# Private and Communication-Efficient Algorithms for Entropy Estimation 

Gecia Bravo-Hermsdorff<br>Department of Statistics<br>University College London<br>gecia.bravo@gmail.com

Róbert Busa-Fekete<br>Google Research<br>busarobi@google.com

Mohammad Ghavamzadeh<br>Google Research<br>ghavamza@google.com

Andres Muñoz Medina<br>Google Research<br>ammedina@google.com

Umar Syed<br>Google Research<br>usyed@google.com


#### Abstract

Modern statistical estimation is often performed in a distributed setting where each sample belongs to a single user who shares their data with a central server. Users are typically concerned with preserving the privacy of their samples, and also with minimizing the amount of data they must transmit to the server. We give improved private and communication-efficient algorithms for estimating several popular measures of the entropy of a distribution. All of our algorithms have constant communication cost and satisfy local differential privacy. For a joint distribution over many variables whose conditional independence is given by a tree, we describe algorithms for estimating Shannon entropy that require a number of samples that is linear in the number of variables, compared to the quadratic sample complexity of prior work. We also describe an algorithm for estimating Gini entropy whose sample complexity has no dependence on the support size of the distribution and can be implemented using a single round of concurrent communication between the users and the server. In contrast, the previously bestknown algorithm has high communication cost and requires the server to facilitate interaction between the users. Finally, we describe an algorithm for estimating collision entropy that matches the space and sample complexity of the best known algorithm but generalizes it to the private and communication-efficient setting.


## 1 Introduction

Statistical estimation has traditionally focused on minimizing the number of samples needed to estimate properties of a distribution. In the 'big data' era, statisticians and computer scientists have also tried to minimize the space complexity of estimation algorithms, particularly in the streaming setting. More recently, the increasing prevalence of mobile computing has led to a focus on the privacy and communication costs of statistical estimation. In this paper, we consider the following setting: a set of users each draw one sample from a distribution, and share information about their sample with a central server. The central server then uses the collected data to estimate a property of the distribution. Users are concerned with preserving the privacy of their sample, and also with minimizing the amount of data that is transmitted to the server.
For example, consider the problem of detecting fingerprinting on the web [31]. Many websites track users across the web without their consent by recording (enough) information about their devices (e.g., installed fonts, operating system, timezone, etc.), a practice known as "browser fingerprinting". Entropy is the standard metric used to quantify the identifiability of the collected fingerprints. So a
private and distributed method for estimating entropy can be used by a browser to warn users that this covert tracking could occur, without ever storing the fingerprints themselves.
The study of entropies has an extensive and rich history in mathematics and sciences. Related quantities called "entropy" appear in many contexts (thermodynamics, information theory, dynamical systems [35], category theory [9], etc.). These may be broadly thought as measures of information of a system or process obeying certain properties, which, in turn, lead to natural measures of disorder, randomness, outcome diversity, information content, uniformity, etc.

In this paper, we study private and communication-efficient algorithms for estimating certain entropies of a distribution. Specifically, we give algorithms for estimating the following entropies, which are widely-used in many scientific fields to quantify the uncertainty, diversity and spread of a discrete distribution:

- Shannon entropy [33], a fundamental quantity in information theory.
- Gini entropy (also known as Tsallis entropy [38] of order 2, or (one minus the) second frequency moment). Some of its applications include measuring ecological diversity [34, 27], market competition among firms [21], effective size of political parties [26], and suitability of features to split on during decision tree learning [30].
- Collision entropy (also known as Rényi entropy [32] of order 2). Some of its applications include measuring the quality of random number generators [36], and determining the number of reads needed to reconstruct a DNA sequence [28].

Our algorithms are implemented in either the non-interactive model (for the Gini and collision entropies), in which all users simultaneously exchange information with the server during a single round of communication, or the (stronger) sequentially interactive model (for the Shannon entropy), in which the server queries users one at a time, possibly in an adaptive manner [22]. When analyzing the communication complexity of an algorithm, we prove bounds on the number of bits that each user transmits to the server. However, the server is allowed to broadcast an arbitrary amount of information to the users (this is also called the blackboard model [16]), including shared random bits (also known as the public coin model [1, 2, 23]). When analyzing the privacy of our algorithms, we use the framework of local differential privacy [17], which ensures that the server learns very little about each user's data.

## Our contributions:

- A sequentially interactive $\alpha$-local differentially private algorithm for estimating the Shannon entropy of a joint distribution on $d$ variables within $\epsilon d$ error using $\tilde{O}\left(d / \alpha^{2} \epsilon^{3}\right)$ samples and $O(1)$ bits per sample. Our analysis assumes that each of the $d$ variables has a constant support size and that their conditional independence graph is a tree. We also describe algorithms that have better dependence on $1 / \epsilon$ in certain special cases, such as when the tree has low diameter or is a chain. Our algorithms achieve $O(1)$ communication complexity by observing only two or three of the $d$ variables in any single sample; we call these pair and triplet observations. The only previously known algorithm for estimating the Shannon entropy of a tree-structured distribution from pair observations is a non-interactive algorithm due to Chow and Liu [12]. We prove that any non-interactive algorithm requires $\Omega\left(d^{2}\right)$ pair observations to achieve $O(d)$ error. We also prove that, for any sequentially interactive algorithm, $\Omega(d / \epsilon)$ pair observations are necessary to achieve $O(\epsilon d)$ error.
- A non-interactive $\alpha$-local differentially private algorithm for estimating the Gini entropy of a distribution within $\epsilon$ error using $O\left(b^{2} \max \left\{1-g, 2^{-b}\right\} /\left(\alpha^{2} \epsilon^{2}\right)\right)$ samples, $b$ bits per sample, and $\tilde{O}(b)$ space, where $g \in[0,1]$ is the Gini entropy of the distribution. The best previous algorithms [10] either have a sample complexity that depends on the support size $k$ of the distribution, or are sequentially interactive, and also require $\Omega(k)$ bits per sample and $\Omega(k)$ space. Also our error bound holds with high probability instead of only in expectation.
- A non-interactive $\alpha$-local differentially private algorithm for estimating the collision entropy of a distribution with support size $k$ within $\epsilon$ error using $\tilde{O}\left(b^{2} k^{2} /\left(\alpha^{2} \epsilon^{2} \min \left\{k, 2^{b}\right\}\right)\right)$ samples, $b$ bits per sample, and $\tilde{O}(b)$ space. Setting $b=\log k$ and $\alpha=O(1)$ recovers the sample and space complexity guarantees of the non-interactive algorithm from [36] up to logarithmic factors, and thus, our algorithm generalizes the previously best known algorithm to the private and communicationefficient setting.


## 2 Related Work

There is a very extensive literature on distributed statistical estimation under communication constraints (see [40] for the paper that appears to have started this thread). Variations on the problem include whether communication is allowed between users, whether communication happens in one or multiple rounds, whether there is a shared source of randomness among the users, and whether communication is limited per-user or only cumulatively across all users.
Many previous results in this area bound the sample and communication complexity of estimating the parameters of a distribution $P_{\theta}$, where $\theta \in \Theta$ (see e.g. [20]). This problem class includes discrete distribution estimation, where the guarantees are usually stated as bounds on the relative entropy or total variation distance between the estimated and true distribution (see e.g. [5]). Other problems of interest are mean estimation [37] and heavy hitter estimation [4].
There has also been significant interest in differentially private statistical estimation, and of particular relevance is the work by [3], who gave private algorithms for estimating certain functionals of a distribution, including the Shannon entropy. However, they used the central model of differential privacy, while in this paper we prove guarantees using the (stronger) local model.

## 3 Entropy Measures

The Shannon, Tsallis, and Rényi entropy of a discrete random variable $X$ are defined as

$$
\begin{array}{ll}
\text { (Shannon) } & H(X)=-\sum_{x} \operatorname{Pr}[X=x] \log \operatorname{Pr}[X=x], \\
\text { (Tsallis) } & T_{\gamma}(X)=\frac{1}{\gamma-1}\left(1-\sum_{x} \operatorname{Pr}[X=x]^{\gamma}\right), \\
\text { (Rényi) } & R_{\gamma}(X)=\frac{1}{1-\gamma} \log \left(\sum_{x}(\operatorname{Pr}[X=x])^{\gamma}\right),
\end{array}
$$

where $\gamma$ in (2) and (3) is a free parameter satisfying $\gamma>0$ and $\gamma \neq 1$. Both Tsallis and Rényi entropy are generalizations of Shannon entropy in the sense that $\lim _{\gamma \rightarrow 1} T_{\gamma}(X)=\lim _{\gamma \rightarrow 1} R_{\gamma}(X)=H(X)$.
In this paper, we describe algorithms for estimating the Shannon entropy and special cases of the Tsallis and the Rényi entropy that are widely used in many scientific fields: $T_{2}(X)$, also known as the Gini entropy, and $R_{2}(X)$, also known as the collision entropy. Substituting $\gamma=2$ into the definitions above and using the abbreviations $G(X) \equiv T_{2}(X)$ and $C(X) \equiv R_{2}(X)$, we have:

$$
\begin{array}{cl}
\text { (Gini) } & G(X) \equiv T_{2}(X)=1-\sum_{x} \operatorname{Pr}[X=x]^{2}, \\
\text { (Collision) } & C(X) \equiv R_{2}(X)=-\log \left(\sum_{x} \operatorname{Pr}[X=x]^{2}\right) .
\end{array}
$$

Gini entropy is so-called because it is equivalent to the Gini diversity index, a statistic proposed by Corrado Gini in 1912 to measure income and wealth inequality [19]. Collision entropy takes its name from the observation that if $X$ and $X^{\prime}$ are independent and identically distributed, then $C(X)=-\log \operatorname{Pr}\left[X=X^{\prime}\right]$.

