Two Views of Constrained Differential Privacy: Belief Revision and Update

Likang Liu[*1], **Keke Sun**[*1], **Chunlai Zhou**[†1], **Yuan Feng**[2]

[1] School of Information, Renmin University of China, Beijing, China
[2] Centre of Quantum Software and Information, University of Technology Sydney, Australia
micahliu2012@gmail.com, skk2020@ruc.edu.cn, czhou@ruc.edu.cn, Yuan.feng@uts.edu.au

Abstract

In this paper, we provide two views of constrained differential private (DP) mechanisms. The first one is as belief revision. A constrained DP mechanism is obtained by standard probabilistic conditioning, and hence can be naturally implemented by Monte Carlo algorithms. The other is as belief update. A constrained DP is defined according to l_2-distance minimization postprocessing or projection and hence can be naturally implemented by optimization algorithms. The main advantage of these two perspectives is that we can make full use of the machinery of belief revision and update to show basic properties for constrained differential privacy especially some important *new* composition properties. Within the framework established in this paper, constrained DP algorithms in the literature can be classified either as belief revision or belief update. At the end of the paper, we demonstrate their differences especially in utility in a couple of scenarios.

Introduction

Theories of belief revision and update have been an important field in AI community, especially in knowledge representation and database systems (van Harmelen, Lifschitz, and Porter 2008). An agent's beliefs about the world may be incorrect or incomplete and she wants to change the beliefs. Such a process is known as *belief revision* (Alchourrón, Gärdenfors, and Makinson 1985). Belief revision is intended to capture changes in belief state reflecting new information about a *static* world. In contrast, belief update is intended to capture changes of belief in response to a changing world. An agent's beliefs may be correct at one time. But as the world changes, for example, other agents take acts and disrupt their environment, certain facts become true and others false. The agent must accommodate these changes to update its state of beliefs. Such a process is called *belief update* (Katsuno and Mendelzon 1991). Besides the traditional symbolic formalism, probability theory can be used to represent an agent's belief state. In the probabilistic setting, an agent's cognitive state is represented by a probability function p over a set Ω of possible worlds. *Conditioning* and *imaging* are two probabilistic versions of belief change that correspond to belief revision and belief update, respectively. Upon learning a sure fact C, on the one hand, conditioning works by suppressing the possible worlds which are inconsistent with C and normalizing the probabilities of the remaining possible worlds. It is a fundamental approach in probabilistic reasoning and statistical inference (Pearl 1988). On the other hand, imaging (or updating) performs by transferring the probabilities of worlds outside C to the closest worlds in C. It is a common method to study intervention and causality (Pearl 2009).

In this paper, we take the two views of belief revision and update to study *constrained differential privacy*. Differential privacy (DP) is a mathematically rigorous definition of privacy which addresses the paradox of learning nothing about an individual while learning useful information about a population (Dwork et al. 2006; Dwork and Roth 2014). Differentially private data releases are often required to satisfy a set of external constraints that reflect the legal, ethical, and logical mandates to which the data curator is obligated (Abowd et al. 2019; Hay et al. 2010). For example, in US Census 2020, the so-called touchstone of DP by Dwork, the Census Bureau is constitutionally mandated to report the total population of each state as *exactly* enumerated without subjecting them to any perturbation protection; in data queries, constraints are often used to improve the accuracy while maintaining the quality of privacy protection of the unconstrained DP mechanisms. The central question in designing DP mechanisms with those constraints (called *constrained DP*) is how to integrate randomized DP privacy mechanisms with *deterministic* constraints while maintaining the standard trade-off between privacy protection and data utility. Our main contribution is to study this integration from the perspectives of belief revision and update.

In this paper, we mainly focus on those constraints that are known to hold also for the original datasets, which are hence called *invariants*. We first give a definition of *data-independent* invariants C (Definition 8), which is a subset of the output space $\mathbb{R}^n$ of the privacy mechanism M. In this paper, Laplace and Gaussian mechanisms are considered. For a given dataset D, $M(D)$ is a (continuous) random vector over $\mathbb{R}^n$. Let $P_{M(D)}$ and $p_{M(D)}$ denote the corresponding probability distribution and (density) function, which is regarded as an agent's belief state. We then design constrained DP mechanisms by performing belief change on $p_{M(D)}$ (or

[*] Equal Contribution.
[†] Corresponding Author.

$P_{M(D)}$). In order to revise $p_{M(D)}$ by conditioning, we have to consider two cases: $P_{M(D)}(C) > 0$ and $P_{M(D)}(C) = 0$. When $P_{M(D)}(C) > 0$, conditioning works as usual. One of our technical contributions is to deal with the challenge when $P_{M(D)}(C) = 0$. We use the techniques of changing variables in multivariate calculus to compute the conditional density. Even though the invariant C is data-independent, conditioning is data-dependent because the denominator in conditional density function depends on the original data D. So conditioning may add more privacy loss as shown in (Gong and Meng 2020). In this paper, we show that, if M is additive and C is represented by a group of linear equalities, conditioning does not incur any extra privacy loss (Lemma 10). In addition to the standard postprocessing and composition properties (Lemmas 11,13), we obtain from the characterizing property of conditioning an interesting form of composition: conditional privacy mechanism on the disjoint union of two invariants is a convex combination of the mechanisms conditioned on individual invariants (Theorem 12). From the perspective of belief update, we perform imaging on $p_{M(D)}$ and show the standard postprocessing and composition properties (Lemmas 15,17). In contrast, imaging is a privacy-preserving postprocessing and does not incur any privacy loss (Lemma 14). Moreover, we obtain a characteristic proposition about composition and show that imaging on mixture of privacy mechanisms is the mixture of imaging privacy mechanisms (Theorem 16). In addition to the above analysis of privacy, we also perform analysis of utility and perform some experiments to show the differences between the two perspectives as belief revision and update. The theory of belief revision and update can guide us in choosing the appropriate constraining approach, thereby clarifying many confusions in the literature regarding these two constrained DPs.

The paper is organized as follows. We first present some background about belief change and differential privacy. Then we provide a detailed analysis of the two views. At the end of the paper, we discuss some related works and future research. The following Figure 1 provides the guide of the paper.

