
Context-Aware Local Differential Privacy

Jayadev Acharya¹ K. A. Bonawitz² Peter Kairouz² Daniel Ramage² Ziteng Sun^{1,3}

Abstract

Local differential privacy (LDP) is a strong notion of privacy that often leads to a significant drop in utility. The original definition of LDP assumes that all the elements in the data domain are equally sensitive. However, in many real-life applications, some elements are more sensitive than others. We propose a context-aware framework for LDP that allows the privacy level to vary across the data domain, enabling system designers to place privacy constraints where they matter without paying the cost where they do not. For binary data domains, we provide a universally optimal privatization scheme and highlight its connections to Warner’s randomized response and Mangat’s improved response. Motivated by geo-location and web search applications, for k -ary data domains, we consider two special cases of context-aware LDP: block-structured LDP and high-low LDP. We study minimax discrete distribution estimation under both cases and provide communication-efficient, sample-optimal schemes and information theoretic lower bounds. We show, using worst-case analyses and experiments on Gowalla’s 3.6 million check-ins to 43,750 locations, that context-aware LDP achieves a far better accuracy under the same number of samples.

1. Introduction

Differential privacy (DP) is a rigorous notion of privacy that enforces a worst-case bound on the privacy loss due to release of query results (Dwork et al., 2006). Its local

version, local differential privacy (LDP) (Definition 1), provides context-free privacy guarantees even in the absence of a trusted data collector (Warner, 1965; Beimel et al., 2008; Kasiviswanathan et al., 2011). Under LDP, all pairs of elements in a data domain are assumed to be equally sensitive, leading to harsh privacy-utility trade-offs in many learning applications. In many real-life settings, however, some elements are more sensitive than others. For example, in URL applications, users may want to differentiate sensitive URLs from non-sensitive ones, and in geo-location applications, users may want to hide their precise location within a city, but not the city itself.

This work introduces a unifying context-aware notion of LDP where different pairs of domain elements can have “arbitrary” sensitivity levels. For binary domains, we provide a universal optimality result and highlight interesting connections to Warner’s response and Mangat’s improved response. For k -ary domains, we look at two canonical examples of context-aware LDP: block-structured LDP and high-low LDP. For block-structured LDP, the domain is partitioned into m blocks and the goal is to hide the identity of elements within the same block but not the block identity. This is motivated by geo-location applications where users can be in different cities, and it is not essential to hide the city of a person but rather where exactly within that city a person is (i.e., which bars, restaurants, or businesses they visit). In other words, in this context, users would like to hide which element in a given block of the domain their data belongs to – and not necessarily which block their data is in, which may be known from side information or application context. For high-low LDP, we assume there is a set of sensitive elements and we only require the information about possessing sensitive elements to be protected. This can be applied in web browsing services, where there are a large number of URLs and not all of them contain sensitive information.

1.1. Related work

Recent works consider relaxations of DP (Chatzikokolakis et al., 2013; Doudalis et al., 2017; Borgs et al., 2018; Asi et al., 2019; Feyisetan et al., 2019; Dong et al., 2019) and LDP (Chatzikokolakis et al., 2017; Alvim et al., 2018; Murakami & Kawamoto, 2018; Kawamoto & Murakami, 2019; Geumlek & Chaudhuri, 2019) that incorporate the context and structure of the underlying data domain. Indeed,

¹ECE, Cornell University, Ithaca, New York, USA ²Google, Seattle, USA ³JA and ZS are supported by NSF-CCF-1846300 (CAREER), NSF-CRII-1657471 and a Google Faculty Research Award. Work partially done while ZS was visiting Google. Correspondence to: Jayadev Acharya <acharya@cornell.edu>, K. A. Bonawitz <bonawitz@google.com>, Peter Kairouz <kairouz@google.com>, Daniel Ramage <dramage@google.com>, Ziteng Sun <zs335@cornell.edu>.

(Doudalis et al., 2017; Murakami & Kawamoto, 2018) investigate settings where a small fraction of the domain elements are secrets and require that the adversary’s knowledge about whether or not a user holds a secret cannot increase much upon the release of data. (Chatzikokolakis et al., 2013; Borgs et al., 2018; Chatzikokolakis et al., 2017; Alvim et al., 2018; Xiang et al., 2019) model the data domain as a metric space and scale the privacy parameter between pairs of elements by their distance. (Schein et al., 2019) considers categorical high dimensional count models and define privacy measures that help privatize counts that are close in ℓ_1 distance.

To investigate the utility gains under this model of privacy, we consider the canonical task of distribution estimation: Estimate an unknown discrete distribution given (privatized) samples from it. The trade-off between utility and privacy in distribution estimation has received recent attention, and optimal rates have been established (Duchi et al., 2013; Dikonikolas et al., 2015; Kairouz et al., 2016; Wang et al., 2016; Ye & Barg, 2018; Acharya et al., 2019c). These recent works show that the sample complexity for k -ary distribution estimation up to accuracy α in total variation distance increases from $\Theta\left(\frac{k}{\alpha^2}\right)$ (without privacy constraints) to $\Theta\left(\frac{k^2}{\alpha^2\varepsilon^2}\right)$ (for $\varepsilon = O(1)$) when we impose ε -LDP constraints. For both the block-structured and high-low models of LDP, we will characterize the precise sample complexity by providing privatization and estimation schemes that are both computation- and communication-efficient.

Another line of work considers a slightly different problem of heavy hitter estimation where there is no distributional assumption under the data samples that the users have and the main focus is on reducing the computational complexity and communication requirements (Bassily & Smith, 2015; Bassily et al., 2017; Hsu et al., 2012; Wang & Blocki, 2017; Bun et al., 2018; Zhu et al., 2019; Acharya & Sun, 2019).

1.2. Our contributions

Motivated by the limitations of LDP, and practical applications where not all domain elements are equally sensitive, we propose a general notion of *context-aware* local differential privacy (see Definition 2) and prove that it satisfies several important properties, such as preservation under post-processing, adaptive composition, and robustness to auxiliary information.

When the underlying data domain is binary, we provide a complete characterization of a universally optimal scheme, which interpolates between Warner’s randomized response and Mangat’s improved response, given in Theorem 2.

We then consider general data domains and investigate two practically relevant models of context-aware LDP. The first is the *block-structured* model and is motivated by applications in geo-location and census data collection. Under

this privacy model, we assume that the underlying data domain is divided into a number of blocks, and the elements within a block are sensitive. For example, when the underlying domain is a set of geographical locations, each block can correspond to the locations within a city. In this case, we would like to privatize the precise value within a particular block, but the privacy of the block ID is not of great concern (Definition 4). In Theorem 4 we characterize the sample complexity of estimating k -ary discrete distributions under this privacy model. We also propose a privatization scheme based on the recently proposed Hadamard Response (HR), which is both computation- and communication-efficient. We then prove the optimality of these bounds by proving a matching information-theoretic lower bound. This is achieved by casting the problem as an information constrained setting and invoking the lower bounds from (Acharya et al., 2019b;a). We show that when all the blocks have roughly the same number of symbols, the sample complexity can be reduced by a factor of the number of blocks compared to that under the classic LDP. See Theorem 4 for a formal statement of this result.