For the problem of estimating Shannon entropy, we specialize to a high-dimensional setting, where we only observe a pair (or triplet) of the dimensions at a time. That is, $X$ is a random-vector of $d$ discrete variables, where $d$ is large, but each $X_{i}$ has a constant support size (e.g., they are binary), and we only observe two (or three) dimensions per sample. Without making any assumption about this joint distribution, the problem is intractable. One of the most common assumptions, which we also adopt in this work, is that the joint distribution is tree-structured. In this case, the distribution can be estimated by the celebrated [12] (and optimal [7]) Chow-Liu algorithm. While the Chow-Liu algorithm requires $\Omega\left(d^{2}\right)$ pairs observations to estimate the Shannon entropy, our sequential algorithm requires only $\mathcal{O}(d)$ pairs observations (see Section 5.2 for more details).
The joint Shannon entropy $H\left(X_{1}, \ldots, X_{d}\right)$ of a set of random variables $X_{1}, \ldots X_{d}$ is the Shannon entropy $H(X)$ of the random variable $X=\left(X_{1}, \ldots, X_{d}\right)$. We write the abbreviated term joint entropy when the use of Shannon entropy is obvious from context.

The mutual information between two random variables $X$ and $Y$ and their conditional mutual information given another random variable $Z$ are defined as:

$$
\begin{align*}
I(X ; Y) & =H(X)+H(Y)-H(X, Y)  \tag{4}\\
I(X ; Y \mid Z) & =H(X, Z)+H(Y, Z)-H(X, Y, Z)-H(Z) \tag{5}
\end{align*}
$$

## 4 Estimation Algorithms and Evaluation Criteria

A set of $n$ users and a central server cooperate according to the following protocol to estimate the entropy of a random variable $X$ :

1. Each user $i \in[n]$ draws an independent sample $x_{i}$ according to the distribution of $X$.
2. For $r$ rounds:
(a) The server sends information to a subset of the users.
(b) Those users send (partial) information about their sample back to the server.
3. The server outputs an estimate of the Shannon entropy (Algorithms 1, 2, and 4) or the Gini or collision (Algorithm5) entropies of $X$.

An estimation algorithm specifies the steps that each user and the server perform to implement the above protocol. The algorithm is non-interactive if the protocol consists of a single round in which all users participate. In a non-interactive algorithm the server cannot adapt its queries to users based on responses from other users, since the server communicates with all the users concurrently. An algorithm is sequentially interactive if each round consists of communication with a single user, who is never contacted again. Sequential interactivity enables the server to query users adaptively [22].
We evaluate estimation algorithms according to the following criteria:

- Sample complexity: The number of users from whom the server requests data.
- Space complexity: The space used by the server when executing the algorithm.
- Communication complexity: The maximum number of bits transmitted by any single user to the server. Note that the amount of information sent by the server to the users is not counted when determining communication complexity.
- Privacy: Let $x_{i}$ be the sample belonging to user $i$ and $o_{i}$ be the data observed by the server from user $i$. We say that an algorithm satisfies $\alpha$-local differential privacy if

$$
\operatorname{Pr}\left[o_{i} \in O \mid x_{i}=x\right] \leq e^{\alpha} \operatorname{Pr}\left[o_{i} \in O \mid x_{i}=x^{\prime}\right]
$$

for any user $i$, measurable set $O$, and possible sample values $x, x^{\prime}$.

- Error: The absolute difference between the true entropy of the distribution and the estimate output by the server.


## 5 Estimating Shannon Entropy of Tree-structured Joint Distributions

In this section we assume that $X=\left(X_{1}, \ldots, X_{d}\right)$ is a vector of $d$ discrete variables, and that the support size of each variable $X_{i}$ is constant (e.g., each variable is Boolean). We also assume that $X$ has a tree-structured distribution, which means that there exists a rooted tree $T$ with $d$ nodes such that for any node $i \in[d]$ we have $\operatorname{Pr}\left[X_{i} \mid X_{-i}\right]=\operatorname{Pr}\left[X_{i} \mid X_{\mathrm{pa}_{T}(i)}\right]$, where $X_{-i}=$ $\left(X_{1}, \ldots X_{i-1}, X_{i+1}, \ldots, X_{d}\right)$ and the node $\mathrm{pa}_{T}(i)$ is the parent of node $i$ in tree $T$. If $i$ is the root node, then we define $\operatorname{Pr}\left[X_{i} \mid X_{\mathrm{pa}_{T}(i)}\right]=\operatorname{Pr}\left[X_{i}\right]$. Equivalently, a tree-structured distribution is a Markov random field with a tree as the underlying undirected graph. Essentially, the tree-structured assumption implies that the only correlations among the $X_{i}$ 's are pairwise correlations. If $T$ is a chain or a star we say that $X$ is chain-structured and star-structured, respectively. We will treat these two special cases at the end of this section (Algorithms 2and 4, respectively).

### 5.1 Estimating Entropy of a Marginal Distribution When the Support Size is Small

Before proceeding to describe algorithms for estimating the Shannon entropy of tree-structured distributions, we use existing results for private distribution estimation to devise a local differentially private estimator for the Shannon entropy that is sample and communication efficient when the support size of the distribution is small (as is the case for the individual marginals). The server will repeatedly invoke this algorithm as a subroutine in the sections below.

First we recall that the difference in Shannon entropy of two random variables can be upper bounded according to Theorem 17.3.3 of [13] as

$$
\left|H\left(X_{\mathbf{p}}\right)-H\left(X_{\mathbf{p}^{\prime}}\right)\right| \leq\left\|\mathbf{p}-\mathbf{p}^{\prime}\right\|_{1} \log \frac{c}{\left\|\mathbf{p}-\mathbf{p}^{\prime}\right\|_{1}}
$$

where $X_{\mathbf{p}}$ and $X_{\mathbf{p}^{\prime}}$ are two discrete random variables with support size $c$ and distributions $\mathbf{p}$ and $\mathbf{p}^{\prime}$. Next, we apply a local differentially private learning algorithm for discrete distribution due to [4] that learns the parameters of a discrete distribution with small $L_{1}$ error. The following theorem combines these two results by using the fact that $x / \log (1 / x) \leq 1$ whenever $0<x \leq 1 / 2$.
Theorem 5.1. For any discrete distribution $X$ with support size $c$ and for any $1 / 2 \geq \epsilon>0$, there exists an estimator satisfying $\alpha$-local differential privacy that estimates $H(X)$ within $\epsilon$ error using $n=O\left(c^{2} \log \frac{1}{\delta} /\left(\epsilon^{2} \alpha^{2}\right)\right)$ samples with probability $1-\delta$ when $\alpha \in(0,1)$.

This algorithm resulting from Theorem 5.1 can be used to privately estimate the entropy $H\left(X_{i}\right)$, mutual information $I\left(X_{i} ; X_{j}\right)$, and conditional mutual information $I\left(X_{i} ; X_{j} \mid X_{k}\right)$ of any variables $X_{i}, X_{j}$ and $X_{k}$. This can be done using $O\left(\frac{\log \frac{1}{8}}{\alpha^{2} \epsilon^{2}}\right)$ samples per estimate and $O(1)$ bits per sample, since each of these variables has constant support size, and both mutual information and conditional mutual information can be expressed in terms of entropies (Eqs. (4) and (5)). We call such an estimate ( $\alpha, \epsilon, \delta$ )-good.

### 5.2 Our Algorithm for Tree-structured Joint Distributions

Note that the support size of $X$ can be exponential in $d$. In the worst case, estimating the entropy of a distribution with support size $k$ within constant error requires $\tilde{\Theta}(k)$ samples [18]. However the tree-structure of $X$ can be exploited to significantly reduce the sample complexity. In their seminal paper, Chow and Liu [12] proved the identity

$$
\begin{equation*}
H(X)=\sum_{i=1}^{d} H\left(X_{i}\right)-\max _{T} \sum_{i=1}^{d} I\left(X_{i} ; X_{\mathrm{pa}_{T}(i)}\right) \tag{6}
\end{equation*}
$$

for any tree-structured random variable $X$, where the maximization is taken over all possible trees connecting the $d$ variables.
Eq. (6) suggests a communication-efficient algorithm for estimating the entropy of $X$, which is known as the Chow-Liu algorithm: First, estimate each marginal entropy $H\left(X_{i}\right)$ and each mutual information $I\left(X_{i} ; X_{j}\right)$. Next, compute a maximum spanning tree on the $d$ variables, where the weight of each edge $\left(X_{i}, X_{j}\right)$ is the estimate of the mutual information $I\left(X_{i} ; X_{j}\right)$. Finally, plug these estimators into Eq. (6).
The Chow-Liu algorithm requires $\Omega\left(d^{2}\right)$ samples, since it computes the mutual information between every pair of variables in order to compute a maximum spanning tree. However, estimating the right-hand side of Eq. (6) only requires estimating the weight of the maximum spanning tree, which is significantly easier than finding the tree itself. Algorithm 1 adapts a technique from [11] that estimates the weight of the maximum spanning tree of a graph in time that is sublinear in the number of edges in the graph. The basic idea is to select nodes of the graph at random and use breadth-first search to determine the size of each of their connected components if we were to drop edges that do not meet a weight threshold, short-circuiting the search when the size becomes too large. These quantities are combined to estimate the weight of the maximum spanning tree. In our case, an edge weight is a mutual information between a pair of variables, which we estimate from pair observations. Theorem 5.2. Algorithm 1 is $\alpha$-locally differentially private and has $O(1)$ communication complexity. The number of samples requested by the server is $O\left(\frac{d \log \left(\frac{1}{\delta}\right)}{\alpha^{2} \epsilon^{3}}\right)$. Let $\hat{H}$ be the entropy estimate output by the algorithm. If $X$ is tree-structured, then $|\hat{H}-H(X)| \leq \epsilon d$ with probability $1-\delta$.

```
Algorithm 1 Shannon entropy estimation for tree-structured distribution
    Let \(M=\left\lceil\frac{2}{\epsilon}\right\rceil\) and \(R=\left\lceil\frac{1}{\epsilon^{2}}\right\rceil\)
    for \(m=1, \ldots, M\) do
        for \(r=1, \ldots, R\) do
            Choose positive integer \(Z\) randomly according to \(\operatorname{Pr}[Z \geq z]=1 / z\).
            Choose \(i^{*}\) uniformly at random from \([d]\).
            if \(Z \geq \frac{2}{\epsilon}\) then \(\gamma_{m r} \leftarrow 0\).
            else \(\quad \triangleright\) Breadth-first search
                    Initialize queue to contain \(i^{*}\) and a set \(V=\left\{i^{*}\right\}\).
                    while queue length is non-zero and shorter than \(Z\) do
                        Remove \(i\) from front of queue and \(V=V \cup\{i\}\).
                for \(j=[d] \backslash V\) do
                    Server computes \((\alpha, \epsilon, \delta)\)-good estimate \(\hat{I}_{i j}\) of \(I\left(X_{i} ; X_{j}\right)\).
                    if \(\hat{I}_{i j} \geq \epsilon m\) then add \(j\) to back of queue.
                    if queue length is zero then \(\gamma_{m r} \leftarrow 0\) else \(\gamma_{m r} \leftarrow 1\).
            \(\hat{c}_{m} \leftarrow \frac{d}{R} \sum_{r=1}^{R} \gamma_{m r}\).
    \(\hat{W} \leftarrow \epsilon M d-\epsilon \sum_{m=1}^{M} \hat{c}_{m}\)
    Server computes \((\alpha, \epsilon, \delta)\)-good estimate of each entropy (for the first sum in Eq. (6)).
    Let \(\hat{S}\) be the sum of the entropy estimates.
    Return \(\hat{H}=\hat{S}-\hat{W}\).
```


