

Dealing with Range Anxiety in Mean Estimation via Statistical Queries

Vitaly Feldman
IBM Research - Almaden

Abstract

We give algorithms for estimating the expectation of a given real-valued function $\phi : X \rightarrow \mathbb{R}$ on a sample drawn randomly from some unknown distribution D over domain X , namely $\mathbf{E}_{\mathbf{x} \sim D}[\phi(\mathbf{x})]$. Our algorithms work in two well-studied models of restricted access to data samples. The first one is the statistical query (SQ) model in which an algorithm has access to an *SQ oracle* for the input distribution D over X instead of i.i.d. samples from D . Given a query function $\phi : X \rightarrow [0, 1]$, the oracle returns an estimate of $\mathbf{E}_{\mathbf{x} \sim D}[\phi(\mathbf{x})]$ within some tolerance τ . The second, is a model in which only a single bit is communicated from each sample. In both of these models the error obtained using a naive implementation would scale polynomially with the range of the random variable $\phi(\mathbf{x})$ (which might even be infinite). In contrast, without restrictions on access to data the expected error scales with the standard deviation of $\phi(\mathbf{x})$. Here we give a simple algorithm whose error scales linearly in standard deviation of $\phi(\mathbf{x})$ and logarithmically with an upper bound on the second moment of $\phi(\mathbf{x})$.

As corollaries, we obtain algorithms for high dimensional mean estimation and stochastic convex optimization in these models that work in more general settings than previously known solutions.

1 Overview

We consider the problem of estimating the expectation $D[\phi] \doteq \mathbf{E}_{\mathbf{x} \sim D}[\phi(\mathbf{x})]$, where D is the unknown input distribution over X and $\phi : X \rightarrow \mathbb{R}$. We study this problem in the statistical query (SQ) model in which an algorithm has access to a statistical query oracle for D in place of i.i.d. samples from D as in the traditional setting of statistics and machine learning. The most commonly studied SQ oracle was introduced by Kearns [Kea98] and gives an estimate of the mean of any bounded function with fixed tolerance.

Definition 1.1. *Let D be a distribution over a domain X and $\tau > 0$. A statistical query oracle $STAT_D(\tau)$ is an oracle that given as input any function $\phi : X \rightarrow [0, 1]$, returns some value v such that $|v - D[\phi]| \leq \tau$.*

A special case of statistical queries are counting or linear queries in which the distribution D is uniform over the elements of a given database $S \in X^n$. In other words the goal is to estimate the empirical mean of ϕ on the given set of data points. This setting is studied extensively in the literature on differential privacy (see [DR14] for an overview) and our discussion applies to this setting as well.

Tolerance τ of statistical queries roughly corresponds to the number of random samples in the traditional setting. Namely, the Chernoff-Hoeffding bounds imply that n i.i.d. samples allow estimation of $D[\phi]$ with tolerance $\tau = \Theta(1/\sqrt{n})$ (with high probability). However, if the variance of $\phi(\mathbf{x})$ is low then n random samples will likely give a more accurate estimate. To address this discrepancy a somewhat stronger oracle was introduced in [FGRVX12].

Definition 1.2. Let D be a distribution over a domain X and $n > 0$. A statistical query oracle $VSTAT_D(n)$ is an oracle that given as input any function $\phi : X \rightarrow [0, 1]$ returns a value v such that $|v - p| \leq \max\left\{\frac{1}{n}, \sqrt{\frac{p(1-p)}{n}}\right\}$, where $p \doteq D[\phi]$.

An important property of the accuracy guarantees of the $VSTAT$ oracle is that the guarantees can be easily implemented in the model which allows only a single bit to be communicated about each sample. Formally, a 1-bit sampling oracle for D is defined as follows [BD98].

Definition 1.3 (1-STAT oracle). Let D be a distribution over the domain X . The 1- $STAT_D$ oracle is the oracle that given any function $h : X \rightarrow \{0, 1\}$, takes an independent random sample \mathbf{x} from D and returns $h(\mathbf{x})$.

Identical and closely related models are often studied in the context of distributed statistical estimation with communication constraints (e.g. [Luo05; RWV06; RG06; ZDJW13; SD15; SVW16; SYMK16]).

As shown by Feldman et al. [FGRVX12], to simulate $VSTAT_D(n)$ on a query $\phi : X \rightarrow [0, 1]$, one can take the mean of $O(n)$ bits where each bit is equal to an independent coin flip with bias $\phi(\mathbf{x})$ for a random sample $\mathbf{x} \sim D$. To do this, for each sample, we define a Boolean function ϕ_θ whose value at a point x is the indicator of the condition $\phi(x) \leq \theta$, where θ is chosen uniformly from $[0, 1]$ and then pass ϕ_θ to the 1- $STAT_D$ oracle. This implies that any algorithm that works in the SQ model with $VSTAT$ can be simulated using the 1-bit sampling oracle (see Thm. 2.1 for the formal statement). The correspondence between these oracles is quite tight since n samples from the 1-bit sampling oracle can be simulated (with high probability) using n queries to $VSTAT_D(O(n))$ [FGRVX12]¹. Hence for most of the discussion we will focus on results for the $VSTAT$ oracle.

