Hiding Identities: Estimation Under Local Differential Privacy

Antonious M. Girgis, Deepesh Data, and Suhas Diggavi
University of California, Los Angeles, USA
Email: amgirgis@ucla.edu, deepesh.data@gmail.com, suhas@ee.ucla.edu

***Abstract*—In this paper, we study an estimation problem under the local differential privacy (LDP) framework: There is an ordered list of d values (e.g., real numbers); a set of n users, where each user observes an element from this list and each value in the list is observed by at least one user; and an untrusted server, who wants to estimate the values that the users possess, without learning (in the sense of LDP) the actual value that each user has and its corresponding index in the list. Towards this, we propose two LDP estimation schemes: The first one is under the assumption that the server knows the number of users that observe each value; and the second one is for the general scenario, in which the server does not have this prior information. We show that the minimax risk decreases with the total number of users under a very mild condition on the number of users observing each value.**

I. Introduction

Differential privacy (DP) [1] has become a standard definition of privacy in privacy-preserving data analysis. DP ensures that the participation of a single person in a database does not change the probability of an outcome by much. In this paper, we focus on a variant of DP, called local differential privacy (LDP), in which the data is distributed among multiple users, and an aggregator wants to compute a statistic on the users' data privately [2], [3]. LDP provides a method in which each user gives a privatized-version of its data to the aggregator so that it can compute the desired statistic without compromising the privacy of individual user's data. LDP has received a considerable amount of attention both in academia [4]–[6] as well as in industrial applications [7], [8]. A classic example of an LDP mechanism is *randomized response* [9].

We consider a novel estimation problem under differential privacy constraints: There is an *ordered* list of d values (e.g., real numbers); a set of n users, each holding an element from this list and each value is held by at least one user; and an untrusted server, who wants to estimate the list, in a manner that the actual values that users possess from the list as well as their corresponding indices are hidden from the server in the sense of LDP. We study the privacy-utility trade-offs for this private estimation problem. To the best of our knowledge, this problem has not been studied before. This model may be useful in several scenarios. In an academic campus, multiple students move from time to time across different classrooms. Suppose there is an administrator who wants to know the number of students present in different classrooms at any given time, maybe to control the temperature of the classrooms. Note that students have a good (or maybe exact) estimate of the number of students present in their classrooms, but may not have much idea of other classrooms. Now, if the students share their raw observation with the administrator, they reveal their locations on the campus, which they may be uncomfortable sharing. Hence, it raises an intriguing question: How can the administrator have a good estimate of the number of students present in each classroom, while preserving privacy of students' locations on the campus?

We study it in the LDP framework, where each user perturbs its own data (independent of other users) and provides it to the server, with the guarantee that any change to the data of a single user leads to a small change in the probability of the outcome of the algorithm. Hence, the untrusted server cannot infer the private data of an individual user from observing the output of the LDP mechanism.

Our contributions. Apart from the problem formulation, our contributions can be summarized as follows.

- We first formulate a minimax risk for the above-described problem under LDP contraints.
- Under the assumption that the server knows the number of users that observe each element of the (unknown) list, we propose an LDP mechanism for each user, and characterize the privacy-utility trade-offs of our proposed mechanism for different loss functions. We show that minimax risk decreases as the number of users n increases under the condition that the number of users that observe any value from the list increase as $\omega(\sqrt{n})$.
- We develop an LDP mechanism for the general case, when the server does not have the above prior knowledge about the frequency of each element of the list. Our proposed mechanism consists of two stages: In the first stage, the server estimates the number of users that observe each value of the list; in the second stage, the server estimates the list of d values given the estimated values from the first stage. We analyze the privacy-utility trade-offs, and prove that the minimax risk converges with the same order as in the case when the server knows the number of users that observes each value from the list. Lower bounds for this formulation is part of (ongoing) future work.

In our proposed algorithm for the general case, we propose a new *optimal* LDP scheme for frequency estimation that is different from the schemes proposed in [5], [6], [10]. Our LDP frequency estimation scheme is designed such that each user sends a binary bit with different mean depending on the

This was supported by the NSF grant #1740047 and by the UC-NL grant LFR-18-548554.

index of its value from the list. The simplicity of our proposed scheme helps us to prove some properties of the frequency estimators used in the main algorithm.

Related work: Statistical estimation under LDP has been studied extensively in different contexts in the literature. In [11]–[13], the authors studied discrete distribution estimation under LDP, wherein a service provider wants to learn a discrete distribution from users' samples without learning the underlying samples. The work in [14]–[16] studied differentially private multi-party computation, where each party is interested in computing a function of the entire dataset of all parties while preserving privacy of its own data. In [17], [18], the authors derived lower and upper bounds for estimation under LDP – their work considers that all users observe i.i.d. samples from the same distribution, and the goal for each user is to preserve privacy of its raw sample. Our work is different from these papers, in the sense that in our setup, there is no underlying distribution from which users draw their data – there is an arbitrary and ordered list of, say, real numbers, and each user observes a different element from this list. The goal in our work is to preserve (in addition to the actual data) privacy of the index of the element that each user observes from the list. We will see in Section II-A that in order to preserve privacy of the index of the value that each user observes, it is required to preserve privacy of both the index and the value. Hence, our proposed LDP mechanisms give more private output in the following sense: The adversary cannot infer from the output of the algorithm not only the raw data but also the index of the data (in the ordered list from which it is drawn) that each user observes.