The second model we consider is the *high-low* model, where there are a few domain elements that are sensitive while the others are not. A form of this high-low notion of privacy was proposed in (Murakami & Kawamoto, 2018), which also considered distribution estimation and proposed an algorithm based on RAPPOR (Erlingsson et al., 2014). We propose a new privatization scheme based on HR that has the same sample complexity as (Murakami & Kawamoto, 2018) with a much smaller communication budget. We also prove a matching information theoretic lower bound showing the optimality of these schemes (Murakami & Kawamoto, 2018). In this case, the sample complexity only depends quadratically in the number of sensitive elements and linearly in the domain size compared to the quadratic dependence on the domain size under the classic LDP. See Theorem 3 for a formal statement of this result.

As a consequence of these results, we observe that the sample complexity of distribution estimation can be significantly less than that in the classical LDP. Thus contextual privacy should be viewed as one possible method to improve the trade-off between the utility and privacy in LDP when different domain elements have different levels of privacy.

To validate our worst-case (minimax) analyses, we conduct experiments on synthetic data sampled from uniform, geometric, and power law distributions in addition to experiments on over 3.6 million check-ins to 43,750 locations from the Gowalla dataset. Our experiments confirm that context-aware LDP achieves a far better accuracy under the same number of samples.

1.3. Organization

In Section 2 we define LDP and the problem of distribution estimation. In Section 3 we define context-aware LDP and provide an operational definition along with some of its specializations. In Section 4 we provide the optimal privatization scheme for binary domains. In Section 5 and Section 6 we derive the optimal sample complexity in the high-low model, and block-structured model respectively. Experimental results are presented in Section 7. We conclude with a few interesting extensions in Section 8.

2. Preliminaries

Let \mathcal{X} be the k -ary underlying data domain, wlog let $\mathcal{X} = [k] := \{1, \dots, k\}$. There are n users, and user i has a (potentially) sensitive data point $X_i \in \mathcal{X}$. A privatization scheme is a mechanism to add noise to mask the true X_i 's, and can be represented by a conditional distribution $Q^1 : \mathcal{X} \rightarrow \mathcal{Y}$, where \mathcal{Y} denotes the output domain of the privatization scheme. For $y \in \mathcal{Y}$, and $x \in \mathcal{X}$, $Q(y | x)$ is the probability that the privatization scheme outputs y , upon observing x . If we let X and Y denote the input and output of a privatization scheme, then $Q(y | x) = \Pr(Y = y | X = x)$. For any set $S \subset \mathcal{Y}$, we have $Q(S | x) := \Pr(Y \in S | X = x)$.

Definition 1 (Local Differential Privacy. (Warner, 1965; Kasiviswanathan et al., 2011)). Let $\varepsilon > 0$. A conditional distribution Q is ε -locally differentially private (ε -LDP) if for all $x, x' \in \mathcal{X}$ and all $S \subset \mathcal{Y}$,

$$Q(S | x) \leq e^\varepsilon Q(S | x'). \quad (1)$$

Let \mathcal{Q}_ε be the set of all ε -LDP conditional distributions with input $[k]$.

Distribution estimation. Let

$$\Delta_k := \{(p_1, \dots, p_k) : \forall x \in [k], p_x \geq 0; \sum p_x = 1\}$$

be all discrete distributions over $[k]$. We assume that the user's data $X^n := X_1, \dots, X_n$ are independent draws from an unknown $p \in \Delta_k$. User i passes X_i through the privatization channel Q and sends the output Y_i to a central server. Upon observing the messages $Y^n := Y_1, \dots, Y_n$, the server then outputs \hat{p} as an estimate of p . Our goal is to select Q from a set of *allowed channels* \mathcal{Q} and to design an estimation scheme $\hat{p} : \mathcal{Y}^n \rightarrow \Delta_k$ that achieves the following min-max risk

$$r(k, n, d, \mathcal{Q}) = \min_{Q \in \mathcal{Q}} \left\{ \min_{\hat{p}} \max_{p \in \Delta_k} \mathbb{E}[d(p, \hat{p})] \right\}, \quad (2)$$

where $d(\cdot, \cdot)$ is a measure of distance between distributions. In this paper we consider the total variation distance,

¹In the remaining parts of this paper, we use terms privatization scheme, conditional distribution, and channel to refer to this Q matrix interchangeably.

$d_{TV}(p, q) := \frac{1}{2} \sum_{i=1}^k |p_i - q_i|$. For a parameter $\alpha > 0$, the sample complexity of distribution estimation to accuracy α is the smallest number of users for the min-max risk to be smaller than α ,

$$n(k, \alpha, \mathcal{Q}) := \arg \min_n \{r(k, n, d_{TV}, \mathcal{Q}) < \alpha\}.$$

When \mathcal{Q} is \mathcal{Q}_ε , the channels satisfying ε -LDP constraints with input domain $[k]$, it is now well established that for $\varepsilon = O(1)$ (Duchi et al., 2013; Kairouz et al., 2016; Ye & Barg, 2018; Wang et al., 2016; Acharya et al., 2019b;c),

$$n(k, \alpha, \mathcal{Q}_\varepsilon) = \Theta\left(\frac{k^2}{\alpha^2 \varepsilon^2}\right). \quad (3)$$

Hadamard matrix (Sylvester's construction). Let $H_1 = [1]$, then Sylvester's construction of Hadamard matrices is a sequence of square matrices of size $2^i \times 2^i$ recursively defined as

$$H_{2m} = \begin{bmatrix} H_m & H_m \\ H_m & -H_m \end{bmatrix}$$

Letting $S_i = \{y \mid y \in [m], H(i+1, y) = +1\}$ be the column indices of +1's in the $(i+1)$ th row of H_m , we have

- $\forall i \in [m-1], |S_i| = \frac{m}{2}$.
- $\forall i \neq j \in [m-1], |S_i \cap S_j| = \frac{m}{4}$

3. Context-Aware LDP

In local differential privacy, all elements in the data domain are assumed to be equally sensitive, and the same privacy constraint is enforced on all pairs of them (see (1)). However, in many settings some domain elements might be more sensitive than others. To capture this, we present a context-aware notion of privacy. Let $E \in \mathbb{R}_{\geq 0}^{k \times k}$ be a matrix of non-negative entries, where for $x, x' \in [k]$, $\varepsilon_{x, x'}$ is the (x, x') th entry of E .

Definition 2 (Context-Aware LDP). A conditional distribution Q is E -LDP if for all $x, x' \in \mathcal{X}$ and $S \subseteq \mathcal{Y}$,

$$Q(S | x) \leq e^{\varepsilon_{x, x'}} Q(S | x'). \quad (4)$$

This definition allows us to have a different sensitivity level between each pair of elements to incorporate context information. For a *privacy matrix* E , the set of all E -LDP mechanisms is denoted by \mathcal{Q}_E . When all the entries of E are ε , we obtain the classical LDP.

