

IBM Research Report

Calibrating Noise to Variance in Adaptive Data Analysis

Vitaly Feldman
Google Research

Thomas Steinke
IBM Research Division
Almaden Research Center
650 Harry Road
San Jose, CA 95120-6099
USA



Calibrating Noise to Variance in Adaptive Data Analysis

Vitaly Feldman*

Thomas Steinke†

Abstract

Datasets are often used multiple times and each successive analysis may depend on the outcome of previous analyses. Standard techniques for ensuring generalization and statistical validity do not account for this adaptive dependence. A recent line of work studies the challenges that arise from such adaptive data reuse by considering the problem of answering a sequence of “queries” about the data distribution where each query may depend arbitrarily on answers to previous queries.

The strongest results obtained for this problem rely on differential privacy – a strong notion of algorithmic stability with the important property that it “composes” well when data is reused. However the notion is rather strict, as it requires stability under replacement of an arbitrary data element. The simplest algorithm is to add Gaussian (or Laplace) noise to distort the empirical answers. However, analysing this technique using differential privacy yields suboptimal accuracy guarantees when the queries have low variance.

Here we propose a relaxed notion of stability that also composes adaptively. We demonstrate that a simple and natural algorithm based on adding noise scaled to the standard deviation of the query provides our notion of stability. This implies an algorithm that can answer statistical queries about the dataset with substantially improved accuracy guarantees for low-variance queries. The only previous approach that provides such accuracy guarantees is based on a more involved differentially private median-of-means algorithm and its analysis exploits stronger “group” stability of the algorithm.

1 Introduction

The central challenge in statistical data analysis is to infer the properties of some unknown population given only a small number of samples from that population. While a plethora of techniques for guaranteeing statistical validity are available, few techniques can account for the effects of *adaptivity*. Namely, if a single dataset is used multiple times, then the choice of which subsequent analyses to perform may depend on the outcomes of previous analyses. This adaptive dependence increases the risk of overfitting — that is, inferring a conclusion that does not generalize to the underlying population.

To formalize this problem, Dwork et al. [DFH⁺14] and subsequent works [HU14, SU15, BNS⁺16, FS17, etc.] study the following question: How many data samples are necessary to accurately answer a sequence of queries about the data distribution when the queries are chosen adaptively – that is,

*Google Research (part of this work was done while at IBM Research – Almaden and while visiting the Simons Institute, UC Berkeley) vitaly@post.harvard.edu

†IBM Research – Almaden alkls@thomas-steinke.net

each query can depend on answers to previous queries? Each query corresponds to a procedure that the analyst wishes to execute on the data. The goal is to design an algorithm that provides answers to these adaptive queries that are close to answers that would have been obtained had each corresponding analysis been run on independent samples freshly drawn from the data distribution.

A common and relatively simple class of queries are statistical queries [Kea98]. A statistical query is specified by a function $\psi : \mathcal{X} \rightarrow [0, 1]$ and corresponds to analyst wishing to compute the true mean $\mathbf{E}_{X \sim \mathcal{P}} [\psi(X)]$ of ψ on the data distribution \mathcal{P} . (This is usually done by using the empirical mean $\frac{1}{n} \sum_{i=1}^n \psi(S_i)$ on a dataset S consisting of n i.i.d. draws from the distribution \mathcal{P} .) For example, such queries can be used to estimate the true loss (or error) of a predictor, the gradient of the loss function, or the moments of the data distribution. Standard concentration results imply that, given n independent samples from \mathcal{P} , k fixed (i.e. not adaptively-chosen) statistical queries can be answered with an additive error of at most $O\left(\sqrt{\log(k)/n}\right)$ with high probability by simply using the empirical mean of each query. At the same time it is not hard to show that, for a variety of simple adaptive sequences of queries, using the empirical mean to estimate the expectation leads to an error of $\Omega(\sqrt{k/n})$ [DFH⁺14]. Equivalently, in the adaptive setting, the number of samples required to ensure fixed error scales linearly (rather than logarithmically in the non-adaptive setting) with the number of queries and, in particular, in the worst case, using empirical estimates gives the same guarantees as using fresh samples for every query (by splitting the dataset into k parts).

Dwork et al. [DFH⁺14] showed that, remarkably, it is possible to quadratically improve the dependence on k in the adaptive setting by simply perturbing the empirical answers. Specifically, let $S \in \mathcal{X}^n$ denote a dataset consisting of n i.i.d. samples from some (unknown) probability distribution \mathcal{P} . Given S , the algorithm receives k adaptively-chosen statistical queries $\psi_1, \dots, \psi_k : \mathcal{X} \rightarrow [0, 1]$ one-by-one and provides k approximate answers $v_1, \dots, v_k \in \mathbb{R}$. Namely, $v_j = \frac{1}{n} \sum_{i=1}^n \psi_j(S_i) + \xi_j$, where each “noise” variable ξ_j is drawn independently from $\mathcal{N}(0, \sigma^2)$. The results of Dwork et al. [DFH⁺14] and subsequent sharper analyses [BNS⁺16, Ste16] show that, with high probability (over the drawing of the sample $S \sim \mathcal{P}^n$, the noise ξ , and the choice of queries), we have the following guarantee

$$\forall j \in \{1, \dots, k\} \quad \left| v_j - \mathbf{E}_{X \sim \mathcal{P}} [\psi_j(X)] \right| \leq O\left(\sqrt{\frac{\sqrt{k \log k}}{n}}\right). \quad (1)$$

This quadratic relationship between n and k was also shown to be optimal in the worst case [HU14, SU15].

The approach of Dwork et al. [DFH⁺14] relies on properties of differential privacy [DMNS06, DKM⁺06] and known differentially private algorithms. Differential privacy is a stability property of an algorithm, namely it requires that replacing any element in the input dataset results in a small change in the output distribution of the algorithm. Specifically, a randomized algorithm $M : \mathcal{X}^n \rightarrow \mathcal{Y}$ is (ϵ, δ) -differentially private if, for all datasets $s, s' \in \mathcal{X}^n$ that differ on a single element and all events $E \subseteq \mathcal{Y}$,

$$\Pr [M(s) \in E] \leq e^\epsilon \Pr [M(s') \in E] + \delta.$$

This stability notion implies that a function output by a differentially private algorithm on a given dataset generalizes to the underlying distribution [DFH⁺14, BNS⁺16]. Specifically, if a differentially private algorithm is run on a dataset drawn i.i.d from any distribution and the algorithm outputs a function, then the empirical mean of that function on the input dataset is close to the expectation of that function on sample from the same distribution.

The second crucial property of differential privacy is that it composes adaptively: running several differentially private algorithms on the same dataset still is differentially private (with somewhat worse parameters) even if each algorithm depends on the output of all the previous algorithms. This property makes it possible to answer adaptively-chosen queries with differential privacy and a number of algorithms have been developed for answering different types of queries. The generalization property of differential privacy then implies that such algorithms can be used to provide answers to adaptively-chosen queries while ensuring generalization [DFH⁺14]. Specifically, the algorithm for answering statistical queries mentioned above is based on the most basic differentially private algorithm: perturbation by adding Laplace or Gaussian noise [DMNS06].

Differential privacy requires that the output distribution of an algorithm does not change much when any element of a dataset is replaced with an arbitrary other element in the domain \mathcal{X} . As a result, the amount of noise that needs to be added to ensure differential privacy scales linearly with the range of the function ψ whose expectation needs to be estimated. If the range of ψ is comparable to the standard deviation of $\psi(x)$ on x drawn from \mathcal{P} (such as then ψ has range $\{0, 1\}$ and mean $1/2$) then the error resulting from addition of noise is comparable to the standard deviation of ψ . However, for queries whose standard deviation is much lower than the range, the error introduced by noise is much worse than the sampling error. Variance is much smaller than the range for a variety of common settings, for example, difference between candidate predictors for the same problem or individual input features when the input is usually sparse.

Achieving error guarantees in the adaptive setting that scale with the standard deviation instead of range is a natural problem. Recently, Feldman and Steinke [FS17] gave a different algorithm that achieves such a guarantee. Specifically, their algorithm ensures that with probability at least $1 - \beta$,

$$\forall j \in \{1, \dots, k\} \quad \left| v_j - \mathbf{E}_{X \sim \mathcal{P}} [\psi_j(X)] \right| \leq \text{sd}(\psi_j(\mathcal{P})) \cdot O \left(\sqrt{\frac{\sqrt{k \log^3(k/\beta)}}{n}} \right) + \beta, \quad (2)$$

where $\text{sd}(\psi_j(\mathcal{P})) = \sqrt{\mathbf{E}_{Y \sim \mathcal{P}} [(\psi_j(Y) - \mathbf{E}_{X \sim \mathcal{P}} [\psi_j(X)])^2]}$ is the standard deviation of ψ_j on the distribution \mathcal{P} and $\beta > 0$ can be chosen arbitrarily. Their algorithm is based on an approximate version of the median of means algorithm and its analysis still relies on differential privacy. (Their results extend beyond statistical queries, but we restrict our attention to statistical queries in this paper.)

In this work, we ask: does the natural approach of perturbing the empirical answers with noise scaled to the standard deviation suffice to answer adaptive queries with accuracy scaling to sampling error? To answer this seemingly simple question, we address a more fundamental problem: does there exist a notion of stability that has the advantages of differential privacy (namely, allows adaptive composition and implies generalization) but avoids the poor dependence on the worst-case sensitivity of the query. This algorithm was analyzed by Bassily and Freund [BF16] via a notion of typical stability they introduced. Their analysis shows that the algorithm will ensure the correct scaling of the error with standard deviation but it does not improve on the naive mechanisms in terms of scaling with k . Several works have considered relaxations of differential privacy in this context. For example, Bassily et al. [BNS⁺16] considered a notion of stability based on using KL divergence or total variation distance in place of differential privacy (which can be defined in terms of approximate max divergence). Wang et al. [WLF16] considered the expected KL divergence between the output of the algorithm when run on a random i.i.d dataset versus the same dataset with one element replaced by a fresh sample; unfortunately, their stability

definition does *not* compose adaptively. Notions based on the mutual information between the dataset and the output of the algorithm and their relationship to differential privacy have also been studied [DFH⁺15, RZ16, RRST16, RRT⁺16, XR17]. However, to the best of our knowledge, these approaches do not give a way to analyze the calibrated noise addition that ensures correct dependence on k .

1.1 Our Contributions

We demonstrate that calibrating the variance of the perturbation to the empirical variance of the query suffices to ensure generalization, as long as the noise rate does not become too small. To ensure this lower bound on the noise rate we simply add a second order term to the variance of the perturbation. Specifically, our algorithm is described in Figure 1. The only difference between our algorithm and previous work [DFH⁺14, BNS⁺16] is that in prior work the variance of the Gaussian perturbation is fixed.




Figure 1: Calibrating noise to variance for answering adaptive queries.

We prove that this algorithm has the following accuracy guarantee.

Theorem 1.1 (Main Theorem). *Let \mathcal{P} be a distribution on \mathcal{X} and let M be our algorithm from Figure 1 instantiated with $T = n^2/k$ and $t = n\sqrt{2\ln(2k)/k}$. Suppose M is given a sample $S \sim \mathcal{P}^n$ and is asked adaptive statistical queries $\psi_1, \dots, \psi_k : \mathcal{X} \rightarrow [0, 1]$. Then M produces answers $v_1, \dots, v_k \in \mathbb{R}$ satisfying the following.*

$$\mathbf{E} \left[\max_{j=1}^k \frac{|v_j - \mathbf{E}_{X \sim \mathcal{P}}[\psi_j(X)]|}{\max\{\tau \cdot \text{sd}(\psi_j(\mathcal{P})), \tau^2\}} \right] \leq 4, \quad \text{where} \quad \tau = \sqrt{\frac{\sqrt{2k \ln(2k)}}{n}}.$$

Intuitively (that is, ignoring the second term in the maximum), the conclusion of Theorem 1.1 states that, with good probability, the error in each answer scales as the standard deviation of the query multiplied by $\tilde{O}\left(\sqrt{\sqrt{k}/n}\right)$ — which is what would be expected if we used n/\sqrt{k} fresh samples for each query. The $\ln k$ factor arises from the fact that we take a union bound over the k queries.