Paper organization. We formulate our problem with and without the assumption that the server knows the frequency of each element of the list in Section II. We state our main results in Section III. We give our algorithm under the above assumption in Section III-A and for the general case in Section III-B. In Section III-C, we give our new LDP mechanism for frequency estimation. Omitted proofs from this paper can be found in the full version [19].

II. Problem Formulation

Let $\Theta \subseteq \mathbb{R}^d$ denote a set of ordered lists of d elements. Consider an arbitrary element $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_d)$ from Θ. In our problem, every user observes a value from $\boldsymbol{\theta}$. For $i \in [n] := \{1, \ldots, n\}$, let the i-th user observe θ_{a_i} for some $a_i \in [d]$[1]. Note that $d < n$. We can divide the n users into d disjoint sets, $\{\mathcal{S}_j\}_{j=1}^d$, where the users in $\mathcal{S}_j$ observe θ_j. Let $\alpha_j := \frac{|\mathcal{S}_j|}{n}$ denote the fraction of users that observe θ_j. As mentioned in Section I, we assume that each θ_j is observed by at least one of the users. So, we have $\sum_{j=1}^d \alpha_j = 1$, where $\alpha_j \geq \frac{1}{n}$ for every $j \in [d]$.

For given $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_d)$ such that $\sum_{j=1}^d \alpha_j = 1$ and $\alpha_j \geq \frac{1}{n}, j = 1, \ldots, d$, we define a set $\mathcal{A}_{\boldsymbol{\alpha}} \subset [d]^n$, where $\mathcal{A}_{\boldsymbol{\alpha}} = \{\mathbf{a} \in [d]^n : \frac{|\{i:a_i=j\}|}{n} = \alpha_j, \forall j \in [d]\}$ denotes the set of assignment vectors $\mathbf{a} = (a_1, \ldots, a_n) \in [d]^n$ such that $\alpha_j n$ users observe θ_j. The i-th user generates a private output Y_i as a function of θ_{a_i}. Then, the untrusted server estimates the vector $\boldsymbol{\theta}$ from observing the private outputs $Y^n = (Y_1, \ldots, Y_n)$.

A. Local Differential Privacy (LDP) Mechanisms

In order for the server to estimate $\boldsymbol{\theta}$, each user $i \in [n]$ shares information about its observed value θ_{a_i}, while keeping the index of its event a_i private. In other words, from Y_i, the untrusted server should not learn whether the i-th user observes θ_j or $\theta_{j'}$ for any $j, j' \in [d]$. For the i-th user, a private mechanism Q_i with input $a_i \in [d]$ and $\theta_{a_i} \in \mathbb{R}$ is used to generate a private output $y_i \in \mathcal{Y}_i$. We consider private mechanisms $\{Q_i\}$ that provide local differential privacy for each user [3].

Definition 1. A private mechanism Q is said to satisfy ϵ-local differential privacy (ϵ-LDP), if for every $\boldsymbol{\theta} \in \Theta$ and every pair $j, j' \in [d]$, we have

$$\sup_{y \in \mathcal{Y}} \frac{Q\left(y|\theta_j, a = j\right)}{Q\left(y|\theta_{j'}, a = j'\right)} \leq e^{\epsilon}, \tag{1}$$

where $Q\left(y|\theta_j, a = j\right) = \Pr\left[Y = y|\theta_j, a = j\right]$ and ϵ captures the privacy level. The smaller the privacy level ϵ, the untrusted server has more difficulty inferring whether $a_i = j$ or $a_i = j'$ for any pair $j, j' \in [d]$.

Let $\mathcal{Q}_\epsilon$ denote the set of private mechanisms that satisfy ϵ-LDP. Thus, each user chooses a private mechanism $Q_i \in \mathcal{Q}_\epsilon$ to generate a private output $Y_i \in \mathcal{Y}_i$ for $i \in [n]$.

B. Minimax Risk Estimation

We first formulate the minimax problem under the following assumption. Then, we formulate the general problem without this assumption.

Assumption 1. For a given (unknown) $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_d)$, we assume that the server knows exactly the number of users that observe each θ_j, i.e., it knows the frequency vector $\boldsymbol{\alpha}$. However, it does not know the exact assignment vector $\mathbf{a} \in \mathcal{A}_{\boldsymbol{\alpha}}$.

The server uses the users' private outputs $Y^n = (Y_1, \ldots, Y_n)$ to estimate the vector $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_d)$. Let $\hat{\boldsymbol{\theta}}\left(Y^n\right) = \left(\hat{\theta}_1, \ldots, \hat{\theta}_d\right)$ denote the estimator of the server that maps Y^n to a vector in the space Θ. For given private mechanisms $Q^n = [Q_1, \ldots, Q_n]$, the performance of the estimator $\hat{\boldsymbol{\theta}}$ is measured by the expected loss:

$$\sup_{\boldsymbol{\theta} \in \Theta} \sup_{\mathbf{a} \in \mathcal{A}_{\boldsymbol{\alpha}}} \mathbb{E}\left[\ell\left(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}\left(Y^n\right)\right)\right],$$

where the expectation is taken over the randomness in the outputs Y^n, and $\ell : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}^+$ denotes the loss function. Now we define the minimax problem under Assumption 1:

$$r_{\epsilon,n,d}^{\ell,\boldsymbol{\alpha}} = \inf_{Q^n \in \mathcal{Q}_\epsilon} \inf_{\hat{\boldsymbol{\theta}}} \sup_{\boldsymbol{\theta} \in \Theta} \sup_{\mathbf{a} \in \mathcal{A}_{\boldsymbol{\alpha}}} \mathbb{E}\left[\ell\left(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}\left(Y^n\right)\right)\right], \tag{2}$$