### 5.3 Our Algorithm for Chain-structured Joint Distributions

Verma and Pearl [39] observed that if $X$ is chain-structured with chain $T$ then for any triplet ( $X_{i}$, $\left.X_{j}, X_{k}\right)$, if $X_{k}$ is on the unique path in $T$ between $X_{i}$ and $X_{j}$, then $I\left(X_{i} ; X_{j} \mid X_{k}\right)=0$. Thus, for any triplet $\left(X_{i}, X_{j}, X_{k}\right)$ in $T$, at least one of $I\left(X_{i} ; X_{j} \mid X_{k}\right), I\left(X_{i} ; X_{k} \mid X_{j}\right)$, or $I\left(X_{j} ; X_{k} \mid X_{i}\right)$ has to be zero. This observation alone is not enough to recover the chain, since the conditional mutual information $I\left(X_{i} ; X_{j} \mid X_{k}\right)$ can also be zero for $X_{i}, X_{j}$ and $X_{k}$ when $X_{k}$ is not on the path between $X_{i}$ and $X_{j}$ in the chain $T$. Nevertheless, under the mild assumption that the mutual information $I\left(X_{i}, X_{j}\right)$ between every pair of variables is distinct, we can recover the chain $T$ by estimating the conditional mutual information of triplets of variables.

Our algorithm is similar to sorting algorithms such as mergesort [24], which require $O\left(d \log _{2} d\right)$ pairwise comparisons over $d$ items. While we cannot compare pairs explicitly like in a sorting problem, for any triplet $\left(X_{i}, X_{j}, X_{k}\right)$, we can use their conditional mutual information estimators to locally decide which "item" is between the other two: i.e., $X_{i} \leftrightarrow X_{j} \leftrightarrow X_{k}, X_{i} \leftrightarrow X_{k} \leftrightarrow X_{j}$ or $X_{k} \leftrightarrow X_{i} \leftrightarrow X_{j}$ in the chain $T$. This suggests our Algorithm 2, which inserts the variables in a chain one by one in a sequential manner. Algorithm 2 calls Algorithm 3 as a subroutine that seeks to find the position where to insert.
Theorem 5.3. Algorithm 2 is $\alpha$-locally differentially private and has $O(1)$ communication complexity. The number of samples requested by the server is $O\left(\frac{d \log \frac{d}{\delta}}{\alpha^{2} \epsilon^{2}}\right)$. Let $\hat{H}$ be the entropy estimate output by the algorithm. If $X$ is chain-structured and $\left|I\left(X_{i} ; X_{j}\right)-I\left(X_{j} ; X_{k}\right)\right| \geq \epsilon$, then $|\hat{H}-H(X)| \leq \epsilon d$ with probability $1-\delta$.

### 5.4 Our Algorithm for Star-structured Joint Distributions

If $X$ is star-structured then recovering the star $T$ is a matter of identifying its center, which can be done by computing the mutual information between only $\tilde{\mathcal{O}}(d)$ pairs of variables. The algorithm picks a random marginal $X_{i}$ and takes a "Prim's step" [29], i.e., chooses the neighboring node (say $X_{k}$ ) that has the largest mutual information with $X_{i}$. Assuming that the mutual information $I\left(X_{i}, X_{j}\right)$ between every pair of variables is distinct, the edge between $X_{i}$ and $X_{k}$ is in the maximal spanning tree. Next, the algorithm estimates $\sum_{j \neq i} I\left(X_{i}, X_{j}\right)$ and $\sum_{j \neq k} I\left(X_{k}, X_{j}\right)$ to decide whether $X_{i}$ or $X_{k}$ is the center node of the star. Algorithm 4 presents the procedure, and Theorem 5.4 gives its sample complexity.

Theorem 5.4. Algorithm 4 is $\alpha$-locally differentially private and has $O(1)$ communication complexity. The number of samples requested by the server is $O\left(\frac{d \log \frac{d}{\delta}}{\alpha^{2} \epsilon^{2}}\right)$. Let $\hat{H}$ be the entropy estimate output by the algorithm. If $X$ is star-structured and $\left|I\left(X_{i} ; X_{j}\right)-I\left(X_{j} ; X_{k}\right)\right| \geq \epsilon$, then $|\hat{H}-H(X)| \leq \epsilon d$ with probability $1-\delta$.

```
Algorithm 2 Shannon entropy estimation for chain-structured distribution
    \(S=[d], C=\emptyset\), pick an arbitrary \(i, j, k \in S\) and set \(S=S \backslash\{i, j, k\}\).
    Server computes \((\alpha, \epsilon, \delta)\)-good estimates \(\hat{I}\left(X_{i} ; X_{j} \mid X_{k}\right), \hat{I}\left(X_{i} ; X_{k} \mid X_{j}\right)\) and \(\hat{I}\left(X_{k} ; X_{j} \mid X_{i}\right)\).
    if \(\hat{I}\left(X_{i} ; X_{j} \mid X_{k}\right)>\epsilon\) then \(\boldsymbol{x}_{1}=(i, k, j)\)
    else if \(\hat{I}\left(X_{i} ; X_{k} \mid X_{j}\right)>\epsilon\) then \(\boldsymbol{x}_{1}=(i, j, k)\)
    else if \(\hat{I}\left(X_{k} ; X_{j} \mid X_{i}\right)>\epsilon\) then \(\boldsymbol{x}_{1}=(j, i, k)\)
    for \(i \in(1, \ldots, d-3)\) do
        Pick item \(j\) from \(S\) and set \(S=S \backslash\{j\}, r=x_{i, 1}\) and \(p=x_{i, i+2}\)
        Server computes \((\alpha, \epsilon, \delta)\)-good estimates \(\hat{I}\left(X_{j} ; X_{p} \mid X_{o}\right), \hat{I}\left(X_{r} ; X_{j} \mid X_{p}\right)\).
        if \(\hat{I}\left(X_{j} ; X_{p} \mid X_{r}\right)>\epsilon\) then \(\boldsymbol{x}_{i+1}=\left(j, \boldsymbol{x}_{i}\right) \quad \triangleright\) Attach \(X_{j}\) to the head of the chain
        else if \(\hat{I}\left(X_{o} ; X_{j} \mid X_{p}\right)>\epsilon\) then \(\boldsymbol{x}_{i+1}=\left(\boldsymbol{x}_{i}, j\right) \quad \triangleright\) Attach \(X_{j}\) to the tail of the chain
        else \(\quad \triangleright\) Insert \(X_{j}\) into the chain defined by \(\boldsymbol{x}_{i}\)
            \(\ell=\operatorname{TernarySEARCH}\left(\boldsymbol{x}_{i}, 1, i+2, j\right) \quad \triangleright\) Defined in Algorithm 3
            \(\boldsymbol{x}_{i+1}=\left(\boldsymbol{x}_{i}[1, \ldots, \ell], j, \boldsymbol{x}_{i}[\ell+1, \ldots, i+2]\right)\)
    Create chain \(T\) according to the order defined by \(\boldsymbol{x}_{d-2}\).
    Server computes \((\alpha, \epsilon, \delta)\)-good estimate of each term in Eq. (6) using \(T\) and returns \(\hat{H}\).
```

```
Algorithm 3 TERNARYSEARCH \(\left(\boldsymbol{x}, \ell_{l}, \ell_{h}, j\right)\)
    if \(\ell_{l}=\ell_{h}-1\) then return \(\ell_{l}\)
    Pick the median element \(k=\left\lceil\left(\ell_{h}+\ell_{l}\right)\right\rceil\), and set \(i=x_{\ell_{l}}\) and \(o=x_{\ell_{h}}\)
    Server computes \((\alpha, \epsilon, \delta)\)-good estimate \(\hat{I}\left(X_{i} ; X_{k} \mid X_{j}\right)\).
    if \(\hat{I}\left(X_{i} ; X_{k} \mid X_{j}\right)>\epsilon\) then return \(\operatorname{TernarySEARCh}(\boldsymbol{x}, i, k, j)\)
    else return TERNARYSEARCH \((\boldsymbol{x}, k, o, j)\)
```

```
Algorithm 4 Shannon entropy estimation for star-structured distribution
    Pick \(i \in[d]\) uniformly at random.
    Server computes \((\alpha, \epsilon, \delta)\)-good estimate \(\hat{I}\left(X_{i}, X_{j}\right)\) for all \(j \in[d] \backslash\{i\}\).
    Find \(k=\arg \max _{j \in[d] \backslash\{i\}} \hat{I}\left(X_{i}, X_{j}\right)\)
    Server computes \((\alpha, \epsilon, \delta)\)-good estimate \(\hat{I}\left(X_{k}, X_{j}\right)\).
    if \(\sum_{j} \hat{I}\left(X_{i}, X_{j}\right)>\sum_{j} \hat{I}\left(X_{k}, X_{j}\right)\) then let \(T\) be a star with \(X_{i}\) as center
    else let \(T\) be a star with \(X_{k}\) as the center.
    Server computes \((\alpha, \epsilon, \delta)\)-good estimate of each term in Eq. (6) using \(T\) and returns \(\hat{H}\).
```