More precisely, applying Markov’s inequality to the conclusion of Theorem 1.1, shows that, with

probability at least 90%,

$$\forall j \quad \left| v_j - \mathbf{E}_{X \sim \mathcal{P}} [\psi_j(X)] \right| \leq 40 \cdot \max \left\{ \tau \cdot \text{sd}(\psi_j(\mathcal{P})), \tau^2 \right\} \leq \text{sd}(\psi_j(\mathcal{P})) \cdot 40 \sqrt{\frac{\sqrt{2k \ln(2k)}}{n}} + 40 \frac{\sqrt{2k \ln(2k)}}{n}. \quad (3)$$

This guarantee is directly comparable to the earlier bound (2) of Feldman and Steinke [FS17] – though it is weaker in two ways: First, Theorem 1.1 is a bound on the expectation and does not readily yield high probability bounds (other than via Markov’s inequality). Second, the second term in the maximum (which we think of as a low-order term) still depends linearly on the sensitivity and is potentially larger. The advantage of our algorithm is that it is substantially simpler than the earlier work.

Now we turn to the analysis of our algorithm. Clearly the *empirical error* of our algorithm — that is $|v_j - \mu_j|$ — scales with the empirical standard deviation σ_j . However, we must bound the *true error*, namely $|v_j - \mathbf{E}_{X \sim \mathcal{P}} [\psi_j(X)]|$. By the triangle inequality, it suffices to bound the *generalization error* $|\mu_j - \mathbf{E}_{X \sim \mathcal{P}} [\psi_j(X)]|$ in terms of standard deviation and to relate the empirical standard deviation σ_j to the true standard deviation $\text{sd}(\psi_j(\mathcal{P}))$.

1.1.1 Average KL leave-one-out stability and generalization

The key to our analysis is the following stability notion.

Definition 1.2 (Average Leave-one-out KL stability). *An algorithm $M : (\mathcal{X}^n \cup \mathcal{X}^{n-1}) \rightarrow \mathcal{Y}$ is ε -ALKL stable if, for all $s \in \mathcal{X}^n$,*

$$\frac{1}{n} \sum_{i \in [n]} D(M(s) \| M(s_{-i})) \leq \varepsilon,$$

where $s_{-i} \in \mathcal{X}^{n-1}$ denotes s with the i^{th} element removed. Here $D(\cdot \| \cdot)$ denotes the Kullback-Leibler divergence.

Our notion differs from differential privacy in three significant ways. First, we use stability to leaving one out (LOO) rather than replacing one element. Second, we average the stability parameter across the n dataset elements. Third, we use KL divergence instead of (approximate) max divergence. This is necessary to obtain stronger bounds for our calibrated noise addition as our algorithm does not satisfy differential privacy with parameters that would be suitable to ensure generalization. We note that average LOO stability is a well-studied way to define algorithmic stability for the loss function (e.g. [BE02, PRMN04]). The use of KL divergence appears to be necessary to ensure adaptive composition of our averaged notion. Specifically, the following composition result is easy to prove.

Lemma 1.3 (Composition). *Suppose $M : (\mathcal{X}^n \cup \mathcal{X}^{n-1}) \rightarrow \mathcal{Y}$ is ε -ALKL stable and $M' : \mathcal{Y} \times (\mathcal{X}^n \cup \mathcal{X}^{n-1}) \rightarrow \mathcal{Z}$ is such that $M'(y, \cdot) : (\mathcal{X}^n \cup \mathcal{X}^{n-1}) \rightarrow \mathcal{Z}$ is ε' -ALKL stable for all $y \in \mathcal{Y}$. Then the composition $s \mapsto M'(M(s), s)$ is $(\varepsilon + \varepsilon')$ -ALKL stable.*

Using composition, we can show that our algorithm (Figure 1, with the parameters set as in Theorem 1.1) is $\frac{kt}{n^2}$ -ALKL stable. In particular, we show that each one of the k answers is computed in a way that is $\frac{t}{n^2}$ -ALKL stable. This follows from the properties of the KL divergence between

Gaussian distributions and the way we calibrate the noise. (Alternatively, we could use Laplace noise to obtain similar results.)

We note that ε -differential privacy [DMNS06] (or $\frac{1}{2}\varepsilon^2$ -concentrated differential privacy [BS16]) implies $\frac{1}{2}\varepsilon^2$ -ALKL stability. Thus we can also compose any ALKL stable algorithm with any differentially private algorithm as well as any KL-stable algorithm [BNS⁺16].

Crucially, average KL-divergence is strong enough to provide a generalization guarantee that scales with the standard deviation of the queries, as we require. Our proof is based on the approach introduced by Dwork et al. [DFH⁺15]. They showed that $(\epsilon, 0)$ -differential privacy implies an upper bound on approximate max-information between the input dataset and output of the algorithm when run on independent samples and used the upper bound to prove generalization bounds. This approach was extended to (ϵ, δ) -differential privacy by Rogers et al. [RRST16]. Here, we demonstrate that ALKL stability implies a bound on the mutual information between the input and output of the algorithm when run on independent samples.¹

Proposition 1.4. *Let $M : (\mathcal{X}^n \cup \mathcal{X}^{n-1}) \rightarrow \mathcal{Q}$ be ε -ALKL stable. Let $S \in \mathcal{X}^n$ consist of n independent samples from some distribution \mathcal{P} . Then*

$$I(S; M(S)) \leq \varepsilon n,$$

where I denotes mutual information.

To prove Proposition 1.4, we introduce an intermediate notion of mutual information stability:

Definition 1.5 (Mutual Information Stability). *A randomized algorithm $M : \mathcal{X}^n \rightarrow \mathcal{Y}$ is ε -MI stable if, for any random variable S distributed over \mathcal{X}^n (including non-product distributions),*

$$\frac{1}{n} \sum_{i=1}^n I(M(S); S_i | S_{-i}) \leq \varepsilon.$$

This notion is based on the notion of stability studied by Raginsky et al. [RRT⁺16] that considers only product distributions over the datasets and, as a result, does not compose adaptively. We show that (i) average leave-one-out KL stability implies mutual information stability and (ii) mutual information stability implies a mutual information bound. Combining these facts will yield Proposition 1.4. Further, we point out that mutual information stability also composes adaptively in the same way as ALKL stability and hence could be useful for understanding adaptive data analysis for more general queries (*e.g.* unlike ALKL stability it does not require $M(S_{-i})$ to be defined).

As shown by Russo and Zou [RZ16], a bound on mutual information implies generalization results. Using a similar technique, we show that if $\mu_j = \frac{1}{n} \sum_{i=1}^n \psi_j(S_i)$ is subgaussian (or, more generally, subexponential) when ψ_j is independent from the sample S (consisting of n i.i.d. draws from \mathcal{P}) and the mutual information $I(S; \psi_j)$ is small, then $\mathbf{E}[\mu_j] \approx \mathbf{E}_{X \sim \mathcal{P}}[\psi_j(X)]$. Moreover, the quality of the approximation scales with the standard deviation (or, rather, the subgaussian constant). We can similarly relate the empirical variance σ_j^2 to the true variance. Thus a bound on mutual information suffices to bound generalization error and, thus, prove Theorem 1.1.

Another known implication of bounded mutual information is that any event that would happen with sufficiently low probability on fresh data will still happen with low probability [RZ16, RRST16]. In particular, if E is some “bad” event – such as overfitting the data or making a false discovery –

¹We thank Adam Smith for pointing to this approach of proving generalization via mutual information.

and we know that we are exponentially unlikely to overfit fresh data S' , then the probability of M overfitting its input data S is also small, provided the mutual information is small. (See Sec. 3.3 for additional details.)

One downside of using KL divergence is that does not allow us to prove high probability bounds, as can be done with differential privacy and the notion of approximate max-information [DFH⁺15]. We note, however, that our analysis still upper bounds the expectation of the largest error and. Using known techniques, the confidence can be amplified at the expense of a somewhat more complicated algorithm. In addition, our algorithm yields stronger stability guarantees than just ALKL stability. For example, the minimum noise level of $1/T$ ensures differential privacy (albeit with relatively large parameters²). The parameters can be improved using the averaging over the indices that we use in ALKL stability but that leads to a notion that does not appear to compose adaptively. Using a different analysis technique it might be possible to exploit the stronger stability properties of our algorithm to prove high probability generalization bounds. We leave this as an open problem. On the other hand, stability with KL divergence is easier to analyze and allows a potentially wider range of algorithms to be used.

2 Preliminaries

We use $X \sim \mathcal{P}$ to denote that X is drawn from the distribution \mathcal{P} . For the most part, we adopt the convention that upper-case letters denote random variables, whereas lower-case letters denote realizations thereof. For $n \in \mathbb{N}$, we denote $[n] = \{1, 2, \dots, n\}$ and $S \sim \mathcal{P}^n$ denotes that $S = (S_1, \dots, S_n)$ consists of n independent draws from the distribution \mathcal{P} . We use S_i to denote the i^{th} element of S and $S_{-i} = (S_1, \dots, S_{i-1}, S_{i+1}, \dots, S_n)$ to denote the other $n - 1$ elements. For two random variables X and Y and a realization x of X , we use the notation $Y|X = x$ to denote the conditional distribution of Y given $X = x$.

For a distribution \mathcal{P} on \mathcal{X} and a function $\psi : \mathcal{X} \rightarrow \mathbb{R}$, we use $\psi(\mathcal{P})$ to denote the distribution on \mathbb{R} obtained by applying ψ to a random sample from \mathcal{P} . The mean of this distribution is denoted $\mathcal{P}[\psi] = \mathbf{E}_{X \sim \mathcal{P}}[\psi(X)]$. We use $\text{sd}(\psi(\mathcal{P})) = \sqrt{\mathbf{Var}[\psi(\mathcal{P})]} = \sqrt{\mathbf{E}_{X \sim \mathcal{P}}[\psi(X)^2] - \mathbf{E}_{X \sim \mathcal{P}}[\psi(X)]^2}$ to denote the standard deviation of this distribution. We also interpret a tuple $s \in \mathcal{X}^n$ as a distribution — namely the distribution obtained by selecting s_i for a random $i \in [n]$ — and we analogously define the empirical mean and standard deviation: $s[\psi] = \frac{1}{n} \sum_{i=1}^n \psi[s_i]$ and $\text{sd}(\psi(s)) = \sqrt{\mathbf{Var}[\psi(s)]} = \sqrt{\frac{1}{n} \sum_{i=1}^n (\psi(s_i) - S[\psi])^2}$.

2.1 KL Divergence

Before continuing, we first establish some relevant properties of the KL divergence. See the textbook by Cover and Thomas [CT12] for an introduction to the properties of KL divergence (a.k.a. relative entropy).

First we state the definition of KL divergence for completeness.