[1] We assume that the i-th user knows the index a_i of the element θ_{a_i}. In our academic campus example, the student knows its location and the overall list of locations enabling it to identify its index in the list.

where $\inf_{\hat{\boldsymbol{\theta}}}$ is to design the best estimator that minimizes the expected risk, and $\inf_{Q^n \in \mathcal{Q}_\epsilon}$ is to design an ϵ-LDP private mechanism for users to minimize the expected risk.

Now we define the minimax risk when Assumption 1 is not satisfied. In this case, the minimax risk problem can be formulated as

$$r^{\ell}_{\epsilon,n,d} = \inf_{Q^n \in \mathcal{Q}_\epsilon} \inf_{\hat{\boldsymbol{\theta}}} \sup_{\boldsymbol{\theta} \in \Theta} \sup_{\mathbf{a} \in [d]^n} \mathbb{E}\left[\ell\left(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}\left(Y^n\right)\right)\right]. \quad (3)$$

The difference between (2) and (3) is that, in (2) we take supremum over the assignment vectors $\mathbf{a} \in \mathcal{A}_{\boldsymbol{\alpha}}$, since the server knows a prior knowledge that $\mathbf{a} \in \mathcal{A}_{\boldsymbol{\alpha}}$, whereas, in (3) we take the supremum over all possible assignment vectors $\mathbf{a} \in [d]^n$, since the server has no prior information about $\mathbf{a}$.

Remark 1. (Deterministic vs probabilistic) In the above problem formulation, we assume that the vector $\boldsymbol{\theta}$ is deterministic, where all users in $\mathcal{S}_j$ have access to the exact value θ_j for $j \in [d]$. We refer to this model as a deterministic model. We can extend our formulations to the probabilistic model as follows. Let $\mathcal{P}$ be a set of distributions on a space $\mathcal{X} = \mathcal{X}_1 \times \cdots \times \mathcal{X}_d \subseteq \mathbb{R}^d$. Let $\mathcal{P}_j$ denote the set of marginal distributions on the sample space $\mathcal{X}_j \subseteq \mathbb{R}$. Let $\boldsymbol{\theta}(P) \in \Theta$ be a function mapping from $\mathcal{P}$ to Θ. Each user $i \in [n]$ observes a random sample $X_{a_i} \in \mathcal{X}_{a_i}$ drawn from unknown distribution $P_{a_i} \in \mathcal{P}_{a_i}$. For example, consider that the server wants to estimate the mean $\boldsymbol{\theta}$ from a family of normal distributions $\mathcal{P} \triangleq \{\mathcal{N}\left(\boldsymbol{\theta}, \sigma^2 \mathbb{I}_d\right) : \boldsymbol{\theta} \in \Theta\}$ with known variance σ^2. In this case, each user does not observe the exact event θ_{a_i}, and hence, a private mechanism Q is said to satisfy ϵ-LDP if for every $\mathbf{x} = (x_1, \ldots, x_d) \in \mathcal{X}$ and every pair $j, j' \in [d]$, we have

$$\sup_{y \in \mathcal{Y}} \frac{Q\left(y | x_j, a = j\right)}{Q\left(y | x_{j'}, a = j'\right)} \leq e^{\epsilon} \quad (4)$$

Furthermore, the minimax risk problem under Assumption 1 is formulated as

$$r^{\ell,\boldsymbol{\alpha}}_{\epsilon,n,d} = \inf_{Q^n \in \mathcal{Q}_\epsilon} \inf_{\hat{\boldsymbol{\theta}}} \sup_{P \in \mathcal{P}} \sup_{\mathbf{a} \in \mathcal{A}_{\boldsymbol{\alpha}}} \mathbb{E}_{P, Q^n}\left[\ell\left(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}\left(Y^n\right)\right)\right], \quad (5)$$

where the expectation is taken over the randomness in Y^n. In this paper, we only present achievable schemes for the deterministic model, with the understanding that it is possible to extend our ideas to the probabilistic model as well.

III. Main Results

In this section we present the main results of this paper. We present two ϵ-LDP algorithms and characterize the minimax risks $r^{\ell,\boldsymbol{\alpha}}_{\epsilon,n,d}$ (which is under Assumption 1) in Theorem 1 and $r^{\ell}_{\epsilon,n,d}$ for the general case in Theorem 2.

Let $\Theta = [C_{\min} : C_{\max}]^d$, i.e., $C_{\min} \leq \theta_j \leq C_{\max}$ for all $j \in [d]$, where $C_{\min}$, $C_{\max}$ are global constants known to all parties and $C_{\min} < C_{\max}$. Define $b := C_{\max} - C_{\min}$.

A. An ϵ-LDP Algorithm and its Bounded Minimax Risk under Assumption 1

In this section, we present an algorithm DIST-EST-AS under Assumption 1, and prove that it is ϵ-LDP and has bounded minimax risks.

