

The Price of Selection in Differential Privacy

Mitali Bafna*

IIT Madras

MITALI.BAFNA@GMAIL.COM

Jonathan Ullman

Northeastern University

JULLMAN@CCS.NEU.EDU

Abstract

In the differentially private top- k selection problem, we are given a dataset $X \in \{\pm 1\}^{n \times d}$, in which each row belongs to an individual and each column corresponds to some binary attribute, and our goal is to find a set of $k \ll d$ columns whose means are approximately as large as possible. Differential privacy requires that our choice of these k columns does not depend too much on any one individual’s dataset. This problem can be solved using the well known exponential mechanism and composition properties of differential privacy. In the high-accuracy regime, where we require the error of the selection procedure to be smaller than the so-called sampling error $\alpha \approx \sqrt{\ln(d)/n}$, this procedure succeeds given a dataset of size $n \gtrsim k \ln(d)$.

We prove a matching lower bound, showing that a dataset of size $n \gtrsim k \ln(d)$ is necessary for private top- k selection in this high-accuracy regime. Our lower bound shows that selecting the k largest columns requires more data than simply estimating the value of those k columns, which can be done using a dataset of size just $n \gtrsim k$.

Keywords: Differential privacy; lower bounds; variable selection

1. Introduction

The goal of privacy-preserving data analysis is to enable rich statistical analysis of a sensitive dataset while protecting the privacy of the individuals who make up that dataset. It is especially desirable to ensure *differential privacy* [Dwork et al. \(2006\)](#), which ensures that no individual’s information has a significant influence on the information released about the dataset. The central problem in differential privacy research is to determine precisely what statistics can be computed by differentially private algorithms and how accurately they can be computed.

The seminal work of [Dinur and Nissim \(2003\)](#) established a “price of privacy”: If we want to release the answer to $\gtrsim n$ statistics on a dataset of n individuals, then we must add noise whose magnitude is at least the same order as the *sampling error* of $\approx 1/\sqrt{n}$, or else an attacker can reconstruct nearly all of the sensitive information in the dataset, violating any reasonable notion of privacy. For example, if we have a dataset $X = (x_1, \dots, x_n) \in \{\pm 1\}^{n \times d}$ and we want to privately approximate its *marginal vector* $q = \frac{1}{n} \sum_{i=1}^n x_i$, then it suffices to introduce error of magnitude $\Theta(\sqrt{d}/n)$ to each entry of q [Dinur and Nissim \(2003\)](#); [Dwork and Nissim \(2004\)](#); [Blum et al. \(2005\)](#); [Dwork et al. \(2006\)](#), and this amount of error is also necessary [Bun et al. \(2014\)](#); [Steinke and Ullman \(2017a\)](#). Thus, when $d \gg n$, the error must be asymptotically larger than the sampling error.

Top- k Selection. In many settings, we release the marginals of the dataset in order to find a small set of “interesting” marginals, and we don’t need the entire vector. For example, we may be interested

* This work was done while the author was a visiting scholar at Northeastern University.

in finding only the attributes that are unusually frequent in the dataset. Thus, an appealing approach to overcome the limitations on computing marginals is to find only the *top-k* (approximately) largest coordinates of the marginal vector q , up to some error α .¹

Once we find these k coordinates, we can approximate the corresponding marginals with additional error $O(\sqrt{k}/n)$. But, how much error must we have in the *top-k* selection itself? The simplest way to solve this problem is to greedily find k coordinates using the differentially private *exponential mechanism* McSherry and Talwar (2007). This approach finds the *top-k* marginals up to error $\tilde{O}(\sqrt{k} \log(d)/n)$. The *sparse vector algorithm* Dwork et al. (2010a); Roth and Roughgarden (2010); Hardt and Rothblum (2010) would provide similar guarantees.

Thus, when $k \ll d$, we can find the *top-k* marginals and approximate their values with much less error than approximating the entire vector of marginals. However, the bottleneck in this approach is the $\sqrt{k} \log(d)/n$ error in the selection procedure, and this $\log(d)$ factor is significant in very high-dimensional datasets. For comparison, the sampling error for *top-k* selection is $\approx \sqrt{\log(d)/n}$ so the error introduced is asymptotically larger than the sampling error when $k \log(d) \gg n$. However, the best known lower bound for *top-k* selection follows by scaling down the lower bounds for releasing the entire marginal vector, and say that the error must be $\Omega(\sqrt{k}/n)$.

Top-k selection is a special case of fundamental data analysis procedures like variable selection and sparse regression. Moreover, private algorithms for selection problems underlie many powerful results in differential privacy: private control of false discovery rate Dwork et al. (2015b), algorithms for answering exponentially many queries Roth and Roughgarden (2010); Hardt and Rothblum (2010); Gupta et al. (2012); Jain and Thakurta (2012); Ullman (2015), approximation algorithms Gupta et al. (2010), frequent itemset mining Bhaskar et al. (2010), sparse regression Smith and Thakurta (2013), and the optimal analysis of the generalization error of differentially private algorithms Bassily et al. (2016). Therefore it is important to precisely understand optimal algorithms for differentially private *top-k* selection.

Our main result says that existing differentially private algorithms for *top-k* selection are essentially optimal in the high-accuracy regime where the error is required to be asymptotically smaller than the sampling error.

Theorem 1 (Lower Bound for Approximate Top-k) *There exist functions $n = \Omega(k \log(d))$ and $\alpha = \Omega(\sqrt{\log(d)/n})$ such that for every d and every $k = d^{o(1)}$ (and larger than some absolute constant), there is no differentially private algorithm M that takes an arbitrary dataset $X \in \{\pm 1\}^{n \times d}$ and (with high probability) outputs an α -accurate *top-k* marginal vector for X .*

Subsequent to our work, Steinke and the second author proved an $n = \Omega(\sqrt{k} \log(d))$ lower bound in the low-accuracy regime where $\alpha = \frac{1}{10}$ Steinke and Ullman (2017b), which is optimal up to constant factors. Their result is incomparable to Theorem 1, since the lower bound on n is (necessarily) quantitatively weaker, but holds for algorithms that add much more noise than the sampling error.

The Role of Distributional Assumptions. The lower bound in Theorem 1 holds even if the dataset X is chosen uniformly at random, and accuracy is only required to hold with high probability over the choice of the dataset X and the coins of the algorithm M . A natural question to ask is whether the lower bound also holds when the data is chosen according to some large family of distributions.

1. Here, the algorithm has error α if it returns a set $S \subseteq \{1, \dots, d\}$ consisting of k coordinates, and for each coordinate $j \in S$, $q_j \geq \tau - \alpha$, where τ is the k -th largest value among all the coordinates $\{q_1, \dots, q_d\}$.

For example, lower bounds for releasing the entire marginal vector [Dwork et al. \(2015a\)](#) hold under the weaker assumption that the columns are independent and each bit of the dataset has some entropy. However, for top- k statistics, some stronger assumptions appear to be necessary, because our lower bound can be circumvented in some natural cases. The simplest such case is when there is a gap γ between the k -th largest marginal and the $(k + 1)$ -st largest marginal, in which case the propose-test-release framework of [Dwork and Lei \(2009\)](#) can be used to find the exact top- k vector using a dataset of size $n = O_\gamma(1)$ (see e.g. [\(Dwork et al., 2014, Proposition 7.2\)](#) for a precise statement). There are several other algorithms that can identify the top- k marginals using $n \ll k \log(d)$ samples under weaker or incomparable assumptions on the dataset, such as the large-margin mechanism of [Chaudhuri et al. \(2014\)](#) and the choosing mechanism [Beimel et al. \(2013\)](#). An interesting future direction is to characterize which distributional assumptions are sufficient to bypass our lower bound.