Definition 2.1. *Let \mathcal{P} and \mathcal{Q} be probability distributions on a space Ω . Suppose \mathcal{P} is absolutely*

²Specifically, with the parameter setting from Theorem 1.1, our algorithm satisfies $\left(O\left(\sqrt{\log(1/\delta)}\right), \delta\right)$ -differential privacy for all $\delta > k^{-\Omega(k)}$.

continuous with respect to \mathcal{Q} . Then the KL divergence from \mathcal{Q} to \mathcal{P} is

$$D(\mathcal{P} \parallel \mathcal{Q}) = \mathbf{E}_{X \sim \mathcal{P}} \left[\ln \left(\frac{\mathcal{P}(X)}{\mathcal{Q}(X)} \right) \right],$$

where $\mathcal{P}(x)$ and $\mathcal{Q}(x)$ denote the probability mass or density functions of \mathcal{P} and \mathcal{Q} respectively evaluated at the point X . (More generally, $\mathcal{P}(x)/\mathcal{Q}(x)$ denotes the Radon-Nikodym derivative of \mathcal{P} with respect to \mathcal{Q} evaluated at x .)

In some cases we will abuse notation and refer to $D(X \parallel Y)$ where X and Y are “random variables” rather than formally-defined distributions. This should be read to be the divergence between the distribution of X and the distribution of Y .

We state the well-known chain rule:

Lemma 2.2 ([CT12, Theorem 2.5.3]). *Let \mathcal{P} and \mathcal{Q} be two distributions over some domain $\mathcal{X} \times \mathcal{Y}$. Then*

$$D(\mathcal{P}(x, y) \parallel \mathcal{Q}(x, y)) = D(\mathcal{P}(x) \parallel \mathcal{Q}(x)) + \mathbf{E}_{x' \sim \mathcal{P}} [D(\mathcal{P}(y|x = x') \parallel \mathcal{Q}(y|x = x'))].$$

Here $\mathcal{P}(x)$ (or $\mathcal{Q}(x)$) denotes the marginal distributions of \mathcal{P} (or \mathcal{Q}) over \mathcal{X} and $\mathcal{P}(y|x = x')$ denotes the marginal distribution of \mathcal{P} on \mathcal{Y} conditioned on $x = x'$.

We begin by looking at the KL divergence between two Gaussian distributions, as this is what our mechanism uses. Recall that the Gaussian (or normal) distribution with mean μ and variance σ^2 — denoted $\mathcal{N}(\mu, \sigma^2)$ — has a probability density at x given by $\frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right)$.

Lemma 2.3 ([GAL13, Table 3]). *Let $\mu, \tilde{\mu}, \sigma, \tilde{\sigma} \in \mathbb{R}$. Then*

$$D(\mathcal{N}(\mu, \sigma^2) \parallel \mathcal{N}(\tilde{\mu}, \tilde{\sigma}^2)) = \frac{(\mu - \tilde{\mu})^2}{2\tilde{\sigma}^2} + \frac{1}{2} \left(\frac{\sigma^2}{\tilde{\sigma}^2} - 1 - \ln \left(\frac{\sigma^2}{\tilde{\sigma}^2} \right) \right).$$

Corollary 2.4. *Let $\mu, \tilde{\mu}, \sigma, \tilde{\sigma}, x \in \mathbb{R}$. If $\sigma\tilde{\sigma} \neq 0$, then*

$$D(\mathcal{N}(\mu, \sigma^2) \parallel \mathcal{N}(\tilde{\mu}, \tilde{\sigma}^2)) \leq \frac{1}{2} \cdot \left(\frac{(\mu - \tilde{\mu})^2}{\sigma^2} + \left(\frac{\tilde{\sigma}^2}{\sigma^2} - 1 \right)^2 \cdot \min \left\{ 1, \frac{1}{6} \left(2 + \frac{\sigma^2}{\tilde{\sigma}^2} \right) \right\} \right) \cdot \frac{\sigma^2}{\tilde{\sigma}^2}.$$

Proof. This follows from Lemma 2.3 and the inequalities $x - 1 - \ln x \leq (1/x - 1)^2 \cdot x$ and $x - 1 - \ln x \leq (1/x - 1)^2 \cdot x \cdot (2 + x)/6$ for all $x > 0$. □

An analogous result holds for the Laplace distribution. Although we do not work this out, it implies that our results can be extended to work for the Laplace distribution (with slightly different constants and a higher power of $\ln k$, since the Laplace distribution has heavier tails). Recall that the Laplace distribution with mean μ and variance $2\sigma^2$ — denoted $\mu + \text{Lap}(\sigma)$ — has a probability density at x given by $\frac{1}{2|\sigma|} \exp\left(\left|\frac{x-\mu}{\sigma}\right|\right)$.

Lemma 2.5 ([GAL13, Table 3]). *Let $\mu, \tilde{\mu}, \sigma, \tilde{\sigma} \in \mathbb{R}$. If $\sigma, \tilde{\sigma} > 0$, then*

$$\begin{aligned} D(\mu + \text{Lap}(\sigma) \parallel \tilde{\mu} + \text{Lap}(\tilde{\sigma})) &= \frac{\sigma}{\tilde{\sigma}} \left(e^{-\frac{|\tilde{\mu}-\mu|}{\sigma}} - \left(1 - \frac{|\tilde{\mu}-\mu|}{\sigma} \right) \right) + \frac{\sigma}{\tilde{\sigma}} - 1 - \ln \left(\frac{\sigma}{\tilde{\sigma}} \right) \\ &\leq \frac{(\tilde{\mu} - \mu)^2}{2\sigma\tilde{\sigma}} + \frac{1}{7} \left(\frac{\tilde{\sigma}^2}{\sigma^2} - 1 \right)^2 \frac{\sigma^2}{\tilde{\sigma}^2}. \end{aligned}$$

Next we have a technical lemma relating expectations to KL divergence.

Lemma 2.6 ([Gra11, Theorem 5.2.1]). *Let \mathcal{P} and \mathcal{Q} be probability distributions on Ω . Then*

$$D(\mathcal{P}\|\mathcal{Q}) = \sup_{f:\Omega\rightarrow\mathbb{R}} \mathbf{E}_{X\sim\mathcal{P}} [f(X)] - \ln \mathbf{E}_{X\sim\mathcal{Q}} [e^{f(X)}].$$

Setting $f(x) = tx$ and rearranging gives the bound we will use:

Corollary 2.7. *Let X and Y be real-valued random variables and $t > 0$. Then*

$$\mathbf{E}[X] \leq \frac{1}{t} \left(D(X\|Y) + \ln \mathbf{E} [e^{tY}] \right).$$

Next we note that KL divergence is a convex function.

Lemma 2.8 ([VEH14, Theorem 11]). *Let $\mathcal{P}_0, \mathcal{P}_1, \mathcal{Q}_0, \mathcal{Q}_1$ be probability distributions on the same space Ω . For $t \in (0, 1)$, let $\mathcal{P}_t = (1-t)\mathcal{P}_0 + t\mathcal{P}_1$ and $\mathcal{Q}_t = (1-t)\mathcal{Q}_0 + t\mathcal{Q}_1$ be the convex combinations interpolating between these distributions. Then, for all $t \in [0, 1]$,*

$$D(\mathcal{P}_t\|\mathcal{Q}_t) \leq (1-t)D(\mathcal{P}_0\|\mathcal{Q}_0) + tD(\mathcal{P}_1\|\mathcal{Q}_1).$$

This lemma immediately extends to convex combinations of more than two distributions.

Next we have a geometric statement about KL divergence:

Lemma 2.9 ([BMDG05, Proposition 1],[FSG08, Theorem II.1]). *Let $\{\mathcal{P}_y\}$ be a family of distributions indexed by $y \in \mathcal{Y}$ and let \mathcal{Q} be a distribution on \mathcal{Y} . Let $\mathcal{P}_{\mathcal{Q}} = \mathbf{E}_{Y\sim\mathcal{Q}} [\mathcal{P}_Y]$ denote the convex combination of the distributions $\{\mathcal{P}_y\}$ weighted by \mathcal{Q} . Then*

$$\inf_{\mathcal{R}} \mathbf{E}_{Y\sim\mathcal{Q}} [D(\mathcal{P}_Y\|\mathcal{R})] = \mathbf{E}_{Y\sim\mathcal{Q}} [D(\mathcal{P}_Y\|\mathcal{P}_{\mathcal{Q}})].$$

Lemma 2.9 shows that the “center” of a collection of probability distributions — as measured by minimizing average KL divergence to one distribution — is none other than the mean of those distributions.

2.2 Mutual Information

A key quantity that we use is mutual information:

Definition 2.10 (Mutual Information). *For two random variables X and Y jointly distributed according to a distribution \mathcal{P} over $\mathcal{X} \times \mathcal{Y}$, the mutual information between X and Y is*

$$I(X; Y) = D(\mathcal{P}(x, y)\|\mathcal{P}(x) \times \mathcal{P}(y)) = \mathbf{E}_{x'\sim\mathcal{P}(x)} [D(\mathcal{P}(y|x = x')\|\mathcal{P}(y))],$$

where $\mathcal{P}(x) \times \mathcal{P}(y)$ denotes the product of the marginal distributions of \mathcal{P} .

Note that mutual information is symmetric – $I(X; Y) = I(Y; X)$.

Definition 2.11 (Conditional Mutual Information). *For three random variables X, Y , and Z . The mutual information between X and Y conditioned on Z is given by*

$$I(X; Y|Z) = \mathbf{E}_{z\sim\mathcal{P}_Z} [I(X|Z = z; Y|Z = z)],$$

where \mathcal{P}_Z is the marginal distribution of Z .

The key property is the chain rule:

Lemma 2.12 (Mutual Information Chain Rule). *For random variables X, Y , and Z , we have*

$$I(X, Y; Z) = I(X; Z) + I(Y; Z|X)$$

3 Average KL Stability & Generalization

In this section, we cover our theoretical tools, which center around our definition of average leave-one-out KL stability, which we restate here. For a randomized algorithm M and input s we use $M(s)$ denote the random variable obtained by running on M on s and the distribution of this random variable (according to the context).

Definition 3.1 (Average Leave-one-out KL stability). *An algorithm $M : (\mathcal{X}^n \cup \mathcal{X}^{n-1}) \rightarrow \mathcal{Y}$ is ε -ALKL stable if, for all $s \in \mathcal{X}^n$,*

$$\frac{1}{n} \sum_{i \in [n]} \mathbb{D}(M(s) \| M(s_{-i})) \leq \varepsilon,$$

where $s_{-i} \in \mathcal{X}^{n-1}$ denotes s with the i^{th} element removed.

An algorithm $M : \mathcal{X}^n \rightarrow \mathcal{Y}$ is ε -ALKL stable if there exists $\tilde{M} : (\mathcal{X}^n \cup \mathcal{X}^{n-1}) \rightarrow \mathcal{Y}$ that is ε -ALKL stable and $M(s) = \tilde{M}(s)$ holds for all $s \in \mathcal{X}^n$.

Note that the second half of the definition extends it to mechanisms that are only defined for inputs in \mathcal{X}^n , but not \mathcal{X}^{n-1} .

The key property of our definition is composition. This lemma allows us to account for the accumulation of information through multiple adaptive queries. The following lemma only considers the composition of two algorithms. Induction allows this to be extended to k algorithms.

