy beyond k-anonymity. In *Data Engineering, 2006. ICDE'06. Proceedings of the 22nd International Conference on*. IEEE, 2006.
- [MH18] Ashwin Machanavajjhala and Xi He. Analyzing your location data with provable privacy guarantees. In *Handbook of Mobile Data Privacy*. Springer, 2018.
- [Mir17] Ilya Mironov. Renyi differential privacy. In *Computer Security Foundations Symposium (CSF), 2017 IEEE 30th*. IEEE, 2017.
- [MK19] Takao Murakami and Yusuke Kawamoto. Utility-optimized local differential privacy mechanisms for distribution estimation. In *28th {USENIX} Security Symposium ({USENIX} Security 19)*, 2019.

- [MKA⁺08] Ashwin Machanavajjhala, Daniel Kifer, John Abowd, Johannes Gehrke, and Lars Vilhuber. Privacy: Theory meets practice on the map. In *Proceedings of the 2008 IEEE 24th International Conference on Data Engineering*. IEEE Computer Society, 2008.
- [MP20] Stephen Suffian Joe Glass Adam Scheer McGee Young Matt Golden Marc Paré, Mariano Teehan. Applying energy differential privacy to enable measurement of the ohmconnect virtual power plant, 2020.
- [MPRV09] Ilya Mironov, Omkant Pandey, Omer Reingold, and Salil Vadhan. Computational differential privacy. In *Advances in Cryptology-CRYPTO 2009*. Springer, 2009.
- [MRTZ18] H Brendan McMahan, Daniel Ramage, Kunal Talwar, and Li Zhang. Learning differentially private recurrent language models. In *International Conference on Learning Representations*, 2018.
- [MTV⁺20] Fatemehsadat Mireshghallah, Mohammadkazem Taram, Praneeth Vepakomma, Abhishek Singh, Ramesh Raskar, and Hadi Esmaeilzadeh. Privacy in deep learning: A survey. *arXiv preprint arXiv:2004.12254*, 2020.
- [MZU⁺21] Tamara T Mueller, Alexander Ziller, Dmitrii Usynin, Moritz Knolle, Friederike Jungmann, Daniel Rueckert, and Georgios Kaissis. Partial sensitivity analysis in differential privacy. *arXiv preprint arXiv:2109.10582*, 2021.
- [NAD14] Nadia Niknami, Mahdi Abadi, and Fatemeh Deldar. Spatialpdp: A personalized differentially private mechanism for range counting queries over spatial databases. In *Computer and Knowledge Engineering (ICCKE), 2014 4th International eConference on*. IEEE, 2014.
- [Noz19] Erfan Nozari. *Networked Dynamical Systems: Privacy, Control, and Cognition*. PhD thesis, UC San Diego, 2019.
- [NR19] Boel Nelson and Jenni Reuben. Sok: Chasing accuracy and privacy, and catching both in differentially private histogram publication. *arXiv*, 2019.
- [NRS07] Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. Smooth sensitivity and sampling in private data analysis. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*. ACM, 2007.
- [NV20] Moni Naor and Neil Vexler. Can two walk together: Privacy enhancing methods and preventing tracking of users. In *1st Symposium on Foundations of Responsible Computing (FORC 2020)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020.
- [NYH⁺18] Yiwen Nie, Wei Yang, Liusheng Huang, Xike Xie, Zhenhua Zhao, and Shaowei Wang. A utility-optimized framework for personalized private histogram estimation. *IEEE Transactions on Knowledge and Data Engineering*, 2018.
- [PAE⁺16] Nicolas Papernot, Martín Abadi, Ulfar Erlingsson, Ian Goodfellow, and Kunal Talwar. Semi-supervised knowledge transfer for deep learning from private training data. *arXiv preprint arXiv:1610.05755*, 2016.

- [PGM14] Davide Proserpio, Sharon Goldberg, and Frank McSherry. Calibrating data to sensitivity in private data analysis: a platform for differentially-private analysis of weighted datasets. *Proceedings of the VLDB Endowment*, 2014.
- [PHB21] Kendall Parker, Matthew Hale, and Prabir Barooah. Spectral differential privacy: Application to smart meter data. *IEEE Internet of Things Journal*, 2021.
- [Pin18] Rafael Pinot. Minimum spanning tree release under differential privacy constraints. *arXiv preprint arXiv:1801.06423*, 2018.
- [PPY19] Sarvar Patel, Giuseppe Persiano, and Kevin Yeo. What storage access privacy is achievable with small overhead? In *Proceedings of the 38th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, 2019.