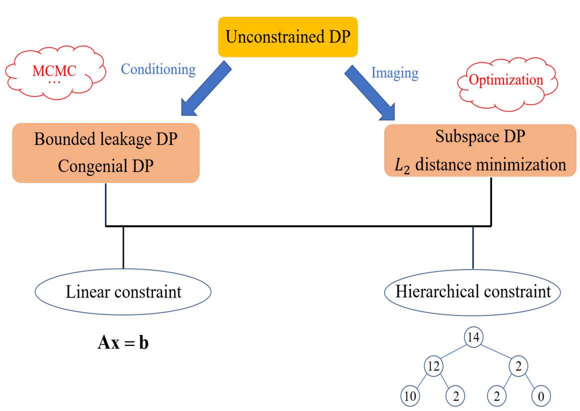

Figure 1. Guideline of the Paper

Preliminaries

In this paper, we are mainly concerned about belief revision and update in the probabilistic setting. So, in this section, we provide some basic knowledge. For details about probabilistic belief revision and update, one may refer to (Gärdenfors 1988) and (Dubois and Prade 1993). In the probabilistic framework, a belief state is represented by a probability measure P (or a probability function p) on the set Ω of possible worlds. In order to motivate the views of constrained differential privacy from the perspective of belief change, we first assume that Ω is finite. If we learn that event E has occurred, i.e., we are certain that E is true, the prior state P is revised according to Bayesian *conditioning*: for any $E' \subseteq \Omega$, $P(E'|E) = \frac{P(E \cap E')}{P(E)}$. It is not well-defined when $P(E) = 0$. Bayesian conditioning is the probabilistic counterpart of belief revision for a *static* world. There is a well-known characterization of Bayesian conditioning: there is no *relative change of beliefs* in the process of conditioning.

Proposition 1 *(Proposition 3.2.1. in (Halpern 2017)) Let $P(E) > 0$. A probability measure P' on Ω is obtained from P according to Bayesian conditioning if and only if P' satisfies the following two conditions:*

1. $P'(\bar{E}) = 0$;
2. $\frac{P'(B)}{P'(B')} = \frac{P(B)}{P(B')}$ *for any $B, B' \subseteq E$ such that $P(B') > 0$.*

There is another more interesting characterization of Bayesian conditioning. It is shown (Gärdenfors 1988) that, for C and C' such that $C \cap C' = \emptyset$, conditioning satisfies the property that $P(\cdot|C \cup C')$ is a convex combination of $P(\cdot|C)$ and $P(\cdot|C')$.

Proposition 2 *For the above defined Bayesian conditioning, and $C, C' \subseteq \Omega$ such that $C \cap C' = \emptyset$, $P(B|C \cup C') = \lambda P(B|C) + (1-\lambda)P(B|C')$ with $\lambda = \frac{P(C)}{P(C)+P(C')}$.*

Another important characterization of Bayesian conditioning is about the minimal change principle through the Kullback-Leibler information distance: for any probability measures P and P' on Ω, $I(P, P') := \sum_{\omega \in \Omega} P'(\omega) \log \frac{P'(\omega)}{P(\omega)}$. The conditional probability $P(\cdot|C)$ minimizes the KL distance $I(P, P')$ from the prior P under the constraint that $P'(C) = 1$.

Now we describe *belief update* in the probabilistic framework. Assume that, for any event $C \subseteq \Omega$ and for any $\omega \in \Omega$, there is a unique $\omega_C \in C$ which is the closest world from ω. According to probabilistic *imaging* (projection), upon learning that C is true, the probability mass $P(\omega)$ assigned to a world $\omega \in \Omega$ is then transferred to ω_C, the closest world in C. In other words, the updated probability after imaging can be written as

$$P_C(\omega) := \sum_{\omega'_C = \omega} P(\omega') \tag{1}$$

Such a process is called *probabilistic belief update*. There is a nice characterization of probabilistic imaging. It is the only updating rule that is homomorphic. Mathematically, for any two probability measures P and P' on Ω and any $\lambda \in [0, 1]$,

$$(\lambda P + (1 - \lambda)P')_C = \lambda P_C + (1 - \lambda)P'_C. \quad (2)$$

In other words, imaging preserves probabilities under the mixture of probabilities. In particular, imaging may turn an impossible world into a possible one, i.e., $P(\omega_C) = 0$ but $P_C(\omega_C) > 0$. On the other hand, imaging may turn a sure event to be uncertain, i.e., $P(B) = 1$ but $P_C(B) < 1$. However, both cases are impossible in the Bayesian conditioning.

Example 3 Here we adopt an example from Section 2.3 in (Dubois and Prade 1993) to illustrate the difference between belief revision and update. There is either an apple (a) or a banana (b) in a box. Let $\omega_1, \omega_2, \omega_3$ and ω_4 denote all the four possible states where $a \wedge b$ is true, $a \wedge \neg b$ is true, $\neg a \wedge b$ is true and $\neg a \wedge \neg b$ is true, respectively. Our current epistemic state p is represented by $p(\omega_1) = p(\omega_4) = 0$, $p(\omega_2) = 0.7$ and $p(\omega_3) = 0.3$. After learning that there is no apple, i.e., $C = \{\omega_3, \omega_4\}$, the epistemic state p changes according to Bayesian conditioning to $p(\omega_3|C) = 1$ and $p(\omega_1|C) = p(\omega_2|C) = p(\omega_4|C) = 0$. In other words, we infer that there is a banana in the box. Next we consider imaging or belief update. In C, ω_3 is the closest world to ω_1 and ω_4 is the closest to ω_2. So $p_C(\omega_3) = p(\omega_1) + p(\omega_3) = 0.3$ and $p_C(\omega_4) = p(\omega_2) + p(\omega_4) = 0.7$. This implies that it is more probable that the box is empty. In belief revision, C is interpreted as "there is no apple in the box" (static world), while, according to belief update, it means "there is no longer any apple" (world change).