Tracing Attacks. Our lower bounds for differential privacy follow from a *tracing attack* [Homer et al. \(2008\)](#); [Sankararaman et al. \(2009\)](#); [Bun et al. \(2014\)](#); [Steinke and Ullman \(2017a\)](#); [Dwork et al. \(2015a, 2017\)](#). In a tracing attack, the dataset X consists of data for n individuals drawn iid from some known distribution over $\{\pm 1\}^d$. The attacker is given data for a *target individual* $y \in \{\pm 1\}^d$ who is either one of the individuals in X (“IN”), or is an independent draw from the same distribution (“OUT”). The attacker is given some statistics about X (e.g. the top- k statistics) and has to determine if the target y is in or out of the dataset. Tracing attacks are a significant privacy violation, as mere presence in the dataset can be sensitive information, for example if the dataset represents the case group in a medical study [Homer et al. \(2008\)](#).

Our results give a tracing attack for top- k statistics in the case where the dataset is drawn uniformly at random. For simplicity, we state the properties of our tracing attack for the case of the exact top- k marginals. We refer the reader to Section 3 for a detailed statement in the case of approximate top- k marginals, which is what we use to establish Theorem 1.

Theorem 2 (Tracing Attack for Exact Top- k) *For every $\rho > 0$, every $n \in \mathbb{N}$, and all parameters $\ln(1/\rho) \ll k \ll d \ll 2^n$ such that $k \log(d/k) \geq O(n \log(1/\rho))$, there exists a tracing attack $\mathcal{A} : \{-1, 1\}^d \times \{0, 1\}^d \rightarrow \{\text{IN}, \text{OUT}\}$ such that the following holds: If we choose $X = (x_1, \dots, x_n) \in \{\pm 1\}^{n \times d}$ uniformly at random, and $t(X)$ is the exact top- k vector² of X , then*

1. *If $y \in \{\pm 1\}^d$ is uniform and independent of X , then $\mathbb{P}[A(y, t(X)) = \text{OUT}] \geq 1 - \rho$, and*
2. *for every $i \in [n]$, $\mathbb{P}[A(x_i, t(X)) = \text{IN}] \geq 1 - \rho$.*

1.1. Proof Overview

Our results use a variant of the *inner product attack* introduced in [Dwork et al. \(2015a\)](#) (and inspired by the work on *fingerprinting codes* [Boneh and Shaw \(1998\)](#); [Tardos \(2008\)](#) and their connection to privacy [Ullman \(2013\)](#); [Bun et al. \(2014\)](#); [Steinke and Ullman \(2017a\)](#)). Given a target individual $y \in \{\pm 1\}^d$, and a top- k vector $t \in \{\pm 1\}^d$, the attack is

$$\mathcal{A}(y, t) = \begin{cases} \text{IN} & \text{if } \langle y, t \rangle \geq \tau \\ \text{OUT} & \text{otherwise} \end{cases}$$

2. Due to the presence of ties, there is typically not a unique top- k . For technical reasons, and for simplicity, we let $t(X)$ denote the unique lexicographically first top- k vector and refer to it as “the” top- k vector.

where $\tau = \Theta(\sqrt{k})$ is an appropriately chosen threshold. The key to the analysis is to show that, when $X = (x_1, \dots, x_n) \in \{\pm 1\}^{n \times d}$ and $y \in \{\pm 1\}^d$ are chosen uniformly at random, $t(X)$ is an accurate top- k vector of X , then

$$\mathbb{E}[\langle y, t(X) \rangle] = 0 \quad \text{and} \quad \forall i \in [n] \quad \mathbb{E}[\langle x_i, t(X) \rangle] > 2\tau.$$

If we can establish these two facts then Theorem 2 will follow from concentration inequalities for the two inner products.

Suppose $t(X)$ is the *exact* top- k vector. Since each coordinate of y is uniform in $\{\pm 1\}$ and independent of X , we can write

$$\mathbb{E}[\langle y, t(X) \rangle] = \sum_j \mathbb{E}[y_j \cdot t(X)_j] = \sum_j \mathbb{E}[y_j] \mathbb{E}[t(X)_j] = 0.$$

Moreover, for every fixed vector $t \in \{\pm 1\}^d$ with k non-zero coordinates, $\langle y, t \rangle$ is a sum of k independent, bounded random variables. Therefore, by Hoeffding's inequality we have that $\langle y, t \rangle = O(\sqrt{k})$ with high probability. Since y, X are independent, this bound also holds with high probability when X is chosen randomly and $t(X)$ is its top- k vector. Thus, for an appropriate $\tau = \Theta(\sqrt{k})$, $\mathcal{A}(y, t(X)) = \text{OUT}$ with high probability.

Now, suppose $y = x_i$ for some row of X , and we want to show that $\mathbb{E}[\langle x_i, t(X) \rangle]$ is sufficiently large. Since X is chosen uniformly at random, one can show that, when $k \ll d \ll 2^n$, the top- k largest marginals of X are all at least $\gamma = \Omega(\sqrt{\log(d/k)/n})$ with high probability. Thus, on average, when $t(X)_j = 1$, we can think of $x_{i,j} \in \{\pm 1\}$ as a random variable with expectation $\geq \gamma$. Therefore,

$$\mathbb{E}[\langle x_i, t(X) \rangle] = \mathbb{E}\left[\sum_{j:t(X)_j=1} x_{i,j}\right] \geq k\gamma = \Omega\left(k\sqrt{\frac{\log(d/k)}{n}}\right)$$

Even though x_i and $t(X)$ are not independent, and do not have independent entries, we show that with high probability over the choice of X ,

$$\langle x_i, t(X) \rangle \geq \Omega\left(k\sqrt{\frac{\log(d/k)}{n}}\right) - O(\sqrt{k})$$

with high probability. Thus, if $k \log(d/k) \gtrsim n$, we have that $\mathcal{A}(x_i, t(X)) = \text{IN}$ with high probability.

Extension to Noisy Top- k . The case of α -approximate top- k statistics does not change the analysis of $\langle y, t \rangle$ in that case that y is independent of x , but does change the analysis of $\langle x_i, t \rangle$ when x_i is a row of X . It is not too difficult to show that for a *random* row x_i , $\mathbb{E}[\langle x_i, \hat{t} \rangle] \gtrsim k(\gamma - \alpha)$, but it is not necessarily true that $\langle x_i, \hat{t} \rangle$ is large for every row x_i . Note that this is not merely a problem with the analysis. Consider an algorithm that simply ignores some row of the dataset x_i and then outputs the exact top- k vector of the remaining rows x_{-i} . Doing so introduces error $\leq 1/n$, but makes x_i independent of the released top- k vector $t(x_{-i})$.

We can, however, show that if \hat{t} is a sufficiently accurate top- k vector, then $\langle x_i, \hat{t} \rangle > \tau$ for at least $(1 - c)n$ rows of X for an arbitrarily small constant $c > 0$. This weaker tracing guarantee is still enough to rule out (ϵ, δ) -differential privacy for any reasonable setting of ϵ, δ (see Lemma 10 for a precise statement), which gives us Theorem 1. The exact statement and parameters are slightly involved, so we refer the reader to Section 3 for a precise statement and analysis of our tracing attack in the case of approximate top- k statistics (Theorem 11).

Paper Outline

In Section 2 we formalize the model in which our results are proven. In Section 3, we prove Theorem 1 via a tracing attack against an approximate top- k vector (Theorem 11). For the case of exact top- k vectors we can get a simpler and slightly stronger statement (Theorem 2), and for space we defer this to Appendix A. The analysis in Section 3 is self-contained.

2. Preliminaries

Definition 3 (Differential Privacy Dwork et al. (2006)) For $\varepsilon \geq 0, \rho \in [0, 1]$ we say that a randomized algorithm $\mathcal{M} : \{\pm 1\}^{n \times d} \rightarrow \mathcal{R}$ is (ε, δ) -differentially private if for every two datasets $X, X' \in \{\pm 1\}^{n \times d}$, such that X, X' differ in at most one row, we have that,

$$\forall S \subseteq \mathcal{R}, \mathbb{P}[\mathcal{M}(X) \in S] \leq e^\varepsilon \cdot \mathbb{P}[\mathcal{M}(X') \in S] + \delta.$$

Definition 4 (Marginals) For a dataset $X \in \{\pm 1\}^{n \times d}$, its marginal vector

$$q(X) = (q_1(X), \dots, q_d(X))$$

is the average of the rows of X . That is, $q_j(X) = \frac{1}{n} \sum_{i=1}^n X_{i,j}$. We use the notation,

$$q_{(1)}(X) \geq q_{(2)}(X) \geq \dots \geq q_{(d)}(X)$$

to refer to the sorted marginals. We will also define $\pi : [d] \rightarrow [d]$ to be the lexicographically first permutation that puts the marginals in sorted order. That is, we define π so that $q_{\pi(j)} = q_{(j)}$ and if $j < j'$ are such that $q_j = q_{j'}$, then $\pi(j) < \pi(j')$.