1.1 The problem

Here we address the case of estimating the expectation of a general real-valued ϕ . Specifically, we want to design an algorithm that given access to $VSTAT_D(n)$ estimates the expectation of ϕ (almost) as well as possible using n random samples of $\phi(\mathbf{x})$. The main that we need to deal with is that the accuracy of estimates in the restricted models we consider depends linearly on the range of ϕ . In contrast, the standard deviation of the mean of ϕ on n i.i.d. samples is equal to $\sqrt{(D[\phi^2] - D[\phi]^2)/n}$ and does not depend on the range of ϕ (which can even be infinite). Equivalently, for a given $\epsilon > 0$, we want to estimate the expectation within ϵ with estimation complexity that is close to the sample complexity of this task. Here the *estimation complexity* of a statistical query algorithm using $VSTAT_D(n)$ refers to the value n ; for an algorithm using $STAT_D(\tau)$ it is $1/\tau^2$.

While mean estimation is the most basic statistical problem, we are not aware of any treatment of the general case for either the SQ or 1-bit sampling model. The closest literature we are aware of considers the problem of mean estimation from sensor networks. In this case each of the sensors can communicate a few bits but does not receive feedback. This corresponds to asking queries non-adaptively in our model. Most of this literature operates under substantially stronger assumptions such as knowing the distribution D up to a shift (or, equivalently, having an unbiased noise from a fixed distribution added to some unknown value) (e.g. [RG06]). In more general treatments we are aware of a bound on the range is assumed with linear dependence of the error on this bound (e.g. [Luo05]). As a result both the nature of algorithms and the resulting bounds appear to be quite different from our setting.

¹Weaker version of this correspondence for $STAT$ was first shown by Ben-David and Dichterman [BD98] and one for a more general communication model by Steinhardt et al. [SVW16].

Another useful property of the VSTAT oracle is that known lower bound techniques against SQ algorithms using $\text{STAT}_D(\tau)$ also apply to the stronger $\text{VSTAT}_D(1/\tau^2)$ oracle. As a result lower bounds for VSTAT usually correspond more tightly to the known computational and sample complexity bounds on the problem (e.g. [FGRVX12; FPV13; FGV15; DKS16]). One could also easily define a stronger statistical query oracle whose accuracy scales with the standard deviation of $\phi(\mathbf{x})$. However, it is unknown how to obtain lower bounds against this stronger oracle using existing analysis techniques (e.g. [FGRVX12; FPV13; Fel16]). From this point of view, our goal is to show that such lower bounds can be obtained indirectly by simulating a stronger oracle using the easier to analyze VSTAT oracle.

The first approach to this general case and the one most commonly used is simply to scale/shift the range of the function ϕ to $[0, 1]$ and then just ask a single query to $\text{VSTAT}_D(n)$ (the answer can then be scaled/shifted back). However that will result in error that is at least R/n , where $R = \sup_{x \in X} \phi(x) - \inf_{x \in X} \phi(x)$ (and in many cases the error will be $\Omega(R/\sqrt{n})$).

Another natural approach would be to truncate ϕ to be in some range $[-R, R]$ such that $|D[\phi] - D[\phi']| \leq \epsilon$, where ϕ' denotes the truncated function. The expectation of ϕ' can then be estimated by scaling it to $[0, 1]$ as before. This approach would provide a reasonably accurate estimate when $\phi(\mathbf{x})$ is strongly concentrated (e.g. sub-gaussian), a good estimate of the standard deviation is known and we know the mean up to a constant multiple of standard deviations (an example of such analysis appears in [FGV15]). But there are many common settings of ϕ and D where these assumptions would not apply.

1.2 Our contribution

We show that it is possible to deal with the case when no upper bound on variance is known and we are only given a loose upper bound on the second moment of $\phi[\mathbf{x}]$. We prove that access to $\text{VSTAT}_D(n)$ suffices to estimate $D[\phi]$ within $\tilde{O}(\sigma/\sqrt{n})$, where $\sigma = \sqrt{D[\phi^2] - D[\phi]^2}$ is the standard deviation of $\phi(\mathbf{x})$. Thus $\text{VSTAT}_D(n)$ can be used to obtain an estimate with accuracy comparable to the accuracy implied by the Bernstein concentration inequality for a sum of i.i.d. random variables combined with the median of means technique (without the median step the accuracy will not be range-independent with high probability). The number of queries needed to implement this algorithm is logarithmic in the upper bound on the second moment and n .

Theorem 1.4. *There exists a statistical query algorithm that given an integer n , $\zeta > 0$, $B > 0$, $\phi : X \rightarrow \mathbb{R}$ such that $D[\phi^2] \leq B^2$ and access to $\text{VSTAT}_D(n)$, outputs a value v such that $|D[\phi] - v| \leq \frac{8\sigma \log(8n)}{\sqrt{n}} + \zeta$, where $\sigma^2 \doteq D[\phi^2] - D[\phi]^2$. The algorithm uses $3 \log(4nB/\zeta^2)$ queries.*

The algorithm is based on several relatively simple ideas. First, the bound on the moment implies that for a suitable choice of quantiles of the distribution of $\phi(\mathbf{x})$ we can truncate the range of ϕ to the values of these quantiles without changing the mean of ϕ too much. We show that $\text{VSTAT}(n)$ oracle allows to find the necessary quantiles. This step requires discretizing the range of ϕ and contributes the additional ζ term in the error. Note that we cannot eliminate this error term without knowing a lower bound on σ but the dependence of the number of queries on $1/\zeta$ is just logarithmic. Further, we show that after this truncation, the length of the range of the function becomes $O(B/\epsilon)$. We then describe an algorithm whose error scales linearly with the square root of the second moment of $\phi(\mathbf{x})$ instead of the square root of the mean as is guaranteed by VSTAT. The algorithm reduces the computation of the mean of ϕ to multiple computations of a mean of ϕ conditioned on being in an interval of constant dynamic range. A logarithmic number of such intervals are sufficient to cover the range of ϕ after the truncation². Finally, to go from an error bounded in terms of

²We remark that this step is different from computing the expectation of each of the (most-significant) bits of the value of ϕ .

the second moment to an error bounded in terms of the standard deviation we show that after shifting the function ϕ by the value of an approximate median, the second moment will be upper-bounded by a constant multiple of the variance.