Let $\mathcal{T}$ be a set of possible *records*. Typically we use t to denote records (or data). A *dataset* is a finite indexed family of records. We use $\mathcal{D}$ to denote the space of all possible datasets. Elements of $\mathcal{D}$ are typically denoted as D or D'. For any $i \leq |D|, D_i$ denotes the i-th record in D and D_{-i} is the dataset D with D_i removed. In other words, $D_{-i} = D \setminus \{D_i\}$. Let Y be the set of *outputs* which are usually denoted by y, y', y_1 or y_2. A *randomized mechanism* $\mathcal{M} : \mathcal{D} \to Y$ maps a dataset D to a random variable $M(D)$ over Y. In other words, for any $D \in \mathcal{D}$ and $E \subseteq Y$, $Pr[\mathcal{M}(D) \in E]$ defines a probability measure over Y. *Differential privacy* is a privacy guarantee that a randomized algorithm behaves similarly on *neighbouring* input databases which differ on at most one record. The two datasets D and D' can differ in two possible ways: either they have the same size and differ only on one record ($|D| = |D'|$, $D_i \neq D'_i$ and, for any $j \neq i$, $D_j = D'_j$), or one is a copy of the other with one extra record ($D' = D_{-i}$ for some i). These two options do not protect the 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. The original definition of differential privacy in (Dwork et al. 2006) takes the second notion of neighbourhood. In this paper, these two notions apply, which are both denoted $D \sim D'$. Usually we use capital letters to denote random variables and lower-case letters to denote their values.

Definition 4 For arbitrary $\epsilon, \delta > 0$, a randomized mechanism $\mathcal{M}$ is called (ϵ, δ)-*differentially private* if, for any $S \subseteq Y$ and D and D' such that $D \sim D'$, the following inequality hold:

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

If $\delta = 0$, we say that $\mathcal{M}$ is ϵ-differentially private. $\lhd$

In this paper, we assume that the output set $Y = \mathbb{R}^n$. For a function $f : \mathcal{D} \to \mathbb{R}^n$ and $k = 1, 2$, the l_k-sensitivity of f is defined as $\Delta_k(f) = \max_{D \sim D'} \|f(D) - f(D')\|_k$. We mainly consider the following two most commonly-used privacy mechanisms.

Definition 5 (Laplace Mechanism) For a deterministic query function $f : \mathcal{D} \to \mathbb{R}^n$, the Laplace mechanism $M : \mathcal{D} \to \mathbb{R}^n$ is given by $M(D) = f(D) + (U_1, U_2, \cdots, U_n)$ where $U_1, U_2, \cdots, U_n$ are i.i.d. Laplace random variables with the probability density function

$$Lap(x|\lambda) = \frac{1}{2\lambda} \exp\left(-\frac{|x|}{\lambda}\right)$$

We will sometimes write $Lap(\lambda)$ to denote the Laplace distribution with the scale λ, and will sometimes abuse notation and write $Lap(\lambda)$ to denote a random variable $U \sim Lap(\lambda)$. $\lhd$

Definition 6 (Gaussian Mechanism) For a deterministic query function $f : \mathcal{D} \to \mathbb{R}^n$, the Gaussian mechanism $M : \mathcal{D} \to \mathbb{R}^n$ is given by $M(D) = f(D) + (U_1, U_2, \cdots, U_n)$. where $U_1, U_2, \cdots, U_n$ are i.i.d. Gaussian random variables with the probability density function $\mathcal{N}(0, \sigma^2)$. Similarly, we will sometimes write $\mathcal{N}(0, \sigma^2)$ to denote a random variable $U \sim \mathcal{N}(0, \sigma^2)$. $\lhd$

To answer queries under differential privacy, we use the Laplace and Gaussian mechanisms, which achieves differential privacy by adding noise to query answers. If we *calibrate* the Laplace and Gaussian noises to the query f, we can show that the above two mechanisms are differentially private.

Proposition 7 *(Dwork et al. 2006) Let* $\lambda = \frac{\Delta_1(f)}{\epsilon}$ *and* $\delta = \frac{\Delta_2(f)(1+\sqrt{1+\ln(1/\delta)})}{\epsilon}$. *We have*

1. *The above Laplace mechanism with* $Lap(\lambda)$ *is* ϵ*-DP;*
2. *The above Gaussian mechanism* $\mathcal{N}(0, \sigma^2)$ *is* (ϵ, δ)*-DP.*

In this paper, the random variables associated with privacy mechanisms are usually continuous. Density functions determine continuous distributions. If a continuous distribution is calculated conditionally on some information, then the density is called a *conditional density* (Applebaum 1996). When the conditioning information involves another random variable with a continuous distribution, the conditional density can be calculated from the joint density for the two random variables. Suppose that two random variables have a joint continuous probability distribution with joint density function $p_{X,Y}(x, y)$ and $p_Y(y)$ is the density function of Y, then the conditional density of the distribution of the random variable X for fixed values y of Y is defined as follows:

$$p_X(x|Y=y) = \frac{p_{X,Y}(x,y)}{p_Y(y)}. \tag{3}$$

It is easy to see that, in this case, conditioning does not change the relative densities.

Two Views of Constrained DP

From the data curator's perspective, in addition to privacy concerns, there often exists external constraints that the privatized output M must meet. These constraints can often be represented as a predicate of $M(D)$ that agrees with what is calculated based on the confidential $f(D)$.

Definition 8 Given a deterministic query $f : \mathcal{D} \to \mathbb{R}^n$ and a privacy mechanism $M : \mathcal{D} \to \mathbb{R}^n$, we call a *convex* (and hence Lebesgue measurable) subset $C \subseteq \mathbb{R}^n$ an *invariant* if, for any $D \in \mathcal{D}$, $M(D) \in C \Leftrightarrow f(D) \in C$ with probability one over the randomness of M. $\lhd$

Our definition of invariant is *independent* of the original dataset D and hence is essentially different from that in (Gong and Meng 2020). The invariants defined there depend on the original dataset. We have not found in the literature yet any practical scenarios with such a dependent invariant. Usually the invariants are represented by a group of linear equalities or inequalities. In other words, $C = \{\mathbf{z} \in \mathbb{R}^n : A\mathbf{z} = \mathbf{b}\}$ or $C = \{\mathbf{z} \in \mathbb{R}^n : A\mathbf{z} \geq \mathbf{b}\}$ for some matrix A and vector $\mathbf{b}$. *Constrained DP* (CDP for short) refers to differential privacy (or differential private mechanism) satisfying some invariant.

Constrained DP as Belief Revision

For a given dataset $D \in \mathbb{R}^n$, $M(D)$ is a random variable. Let $P_{M(D)}$ and $p_{M(D)}$ be the corresponding probability distribution and density function over $\mathbb{R}^n$. Now we give a definition of *constrained DP* by employing the technique of belief revision to the probability distribution and density function associated with the random variable $M(D)$. Given an invariant C, we construct the conditional random variable $M(D)|C$ in two cases.