Definition 5 (Accurate Top- k Vector) Given a dataset $X \in \{\pm 1\}^{n \times d}$ and a parameter $\alpha \geq 0$, a vector $\hat{t} \in \{0, 1\}^d$ is an α -accurate top- k vector of X if

1. \hat{t} has exactly k non-zero coordinates, and
2. $(\hat{t}_i = 1) \Rightarrow (q_i(X) \geq q_{(k)}(X) - \alpha)$.

When $\alpha = 0$, we define the exact top- k vector of X as $t(X) \in \{0, 1\}^d$ to be the lexicographically first 0-accurate top- k vector.³ Specifically, we define $t(X)$ so that $(t(X)_j = 1) \iff j \in \{\pi(1), \dots, \pi(k)\}$. We refer to these set of columns as the top- k columns of X .

For comparison with our results, we state a positive result for privately releasing an α -approximate top- k vector, which is an easy consequence of the exponential mechanism McSherry and Talwar (2007) and the ‘‘advanced composition theorem’’ for differential privacy Dwork et al. (2010b)

Theorem 6 For every $n, d, k \in \mathbb{N}$, and $\varepsilon, \delta, \beta \in (0, 1)$, there is an (ε, δ) -differentially private algorithm that takes as input a dataset $X \in \{\pm 1\}^{n \times d}$, and with probability at least $1 - \beta$, outputs an α -accurate top- k vector of X , for

$$\alpha = O\left(\frac{\sqrt{k \ln(1/\delta)} \ln(kd/\beta)}{\varepsilon n}\right)$$

3. Due to ties, there may not be a unique 0-accurate top- k vector of X . Thus, to avoid technical nuisances, we let $t(X)$ be the unique lexicographically first 0-accurate top- k vector. Therefore we are justified in treating $t(X)$ as a function of X and referring to the top- k vector with the definite article.

2.1. Probabilistic Inequalities

We need several bounds on the tails of sums of independent random variables.

Lemma 7 (Hoeffding Bound) *Let Z_1, \dots, Z_n be independent random variables supported on $\{\pm 1\}$, and let $Z = \frac{1}{n} \sum_{i=1}^n Z_i$. Then*

$$\forall \nu > 0 \quad \mathbb{P}[Z - \mathbb{E}[Z] \geq \nu] \leq e^{-\frac{1}{2}\nu^2 n}.$$

Hoeffding's bound on the upper tail also applies to random variables that are *negative-dependent*, which in this case means that setting any set of the variables B to $+1$ only makes the variables in $[n] \setminus B$ more likely to be -1 [Panconesi and Srinivasan \(1997\)](#). Similarly, if the random variables are *positive-dependent* (their negations are negative-dependent), then Hoeffding's bound applies to the lower tail.

Theorem 8 (Chernoff Bound) *Let Z_1, \dots, Z_n be a sequence of independent $\{0, 1\}$ -valued random variables, let $Z = \sum_{i=1}^n Z_i$, and let $\mu = \mathbb{E}[Z]$. Then*

1. (Upper Tail) $\forall \nu > 0 \quad \mathbb{P}[Z \geq (1 + \nu)\mu] \leq e^{-\frac{\nu^2}{2+\nu}\mu}$, and
2. (Lower Tail) $\forall \nu \in (0, 1) \quad \mathbb{P}[Z \leq (1 - \nu)\mu] \leq e^{-\frac{1}{2}\nu^2\mu}$.

We use the following anti-concentration inequality for sums of independent Rademacher random variables (see e.g. [Ledoux and Talagrand \(2013\)](#)).

Theorem 9 (Anticoncentration) *Let Z_1, \dots, Z_n be independent and uniform in $\{\pm 1\}$, and let $Z = \frac{1}{n} \sum_{i=1}^n Z_i$. Then for every $\beta > 0$, there exists $K_\beta > 1$ such that for every $n \in \mathbb{N}$,*

$$\forall \nu \in \left[\frac{K_\beta}{\sqrt{n}}, \frac{1}{K_\beta} \right] \quad \mathbb{P}[Z \geq \nu] \geq e^{-\frac{1+\beta}{2}\nu^2 n}.$$

2.2. Tracing Attacks

Intuitively, tracing attacks violate differential privacy because if the target individual y is outside the dataset, then $\mathcal{A}(y, M(X))$ reports OUT with high probability, whereas if y were added to the dataset to obtain X' , $\mathcal{A}(y, M(X'))$ reports IN with high probability. Therefore $M(X), M(X')$ must have very different distributions, which implies that M is not differentially private. The next lemma formalizes and quantifies this property.

Lemma 10 (Tracing Violates DP) *Let $M : \{\pm 1\}^{n \times d} \rightarrow \mathcal{R}$ be a (possibly randomized) algorithm. Suppose there exists an algorithm $\mathcal{A} : \{\pm 1\}^d \times \mathcal{R} \rightarrow \{\text{IN}, \text{OUT}\}$ such that when $X \sim \{\pm 1\}^{n \times d}$ and $y \sim \{\pm 1\}^d$ are independent and uniformly random,*

1. (Soundness) $\mathbb{P}[\mathcal{A}(y, M(X)) = \text{IN}] \leq \rho$
2. (Completeness) $\mathbb{P}[\#\{i \mid \mathcal{A}(x_i, M(X)) = \text{IN}\} \geq n - m] \geq 1 - \rho$.

Then M is not (ε, δ) -differentially private for any ε, δ such that $e^\varepsilon \rho + \delta < 1 - \rho - \frac{m}{n}$. If $\rho < 1/4$ and $m > 3n/4$, then M is not $(\varepsilon_0, \delta_0)$ -differentially private for some absolute constants $\varepsilon_0, \delta_0 > 0$.

By setting ρ and m/n small enough, we can make the constant ε_0 arbitrarily large and the constant δ_0 arbitrarily close to 1. Typically differentially private algorithms typically satisfy (ε, δ) -differential privacy where $\varepsilon = o(1), \delta = o(1/n)$, so ruling out differential privacy with constant (ε, δ) is a strong lower bound.

3. Tracing Using an Approximate Top-k Vector

We use the following *inner product attack*, which takes a top- k vector t of a dataset X , and a target individual y , and attempts to determine if y is an element of X or is independent from X .

$$\mathcal{A}_{\rho,d,k}(y, t) :$$

Input: $y \in \{\pm 1\}^d$ and $t \in \{0, 1\}^d$. Let $\tau = \sqrt{2k \ln(1/\rho)}$.

If $\langle y, t \rangle > \tau$, output IN; else output OUT.

The goal of this section is to prove the following theorem about the inner product attack.