An immediate corollary of this result and Theorem 2.1 is the following upper bound on the accuracy that can be obtained using n queries to 1-bit sampling oracle.

Theorem 1.5. *There exists an algorithm that given $B > 0, \delta > 0, \zeta > 0, \phi : X \rightarrow \mathbb{R}$ such that $D[\phi^2] \leq B^2$ and access to n queries to 1-STAT $_D$, with probability $\geq 1 - \delta$, outputs a value v such that $|D[\phi] - v| = \frac{\sigma \cdot \tilde{O}(\log n \cdot \log(nB/\zeta) \cdot \log(1/\delta))}{\sqrt{n}} + \zeta$, where $\sigma^2 \doteq D[\phi^2] - D[\phi]^2$.*

We also give a simpler algorithm and bounds for the case when we have an upper bound on σ and/or already have an estimate of $D[\phi]$ within $O(\sigma)$ (see Lemma 3.2). In this case the algorithm's queries are non-adaptive, that is, query functions do not depend on answers to previous queries. Adaptivity of queries is particularly important for the distributed setting that motivates the 1-bit sampling model since queries need to be communicated back to clients that hold the samples. This step is likely to be slow or even impossible in some systems (such as the sensor systems considered in [Luo05]).

We remark that while the bounds we give are stated primarily for VSTAT $_D(n)$, they imply identical upper bounds for STAT $_D(1/n)$. However STAT $_D(1/n)$ has estimation complexity of n^2 and better bounds for it might be possible in some cases. We also note that we have not attempted to optimize the constants or even the log factors for the 1-bit sampling case. Deriving tighter bounds for this case directly is a natural open problem.

1.3 Applications

As an example application we consider the problems of mean vector estimation and stochastic optimization in the SQ model [FGV15]. In the ℓ_2 mean vector estimation problem the goal is to estimate $\bar{x} \doteq \mathbf{E}_{\mathbf{x} \sim D}[\mathbf{x}]$ within ϵ in the ℓ_2 norm, where D is a distribution over vectors in \mathbb{R}^d . This problem is central to implementation of gradient-based (or first-order) optimization methods in the stochastic setting. In [FGV15] it was showed that using either a randomly rotated basis or Kashin's representation [LV10] this problem can be solved using $O(d)$ queries to STAT $_D(\Omega(\epsilon))$ when D is supported on the unit ball. A related technique is also used in [SYMK16] for a low communication model. Our results give a simpler and more general algorithm for the problem using the VSTAT $_D(\tilde{O}(1/\epsilon^2))$ oracle or queries to 1-STAT $_D$.

Corollary 1.6. *There exists a statistical query algorithm that given $\epsilon \in (0, 1)$ and $B > \epsilon$, for any distribution D over \mathbb{R}^d such that $\mathbf{E}_{\mathbf{x} \sim D}[\|\mathbf{x}\|_2^2] \leq B^2$ and $\mathbf{E}_{\mathbf{x} \sim D}[\|\mathbf{x} - \bar{x}\|_2^2] \leq 1$ outputs a vector \hat{x} such that $\|\hat{x} - \bar{x}\|_2 \leq \epsilon$, where $\bar{x} \doteq \mathbf{E}_{\mathbf{x} \sim D}[\mathbf{x}]$. The algorithm uses $O(d(\log(dB/\epsilon)))$ queries to VSTAT $_D(O((\log(1/\epsilon)/\epsilon)^2))$.*

High dimensional mean estimation is the key component of the SQ algorithms for stochastic convex optimization and learning of linear threshold functions. Specifically, they are used to estimate the gradient of the expected objective function in a variety of gradient descent-based algorithms. Without constraints on data access, it is well-known that stochastic gradient descent will achieve low expected error as long as the variance of the gradients is bounded (e.g. [NJLS09]). It is natural to ask whether this case can also be handled in the SQ and 1-bit sampling models. Corollary 4.1 immediately implies that the optimization algorithms in [FGV15] can be extended from the setting in which the (sub-)gradients of all the functions in the support of the input distribution are uniformly upper-bounded to the setting in which only variance of (sub-)gradient vectors is upper-bounded (additional details can be found in [FGV15]). These corollaries can in turn be easily translated to the 1-bit sampling model and we describe an example of such corollary in Section 4.

2 Preliminaries

For integer $n \geq 1$ let $[n] \doteq \{1, \dots, n\}$. Random variables are denoted by bold letters, e.g., \mathbf{x} . We denote the indicator function of an event A (i.e., the function taking value zero outside of A , and one on A) by $\mathbf{1}_{\{A\}}$.

We will use the following notation for a truncation operation. For a real value z and $a \in \mathbb{R}^+$, let

$$m_a(z) := \begin{cases} z & \text{if } |z| \leq a \\ a & \text{if } z > a \\ -a & \text{if } z < -a. \end{cases}$$

Let $r_a(z) \doteq z - m_a(z)$.