1. Case 1: $P_{M(D)}(C) > 0$. For example, if C is represented by a group of linear inequalities, then usually $P_{M(D)}(C) > 0$. Now we define the conditional random variable $M(D)|C$ (and its probability density function $p_{M(D)|C}$). If $u \in C$, then $p_{M(D)}(u) = \frac{p_{M(D)(u)}}{P_{M(D)}(C)}$; otherwise, $p_{M(D)|C}(u) = 0$.
2. Case 2: $P_{M(D)}(C) = 0$. Here we consider as an illustration a simple case when M is additive and the invariant C is defined by a group of linear inequalities $A\mathbf{z} = \mathbf{b}$ where A is a $(n' \times n)$ matrix $(n' < n)$ and $\mathbf{b}$ is a $(n' \times 1)$ column vector. It follows $P_{M(D)}(C) = 0$. Let $M(D) = f(D) + U$ where f is a deterministic query and U is a random vector $(U_1, U_2, \cdots, U_n)$ with probability density function p_U. Since $Af(D) = AM(D)$ with probability 1, $AU = 0$ with probability 1. So, in this case, the randomness in $M(D)$ comes from the random vector U and hence is independent of the original dataset D. Without loss of generality, we assume that the rank of A is n', i.e., A is of full rank and, by solving the group of linear equations $AU = 0$ of unknowns $U_1, U_1, \cdots, U_n$, we get

$$\begin{cases} U_{n'+1} = U_{n'+1}(U_1, U_2, \cdots, U_{n'}) \\ U_{n'+2} = U_{n'+2}(U_1, U_2, \cdots, U_{n'}) \\ \cdots \\ U_n = U_n(U_1, U_2, \cdots, U_{n'}) \end{cases}$$

In other words, $U_1, \cdots, U_{n'}$ are the n' free variables. Now we define the conditional random variable $M(D)|C$ and its probability density function $p_{M(D)|C}$. If $M(D) = f(D) + u = f(D) + (u_1, \cdots, u_n) \in C$, then $p_{M(D)|C}(f(D) + (u_1, \cdots, u_n)) = p_U(u) = \frac{p_U(u_1,\cdots,u_{n'},U_{n'+1}(u_1,\cdots,u_{n'}),\cdots,U_n(u_1,\cdots,u_{n'}))}{\int_{\mathbb{R}^{n'}} p_U(u_1,\cdots,u_{n'},U_{n'+1}(u_1,\cdots,u_{n'}),\cdots,U_n(u_1,\cdots,u_{n'}))du_1\cdots d_{u_{n'}}}$ (let K_C denotes the denominator); if $M(D) = f(D) + u \notin C$, then $p_{M(D)|C}(f(D) + u) = 0$. In summary, if $(v_1, v_2, \cdots, v_n) \in C$, then $p_{M(D)|C}(v_1, v_2, \cdots, v_n) = \frac{p_{M(D)}(v_1,\cdots,v_n)}{K_C}$; if $(v_1, \cdots, v_n) \notin C$, then $p_{M(D)|C}(v_1, \cdots, v_n) = 0$. Note that K_C depends only on C and the noise-adding random vector $(U_1, \cdots, U_n)$.

Definition 9 For a privacy mechanism $M : \mathcal{D} \to \mathbb{R}^n$ and an invariant $C \subseteq \mathbb{R}^n$, we define the *constrained privacy mechanism* $M(\cdot|C)$ satisfying the invariant C as belief revision according to probabilistic conditioning as follows:

$$M(\cdot|C)(D) := M(D)|C \tag{4}$$

For short, we call $M(\cdot|C)$ a *conditional privacy mechanism* on the invariant C.

$\lhd$

Congenial DP under mandated disclosure considered in (Gong and Meng 2020) is our conditional DP for the first case, i.e., $M(D)(C) > 0$. In this case, define $c_D = Pr[M(D) \in C]$ and $c_{D'} = Pr[M(D') \in C]$. Set $\gamma = \frac{1}{\epsilon} \max_{D \sim D'} \log \frac{c_D}{c_{D'}}$. From a similar argument to Theorem 2.1 in (Gong and Meng 2020), we know that, in this case, if M is ϵ-differentially private, then $M(\cdot|C)$ is $(1+\gamma)\epsilon$-differentially private for some $\gamma \in [-1, 1]$. So the conditioning may incur an additional privacy loss with a factor γ. The following proposition shows a similar proposition for the second case when $M(D)(C) = 0$ but with $\gamma = 0$.

Lemma 10 *Let M be additive and C be represented by a group of linear equalities $A\mathbf{z} = \mathbf{b}$ as above. For an invariant $C \subseteq \mathbb{R}^n$, if $M : \mathcal{D} \to \mathbb{R}^n$ is ϵ-differentially private, then $M(\cdot|C)$ is also ϵ-differentially private.*

Proof. Assume that M is additive and C is defined by $A\mathbf{z} = \mathbf{b}$. For any two neighbouring datasets D and D', since M is ϵ-differentially private, $e^{-\epsilon} \leq \frac{p_{M(D)}(v_1,\cdots,v_n)}{p_{M(D')}(v_1,\cdots,v_n)} \leq e^{\epsilon}$ for any $(v_1, \cdots, v_n) \in \mathbb{R}^n$. For any $(v_1, \cdots, v_n) \in C$, $p_{M(D)|C}(v_1, \cdots, v_n) = \frac{p_{M(D)}(v_1,\cdots,v_n)}{K_C}$

and $p_{M(D')|C}(v_1, \cdots, v_n) = \frac{p_{M(D')}(v_1,\cdots,v_n)}{K_C}$. So, it follows that, for any $(v_1, \cdots, v_n) \in C$, $e^{-\epsilon} \leq \frac{p_{M(D)|C}(v_1,\cdots,v_n)}{p_{M(D')|C}(v_1,\cdots,v_n)} \leq e^{\epsilon}$. This implies that $e^{-\epsilon} \leq \frac{Pr[M(D|C) \in E]}{Pr[M(D'|C) \in E]} \leq e^{\epsilon}$ for any $E \subseteq C$. We have shown that $M(\cdot|C)$ is ϵ-differentially private.