Theorem 11 *For every $\rho > 0$, there exist universal constants $C, C' \in (0, 1)$ (depending only on ρ) such that if $d \in \mathbb{N}$ is sufficiently large and $n, d, k \in \mathbb{N}$ and $\alpha \in (0, 1)$ satisfy,*

$$d \leq 2^{Cn}, \quad k \leq d^C, \quad n = C'k \ln(d/2k), \quad \text{and} \quad \alpha \leq C \sqrt{\frac{\ln(d/2k)}{n}},$$

and $\hat{t} : \{\pm 1\}^{n \times d} \rightarrow \{0, 1\}^d$ is any randomized algorithm such that,

$$\forall X \in \{\pm 1\}^{n \times d} \quad \mathbb{P}[\hat{t}(X) \text{ is an } \alpha\text{-approximate top-}k \text{ vector for } X] \geq 1 - \rho,$$

then $\mathcal{A}_{\rho,d,k}$ (Section A) has the following properties: If $X \sim \{\pm 1\}^{n \times d}$, $y \in \{\pm 1\}^d$ are independent and uniform, then,

1. (Soundness) $\mathbb{P}[\mathcal{A}_{\rho,d,k}(y, \hat{t}(X)) = \text{IN}] \leq \rho$, and
2. (Completeness) $\mathbb{P}[\#\{i \in [n] \mid \mathcal{A}_{\rho,d,k}(x_i, \hat{t}(X)) = \text{IN}\} < (1 - e^2 \rho)n] < 2\rho + 2e^{-k/6}$.

We prove the soundness and completeness properties separately. The proof of soundness is straightforward and follows immediately from the next lemma.

Lemma 12 (Soundness) *For every top- k vector $t \in \{0, 1\}^d$ with at most k non-zeros, if $y \in \{\pm 1\}^d$ is independent of t , then, $\mathbb{P}[\langle y, t \rangle \geq \sqrt{2k \ln(1/\rho)}] \leq \rho$.*

Proof [Lemma 12] Observe that $\langle y, t \rangle = \sum_{j:t_j=1} y_j$ is a sum of k independent, uniform random variables in $\{\pm 1\}$. By Hoeffding's inequality, $\mathbb{P}[\sum_{j:t_j=1} y_j \geq \sqrt{2k \ln(1/\rho)}] \leq \rho$. This completes the proof of the lemma. ■

Intuitively, the proof of completeness takes the same general form as it did for the case of exact top- k statistics, which we described in the introduction. First, for some parameter $\gamma > 0$, with high probability X has at least k marginals that are at least γ . Therefore, any marginal j contained in any α -accurate top- k vector has value $q_j \geq \lambda := \gamma - \alpha$. From this fact, we can conclude that for all α -accurate $\hat{t}(X)$, the expectation of $\langle x_i, \hat{t}(X) \rangle \geq k(\gamma - \alpha)$ where the expectation is taken over choices of X and i . However, the k columns selected by \hat{t} may be significantly correlated so that for

some choices of i , $\langle x_i, \hat{t}(X) \rangle$ is small with high probability. At a high level we solve this problem as follows: first, we restrict to the set of d_λ columns j such that $q_j \geq \gamma - \alpha$, which remain mutually independent. Then we argue that for every fixed α -accurate top- k vector specifying a subset of k of these columns, with overwhelming probability the inner product is large for most choices of $i \in [n]$. Finally, we take a union bound over all $\binom{d_\lambda}{k}$ possible choices of α -accurate top- k vector. To make the union bound tolerable, we need that with high probability d_λ is not too big. Our choice of γ was such that only about k columns are above γ , therefore if we take λ very close to γ , we will also be able to say that d_λ is not too much bigger than k . By assuming that the top- k vector is α -accurate for $\alpha \ll \gamma$, we get that $\lambda = \gamma - \alpha$ is very close to γ .

Before stating the exact parameters and conditions in Lemma 15, we will need to state and prove a few claims about random matrices.

Claim 13 *For every $\beta > 0$, there is a universal constant $C \in (0, 1)$ (depending only on β), such that for every $n \in \mathbb{N}$, $d \leq 2^{Cn}$ and $k \leq Cd$, if $X \in \{\pm 1\}^{n \times d}$ is drawn uniformly at random, then for $\gamma := \sqrt{\frac{2}{1+\beta} \cdot \frac{\ln(d/2k)}{n}}$, we have,*

$$\mathbb{P} [q_{(k)}(X) < \gamma] \leq e^{-k/4}.$$

The proof is a straightforward application of the anticoncentration properties of Rademacher random variables (Lemma 9). See Claim 20 for the proof.

Claim 14 *For every $n, d \in \mathbb{N}$ and every $\lambda \in (0, 1)$, if $X \in \{\pm 1\}^{n \times d}$ is drawn uniformly at random, then for $d_\lambda := 2d \exp(-\frac{1}{2}\lambda^2 n)$,*

$$\mathbb{P} [|\{j \mid q_j > \lambda\}| > d_\lambda] \leq e^{-d_\lambda/6}.$$

Proof [Claim 14] For every $j \in [d]$, define E_j to be the event that $q_j = \frac{1}{n} \sum_i x_{ij} > \lambda$. Since the x_{ij} 's are independent, applying Hoeffding's bound to $\sum_i x_{ij}$ gives,

$$\forall j \in [d] \quad \mathbb{P} [E_j] = \mathbb{P} [q_j > \lambda] \leq e^{-\lambda^2 n/2}.$$

By linearity of expectation, we have that $\mathbb{E}[\sum_j E_j] \leq de^{-\lambda^2 n/2} = \frac{1}{2}d_\lambda$. Since the columns of X are independent, we can apply a Chernoff bound (Theorem 8) to $\sum_j E_j$, which gives,

$$\mathbb{P} \left[\sum_{j=1}^d E_j > d_\lambda \right] \leq e^{-d_\lambda/6}.$$

This completes the proof of the claim. ■

Now we are ready to state our exact claim about the completeness of the attack when given an α -accurate top- k vector.

Lemma 15 (Completeness) *For every $\rho > 0$, there exist universal constants $C_2, C_3, C_4, C_5 \in (0, 1)$ (depending only on ρ) such that if $n, d, k \in \mathbb{N}$ and $\alpha \in (0, 1)$ satisfy,*

$$4k \leq \min\{(2d)^{C_2}, 4C_4 d\}, \quad 8n \ln(1/\rho) = C_3^2 k \ln(2d), \quad d \leq 2^{C_4 n}, \quad \alpha \leq C_5 \sqrt{\frac{\ln(2d)}{n}},$$

and \hat{t} is an algorithm that, for every $X \in \{\pm 1\}^{n \times d}$, outputs an α -accurate top- k vector with probability at least $1 - \rho$, then for a uniformly random $X \in \{\pm 1\}^{n \times d}$, we have

$$\mathbb{P} [\#\{i \in [n] \mid \langle x_i, \hat{t}(X) \rangle \geq \tau_c\} < (1 - e^2 \rho)n] < 2\rho + e^{-k/4} + e^{-k/6},$$

where $\tau_c := C_3 k \sqrt{\frac{\ln(2d)}{n}} - \sqrt{2k \ln(1/\rho)}$.

To see how the completeness property of Theorem 11 follows from the lemma, observe that if $8n \ln(1/\rho) = C_3^2 k \ln(2d)$, then,

$$\tau_c = C_3 k \sqrt{\frac{\ln(2d)}{n}} - \sqrt{2k \ln(1/\rho)} = \sqrt{2k \ln(1/\rho)} = \tau,$$

where τ is the threshold in $\mathcal{A}_{\rho, d, k}$. Therefore Lemma 15 implies that,

$$\mathbb{P} [\#\{i \in [n] \mid \mathcal{A}_{\rho, d, k}(x_i, \hat{t}(X)) = \text{IN}\} < (1 - e^2 \rho)n] < \rho + e^{-k/4} + e^{-k/6}.$$

The universal constants C, C' will be $C = \min\{C_2, C_4, C_5\} - \delta$ for an arbitrarily small $\delta > 0$, and $C' = C_3^2$. As long as d is sufficiently large the conditions $k \leq d^C$ in Theorem 11 will imply the corresponding condition in the above lemma.

Proof [Lemma 15] First, we will condition everything on the event,

$$G_\alpha := \{\hat{t} = \hat{t}(X) \text{ is an } \alpha\text{-accurate top-}k \text{ vector of } X\}.$$

By assumption, for every $X \in \{\pm 1\}^{n \times d}$, $\mathbb{P}[G_\alpha] \geq 1 - \rho$.