We consider mean estimation with two related oracles. The first one is the VSTAT oracle from [FGRVX12] (see Def. 1.2) that strengthens the STAT_D oracle of Kearns [Kea98] (see Def. 1.1). Clearly, $\text{VSTAT}_D(n)$ is at least as strong as $\text{STAT}_D(1/\sqrt{n})$ (but no stronger than $\text{STAT}_D(1/n)$). When the range of ϕ is $\{0, 1\}$, one way to think about VSTAT is as providing a confidence interval for the bias p , namely $[v - \tau_v, v + \tau_v]$, where $\tau_v \approx \max\{1/n, \sqrt{(v(1-v)/n)}\}$. The accuracy τ_v that VSTAT ensures corresponds (up to a small constant factor) to the width of the standard confidence interval (say, with 95% coverage) for the bias p of a Bernoulli random variable given n independent samples (e.g. Clopper-Pearson interval [CP34]). Therefore, at least for Boolean queries, it captures precisely the accuracy that can be inferred when estimating the mean using n random samples. For more general real-valued queries (even those with range in $[0, 1]$) the accuracy of $\text{VSTAT}_D(n)$ might no longer reflect the accuracy that will be achieved given n i.i.d. samples.

The second oracle we consider is the 1-bit sampling oracle (see Def. 1.3). In the context of learning theory it was first studied by Ben-David and Dichterman [BD98]. Our results for this oracle will be obtained via the following simulation of the VSTAT_D oracle using the 1- STAT_D oracle.

Theorem 2.1 ([FGRVX12]). *Let $n, q > 0$ be any integers and $\delta > 0$. For any algorithm \mathcal{A} that asks at most q queries to $\text{VSTAT}_D(n)$ there exists an algorithm \mathcal{B} that provides answers that satisfy the accuracy guarantees of $\text{VSTAT}_D(n)$ to all the queries of \mathcal{A} with probability at least $1 - \delta$. \mathcal{B} uses $O(qn \cdot \log(q/\delta))$ queries to 1- STAT_D .*

3 Mean estimation using statistical queries

3.1 Non-adaptive algorithm

We start by showing how the accuracy of mean estimation can be improved given a reasonably tight bound on the second moment of ϕ . The algorithm for this problem has the advantage of being non-adaptive.

In the first step we show that the dependence on $\sqrt{D[\phi]}$ in the accuracy guarantees of VSTAT can be strengthened to $\sqrt{D[\phi^2]}$ at the expense of logarithmic factors in the complexity. We will use \log to denote the logarithm to base 2.

Lemma 3.1. *There exists a statistical query algorithm that, for $n, R > 0$, any function $\phi : X \rightarrow [0, R]$ and any input distribution D over X , outputs a value v such that $|D[\phi] - v| \leq \frac{4R}{n} + \frac{2s \cdot \log n}{\sqrt{n}}$, where $s \doteq \sqrt{D[\phi^2]}$. The algorithm uses at most $\log n$ (non-adaptive) queries to $\text{VSTAT}_D(n)$.*

Proof. We assume for simplicity that $R = 1$ since we can always scale ϕ to this setting and then scale back the result. We let $t = \lfloor \log n \rfloor$ and observe that

$$D[\phi] = \mathbf{E}_D [\phi(\mathbf{x}) \cdot \mathbf{1}_{\{\phi(\mathbf{x}) \in [0, 2^{-t}]\}}] + \sum_{i \in [t]} \mathbf{E}_D [\phi(\mathbf{x}) \cdot \mathbf{1}_{\{\phi(\mathbf{x}) \in (2^{-i}, 2^{-i+1}]\}}]. \quad (1)$$

For every $i \in [t]$, we define ϕ_i to be the restriction of ϕ to values in the interval $(2^{-i}, 2^{-i+1}]$, scaled and shifted to the range $[0, 1]$. Namely

$$\phi_i(x) \doteq 2^{i-1} \cdot \phi(x) \cdot \mathbf{1}_{\{\phi(\mathbf{x}) \in (2^{-i}, 2^{-i+1}]\}}.$$

Note that for every x , $\phi_i(x) \in [0, 1]$. Using this definition, we can rewrite eq.(1) as

$$D[\phi] = \mathbf{E}_D [\phi(\mathbf{x}) \cdot \mathbf{1}_{\{\phi(\mathbf{x}) \in [0, 2^{-t}]\}}] + \sum_{i \in [t]} 2^{-i+1} \cdot D[\phi_i]. \quad (2)$$

For every $i \in [t]$, we query function ϕ_i to $\text{VSTAT}_D(n)$ to get an estimate v_i of $D[\phi_i]$. By Chebyshev's inequality,

$$D[\phi_i] \leq \mathbf{Pr}_D[\phi(\mathbf{x}) > 2^{-i}] \leq \frac{D[\phi^2]}{2^{-2i}} = 2^{2i} \cdot s^2.$$

Therefore, by the definition of $\text{VSTAT}_D(n)$,

$$|v_i - D[\phi_i]| \leq \max \left\{ \frac{1}{n}, \sqrt{\frac{D[\phi_i]}{n}} \right\} \leq \frac{1}{n} + \sqrt{\frac{D[\phi_i]}{n}} \leq \frac{1}{n} + \frac{2^i s}{\sqrt{n}}. \quad (3)$$

Let $v \doteq \sum_{i \in [t]} 2^{-i+1} \cdot v_i$.