Theorem 1. *Fix an arbitrary $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_d)$ such that $\sum_{j=1}^d \alpha_j = 1$ and $\alpha_j \geq \frac{1}{n}$ for $j \in [d]$. Let $\epsilon = \mathcal{O}(1)$. Consider an arbitrary $\boldsymbol{\theta} \in \Theta$. For any assignment vector $\mathbf{a} \in \mathcal{A}_{\boldsymbol{\alpha}}$,* DIST-EST-AS *is an ϵ-LDP algorithm that achieves the following bounds for estimating $\boldsymbol{\theta} \in \Theta$:*

$$r^{\ell_2^2,\boldsymbol{\alpha}}_{\epsilon,n,d} = \mathbb{E}\left[||\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}||_2^2\right] = \mathcal{O}\left(\frac{b^2}{\epsilon^2 n} \sum_{j=1}^d \frac{1}{\alpha_j^2}\right) \quad (6)$$

$$r^{\ell_1,\boldsymbol{\alpha}}_{\epsilon,n,d} = \mathbb{E}\left[||\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}||_1\right] = \mathcal{O}\left(\sqrt{\frac{db^2}{\epsilon^2 n} \sum_{j=1}^d \frac{1}{\alpha_j^2}}\right) \quad (7)$$

$$r^{\ell_\infty,\boldsymbol{\alpha}}_{\epsilon,n,d} = \mathbb{E}\left[\max_{j \in [d]} |\hat{\theta}_j - \theta_j|\right] = \mathcal{O}\left(\frac{b}{\epsilon \alpha^*} \sqrt{\frac{\log(d)}{n}}\right), \quad (8)$$

where $\alpha^* = \min_{j \in [d]} \alpha_j$.

Algorithm 1 DIST-EST-AS: ϵ-LDP distributed estimation under Assumption 1

Inputs: Frequencies $\boldsymbol{\alpha} \succ \mathbf{0}$, vector assignment $\mathbf{a} \in \mathcal{A}_{\boldsymbol{\alpha}}$, event vector $\boldsymbol{\theta} \in \Theta$, and privacy level ϵ.
for user $i = 1$ to n **do**
  user i generates vector $\mathbf{Y}_i = [Y_{i1}, \ldots, Y_{id}] \in \mathbb{R}^d$
  for $j = 1$ to d **do**
    $Z_{ij} \sim \text{Lap}\left(\frac{2b}{\epsilon}\right)$
    if $a_i \neq j$ **then**
      $Y_{ij} = C_{\min} + Z_{ij}$
    else
      $Y_{ij} = \theta_j + Z_{ij}$
    end if
  end for
end for
Server computes $\overline{\mathbf{Y}} = \frac{1}{n} \sum_{i=1}^n \mathbf{Y}_i$
$\hat{\boldsymbol{\theta}} \leftarrow \left(\overline{\mathbf{Y}} - (\mathbf{1} - \boldsymbol{\alpha}) C_{\min}\right) \oslash \boldsymbol{\alpha}$, where for any two vectors $\mathbf{x} = (x_1, \ldots, x_d)$, $\mathbf{y} = (y_1, \ldots, y_d)$, $\mathbf{x} \oslash \mathbf{y} := \left(\frac{x_1}{y_1}, \ldots, \frac{x_d}{y_d}\right)$ is defined as the component-wise division.

We prove Theorem 1 in Section IV. Theorem 1 implies that if the frequencies satisfy $1/\alpha_j^2 \in o(n)$ (which can be equivalently written as $n\alpha_j = \omega(\sqrt{n})$), then the minimax risk goes to zero as the number of users become arbitrarily large. The main idea of our proposed algorithm (Algorithm 1)[2] is that the i-th user generates a vector of d Laplace random variables, where the a_i-th random variable has mean θ_{a_i}, while the other random variables have mean $C_{\min}$. The variance of the noise is chosen to preserve privacy of the output. The server aggregates the response of all users for each event. Hence, for the j-th value θ_j, there are $n\alpha_j$ Laplace random variables with mean θ_j, while there are $(1 - \alpha_j) n$ Laplace random variables with mean $C_{\min}$. As a result, the server can

[2]Note that Algorithm 1 represents the ϵ-LDP mechanisms of all users as well as the estimator of the server, where each user has only access to her observation θ_{a_i}.

estimate $\boldsymbol{\theta}$ from knowing the frequencies $\boldsymbol{\alpha}$ and $C_{\min}$ without disclosing the identity of the users that observe this event.

Now we state our main result for the general case.

B. An ϵ-LDP Algorithm and its Bounded Minimax Risk for the General Case

In this section, we present an algorithm for the general case (where the server does not have any prior information about the frequency vector $\boldsymbol{\alpha}$), and prove that it is ϵ-LDP and has bounded minimax risks.

As mentioned in Section II, we assume, without loss of generality, that $\alpha_j \geq \frac{1}{n}$, i.e., each value θ_j is observed by at least one user. To be concrete, assume that all frequencies are bounded by $\alpha_j \geq \delta, j = 1, 2, \ldots$, for some $\delta \geq \frac{1}{n}$. Here δ can be treated as a confidence bound known to all parties.