For convenience define the constant $c := e^2 \rho$, so that the lemma asserts that, with high probability, $\mathcal{A}(x_i, \hat{t}(X)) = \text{IN}$ for at least $(1 - c)n$ rows x_i . We will first condition on the event that at least k marginals are above the threshold γ . Now, by Claim 13, with an appropriate choice of,

$$\beta := \frac{c}{16 \ln(1/\rho)} \quad \gamma := \sqrt{\frac{2}{1 + \frac{c}{16 \ln(1/\rho)}}} \cdot \sqrt{\frac{\ln(d/2k)}{n}},$$

and the assumptions that $k \leq C_4 d$ and $d \leq 2^{C_4 n}$ for some universal constant C_4 depending only on β , the event,

$$G_\gamma := \left\{ q_{(k)}(X) \geq \gamma = C_1 \sqrt{\frac{\ln(2d)}{n}} \right\},$$

will hold with probability $1 - e^{-k/4}$. Here we define the universal constants,

$$C_1 := \sqrt{\frac{2}{1 + \frac{c}{8 \ln(1/\rho)}}} \quad C_2 := \frac{c}{2c + 4 \ln(1/\rho)},$$

depending only on ρ . These constants were chosen so that, provided $4k \leq (2d)^{C_2}$, the inequality in the definition of G_γ will be satisfied.

In light of the above analysis, we condition the rest of the analysis on the event G_γ , which satisfies $\mathbb{P}[G_\gamma] \geq 1 - e^{-k/4}$.

If we condition on G_α and G_γ , then for any marginal j chosen by \hat{t} (i.e. $\hat{t}_j = 1$), we can say that $q_j \geq \lambda$ for any $\lambda \leq \gamma - \alpha$. Now, we define the constants,

$$C_3 := \sqrt{\frac{2}{1 + \frac{c}{4\ln(1/\rho)}}} \quad C_5 := C_1 - C_3 > 0,$$

where one can verify that the inequality $C_1 - C_3 > 0$ holds for all choices of c . Now by our assumption that $\alpha < C_5 \sqrt{\frac{\ln(2d)}{n}}$, we can define $\lambda := C_3 \sqrt{\frac{\ln(2d)}{n}}$.

For any matrix $X \in \{\pm 1\}^{n \times d}$, we define $S_\lambda = S_\lambda(X) \subseteq \{1, \dots, d\}$ to be the set of columns of X whose marginals are greater than λ . The analysis above says that, conditioned on G_γ and G_α , if $\hat{t}_j = 1$, then $j \in S_\lambda$. Note that, if X is chosen uniformly at random, and we define $X_{\geq \lambda} \in \{\pm 1\}^{n \times |S_\lambda|}$ to be the restriction of X to the columns contained in S_λ , then the columns of $X_{\geq \lambda}$ remain independent.

The size of S_λ is a random variable supported on $\{0, 1, \dots, d\}$. In our analysis we will need to condition on the event that $|S_\lambda| \ll d$. Using Claim 14 we have that if $d_\lambda := 2de^{-\frac{\lambda^2 n}{2}}$, then the event, $G_S := \{|S_\lambda(X)| \leq d_\lambda\}$, satisfies, $\mathbb{P}[G_S] \geq 1 - e^{-d_\lambda/6} \geq 1 - e^{-k/6}$, where we have used the fact that $d_\lambda \geq k$. This fact is not difficult to verify from our choice of parameters. Intuitively, since $\lambda \leq \gamma$, and there are at least k marginals larger than γ , there must also typically be at least k marginals larger than λ . We condition the remainder of the analysis on the event G_S .

For some vector t , we will say that the attack \mathcal{A} fails on t when we fail to trace more than cn rows, i.e. \mathcal{A} fails when, $|\{i : \langle x_i, t \rangle < k\lambda - \sqrt{2k \ln(1/\rho)}\}| > cn = e^2 \rho n$. Formally we have that,

$$\begin{aligned} \mathbb{P}[\mathcal{A} \text{ fails on } \hat{t}] &\leq \mathbb{P}[\mathcal{A} \text{ fails on } \hat{t} \wedge G_\alpha \wedge G_\gamma \wedge G_S] + \mathbb{P}[\neg G_\alpha \vee \neg G_\gamma \vee \neg G_S] \\ &\leq \mathbb{P}[\mathcal{A} \text{ fails on } \hat{t} \wedge G_\alpha \wedge G_\gamma \wedge G_S] + \rho + e^{-k/4} + e^{-k/6}. \end{aligned} \quad (1)$$

Thus, to complete the proof, it suffices to show that,

$$\begin{aligned} &\mathbb{P}[\mathcal{A} \text{ fails on } \hat{t} \wedge G_\alpha \wedge G_\gamma \wedge G_S] \\ &= \mathbb{P}[(\mathcal{A} \text{ fails on } \hat{t}) \wedge (\hat{t} \text{ is } \alpha\text{-accurate}) \wedge (q_{(k)} \geq \gamma) \wedge (|S_\lambda| \leq d_\lambda)] \\ &\leq \mathbb{P}[(\mathcal{A} \text{ fails on } \hat{t}) \wedge (\hat{t} \subseteq S_\lambda) \wedge (|S_\lambda| \leq d_\lambda)] \\ &\leq \mathbb{P}\left[\left(\exists v \in \binom{S_\lambda}{k} \mathcal{A} \text{ fails on } v\right) \wedge (|S_\lambda| \leq d_\lambda)\right], \end{aligned}$$

where we have abused notation and written $\hat{t} \subseteq S_\lambda$ to mean that $\hat{t}_j = 1 \implies j \in S_\lambda$, and used $v \in \binom{S_\lambda}{k}$ to mean that v is a subset of S_λ of size exactly k .

We will now upper bound $\mathbb{P}\left[\left(\exists v \in \binom{S_\lambda}{k} \mathcal{A} \text{ fails on } v\right) \wedge (|S_\lambda| \leq d_\lambda)\right]$. Observe that, since the columns of X are identically distributed, this probability is independent of the specific choice of S_λ and depends only on $|S_\lambda|$. Further, decreasing the size of S_λ only decreases the probability. Thus, we will fix a set S of size exactly d_λ and assume $S_\lambda = S$. Thus, for our canonical choice of set $S = \{1, \dots, d_\lambda\}$, we need to bound $\mathbb{P}\left[\exists v \in \binom{S}{k} \mathcal{A} \text{ fails on } v\right]$.

Consider a fixed vector $v \subseteq S$. That is, a vector $v \in \{0, 1\}^d$ such that $v_j = 1 \implies j \in S$. Define the event $E_{i,v}$ to be the event that $\langle x_i, v \rangle$ is too small for some specific row i and some specific vector $v \subseteq S$. That is,

$$E_{i,v} := \left\{ \langle x_i, v \rangle < \tau_c := k\lambda - \sqrt{2k \ln(1/\rho)} \right\}.$$

Since the columns of X_S are independent, for a fixed i and v , by Hoeffding's inequality gives,

$$\mathbb{P}[E_{i,v}] = \mathbb{P}\left[\sum_{j:v_j=1} x_{i,j} < \tau_c\right] \leq \rho.$$

We have proved that the probability that $\langle x_i, v \rangle$ is small, is small for a given row. We want to bound the probability that $\langle x_i, v \rangle$ is small for an entire set of rows $R \subseteq [n]$. Unfortunately, since we require that $q_j \geq \lambda$ for every column $j \in S$, the rows x_i are no longer independent. However, the rows satisfy a negative-dependence condition, captured in the following claim.

Claim 16 For every $R \subseteq [n]$, $\mathbb{P}[\bigwedge_{i \in R} E_{i,v}] \leq \rho^{|R|}$.