Then, by eq. (2) and eq. (3) we get that

$$\begin{aligned} |D[\phi] - v| &= \left| \mathbf{E}_D [\phi(\mathbf{x}) \cdot \mathbf{1}_{\{\phi(\mathbf{x}) \in [0, 2^{-t}]\}}] + \sum_{i \in [t]} 2^{-i+1} \cdot (D[\phi_i] - v_i) \right| \\ &\leq \mathbf{E}_D [\phi(\mathbf{x}) \cdot \mathbf{1}_{\{\phi(\mathbf{x}) \in [0, 2^{-t}]\}}] + \sum_{i \in [t]} 2^{-i+1} \cdot |D[\phi_i] - v_i| \\ &\leq 2^{-t} + \sum_{i \in [t]} 2^{-i+1} \cdot \left(\frac{1}{n} + \frac{2^i s}{\sqrt{n}} \right) \\ &\leq \frac{2}{n} + \frac{2}{n} + \frac{2t \cdot s}{\sqrt{n}} \leq \frac{4}{n} + \frac{2s \cdot \log n}{\sqrt{n}}. \end{aligned}$$

□

In the algorithm above the parameter of VSTAT scales linearly with the range R . However given a bound on $D[\phi^2]$ we can truncate the range of ϕ with little change in expectation. We use this in the lemma below.

Lemma 3.2. *There exists a statistical query algorithm that given $B > 0$ and $\epsilon \in (0, B/16]$, for any distribution D and function $\phi : X \rightarrow \mathbb{R}^+$ such that $D[\phi^2] \leq B^2$ outputs a value v such that $|D[\phi] - v| \leq \epsilon$. The algorithm uses at most $3 \log(B/\epsilon)$ (non-adaptive) queries to $\text{VSTAT}_D((8B \log(B/\epsilon)/\epsilon)^2)$.*

Proof. We first assume that $B = 1$. Observe that we can truncate the range of ϕ to $[0, a]$ for $a \doteq 4/\epsilon$ without a significant change in the expectation. Namely, we claim that $|D[m_a(\phi)] - D[\phi]| \leq \epsilon/4$. To see this when $a = 4/\epsilon$ note that $\phi - m_a(\phi) = r_a(\phi)$ and

$$\mathbf{E}[r_a(\phi)] = \int_0^\infty \Pr[r_a(\phi) \geq t] dt = \int_0^\infty \Pr[\phi \geq t + a] dt = \int_a^\infty \Pr[\phi \geq t] dt \leq \int_a^\infty \frac{1}{t^2} dt = \frac{1}{a} \leq \epsilon/4,$$

where we used Chebyshev's inequality to obtain a bound on $\Pr[\phi \geq t]$.

Now to estimate $D[m_a(\phi)]$ we use Lemma 3.1 with $n = (4 \log(1/\epsilon)/\epsilon)^2$. For the given range and our assumption that $D[\phi^2] \leq 1$, this leads to an error of at most

$$\frac{16}{(8 \log(1/\epsilon)/\epsilon)^2 \cdot \epsilon} + \frac{2 \log(8 \log(1/\epsilon)/\epsilon)}{8 \log(1/\epsilon)/\epsilon} \leq \frac{\epsilon}{4 \log(1/\epsilon)^2} + \frac{\epsilon \cdot \log(8 \log(1/\epsilon)/\epsilon)}{4 \log(1/\epsilon)} \leq \frac{3\epsilon}{4},$$

where we used the assumption that $\epsilon \leq 1/16$ to obtain the last inequality. Altogether this implies that the output of the lemma on the truncated function will have an error of at most ϵ . Finally, to generalize the analysis to any $B > 0$ we simply scale the random variable by $1/B$ and estimate its mean within ϵ/B . \square

3.2 Quantile estimation

To deal with the case where we do not have a tight upper bound on the second moment we will need the following lemma about the estimation complexity of finding approximate distribution quantiles using the VSTAT oracle. The algorithm itself is just the folklore algorithm for finding an (approximate) quantile of a distribution using binary search.

Lemma 3.3. *Let ϕ be a function with a finite range $Z \subset \mathbb{R}$. There exists an SQ algorithm that given p and δ such that $1 \geq p \geq 2\delta > 0$, outputs a point $a \in Z$ such that $\Pr_D[\phi \geq a] \geq p - \delta$ and $\Pr_D[\phi > a] < p$. The algorithm uses $\log(|Z|)$ queries to $VSTAT_D(4p/\delta^2)$.*

Proof. For any $z \in Z$, let $p_z \doteq \Pr_D[\phi \geq z]$. We perform a binary search to find the largest point z such that the estimate of $\Pr_D[\phi \geq z]$ given by $VSTAT_D(4p/\delta^2)$ is at least $p - \delta/2$. We denote the point by a and refer to estimates we have obtained by \tilde{p}_z . By definition, $\tilde{p}_a \geq p - \delta/2$ and therefore $p_a \geq p - \delta$, since otherwise

$$\tilde{p}_a < p - \delta + \max \left\{ \sqrt{\frac{p - \delta}{4p/\delta^2}}, \frac{\delta^2}{4p} \right\} < p - \delta + \delta/2 = p - \delta/2.$$

On the other hand for the smallest point in Z larger than a (denote it by a_+) we know that $\tilde{p}_{a_+} < p - \delta/2$. This implies that $p_{a_+} < p$ since otherwise

$$\tilde{p}_{a_+} \geq p_{a_+} - \max \left\{ \sqrt{\frac{p_{a_+}}{4p/\delta^2}}, \frac{\delta^2}{4p} \right\} = p_{a_+} - \frac{\delta}{2} \cdot \sqrt{\frac{p_{a_+}}{p}} \geq \sqrt{p \cdot p_{a_+}} - \delta/2 \geq p - \delta/2,$$

where in the penultimate step we multiplied the inequality by $\sqrt{\frac{p}{p_{a_+}}}$ which is at most 1 by our contrapositive assumption. This means that $\Pr_D[\phi > a] = p_{a_+} < p$. \square

As a special case we obtain the following corollary.