### 5.5 Lower Bounds

We prove sample complexity lower bounds for estimating the Shannon entropy of a tree-structured joint distribution from pair observations. Our first lower bound focuses on the non-interactive case, when the algorithm must select all the pairs in advance. The second claim is more general, and holds for all sequentially interactive algorithms.
Theorem 5.5. For any non-interactive algorithm that uses $o\left(d^{2}\right)$ pair observations to estimate Shannon entropy, there exists a tree-structured distribution over $\{0,1\}^{d}$ such that the error of the algorithm is $\Omega(d)$ with constant probability.
Theorem 5.6. For any $\epsilon>0$ and for any sequentially interactive algorithm that uses $o(d / \epsilon)$ pair observations to estimate Shannon entropy, there exists a tree-structured distribution on $\{0,1\}^{d}$ such that the error of the algorithm is $\Omega(\epsilon \cdot d)$ with constant probability.

The lower bound given in Theorem 5.5 is based on Turán's theorem [8], which we use to show that for any algorithm with sub-quadratic sample complexity and for any constant $C \in(0,1)$, there is a graph with $d$ nodes containing a $C \cdot d$-clique - when $d$ is large enough - such that the algorithm does not observe any edge of that clique. This implies that the additive error of the algorithm is linear in $d$. The lower bound for sequentially interactive algorithms in Theorem 5.6 is based on a information theoretical approach. Interestingly, our construction of problem instances for which we applied Le Cam's theorem is fairly simple, since it contains $d$ independent random variables in every case. Nevertheless, this lower bound shows that Algorithm 1 is optimal in $d$.

### 5.6 Comparison to Prior Work

To the best of our knowledge, the Chow-Liu algorithm is the only published method for estimating the entropy of a distribution that takes advantage of its tree structure. Since the algorithm is noninteractive, the lower bound in Theorem 5.5 shows that our algorithms have provably better sample complexity when the number of variables $d$ is large (note that the dependence on $d$ in each of Theorems [5.2, 5.3 and 5.4 is sub-quadratic). The Chow-Liu algorithm can also be used to estimate the distribution itself, not just its entropy, and it has recently been shown [7, 15] that the algorithm has optimal sample complexity when given full observations (i.e., samples of the entire vector ( $X_{1}, \ldots, X_{d}$ ) and not just pairs or triplets of the variables). Thus the Chow-Liu algorithm is optimal for estimating a tree-structured distribution, but suboptimal for estimating the entropy of a treestructured distribution. The root cause of this difference appears to be the fact that it is significantly easier to estimate the weight of the maximum spanning tree than finding the tree itself.

## 6 Estimating Gini and Collision Entropies

In this section, we describe a non-iterative protocol (Algorithm 5) that estimates both the Gini and collision entropies of a discrete random variable $X$ while observing only $b$ bits per sample from its distribution, and does not require any extra assumption on the distribution of $X$ (such as assuming it is tree-structured). First, the server partitions all users into pairs (assume for simplicity that the number of users is even). The server then distributes a $b$-bit hash function to each user, along with a distinct salt to each user pair. Each user then hashes their sample along with their salt, and returns the hash value to the server. The server computes entropy estimates based on the number of observed hash collisions across all pairs. In Algorithm5, we let $\langle x, y\rangle$ denote a binary string that encodes $x$, followed by a delimiter, and by $y$.

```
Algorithm 5 Gini and collision entropy estimation
    Each user \(i \in[n]\) draws a sample \(x_{i}\) independently from the distribution of \(X\).
    2: Server partitions the \(n\) users into \(\frac{n}{2}\) disjoint pairs.
    3: Let \(q_{i} \in\left[\frac{n}{2}\right]\) be the index of the pair containing user \(i\).
    4: Server transmits \(q_{i}\) and hash function \(h:\{0,1\}^{*} \mapsto\{0,1\}^{b}\) to each user \(i\).
    5: Each user \(i\) generates a \(b\)-bit hash value \(h_{i}=h\left(\left\langle q_{i}, x_{i}\right\rangle\right)\) for their sample \(x_{i}\).
    6: Each user \(i\) lets \(\hat{h}_{i}=h_{i}\) with probability \(\lambda=\frac{e^{\alpha}}{2^{b}+e^{\alpha}}\) and otherwise draws \(\hat{h}_{i}\) uniformly from [2 \(\left.2^{b}\right]\).
    7: Server receives \(\hat{h}_{i}\) from each user \(i\).
    8: If pair \(q\) contains users \(i\) and \(j\) then let \(c_{q}=\mathbf{1}\left[\hat{h}_{i}=\hat{h}_{j}\right]\) indicate whether a hash collision was
    observed for pair \(q\).
    9: Server computes \(\bar{c}=\left(\frac{2}{(1-\lambda) n} \sum_{q} c_{q}\right)-\frac{\lambda}{2^{b}}\).
10: Server outputs \(\hat{G}=\frac{2^{b}}{2^{b}-1} \bar{c}-\frac{1}{2^{b}-1}\) and \(\hat{C}=-\log (1-\hat{G})\).
```

Algorithm 5 is based on the observation that if $X$ and $X^{\prime}$ are independent and identically distributed then the Gini entropy is equal to $1-\operatorname{Pr}\left[X=X^{\prime}\right]$ and the collision entropy is equal to $-\log \operatorname{Pr}[X=$ $\left.X^{\prime}\right]$. If the server observed each sample directly then it could estimate $\operatorname{Pr}\left[X=X^{\prime}\right]$ using the collision frequency, i.e., the fraction of sample pairs $\left(x_{i}, x_{j}\right)$ such that $x_{i}=x_{j}$. However, the server only observes a $b$-bit hash of each sample. Among sample pairs in which there is a true collision, all of them also produce a hash collision. Among samples pairs in which there is not a true collision, about a $\frac{1}{2^{b}}$ fraction of them produce a hash collision. Therefore the true collision frequency can be
estimated using an appropriately bias-corrected hash collision frequency, and the server uses this estimate to approximate the Gini and collision entropies.
The analysis of Algorithm 5 is given in Theorem 6.1 below. As is customary, for the analysis we assume that the hash function $h$ is constructed by assigning each element of its domain to a uniform random element of its range [6]. See the Appendix G]for the proof of the theorem.
Theorem 6.1. Algorithm 5 is $\alpha$-locally differentially private, has $\tilde{O}(b)$ communication complexity and $\tilde{O}(b)$ space complexity. Let $\hat{G}$ and $\hat{C}$ be the outputs of the algorithm. Let $\alpha, \epsilon, \delta \in(0,1)$. If $n=\Omega\left(\frac{b^{2} \max \left\{1-G(X), 2^{-b}\right\} \log \frac{1}{\delta}}{\alpha^{2} \epsilon^{2}}\right)$, then $|\hat{G}-G(X)| \leq \epsilon$ with probability at least $1-\delta$. Also, if $X$ has support size $k$ and $n=\Omega\left(\frac{b^{2} k^{2} \log \frac{1}{\delta}}{\alpha^{2} \epsilon^{2} \min \left\{k, 2^{b}\right\}}\right)$, then $|\hat{C}-C(X)| \leq \epsilon$ with probability at least $1-\delta$.

### 6.1 Comparison to Prior Work

Recall that Gini entropy is proportional to the second frequency moment. Local differentially private algorithms for estimating frequency moments were recently studied in [10]. Letting $b=1$ in Algorithm 5 yields a sample complexity of $\tilde{O}\left(1 / \alpha^{2} \epsilon^{2}\right)$, which is independent of the distribution's support size, unlike the sample complexity of the non-interactive algorithm for estimating the second frequency moment from [10]. Also our algorithm only uses 1 bit per sample and $\tilde{O}(1)$ space, while the previous algorithm uses $\Omega(k)$ bits per sample and $\Omega(k)$ space, where $k$ is the support size of the distribution. The authors in [10] asked whether there is a non-interactive algorithm for privately estimating frequency moments with a sample complexity that is independent of the distribution's support size. Here we affirmatively answer this open question for the second frequency moment.

The best known algorithm for estimating collision entropy using $\tilde{O}(1)$ space is due to [36]. The sample complexity of their algorithm is $\tilde{O}\left(k / \epsilon^{2}\right)$ and its communication complexity is $O(\log k)$ bits per user. Letting $b=\log k$ and $\alpha=O(1)$ in Algorithm 5 recovers these results (up to logarithmic factors), and using smaller values for $b$ or $\alpha$ generalizes the previous algorithm to the private and communication-efficient setting. Also, it was shown in [14] that (conditioned on a plausible conjecture) any algorithm that estimates collision entropy to within $O(1)$ error using $O(1)$ space requires $\Omega(k)$ samples. Therefore our algorithm is likely to be Pareto optimal with respect to the sample complexity-space complexity trade-off.