- [PTB19] Balazs Pejo, Qiang Tang, and Gergely Biczók. Together or alone: The price of privacy in collaborative learning. *Proceedings on Privacy Enhancing Technologies*, 2019.
- [PYGPA19] Rafael Pinot, Florian Yger, Cédric Gouy-Pailler, and Jamal Atif. A unified view on differential privacy and robustness to adversarial examples. *arXiv preprint arXiv:1906.07982*, 2019.
- [PZJL20] Alireza Partovi, Wei Zheng, Taeho Jung, and Hai Lin. Ensuring privacy in location-based services: a model-based approach. *arXiv preprint arXiv:2002.10055*, 2020.
- [RA17] Benjamin IP Rubinstein and Francesco Aldà. Pain-free random differential privacy with sensitivity sampling. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 2017.
- [RCM⁺20] Ryan Rogers, Adrian Rivera Cardoso, Koray Mancuhan, Akash Kaura, Nikhil Gahlawat, Neha Jain, Paul Ko, and Parvez Ahammad. A members first approach to enabling linkedin’s labor market insights at scale, 2020.
- [Reu18] Jenni Reuben. Towards a differential privacy theory for edge-labeled directed graphs. *SICHERHEIT 2018*, 2018.
- [RHMS09] Vibhor Rastogi, Michael Hay, Gerome Miklau, and Dan Suciu. Relationship privacy: output perturbation for queries with joins. In *Proceedings of the twenty-eighth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*. ACM, 2009.
- [Rot10] Aaron Roth. New algorithms for preserving differential privacy. *Microsoft Research*, 2010.
- [RP21] Sukumar Rajendran and J Prabhu. A novel study of different privacy frameworks metrics and patterns. In *Advances in Distributed Computing and Machine Learning*, pages 181–196. Springer, 2021.
- [RSP⁺20] Ryan Rogers, Subbu Subramaniam, Sean Peng, David Durfee, Seunghyun Lee, Santosh Kumar Kancha, Shraddha Sahay, and Parvez Ahammad. LinkedIn’s audience engagements api: A privacy preserving data analytics system at scale, 2020.

- [RW21] Rachel Redberg and Yu-Xiang Wang. Privately publishable per-instance privacy. In *NeurIPS 2020 Competition and Demonstration Track*. PMLR, 2021.
- [Sam01] Pierangela Samarati. Protecting respondents identities in microdata release. *IEEE transactions on Knowledge and Data Engineering*, 2001.
- [SCDFSM17] Jordi Soria-Comas, Josep Domingo-Ferrer, David Sánchez, and David Megías. Individual differential privacy: A utility-preserving formulation of differential privacy guarantees. *IEEE Transactions on Information Forensics and Security*, 2017.
- [SCR⁺11] Elaine Shi, HTH Chan, Eleanor Rieffel, Richard Chow, and Dawn Song. Privacy-preserving aggregation of time-series data. In *Annual Network & Distributed System Security Symposium (NDSS)*. Internet Society., 2011.
- [SDS⁺19] Alexandre Sablayrolles, Matthijs Douze, Cordelia Schmid, Yann Ollivier, and Hervé Jégou. White-box vs black-box: Bayes optimal strategies for membership inference. In *International Conference on Machine Learning*, pages 5558–5567. PMLR, 2019.
- [Sea16] Adam Sealon. Shortest paths and distances with differential privacy. In *Proceedings of the 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*. ACM, 2016.
- [SM84] Goldwasser Shafi and Silvio Micali. Probabilistic encryption. *Journal of computer and system sciences*, 1984.
- [SMM19] David M Sommer, Sebastian Meiser, and Esfandiar Mohammadi. Privacy loss classes: The central limit theorem in differential privacy. *Proceedings on Privacy Enhancing Technologies*, 2019.
- [SPB21] István András Seres, Balázs Pejó, and Péter Burcsi. The effect of false positives: Why fuzzy message detection leads to fuzzy privacy guarantees? *arXiv preprint arXiv:2109.06576*, 2021.
- [SSB16] Sean Simmons, Cenk Sahinalp, and Bonnie Berger. Enabling privacy-preserving gwass in heterogeneous human populations. *Cell systems*, 2016.
- [ST12] Klara Stokes and Vicenç Torra. n-confusion: a generalization of k-anonymity. In *Proceedings of the 2012 Joint EDBT/ICDT Workshops*. ACM, 2012.
- [Swe02] Latanya Sweeney. k-anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 2002.
- [S XK⁺19] Haipei Sun, Xiaokui Xiao, Issa Khalil, Yin Yang, Zhan Qin, Hui Wendy Wang, and Ting Yu. Analyzing subgraph statistics from extended local views with decentralized differential privacy. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2019.