Symmetric E matrices. When the E matrix is symmetric, context-aware LDP introduces a similar structure as Metric-based LDP (Alvim et al., 2018), which considers \mathcal{X} to be a metric space endowed with a distance $d_{\mathcal{X}}$. In this case, consider

$$\varepsilon_{x, x'} = \varepsilon_{x', x} = d_{\mathcal{X}}(x, x').$$

It is required that it is harder to distinguish close-by symbols compared to symbols that are relatively far from each other.

As a special case, we introduce (later in this section) the notion of Block-Structured LDP as an example which only requires the elements to be indistinguishable if they are in the same block.

Asymmetric E matrices. An advantage of our framework is that it allows the matrix E to be asymmetric. Consider a binary setting where we ask a user whether or not they have used drugs before. Here, people whose answer is **no** can occasionally say **yes** to protect those whose answer is **yes**. Thus, when the data collector sees a **yes**, they are lost as to whether it came from a true **yes** answer or a **no** that was flipped to a **yes**. Further, there is no need to randomize **yes** answers because it is okay for the data collector to know who did not do drugs. This is captured in our framework by allowing $Q(\mathbf{no} \mid \mathbf{no})/Q(\mathbf{no} \mid \mathbf{yes})$ to be arbitrarily large and placing an upper bound on $Q(\mathbf{yes} \mid \mathbf{yes})/Q(\mathbf{yes} \mid \mathbf{no})$. A well-known scheme that satisfies this requirement is Mangat’s improved randomized response (Mangat, 1994). See Section 4 for more details. This intuition is generalized to k -ary alphabet in the High-Low LDP model introduced later in this section.

We provide an operational (hypothesis testing) interpretation for context-aware LDP in Section 3.1. We show that context-aware LDP is robust to post-processing, robust to auxiliary information, and is adaptively composable in Section 3.2.

For binary domains, namely when $k = 2$, we give a single optimal privatization scheme for all E matrices in Section 4. For general k , it is unclear whether there is a simple characterization of the optimal schemes for all E matrices. First note that if $\varepsilon^* = \min_{i,j} \varepsilon_{i,j}$ is the smallest entry of E , then any ε^* -LDP algorithm is also E -LDP. However, this is not helpful since it does not help us capture the context of the application and consequently get rid of the stringent requirements of standard LDP. Motivated by applications in geo-location and web search, we consider structured E matrices that are both practically motivating, and are parameterized by a few parameters (or a single parameter) so that they are easier to analyze and implement.

High-Low LDP (HLLDP). The high-low model captures applications where there are certain domain elements that are private, and the remaining elements are non-private. We only want to protect the privacy of private elements. This is formalized below.

Definition 3. Let $A = \{x_1, \dots, x_s\} \subset \mathcal{X}$ denote the set of sensitive domain elements, and all symbols in $B := \mathcal{X} \setminus A$ are non-sensitive. A privatization scheme is said to be (A, ε) -HLLDP if $\forall S \subset \mathcal{Y}$, and $x \in A, x' \in \mathcal{X}$,

$$Q(S \mid x) \leq e^\varepsilon Q(S \mid x'), \quad (5)$$

which corresponds to the following E matrix:

$$\varepsilon_{x,x'} = \begin{cases} \varepsilon, & x \in A, \\ \infty, & x \in B. \end{cases}$$

This implies that when the input symbol is in A , the output distribution cannot be multiplicatively larger than the output distribution for any other symbol, but there is no such restriction for symbols in B . HLLDP was also defined in (Murakami & Kawamoto, 2018). We solve the problem of minimax distribution estimation under this privacy model in Section 5.

Block-Structured LDP (BSLDP). In applications such as geo-location, it is important to preserve the privacy of symbols that are close to each other. We consider this model where the data domain is divided into various blocks (e.g., the cities), and we would like the symbols within a block (e.g., the various locations within a city) to be hard to distinguish. This is formalized below.

Definition 4. Suppose there is a partition of \mathcal{X} , which is $\Pi = \{\mathcal{X}_1, \mathcal{X}_2, \dots, \mathcal{X}_m\}$. With a slight abuse of notation, we define $\forall x \in \mathcal{X}_i, \Pi(x) := i$. Then $\forall \varepsilon > 0$, a privatization scheme is said to be (Π, ε) -BSLDP if it satisfies $\forall x, x' \in \mathcal{X}_i$ such that $\Pi(x) = \Pi(x')$, and any $S \subset \mathcal{Y}$, we have

$$Q(S \mid x) \leq e^\varepsilon Q(S \mid x'), \quad (6)$$

which corresponds to the following E matrix:

$$\varepsilon_{x,x'} = \begin{cases} \varepsilon, & \Pi(x) = \Pi(x'), \\ \infty, & \Pi(x) \neq \Pi(x'). \end{cases}$$

This definition relaxes the local version of differential privacy in the following way. Given a partition of the input set $\Pi = \{\mathcal{X}_1, \mathcal{X}_2, \dots, \mathcal{X}_m\}$, it requires different levels of indistinguishability for element pairs in the same block and those in different blocks. We solve the problem of minimax distribution estimation under this model in Section 6.

3.1. Operational interpretation of context-aware LDP

Recall that the entries of E denote the amount of privacy across the corresponding row and column symbols. We provide an operational interpretation of E -LDP by considering a natural hypothesis testing problem. Suppose we are promised that the input is in $\{x, x'\}$ for some symbols $x, x' \in [k]$, and an E -LDP scheme outputs a symbol $Y \in \mathcal{Y}$. Given Y , the goal is to test whether the input is x or x' .

Theorem 1 (Operational Interpretation of Context-Aware LDP). *A conditional distribution Q is E -locally differentially private if and only if for all $x, x' \in \mathcal{X}$ and all decision rules $\hat{X} : \mathcal{Y} \rightarrow \{x, x'\}$,*

$$P_{\text{FA}}(x, x') + e^{\varepsilon_{x',x}} P_{\text{MD}}(x, x') \geq 1, \quad (7)$$

$$e^{\varepsilon_{x,x'}} P_{\text{FA}}(x, x') + P_{\text{MD}}(x, x') \geq 1, \quad (8)$$




Figure 1. Error region for hypothesis testing between x and x' under DP constraints when $\varepsilon_{x,x'} < \varepsilon_{x',x}$. (Blue for Context Aware LDP, Red for standard LDP)

where $P_{FA}(x, x') = \Pr(\hat{X} = x \mid X = x')$ is the false alarm rate and $P_{MD}(x, x') = \Pr(\hat{X} = x' \mid X = x)$ is the miss detection rate.

The proof is provided in Section A.4. Consider a test for distinguishing x and x' given above in Theorem 1. Figure 1 shows the effective error regions for any estimator \hat{X} under the privacy constraints $\varepsilon_{x,x'}$ and $\varepsilon_{x',x}$. We can see that unlike the symmetric region under LDP, we are pushing the miss detection rate to be higher when $\varepsilon_{x,x'} < \varepsilon_{x',x}$. This shows that symbol x is more private than symbol x' , namely we want to protect the identity of symbol being x more than we want to protect the identity being x' .