Lemma 3.2 (Composition). *Suppose $M : (\mathcal{X}^n \cup \mathcal{X}^{n-1}) \rightarrow \mathcal{Y}$ is ε -ALKL stable and $M' : \mathcal{Y} \times (\mathcal{X}^n \cup \mathcal{X}^{n-1}) \rightarrow \mathcal{Z}$ is such that $M'(y, \cdot) : (\mathcal{X}^n \cup \mathcal{X}^{n-1}) \rightarrow \mathcal{Z}$ is ε' -ALKL stable for all $y \in \mathcal{Y}$. Then the composition $s \mapsto M'(M(s), s)$ is $(\varepsilon + \varepsilon')$ -ALKL stable.*

Proof. Fix $s \in \mathcal{X}^n$. By the chain rule for KL divergence (Lemma 2.2),

$$\begin{aligned} & \frac{1}{n} \sum_{i \in [n]} \mathbb{D}(M'(M(s), s) \| M'(M(s_{-i}), s_{-i})) \\ & \leq \frac{1}{n} \sum_{i \in [n]} \left(\mathbb{D}(M(s) \| M(s_{-i})) + \mathbf{E}_{y \sim M(s)} [\mathbb{D}(M'(y, s) \| M'(y, s_{-i}))] \right) \\ & = \frac{1}{n} \sum_{i \in [n]} \mathbb{D}(M(s) \| M(s_{-i})) + \mathbf{E}_{y \sim M(s)} \left[\frac{1}{n} \sum_{i \in [n]} \mathbb{D}(M'(y, s) \| M'(y, s_{-i})) \right] \\ & \leq \varepsilon + \varepsilon'. \end{aligned}$$

□

Another key property of our definition of average leave-one-out KL stability is *postprocessing*. That is, if M is ε -ALKL stable, then applying an arbitrary function to the output of M continues to be ε -ALKL stable. This can be seen by taking $\varepsilon' = 0$ in the above composition lemma or by using the data processing inequality for KL divergence [VEH14, Theorem 1].

3.1 Mutual Information

In order to show that our notion of average leave-one-out KL stability implies generalization, we first show that it implies a bound on mutual information:

Proposition 3.3. *Let $M : \mathcal{X}^n \rightarrow \mathcal{Y}$ be ε -ALKL stable. Let $S \in \mathcal{X}^n$ be a product distribution. Then $I(M(S); S) \leq \varepsilon \cdot n$.*

To prove Proposition 3.3, we introduce an intermediate notion of stability that is based on that of Raginsky et al. [RRT⁺16]. Specifically, mutual information stability is defined as follows.

Definition 3.4 (Restating Definition 1.5). *A randomized algorithm $M : \mathcal{X}^n \rightarrow \mathcal{Y}$ is ε -MI stable if, for any random variable S distributed over \mathcal{X}^n (including non-product distributions),*

$$\frac{1}{n} \sum_{i=1}^n I(M(S); S_i | S_{-i}) \leq \varepsilon.$$

We show that mutual information stability has the following properties.

1. Average leave-one-out KL stability implies mutual information stability.
2. Mutual information stability implies a mutual information bound.
3. Mutual information stability composes adaptively.

Combining properties (1) and (2) yields Proposition 3.3. The adaptive composition property of mutual information stability implies that it might be useful for analysis of adaptive procedures which are not ALKL stable (although we do not use this property since ALKL stability itself composes adaptively).

Lemma 3.5. *If $M : \mathcal{X}^n \rightarrow \mathcal{Y}$ is ε -ALKL stable, then it also is ε -MI stable.*

Proof. Let S be a random variable distributed according to some distribution \mathcal{P} on \mathcal{X}^n . Let \mathcal{P}_i and \mathcal{P}_{-i} denote the marginal distribution of S_i and S_{-i} , respectively. For $z \in \mathcal{X}^{n-1}$ we use $\mathcal{P}(s_i|z)$ to denote the distribution of S_i conditioned on $S_{-i} = z$. Now, by the definition of (conditional) mutual information,

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n I(M(S); S_i | S_{-i}) &= \frac{1}{n} \mathbf{E}_{z \sim \mathcal{P}_{-i}} [I(M(S) | S_{-i} = z; S_i | S_{-i} = z)] \\ &= \frac{1}{n} \mathbf{E}_{z \sim \mathcal{P}_{-i}} \left[\mathbf{E}_{x \sim \mathcal{P}(s_i|z)} [\mathbf{D}(M(S) | S_{-i} = z, S_i = x \| M(S) | S_{-i} = z)] \right] \\ &= \frac{1}{n} \mathbf{E}_{z \sim \mathcal{P}_{-i}} \left[\mathbf{E}_{x \sim \mathcal{P}(s_i|z)} [\mathbf{D}(M(z \circ_i x) \| M(S) | S_{-i} = z)] \right]. \end{aligned}$$

Here $z \circ_i x$ refers to the vector $s \in \mathcal{X}^n$ such that $s_{-i} = z$ and $s_i = x$. Here the inner expectation is over x drawn from the distribution of S_i conditioned on $S_{-i} = z$ — of the KL divergence from the distribution of $M(z \circ_i x)$ to the distribution of $M(S)$ conditioned on $S_{-i} = z$. The latter distribution is exactly the convex combination of the distribution of $M(z \circ_i x)$ weighted by the distribution of $x \sim \mathcal{P}(s_i|z)$.

Now the key observation: the convex combination $M(S)|_{S_{-i} = z}$ is the distribution that minimizes the inner expectation. Hence, we can replace it by $M(z)$ and only increase the expression. Formally, by Lemma 2.9, for all $z \in \mathcal{X}^{n-1}$,

$$\mathbf{E}_{x \sim \mathcal{P}(s_i|z)} [\mathbf{D}(M(z \circ_i x) \| M(S)|_{S_{-i} = z})] \leq \mathbf{E}_{x \sim \mathcal{P}(s_i|z)} [\mathbf{D}(M(z \circ_i x) \| M(z))]$$

Note that $M(z)$ refers either to execution of M itself or (if M is not defined over inputs of length $n-1$) to M 's extension to \mathcal{X}^{n-1} promised by the second half of Definition 3.1. The result now follows, as we have established that

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n I(M(S); S_i | S_{-i}) &\leq \frac{1}{n} \mathbf{E}_{z \sim \mathcal{P}_{-i}} \left[\mathbf{E}_{x \sim \mathcal{P}(s_i|z)} [\mathbf{D}(M(z \circ_i x) \| M(z))] \right] \\ &= \frac{1}{n} \sum_{i=1}^n \mathbf{E}_{s \sim \mathcal{P}} [\mathbf{D}(M(s) \| M(s_{-i}))]. \end{aligned}$$

□

Lemma 3.6. *Suppose $M : \mathcal{X}^n \rightarrow \mathcal{Y}$ is ε -MI stable. Let \mathcal{P} be a distribution over \mathcal{X} and S be distributed according to \mathcal{P}^n . Then $I(M(S); S) \leq \varepsilon n$.*

Proof. Denote $S_{<i} = (S_1, \dots, S_{i-1})$, $S_{>i} = (S_{i+1}, \dots, S_n)$ and $S_{\leq i} = (S_1, \dots, S_i)$. By the chain rule for mutual information (Lemma 2.12 and induction),

$$I(M(S); S) = \sum_{i=1}^n I(M(S); S_i | S_{<i}).$$

By the definition of (conditional) mutual information,

$$\begin{aligned} I(M(S); S_i | S_{<i}) &= \mathbf{E}_{z \sim \mathcal{P}^{i-1}} [I(M(S) | S_{<i} = z; S_i | S_{<i} = z)] \\ &= \mathbf{E}_{z \sim \mathcal{P}^{i-1}} \left[\mathbf{E}_{x \sim \mathcal{P}} [\mathbf{D}(M(S) | S_{\leq i} = z \circ x \| M(S) | S_{<i} = z)] \right] \end{aligned} \quad (4)$$

Here $z \circ x \in \mathcal{X}^i$ denotes the concatenation of $z \in \mathcal{X}^{i-1}$ with $x \in \mathcal{X}$. Now, by the convexity of KL divergence (Lemma 2.8), we can move the randomness of $S_{>i}$ from the divergence and into the expectation. Namely,

$$\mathbf{D}(M(S) | S_{\leq i} = z \circ x \| M(S) | S_{<i} = z) \leq \mathbf{E}_{z' \sim \mathcal{P}^{n-i}} [\mathbf{D}(M(S) | S = z \circ x \circ z' \| M(S) | S_{-i} = z \circ z')].$$

Here we use the fact that $M(S) | S_{\leq i} = z \circ x$ and $M(S) | S_{<i} = z$ are convex combinations of the distribution of $M(s)$ weighted by $S_{>i}$ (note that independence is crucial here). Plugging this into eq. (4) and using the definition of (conditional) mutual information we get

$$\begin{aligned} I(M(S); S_i | S_{<i}) &\leq \mathbf{E}_{z \sim \mathcal{P}^{i-1}} \left[\mathbf{E}_{x \sim \mathcal{P}} \left[\mathbf{E}_{z' \sim \mathcal{P}^{n-i}} [\mathbf{D}(M(S) | S = z \circ x \circ z' \| M(S) | S_{-i} = z \circ z')] \right] \right] \\ &= \mathbf{E}_{s \sim \mathcal{P}} [\mathbf{D}(M(S) | S = s \| M(S) | S_{-i} = s_{-i})] \\ &= I(M(S); S_i | S_{-i}) \end{aligned}$$

Combining these (in)equalities yields the result:

$$I(M(S); S) \leq \sum_{i=1}^n I(M(S); S_i | S_{-i}).$$

□

Lemma 3.7. *Suppose $M : \mathcal{X}^n \rightarrow \mathcal{Y}$ is ε -MI stable and $M' : \mathcal{Y} \times \mathcal{X}^n \rightarrow \mathcal{Z}$ is such that $M'(y, \cdot) : \mathcal{X}^n \rightarrow \mathcal{Z}$ is ε' -MI stable for all $y \in \mathcal{Y}$. Then the composition $s \mapsto M'(M(s), s)$ is $(\varepsilon + \varepsilon')$ -MI stable.*

Proof. Let S be a random variable on \mathcal{X}^n and let $\mathcal{P}_{M(S)}$ denote the probability distribution of $M(S)$. By the chain rule,

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n I(M'(M(S), S); S_i | S_{-i}) &\leq \frac{1}{n} \sum_{i=1}^n I(M(S); S_i | S_{-i}) + I(M'(M(S), S); S_i | S_{-i}, M(S)) \\ &= \frac{1}{n} \sum_{i=1}^n I(M(S); S_i | S_{-i}) + \mathbf{E}_{y \sim \mathcal{P}_{M(S)}} [I(M'(y, S); S_i | S_{-i}, M(S) = y)] \\ &= \frac{1}{n} \sum_{i=1}^n I(M(S); S_i | S_{-i}) + \mathbf{E}_{y \sim \mathcal{P}_{M(S)}} \left[\frac{1}{n} \sum_{i=1}^n I(M'(y, S); S_i | S_{-i}, M(S) = y) \right] \\ &\leq \varepsilon + \varepsilon'. \end{aligned}$$

The key is that the stability property holds for all distributions, which means it holds for the distribution of S conditioned on $M(S)$. Note that if we defined mutual information stability only to quantify over product distributions, then this proof would not carry through, as S conditioned on $M(S)$ is not necessarily a product distribution anymore. □

3.2 Generalization in expectation

In this section we translate an upper bound on $I(S; M(S))$ into an upper bound on the expectation of the generalization error. As in earlier work [RZ16], our main technical tool is Corollary 2.7. However we deal with more general random variables (not just subgaussian) and also prove bounds that are scaled to standard deviation of the random variable as opposed to the subgaussian constant. In Section 3.3 we describe an alternative approach to generalization which is based on bounding the probability of any “bad” event.

The following proposition bounds the expected generalization error.