Theorem 2. *Let $\epsilon = \mathcal{O}(1)$. Consider an arbitrary $\boldsymbol{\theta} \in \Theta$ and $\delta > 0$. Let $\alpha_j \geq \delta$ for all $j \in [d]$. Then* DIST-EST *is an ϵ-LDP algorithm that achieves the following bounds for estimating $\boldsymbol{\theta} \in \Theta$:*

$$r^{\ell_2^2}_{\epsilon,n,d} = \sup_{\boldsymbol{\alpha}} \mathcal{O}\left(\frac{b^2}{\epsilon^2 n}\sum_{j=1}^{d}\frac{1}{\alpha_j^2}\right) = \mathcal{O}\left(\frac{db^2}{\delta^2\epsilon^2 n}\right) \tag{9}$$

$$r^{\ell_1}_{\epsilon,n,d} = \sup_{\boldsymbol{\alpha}} \mathcal{O}\left(\sqrt{\frac{db^2}{\epsilon^2 n}\sum_{j=1}^{d}\frac{1}{\alpha_j^2}}\right) = \mathcal{O}\left(\sqrt{\frac{d^2b^2}{\delta^2\epsilon^2 n}}\right) \tag{10}$$

Algorithm 2 DIST-EST: ϵ-LDP distributed estimation

Inputs: Confidence bound $\delta > 0$, vector assignment $\mathbf{a} \in [d]^n$ such that $\alpha_j \geq \delta$ for $j \in [d]$, event vector $\boldsymbol{\theta} \in \Theta$, and privacy level ϵ.
$\hat{\boldsymbol{\alpha}} \leftarrow$ FREQ-EST $(\mathbf{a}, \delta, \epsilon/2)$ (run Algorithm 3 to obtain the estimate $\hat{\boldsymbol{\alpha}}$)
$\hat{\boldsymbol{\theta}} \leftarrow$ DIST-EST-AS $(\hat{\boldsymbol{\alpha}}, \mathbf{a}, \theta, \epsilon/2)$ (run Algorithm 1 to obtain the estimate $\hat{\boldsymbol{\theta}}$)

To prove Theorem 2, we propose an ϵ-LDP scheme that consists of two stages. In the first stage, the server estimates the frequencies $\boldsymbol{\alpha}$ using LDP frequency estimation scheme FREQ-EST proposed in Section III-C. In the second stage, the server estimates the vector $\boldsymbol{\theta}$ given the estimated frequencies from the first stage. The main challenge here is to analyse how the error in estimating frequencies $\boldsymbol{\alpha}$ affects the error in estimating the event vector $\boldsymbol{\theta}$. Observe that for fixed frequencies $\boldsymbol{\alpha}$, the minimax risk for the general case given in Theorem 2 has the same order as the minimax risk given in Theorem 1, which is for the case when the server knows the frequencies $\boldsymbol{\alpha}$. To obtain this result, we fist prove some properties of FREQ-EST (See Theorems 4 in Section III-C). These properties help us to prove that the error in privately estimating the frequencies $\boldsymbol{\alpha}$ does not hurt the minimax risk of estimating $\boldsymbol{\theta}$. Observe that once $1/\delta^2 \in o(n)$, the minimax risk of estimating the d events goes to zero as the number of users is arbitrary large ($n \to \infty$). We prove Theorem 2 in Section IV.

C. Frequency Estimation under LDP

In this section, we present a new LDP mechanism FREQ-EST for estimating the frequencies $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_d)$ while preserving the privacy of each user. As mentioned in Section I, the frequency estimation under an LDP constraint has been studied before, and a number of schemes have been proposed [5], [6], [10] so far. Our algorithm FREQ-EST is designed such that each user sends a binary bit with different mean depending on the index of the value it has. Our algorithm is arguably simpler than the ones in [5], [6], [10], and this simplicity helps us to prove some properties which we used in the main algorithm DIST-EST of this paper.

Our private frequency estimator FREQ-EST is described in Algorithm 3. We assume an input parameter $\delta \geq 0$, which denotes a confidence bound representing a prior information about the frequencies, such that $\alpha_j \geq \delta$ for all $j \in [d]$. Let

Algorithm 3 FREQ-EST: ϵ-LDP frequency estimation

Inputs: Vector assignment $\mathbf{a} \in [d]^n$, Confidence bound $\delta \geq 0$, and privacy level ϵ.
for u **do**ser $i = 1$ to n
  user i generates vector $\mathbf{Y}_i = [Y_{i1}, \ldots, Y_{id}]$ of d bits
  for $j = 1$ to d **do**
    if $a_i \neq j$ **then**
      $Y_{ij} \sim \text{Bern}\left(\frac{1}{e^{\epsilon/2}+1}\right)$
    else
      $Y_{ij} \sim \text{Bern}\left(\frac{e^{\epsilon/2}}{e^{\epsilon/2}+1}\right)$
    end if
  end for
end for
Server computes $\overline{\mathbf{Y}} = \frac{1}{n}\sum_{i=1}^{n}\mathbf{Y}_i$.
$\mathbf{T} \leftarrow \frac{e^{\epsilon/2}+1}{e^{\epsilon/2}-1}\left(\overline{\mathbf{Y}} - \frac{1}{e^{\epsilon/2}+1}\right)$
$\hat{\boldsymbol{\alpha}} \leftarrow \langle \mathbf{T} \rangle_\delta$

$X \sim \text{Bern}(p)$ denote a Bernoulli random variable such that $X = 1$ with probability p and $X = 0$ with probability $1 - p$. For vector $\mathbf{x}$, the operation $\langle \mathbf{x} \rangle_\delta$ denotes elementwise truncation from below δ. In Algorithm 3, the i-th user generates a vector of d binary random variables, where the a_i-th bit is a Bernoulli random variable with parameter $\frac{e^{\epsilon/2}}{e^{\epsilon/2}+1}$, while the other bits are Bernoulli random variable with $p = \frac{1}{e^{\epsilon/2}+1}$.