To maintain the flow of the analysis, we defer the proof of this claim to Section 3.1

By definition, \mathcal{A} fails on v only if there exists a set R of exactly $cn = e^2 \rho n$ rows such that $\bigwedge_{i \in R} E_{i,v}$. Taking a union bound over all such sets R and all v , we have,

$$\begin{aligned} \mathbb{P}\left[\exists v \in \binom{S}{k} \mathcal{A} \text{ fails on } v\right] &\leq \binom{d_\lambda}{k} \cdot \binom{n}{cn} \cdot \rho^{cn} \\ &\leq \left(\frac{ed_\lambda}{k}\right)^k \cdot \left(\frac{en\rho}{cn}\right)^{cn} \\ &\leq d_\lambda^k \cdot e^{-cn}, \end{aligned}$$

where we have used the identity $\binom{a}{b} \leq \left(\frac{ea}{b}\right)^b$. We have already set the parameter λ , and set $d_\lambda = 2de^{-\frac{\lambda^2 n}{2}}$. Thus, all that remains is to show that for our choice of parameters, $d_\lambda^k \cdot e^{-cn} \leq \rho$, which is equivalent to $cn \geq \ln(1/\rho) + k \ln(d_\lambda)$. Substituting our choice of λ gives the condition,

$$\frac{k\lambda^2 n}{2} \geq \ln(1/\rho) + k \ln(2d) - cn.$$

One can check that, for our choice of $n = \frac{C_3^2 k \ln(2d)}{8 \ln(1/\rho)}$, and our choice of $\lambda = C_3 \sqrt{\frac{\ln(2d)}{n}}$ where C_3 has been defined above, the preceding equation is satisfied.

Thus, we have established that,

$$\mathbb{P}\left[\exists v \in \binom{S}{k} \mathcal{A} \text{ fails on } v\right] \leq d_\lambda^k \cdot e^{-cn} \leq \rho.$$

As we have argued above, this implies that,

$$\mathbb{P}[\mathcal{A} \text{ fails on } \hat{t}] \leq 2\rho + e^{-k/4} + e^{-k/6}.$$

This completes the proof of the completeness lemma. ■

3.1. Proof of Claim 16

Recall that, for a given $X \in \{\pm 1\}^{n \times d}$, $E_{i,v}$ is the event that $\langle x_i, v \rangle < \tau_c$ for a specific row i and a specific vector $v \subseteq S$, where $S = S_\lambda$ is the set of columns j of X such that $q_j \geq \lambda$. Thus, we can think of $X_S \in \{\pm 1\}^{n \times |S|}$ as a matrix with $|S|$ independent columns that are uniformly random subject to the constraint that each column's mean is at least λ . Since, flipping some entries of X_S from -1 to $+1$ can only increase $\langle x_i, v \rangle$, we will in fact use the distribution \tilde{X}_S in which each column's mean is exactly λn . Thus, when we refer to the probabilities of events involving random variables x_{ij} , we will use this distribution on X_S as the probability space. Additionally, since v is fixed, and the probability is the same for all v , we will simply write E_i to cut down on notational clutter.

For a specific set $R \subseteq [n]$, we need to calculate $\mathbb{P}_{\tilde{X}_S} [\bigwedge_{i \in R} E_i]$. We can write,

$$\begin{aligned} \mathbb{P}_{\tilde{X}_S} \left[\bigwedge_{i \in R} E_i \right] &= \mathbb{P}_{\tilde{X}_S} \left[\bigwedge_{i \in R} \left(\sum_{j \in v} x_{ij} < \tau_c \right) \right] \\ &\leq \mathbb{P}_{\tilde{X}_S} \left[\sum_{i \in R, j \in v} x_{ij} < |R| \tau_c \right]. \end{aligned} \quad (2)$$

The key property of X_S is that its entries x_{ij} are positively correlated. That is, for every set $I \subset [n] \times S$ of variables x_{ij} , we have,

$$\mathbb{P}_{\tilde{X}_S} [\forall (i, j) \in I \quad x_{ij} = -1] \leq \prod_{(i, j) \in I} \mathbb{P}_{\tilde{X}_S} [x_{ij} = -1]. \quad (3)$$

Since the columns of \tilde{X}_S are independent if we partition the elements of I into sets I_1, \dots, I_k , where each set I_l has pairs of I which come from the column l , then,

$$\mathbb{P}_{\tilde{X}_S} [\forall (i, j) \in I \quad x_{ij} = -1] = \prod_{l \in [k]} \mathbb{P}_{\tilde{X}_S} [\forall (i, j) \in I_l \quad x_{ij} = -1].$$

So it is enough to show that equation 3 holds when $I = \{(i_1, l), \dots, (i_p, l)\}$. For simplicity of notation we will refer to these elements of I as $\{1, \dots, p\}$. We have that,

$$\mathbb{P} [\forall a \in I \quad x_a = -1] = \prod_{a=1}^p \mathbb{P} [x_a = -1 \mid (\forall b \in \{1, \dots, a-1\}, x_b = -1)]. \quad (4)$$

We will show that each of the terms in the product is smaller than $\mathbb{P} [x_a = -1]$. For a fixed $a \in I$, let B be the set $\{1, \dots, a-1\}$ and let E be the event that $(\forall b \in B, x_b = -1)$. Since every column of X_S sums to $n\lambda$, we have,

$$\mathbb{E} \left[\sum_{i \in [n]} x_{il} \mid B \right] = n\lambda.$$

On the other hand, since the bits in B are all set to -1 and all the other bits in column l are equal in expectation,

$$\mathbb{E} \left[\sum_{i \in [n]} x_{il} \mid B \right] = -|B| + (n - |B|) \cdot (\mathbb{E} [x_a \mid B]),$$

which means that,

$$\mathbb{E}[x_a \mid B] \geq \lambda = \mathbb{E}[x_a].$$

Since $\mathbb{P}[x_a = -1] = (1 - \mathbb{E}[x_a])/2$, we get that, $\mathbb{P}[x_a = -1 \mid B] \leq \mathbb{P}[x_a = -1]$. Substituting this back into (4), we get that the variables are positively correlated.

We have that, $\mathbb{E}\left[\sum_{i \in R, j \in v} x_{ij}\right] = |R|k\lambda$, and since Hoeffding's inequality applies equally well to positively-correlated random variables [Panconesi and Srinivasan \(1997\)](#), we also have,

$$\begin{aligned} & \mathbb{P}_{\tilde{X}_S} \left[\sum_{i \in R, j \in v} x_{ij} \leq |R|\tau_c \right] \\ & \leq \mathbb{P}_{\tilde{X}_S} \left[\sum_{i \in R, j \in V} x_{ij} < |R|k\lambda - |R|\sqrt{2k \ln(1/\rho)} \right] \\ & \leq \exp \left(-\frac{\left(|R|\sqrt{2k \ln(1/\rho)}\right)^2}{2|R|k} \right) = \rho^{|R|}. \end{aligned}$$

Substituting this in equation 2, we get that, $\mathbb{P}_{\tilde{X}_S} \left[\bigwedge_{i \in R} E_i \right] \leq \rho^{|R|}$. Finally, we use the fact that, by our definition of the distributions, X_S, \tilde{X}_S , we have, $\mathbb{P}_{X_S} \left[\bigwedge_{i \in R} E_i \right] \leq \mathbb{P}_{\tilde{X}_S} \left[\bigwedge_{i \in R} E_i \right] \leq \rho^{|R|}$. This completes the proof.

Acknowledgments

We are grateful to Adam Smith and Thomas Steinke for many helpful discussions about tracing attacks and private top- k selection.

References

- Raef Bassily, Kobbi Nissim, Adam Smith, Thomas Steinke, Uri Stemmer, and Jonathan Ullman. Algorithmic stability for adaptive data analysis. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1046–1059. ACM, 2016.
- Amos Beimel, Kobbi Nissim, and Uri Stemmer. Private learning and sanitization: Pure vs. approximate differential privacy. In *RANDOM*, 2013.
- Raghav Bhaskar, Srivatsan Laxman, Adam Smith, and Abhradeep Thakurta. Discovering frequent patterns in sensitive data. In *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 503–512. ACM, 2010.
- Avrim Blum, Cynthia Dwork, Frank McSherry, and Kobbi Nissim. Practical privacy: the SuLQ framework. In *PODS*, 2005.
- Dan Boneh and James Shaw. Collusion-secure fingerprinting for digital data. *IEEE Transactions on Information Theory*, 44(5):1897–1905, 1998.