3.2. Properties of context-aware LDP

Context-aware LDP also satisfies several important properties held by classical LDP, including post-processing, adaptive composition and robustness to auxiliary information. We provide their proofs in Section A.

Proposition 1 (Preservation under post-processing). *Let $\mathcal{A}_1 : \mathcal{X} \rightarrow \mathcal{Y}_1$ be an E -LDP scheme and $\mathcal{A}_2 : \mathcal{Y}_1 \rightarrow \mathcal{Y}_2$ is any algorithm that processes the output of \mathcal{A}_1 , then the scheme $\mathcal{A} = \mathcal{A}_2 \circ \mathcal{A}_1$ is also E -LDP.*

Proposition 2 (Adaptive composition). *Let $\mathcal{A}_1 : \mathcal{X} \rightarrow \mathcal{Y}_1$ be an E_1 -LDP scheme and $\mathcal{A}_2 : \mathcal{X} \times \mathcal{Y}_1 \rightarrow \mathcal{Y}_2$ be an E_2 -LDP scheme, then the scheme \mathcal{A} defined as $(\mathcal{A}_1, \mathcal{A}_2)$ is $(E_1 + E_2)$ -LDP.*

Proposition 3 (Robustness to auxiliary information). *Let p^* be a prior we have over \mathcal{X} and $\mathcal{A} : \mathcal{X} \rightarrow \mathcal{Y}$ be an E -LDP scheme. Then $\forall x_1, x_2 \in \mathcal{X}$ and $y \in \mathcal{Y}$,*

$$\frac{\Pr(X = x_1 \mid Y = y)}{\Pr(X = x_2 \mid Y = y)} \leq e^{\varepsilon_{x_1, x_2}} \frac{p^*(x_1)}{p^*(x_2)}.$$

4. Binary Domains

Consider a binary domain, namely $k = 2$ and domain elements $\{1, 2\}$. In this case, we have

$$E = \begin{bmatrix} 0 & \varepsilon_{1,2} \\ \varepsilon_{2,1} & 0 \end{bmatrix}.$$

Perhaps the oldest and simplest privatization mechanism is Warner's randomized response for confidential survey interviews (Warner, 1965). In this section, we give the optimal scheme for all utility functions that obey the *data processing inequality* under all possible binary constraints. We prove that when $\varepsilon_{1,2} = \varepsilon_{2,1}$, the optimal scheme is Warner's randomized response.

Consider the composition of two privatization mechanisms $Q \circ W$ where the output of the first mechanism Q is applied to another mechanism W . We say that a utility function $U(\cdot)$ obeys the data processing inequality if for all Q and W

$$U(Q \circ W) \leq U(Q).$$

In other words, further processing of the data can only reduce the utility. Such utility functions are ubiquitous. For example, in the minimax distribution learning context of this paper, $U(Q)$ may be chosen as $-\min_{\hat{p}} \max_p \mathbb{E}[d(p, \hat{p})]$ (i.e., the negative of the minimax risk under a fixed mechanism Q) with d being any ℓ_p distance or f -divergence.

Theorem 2 (Optimal Mechanism for Binary Domains). *Let \mathcal{X} be a binary input alphabet and $U(Q)$ be any utility function that obeys the data processing inequality. Then for any $\varepsilon_{1,2}, \varepsilon_{2,1} \geq 0$, the following privatization mechanism*

$$Q^* = \frac{\begin{bmatrix} e^{\varepsilon_{2,1}} - 1 & 1 - e^{-\varepsilon_{1,2}} \\ e^{-\varepsilon_{1,2}}(e^{\varepsilon_{2,1}} - 1) & e^{\varepsilon_{2,1}}(1 - e^{-\varepsilon_{1,2}}) \end{bmatrix}}{e^{\varepsilon_{2,1}} - e^{-\varepsilon_{1,2}}} \quad (9)$$

solves

$$\underset{Q}{\text{maximize}} \quad U(Q) \quad \text{subject to} \quad Q \in \mathcal{Q}_E.$$

Here $\forall x, y \in \{1, 2\}$, $Q^*(x, y) := Q^*(y \mid x)$.

As a special case of the above theorem, if we consider the original local differential privacy setup where $\varepsilon_{1,2} = \varepsilon_{2,1} = \varepsilon$, then the optimal mechanism for binary alphabets is

$$Q^* = \frac{1}{e^\varepsilon + 1} \begin{bmatrix} e^\varepsilon & 1 \\ 1 & e^\varepsilon \end{bmatrix}. \quad (10)$$

This is Warner's randomized response model in confidential structured survey interview with $p = e^\varepsilon / (e^\varepsilon + 1)$ (Warner, 1965). Warner's randomized response was shown to be optimal for binary alphabets in (Kairouz et al., 2014). Another interesting special case of the above theorem is Mangat's

improved randomized response strategy (Mangat, 1994). To see this, let $\varepsilon_{1,2} = \infty$ and $p = e^{-\varepsilon_{2,1}}$. Then

$$Q^* = \begin{bmatrix} 1-p & p \\ 0 & 1 \end{bmatrix}. \quad (11)$$

This is exactly Mangat's improved randomized response strategy. Thus Mangat's randomized response with $p = e^{-\varepsilon_{2,1}}$ is optimal for all utility functions obeying the data processing inequality under this generalized differential privacy framework with $\varepsilon_{1,2} = \infty$, which corresponds to the case where element 1 is not sensitive.

5. Distribution Estimation under HLLDP

We characterize the optimal sample complexity of distribution estimation under the high-low model (see Definition 3).

Theorem 3. *Let $A \subset \mathcal{X}$, with $|A| = s < k/2$, and $\varepsilon = O(1)$. Let $\mathcal{Q}_{A,\varepsilon}$ be the set of all possible channels satisfying (A, ε) -HLLDP, then:*

$$n(k, \alpha, \mathcal{Q}_{A,\varepsilon}) = \Theta\left(\frac{s^2}{\alpha^2 \varepsilon^2} + \frac{k}{\alpha^2 \varepsilon}\right). \quad (12)$$

When the size of the sensitive set is relatively large, e.g. $s > \sqrt{k}$, the sample complexity is $\Theta(s^2/(\alpha^2 \varepsilon^2))$, which corresponds to classic LDP with alphabet size s . This question was considered in (Murakami & Kawamoto, 2018), which gave an algorithm based on RAPPOR (Erlingsson et al., 2014) that has the optimal sample complexity, but requires $\Omega(k)$ bits of communication from each user, which is prohibitive in settings where the uplink capacity is limited for users. We design a scheme based on Hadamard Response (HR), which is also sample-optimal but requires only $\log k$ bits of communication from each user.

(Murakami & Kawamoto, 2018) noted that a lower bound of $\frac{s^2}{\alpha^2 \varepsilon^2}$ (the first term in (12)) is immediately implied by previously known results on distribution estimation under standard LDP ((3) for $k = s$). However, obtaining a lower bound equalling the second term was still open. Using the recently proposed technique in (Acharya et al., 2019b), we prove this lower bound in Section 5.2.