- [SXY21] Zixuan Shen, Zhihua Xia, and Peipeng Yu. Pldp: Personalized local differential privacy for multidimensional data aggregation. *Security and Communication Networks*, 2021, 2021.

- [SZH⁺20] Mingjing Sun, Chengcheng Zhao, Jianping He, Peng Cheng, and Daniel Quevedo. Privacy-preserving correlated data publication: Privacy analysis and optimal noise design. *IEEE Transactions on Network Science and Engineering*, 2020.
- [TC12] Christine Task and Chris Clifton. A guide to differential privacy theory in social network analysis. In *Proceedings of the 2012 International Conference on Advances in Social Networks Analysis and Mining (ASONAM 2012)*. IEEE Computer Society, 2012.
- [TCAY19] Shun Takagi, Yang Cao, Yasuhito Asano, and Masatoshi Yoshikawa. Geo-graph-indistinguishability: Protecting location privacy for lbs over road networks. In *IFIP Annual Conference on Data and Applications Security and Privacy*. Springer, 2019.
- [TCAY20] Shun Takagi, Yang Cao, Yasuhito Asano, and Masatoshi Yoshikawa. Geo-graph-indistinguishability: Location privacy on road networks based on differential privacy. *arXiv preprint arXiv:2010.13449*, 2020.
- [TCY21] Shun Takagi, Yang Cao, and Masatoshi Yoshikawa. Asymmetric differential privacy. *arXiv preprint arXiv:2103.00996*, 2021.
- [TD16] Aristide CY Tossou and Christos Dimitrakakis. Algorithms for differentially private multi-armed bandits. In *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.
- [TDG16] Raphael R Toledo, George Danezis, and Ian Goldberg. Lower-cost e-private information retrieval. *Proceedings on Privacy Enhancing Technologies*, 2016.
- [Tea17] Differential Privacy Team. Learning with privacy at scale. <https://machinelearning.apple.com/docs/learning-with-privacy-at-scale/appledifferentialprivacysystem.pdf>, 2017.
- [TF20] Aleksei Triastcyn and Boi Faltings. Bayesian differential privacy for machine learning. In *International Conference on Machine Learning*, pages 9583–9592. PMLR, 2020.
- [TNA16] Vicenç Torra and Guillermo Navarro-Arribas. Integral privacy. In *International Conference on Cryptology and Network Security*, pages 661–669. Springer, 2016.
- [TSD20] Michael Carl Tschantz, Shayak Sen, and Anupam Datta. Sok: Differential privacy as a causal property. In *2020 IEEE Symposium on Security and Privacy (SP)*, 2020.
- [War65] Stanley L Warner. Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 1965.
- [WBK19] Yu-Xiang Wang, Borja Balle, and Shiva Prasad Kasiviswanathan. Subsampled rényi differential privacy and analytical moments accountant. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1226–1235. PMLR, 2019.
- [WCM18] Sameer Wagh, Paul Cuff, and Prateek Mittal. Differentially private oblivious ram. *Proceedings on Privacy Enhancing Technologies*, 2018.
- [WDN17] Xiaotong Wu, Wanchun Dou, and Qiang Ni. Game theory based privacy preserving analysis in correlated data publication. In *Proceedings of the Australasian Computer Science Week Multiconference*. ACM, 2017.

- [WE18] Isabel Wagner and David Eckhoff. Technical privacy metrics: a systematic survey. *ACM Computing Surveys (CSUR)*, 2018.
- [WHWX16] Genqiang Wu, Yeping He, Jingzheng Wu, and Xianyao Xia. Inherit differential privacy in distributed setting: Multipartly randomized function computation. In *Trust-com/BigDataSE/I SPA, 2016 IEEE*. IEEE, 2016.
- [WLF16] Yu-Xiang Wang, Jing Lei, and Stephen E Fienberg. On-average kl-privacy and its equivalence to generalization for max-entropy mechanisms. In *International Conference on Privacy in Statistical Databases*. Springer, 2016.
- [WSMD20] Yu Wang, Hussein Sibai, Sayan Mitra, and Geir E Dullerud. Differential privacy for sequential algorithms. *arXiv preprint arXiv:2004.00275*, 2020.
- [WW21a] Hao Wang and Huan Wang. Correlated tuple data release via differential privacy. *Information Sciences*, 560:347–369, 2021.
- [WW21b] Hao Wang and Huan Wang. Differentially private publication for correlated non-numerical data. *The Computer Journal*, 2021.