## 7 Experiments

In this section we present two sets of experiments to support our theoretical findings. First, we demonstrate that Algorithm 1 is indeed able to estimate the Shannon entropy of tree-structured distributions with linear sample complexity in $d$. Thus it has a superior sample complexity comparing to the state-of-the-art non-interactive method [12, 7], which has a quadratic sample complexity in $d$. The sample complexity is defined here in terms of number of observations from pairs of marginals. In the second set of experiments, we use our Algorithm 5 to estimate the collision entropy of discrete distributions, and compare its performance to that of the best-known communication efficient, non-private algorithm for this task (which we refer to as Skorski's algorithm [36]).
Estimating Shannon entropy: To estimate the Shannon entropy of a tree-structured distribution as given by eq. (6), the marginal entropies and the mutual information between certain or all pairs of marginals have to be estimated. The Chow-Liu algorithm estimates the mutual information between all pairs of marginals which results in quadratic sample complexity, whereas Algorithm 1 estimates the mutual information only for a linear fraction of pairs. Both algorithms estimate the marginal entropy values by sampling the marginals independently from the mutual information estimations, and they both use the same $\epsilon$ additive error and privacy budget $\alpha$ to estimate the mutual information between pairs of marginals. Thus it is fair to compare their performance in terms of number of pairs for which they estimate the mutual information. For the Chow-Liu algorithm this is always $d^{2}$, whereas our Algorithm 1 is randomized, thus we evaluate it over 100 repetitions and report the average.
We ran this experiment on random tree-structured joint distributions over $\{0,1\}^{d}$. To create a random tree-structured distribution, we first sample the skeleton of the distribution by taking the maximum spanning tree of a complete graph with $d$ nodes and edge weights distributed according to a standard normal. Next, we sample the parameters for each marginal uniformly at random from $[0,1]$, and
then we achieve the tree-structure by inducing dependence between pairs of variables while preserving the marginals. Specifically, for two marginals $X_{i}$ and $X_{j}$ with sampled parameters $p_{i}$ and $p_{j}$, we set: $P\left(X_{i}=0, X_{j}=0\right)=\left(1-p_{i}\right) \cdot\left(1-p_{j}\right)+r_{i j}, P\left(X_{i}=1, X_{j}=1\right)=p_{i} \cdot p_{j}+r_{i j}$, $P\left(X_{i}=0, X_{j}=1\right)=\left(1-p_{i}\right) \cdot p_{j}-r_{i j}$, and $P\left(X_{i}=1, X_{j}=0\right)=p_{i} \cdot\left(1-p_{j}\right)-r_{i j}$, where $r_{i j}$ is sampled uniformly at random from a range of values such that each probability stays positive. The results are displayed in Figure 1a It is clear that the sample size (i.e., number of mutual information estimate) is close to linear for our algorithm, whereas the Chow-Liu algorithm requires a larger sample size.


(b) Absolute error in estimating the collision entropy of an exponential distribution with domain size $k=1000$ for the Skorski's algorithm [36] and for our Algorithm 5

Estimating collision entropy: In this set of experiments, we drew samples from a discrete exponential distribution $p_{i} \propto e^{-i}$ with support size $k=1000$ and estimated the collision entropy using algorithm [36] (the previously best-known communication-efficient algorithm for this task) and our Algorithm5 (with and without local differential privacy). Algorithm [36] requires $O(\log k)$ bits per sample and is not private, while our algorithm only requires 1 bit per sample and is differentially private. The results are displayed in Figure 1 b . The previous algorithm has $5 \%$ estimation error after observing 10000 bits, while our algorithm has less than $3.5 \%$ estimation error. Thus our algorithm has lower error for the same communication cost while also being local differentially private.

## 8 Conclusion and future work

Estimating entropy is of importance in many practical applications. In this paper, we studied three widely used entropy measures: Shannon, Gini and collision entropy. We described estimation algorithms for each entropy that require minimal communication and satisfy local differential privacy. We also validated our theoretical results with simulations.

Our sequentially interactive algorithm for estimating Shannon entropy of high-dimensional treestructured distributions observes only two of these dimensions per sample and has a sample complexity $O\left(d / \epsilon^{3}\right)$. Our approach relies on the celebrated Chow-Liu approximation [12], providing a substantial improvement on the $\Omega\left(d^{2}\right)$ sample complexity of the original non-interactive Chow-Liu algorithm. We also identified two special cases (viz., when the underlying graphical model of the joint distribution is either a chain or star graph) and provided algorithms with a sample complexity of $\tilde{O}\left(d \log d / \epsilon^{2}\right)$. Our algorithm for Gini and collision entropy estimation also improved on the state-of-the-art, either by improving the sample complexity and communication complexity of previous work, or by generalizing the best known algorithm to the private and communication-efficient setting.
A natural extension of our work on Shannon entropy estimation is to consider higher-order correlations in the Chow-Liu decomposition [25]. In contrast to the second-order case which reduces to a maximum spanning tree problem, discovering the underlying structure of the joint distribution is already computationally challenging. However, efficiently estimating the entropy of the resulting distribution might still be possible.

## References

[1] Jayadev Acharya, Clement Canonne, Cody Freitag, and Himanshu Tyagi. Test without trust: Optimal locally private distribution testing. In Proceedings of the Twenty-Second International

Conference on Artificial Intelligence and Statistics, volume 89 of Proceedings of Machine Learning Research, pages 2067-2076. PMLR, 2019.
[2] Jayadev Acharya, Clément L Canonne, and Himanshu Tyagi. Inference under information constraints: Lower bounds from chi-square contraction. In Proceedings of the Thirty-Second Conference on Learning Theory, volume 99 of Proceedings of Machine Learning Research, pages 3-17. PMLR, 2019.
[3] Jayadev Acharya, Gautam Kamath, Ziteng Sun, and Huanyu Zhang. Inspectre: Privately estimating the unseen. In International Conference on Machine Learning, pages 30-39. PMLR, 2018.
[4] Jayadev Acharya and Ziteng Sun. Communication complexity in locally private distribution estimation and heavy hitters. In International Conference on Machine Learning, pages 51-60. PMLR, 2019.
[5] Jayadev Acharya, Ziteng Sun, and Huanyu Zhang. Hadamard response: Estimating distributions privately, efficiently, and with little communication. In The 22nd International Conference on Artificial Intelligence and Statistics, pages 1120-1129. PMLR, 2019.
[6] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In Proceedings of the 1st ACM Conference on Computer and Communications Security, pages 62-73, 1993.
[7] Arnab Bhattacharyya, Sutanu Gayen, Eric Price, and NV Vinodchandran. Near-optimal learning of tree-structured distributions by Chow-Liu. In Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing, pages 147-160, 2021.
[8] Béla Bollobás. Modern Graph Theory. Graduate Texts in Mathematics 184. 1998.
[9] Tai-Danae Bradley. Entropy as a topological operad derivation. Entropy, 23(9), 2021.
[10] Cristina Butucea and Yann Issartel. Locally differentially private estimation of functionals of discrete distributions. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, Advances in Neural Information Processing Systems, volume 34, pages 24753-24764. Curran Associates, Inc., 2021.
[11] Bernard Chazelle, Ronitt Rubinfeld, and Luca Trevisan. Approximating the minimum spanning tree weight in sublinear time. SIAM Journal on computing, 34(6):1370-1379, 2005.
[12] C. Chow and C. Liu. Approximating discrete probability distributions with dependence trees. IEEE Transactions on Information Theory, 14(3):462-467, 1968.
[13] Thomas M. Cover and Joy A. Thomas. Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing). Wiley-Interscience, New York, NY, USA, 2006.
[14] Michael Crouch, Andrew McGregor, Gregory Valiant, and David P Woodruff. Stochastic streams: Sample complexity vs. space complexity. In 24th Annual European Symposium on Algorithms (ESA 2016). Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2016.
[15] Constantinos Daskalakis and Qinxuan Pan. Tree-structured Ising models can be learned efficiently. arXiv e-prints, pages arXiv-2010, 2020.
[16] John Duchi and Ryan Rogers. Lower bounds for locally private estimation via communication complexity. In Conference on Learning Theory, pages 1161-1191. PMLR, 2019.
[17] Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. Found. Trends Theor. Comput. Sci., 9(3-4):211-407, 2014.
[18] Kazuto Fukuchi and Jun Sakuma. Minimax optimal estimators for additive scalar functionals of discrete distributions. In 2017 IEEE International Symposium on Information Theory (ISIT), pages 2103-2107. IEEE, 2017.
[19] C. Gini. Variabilitá e Mutuabilitá. Contributo allo Studio delle Distribuzioni e delle Relazioni Statistiche. Bologna: C. Cuppini., 1912.
[20] Yanjun Han, Ayfer Özgür, and Tsachy Weissman. Geometric lower bounds for distributed parameter estimation under communication constraints. In Conference On Learning Theory, pages 3163-3188. PMLR, 2018.
[21] Oc Herfindahl. Concentration in the US steel industry. 1950.
[22] Matthew Joseph, Jieming Mao, Seth Neel, and Aaron Roth. The role of interactivity in local differential privacy. In 2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS), pages 94-105. IEEE, 2019
[23] Matthew Joseph, Jieming Mao, Seth Neel, and Aaron Roth. The role of interactivity in local differential privacy. In David Zuckerman, editor, 60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9-12, 2019, pages 94-105. IEEE Computer Society, 2019.
[24] Jyrki Katajainen and Jesper Larsson Träff. A meticulous analysis of mergesort programs. In Italian Conference on Algorithms and Complexity, pages 217-228. Springer, 1997.
[25] Edith Kovács and Tamás Szántai. On the Approximation of a Discrete Multivariate Probability Distribution Using the New Concept of t-Cherry Junction Tree, pages 39-56. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010.
[26] Markku Laakso and Rein Taagepera. "Effective" number of parties: A measure with application to west europe. Comparative Political Studies, 12(1):3-27, 1979.
[27] Tom Leinster. Entropy and Diversity: The Axiomatic Approach. Cambridge University Press, 2021.
[28] Abolfazl S Motahari, Guy Bresler, and NC David. Information theory of DNA shotgun sequencing. IEEE Transactions on Information Theory, 59(10):6273-6289, 2013.
[29] Robert Clay Prim. Shortest connection networks and some generalizations. The Bell System Technical Journal, 36(6):1389-1401, 1957.
[30] Laura Elena Raileanu and Kilian Stoffel. Theoretical comparison between the Gini index and information gain criteria. Annals of Mathematics and Artificial Intelligence, 41(1):77-93, 2004.
[31] Florencia Rao. https://www.secureauth.com/blog/ the-art-of-war-of-browser-fingerprinting/, 2021.
[32] Alfréd Rényi. On measures of information and entropy. In Proceedings of the fourth Berkeley Symposium on Mathematics, Statistics and Probability, pages 547-561, 1960.
[33] C.E. Shannon. A mathematical theory of communication. Bell Syst. Tech. J., 27:379-423, 1948.
[34] Edward H Simpson. Measurement of diversity. nature, 163(4148):688-688, 1949.
[35] Yakov G. Sinai. On the notion of entropy of a dynamical system. 2010.
[36] Maciej Skorski. Evaluating entropy for true random number generators: Efficient, robust and provably secure. In International Conference on Information Security and Cryptology, pages 526-541, 032017.
[37] Ananda Theertha Suresh, X Yu Felix, Sanjiv Kumar, and H Brendan McMahan. Distributed mean estimation with limited communication. In International Conference on Machine Learning, pages 3329-3337. PMLR, 2017.
[38] Constantino Tsallis. Introduction to Nonextensive Statistical Mechanics. Springer New York, NY, Greece, 2009.
[39] Thomas Verma and Judea Pearl. Equivalence and synthesis of causal models. In Proceedings of the Sixth Annual Conference on Uncertainty in Artificial Intelligence, UAI '90, page 255-270, USA, 1990. Elsevier Science Inc.
[40] Yuchen Zhang, John C Duchi, Michael I Jordan, and Martin J Wainwright. Information-theoretic lower bounds for distributed statistical estimation with communication constraints. In NIPS, pages 2328-2336. Citeseer, 2013.