5.1. Achievability using a variant of HR

We propose an algorithm based on Hadamard response (Acharya et al., 2019c), which gives us a tight upper bound for k -ary distribution estimation under (A, ε) -high-low LDP. Let $s = |A|$ and S be the smallest power of 2 larger than s , i.e., $S := 2^{\lceil \log(s+1) \rceil}$. Let $t := k - s$ be the number of all non-sensitive elements. Then we have $S + t \leq 2(s + t) = 2k$. Let H_S be the $S \times S$ Hadamard matrix using Sylvester's construction. Define the output alphabet to be $[S + t] = \{1, \dots, S + t\}$. then the channel is

defined as the following: When $x \in A = [s]$, we have

$$Q(y | x) = \begin{cases} \frac{2e^\varepsilon}{S(e^\varepsilon+1)} & \text{if } y \in [S] \text{ s.t. } H_S(x, y) = 1, \\ \frac{2}{S(e^\varepsilon+1)} & \text{if } y \in [S] \text{ s.t. } H_S(x, y) = -1, \\ 0, & \text{if } y \notin [S]. \end{cases} \quad (13)$$

Else if $x \notin A$, we have

$$Q(y | x) = \begin{cases} \frac{2}{S(e^\varepsilon+1)}, & \text{if } y \in [S], \\ \frac{e^\varepsilon-1}{e^\varepsilon+1}, & \text{if } y = x + S - s, \\ 0, & \text{otherwise.} \end{cases} \quad (14)$$

It is easy to verify that this scheme satisfies (A, ε) -high-low LDP. Next we construct estimators based on the observations Y_1, Y_2, \dots, Y_n and states its performance in Proposition 4. For all $i \in [s]$,

$$\hat{p}_i = \frac{2(e^\varepsilon + 1)}{e^\varepsilon - 1} \left(\widehat{p(S_i)} - \frac{1}{e^\varepsilon + 1} \right) - \widehat{p(A)}, \quad (15)$$

where

$$\widehat{p(A)} = \frac{e^\varepsilon + 1}{e^\varepsilon - 1} \left(\frac{1}{n} \sum_{m=1}^n \mathbf{1}\{Y_m \in S\} - \frac{2}{e^\varepsilon + 1} \right),$$

$$\widehat{p(S_i)} = \frac{1}{n} \left(\sum_{m=1}^n \mathbf{1}\{Y_m \in S_i\} \right).$$

For all $i \notin [s]$, we simply use the empirical estimates

$$\hat{p}_i = \frac{e^\varepsilon + 1}{e^\varepsilon - 1} \frac{1}{n} \left(\sum_{m=1}^n \mathbf{1}\{Y_m = i + S - s\} \right). \quad (16)$$

Proposition 4. *The estimators defined in (15) and (16) satisfy the following:*

$$\mathbb{E} [d_{TV}(\hat{p}, p)] \leq \sqrt{\frac{3s^2}{n} \left(\frac{e^\varepsilon + 1}{e^\varepsilon - 1} \right)^2} + \sqrt{\frac{e^\varepsilon + 1}{e^\varepsilon - 1} \frac{k}{n}}. \quad (17)$$

Since $\frac{e^\varepsilon + 1}{e^\varepsilon - 1} = O\left(\frac{1}{\varepsilon}\right)$ when $\varepsilon = O(1)$, setting the right hand side to be smaller than α gives us the upper bound part of Theorem 3. For the proof of (17), please see Section C.2. In this scheme, each user only needs to communicate at most $\log k + 1$ bits since $S + t \leq 2k$ while the scheme proposed in (Murakami & Kawamoto, 2018) needs $\Omega(k)$ bits.

5.2. Lower bound

We now prove the lower bound part of Theorem 3. A lower bound of $\Omega(s^2/\varepsilon^2 \alpha^2)$ follows from (3), which follows directly from the lower bounds on sample complexity of distribution under standard ε -LDP (e.g., Theorem IV.1 in (Ye & Barg, 2018)).

To prove a lower bound equalling the second term, we use the framework developed in (Acharya et al., 2019b) to prove lower bounds for distributed inference under various information constraints (e.g., privacy and communication constraints). Their key idea is to design a packing of distributions around the uniform distribution and show that the amount of information that can be gleaned from the output of these schemes is insufficient for distribution estimation. In particular, we will use their following result.

Lemma 1. [Lemma 13](Acharya et al., 2019b) Let \mathbf{u} be the uniform distribution over $[k]$ and \mathcal{P} be a family of distributions satisfying the following two conditions.

1. $\forall \mathbf{p} \in \mathcal{P}$, we have $d_{TV}(\mathbf{p}, \mathbf{u}) \geq \alpha$.
2. $\forall \mathbf{p}_1 \in \mathcal{P}$, we have

$$|\{\mathbf{p}_2 \in \mathcal{P} | d_{TV}(\mathbf{p}_1, \mathbf{p}_2) \leq \frac{\alpha}{3}\}| \leq C_\alpha.$$

Suppose \mathcal{Q} is the set of all channels we can use to get information about X , then we have the sample complexity of k -ary distribution learning up to TV distance $\pm\alpha/3$ under channel constraints \mathcal{Q} is at least

$$\Omega\left(\frac{\log |\mathcal{P}| - \log C_\alpha}{\max_{Q \in \mathcal{Q}} \chi^2(Q|\mathcal{P})}\right),$$

where $\mathbf{p}^Q(\mathbf{u}^Q)$ is the distribution of Y when $X \sim \mathbf{p}(\mathbf{u})$, and

$$\chi^2(\mathcal{P}) := \frac{1}{|\mathcal{P}|} \sum_{\mathbf{p} \in \mathcal{P}} d_{\chi^2}(\mathbf{p}^Q, \mathbf{u}^Q),$$

$$d_{\chi^2}(\mathbf{p}, \mathbf{q}) := \sum_{y \in \mathcal{Y}} \frac{(\mathbf{p}(y) - \mathbf{q}(y))^2}{\mathbf{q}(y)}.$$

Let $k' = k - s \geq k/2$ be the number of non-sensitive elements. Let $\mathcal{Z} = \{+1, -1\}^{k'}$ be the set of k' bits. For all $z \in \mathcal{Z}$, define p_z as the following

$$p_z(i) = \begin{cases} \frac{1}{k} + \frac{\alpha \sum_{j=1}^{k'} z_j}{k'} & i = 1, \\ \frac{1}{k} & i = 2, 3, \dots, s, \\ \frac{1}{k} + \frac{\alpha z_{i-s}}{k'} & i = s+1, s+2, \dots, k. \end{cases}$$