Corollary 3.4. *Let ϕ be a function over X with a finite range $Z \subset \mathbb{R}$. There exists an SQ algorithm that given an integer n outputs a point $a \in Z$ such that $\Pr_D[\phi \geq a] \geq 8/n$ and $\Pr_D[\phi > a] < 16/n$. The algorithm uses $\log(|Z|)$ queries to $VSTAT_D(n)$.*

One natural way to apply this lemma in the continuous setting is to first discretize the range with some step size ζ and then use Lemma 3.3. This leads to the following version of Lemma 3.3:

Corollary 3.5. *Let $\phi : X \rightarrow [-B, B]$ be a function and $\zeta > 0$. There exists an SQ algorithm that given p and δ such that $1 \geq p \geq 2\delta > 0$, outputs a point a that is an integer multiple of ζ such that $\Pr_D[\phi \geq a] \geq p - \delta$ and $\Pr_D[\phi \geq a + \zeta] < p$. The algorithm uses $\lceil \log(2B/\zeta) \rceil$ queries to $VSTAT_D(4p/\delta^2)$.*

3.3 General case

Lemma 3.2 has estimation complexity that scales linearly with the bound on the second moment and is the strongest statement we give without relying on adaptive queries. Next, using Lemma 3.3, we describe a procedure that can estimate the mean with just logarithmic dependence on a given upper bound on the second moment.

Lemma 3.6. *There exists a statistical query algorithm that given an integer n , $\zeta > 0$, $B > 0$ and $\phi : X \rightarrow \mathbb{R}^+$ such that $D[\phi^2] \leq B^2$, outputs a value v such that $|D[\phi] - v| \leq \frac{2s \log(8n)}{\sqrt{n}} + \zeta$, where $s \doteq \sqrt{D[\phi^2]}$. The algorithm uses $\log(4Bn/\zeta^2)$ queries to $VSTAT_D(n)$.*

Proof. We start by truncating and discretizing the range of ϕ , namely let ψ be equal to ϕ rounded down to the closest multiple of $\zeta/2$ and truncated at $2B/\zeta$. As in Lemma 3.2, we note that the condition $D[\phi^2] \leq B^2$ implies that truncation to range $[0, 2B/\zeta]$ step can affect the expectation by at most $\zeta/2$. Clearly, rounding down to the closest multiple of $\zeta/2$ also affects the expectation by at most $\zeta/2$. This means that ψ has range Z of size at most $4B/\zeta^2$, $|D[\phi] - D[\psi]| \leq \zeta$ and $D[\psi^2] \leq s^2$.

We now further truncate the range of ψ to exclude values that are in the top $8/n$ quantile. Namely, by Corollary 3.4, we can find a value $a \in Z$ such that $\Pr_D[\psi \geq a] \geq 8/n$ and $\Pr_D[\psi > a] < 16/n$. Let $\psi_a(x)$ be defined as $m_a(\psi(x))$.

We first observe that $s^2 = D[\psi^2] \geq a^2 \cdot 8/n$. Therefore $a \leq \frac{s\sqrt{n}}{\sqrt{8}}$. Next note that

$$\begin{aligned} |D[\psi] - D[\psi_a]| &\leq \sum_{z \in Z, z > a} z \cdot \Pr_D[\psi = z] \\ &\leq \sqrt{\sum_{z \in Z, z > a} \Pr_D[\psi = z]} \cdot \sqrt{\sum_{z \in Z, z > a} z^2 \cdot \Pr_D[\psi = z]} \\ &\leq \sqrt{\Pr_D[\psi > a]} \cdot \sqrt{D[\psi^2]} \leq \frac{4s}{\sqrt{n}}, \end{aligned} \tag{4}$$

where we used the Cauchy-Schwartz inequality to obtain the second line.

We can now apply Lemma 3.1 to ψ_a and obtain a value v such that

$$|D[\psi_a] - v| \leq \frac{2s \log n}{\sqrt{n}} + \frac{4a}{n} \leq \frac{2s \log n + \sqrt{2}s}{\sqrt{n}}.$$

Combining this with eq. (4) and the properties of ψ we get the claim. \square

We can easily extend these results to variables with possibly negative range by estimating the mean in the positive and the negative range separately. Further, one does not necessarily need to split the range at 0. Any value a can be used and the resulting bound will be in terms of $s_a \doteq \sqrt{D[(\phi - a)^2]}$ instead of $s_0 = \sqrt{D[\phi^2]}$. Naturally, in order to reduce the error we should use $a = D[\phi]$ which would give an error in terms of variance $\sigma^2 = D[(\phi - a)^2] - D[\phi]^2$. The true mean is not known to us but, as we show below, we can use an (approximate) median instead.