QED

Lemma 11 *Let $M(\cdot|C)$ be a (ϵ, δ) conditional differential private for some privacy mechanism $M : \mathcal{D} \to \mathbb{R}^n$ and invariant C. If $h : \mathbb{R}^n \to \mathbb{R}^{n'}$ is measurable, then*

1. *$h \circ M(\cdot|C)$ is also (ϵ, δ)- differentially private.*
2. *$(h \circ M)(\cdot|C)$ is also (ϵ, δ)-differentially private, and $h \circ (M(\cdot|C)) = (h \circ M)(\cdot|C)$.*

Proof. The proof of the first part follows from the observation that $h \circ M(D|C) \in Z$ iff $M(D|C) \in h^{-1}(Z)$ for any measurable $Z \subseteq \mathbb{R}^{n'}$. And the second part follows from the fact

$$\frac{Pr[(h \circ M)(\cdot|C)(D) \in B]}{Pr[(h \circ M)(\cdot|C)(D') \in B]} = \frac{Pr[M(\cdot|C))(D) \in h^{-1}(B)]}{Pr[M(\cdot|C)(D') \in h^{-1}(B)]}$$

QED

The following two propositions are about composition. The first one comes from the characteristic property in Lemma 2 about conditioning. It is a *new* composition property. It tells us that, constrained DP mechanism by conditioning on the disjoint union of two invariants can be obtained by the convex combination of constrained DP on these two individual invariants.

Theorem 12 *(Disjoint-union Composition) Let C and C' be two invariants such that $C \cap C' = \emptyset$. We have*

1. *The conditional privacy mechanism on $C \cup C'$ is a convex combination of the conditional privacy mechanisms on C and on C', i.e., for any $D \in \mathcal{D}$, $Pr[M(D|C \cup C') \in B] = \lambda Pr[M(D|C) \in B] + (1 - \lambda) Pr[M(D|C') \in B]$ for some $\lambda \in [0, 1]$;*
2. *If the conditional privacy mechanism $M(\cdot|C)$ and $M(\cdot|C')$ are both (ϵ, δ)-differential private, then $M(\cdot|C \cup C')$ is (ϵ, δ)-differentially private.*

Proof. The first part is from Lemma 2 with $\lambda = \frac{P_{M(D)}(C)}{P_{M(D)}(C) + P_{M(D)}(C')}$ and the second follows from the first part. Note that λ here depends on D and hence is data-dependent. QED

Lemma 13 *Given $\epsilon_1 \geq 0$ and $\epsilon_2 > 0$, if $M_1(\cdot|C_1)$ is ϵ_1-differentially private and $M_2(\cdot|C_2)$ is ϵ_2-differentially private, then $(M_1, M_2)(\cdot|C_{12})$ such that $C_{12} = (C_1, C_2)$ is $(\epsilon_1 + \epsilon_2)$-differentially private.*

The probability distribution $P' = P_{M(D|C)}$ of the conditional privacy mechanism minimizes the KL-divergence $I(P', P_{M(D)})$ with the requirement that $P'(C) = 1$.

Constrained DP as Belief Update

Now we define how to update a privacy mechanism $M : \mathcal{D} \to \mathbb{R}^n$ with an invariant C according to the following probabilistic imaging rule of belief update:

$$\bar{\mathbf{y}} = \arg\min_{\mathbf{y} \in C} \|\mathbf{y} - \tilde{\mathbf{y}}\|_2 \tag{5}$$

where $\tilde{\mathbf{y}}$ denotes the noisy output of the mechanism M, i.e., $M(D) = \tilde{\mathbf{y}}$. So $\bar{\mathbf{y}}$ is the "closest world" in the invariant C from the noisy $\tilde{\mathbf{y}}$. Let f_{L_2} denote the deterministic function of *postprocessing with L_2 minimization*, i.e., $f_{L_2}(\tilde{\mathbf{y}}) = \bar{\mathbf{y}}$. Let $p_{M(D)}$ and $p_{f_{L_2} M(D)}$ denote the corresponding probability function of the two random vectors $M(D)$ and $f_{L_2}(M(D))$, respectively. For $\mathbf{y} \in \mathbb{R}^n$,

$$p_{f_{L_2} M(D)}(\mathbf{y}) = \int_{f_{L_2}(\mathbf{y}') = \mathbf{y}} p_{M(D)}(\mathbf{y}') d\mathbf{y}' \tag{6}$$

From the above Eq. (6), we see that $p_{f_{L_2} M(D)}$ is obtained from $p_{M(D)}$ according to probabilistic imaging in the sense of Eq. (1). Let $P_{M(D)}$ and $P_{f_{L_2}(M(D))}$ denote the corresponding probability measures of $p_{M(D)}$ and $p_{f_{L_2}(M(D))}$, respectively. In other words, $P_{M(D)}(B) = Pr[M(D) \in B]$ and $P_{f_{L_2}(M(D))}(B) = Pr[f_{L_2} M(D) \in B]$. From Eq, (6), we have $P_{f_{L_2}(M(D))} = (P_{M(D)})_C$. Let M_C denote the corresponding privacy mechanism updated according to the invariant C, i.e., $M_C(D) = f_{L_2}(M(D))$.

Generally, when M is (ϵ, δ)-differentially private, $f_{L_2} \circ M$ is not necessarily (ϵ, δ)-differentially private (Gong and Meng 2020). However, in this paper, we consider only *data-independent* invariants (Definition (8)). So, f_{L_2} as belief update does preserve privacy. Indeed, for any $B \subseteq \mathbb{R}^n$,

$$\begin{aligned} \frac{Pr[M_C(D) \in B]}{Pr[M_C(D') \in B]} &= \frac{Pr[M_C(D) \in B \cap C]}{Pr[M_C(D') \in B \cap C]} \\ &= \frac{Pr[M(D) \in f_{L_2}^{-1}(B \cap C)]}{Pr[M(D') \in f_{L_2}^{-1}(B \cap C)]} \end{aligned}$$

So, if M is (ϵ, δ)-DP, then so is M_C. In particular, when M is an additive privacy mechanism and the invariant can be represented by a group of linear equalities, f_{L_2} as postprocessing preserves privacy.