Let $\mathcal{P}_{\mathcal{Z}} = \{p_z | z \in \mathcal{Z}\}$ be the set of all distributions defined by $z \in \mathcal{Z}$. Let $U_{\mathcal{Z}}$ be a uniform distribution over set \mathcal{Z} . Then we have

$$\begin{aligned} & \chi^2(Q|\mathcal{P}_{\mathcal{Z}}) \\ &= \mathbb{E}_{z \sim U_{\mathcal{Z}}} \left[\sum_{y \in \mathcal{Y}} \frac{(p_z^Q(y) - u^Q(y))^2}{u^Q(y)} \right] \\ &= \frac{k\alpha^2}{k'^2} \mathbb{E}_{z \sim U_{\mathcal{Z}}} \left[\sum_{y \in \mathcal{Y}} \frac{(\sum_{i=1}^{k'} (Q(y|1) - Q(y|s+i)) z_i)^2}{\sum_{j \in [k]} Q(y|j)} \right] \\ &\leq \frac{4\alpha^2}{k} \sum_{y \in \mathcal{Y}} \frac{\sum_{i=1}^{k'} (Q(y|1) - Q(y|s+i))^2}{\sum_{j \in [k]} Q(y|j)}. \end{aligned}$$

To bound this quantity, we have the following claim, which we prove in Section D.1.

Claim 1. If $\forall Q \in \mathcal{Q}$, we have: $\forall i \in \{s+1, s+2, \dots, s+k'\}, y \in \mathcal{Y}$,

$$Q(y|1) \leq e^\varepsilon Q(y|i),$$

$$\sum_{i=1}^{k'} \sum_{y \in \mathcal{Y}} \frac{(Q(y|s+i) - Q(y|1))^2}{\sum_{j \in [k]} Q(y|j)} = O(\varepsilon k). \quad (18)$$

Moreover, we have $|\mathcal{P}| = 2^{k'}$ and

$$C_\alpha \leq 2^{(1-h(1/3))k'},$$

where $h(x) = x \log(1/x) + (1-x) \log(1/(1-x))$. Combining these results and using Lemma 1, we get the sample complexity is at least

$$\Omega\left(\frac{\log |\mathcal{P}| - \log C_\alpha}{\max_{Q \in \mathcal{Q}} \chi^2(Q|\mathcal{P})}\right) = \Omega\left(\frac{k}{\alpha^2 \varepsilon}\right).$$

6. Distribution Estimation under BSLDP

For distribution estimation under block-structured LDP constraints we prove the following theorem.

Theorem 4. Let $\varepsilon = O(1)$, and $\Pi = \{\mathcal{X}_1, \mathcal{X}_2, \dots, \mathcal{X}_m\}$ be a partition of \mathcal{X} and $\mathcal{Q}_{\Pi, \varepsilon}$ be the set of all possible channels that satisfy (Π, ε) -BSLDP, then

$$n(k, \alpha, \mathcal{Q}_{\Pi, \varepsilon}) = \Theta\left(\frac{\sum_{i=1}^m k_i^2}{\alpha^2 \varepsilon^2}\right),$$

where $\forall i \in [m], |\mathcal{X}_i| = k_i$ and $|\mathcal{X}| = k = \sum_{i=1}^m k_i$.

In the special case when all the blocks have the same size k/m , the sample complexity is $\Theta(k^2/(m\alpha^2\varepsilon^2))$, which saves a factor of m compared classic LDP. In Section 6.1, we describe an algorithm based on HR that achieves this error. Moreover, it only uses $O(\log k)$ bits of communication from each user. A matching lower is proved bound in Section 6.2, which shows that our algorithm is information theoretically optimal.

6.1. Achievability using a variant of HR

The idea of the algorithm is to perform Hadamard Response proposed in (Acharya et al., 2019c) within each block. Without loss of generality we assume each block $\mathcal{X}_j = \{(j, i) | i \in [k_j]\}$. For each block $\mathcal{X}_j, j \in [m]$, we associate a Hadamard matrix H_{K_j} with $K_j = 2^{\lceil \log(k_j+1) \rceil}$. Let $\mathcal{Y}_j = \{(j, i) | i \in [K_j]\}$. For each $x = (j, i) \in \mathcal{X}_j$, we assign the $(i+1)$ th row of H_{K_j} to x . Define the set of locations of '+1's at the $(i+1)$ th row of H_{K_j} to be S_x .

Then the output domain is $\mathcal{Y} := \cup_{j=1}^m \mathcal{Y}_j$. The privatization scheme is given as

$$Q(Y = (j, i) | X) = \begin{cases} \frac{2e^\varepsilon}{K_j(1+e^\varepsilon)}, & X \in \mathcal{X}_j, i \in S_X, \\ \frac{2}{K_j(1+e^\varepsilon)}, & X \in \mathcal{X}_j, i \notin S_X, \\ 0, & \text{elsewhere.} \end{cases}$$

It is easy to verify that this scheme satisfies the privacy constraints. Next we construct estimators based on Y_1, Y_2, \dots, Y_n and state the performance in Proposition 5. Let $Y(1), Y(2)$ be the two coordinates of each output Y . For each $j \in [m]$ and $x \in \mathcal{X}_j$,

$$\hat{p}_x = \frac{2(e^\varepsilon + 1)}{e^\varepsilon - 1} \left(\widehat{p(S_x)} - \frac{\widehat{p(\mathcal{X}_j)}}{2} \right), \quad (19)$$

where

$$\widehat{p(\mathcal{X}_j)} = \frac{1}{n} \sum_{t=1}^n \mathbf{1}\{Y_t(1) = j\},$$

$$\widehat{p(S_x)} = \frac{1}{n} \sum_{t=1}^n \mathbf{1}\{Y_t(1) = j, Y_t(2) \in S_x\}.$$

Proposition 5. *Under the unbiased estimator \hat{p} described in (19), we have:*

$$\mathbb{E} [\ell_2^2(\hat{p}, p)] \leq \frac{12 \max_i k_i}{n} \left(\frac{e^\varepsilon + 1}{e^\varepsilon - 1} \right)^2,$$

$$\mathbb{E} [d_{TV}(\hat{p}, p)] \leq \frac{2(e^\varepsilon + 1)}{e^\varepsilon - 1} \sqrt{\frac{3 \sum_{j=1}^m k_i^2}{n}}. \quad (20)$$

Since $\frac{e^\varepsilon + 1}{e^\varepsilon - 1} = O(\frac{1}{\varepsilon})$ when $\varepsilon = O(1)$, we get desired bounds in Theorem 4. For the proof of (20), see Section C.4. We note here that our algorithm also gives the optimal bound in terms of ℓ_2 distance. A matching lower bound can shown using well established results (Duchi et al., 2013) on LDP by considering the maximum of expected loss if we put all the mass on each single block.