- Mark Bun, Jonathan Ullman, and Salil P. Vadhan. Fingerprinting codes and the price of approximate differential privacy. In *STOC*, 2014.
- Kamalika Chaudhuri, Daniel J Hsu, and Shuang Song. The large margin mechanism for differentially private maximization. In *Advances in Neural Information Processing Systems*, pages 1287–1295, 2014.
- Irit Dinur and Kobbi Nissim. Revealing information while preserving privacy. In *PODS*, 2003.
- Cynthia Dwork and Jing Lei. Differential privacy and robust statistics. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 371–380. ACM, 2009.
- Cynthia Dwork and Kobbi Nissim. Privacy-preserving datamining on vertically partitioned databases. In *CRYPTO*, 2004.
- Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *TCC*, 2006.
- Cynthia Dwork, Moni Naor, Toniann Pitassi, and Guy N Rothblum. Differential privacy under continual observation. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 715–724. ACM, 2010a.
- Cynthia Dwork, Guy N. Rothblum, and Salil P. Vadhan. Boosting and differential privacy. In *FOCS*. IEEE, 2010b.
- Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.
- Cynthia Dwork, Adam D. Smith, Thomas Steinke, Jonathan Ullman, and Salil P. Vadhan. Robust traceability from trace amounts. In *FOCS*, 2015a.
- Cynthia Dwork, Weijie Su, and Li Zhang. Private false discovery rate control. *arXiv preprint arXiv:1511.03803*, 2015b.
- Cynthia Dwork, Adam Smith, Thomas Steinke, and Jonathan Ullman. Exposed! A survey of attacks on private data. *Annual Review of Statistics and Its Application*, 2017.
- Anupam Gupta, Katrina Ligett, Frank McSherry, Aaron Roth, and Kunal Talwar. Differentially private combinatorial optimization. In *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms*, pages 1106–1125. Society for Industrial and Applied Mathematics, 2010.
- Anupam Gupta, Aaron Roth, and Jonathan Ullman. Iterative constructions and private data release. In *TCC*, 2012.
- Moritz Hardt and Guy N. Rothblum. A multiplicative weights mechanism for privacy-preserving data analysis. In *FOCS*, 2010.
- Nils Homer, Szabolcs Szelinger, Margot Redman, David Duggan, Waibhav Tembe, Jill Muehling, John V Pearson, Dietrich A Stephan, Stanley F Nelson, and David W Craig. Resolving individuals contributing trace amounts of dna to highly complex mixtures using high-density snp genotyping microarrays. *PLoS genetics*, 4(8):e1000167, 2008.

- Prateek Jain and Abhradeep Thakurta. Mirror descent based database privacy. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 579–590. Springer, 2012.
- Michel Ledoux and Michel Talagrand. *Probability in Banach Spaces: isoperimetry and processes*. Springer Science & Business Media, 2013.
- Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *Foundations of Computer Science, 2007. FOCS'07. 48th Annual IEEE Symposium on*, pages 94–103. IEEE, 2007.
- Alessandro Panconesi and Aravind Srinivasan. Randomized distributed edge coloring via an extension of the chernoff–hoeffding bounds. *SIAM Journal on Computing*, 26(2):350–368, 1997.
- Aaron Roth and Tim Roughgarden. Interactive privacy via the median mechanism. In *STOC*, pages 765–774. ACM, June 5–8 2010.
- Sriram Sankararaman, Guillaume Obozinski, Michael I Jordan, and Eran Halperin. Genomic privacy and limits of individual detection in a pool. *Nature genetics*, 41(9):965–967, 2009.
- Adam Smith and Abhradeep Thakurta. Differentially private model selection via stability arguments and the robustness of the lasso. *J Mach Learn Res Proc Track*, 30:819–850, 2013.
- Thomas Steinke and Jonathan Ullman. Between pure and approximate differential privacy. *Journal of Privacy and Confidentiality*, 7, 2017a.
- Thomas Steinke and Jonathan Ullman. Tight lower bounds for differentially private selection. *arXiv preprint arXiv:1704.03024*, 2017b.
- Gábor Tardos. Optimal probabilistic fingerprint codes. *Journal of the ACM (JACM)*, 55(2):10, 2008.
- Jonathan Ullman. Answering $n^{2+o(1)}$ counting queries with differential privacy is hard. In *STOC*, 2013.
- Jonathan Ullman. Private multiplicative weights beyond linear queries. In *PODS*, 2015.

Appendix A. Tracing Using the Top-k Vector

Recall the inner product attack, which takes a top- k vector t of a dataset X , and a target individual y , and attempts to determine if y is an element of X or is independent from X .

$\mathcal{A}_{\rho,d,k}(y,t) :$ <p style="text-align: center;">Input: $y \in \{\pm 1\}^d$ and $t \in \{0, 1\}^d$. Let $\tau = \sqrt{2k \ln(1/\rho)}$.</p> <p style="text-align: center;">If $\langle y, t \rangle > \tau$, output IN; else output OUT.</p>

In this section we will analyze this attack when $X \in \{\pm 1\}^{n \times d}$ is a uniformly random matrix, and $t = t(X)$ is the exact top- k vector of X . In this case, we have the following theorem.

Theorem 17 *There is a universal constant $C \in (0, 1)$ such that if $\rho > 0$ is any parameter and $n, d, k \in \mathbb{N}$ satisfy $d \leq 2^{Cn}$, $k \leq Cd$ and $k \ln(d/2k) \geq 8n \ln(1/\rho)$, then, $\mathcal{A}_{\rho, d, k}$ has the following properties: If $X \sim \{\pm 1\}^{n \times d}$, $y \sim \{\pm 1\}^d$ are independent and uniform, and $t(X)$ is the exact top- k vector of X , then,*

1. (Soundness) $\mathbb{P}[\mathcal{A}_{\rho, d, k}(y, t(X)) = \text{IN}] \leq \rho$, and
2. (Completeness) for every $i \in [n]$, $\mathbb{P}[\mathcal{A}_{\rho, d, k}(x_i, t(X)) = \text{OUT}] < \rho + e^{-k/4}$.

We will prove the soundness and completeness properties separately in Lemmas 18 and 19, respectively. The proof of soundness is straightforward.

Lemma 18 (Soundness) *For every $\rho > 0$, $n \in \mathbb{N}$, and $k \leq d \in \mathbb{N}$, if $X \sim \{\pm 1\}^{n \times d}$, $y \sim \{\pm 1\}^d$ are independent and uniformly random, and $t(X)$ is the exact top- k vector, then,*

$$\mathbb{P} \left[\langle y, t(X) \rangle \geq \sqrt{2k \ln(1/\rho)} \right] \leq \rho.$$

Proof Recall that $\tau := \sqrt{2k \ln(1/\rho)}$. Since X, y are independent, we have,

$$\begin{aligned} & \mathbb{P}_{X, y} [\langle y, t(X) \rangle \geq \tau] \\ = & \sum_{T \subseteq [d]: |T|=k} \mathbb{P}_{X, y} [\langle y, t(X) \rangle \geq \tau \mid t(X) = \mathbb{I}_T] \cdot \mathbb{P}_X [t(X) = \mathbb{I}_T] \\ = & \sum_{T \subseteq [d]: |T|=k} \mathbb{P}_y \left[\sum_{j \in T} y_j \geq \tau \right] \cdot \mathbb{P}_X [t(X) = \mathbb{I}_T] \quad (X, y \text{ are independent}) \\ \leq & \max_{T \subseteq [d]: |T|=k} \mathbb{P}_y \left[\sum_{j \in T} y_j \geq \tau \right]. \end{aligned}$$

For every fixed T , the random variables $\{y_j\}_{j \in T}$ are independent and uniform on $\{\pm 1\}$, so by Hoeffding's inequality,

$$\mathbb{P} \left[\sum_{j \in T} y_j \geq \sqrt{2k \ln(1/\rho)} \right] \leq \rho.$$

This completes the proof of the lemma. ■

We now turn to proving the completeness property, which will follow immediately from the following lemma.