Lemma 14 *If M is (ϵ, δ)-diferentially private, then M_C is also (ϵ, δ)-differentially private.*

Lemma 15 *(Postprocessing) If M_C is (ϵ, δ)-differentially private, then, for any measuable function h from $\mathbb{R}^n$ to $\mathbb{R}^{n'}$, $(h \circ M)_C$ is also (ϵ, δ)-differentially private.*

Proof. Note that the randomness comes not from h but from M. So the lemma follows from the observation:

$$\frac{Pr[(h \circ M)_C(D) \in B]}{Pr[(h \circ M)_C(D') \in B]} = \frac{Pr[M_C(D) \in h^{-1}(B)]}{Pr[M_C(D') \in h^{-1}(B)]}$$

QED

The following composition property is *new* which says that the constrained differential privacy as belief update is also preserved under convex combination. A convex combination of the privacy mechanisms is such a randomized mechanism M that it outputs the mechanism M_i with probability a_i where $\sum_i a_i = 1$. Formally, suppose that we have some privacy mechanisms $M_i : \mathcal{D} \to Y (1 \le i \le n)$. Now we define a new privacy mechanism M such that $Pr[M = M_i] = a_i \ge 0 (1 \le i \le n)$ and $a_1 + a_2 + \cdots + a_n = 1$.

Theorem 16 *(Mixture Composition) Let $a_1, a_2, \cdots, a_n \ge 0$ such that $a_1 + a_2 + \cdots, a_n = 1$ and M be the above defined mixed mechanism. We have that*

1. $P_{M_C}(B) = a_1 P_{(M_1)_C}(B) + a_2 P_{(M_2)_C}(B) + \cdots + a_n P_{(M_n)_C}(B)$ *for any* $B \subseteq Y$;
2. *if* $(M_1)_C, (M_2)_C, \cdots, (M_n)_C$ *are* (ϵ, δ)*-differentially private, then* M_C *is also* (ϵ, δ)*-differentially private.*

Proof. The proof of the first part follows directly from the characterizing Eq. (2) for probabilistic update. And the second part is immediate from the first. QED

The theorem implies that, if we want to find the imaging of the mixture of some privacy mechanisms, we can find the imaging of those privacy mechanisms first, which may be much easier, and then get their mixture. It is interesting to note that constrained DP as belief update satisfies the Axiom of Choice in (Kifer and Lin 2010).

Lemma 17 *(Basic Composition) Given $\epsilon_1 > 0$ and $\epsilon_2 > 0$, if M_1 is ϵ_1-differentially private with invariant set C_1 and M_2 is ϵ_2-differentially private with invariant set C_2, then the joint mechanism (M_1, M_2) is $(\epsilon_1 + \epsilon_2)$-differentially private for the invariant set $C_{12} := (C_1, C_2)$.*

Now we describe a mechanism design for imaging. As usual, let $M(D) = f(D) + (U_1, \cdots, U_n)$. Then $M_C(D) = f(D) + \Pi_C(U_1, \cdots, U_n)$ where $\Pi_C(U_1, \cdots, U_n)$ is a random vector with the following density function: for $(u_1', \cdots, u_n') \in C$, $p_{\Pi_C(U_1,\cdots,U_n)}(u_1', \cdots, u_n') = \int_{S(u_1',\cdots,u_n')} p_{(U_1,\cdots,U_n)}(u_1, u_2, \cdots, u_n) du_1 \cdots du_n$ where $S(u_1', \cdots, u_n') = \{(u_1, \cdots, u_n) : \Pi_C(u_1, \cdots, u_n) = (u_1', \cdots, u_n')\}$. This mechanism design covers the projection mechanism for subspace DP (Gao, Gong, and Yu 2022).

In most cases, conditioning and imaging are two different methods to achieve the constraint differential privacy. But, in the special case when the mechanism is spherical Gaussian mechanism (Definition 6) and invariant is represented by linear-equality constraint, then these two methods achieve the same results (Gao, Gong, and Yu 2022).

Utility Analysis

In this section, we analyze the utility in two scenarios: one is a simple linear constraint; the other is hierarchical constraint. It seems that, in both scenarios, conditioning achieves better in utility than imaging. Here we consider the Laplace mechanism in Definition 5. We first compare the utilities of conditioning and imaging on the linear constraint $C = \{(z_1, \cdots, z_n) \in \mathbb{R}^n : z_1 + z_2 + \cdots + z_n = b\}$ for some real number b. The following proposition (Theorem 12 from (Zhu, Van Hentenryck, and Fioretto 2021)) characterizes the variance of the marginal distribution of the postprocessed noise $f_{L_2}(\tilde{\mathbf{x}}) - \mathbf{x}$.

Proposition 18 $Var(f_{L_2}(\tilde{\mathbf{x}}) - \mathbf{x})_i = 2\lambda^2(1 - \frac{1}{n})$ *for* $i = 1, \cdots, n$.

Here we consider the simplest case where $n = 3$ and the invariant is $C = \{(z_1, z_2, z_3) \in \mathbb{R}^3 : z_1 + z_2 + z_3 = b\}$ for some constant b. The Laplace mechanism $M(D) = f(D) + (U_1, U_2, U_3)$ where U_1, U_2 and U_3 are identically independent Laplace random variables. Under the constraint $z_1 + z_2 + z_3 = b$, it is easy to see that $U_1 + U_2 + U_3 = 0$. Let $p_{(U_1,U_2,U_3)}$ denote the probability density function (p.d.f) of the random vector (U_1, U_2, U_3). It follows that $p_{U_1,U_2,U_3}(u_1, u_2, u_3) = (\frac{1}{2\lambda})^3 \exp(-\frac{|u_1|+|u_2|+|u_3|}{\lambda})$. Let (U_1^*, U_2^*, U_3^*) denote the conditional random vector $(U_1, U_2, U_3)|(U_1+U_2+U_3 = 0)$. So, if $M(D) = f(D)+(u_1, u_2, u_3)$ such that $u_1+u_2+u_3 = 0$, then the probability density of $M(D)|C$ at (u_1, u_2, u_3) is $\frac{p_{U_1,U_2,U_3}(u_1,u_2,-(u_1+u_2))}{\int_{\mathbb{R}^2} p_{U_1,U_2,U_3}(u_1,u_2,-(u_1+u_2))du_1du_2}$; if $u_1 + u_2 + u_3 \neq 0$, then the probability mass of $M(D)|C$ at (u_1, u_2, u_3) is 0. So the conditional privacy mechanism $M(\cdot|C)(D) = f(D) + (U_1, U_2, U_3)|C$. We formulate the probability density of $M(D)|C$ at (u_1, u_2, u_3), defined as $h(u_1, u_2)$

$$h(u_1, u_2) = \frac{\exp\left(-\frac{|u_1|+|u_2|+|u_1+u_2|}{\lambda}\right)}{\iint\limits_{u_1,u_2} \exp\left(-\frac{|u_1|+|u_2|+|u_1+u_2|}{\lambda}\right) du_1 du_2}$$