Proposition 3.8. *Let M be a randomized algorithm with input from \mathcal{X}^n and output in \mathcal{Q} , where \mathcal{Q} is the set of functions $\psi : \mathcal{X} \rightarrow [0, 1]$. Let \mathcal{P} be a distribution on \mathcal{X} and $S \sim \mathcal{P}^n$. Let $\tau > 0$. Suppose $I(S; M(S)) \leq \varepsilon n$. Then*

$$\mathbf{E}_{\substack{S \sim \mathcal{P}^n \\ \psi \sim M(S)}} \left[\frac{S[\psi] - \mathcal{P}[\psi]}{\max\{\text{sd}(\psi(\mathcal{P})), \tau\}} \right] \leq \left\{ \begin{array}{ll} 2\sqrt{\varepsilon} & \text{if } \sqrt{\varepsilon} \leq \tau \\ \varepsilon/\tau + \tau & \text{if } \sqrt{\varepsilon} \geq \tau \end{array} \right\} \leq 2\sqrt{\varepsilon} + \varepsilon/\tau.$$

Proof. Define a random variable $X = \frac{S[\psi] - \mathcal{P}[\psi]}{\max\{\text{sd}(\psi(\mathcal{P})), \tau\}}$ for $S \sim \mathcal{P}^n$ and $\psi \sim M(S)$. Our goal is to bound $\mathbf{E}[X]$. Let $Y = \frac{S[\psi] - \mathcal{P}[\psi]}{\max\{\text{sd}(\psi(\mathcal{P})), \tau\}}$ for $(S, S') \sim \mathcal{P}^n \times \mathcal{P}^n$ and $\psi \sim M(S')$. That is, Y is X altered so that the query ψ is independent of the data S , but has the same marginal distribution.

Since $I(S; M(S)) = D(S, M(S) \| S, M(S')) \leq \varepsilon n$, we have $D(X \| Y) \leq \varepsilon n$ by the data processing inequality. By Corollary 2.7,

$$\mathbf{E}[X] \leq \inf_{\lambda > 0} \frac{1}{\lambda} \left(D(X \| Y) + \ln \mathbf{E} \left[e^{\lambda Y} \right] \right). \quad (5)$$

Thus it only remains to bound $\mathbf{E} \left[e^{\lambda Y} \right]$. We have

$$\begin{aligned} \mathbf{E} \left[e^{\lambda Y} \right] &= \mathbf{E}_{\substack{(S, S') \sim \mathcal{P}^n \times \mathcal{P}^n \\ \psi \sim M(S')}} \left[\exp \left(\frac{\lambda}{n} \sum_{i=1}^n \frac{\psi(S_i) - \mathcal{P}[\psi]}{\max \{ \text{sd}(\psi(\mathcal{P})), \tau \}} \right) \right] \\ &= \mathbf{E}_{\substack{S' \sim \mathcal{P}^n \\ \psi \sim M(S')}} \left[\prod_{i=1}^n \mathbf{E}_{S_i \sim \mathcal{P}} \left[\exp \left(\frac{\lambda}{n} \frac{\psi(S_i) - \mathcal{P}[\psi]}{\max \{ \text{sd}(\psi(\mathcal{P})), \tau \}} \right) \right] \right]. \end{aligned}$$

Thus it suffices to bound $\mathbf{E}_{S_i \sim \mathcal{P}} \left[\exp \left(\frac{\lambda}{n} \frac{\psi(S_i) - \mathcal{P}[\psi]}{\max \{ \text{sd}(\psi(\mathcal{P})), \tau \}} \right) \right]$ for a fixed i and a fixed ψ . The random variable $Y_i = \frac{\psi(S_i) - \mathcal{P}[\psi]}{\max \{ \text{sd}(\psi(\mathcal{P})), \tau \}}$ has mean 0 and variance at most 1 (since $S_i \sim \mathcal{P}$). Also $|Y_i| \leq 1/\tau$. Thus by Lemma 3.9 (stated below), we have $\mathbf{E} \left[e^{\frac{\lambda}{n} Y_i} \right] \leq e^{\frac{\lambda^2}{n^2}}$ for all $\lambda/n \leq \tau$. Hence $\mathbf{E} \left[e^{\lambda Y} \right] \leq e^{\lambda^2/n}$ for all $\lambda \leq n\tau$. Plugging this into eq. (5), we get

$$\mathbf{E}[X] \leq \inf_{0 < \lambda \leq n\tau} \frac{\varepsilon n}{\lambda} + \frac{\lambda}{n} = \begin{cases} 2\sqrt{\varepsilon} & \text{if } \sqrt{\varepsilon} \leq \tau \\ \varepsilon/\tau + \tau & \text{if } \sqrt{\varepsilon} \geq \tau \end{cases}.$$

□

Lemma 3.9. *Let Y be a random variable supported on $[-1/\tau, 1/\tau]$. Suppose $\mathbf{E}[Y] = 0$ and $\mathbf{E}[Y^2] \leq 1$. Then for $\lambda \in [0, \tau]$, $\mathbf{E} \left[e^{\lambda Y} \right] \leq e^{\lambda^2}$.*

This lemma is similar to the proof of Bernstein's inequality [Ber24].

Proof. Since $|Y| \leq 1/\tau$, we have $|Y|^k \leq (1/\tau)^{k-2} Y^2$ for all $k \geq 2$. Thus, for all $\lambda \geq 0$, we have

$$\begin{aligned} \mathbf{E} \left[e^{\lambda Y} \right] &= 1 + \lambda \mathbf{E}[Y] + \sum_{k=2}^{\infty} \frac{\lambda^k}{k!} \mathbf{E} \left[Y^k \right] \leq 1 + \sum_{k=2}^{\infty} \frac{\lambda^k}{k!} (1/\tau)^{k-2} \mathbf{E} \left[Y^2 \right] \\ &= 1 + \left(e^{\lambda/\tau} - 1 - \lambda/\tau \right) \tau^2 \mathbf{E} \left[Y^2 \right] \leq e^{(e^{\lambda/\tau} - 1 - \lambda/\tau) \cdot \tau^2}. \end{aligned}$$

If $\lambda/\tau \leq 1$, then $e^{\lambda/\tau} - 1 - \lambda/\tau \leq (\lambda/\tau)^2$, which yields the result. □

To analyse our algorithm, we also need to bound the empirical error in terms of the standard deviation. Note that the empirical error – the noise we add – scales with the empirical standard deviation. Thus we must bound the empirical variance in terms of the true variance:

Proposition 3.10. *Let M be a randomized algorithm with input from \mathcal{X}^n and output in \mathcal{Q} , where \mathcal{Q} is the set of functions $\psi : \mathcal{X} \rightarrow [0, 1]$. Let \mathcal{P} be a distribution on \mathcal{X} and $S \sim \mathcal{P}^n$. Let $\tau > 0$. Suppose $I(S; M(S)) \leq \varepsilon n$. Then*

$$\mathbf{E}_{\substack{S \sim \mathcal{P} \\ \psi \sim M(S)}} \left[\left(\frac{\text{sd}(\psi(S))}{\max \{ \text{sd}(\psi(\mathcal{P})), \tau \}} \right)^2 \right] = \mathbf{E}_{\substack{S \sim \mathcal{P} \\ \psi \sim M(S)}} \left[\left(\frac{\sqrt{\frac{1}{n} \sum_{i=1}^n (\psi(S_i) - S[\psi])^2}}{\max \{ \text{sd}(\psi(\mathcal{P})), \tau \}} \right)^2 \right] \leq 2 + \varepsilon/\tau^2.$$

To prove Proposition 3.10 we make use of the following two standard facts.

Let $\psi : \mathcal{X} \rightarrow \mathbb{R}$. Let \mathcal{P} be a distribution on \mathcal{X} and $S \in \mathcal{X}^n$. Then

$$\frac{1}{n} \sum_{i=1}^n (\psi(S_i) - S[\psi])^2 \leq \frac{1}{n} \sum_{i=1}^n (\psi(S_i) - \mathcal{P}[\psi])^2. \quad (6)$$

Let X be a random variable supported on $[0, 1]$. Suppose $\mathbf{E}[X] = \sigma$. Then for $s \in [0, 1]$,

$$\mathbf{E} \left[e^{sX} \right] \leq \mathbf{E} [1 + 2sX] = 1 + 2\sigma s \leq e^{2\sigma s}. \quad (7)$$

Proof of Proposition 3.10. By (6),

$$\mathbf{E}_{\substack{S \sim \mathcal{P} \\ \psi \sim M(S)}} \left[\left(\frac{\sqrt{\frac{1}{n} \sum_{i=1}^n (\psi(S_i) - S[\psi])^2}}{\max\{\text{sd}(\psi(\mathcal{P})), \tau\}} \right)^2 \right] \leq \mathbf{E}_{\substack{S \sim \mathcal{P} \\ \psi \sim M(S)}} \left[\frac{\frac{1}{n} \sum_{i=1}^n (\psi(S_i) - \mathcal{P}[\psi])^2}{\max\{\mathbf{Var}[\psi(\mathcal{P})], \tau^2\}} \right].$$

Define a random variable $X = \sum_{i=1}^n \frac{(\psi(S_i) - \mathcal{P}[\psi])^2}{\max\{\mathbf{Var}[\psi(\mathcal{P})], \tau^2\}}$ for $S \sim \mathcal{P}^n$ and $\psi \sim M(S)$. Our goal is thus to bound $\frac{1}{n} \mathbf{E}[X]$. Define another random variable $Y = \sum_{i=1}^n \frac{(\psi(S_i) - \mathcal{P}[\psi])^2}{\max\{\mathbf{Var}[\psi(\mathcal{P})], \tau^2\}}$ for $(S, S') \sim \mathcal{P}^n \times \mathcal{P}^n$ and $\psi \sim M(S')$. That is, Y is defined for S and ψ being independent, whereas X has them being dependent through M . By the data processing inequality, $D(X||Y) \leq D(S, M(S)||S, M(S')) = I(S; M(S)) \leq \varepsilon n$. Now, by Corollary 2.7,

$$\mathbf{E}[X] \leq \inf_{\lambda > 0} \frac{1}{\lambda} \left(D(X||Y) + \ln \mathbf{E} \left[e^{\lambda Y} \right] \right). \quad (8)$$

To bound $\mathbf{E} \left[e^{\lambda Y} \right]$ we note that Y is determined by S and ψ . Since these are independent, we may consider an arbitrary fixed ψ . We let Z denote Y conditioned on ψ being equal to a fixed $\phi \in \mathcal{Q}$. We can write Z as a sum of n independent terms $Z_i = \frac{(\phi(S_i) - \mathcal{P}[\phi])^2}{\max\{\mathbf{Var}[\phi(\mathcal{P})], \tau^2\}}$, and hence $\mathbf{E} \left[e^{\lambda Z} \right] = \prod_i^n \mathbf{E} \left[e^{\lambda Z_i} \right]$. For each i , we have $\mathbf{E}[Z_i] = \frac{\mathbf{Var}[\phi(\mathcal{P})]}{\max\{\mathbf{Var}[\phi(\mathcal{P})], \tau^2\}} \leq 1$ and $0 \leq Z_i \leq 1/\tau^2$. Thus, by (7) (with $X = Z_i \tau^2$, $\sigma \leq \tau^2$, and $s = \lambda/\tau^2$), $\mathbf{E} \left[e^{\lambda Z_i} \right] \leq e^{2\lambda}$ for $\lambda \in [0, \tau^2]$.