Theorem 3. *Let $\epsilon = \mathcal{O}(1)$. For any assignment vector $\mathbf{a} \in [d]^n$,* FREQ-EST *is an ϵ-LDP algorithm that achieves the following bounds for estimating the frequency vector $\boldsymbol{\alpha}$:*

$$r^{\ell_2^2}_{\epsilon,n,d} = \mathbb{E}\left[||\hat{\boldsymbol{\alpha}} - \boldsymbol{\alpha}||_2^2\right] = \mathcal{O}\left(\frac{d}{\epsilon^2 n}\right) \tag{11}$$

$$r^{\ell_1}_{\epsilon,n,d} = \mathbb{E}\left[||\hat{\boldsymbol{\alpha}} - \boldsymbol{\alpha}||_1\right] = \mathcal{O}\left(\sqrt{\frac{d^2}{\epsilon^2 n}}\right) \tag{12}$$

$$r^{\ell_\infty}_{\epsilon,n,d} = \mathbb{E}\left[\max_{j\in[d]}|\hat{\alpha}_j - \alpha_j|\right] = \mathcal{O}\left(\frac{1}{\epsilon}\sqrt{\frac{\log(d)}{n}}\right) \tag{13}$$

Theorem 3 is valid even when $\delta = 0$. Observe that FREQ-EST is order optimal for heavy hitter estimation for the ℓ_∞ loss, matching the bounds in [6], [10].

In the following we bound the expectation $\mathbb{E}\left[\frac{1}{\hat{\alpha}_j^2}\right]$ for given confidence bound $\delta > 0$. In addition, we also bound the expectation $\mathbb{E}\left[\left(\frac{\hat{\alpha}_j-\alpha_j}{\hat{\alpha}_j}\right)^2\right]$. These two values will be used in the proof of Theorem 2 to analyze the performance of Algorithm 2 DIST-EST.

Theorem 4. *For $\delta > 0$, let $\alpha_j \geq \delta$ for $j \in [d]$. We have*

$$\mathbb{E}\left[\frac{1}{\hat{\alpha}_j^2}\right] \leq \frac{4}{\alpha_j^2} + C_1 e^{-2n\left(\frac{e^{\epsilon/2}-1}{e^{\epsilon/2}+1}\right)^2 \frac{{\alpha_j}^2}{4}},$$

$$\mathbb{E}\left[\frac{1}{\hat{\alpha}_j^4}\right] \leq \frac{16}{\alpha_j^4} + C_2 e^{-2n\left(\frac{e^{\epsilon/2}-1}{e^{\epsilon/2}+1}\right)^2 \frac{{\alpha_j}^2}{4}},$$

$$\mathbb{E}\left[\left(\frac{\hat{\alpha}_j-\alpha_j}{\hat{\alpha}_j}\right)^2\right] \leq C_{\epsilon,n}\sqrt{\left(\frac{16}{\alpha_j^4} + Ce^{-2n\left(\frac{e^{\epsilon/2}-1}{e^{\epsilon/2}+1}\right)^2 \frac{{\alpha_j}^2}{4}}\right)},$$

where $C_1 = \frac{1}{\delta^2} - \frac{4}{9}$, $C_2 = \frac{1}{\delta^4} - \frac{16}{81}$ are global constants, and $C_{\epsilon,n} = \sqrt{\left(\frac{\left(e^{\epsilon/2}\right)^4 + e^{\epsilon/2}}{\left(e^{\epsilon/2}-1\right)^4\left(e^{\epsilon/2}+1\right)n^3}\right)}$. The bound in the third inequality is $\mathcal{O}\left(\frac{1}{\alpha_j^2\epsilon^2 n^{3/2}}\right)$.

IV. Proof Outlines

In this section, we give prove outlines of our main theorems and defer the complete proofs to the full version [19].

Proof of Theorem 1. **Proof of ϵ-LDP.** Consider an arbitrary user $i \in [n]$. Let $\mathbf{y} = (y_1, \ldots, y_d) \in \mathbb{R}^d$ be one of the outputs of the i-th user. Let $f(y) = \frac{\epsilon}{4b} e^{-\frac{\epsilon|y|}{2b}}$ denote the probability density function of the Laplace distribution with zero mean and variance $\frac{2b}{\epsilon}$. Thus, we get

$$\begin{aligned}\frac{Q(\mathbf{y}|\theta_j, a_i = j)}{Q(\mathbf{y}|\theta_{j'}, a_i = j')} &= \frac{f(y_j - \theta_j) f(y_{j'} - C_{\min})}{f(y_j - C_{\min}) f(y_{j'} - \theta_{j'})} \\ &= \exp\left(\frac{\epsilon}{2b}\left(|y_j - C_{\min}| - |y_j - \theta_j|\right)\right) \times \\ &\qquad \exp\left(\frac{\epsilon}{2b}\left(|y_{j'} - \theta_{j'}| - |y_{j'} - C_{\min}|\right)\right) \\ &\leq \exp\left(\frac{\epsilon}{2b}\left(|\theta_j - C_{\min}| + |\theta_{j'} - C_{\min}|\right)\right) \leq e^{\epsilon}.\end{aligned}$$