Define $h(u_1) = \int\limits_{u_2} h(u_1, u_2) du_2$, the marginal variance of u_1 is then given by $Var(u_1) = \int\limits_{u_1} u_1^2 h(u_1) du_1$. Now, consider computing $h(u_1, u_2)$, $h(u_1)$ and $Var(u_1)$. The denominator of $h(u_1, u_2)$ is a constant $K = \iint\limits_{u_1,u_2} \exp\left(-\frac{|u_1|+|u_2|+|u_1+u_2|}{\lambda}\right) du_1 du_2$. Since the integration region is symmetric about the origin and $p(u_1, u_2, -u_1 - u_2) = p(-u_1, -u_2, u_1 + u_2)$, we get $K = \frac{3}{2}\lambda^2$.To compute $h(u_1)$, we firstly consider the case $u_1 \ge 0$

$$\begin{aligned} h(u_1) &= \tfrac{1}{K} \int_{-\infty}^{+\infty} \exp\left(-\tfrac{|u_1|+|u_2|+|u_1+u_2|}{\lambda}\right) du_2 \\ &= \tfrac{1}{K} \left((\lambda + u_1) \exp\left(-\tfrac{2u_1}{\lambda}\right)\right) \end{aligned}$$

Similar derivation can be performed on case $u_1 < 0$. Thus for any u_1, it follows that $h(u_1) = \frac{1}{K}\left((\lambda + |u_1|) \exp\left(-\frac{2|u_1|}{\lambda}\right)\right)$. At last, we get marginal variance of u_1, $Var(u_1) = \int\limits_{u_1} u_1^2 h(u_1) du_1 = \frac{1}{K} \int_{-\infty}^{+\infty} u_1^2 \left((\lambda + |u_1|) \exp\left(-\frac{2|u_1|}{\lambda}\right)\right) du_1 = \frac{5}{6}\lambda^2$. So when $n = 3$, the variance $Var(u_1)$ of the marginal distribution of conditioning $M(D)$ is smaller than the above variance $Var(f_{L_2}(\tilde{\mathbf{x}}) - \mathbf{x})$ of marginal distribution by

CDP	Privacy preserving	Privacy preserving (linear eq.)	Post-processing	Basic composition	Disjoint union composition	Mixture composition	Minimum principle	Approximate computation
Belief revision	✘	✔	✔	✔	✔	✘	KL-divergence	Monte Carlo
Belief update	✔	✔	✔	✔	✘	✔	L_2-distance L_1-distance	Optimization

Table 1. Summary of Main Results

imaging (Proposition 18). This is also true for the case when $n = 2$ (Example 4.1 in (Gong and Meng 2020)). We performed some simulation experiments for larger n which showed similar results. We conjecture that this holds generally for any n and hence the variance of the marginal distribution by conditioning is smaller than that by imaging on this invariant for simple counting query.

MCMC Method and Comparative Experiment

In this part, we experimentally compare accuracy between conditioning approach and imaging approach in processing data based on region hierarchy. In the experiment, we choose the improved MCMC method to obtain samples of the consistency constraint privacy mechanism, and compare it with the classic post-processing projection technique such as TopDown algorithm. We choose New York City Taxi Dataset for the experiment. The specific selection is the yellow taxi trip dataset in February 2022. The relevant document is called "yellow_tripdata_2022-02.parquet" while records all trip data of the iconic yellow taxi in New York City in February 2022. The dataset has 19 attribute columns, 2979431 record rows, where each row represents a taxi trip. We only use one attribute "PULocationID" in this experiment, which ranging from 1 to 263, indicates TLC Taxi Zone in which the taximeter was engaged. We treat each taxi as a group and build a 3-level hierarchy of trip record frequency in each zone. New York city, abbreviated as NYK, is at Level 1, six boroughs, i.e., Bronx (Bx), Brooklyn (Bl), EWR, Manhattan (M), Queens (Q) and Staten Island (SI), is at Level 2 and Level 3 includes 263 zones corresponding to "PULocationID". Here we provide an improved metropolis Hastings (MH) algorithm M_{MH}. Our experiment shows the advantage in accuracy by comparing the conditioning algorithm with the imaging algorithm. Trip frequency distribution in all zones is taken as the confidential query $\mathbf{x}$, and the Laplace mechanism is selected to perturb $\mathbf{x}$. Finally, the output $\tilde{\mathbf{x}}$ satisfying the differential privacy and consistency constraints is obtained. In this experiment we select L_1- distance between $\mathbf{x}$ and $\tilde{\mathbf{x}}$ as the performance evaluation criteria. For comparison, we normalized the L_1-distance. i.e., $\frac{1}{m}|\mathrm{x} - \tilde{\mathrm{x}}|$, where m is the dimension of $\mathbf{x}$ and $\tilde{\mathbf{x}}$. Algorithm's running efficiency at different levels of privacy budget is shown in Table 2.

Through the comparison of the two algorithms under different privacy budget conditions and different hierarchy levels, it can be seen that in most cases, the conditioning algorithm M_{MH} will be more accurate than the classic imaging or projection algorithms. And since the noise decreases as the privacy budget increases, the errors decrease too.

ϵ	Level	M_{MH}	TopDown
0.5	1	0.013352	0.036806
	2	0.028890	0.162698
	3	1.680823	2.345461
1	1	0.018244	0.023974
	2	0.057345	0.091148
	3	1.534053	1.526361
2	1	0.003445	0.005173
	2	0.015032	0.027614
	3	1.052862	1.260267

Table 2. Accuracy Comparison of Algorithms Running on NY City Taxi Dataset at L_1-distance

Related Works and Conclusion

The main contributions and the comparisons between these two approaches are summarized in Table 1 (✔ there means "true" and ✘ "not necessarily true"). Belief change may explain why almost all constrained DP mechanisms in the literature are essentially classified either as belief revision or as belief update. There is a long tradition of designing constrained DP by imaging in the database-system community (Zhang, Xiao, and Xie 2016; Hay et al. 2010; Wang et al. 2020; Lee, Wang, and Kifer 2015; Zhu, Van Hentenryck, and Fioretto 2021; Gao, Gong, and Yu 2022) and later in US Census (Abowd et al. 2022). Constrained DP as belief revision has appeared quite recently and mainly from the statistics community. For example, congenial DP in (Gong and Meng 2020; Gong 2022) and bounded leakage DP (Ligett, Peale, and Reingold 2020) are essentially as belief revision. None of these papers relates their ideas to the notions of belief revision and update. With this connection, we contribute two interesting *new* theorems about constrained DP (Theorems 12 and 16). We expect to obtain more important new properties about constrained DP from the well-established perspectives of belief revision and update. Also we will consider the models of "screened revision" and "credibility-limited revision" in the full version. Constrained DP by conditioning on invariant C can be regarded as a special bounded leakage DP when the invariant C can be represented by $[M'(D) = o]$ for some o and some randomized algorithm M'. We may consider to extend constrained DP as belief update to a similar more general setting. Conditioning and imaging are two important approaches in statistical and causal inference. It may be an interesting research topic to explore the relationships between constrained DP and causality (Tschantz, Sen, and Datta 2020).