This implies that $\mathbf{E} \left[e^{\lambda Z} \right] \leq e^{2\lambda n}$ for $\lambda \in [0, \tau^2]$ for every ϕ and hence $\mathbf{E} \left[e^{\lambda Y} \right] \leq e^{2\lambda n}$ under the same condition. Substituting this into eq. (8) yields

$$\mathbf{E}[X] \leq \inf_{0 < \lambda \leq \tau^2} \frac{1}{\lambda} \left(\varepsilon n + \ln \left(e^{2\lambda n} \right) \right) = \varepsilon n / \tau^2 + 2n.$$

□

3.3 Preservation of low-probability events

Propositions 3.8 and 3.10 bound the expected generalization error given a bound on mutual information. An alternative approach to analysis of generalization is to use a bound on mutual information to upper bound the increase in probability of any “bad” event that results from the dependence between the dataset and algorithm’s output. Specifically, we prove the following simple lemma:

Lemma 3.11. *Let S consist of n independent samples from some distribution \mathcal{P} . Let S' be an independent copy of S . Let $M : \mathcal{X}^n \rightarrow \mathcal{Y}$ and let E be an event on $\mathcal{X}^n \times \mathcal{Y}$ satisfying*

$$\Pr [(S', M(S)) \in E] \leq \delta.$$

Then

$$\Pr [(S, M(S)) \in E] \leq \frac{I(S; M(S)) + \ln 2}{\ln(1/\delta)}.$$

Intuitively, Lemma 3.11 says that if an event happens with very low probability on fresh data, then it happens with somewhat low probability on non-fresh data, as long as the mutual information between the event and the data is low. Note however, that the probability grows from δ to $\frac{I(S; M(S)) + \ln 2}{\ln(1/\delta)}$. In particular, the inverse of the probability decreases exponentially. For example, Lemma 3.11 can be used to correct a p -value obtained under the assumption that the data is independent from the choice of the test (since p -value is the probability that a test statistic satisfies a chosen condition) [RZ16, RRST16].

The same approach to generalization is used in [DFH⁺14, DFH⁺15, RRST16] for differential privacy and max-information and in [RZ16, RRST16] for mutual information. The bound implicit in [RZ16] is $\Pr [(S, M(S)) \in E] \leq \delta + \sqrt{\frac{I(S; M(S))}{\ln(1/(2\delta))}}$ which is asymptotically worse than our bound. The bound in [RRST16] is derived by first using mutual information to bound approximate max-information [DFH⁺15]. Their approach yields the following bound

$$\Pr [(S, M(S)) \in E] \leq \inf_{k \geq 0} \left(2^k \cdot \delta + \frac{I(S; M(S)) + 0.54}{k} \right)$$

which is comparable to the bound in Lemma 3.11.

As a more concrete application, we demonstrate how Lemma 3.11 can be used to derive a bound on the probability of generalization error being large (or, equivalently, to construct a valid confidence interval for the true expectation of a real-valued function).

Recall the setting of Proposition 3.8. Here S consists of n independent samples from \mathcal{P} and M outputs a function $\psi : \mathcal{X} \rightarrow [0, 1]$ and has $I(S; M(S)) \leq \varepsilon n$. By Bernstein's inequality, for n fresh samples S' (independent from $M(S)$), we have

$$\Pr_{\substack{(S, S') \sim \mathcal{P}^n \times \mathcal{P}^n \\ \psi \sim M(S)}} \left[\frac{S'[\psi] - \mathcal{P}[\psi]}{\max\{\text{sd}(\psi(\mathcal{P}), \tau)\}} > t \right] \leq \exp\left(\frac{-t^2 n}{2 + \frac{2}{3} \frac{t}{\tau}}\right)$$

for all $t > 0$. Thus, by Lemma 3.11, for all $t > 0$,

$$\Pr_{\substack{S \sim \mathcal{P}^n \\ \psi \sim M(S)}} \left[\frac{S[\psi] - \mathcal{P}[\psi]}{\max\{\text{sd}(\psi(\mathcal{P}), \tau)\}} > t \right] \leq \frac{2 + \frac{2}{3} \frac{t}{\tau}}{t^2} \cdot \left(\varepsilon + \frac{\ln 2}{n} \right).$$

In our application of Proposition 3.8, we have $\varepsilon = \tau^2 \geq 1/n$; setting $t = 3\tau/\beta$ and simplifying yields

$$\Pr_{\substack{S \sim \mathcal{P}^n \\ \psi \sim M(S)}} \left[\frac{S[\psi] - \mathcal{P}[\psi]}{\max\{\tau \cdot \text{sd}(\psi(\mathcal{P}), \tau^2)\}} > \frac{3}{\beta} \right] \leq \beta.$$

Note that this bound appears to correspond to an application of Markov's inequality to the conclusion Proposition 3.8. However, since the random variable in question may take both positive and negative

values, Markov’s inequality cannot be applied. Namely, Lemma 3.11 corresponds to a strengthening of Proposition 3.8 that bounds the expectation of the absolute value of the random variable. This approach can also be easily used to get a bound (based on Markov’s inequality) on the tail of the largest error we state in Theorem 1.1. This follows from the fact that a high probability bound on this tail is easy to prove when the dataset is independent from the algorithm’s answers.

To prove Lemma 3.11, we observe that for any random variable S and any randomized algorithm M ,

$$D(\mathbb{1}_E(S, M(S)) \parallel \mathbb{1}_E(S', M(S))) \leq D(S, M(S) \parallel S', M(S)) = I(S; M(S))$$

where S' is an independent copy of S , E is an arbitrary event on $\mathcal{X}^n \times \mathcal{Y}$, and $\mathbb{1}_E$ is the indicator function of the event. Note that $\mathbb{1}_E(S, M(S))$ is a Bernoulli random variable with bias equal to $\Pr[(S, M(S)) \in E]$. Now the proof of Lemma 3.11 follows from the following lemma.

Lemma 3.12. *Let $B(p)$ denote the Bernoulli random variable with bias p . Then for any $p, q \in (0, 1]$,*

$$p \leq \frac{D(B(p) \parallel B(q)) + \ln 2}{\ln(1/q)}.$$

Proof. We have

$$\begin{aligned} D(B(p) \parallel B(q)) &= p \ln\left(\frac{p}{q}\right) + (1-p) \ln\left(\frac{1-p}{1-q}\right) \\ &= p \ln(1/q) - H(p) + (1-p) \ln(1/(1-q)) \\ &\geq p \ln(1/q) - \ln 2 + 0, \end{aligned}$$

where $H(p) = p \ln(1/p) + (1-p) \ln(1/(1-p))$ is the binary entropy function. Rearranging yields the result. \square

4 Analysis of our Algorithm

Now we assemble the tools developed in the previous section to analyse our algorithm (Figure 1). To do this we must introduce some formalisms for dealing with adaptive algorithms.

Our algorithm M answers adaptively-chosen queries. We call the entity A choosing these queries the *analyst* (or *adversary* to connote worst-case behaviour). The interaction between A and M defines a function mapping inputs (the sample) to a transcript of queries and answers. Figure 2 defines how this function is computed.




Figure 2: $A \xrightarrow{\quad} M : \mathcal{X}^m \rightarrow \mathcal{Q}^k \times \mathcal{R}^k$

With this formalism in hand, we can extend our definition of average leave-one-out KL stability from non-interactive algorithms (Definition 1.2) to interactive algorithms:

Definition 4.1 (Interactive ALKL stable). *An interactive algorithm M is ε -ALKL stable if $A \xrightarrow{\leftarrow} M$ (as defined in Figure 2) is ε -ALKL stable for all interactive algorithms A .*

4.1 Stability of our algorithm

We now show that our algorithm is (interactive) average leave-one-out KL stable:

Theorem 4.2. *Our algorithm (Figure 1) is $\frac{kt}{n^2}$ -ALKL stable for any $n \geq 20$ and $T \leq \min\{t^2, tn/10\}$.*

To establish that our algorithm is average leave-one-out KL stable, we first only consider one query $\psi = \psi_j$. We show that, for each query ψ , the answer given by our algorithm is $\frac{t}{n^2}$ -ALKL stable. Using composition (Lemma 1.3), we can extend this to k queries ψ_1, \dots, ψ_k . That is, we prove that our algorithm is $\frac{kt}{n^2}$ -ALKL stable.

First we recall how our algorithm answers a query $\psi : \mathcal{X} \rightarrow [0, 1]$: The algorithm is given as input a sample $s \in \mathcal{X}^n$ and, for each statistical query $\psi = \psi_j$ the algorithm M outputs a sample from $\mathcal{N}(\mu, \max\{\sigma^2/t, 1/T\})$ where

$$\mu = s[\psi] = \frac{1}{n} \sum_{i \in [n]} \psi(s_i) \quad \text{and} \quad \sigma^2 = \mathbf{Var}[\psi(s)] = \frac{1}{n} \sum_{i \in [n]} (\psi(s_i) - \mu)^2.$$

Here $t, T > 0$ are parameters controlling the accuracy-stability tradeoff.

Theorem 4.2 is implied by the following result.

Proposition 4.3. *Let $s \in \mathcal{X}^n$, $n \geq 2$, $t, T > 0$, and $\psi : \mathcal{X} \rightarrow [0, 1]$. For $i \in [n]$, define*

$$\begin{aligned} \mu &= s[\psi] = \frac{1}{n} \sum_{i \in [n]} \psi(s_i), & \mu_{-i} &= s_{-i}[\psi] = \frac{1}{n-1} \sum_{j \in [n] \setminus \{i\}} \psi(s_j), \\ \sigma^2 &= \mathbf{Var}[\psi(s)] = \frac{1}{n} \sum_{i \in [n]} (\psi(s_i) - \mu)^2, & \sigma_{-i}^2 &= \mathbf{Var}[\psi(s_{-i})] = \frac{1}{n-1} \sum_{j \in [n] \setminus \{i\}} (\psi(s_j) - \mu_{-i})^2. \end{aligned}$$

Then

$$\frac{1}{n} \sum_{i \in [n]} \mathbb{D} \left(\mathcal{N} \left(\mu, \max \left\{ \frac{\sigma^2}{t}, \frac{1}{T} \right\} \right) \left\| \mathcal{N} \left(\mu_{-i}, \max \left\{ \frac{\sigma_{-i}^2}{t}, \frac{1}{T} \right\} \right) \right) \leq \frac{1}{4n^2} \left(2t + \frac{T}{t} \cdot (1 + \zeta) \right) \cdot (1 + \zeta), \quad (9)$$

where $1 + \zeta = \left(1 + \frac{1}{n-1}\right)^2 \left(1 + \frac{T}{tn} \left(1 + \frac{1}{n-1}\right)^2\right)$.

In particular, $\zeta = O\left(\frac{1}{n} + \frac{T}{tn}\right)$ and, if $n \geq 20$ and $\frac{T}{t} \leq \frac{n}{10}$, then (9) $\leq \frac{1}{n^2} \max\{t, T/t\}$.

The proof of Proposition 4.3 is technical and is deferred to Appendix A.

To gain some intuition, we present a proof of a weaker result which captures the main points of the proof.

Simplified Proof Sketch. We make three simplifications:

- Ignore constant factors. (Take n and $\frac{t}{T}n$ to be sufficiently large.)
- Consider $\psi : \mathcal{X} \rightarrow \{0, 1\}$ instead of $\psi : \mathcal{X} \rightarrow [0, 1]$.

- Assume $\frac{\sigma^2}{t} \geq \frac{1}{T}$ and $\frac{\sigma_{-i}^2}{t} \geq \frac{1}{T}$ for all i .

Intuitively, the last assumption is not really an assumption, since we perforce ensure this by using $\max\left\{\frac{\sigma^2}{t}, \frac{1}{T}\right\}$ in place of $\frac{\sigma^2}{t}$ and likewise for σ_{-i} . This assumption is only to simplify notation.