## Checklist

1. For all authors...
(a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
(b) Did you describe the limitations of your work? [Yes] Please see our Comparison to Prior Work and Discussion Sections.
(c) Did you discuss any potential negative societal impacts of your work? [N/A] Due to the theoretical nature of our work, we don't foresee any relevant societal impact.
(d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
(a) Did you state the full set of assumptions of all theoretical results? [Yes]
(b) Did you include complete proofs of all theoretical results? [Yes] All theorems are proved in the appendix.
3. If you ran experiments...
(a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [No] The main contribution of our work is theoretical. Our simulations are easy to reproduce and are meant to be a simple validation of our theoretical claims.
(b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
(c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
(d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
(a) If your work uses existing assets, did you cite the creators? [N/A] Our work builds on previous peer-reviewed results, which we properly cite.
(b) Did you mention the license of the assets? [N/A]
(c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

## A Proof of Theorem 5.1

Proof. [4, 5] showed that any discrete distribution can be learnt in total variation distance based on $O\left(c^{2} \log 1 / \delta /\left(\epsilon^{2} \alpha^{2}\right)\right.$ when $\alpha<1$. This result can be plug-in into Theorem 17.3.3 of [13].

## B Proof of Theorem 5.2

The result follows from simply combining the algorithm from [11] for estimating the weight of the maximum spanning tree in sublinear time, which assumes that each edge weight can be computed in $O(1)$ time, with Theorem 5.1 , which gives the number of samples needed to privately estimate an edge weight (i.e., the mutual information between two variables). The main technical complication we must overcome is that the result from [11] assumes that the edge weights are integers with a bounded ratio, while we instead discretize the edge weights with resolution $\epsilon$. To complete the proof, we need the Lemma B. 1 below, which concerns a graph $G$ whose edge weights are multiples of $\epsilon$. The lemma relates the weight of the maximum spanning tree of $G$ to the number of connected components in various subgraphs of $G$. This lemma replaces Claim 5 from [11].
Let $G$ be a connected graph with $n$ vertices and edge weights belonging to the set

$$
\{\epsilon, 2 \epsilon, \ldots, w \epsilon\}
$$

where $\epsilon>0$ and $w$ is a positive integer.
For each $i \in\{1, \ldots, w\}$ let $G_{i}$ be the subgraph of $G$ containing the same vertices as $G$ and only those edges of $G$ whose weight is at least $i \epsilon$. Thus $G_{1}=G$. Let $c_{i}$ be the number of connected components in $G_{i}$. Let $M$ be the weight of a maximum spanning tree of $G$.
Lemma B.1. $M=\epsilon w n-\epsilon \sum_{i=1}^{w} c_{i}$.
Proof. For each $i \in\{1, \ldots, w\}$ let $\gamma_{i}$ be the number of edges with weight $i \epsilon$ in a maximum spanning tree of $G$. We have

$$
\sum_{i<\ell} \gamma_{i}=c_{\ell}-1
$$

for any $\ell \in\{1, \ldots, w\}$, where the empty sum is defined to be zero. This equality can be established by considering Kruskal's greedy algorithm for constructing a maximum spanning tree, which adds edges in decreasing order of weight as long as they do not induce a cycle. Since $G_{\ell}$ has $c_{\ell}$ connected components, and all the edges in $G_{\ell}$ are heavier than all the edges not in $G_{\ell}$, the greedy algorithm must first connect the vertices within each component of $G_{\ell}$ and then use exactly $c_{\ell}-1$ edges not in $G_{\ell}$ to connect the components to each other. Thus we have

$$
\begin{aligned}
M & =\sum_{i=1}^{w} i \epsilon \gamma_{i} \\
& =\epsilon\left(\sum_{i \geq 1} \gamma_{i}+\sum_{i \geq 2} \gamma_{i}+\cdots+\sum_{i \geq w} \gamma_{i}\right) \\
& =\epsilon\left(n-1-\sum_{i<1} \gamma_{i}+n-1-\sum_{i<2} \gamma_{i}+\cdots+n-1-\sum_{i<w} \gamma_{i}\right) \\
& =\epsilon\left(w(n-1)-\sum_{i=1}^{w}\left(c_{i}-1\right)\right) \\
& =\epsilon w n-\epsilon \sum_{i=1}^{w} c_{i}
\end{aligned}
$$

## C Proof of Theorem 5.3

Proof. We start by recalling a lemma that applies to tree-structured distributions.

Lemma C.1. [39] Let $\mathbf{X}=\left(X_{1}, \ldots, X_{d}\right)$ be tree decomposable with tree $T$. Then for any triplets $i, j$ and $k$, if $k$ is on the unique path in $T$ between $i$ and $j$, then

$$
I\left(X_{i} ; X_{j} \mid X_{k}\right)=0
$$

Next we show that the output Algorithm 2 is correct with high probability. We make use of the Conditional Mutual Information Tester of [7]. This testing algorithm consists of estimating the CMI using the plug-in estimator and then applying a $\epsilon$ threshold on the estimate, i.e., if the estimate is smaller than $\epsilon$ then accept, otherwise reject. The sample complexity of Conditional Mutual Information Tester is $O\left(\frac{|\Sigma|^{3}}{\epsilon} \log \frac{d|\Sigma|}{\delta} \log \frac{|\Sigma| \log d / \delta}{\epsilon}\right)$ according to Theorem 1.3 of [7]. Thus if we apply this tester with adjusted confidence parameter, i.e., $\delta / d \log _{3} d$, then the union bound implies that the output of all tests is correct with probability at least $1-\delta$.
Next, note that for any triplet $X_{i}, X_{j}, X_{k}$ such that $X_{i}$ is between $X_{j}$ and $X_{k}$ in the chain, it holds that

$$
I\left(X_{j} ; X_{i}\right)-I\left(X_{j} ; X_{k}\right)=\underbrace{I\left(X_{j} ; X_{i} \mid X_{k}\right)}_{>\epsilon}-\underbrace{I\left(X_{j} ; X_{k} \mid X_{i}\right)}_{=0 \text { due to Lemma C.1 }}
$$

as the edges are different with a margin of $\epsilon$. The same argument implies that $I\left(X_{k} ; X_{i} \mid X_{j}\right)>\epsilon$. Thus, Algorithm 2 divides the nodes correctly in Line 4-10 which along with the testers' correctness with high probability implies the correctness of the algorithm.
Using the result from [7] that the sample complexity of Conditional Mutual Information Tester is $O\left(\frac{|\Sigma|^{3}}{\epsilon} \log \frac{d|\Sigma|}{\delta} \log \frac{|\Sigma| \log d / \delta}{\epsilon}\right)$, we only need to upper bound the number of tests for deciding whether $I\left(X_{i} ; X_{j} \mid X_{k}\right)>\epsilon$ that is carried out by Algorithm 2. It is indeed $3 \sum_{i=3}^{d} \log _{3} i \in$ $O\left(d \log _{3} d\right)$ which concludes the proof.

## D Proof of Theorem 5.4

Proof. First note that the Prim step in Line 3 of Algorithm 4 indeed finds an edge that is in the maximum spanning tree due to the assumption $\left|I\left(X_{i} ; X_{j}\right)-I\left(X_{j} ; X_{k}\right)\right| \geq \epsilon$. Say this edge is between $X_{i}$ and $X_{k}$, what remains to be decided is whether $X_{i}$ or $X_{k}$ is the center of the graph. This can be done by comparing $\sum_{j} \hat{I}\left(X_{i}, X_{j}\right)$ and $\sum_{j} \hat{I}\left(X_{k}, X_{j}\right)$.
So Algorithm 4 estimates $2 d$ mutual information, which requires $4 d$ marginal and $2 d$ pairwise marginal entropy estimation. Recall from Theorem 5.1 that entropy can estimated $\alpha$-locally differential private with $\epsilon$ error using $O\left(c^{2} \log \frac{1}{\delta} /\left(\epsilon^{2} \alpha^{2}\right)\right)$ samples. Thus by the union bound the algorithm is correct if the confidence parameter is set to $\delta / 6 d$, and accordingly has sample complexity $O\left(\frac{6 d c^{2}}{\epsilon^{2} \alpha^{2}} \log \frac{6 d}{\delta}\right)$.