- [WWK⁺17] Xiaotong Wu, Taotao Wu, Maqbool Khan, Qiang Ni, and Wanchun Dou. Game theory based correlated privacy preserving analysis in big data. *IEEE Transactions on Big Data*, 2017.
- [WX17] Hao Wang and Zhengquan Xu. Cts-dp: publishing correlated time-series data via differential privacy. *Knowledge-Based Systems*, 122:167–179, 2017.
- [WXH17] Genqiang Wu, Xianyao Xia, and Yeping He. Information theory of data privacy. *arXiv preprint arXiv:1703.07474*, 2017.
- [WYZ16] Weina Wang, Lei Ying, and Junshan Zhang. On the relation between identifiability, differential privacy, and mutual-information privacy. *IEEE Transactions on Information Theory*, 2016.
- [WZ10] Larry Wasserman and Shuheng Zhou. A statistical framework for differential privacy. *Journal of the American Statistical Association*, 2010.
- [WZL⁺20] Royce J Wilson, Celia Yuxin Zhang, William Lam, Damien Desfontaines, Daniel Simmons-Marengo, and Bryant Gipson. Differentially private sql with bounded user contribution. *Proceedings on Privacy Enhancing Technologies*, 2:230–250, 2020.
- [XX15] Yonghui Xiao and Li Xiong. Protecting locations with differential privacy under temporal correlations. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2015.
- [YLL⁺17] Ziqi Yan, Jiqiang Liu, Gang Li, Zhen Han, and Shuo Qiu. Privmin: Differentially private minhash for jaccard similarity computation. *arXiv preprint arXiv:1705.07258*, 2017.
- [YLZ⁺20] Mengmeng Yang, Lingjuan Lyu, Jun Zhao, Tianqing Zhu, and Kwok-Yan Lam. Local differential privacy and its applications: A comprehensive survey. *arXiv preprint arXiv:2008.03686*, 2020.

- [YSN15] Bin Yang, Issei Sato, and Hiroshi Nakagawa. Bayesian differential privacy on correlated data. In *Proceedings of the 2015 ACM SIGMOD international conference on Management of Data*. ACM, 2015.
- [YWW13] Xiaowei Ying, Xintao Wu, and Yue Wang. On linear refinement of differential privacy-preserving query answering. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*. Springer, 2013.
- [ZCLS21] Qinqing Zheng, Shuxiao Chen, Qi Long, and Weijie Su. Federated f-differential privacy. In *International Conference on Artificial Intelligence and Statistics*, pages 2251–2259. PMLR, 2021.
- [ZLBR20] Hailong Zhang, Sufian Latif, Raef Bassily, and Atanas Rountev. Differentially-private control-flow node coverage for software usage analysis. In *USENIX Security Symposium*, pages 1021–1038, 2020.
- [ZLR⁺13] Tianqing Zhu, Gang Li, Yongli Ren, Wanlei Zhou, and Ping Xiong. Differential privacy for neighborhood-based collaborative filtering. In *Proceedings of the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*. ACM, 2013.
- [ZLW09] Shuheng Zhou, Katrina Ligett, and Larry Wasserman. Differential privacy with compression. In *Information Theory, 2009. ISIT 2009. IEEE International Symposium on*. IEEE, 2009.
- [ZQZ⁺15] Zijian Zhang, Zhan Qin, Liehuang Zhu, Wei Jiang, Chen Xu, and Kui Ren. Toward practical differential privacy in smart grid with capacity-limited rechargeable batteries, 2015.
- [ZSZ⁺18] Jinxue Zhang, Jingchao Sun, Rui Zhang, Yanchao Zhang, and Xia Hu. Privacy-preserving social media data outsourcing. In *IEEE INFOCOM 2018-IEEE Conference on Computer Communications*. IEEE, 2018.
- [ZWB⁺19] Jun Zhao, Teng Wang, Tao Bai, Kwok-Yan Lam, Zhiying Xu, Shuyu Shi, Xuebin Ren, Xinyu Yang, Yang Liu, and Han Yu. Reviewing and improving the gaussian mechanism for differential privacy. *arXiv preprint arXiv:1911.12060*, 2019.
- [ZXLZ15] Tianqing Zhu, Ping Xiong, Gang Li, and Wanlei Zhou. Correlated differential privacy: hiding information in non-iid data set. *IEEE Transactions on Information Forensics and Security*, 2015.
- [ZZLZ20] Tao Zhang, Tianqing Zhu, Renping Liu, and Wanlei Zhou. Correlated data in differential privacy: Definition and analysis. *Concurrency and Computation: Practice and Experience*, page e6015, 2020.