Begin by considering a fixed index i . By Corollary 2.4,

$$\begin{aligned}
& \mathrm{D}\left(\mathcal{N}\left(\mu, \max\left\{\frac{\sigma^2}{t}, \frac{1}{T}\right\}\right)\left\|\mathcal{N}\left(\mu_{-i}, \max\left\{\frac{\sigma_{-i}^2}{t}, \frac{1}{T}\right\}\right)\right)\right) \\
&= \mathrm{D}\left(\mathcal{N}\left(\mu, \frac{\sigma^2}{t}\right)\left\|\mathcal{N}\left(\mu_{-i}, \frac{\sigma_{-i}^2}{t}\right)\right)\right) \\
&\leq \frac{1}{2} \cdot \left(\frac{(\mu - \mu_{-i})^2}{\sigma^2/t} + \left(\frac{\sigma_{-i}^2}{\sigma^2} - 1\right)^2 \cdot \min\left\{1, \frac{1}{6}\left(2 + \frac{\sigma^2}{\sigma_{-i}^2}\right)\right\}\right) \cdot \frac{\sigma^2}{\sigma_{-i}^2} \\
&\leq t \cdot \frac{(\mu - \mu_{-i})^2}{\sigma^2} + \left(\frac{\sigma_{-i}^2 - \sigma^2}{\sigma^2}\right)^2. \quad (\text{Ignoring constant factors.})
\end{aligned}$$

We use $\frac{1}{2} \frac{\sigma^2}{\sigma_{-i}^2} \leq 1$, which holds as long as $T \leq tn$.

Since we assumed (for simplicity) that ψ takes only values 0 and 1, the entire vector $\psi(s)$ is characterized by the mean $\mu = s[\psi]$. In particular, we have $\sigma^2 = \mu(1 - \mu)$ and $\sigma_{-i}^2 = \mu_{-i}(1 - \mu_{-i})$. Thus, we have $|\sigma_{-i}^2 - \sigma^2| \leq |\mu - \mu_{-i}|$ and

$$\mathrm{D}\left(\mathcal{N}\left(\mu, \max\left\{\frac{\sigma^2}{t}, \frac{1}{T}\right\}\right)\left\|\mathcal{N}\left(\mu_{-i}, \max\left\{\frac{\sigma_{-i}^2}{t}, \frac{1}{T}\right\}\right)\right)\right) \leq t \cdot \frac{(\mu - \mu_{-i})^2}{\sigma^2} + \left(\frac{\mu - \mu_{-i}}{\sigma^2}\right)^2 = \left(t + \frac{1}{\sigma^2}\right) \cdot \frac{(\mu - \mu_{-i})^2}{\sigma^2}.$$

Since $\frac{\sigma^2}{t} \geq \frac{1}{T}$, we have $t + \frac{1}{\sigma^2} \leq t + \frac{T}{t}$. Also

$$\mu - \mu_{-i} = \frac{1}{n} \left(\sum_{i'} \psi(s_{i'})\right) - \frac{1}{n-1} \left(-\psi(s_i) + \sum_{i'} \psi(s_{i'})\right) = \left(\frac{1}{n} - \frac{1}{n-1}\right) \cdot n\mu + \frac{\psi(s_i)}{n-1} = \frac{\psi(s_i) - \mu}{n-1}.$$

We now have

$$\mathrm{D}\left(\mathcal{N}\left(\mu, \max\left\{\frac{\sigma^2}{t}, \frac{1}{T}\right\}\right)\left\|\mathcal{N}\left(\mu_{-i}, \max\left\{\frac{\sigma_{-i}^2}{t}, \frac{1}{T}\right\}\right)\right)\right) \leq \left(t + \frac{T}{t}\right) \cdot \frac{(\psi(s_i) - \mu)^2}{\sigma^2(n-1)^2}.$$

Averaging over all $i \in [n]$ gives

$$\begin{aligned}
\frac{1}{n} \sum_{i \in [n]} \mathrm{D}\left(\mathcal{N}\left(\mu, \max\left\{\frac{\sigma^2}{t}, \frac{1}{T}\right\}\right)\left\|\mathcal{N}\left(\mu_{-i}, \max\left\{\frac{\sigma_{-i}^2}{t}, \frac{1}{T}\right\}\right)\right)\right) &\leq \left(t + \frac{T}{t}\right) \cdot \frac{\frac{1}{n} \sum_{i \in [n]} (\psi(s_i) - \mu)^2}{\sigma^2(n-1)^2} \\
&= \left(t + \frac{T}{t}\right) \cdot \frac{\sigma^2}{\sigma^2(n-1)^2} \\
&\leq \left(t + \frac{T}{t}\right) \cdot \frac{2}{n^2},
\end{aligned}$$

as desired. □

4.2 Accuracy guarantees

Note that, by the postprocessing property of average leave-one-out KL stability, applying any function f to the transcript of an average leave-one-out KL stable algorithm still is average leave-one-out KL stable. More precisely, for a ε -ALKL stable interactive algorithm M and any interactive algorithm A , the composed algorithm mapping input s to output $f(A \overset{\rightarrow}{\leftarrow} M(s))$ is ε -ALKL stable. We now invoke the generalization properties of average leave-one-out KL stability:

Lemma 4.4. *Fix a distribution \mathcal{P} on \mathcal{X} . Let M be a ε -ALKL stable interactive algorithm that answers k statistical queries. Let A be an arbitrary interactive algorithm that asks k statistical queries. Let $f : \mathcal{Q}^k \times \mathbb{R}^k \rightarrow \mathcal{Q}$, where \mathcal{Q} denotes the set of statistical queries $\psi : \mathcal{X} \rightarrow [0, 1]$. Let $\tau = \sqrt{\varepsilon}$. Then*

$$\left| \mathbf{E}_{\substack{S \sim \mathcal{P}^n \\ \psi \sim f(A \overset{\rightarrow}{\leftarrow} M(S))}} \left[\frac{S[\psi] - \mathcal{P}[\psi]}{\max\{\text{sd}(\psi(\mathcal{P})), \tau\}} \right] \right| \leq 2\sqrt{\varepsilon} \quad (10)$$

and

$$\mathbf{E}_{\substack{S \sim \mathcal{P}^n \\ \psi \sim f(A \overset{\rightarrow}{\leftarrow} M(S))}} \left[\left(\frac{\text{sd}(\psi(S))}{\max\{\text{sd}(\psi(\mathcal{P})), \tau\}} \right)^2 \right] \leq 3. \quad (11)$$

Proof. This follows from our generalization results (Propositions 3.8 and 3.10), postprocessing, and the connection between average leave-one-out KL stability and mutual information (Proposition 3.3). \square

A natural choice of function f is to simply pick out one of the queries — that is, $f(\psi_1, \dots, \psi_k, v_1, \dots, v_k) = \psi_j$ for some fixed j . However, f can also pick out the “worst” query. For example, the monitor technique of Bassily et al. [BNS⁺16] takes

$$f_*(\psi_1, \dots, \psi_k, v_1, \dots, v_k) = \psi_{j_*}, \quad \text{where} \quad j_* = \operatorname{argmax}_{j \in [k]} |v_j - \mathcal{P}[\psi_j]|.$$

The monitor technique allows us to reason about a single query and derive bounds that apply to all queries simultaneously as this query is the worst query. Since we have a more refined notion of error, we must use a slightly different argument.

We use the following function $f_{\tau, \mathcal{P}} : \mathcal{Q}^k \times \mathbb{R}^k \rightarrow \mathcal{Q}$ to pick out the query with the worst scaled error.

$$f_{\tau, \mathcal{P}}(\psi_1, \dots, \psi_k, v_1, \dots, v_k) = \begin{cases} \psi_{j_*} & \text{if } v_{j_*} \geq \mathcal{P}[\psi_{j_*}] \\ 1 - \psi_{j_*} & \text{if } v_{j_*} < \mathcal{P}[\psi_{j_*}] \end{cases}, \quad \text{where} \quad j_* = \operatorname{argmax}_{j \in [k]} \frac{|v_j - \mathcal{P}[\psi_j]|}{\max\{\text{sd}(\psi_j(\mathcal{P})), \tau\}}. \quad (12)$$

Note that we flip the sign to ensure that the error is always positive.

Now we give a bound on the expected scaled error of our algorithm. We use the following simple technical lemma bounding the maximum of standard Gaussians.

Lemma 4.5. *Let ξ_1, \dots, ξ_k be independent samples from $\mathcal{N}(0, 1)$. Then*

$$\mathbf{E} \left[\max\{\xi_1^2, \dots, \xi_k^2\} \right] \leq 2 \ln(2k).$$

Proof. Let $X = \max_{i \in [k]} \xi_i^2$ and $t > 0$. Since $\cosh(\sqrt{z}) = \frac{1}{2}(e^{\sqrt{z}} + e^{-\sqrt{z}})$ is a convex function of $z \geq 0$, Jensen's inequality gives

$$\cosh\left(t\sqrt{\mathbf{E}[X]}\right) \leq \mathbf{E}\left[\cosh\left(t\sqrt{X}\right)\right] = \mathbf{E}\left[\max_j^k \cosh(t\xi_j)\right] \leq \mathbf{E}\left[\sum_j^k \cosh(t\xi_j)\right] = k \cdot e^{t^2/2}.$$

Rearranging yields

$$\mathbf{E}[X] \leq \left(\frac{1}{t} \cosh^{-1}\left(k \cdot e^{t^2/2}\right)\right)^2 \leq \left(\frac{1}{t} \ln(2k \cdot e^{t^2/2})\right)^2 = \left(\frac{\ln(2k)}{t} + \frac{t}{2}\right)^2,$$

as $\cosh(v) \geq e^v/2$ and, hence, $\cosh^{-1}(u) \leq \ln(2u)$. Setting $t = \sqrt{2\ln(2k)}$ completes the proof. \square

Theorem 4.6 (Main Theorem). *Fix $n, k \geq 20$. Let M be our algorithm from Figure 1 with $T = n^2/k$ and $t = n\sqrt{2\ln(2k)/k}$ that answers k statistical queries given n samples.*

Let \mathcal{P} be a distribution on \mathcal{X} and let A be an interactive algorithm that asks k statistical queries. Then

$$\mathbf{E}_{\substack{S \sim \mathcal{P}^n \\ (\psi, v) \sim A \xrightarrow{\leftarrow} M(S)}} \left[\max_{j \in [k]} \frac{|v_j - \mathcal{P}[\psi_j]|}{\max\{\tau \cdot \text{sd}(\psi_j(\mathcal{P})), \tau^2\}} \right] \leq 4,$$

where $\tau = \sqrt{\frac{\sqrt{2k \ln(2k)}}{n}}$.