## E Proof of Theorem 5.5

Proof. Assume a deterministic algorithm $\mathcal{A}$ that takes sub-quadratic samples from $\mathbf{X}$ and estimates $H(\mathbf{X})$. In addition, assume that its sample complexity is $o\left(d^{2}\right)$. Thus for any constant $C>0$, there exists $d_{0}$ such that for any $d>d_{0}$, the sample complexity of the algorithm is $<C d^{2}$. This implies that, if $d$ is large enough, the algorithm needs $\leq d^{2-\kappa}$. In addition, we can pick $d$ so as $d^{-\kappa}<\kappa$ which implies $d^{2}-d^{2-\kappa}>(1-\kappa) d^{2}$. Thus any deterministic algorithm which takes sub-quadratic sample size never observes $(1-\kappa) d^{2}$ edges for $\forall \kappa$ when $d$ is large enough.
Let us recall that Turán's theorem
Theorem E. 1 ([8]). Let $G(V, E)$ be a graph with d vertices (i.e, $|V|=d$ ) that does not contain a $(\ell+1)$-clique as a subgraph. Then $G$ has at most $\frac{(\ell-1) d^{2}}{2 \ell}$ edges.

In addition, recall that a Turán's graph $G_{T}(d, \ell)[8]$ is defined as the unique graph with $d$ nodes that does not contain a $(\ell+1)$-clique and has the maximum possible number of edges which is $\left\lfloor\frac{(\ell-1) d^{2}}{2 \ell}\right\rfloor$. Let $t(d, \ell)$ denote the number of edges in $G_{T}(d, \ell)$.

Thus for any $d>0$ there exists a graph $G=(V, E)$ such that $|V|=d$ and $|E|>t(d, k)$ which contains a $(k+1)$-clique. This implies that if algorithm $\mathcal{A}$ does not observes at least $t(d, k)$ edges of $G$, then $G$ contains a $k+1$-clique whose edges are never observed by algorithm $\mathcal{A}$. Now we apply Turan's result with $\ell=d / 2$ which implies that

$$
t(d, d / 2)=\left\lfloor(1 / 2-1 / d) d^{2}\right\rfloor
$$

Thus for any $\kappa \in(0,1 / 2)$ and any algorithm with sample complexity $o\left(d^{2}\right)$, if $d$ is large enough, then there will be a $\Theta(d / 2)$-clique for which the algorithm does not observe any edge within this clique.

Finally it is easy to construct two $d / 2$-dimensional problem instances, denoted it $S$ and $S^{\prime}$, with joint entropy that differs by $\Omega(d)$ : take $d / 2$ Bernoulli with parameter $1 / 2$ and take the copy of the same Bernoulli $d / 2$ times. The entropy for these two joint distributions is $d / 2$ and 1 , respectively. This also implies that for any deterministic algorithm we can construct two problem instances which contains $S$ and $S^{\prime}$ so as they are independent from the rest of the marginals and the algorithm does not observe any sample from them, and hence it cannot achieve $o(d)$ additive error.

## F Proof of Theorem 5.6

Proof. Let $\hat{\theta}_{n}=\hat{\theta}\left(x_{1}, \ldots, x_{n}\right)$ such that $\hat{\theta}_{n}:\left(\Sigma^{d}\right)^{n} \mapsto \mathbb{R}$ be an estimator using $n$ samples.
Theorem F.1. [Le Cam's theorem] Let $\mathcal{P}$ be a set of distributions. Then, for any pair of distributions $P_{0}, P_{1} \in \mathcal{P}$, we have

$$
\inf _{\hat{\theta}} \max _{P \in \mathcal{P}} \mathbb{E}_{P}\left[d\left(\hat{\theta}_{n}(P), \theta(P)\right)\right] \geq \frac{d\left(\theta\left(P_{0}\right), \theta\left(P_{1}\right)\right)}{8} e^{-n d_{K L}\left(P_{0}, P_{1}\right)}
$$

where $\theta(P)$ is a parameter taking values in a metric space with metric $d$, and $\hat{\theta}_{n}$ is the estimator of $\theta$ based on $n$ samples.

Let us consider two Bernoulli distributions $P_{0}$ and $P_{1}$ with parameters $p_{0}=1 / 2$ and $p_{1}=1 / 2-\epsilon$, where $\epsilon \in(0,1 / 2)$. The entropy of random variables $X_{0}$ and $X_{1}$ distributed according to $P_{0}$ and $P_{1}$ are $H\left(X_{0}\right)=1$ and

$$
H\left(X_{1}\right)=-\left(\frac{1}{2}-\epsilon\right) \log _{2}\left(\frac{1}{2}-\epsilon\right)-\left(\frac{1}{2}+\epsilon\right) \log _{2}\left(\frac{1}{2}+\epsilon\right)
$$

Thus,

$$
\begin{aligned}
\left|H\left(X_{0}\right)-H\left(X_{1}\right)\right|=H\left(X_{0}\right)-H\left(X_{1}\right) & =\left(\epsilon+\frac{1}{2}\right) \log _{2}(1+2 \epsilon)-\left(\epsilon-\frac{1}{2}\right) \log _{2}(1-2 \epsilon) \\
& \geq\left(\epsilon+\frac{1}{2}\right) \frac{2 \epsilon}{1+2 \epsilon}-\left(\frac{1}{2}-\epsilon\right) 2 \epsilon \\
& \geq 2 \epsilon^{2}
\end{aligned}
$$

where we used that $\log (1-2 \epsilon) \leq-\epsilon$ for $0<\epsilon<1$ and $\frac{\epsilon}{1+\epsilon} \leq \log (1+\epsilon)$ for $\epsilon>-1$. The KL divergence can be upper bounded as

$$
d_{\mathrm{KL}}\left(P_{0}, P_{1}\right)=-\frac{1}{2} \log _{2}\left(1-4 \epsilon^{2}\right) \leq 2 \epsilon^{2}
$$

We can now apply the Le Cam's theorem for the set of Bernoulli distributions with metric $d$ being the $\ell_{1}$-norm as

$$
\inf _{\hat{\theta}} \max _{P \in \mathcal{P}} \mathbb{E}_{P}\left[\left|\hat{\theta}_{n}(P)-H(P)\right|\right] \geq \frac{d\left(\theta\left(P_{0}\right), \theta\left(P_{1}\right)\right)}{8} e^{-n d_{\mathrm{KL}}\left(P_{0}, P_{1}\right)} \geq \frac{\epsilon^{2}}{4} e^{-2 n \epsilon^{2}}
$$

Using this result with $\epsilon^{\prime}=\sqrt{\epsilon}$, the following sample complexity can be obtained for estimating Shannon entropy.
Corollary F.2. For any $\hat{\theta}_{n}$ such that $n \in o(1 / \epsilon)$, there exists a Bernoulli distribution $P$ for which

$$
\mathbb{E}_{P}\left[\left|\hat{\theta}_{n}(P)-H(P)\right|\right] \geq C \cdot \epsilon
$$

with $C>0$ that does not depend on $\epsilon$.

First, note that there is some bound on error $r(\delta)$, either lower or upper, that holds with probability $1-\delta$. This translate into the bound $r(\delta)+\delta$ on the expected error in a straightforward manner. Thus the lower bound presented in Corollary F. 2 also implies that there is no high probability estimator for entropy with $o(1 / \epsilon)$ sample complexity for discrete distributions. This can be used to lower bound of the entropy estimator for joint distribution as follows.
Let $\mathcal{B}=\left\{\mathbf{b}=\left(b_{1}, \ldots, b_{d}\right): b_{j} \in\{0,1\}\right\}$ be the vertices of the $d$ dimensional hypercube, and define a set of $d$-dimensional distribution $\mathcal{P}_{\mathbf{b}}$ indexed by the element of $\mathcal{B}$. Each $P_{\mathbf{b}} \in \mathcal{P}$ contains $X_{0} \sim \operatorname{Bern}(1 / 2)$ if $b_{i}=0$ and $X_{1} \sim \operatorname{Bern}(1 / 2-\epsilon)$ if $b_{i}=1$, i.e.

$$
P_{\mathbf{b}}=X_{b_{1}} \oplus \cdots \oplus X_{b_{d}}
$$

and

$$
\mathcal{P}=\left\{P_{\mathbf{b}}: \mathbf{b} \in\{0,1\}^{d}\right\} .
$$

It is clear that $\mathcal{P}$ is a subset of the tree-structured distributions and each distribution contains $d$ independent Bernoulli random variables, thus

$$
H\left(P_{\mathbf{b}}\right)=\sum_{i=1}^{d} H\left(X_{b_{i}}\right)
$$

Therefore any estimator that achieves at most $\epsilon \cdot d$ additive error for $H\left(P_{\mathbf{b}}\right)$ has to estimate each individual Bernoulli distribution with at most $\epsilon$ error. The sample complexity of any estimator of $H\left(P_{\mathbf{b}}\right)$ with an additive error $O(\epsilon d)$ is $\Omega(d / \epsilon)$.