6.2. Lower bound

We now prove the lower bound part of Theorem 4. The general idea is similar to the proof in Section 5.2, which is based on Lemma 1. Without loss of generality, we assume all the k_i 's are even numbers². We construct a family of distributions as following: Let $\mathcal{Z} = \mathcal{Z}_1 \times \mathcal{Z}_2 \times \dots \times \mathcal{Z}_m$ and $\forall j \in [m], \mathcal{Z}_j = \{+1, -1\}^{\frac{k_j}{2}}$. $\forall z \in \mathcal{Z}$, we denote the j th entry of z as z_j where $z_j \in \mathcal{Z}_j$. Define $z_{j,i}$ to be the i th bit of z_j . $\forall j \in [m]$ and $i \in [k_j/2]$, we have

²If one of the k_i 's is odd, we can remove one element from \mathcal{X}_j , which will make the problem simpler and the sample complexity remains unchanged up to constant factors

$$p_z((j, 2i - 1)) = \frac{1}{k} + \frac{2z_{j,i}k_j\alpha}{\sum_{t=1}^m k_t^2},$$

$$p_z((j, 2i)) = \frac{1}{k} - \frac{2z_{j,i}k_j\alpha}{\sum_{t=1}^m k_t^2}.$$

Note that $\forall z \in \mathcal{Z}, d_{TV}(p_z, u) = \alpha$. Moreover, $|\mathcal{P}| = 2^{\frac{k}{2}}$ and $|\mathcal{C}_\alpha| \leq 2^{\frac{k}{2}h(1/3)}$ where h is the binary entropy function. By Lemma 1, let \mathcal{Q} be the set of channels that satisfy $(\mathbb{P}, \varepsilon)$ -LDP and $\mathcal{P}_\mathcal{Z} = \{p_z | z \in \mathcal{Z}\}$, it would suffice to show that

$$\max_{Q \in \mathcal{W}} \chi^2(Q | \mathcal{P}_\mathcal{Z}) = O\left(\frac{k\alpha^2\varepsilon^2}{\sum_{i=1}^m k_i^2}\right). \quad (21)$$

The proof of (21) is technical and presented in Section D.2.

7. Experiments

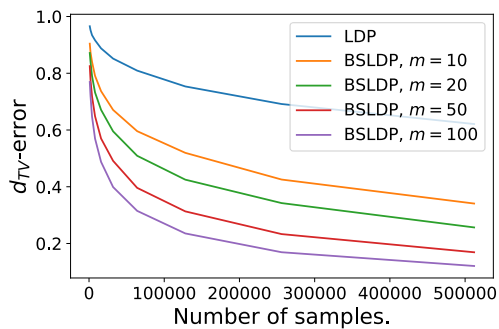
We perform experiments on both synthetic data and real data to empirically validate how the new notion of context-aware LDP and associated algorithms would affect the accuracy of k -ary distribution estimation. Specifically, we choose the special case of block-structured LDP.

For synthetic data, we set $k = 1000, \varepsilon = 1$. We assume all the blocks to have the same size and $m \in \{10, 20, 50, 100\}$. For traditional LDP, we use Hadamard Response (Acharya et al., 2019c), one of the state-of-the-art sample-optimal algorithms. We take $n = 1000 \times 2^i, i \in \{0, 1, \dots, 9\}$ and generate samples from the following three distributions: Geometric distribution $Geo(\lambda)$ where $p(i) \propto (1 - \lambda)^i \lambda$, Zipf distribution $Zipf(\lambda)$ where $p(i) \propto (i + 1)^{-\lambda}$ and uniform distribution. We assume the blocks are partitioned based on their indices (the first block is $\{0, 1, \dots, k/m - 1\}$, the second is $\{k/m, k/m + 1, \dots, 2k/m - 1\}$ and so on). For Geometric distribution, we consider another case where we permute the mass function over the symbols to spread the heavy elements into multiple blocks, denoted by $Geo(\lambda)^*$. The results are shown in Figure 2 and each point is the average of 10 repetitions of the same experiment. We can see that we get consistently better accuracy under the notion of block-structured LDP compared to the classical notion. Moreover, the larger m we have, the better accuracy we get, which is consistent with our analysis.


To validate our algorithm on real datasets, we take the Gowalla user check-in dataset (Cho et al., 2011), which consists of user check-in records with location information (latitudes and longitudes) on the Gowalla website. We take 3671812 check-in records with locations within 25N and 50N for latitude and 130W and 60W for longitude (mostly

(m_1, m_2)	LDP	(5,7)	(25, 35)	(25, 70)
d_{TV} -error	0.591	0.298	0.108	0.082


Table 1. d_{TV} estimation error under different (m_1, m_2) pairs.



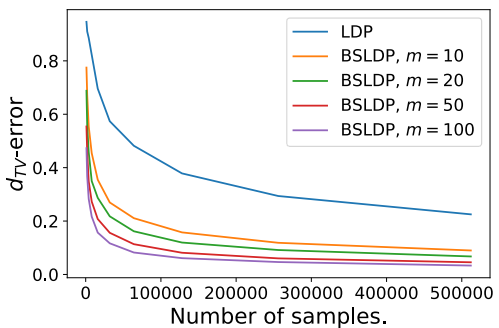
(a) Uniform



(b) Geo(0.95)



(c) Zipf(1)



(d) Geo(0.95)*

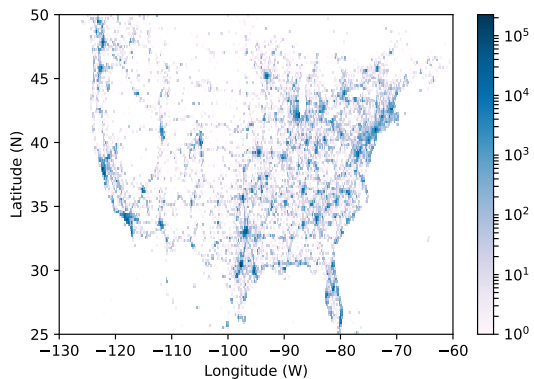
Figure 2. d_{TV} -error comparison under different distributions

Figure 3. Heatmap (distribution) of the check-in records.

within continental US). We round the latitude and longitude for each record up to accuracy 0.2 and regard records with the same latitude and longitude as the same location. By doing so, we get a dataset with 43,750 possible values. Figure 3 shows the empirical distribution of the check-in records of the dataset. We take the empirical distribution of the records as our ground truth and try to estimate it while preserving the privacy of each record. We partition latitudes into m_1 equal parts and longitudes into m_2 equal parts. The resulting grid will be used as the blocks ($m_1 m_2$ blocks in total). Table 1 shows the average d_{TV} error over 100 runs of the experiment for LDP and BS-LDP with different (m_1, m_2) pairs. From the table we can see that by switching to block-structured LDP, we can get more meaningful estimation accuracy compared to classical LDP.

8. Conclusion

We presented a unifying context-aware LDP framework and investigated communication, privacy, and utility trade-offs under both binary and k -ary data domains for minimax distribution estimation. Our theory and experiments on synthetic and real-world data suggest that context-aware LDP leads to substantial utility gains compared to vanilla LDP. In order to examine the effect of the number of data partitions in block-structured LDP, our experiments focused on uniform partitioning of geo-location data and examined the utility gains and various partition sizes. In practice however, non-uniform partitioning can better reflect the different topologies of cities. Thus, more careful experiments with non-uniform partitions need to be performed to better quantify the utility gains. More broadly, more experiments should be conducted to verify that the gains we see on the Gowalla dataset are also applicable in other data domains.