Acknowledgments

This work is partially supported by by NSFC (61732006) and Public Computing Cloud, Renmin University of China. YF is supported by the National Key R&D Program of China (2018YFA0306704) and the Australian Research Council (DP180100691 and DP220102059). We would like to thank Yuxiao Guo for discussion about some techniques about change of variables.

References

Abowd, J.; Ashmead, R.; Cumings-Menon, R.; Garfinkel, S.; Heineck, M.; Heiss, C.; Johns, R.; Kifer, D.; Leclerc, P.; Machanavajjhala, A.; Moran, B.; Sexton, W.; Spence, M.; and Zhuravlev, P. 2022. The 2020 Census Disclosure Avoidance System TopDown Algorithm. *Harvard Data Science Review*, (Special Issue 2). Https://hdsr.mitpress.mit.edu/pub/7evz361i.

Abowd, J.; Ashmead, R.; Simson, G.; Kifer, D.; Leclerc, P.; Machanavajjhala, A.; and Sexton, W. 2019. Census topdown: Differentially private data, incremental schemas, and consistency with public knowledge. *Tech. Rep., US Census Bureau*.

Alchourrón, C. E.; Gärdenfors, P.; and Makinson, D. 1985. On the Logic of Theory Change: Partial Meet Contraction and Revision Functions. *J. Symb. Log.*, 50(2): 510–530.

Applebaum, D. 1996. *Probability and information: An integrated approach*. Cambridge University Press.

Dubois, D.; and Prade, H. 1993. Belief Revision and Updates in Numerical Formalisms: An Overview, with new Results for the Possibilistic Framework. In Bajcsy, R., ed., *Proceedings of the 13th International Joint Conference on Artificial Intelligence. Chambéry, France, August 28 - September 3, 1993*, 620–625. Morgan Kaufmann. ISBN 1-55860-300-X.

Dwork, C.; McSherry, F.; Nissim, K.; and Smith, A. D. 2006. Calibrating Noise to Sensitivity in Private Data Analysis. In Halevi, S.; and Rabin, T., eds., *Theory of Cryptography, Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006, Proceedings*, volume 3876 of *Lecture Notes in Computer Science*, 265–284. Springer.

Dwork, C.; and Roth, A. 2014. The Algorithmic Foundations of Differential Privacy. *Found. Trends Theor. Comput. Sci.*, 9(3-4): 211–407.

Gao, J.; Gong, R.; and Yu, F.-Y. 2022. Subspace differential privacy. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 3986–3995.

Gärdenfors, P. 1988. *Knowledge in flux: Modeling the dynamics of epistemic states*. The MIT press.

Gong, R. 2022. Transparent Privacy is Principled Privacy. *Harvard Data Science Review*, (Special Issue 2). Https://hdsr.mitpress.mit.edu/pub/ld4smnnf.

Gong, R.; and Meng, X. 2020. Congenial Differential Privacy under Mandated Disclosure. In *FODS '20: ACM-IMS Foundations of Data Science Conference*, 59–70. ACM.

Halpern, J. Y. 2017. *Reasoning about uncertainty*. MIT Press.

Hay, M.; Rastogi, V.; Miklau, G.; and Suciu, D. 2010. Boosting the Accuracy of Differentially Private Histograms Through Consistency. *Proceedings of the VLDB Endowment*, 3(1).

Katsuno, H.; and Mendelzon, A. O. 1991. On the Difference between Updating a Knowledge Base and Revising It. In Allen, J. F.; Fikes, R.; and Sandewall, E., eds., *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning (KR'91). Cambridge, MA, USA, April 22-25, 1991*, 387–394. Morgan Kaufmann. ISBN 1-55860-165-1.

Kifer, D.; and Lin, B.-R. 2010. Towards an axiomatization of statistical privacy and utility. In *Proceedings of the twenty-ninth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, 147–158.

Lee, J.; Wang, Y.; and Kifer, D. 2015. Maximum likelihood postprocessing for differential privacy under consistency constraints. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 635–644.

Ligett, K.; Peale, C.; and Reingold, O. 2020. Bounded-Leakage Differential Privacy. In Roth, A., ed., *1st Symposium on Foundations of Responsible Computing, FORC 2020, June 1-3, 2020, Harvard University, Cambridge, MA, USA (virtual conference)*, volume 156 of *LIPIcs*, 10:1–10:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.

Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems - Networks of Plausible Inference*. Morgan Kaufmann series in representation and reasoning. Morgan Kaufmann.

Pearl, J. 2009. *Causality*. Cambridge university press.

Tschantz, M. C.; Sen, S.; and Datta, A. 2020. SoK: Differential privacy as a causal property. In *2020 IEEE Symposium on Security and Privacy (SP)*, 354–371. IEEE.

van Harmelen, F.; Lifschitz, V.; and Porter, B. W., eds. 2008. *Handbook of Knowledge Representation*, volume 3 of *Foundations of Artificial Intelligence*. Elsevier. ISBN 978-0-444-52211-5.

Wang, T.; Lopuhaa-Zwakenberg, M.; Li, Z.; Skoric, B.; and Li, N. 2020. Locally Differentially Private Frequency Estimation with Consistency. In *NDSS'20: Proceedings of the NDSS Symposium*.

Zhang, J.; Xiao, X.; and Xie, X. 2016. Privtree: A differentially private algorithm for hierarchical decompositions. In *Proceedings of the 2016 International Conference on Management of Data*, 155–170.

Zhu, K.; Van Hentenryck, P.; and Fioretto, F. 2021. Bias and variance of post-processing in differential privacy. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 11177–11184.