Lemma 19 (Completeness) *There is a universal constant $C \in (0, 1)$ such that for every $\rho > 0$, $n \in \mathbb{N}$, $d \leq 2^{Cn}$, and $k \leq Cd$, if $X \sim \{\pm 1\}^{n \times d}$ is chosen uniformly at random, $t(X)$ is the exact top- k vector, and x_i is any row of X ,*

$$\mathbb{P} \left[\langle x_i, t(X) \rangle \leq k \sqrt{\frac{\ln(d/2k)}{n}} - \sqrt{2k \ln(1/\rho)} \right] \leq \rho + e^{-k/4}.$$

To see how the completeness property of Theorem 17 follows from the lemma, observe that, if $k \ln(d/2k) \geq 8n \ln(1/\rho)$, then $k \sqrt{\ln(d/2k)/n} - \sqrt{2k \ln(1/\rho)} \geq \tau$. Therefore Lemma 19 implies that $\mathbb{P}[\langle x_i, t(X) \rangle < \tau] \leq \rho + e^{-k/4}$, so $\mathbb{P}[\mathcal{A}_{\rho, d, k}(x_i, t(X)) = \text{IN}] \geq 1 - \rho - e^{-k/4}$.

Before proving the lemma, we will need a few claims about the distribution of $\langle x_i, t(X) \rangle$. The first claim asserts that, although $X \in \{\pm 1\}^{n \times d}$ is uniform, the k columns of X with the largest marginals are significantly biased.

Claim 20 *There is a universal constant $C \in (0, 1)$, such that for every $n \in \mathbb{N}$, $d \leq 2^{Cn}$ and $k \leq Cd$, if $X \in \{\pm 1\}^{n \times d}$ is drawn uniformly at random, then,*

$$\mathbb{P} \left[q_{(k)}(X) < \sqrt{\frac{\ln(d/2k)}{n}} \right] \leq e^{-k/4}.$$

Proof [Proof of Claim 20] For every $j \in [d]$, define E_j to be the event that,

$$q_j = \frac{1}{n} \sum_{i \in [n]} x_{ij} > \sqrt{\frac{\ln(d/2k)}{n}}.$$

We would like to apply Theorem 9 to the random variable $\frac{1}{n} \sum_i x_{ij}$. To do so, we need,

$$\sqrt{\ln(d/2k)/n} \in \left[\frac{K_1}{\sqrt{n}}, \frac{1}{K_1} \right]$$

where K_1 is the universal constant from that theorem (applied with $\beta = 1$). These inequalities will be satisfied as long as $d \leq 2^{Cn}$, and $k \leq Cd$ for a suitable universal constant $C \in (0, 1)$. Applying Theorem 9 gives,

$$\forall j \in [d] \quad \mathbb{P}[E_j] = \mathbb{P} \left[q_j > \sqrt{\frac{\ln(d/2k)}{n}} \right] \geq \frac{2k}{d}.$$

By linearity of expectation, we have that $\mathbb{E}[\sum_j E_j] \geq 2k$. Since the columns of X are independent, and the events E_j only depend on a single column of X , the events E_j are also independent. Therefore, we can apply a Chernoff bound (Theorem 8) to $\sum_j E_j$ to get,

$$\mathbb{P} \left[\sum_{j=1}^d E_j < k \right] \leq e^{-k/4}.$$

If $\sum_j E_j \geq k$, then there exist k values q_j that are larger than $\sqrt{\frac{\ln(d/2k)}{n}}$, so $q_{(k)}$ is also at least this value. This completes the proof of the claim. \blacksquare

The previous claim establishes that if we restrict X to its top- k columns, the resulting matrix $X_t \in \{\pm 1\}^{n \times k}$ is a random matrix whose mean entry is significantly larger than 0. This claim is enough to establish that the inner product $\langle x_i, t(X) \rangle$ is large in expectation over X . However, since the columns of X_t are not necessarily independent, we cannot directly apply concentration to get the high probability statement we need. However, the columns of X_t are independent if we condition on the value and location of the $(k+1)$ -st marginal.

Claim 21 *Let $X \in \{\pm 1\}^{n \times d}$ be a random matrix from a distribution with independent columns, and let $t(X)$ be its marginals. For every $q \in [-1, 1]$, $k, j \in [d]$, $T \in \binom{[d]}{k}$, the conditional distribution,*

$$X \mid (q_{(k+1)} = q) \wedge (\pi(k+1) = j) \wedge (t(X) = \mathbb{I}_T)$$

also has independent columns.

Proof [Proof of Claim 21] Suppose we condition on the value of the $(k+1)$ -st marginal, $q_{(k+1)} = q$, its location, $\pi(k+1) = j$, and the set of top- k marginals $t = \mathbb{I}_T$. By definition of the (exact, lexicographically first) top- k vector, we have that if $\ell < j$, then $\ell \in T$ if and only if $q_\ell \geq q$. Similarly, if $\ell > j$, then $\ell \in T$ if and only if $q_\ell > q$. Since we have conditioned on a fixed tuple (q, j, T) , the statements $q_\ell > q$ and $q_\ell \geq q$ now depend only on the ℓ -th column. Thus, since the columns of X are independent, they remain independent even when condition on any tuple (q, j, T) . Specifically, if $\ell < j$ and $\ell \in T$, then column ℓ is drawn independently from the conditional distribution $((u_1, \dots, u_n) \mid \frac{1}{n} \sum_i u_i \geq q)$, where $(u_1, \dots, u_n) \in \{\pm 1\}^n$ are chosen independently and uniformly at random. Similarly, if $\ell > j$ and $\ell \in T$, then column ℓ is drawn independently from $((u_1, \dots, u_n) \mid \frac{1}{n} \sum_i u_i > q)$. ■

Proof [Proof of Lemma 19] For convenience, define $\gamma = \sqrt{\frac{\ln(d/2k)}{n}}$ and $\tau_c = k\gamma - \sqrt{2k \ln(1/\rho)}$. Fix any row x_i of X . We can write,

$$\begin{aligned} \mathbb{P}[\langle x_i, t(X) \rangle < \tau_c] &\leq \mathbb{P}[\langle x_i, t(X) \rangle < \tau_c \mid q_{(k+1)} \geq \gamma] + \mathbb{P}[q_{(k+1)} < \gamma] \\ &\leq \mathbb{P}[\langle x_i, t(X) \rangle < \tau_c \mid q_{(k+1)} \geq \gamma] + e^{-k/4} && \text{(Claim 20)} \\ &\leq \max_{q \geq \gamma, j \in [d], T \in \binom{[d]}{k}} \mathbb{P}[\langle x_i, t(X) \rangle < \tau_c \mid (q_{(k+1)} = q) \wedge (\pi(k+1) = j) \wedge (t(X) = \mathbb{I}_T)] + e^{-k/4} \end{aligned} \tag{5}$$

Let $G_{q,j,T}$ be the event $(q_{(k+1)} = q) \wedge (\pi(k+1) = j) \wedge (t(X) = \mathbb{I}_T)$. By linearity of expectation, we can write $\mathbb{E}[\langle x_i, t(X) \rangle < \tau_c \mid G_{q,j,T}] \geq kq \geq k\gamma$. Using Claim 21, we have that $\sum_{\ell: t(X)_\ell = 1} x_{i\ell}$ conditioned on $G_{\delta,j,T}$ is a sum of independent $\{\pm 1\}$ -valued random variables. Thus,

$$\begin{aligned} \mathbb{P}[\langle x_i, t(X) \rangle < \tau_c \mid G_{q,j,T}] &= \mathbb{P}[\langle x_i, t(X) \rangle < k\gamma - \sqrt{2k \ln(1/\rho)} \mid G_{q,j,T}] \\ &\leq \mathbb{P}[\langle x_i, t(X) \rangle < kq - \sqrt{2k \ln(1/\rho)} \mid G_{q,j,T}] && (q \geq \gamma) \\ &\leq \rho && \text{(Hoeffding)} \end{aligned}$$

Combining with (5) completes the proof. ■