Lemma 3.7. *Let \mathbf{z} be a random variable over \mathbb{R} and let a be any point such that $\Pr[\mathbf{z} \geq a] \geq 1/3$ and $\Pr[\mathbf{z} \leq a] \geq 1/3$. Then $\mathbf{E}[(\mathbf{z} - a)^2] \leq 4(\mathbf{E}[\mathbf{z}^2] - \mathbf{E}[\mathbf{z}]^2)$.*

Proof. Let $\bar{z} \doteq \mathbf{E}[\mathbf{z}]$ and $\sigma^2 \doteq \mathbf{E}[(\mathbf{z} - \bar{z})^2]$. Observe that

$$\sigma^2 = \mathbf{E}[(\mathbf{z} - \bar{z})^2] \geq (\bar{z} - a)^2 \cdot \Pr[\mathbf{z} \geq |(\bar{z} - a)|] \geq (\bar{z} - a)^2/3.$$

Now

$$\mathbf{E}[(\mathbf{z} - a)^2] \leq \sigma^2 + (\bar{z} - a)^2 \leq 4\sigma^2.$$

□

Note that for a function with range Z such an approximate median can be found using $\log(|Z|)$ queries to $\text{VSTAT}_D(6)$. Therefore we immediately get the following results:

Theorem 3.8. *There exists a statistical query algorithm that given an integer n , $\zeta > 0$, $B > 0$ and $\phi : X \rightarrow \mathbb{R}$ such that $D[\phi^2] \leq B^2$, outputs a value v such that $|D[\phi] - v| \leq \frac{8\sigma \log(8n)}{\sqrt{n}} + \zeta$, where $\sigma^2 \doteq D[\phi^2] - D[\phi]^2$. The algorithm uses $3 \log(4nB/\zeta^2)$ queries to $\text{VSTAT}_D(n)$.*

An alternative way to state essentially the same result is as follows.

Corollary 3.9. *There exists a statistical query algorithm that given $B > \zeta > 0$, $\epsilon \in (0, 1)$, for any distribution D and function $\phi : X \rightarrow \mathbb{R}$ such that $D[\phi^2] \leq B^2$ outputs a value v such that $|D[\phi] - v| \leq \epsilon\sigma + \zeta$. The algorithm uses $O(\log(B/(\epsilon\zeta)))$ queries to $\text{VSTAT}_D(O((\log(1/\epsilon)/\epsilon)^2))$, where $\sigma^2 \doteq D[\phi^2] - D[\phi]^2$.*

4 Applications

As application we consider the problems of mean vector estimation [FGV15]. In the ℓ_2 mean vector estimation problem the goal is to estimate $\bar{\mathbf{x}} \doteq \mathbf{E}_{\mathbf{x} \sim D}[\mathbf{x}]$ within ϵ in ℓ_2 norm, where D is a distribution over vectors in \mathbb{R}^d . Our results give a simple and general algorithm for solving this problem using the VSTAT oracle.

Corollary 4.1. *There exists a statistical query algorithm that given $\epsilon \in (0, 1)$ and $B > \epsilon$, for any distribution D over \mathbb{R}^d such that $\mathbf{E}_{\mathbf{x} \sim D}[\|\mathbf{x}\|_2^2] \leq B^2$ and $\mathbf{E}_{\mathbf{x} \sim D}[\|\mathbf{x} - \bar{\mathbf{x}}\|_2^2] \leq 1$ outputs a vector $\hat{\mathbf{x}}$ such that $\|\hat{\mathbf{x}} - \bar{\mathbf{x}}\|_2 \leq \epsilon$, where $\bar{\mathbf{x}} \doteq \mathbf{E}_{\mathbf{x} \sim D}[\mathbf{x}]$. The algorithm uses $O(d \log(dB/\epsilon))$ queries to $\text{VSTAT}_D(O((\log(1/\epsilon)/\epsilon)^2))$.*

Proof. We apply the algorithm given in Thm. 1.4 to each of the coordinates of \mathbf{x} with $\zeta = \epsilon/\sqrt{2d}$. Namely, for each $i \in [d]$ we use Thm. 1.4 to estimate the expectation of function $\phi_i(x) = x_i$ and let \hat{x}_i be the result. For every $i \in [d]$, $\mathbf{E}_{\mathbf{x} \sim D}[\|\mathbf{x}_i\|_2^2] \leq B^2$ and hence, by setting $n = c(\log(1/\epsilon)/\epsilon)^2$ for an appropriately chosen constant c , we will obtain that the error in the estimation of coordinate i is at most

$$\frac{\epsilon}{\sqrt{2}} \cdot \sqrt{\mathbf{E}_{\mathbf{x} \sim D}[\|\mathbf{x}_i\|_2^2] - \bar{x}_i^2} + \frac{\epsilon}{\sqrt{2d}}.$$

By observing that

$$\sum_{i \in [d]} \mathbf{E}_{\mathbf{x} \sim D}[\|\mathbf{x}_i\|_2^2] - \bar{x}_i^2 = \mathbf{E}_{\mathbf{x} \sim D}[\|\mathbf{x} - \bar{\mathbf{x}}\|_2^2] \leq 1$$

we obtain the claim. □

We remark that if the SQ algorithm \mathcal{A} is non-adaptive then the simulation in Theorem 2.1 produces a non-adaptive algorithm. This is particularly useful in the distributed setting since it allows the bits from each of the samples to be communicated in parallel and does not require any communication back to the clients.