References

- Acharya, J. and Sun, Z. Communication complexity in locally private distribution estimation and heavy hitters. In Chaudhuri, K. and Salakhutdinov, R. (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 51–60, Long Beach, California, USA, 09–15 Jun 2019. PMLR.
- Acharya, J., Canonne, C., Freitag, C., and Tyagi, H. Test without trust: Optimal locally private distribution testing. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 2067–2076, 2019a.
- Acharya, J., Canonne, C. L., and Tyagi, H. Inference under information constraints: Lower bounds from chi-square contraction. In Beygelzimer, A. and Hsu, D. (eds.), *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99 of *Proceedings of Machine Learning Research*, pp. 3–17, Phoenix, USA, 2019b. PMLR.
- Acharya, J., Sun, Z., and Zhang, H. Hadamard response: Estimating distributions privately, efficiently, and with little communication. In Chaudhuri, K. and Sugiyama, M. (eds.), *Proceedings of Machine Learning Research*, volume 89 of *Proceedings of Machine Learning Research*, pp. 1120–1129. PMLR, 16–18 Apr 2019c.
- Alvim, M. S., Chatzikokolakis, K., Palamidessi, C., and Pazzi, A. Local differential privacy on metric spaces: Optimizing the trade-off with utility. *2018 IEEE 31st Computer Security Foundations Symposium (CSF)*, pp. 262–267, 2018.
- Asi, H., Duchi, J., and Javidbakht, O. Element level differential privacy: The right granularity of privacy, 2019.
- Bassily, R. and Smith, A. Local, private, efficient protocols for succinct histograms. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pp. 127–135. ACM, 2015.
- Bassily, R., Nissim, K., Stemmer, U., and Thakurta, A. G. Practical locally private heavy hitters. In *Advances in Neural Information Processing Systems*, pp. 2285–2293, 2017.
- Beimel, A., Nissim, K., and Omri, E. Distributed private data analysis: Simultaneously solving how and what. In *Proceedings of the 28th Annual International Cryptology Conference, CRYPTO '08*, pp. 451–468, Berlin, Heidelberg, 2008. Springer.
- Borgs, C., Chayes, J., Smith, A., and Zadik, I. Revealing network structure, confidentially: Improved rates for node-private graphon estimation. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 533–543. IEEE, 2018.
- Bun, M., Nelson, J., and Stemmer, U. Heavy hitters and the structure of local privacy. In *Proceedings of the 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, pp. 435–447. ACM, 2018.
- Chatzikokolakis, K., Andrés, M. E., Bordenabe, N. E., and Palamidessi, C. Broadening the scope of differential privacy using metrics. In *International Symposium on Privacy Enhancing Technologies Symposium*, pp. 82–102. Springer, 2013.
- Chatzikokolakis, K., Elsalamouny, E., and Palamidessi, C. Efficient utility improvement for location privacy. *Proceedings on Privacy Enhancing Technologies*, 2017(4): 308–328, 2017.
- Cho, E., Myers, S. A., and Leskovec, J. Friendship and mobility: user movement in location-based social networks. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 1082–1090, 2011.
- Diakonikolas, I., Hardt, M., and Schmidt, L. Differentially private learning of structured discrete distributions. In *Advances in Neural Information Processing Systems 28*, pp. 2566–2574. 2015.
- Dong, J., Roth, A., and Su, W. J. Gaussian differential privacy. *arXiv preprint arXiv:1905.02383*, 2019.
- Doudalis, S., Kotsoginannis, I., Haney, S., Machanavajjhala, A., and Mehrotra, S. One-sided differential privacy. *Proceedings of the VLDB Endowment*, 10(11), 2017.
- Duchi, J. C., Jordan, M. I., and Wainwright, M. J. Local privacy and statistical minimax rates. In *FOCS*, pp. 429–438. IEEE, 2013.
- Dwork, C., Mcsherry, F., Nissim, K., and Smith, A. Calibrating noise to sensitivity in private data analysis. In *In Proceedings of the 3rd Theory of Cryptography Conference*, 2006.
- Erlingsson, Ú., Pihur, V., and Korolova, A. Rappor: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security*, pp. 1054–1067. ACM, 2014.
- Feyisetan, O., Diethe, T., and Drake, T. Leveraging hierarchical representations for preserving privacy and utility in text, 2019.
- Geumlek, J. and Chaudhuri, K. Profile-based privacy for locally private computations. In *2019 IEEE International Symposium on Information Theory (ISIT)*, pp. 537–541. IEEE, 2019.

- Hsu, J., Khanna, S., and Roth, A. Distributed private heavy hitters. In *International Colloquium on Automata, Languages, and Programming*, pp. 461–472. Springer, 2012.
- Kairouz, P., Oh, S., and Viswanath, P. Extremal mechanisms for local differential privacy. In *Advances in Neural Information Processing Systems*, pp. 2879–2887, 2014.
- Kairouz, P., Bonawitz, K., and Ramage, D. Discrete distribution estimation under local privacy. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48, ICML'16*, pp. 2436–2444, 2016.
- Kasiviswanathan, S. P., Lee, H. K., Nissim, K., Raskhodnikova, S., and Smith, A. What can we learn privately? *SIAM Journal on Computing*, 40(3):793–826, 2011.
- Kawamoto, Y. and Murakami, T. Local obfuscation mechanisms for hiding probability distributions. In *European Symposium on Research in Computer Security*, pp. 128–148. Springer, 2019.
- Mangat, N. S. An improved randomized response strategy. *Journal of the Royal Statistical Society: Series B (Methodological)*, 56(1):93–95, 1994.
- Murakami, T. and Kawamoto, Y. Restricted local differential privacy for distribution estimation with high data utility. *arXiv preprint arXiv:1808.01394*, 2018.
- Schein, A., Wu, Z. S., Schofield, A., Zhou, M., and Wallach, H. Locally private bayesian inference for count models. In *International Conference on Machine Learning*, pp. 5638–5648, 2019.
- Wang, S., Huang, L., Wang, P., Nie, Y., Xu, H., Yang, W., Li, X.-Y., and Qiao, C. Mutual information optimally local private discrete distribution estimation. *arXiv preprint arXiv:1607.08025*, 2016.
- Wang, T. and Blocki, J. Locally differentially private protocols for frequency estimation. In *Proceedings of the 26th USENIX Security Symposium*, 2017.
- Warner, S. L. Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309):63–69, 1965.
- Xiang, Z., Ding, B., He, X., and Zhou, J. Linear and range counting under metric-based local differential privacy, 2019.
- Ye, M. and Barg, A. Optimal schemes for discrete distribution estimation under locally differential privacy. *IEEE Transactions on Information Theory*, 64:5662–5676, 2018.
- Zhu, W., Kairouz, P., Sun, H., McMahan, B., and Li, W. Federated heavy hitters discovery with differential privacy. *arXiv preprint arXiv:1902.08534*, 2019.