Proof of bounded minimax risks. First, observe that $\hat{\theta}_j = \frac{1}{\alpha_j}\left(\overline{y}_j - (1-\alpha_j) C_{\min}\right)$, where $\overline{y}_j = \frac{1}{n}\sum_{i=1}^n y_{ij}$. Thus, $\mathbb{E}\left[\hat{\theta}_j\right] = \theta_j$. Hence, the error to estimate $\boldsymbol{\theta}$ under ℓ_2^2 loss is bounded by

$$\begin{aligned}\mathbb{E}\left[||\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}||_2^2\right] &= \sum_{j=1}^d \mathbb{E}\left[\left(\hat{\theta}_j - \theta_j\right)^2\right] \\ &= \sum_{j=1}^d \frac{1}{\alpha_j^2}\mathbb{E}\left[\left(\overline{y}_j - \mathbb{E}\left(\overline{y}_j\right)\right)^2\right] = \sum_{j=1}^d \frac{1}{\alpha_j^2}\frac{\sum_{i=1}^n \mathrm{Var}(Z_{ij})}{n^2} \\ &= \sum_{j=1}^d \frac{1}{\alpha_j^2}\frac{8b^2}{\epsilon^2 n} = \mathcal{O}\left(\frac{b^2}{\epsilon^2 n}\sum_{j=1}^d \frac{1}{\alpha_j^2}\right)\end{aligned}$$

We can easily get the results in (7) for the ℓ_1 loss from the inequality $\mathbb{E}\left[||\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}||_1\right] \leq \mathbb{E}\left[\sqrt{d||\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}||_2^2}\right] \leq \sqrt{d\,\mathbb{E}\left[||\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}||_2^2\right]}$, where the last inequality obtained from Jensen's inequality and concavity of the function $f(x) = \sqrt{x}$. Now, we prove the results in (8). Observe that

$$\left(\hat{\theta}_j - \theta_j\right) = \frac{1}{\alpha_j n}\sum_{i=1}^n \left(Y_{ij} - \mathbb{E}\left[Y_{ij}\right]\right) = \frac{1}{\alpha_j n}\sum_{i=1}^n Z_{ij},$$

where $Z_{ij} \sim \mathrm{Lap}\left(\frac{2b}{\epsilon}\right)$. Thus, Z_{ij} is a sub-exponential random variable with parameter $\lambda = \frac{4b}{\epsilon}$ and its moment generating function is $\mathbb{E}\left[e^{sZ_{ij}}\right] \leq e^{s^2\lambda^2/2}$ for $|s| \leq 1/\lambda$. Thus, we get

$$\begin{aligned}\mathbb{E}\left[e^{s\frac{1}{\alpha_j n}\sum_{i=1}^n Z_{ij}}\right] &= \prod_{i=1}^n \mathbb{E}\left[e^{\frac{s}{\alpha_j n}Z_{ij}}\right] \\ &\leq \prod_{i=1}^n e^{\frac{s^2}{\alpha_j^2 n^2}\frac{4b^2}{\epsilon^2}} = e^{\frac{s^2}{\alpha_j^2 n}\frac{4b^2}{\epsilon^2}}, \qquad \forall\, |\frac{s}{\alpha_j n}| \leq \frac{\epsilon}{4b}\end{aligned} \tag{14}$$

The remaining of the proof follows similar steps as [20, Theorem 1.14] and is provided in [19]. ∎

Proof of Theorem 2. **Proof of ϵ-LDP.** Observe that Algorithm 2 DIST-EST first runs FREQ-EST with the privacy parameter $\epsilon/2$, which is $\epsilon/2$-LDP, and then it runs DIST-EST-AS with the privacy parameter $\epsilon/2$, which is $\epsilon/2$-LDP. Furthermore, these two mechanisms are independent. Hence, from the composition theorem [21], DIST-EST is ϵ-LDP.

Proof of bounded minimax risks. First, observe that $\hat{\theta}_j = \frac{1}{\hat{\alpha}_j}\left(\overline{y}_j - (1-\hat{\alpha}_j) C_{\min}\right)$, where $\overline{y}_j = \frac{1}{n}\sum_{i=1}^n y_{ij}$, and $\hat{\alpha}_j$ is the estimate of α_j obtained from Algorithm 3. Thus, the error in estimating $\boldsymbol{\theta}$ under ℓ_2^2 loss is given by