These corollaries can in turn be translated to the 1-bit sampling model.

Combining Corollary 3.9 and Theorem 2.1 we obtain the following algorithm for estimating the mean.

Corollary 4.2. *There exists an algorithm that given $\epsilon, \delta \in (0, 1)$ and $B > \epsilon$, for any distribution D over \mathbb{R}^d such that $\mathbf{E}_{\mathbf{x} \sim D}[\|\mathbf{x}\|_2^2] \leq B^2$ and $\mathbf{E}_{\mathbf{x} \sim D}[\|\mathbf{x} - \bar{\mathbf{x}}\|_2^2] \leq 1$ outputs a vector $\hat{\mathbf{x}}$ such that, with probability at least $1 - \delta$, $\|\hat{\mathbf{x}} - \bar{\mathbf{x}}\|_2 \leq \epsilon$, where $\bar{\mathbf{x}} \doteq \mathbf{E}_{\mathbf{x} \sim D}[\mathbf{x}]$. The algorithm uses $\tilde{O}(d/\epsilon^2 \cdot \log(B/(\delta\epsilon)))$ queries to 1-STAT $_D$.*

We remark that in the high dimensional setting linear dependence on d is unavoidable for the 1-bit sampling model [ZDJW13] whereas given entire samples, the achievable accuracy is dimension independent.

Acknowledgements

I thank the anonymous reviewers of Algorithmic Learning Theory 2017 conference for insightful comments and careful proofreading of this work.

References

- [BD98] S. Ben-David and E. Dichterman. “Learning with Restricted Focus of Attention”. In: *J. Comput. Syst. Sci.* 56.3 (1998), pp. 277–298.
- [CP34] C. J. Clopper and E. S. Pearson. “The Use of Confidence or Fiducial Limits Illustrated in the Case of the Binomial”. In: *Biometrika* 26.4 (1934), pp. 404–413.
- [DKS16] I. Diakonikolas, D. M. Kane, and A. Stewart. “Statistical Query Lower Bounds for Robust Estimation of High-dimensional Gaussians and Gaussian Mixtures”. In: *CoRR* abs/1611.03473 (2016).
- [DR14] C. Dwork and A. Roth. *The Algorithmic Foundations of Differential Privacy*. Vol. 9. 3-4. 2014, pp. 211–407.
- [Fel16] V. Feldman. “A General Characterization of the Statistical Query Complexity”. In: *CoRR* abs/1608.02198 (2016). Extended abstract in COLT 2017.
- [FGRVX12] V. Feldman, E. Grigorescu, L. Reyzin, S. Vempala, and Y. Xiao. “Statistical Algorithms and a Lower Bound for Detecting Planted Cliques”. In: *arXiv, CoRR* abs/1201.1214 (2012). Extended abstract in STOC 2013.
- [FGV15] V. Feldman, C. Guzman, and S. Vempala. “Statistical Query Algorithms for Mean Vector Estimation and Stochastic Convex Optimization”. In: *CoRR* abs/1512.09170 (2015). Extended abstract in SODA 2017.
- [FPV13] V. Feldman, W. Perkins, and S. Vempala. “On the Complexity of Random Satisfiability Problems with Planted Solutions”. In: *CoRR* abs/1311.4821 (2013). Extended abstract in STOC 2015.
- [Kea98] M. Kearns. “Efficient noise-tolerant Learning from statistical queries”. In: *Journal of the ACM* 45.6 (1998), pp. 983–1006.
- [Luo05] Z.-Q. Luo. “Universal decentralized estimation in a bandwidth constrained sensor network”. In: *IEEE Transactions on information theory* 51.6 (2005), pp. 2210–2219.

- [LV10] Y. Lyubarskii and R. Vershynin. “Uncertainty principles and vector quantization”. In: *Information Theory, IEEE Transactions on* 56.7 (2010), pp. 3491–3501.
- [NJLS09] A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro. “Robust stochastic approximation approach to stochastic programming.” In: *SIAM J. Optim.* 19.4 (2009), pp. 1574–1609. ISSN: 1052-6234; 1095-7189/e.
- [RG06] A. Ribeiro and G. B. Giannakis. “Bandwidth-constrained distributed estimation for wireless sensor networks-part I: Gaussian case”. In: *IEEE transactions on signal processing* 54.3 (2006), pp. 1131–1143.
- [RWV06] R. Rajagopal, M. J. Wainwright, and P. Varaiya. “Universal quantile estimation with feedback in the communication-constrained setting”. In: *Information Theory, 2006 IEEE International Symposium on*. IEEE. 2006, pp. 836–840.
- [SD15] J. Steinhardt and J. C. Duchi. “Minimax rates for memory-bounded sparse linear regression”. In: *COLT*. 2015, pp. 1564–1587.
- [SVW16] J. Steinhardt, G. Valiant, and S. Wager. “Memory, Communication, and Statistical Queries”. In: *COLT*. 2016, pp. 1490–1516.
- [SYMK16] A. T. Suresh, F. X. Yu, H. B. McMahan, and S. Kumar. “Distributed mean estimation with limited communication”. In: *arXiv preprint arXiv:1611.00429* (2016).
- [ZDJW13] Y. Zhang, J. C. Duchi, M. I. Jordan, and M. J. Wainwright. “Information-theoretic lower bounds for distributed statistical estimation with communication constraints”. In: *NIPS*. 2013, pp. 2328–2336.