## G Proof of Theorem 6.1

Proof. Algorithm 5 clearly has $\tilde{O}(b)$ communication complexity and $\tilde{O}(b)$ space complexity, since it only has to maintain a counter of collisions between $b$-bit hashes. Each user replaces their hash with a random hash with probability $\lambda$, and therefore the algorithm is $\alpha$-local differentially private, since $\log \left(\frac{\lambda}{(1-\lambda) / 2^{b}}\right)=\log \left(\frac{\lambda 2^{b}}{1-\lambda}\right)=\alpha$ where we used $\lambda=\frac{e^{\alpha}}{e^{\alpha}+2^{b}}$. Before proving the sample complexity we will first prove the following result.
Lemma G.1. Let $\hat{G}$ be output by Algorithm 5 Let $\epsilon, \delta \in(0,1)$. If

$$
n \geq \frac{6 b^{2} \log \frac{2}{\delta}}{\alpha^{2} \epsilon^{2}\left(\left(1-\frac{1}{2^{b}}\right) G(X)+\frac{1}{2^{b}}\right)}
$$

then $|\hat{G}-G(X)| \leq \epsilon\left(G(X)+\frac{1}{2^{b}-1}\right)$ with probability at least $1-\delta$.
Proof. Recall that if $X$ and $X^{\prime}$ are independent and identically distributed then

$$
G(X)=1-\operatorname{Pr}\left[X=X^{\prime}\right]
$$

We will calculate the expected value of each $c_{q}$. Suppose pair $q$ contains samples $x_{i}$ and $x_{j}$. If $x_{i}=x_{j}$ then $c_{q}=1$ with probability $1-\lambda+\frac{\lambda}{2^{b}}$, and otherwise $c_{q}=1$ with probability $\frac{1}{2^{b}}$. Thus

$$
\begin{aligned}
\mathrm{E}\left[c_{q}\right] & =\operatorname{Pr}\left[c_{q}=1 \mid x_{i}=x_{j}\right] \operatorname{Pr}\left[x_{i}=x_{j}\right]+\operatorname{Pr}\left[c_{q}=1 \mid x_{i} \neq x_{j}\right] \operatorname{Pr}\left[x_{i} \neq x_{j}\right] \\
& =\left(1-\lambda+\frac{\lambda}{2^{b}}\right) \operatorname{Pr}\left[x_{i}=x_{j}\right]+\frac{1}{2^{b}} \operatorname{Pr}\left[x_{i} \neq x_{j}\right] \\
& =\left(1-\lambda+\frac{\lambda}{2^{b}}\right) \operatorname{Pr}\left[x_{i}=x_{j}\right]+\frac{1}{2^{b}}\left(1-\operatorname{Pr}\left[x_{i}=x_{j}\right]\right) \\
& =\left(1-\lambda+\frac{\lambda}{2^{b}}-\frac{1}{2^{b}}\right) \operatorname{Pr}\left[x_{i}=x_{j}\right]+\frac{1}{2^{b}} \\
& =\left(1-\lambda+\frac{\lambda}{2^{b}}-\frac{1}{2^{b}}\right)(1-G(X))+\frac{1}{2^{b}}
\end{aligned}
$$

where the last line follows because $x_{i}$ and $x_{j}$ independent samples from the distribution of $X$.

Recall that by the Chernoff bound if $z_{1}, \ldots, z_{m}$ are independent random variables such that $z_{i} \in$ $\{0,1\}$ then for all $\epsilon \in(0,1)$ the average $\bar{z}=\left(z_{1}+\cdots z_{m}\right) / m$ satisfies

$$
\begin{aligned}
& \operatorname{Pr}[\bar{z} \geq(1+\epsilon) \mathrm{E}[\bar{z}]] \leq \exp \left(-\frac{\epsilon^{2} m}{3} \mathrm{E}[\bar{z}]\right), \text { and } \\
& \operatorname{Pr}[\bar{z} \leq(1-\epsilon) \mathrm{E}[\bar{z}]] \leq \exp \left(-\frac{\epsilon^{2} m}{3} \mathrm{E}[\bar{z}]\right)
\end{aligned}
$$

The $c_{q}$ 's are independent random variables because each $c_{q}$ is defined using a distinct pair of samples and distinct pair index. Also, each $c_{q} \in\{0,1\}$. Note that we proved above that the average of the $c_{q}$ 's is $\frac{2}{n} \sum_{q} \mathrm{E}\left[c_{q}\right]=\left(1-\lambda+\frac{\lambda}{2^{b}}-\frac{1}{2^{b}}\right)(1-G(X))+\frac{1}{2^{b}}$. Therefore

$$
\begin{aligned}
\operatorname{Pr}\left[\hat{G} \geq G(X)+\epsilon\left(G(X)+\frac{1}{2^{b}-1}\right)\right] & =\operatorname{Pr}\left[\frac{2^{b}}{2^{b}-1} \bar{c}-\frac{1}{2^{b}-1} \geq G(X)+\epsilon\left(G(X)+\frac{1}{2^{b}-1}\right)\right] \\
& =\operatorname{Pr}\left[\bar{c} \geq(1+\epsilon)\left(\left(1-\frac{1}{2^{b}}\right) G(X)+\frac{1}{2^{b}}\right)\right] \\
& =\operatorname{Pr}[\bar{c} \geq(1+\epsilon) \mathrm{E}[\bar{c}]] \\
& \leq \exp \left(-\frac{\epsilon^{2} n}{6} \mathrm{E}[\bar{c}]\right) \\
& =\exp \left(-\frac{\alpha^{2} \epsilon^{2} n}{6 b^{2}}\left(\left(1-\frac{1}{2^{b}}\right)(1-G(X))+\frac{1}{2^{b}}\right)\right)
\end{aligned}
$$

where the inequality follows from the Chernoff upper bound. By a very similar calculation

$$
\begin{aligned}
\operatorname{Pr}\left[\hat{G} \leq G(X)-\epsilon\left(G(X)+\frac{1}{2^{b}-1}\right)\right] & =\operatorname{Pr}\left[\frac{2^{b}}{2^{b}-1} \bar{c}-\frac{1}{2^{b}-1} \leq G(X)-\epsilon\left(G(X)+\frac{1}{2^{b}-1}\right)\right] \\
& =\operatorname{Pr}\left[\bar{c} \leq(1-\epsilon)\left(\left(1-\frac{1}{2^{b}}\right) G(X)+\frac{1}{2^{b}}\right)\right] \\
& =\operatorname{Pr}[\bar{c} \leq(1-\epsilon) \mathrm{E}[\bar{c}]] \\
& \leq \exp \left(-\frac{\epsilon^{2} n}{6} \mathrm{E}[\bar{c}]\right) \\
& =\exp \left(-\frac{\alpha^{2} \epsilon^{2} n}{6 b^{2}}\left(\left(1-\frac{1}{2^{b}}\right)(1-G(X))+\frac{1}{2^{b}}\right)\right)
\end{aligned}
$$

where the inequality follows from the Chernoff lower bound. Combining the above we have
$\operatorname{Pr}\left[|\hat{G}-G(X)| \geq \epsilon\left(1-G(X)+\frac{1}{2^{b}-1}\right)\right] \leq 2 \exp \left(-\frac{\alpha^{2} \epsilon^{2} n}{6 b^{2}}\left(\left(1-\frac{1}{2^{b}}\right)(1-G(X))+\frac{1}{2^{b}}\right)\right)$
and rearranging proves the lemma.
Now the first sample complexity bound in the theorem follows immediately from Lemma G. 1 As for the second sample complexity bound, since $C(X)=-\log (1-G(X))$ we have

$$
\begin{align*}
|C(X)-\hat{C}|=|\log (1-\hat{G})-\log (1-G(X))| & =\left|\log \frac{1-\hat{G}}{1-G(X)}\right| \\
& \leq \log \left(1+\frac{|\hat{G}-G(X)|}{1-G(X)}\right) \leq \frac{|\hat{G}-G(X)|}{1-G(X)} \tag{7}
\end{align*}
$$

Observe that $G(X) \in\left[\frac{1}{k}, 1\right]$ for any random variable $X$ with support size $k$. Now we consider two cases. First, assume $2^{b}>k$. Since
$n \geq \frac{24 k^{2} b^{2} \log \frac{2}{\delta}}{\alpha^{2} \epsilon^{2} \min \left\{k, 2^{b}\right\}}=\frac{6 b^{2} \log \frac{2}{\delta}}{\alpha^{2}\left(\frac{\epsilon}{2}\right)^{2} \frac{1}{k}} \geq \frac{6 b^{2} \log \frac{2}{\delta}}{\alpha^{2}\left(\frac{\epsilon}{2}\right)^{2}\left(\left(1-\frac{1}{2^{b}}\right) \frac{1}{k}+\frac{1}{2^{b}}\right)} \geq \frac{6 b^{2} \log \frac{2}{\delta}}{\alpha^{2}\left(\frac{\epsilon}{2}\right)^{2}\left(\left(1-\frac{1}{2^{b}}\right)(1-G(X))+\frac{1}{2^{b}}\right)}$
we have by Lemma G. 1 that

$$
|\hat{G}-G(X)| \leq \frac{\epsilon}{2}\left(1-G(X)+\frac{1}{2^{b}-1}\right) \leq \epsilon(1-G(X))
$$

where the second inequality uses the fact that $2^{b}>k$ implies $\frac{1}{2^{b}-1} \leq \frac{1}{k} \leq 1-G(X)$. Combining with Eq. (7) we have $|\hat{C}-C(X)| \leq \frac{|\hat{G}-G(X)|}{1-G(X)} \leq \epsilon$.

Next, assume $2^{b} \leq k$. We have

$$
n \geq \frac{24 b^{2} k^{2} \log \frac{2}{\delta}}{\alpha^{2} \epsilon^{2} \min \left\{k, 2^{b}\right\}}=\frac{6 b^{2} \log \frac{2}{\delta}}{\alpha^{2}\left(\frac{\epsilon}{2 k}\right)^{2} 2^{b}} \geq \frac{6 b^{2} \log \frac{2}{\delta}}{\alpha^{2}\left(\frac{\epsilon}{2 k}\right)^{2}\left(\left(1-\frac{1}{2^{b}}\right)(1-G(X))+\frac{1}{2^{b}}\right)}
$$

where the second inequality uses $2^{b} \geq 1$ and $1-G(X) \leq 1$. Thus by Lemma G. 1

$$
|\hat{G}-G(X)| \leq \frac{\epsilon}{2 k}\left(1-G(X)+\frac{1}{2^{b}-1}\right) \leq \frac{\epsilon}{k}
$$

Combining with Eq. (7) we have

$$
|\hat{C}-C(X)| \leq \frac{|\hat{G}-G(X)|}{1-G(X)} \leq k|\hat{G}-G(X)| \leq \epsilon
$$