$$\begin{aligned}\mathbb{E}\left[||\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}||_2^2\right] &= \sum_{j=1}^d \mathbb{E}\left[\left(\hat{\theta}_j - \theta_j\right)^2\right] \\ &= \sum_{j=1}^d \mathbb{E}\left[\frac{1}{\hat{\alpha}_j^2}\left(\overline{y}_j - \mathbb{E}\left(\overline{y}_j\right) - (\hat{\alpha}_j - \alpha_j)(\theta_j - C_{\min})\right)^2\right] \\ &\leq \sum_{j=1}^d \mathbb{E}\left[\frac{1}{\hat{\alpha}_j^2}\right]\frac{\sum_{i=1}^n \mathrm{Var}(Z_{ij})}{n^2} + b^2\mathbb{E}\left[\left(\frac{\hat{\alpha}_j - \alpha_j}{\hat{\alpha}_j}\right)^2\right] \\ &= \sum_{j=1}^d \mathbb{E}\left[\frac{1}{\hat{\alpha}_j^2}\right]\frac{8b^2}{\epsilon^2 n} + b^2\mathbb{E}\left[\left(\frac{\hat{\alpha}_j - \alpha_j}{\hat{\alpha}_j}\right)^2\right] \\ &\overset{(a)}{\leq} \mathcal{O}\left(\frac{b^2}{\epsilon^2 n}\sum_{j=1}^d \frac{1}{\alpha_j^2}\right) + \mathcal{O}\left(\frac{b^2}{\epsilon^2 n^{3/2}}\sum_{j=1}^d \frac{1}{\alpha_j^2}\right),\end{aligned}$$

where step (a) follows from Theorem 4. We can obtain the results in (10) for the ℓ_1 loss from the inequality $\mathbb{E}\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}\|_1 \leq \mathbb{E}\sqrt{d\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}\|_2^2} \leq \sqrt{d\,\mathbb{E}\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}\|_2^2}$, where the last inequality obtained from Jensen's inequality and concavity of the function $f(x) = \sqrt{x}$. This completes the proof of Theorem 2. ∎

REFERENCES

[1] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Theory of cryptography conference*. Springer, 2006, pp. 265–284.

[2] J. C. Duchi, M. I. Jordan, and M. J. Wainwright, "Privacy aware learning," *J. ACM*, vol. 61, no. 6, pp. 38:1–38:57, 2014.

[3] ——, "Local privacy and statistical minimax rates," in *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*. IEEE, 2013, pp. 429–438.

[4] C. Dwork, "Differential privacy and the us census," in *Proceedings of the 38th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, ser. PODS '19. New York, NY, USA: ACM, 2019, pp. 1–1. [Online]. Available: http://doi.acm.org/10.1145/3294052.3322188

[5] J. Hsu, S. Khanna, and A. Roth, "Distributed private heavy hitters," in *International Colloquium on Automata, Languages, and Programming*. Springer, 2012, pp. 461–472.

[6] R. Bassily and A. Smith, "Local, private, efficient protocols for succinct histograms," in *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*. ACM, 2015, pp. 127–135.

[7] Ú. Erlingsson, V. Pihur, and A. Korolova, "RAPPOR: randomized aggregatable privacy-preserving ordinal response," in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, Scottsdale, AZ, USA, November 3-7, 2014*, 2014, pp. 1054–1067.

[8] Apple, "Differential privavy," 2017. [Online]. Available: \url{https://www.apple.com/privacy/docs/Differential_Privacy_Overview.pdf}

[9] S. L. Warner, "Randomized response: A survey technique for eliminating evasive answer bias," *Journal of the American Statistical Association*, vol. 60, no. 309, pp. 63–69, 1965.

[10] J. Acharya and Z. Sun, "Communication complexity in locally private distribution estimation and heavy hitters," *arXiv preprint arXiv:1905.11888*, 2019.

[11] P. Kairouz, K. Bonawitz, and D. Ramage, "Discrete distribution estimation under local privacy," *arXiv preprint arXiv:1602.07387*, 2016.

[12] M. Ye and A. Barg, "Optimal schemes for discrete distribution estimation under locally differential privacy," *IEEE Transactions on Information Theory*, vol. 64, no. 8, pp. 5662–5676, Aug 2018.

[13] J. Acharya, Z. Sun, and H. Zhang, "Hadamard response: Estimating distributions privately, efficiently, and with little communication," *arXiv preprint arXiv:1802.04705*, 2018.

[14] A. McGregor, I. Mironov, T. Pitassi, O. Reingold, K. Talwar, and S. Vadhan, "The limits of two-party differential privacy," in *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*. IEEE, 2010, pp. 81–90.

[15] P. Kairouz, S. Oh, and P. Viswanath, "Differentially private multi-party computation: Optimality of non-interactive randomized response," *arXiv preprint arXiv:1407.1546*, 2014.

[16] G. Wu, Y. He, J. Wu, and X. Xia, "Inherit differential privacy in distributed setting: Multiparty randomized function computation," in *2016 IEEE Trustcom/BigDataSE/ISPA*. IEEE, 2016, pp. 921–928.

[17] J. C. Duchi, M. I. Jordan, and M. J. Wainwright, "Minimax optimal procedures for locally private estimation," *Journal of the American Statistical Association*, vol. 113, no. 521, pp. 182–201, 2018.

[18] J. Duchi and R. Rogers, "Lower bounds for locally private estimation via communication complexity," *arXiv preprint arXiv:1902.00582*, 2019.

[19] A. M. Girgis, D. Data, and S. Diggavi, "Hiding identities: Estimation under local differential privacy," 2020, Available on arXiv.

[20] P. Rigollet, "18. s997: High dimensional statistics," *Lecture Notes, Cambridge, MA, USA: MIT Open-CourseWare*, 2015.

[21] C. Dwork and A. Roth, "The algorithmic foundations of differential privacy," *Foundations and Trends in Theoretical Computer Science*, vol. 9, no. 3–4, pp. 211–407, 2014.