Proof. Set $\varepsilon = \tau^2 = \frac{kt}{n^2} = \frac{t}{T} = \frac{\sqrt{2k \ln(2k)}}{n}$. Since $n \geq 20$ and $T \leq \min\{t^2, nt/10\}$, $A \xrightarrow{\leftarrow} M$ is ε -ALKL stable by Theorem 4.2. Let $S \sim \mathcal{P}^n$ and $(\psi, v) \sim A \xrightarrow{\leftarrow} M(S)$. Define

$$j_* = \operatorname{argmax}_{j \in [k]} \frac{|v_j - \mathcal{P}[\psi_j]|}{\max\{\text{sd}(\psi_j(\mathcal{P})), \tau\}}$$

and

$$(\psi_*, v_*) = \begin{cases} (\psi_{j_*}, v_{j_*}) & \text{if } v_{j_*} \geq \mathcal{P}[\psi_{j_*}] \\ (1 - \psi_{j_*}, 1 - v_{j_*}) & \text{if } v_{j_*} < \mathcal{P}[\psi_{j_*}] \end{cases}.$$

Thus

$$\mathbf{E}_{\substack{S \sim \mathcal{P}^n \\ (\psi, v) \sim A \xrightarrow{\leftarrow} M(S)}} \left[\max_{j \in [k]} \frac{|v_j - \mathcal{P}[\psi_j]|}{\max\{\tau \cdot \text{sd}(\psi_j(\mathcal{P})), \tau^2\}} \right] = \frac{1}{\tau} \mathbf{E} \left[\frac{v_* - \mathcal{P}[\psi_*]}{\max\{\text{sd}(\psi_*(\mathcal{P})), \tau\}} \right]. \quad (13)$$

Let $\sigma_* = \sigma_{j_*} = \sqrt{\frac{1}{n} \sum_{i=1}^n (\psi_{j_*}(S_i) - S[\psi_{j_*}])^2}$ be the empirical standard deviation corresponding to the query ψ_* and the sample S . By Lemma 4.4,

$$\mathbf{E} \left[\frac{S[\psi_*] - \mathcal{P}[\psi_*]}{\max\{\text{sd}(\psi_*(\mathcal{P})), \tau\}} \right] \leq 2\tau \quad (14)$$

and

$$\mathbf{E} \left[\left(\frac{\sigma_*}{\max\{\text{sd}(\psi_*(\mathcal{P})), \tau\}} \right)^2 \right] \leq 3. \quad (15)$$

Let ξ_1, \dots, ξ_k be the independent standard Gaussians sampled by M (Figure 1). Let $\xi_* = \xi_{j_*}$ if $v_{j_*} \geq \mathcal{P}[\psi_{j_*}]$ and $\xi_* = -\xi_{j_*}$ if $v_{j_*} < \mathcal{P}[\psi_{j_*}]$. By the definition of our algorithm,

$$v_* = S[\psi_*] + \xi_* \cdot \sqrt{\max\{\sigma_*^2/t, 1/T\}} = S[\psi_*] + \frac{1}{\sqrt{t}} \cdot \xi_* \cdot \max\{\sigma_*, \tau\}.$$

Thus, by Cauchy-Schwartz,

$$\begin{aligned} \mathbf{E} \left[\frac{v_* - \mathcal{P}[\psi_*]}{\max\{\text{sd}(\psi_*(\mathcal{P})), \tau\}} \right] &= \mathbf{E} \left[\frac{S[\psi_*] - \mathcal{P}[\psi_*]}{\max\{\text{sd}(\psi_*(\mathcal{P})), \tau\}} \right] + \frac{1}{\sqrt{t}} \cdot \mathbf{E} \left[\frac{\xi_* \cdot \max\{\sigma_*, \tau\}}{\max\{\text{sd}(\psi_*(\mathcal{P})), \tau\}} \right] \\ &\leq \mathbf{E} \left[\frac{S[\psi_*] - \mathcal{P}[\psi_*]}{\max\{\text{sd}(\psi_*(\mathcal{P})), \tau\}} \right] + \frac{1}{\sqrt{t}} \cdot \sqrt{\mathbf{E}[\xi_*^2] \cdot \mathbf{E} \left[\left(\frac{\max\{\sigma_*, \tau\}}{\max\{\text{sd}(\psi_*(\mathcal{P})), \tau\}} \right)^2 \right]} \\ \text{(by (14))} \quad &\leq 2\tau + \frac{1}{\sqrt{t}} \cdot \sqrt{\mathbf{E}[\xi_*^2] \cdot \mathbf{E} \left[\left(\frac{\sigma_*}{\max\{\text{sd}(\psi_*(\mathcal{P})), \tau\}} \right)^2 + \left(\frac{\tau}{\max\{\text{sd}(\psi_*(\mathcal{P})), \tau\}} \right)^2 \right]} \\ \text{(by (15))} \quad &\leq 2\tau + \frac{1}{\sqrt{t}} \cdot \sqrt{\mathbf{E}[\xi_*^2] \cdot (3+1)} \\ \text{(by Lemma 4.5)} \quad &\leq 2\tau + \frac{1}{\sqrt{t}} \cdot \sqrt{2 \ln(2k) \cdot 4} \\ &= 4\tau = 4\sqrt{\frac{\sqrt{2k \ln(2k)}}{n}}. \end{aligned}$$

Combining with (13) completes the proof. \square

Acknowledgements

We thank Adam Smith for encouraging us to prove our generalization bounds via the mutual information between the input and output of the ALKL stable algorithm. This suggestion greatly improved our results and simplified the analysis.

References

- [BE02] Olivier Bousquet and André Elisseeff. Stability and generalization. *JMLR*, 2:499–526, 2002.
- [Ber24] Sergei Bernstein. On a modification of chebyshev’s inequality and of the error formula of laplace. *Ann. Sci. Inst. Sav. Ukraine, Sect. Math*, 1(4):38–49, 1924.
- [BF16] Raef Bassily and Yoav Freund. Typicality-based stability and privacy. *CoRR*, abs/1604.03336, 2016.
- [BMDG05] Arindam Banerjee, Srujana Merugu, Inderjit S Dhillon, and Joydeep Ghosh. Clustering with bregman divergences. *Journal of machine learning research*, 6(Oct):1705–1749, 2005.
- [BNS⁺16] Raef Bassily, Kobbi Nissim, Adam D. Smith, Thomas Steinke, Uri Stemmer, and Jonathan Ullman. Algorithmic stability for adaptive data analysis. In *STOC*, pages 1046–1059, 2016.

- [BS16] Mark Bun and Thomas Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of Cryptography Conference*, pages 635–658. Springer Berlin Heidelberg, 2016.
- [CT12] Thomas M Cover and Joy A Thomas. *Elements of information theory*. John Wiley & Sons, 2 edition, 2012.
- [DFH⁺14] Cynthia Dwork, Vitaly Feldman, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Aaron Roth. Preserving statistical validity in adaptive data analysis. *CoRR*, abs/1411.2664, 2014. Extended abstract in STOC 2015.
- [DFH⁺15] Cynthia Dwork, Vitaly Feldman, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Aaron Roth. Generalization in adaptive data analysis and holdout reuse. *CoRR*, abs/1506, 2015. Extended abstract in NIPS 2015.
- [DKM⁺06] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 486–503. Springer Berlin Heidelberg, 2006.
- [DMNS06] C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *TCC*, pages 265–284, 2006.
- [FS17] Vitaly Feldman and Thomas Steinke. Generalization for adaptively-chosen estimators via stable median. In *Conference on Learning Theory (COLT)*, 2017.
- [FSG08] Béla A Frigyük, Santosh Srivastava, and Maya R Gupta. Functional bregman divergence and bayesian estimation of distributions. *IEEE Transactions on Information Theory*, 54(11):5130–5139, 2008.
- [GAL13] Manuel Gil, Fady Alajaji, and Tamas Linder. Rényi divergence measures for commonly used univariate continuous distributions. *Information Sciences*, 249:124–131, 2013.
- [Gra11] Robert M Gray. *Entropy and information theory*. Springer, 2011.
- [HU14] M. Hardt and J. Ullman. Preventing false discovery in interactive data analysis is hard. In *FOCS*, pages 454–463, 2014.
- [Kea98] M. Kearns. Efficient noise-tolerant learning from statistical queries. *Journal of the ACM*, 45(6):983–1006, 1998.
- [PRMN04] Tomaso Poggio, Ryan Rifkin, Sayan Mukherjee, and Partha Niyogi. General conditions for predictivity in learning theory. *Nature*, 428(6981):419–422, 2004.
- [RRST16] Ryan Rogers, Aaron Roth, Adam Smith, and Om Thakkar. Max-information, differential privacy, and post-selection hypothesis testing. In *Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on*, pages 487–494. IEEE, 2016.
- [RRT⁺16] Maxim Raginsky, Alexander Rakhlin, Matthew Tsao, Yihong Wu, and Aolin Xu. Information-theoretic analysis of stability and bias of learning algorithms. In *2016 IEEE Information Theory Workshop, ITW 2016, Cambridge, United Kingdom, September 11-14, 2016*, pages 26–30, 2016.

- [RZ16] Daniel Russo and James Zou. Controlling bias in adaptive data analysis using information theory. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics, AISTATS*, 2016.
- [Ste16] Thomas Steinke. Adaptive data analysis. 2016. Lecture Notes <http://people.seas.harvard.edu/~madhusudan/courses/Spring2016/notes/thomas-notes-ada.pdf>.
- [SU15] Thomas Steinke and Jonathan Ullman. Interactive fingerprinting codes and the hardness of preventing false discovery. In *COLT*, pages 1588–1628, 2015.
- [VEH14] Tim Van Erven and Peter Harremoos. Rényi divergence and kullback-leibler divergence. *IEEE Transactions on Information Theory*, 60(7):3797–3820, 2014.
- [WLF16] Yu-Xiang Wang, Jing Lei, and Stephen E Fienberg. On-average kl-privacy and its equivalence to generalization for max-entropy mechanisms. In *International Conference on Privacy in Statistical Databases*, pages 121–134. Springer, 2016.
- [XR17] Aolin Xu and Maxim Raginsky. Information-theoretic analysis of generalization capability of learning algorithms. *CoRR*, abs/1705.07809, 2017.

A Average Leave-One-Out KL Stability of Scaled Gaussian Noise

Proposition A.1 (Restating Proposition 4.3). *Let $s \in \mathcal{X}^n$, $n \geq 2$, $t, T > 0$, and $\psi : \mathcal{X} \rightarrow [0, 1]$. For $i \in [n]$, define*

$$\begin{aligned} \mu &= s[\psi] = \frac{1}{n} \sum_{i \in [n]} \psi(s_i), & \mu_{-i} &= s_{-i}[\psi] = \frac{1}{n-1} \sum_{j \in [n] \setminus \{i\}} \psi(s_j), \\ \sigma^2 &= \mathbf{Var}[\psi(s)] = \frac{1}{n} \sum_{i \in [n]} (\psi(s_i) - \mu)^2, & \sigma_{-i}^2 &= \mathbf{Var}[\psi(s_{-i})] = \frac{1}{n-1} \sum_{j \in [n] \setminus \{i\}} (\psi(s_j) - \mu_{-i})^2. \end{aligned}$$

Then

$$\frac{1}{n} \sum_{i \in [n]} \mathbf{D} \left(\mathcal{N} \left(\mu, \max \left\{ \frac{\sigma^2}{t}, \frac{1}{T} \right\} \right) \middle\| \mathcal{N} \left(\mu_{-i}, \max \left\{ \frac{\sigma_{-i}^2}{t}, \frac{1}{T} \right\} \right) \right) \leq \frac{1}{4n^2} \left(2t + \frac{T}{t} \cdot (1 + \zeta) \right) \cdot (1 + \zeta),$$

where $1 + \zeta = \left(1 + \frac{1}{n-1} \right)^2 \left(1 + \frac{T}{tn} \left(1 + \frac{1}{n-1} \right)^2 \right)$.

Proof. We have

$$\mu - \mu_{-i} = \mu - \frac{n\mu - \psi(s_i)}{n-1} = \frac{\psi(s_i) - \mu}{n-1} \tag{16}$$

and

$$\frac{1}{n} \sum_{i \in [n]} (\mu - \mu_{-i})^2 = \frac{1}{n(n-1)^2} \sum_{i \in [n]} (\psi(s_i) - \mu)^2 = \frac{\sigma^2}{(n-1)^2}. \tag{17}$$

Furthermore,

$$\begin{aligned}
\sigma^2 - \sigma_{-i}^2 &= \sigma^2 - \left(\frac{1}{n-1} \left(-\psi(s_i)^2 + \sum_{j \in [n]} \psi(s_j)^2 \right) - \mu_{-i}^2 \right) \\
&= \sigma^2 - \left(\frac{1}{n-1} \left(-\psi(s_i)^2 + n(\sigma^2 + \mu^2) \right) - \mu_{-i}^2 \right) \\
&= \sigma^2 \left(1 - \frac{n}{n-1} \right) + \frac{\psi(s_i)^2}{n-1} - \frac{n}{n-1} \mu^2 + \mu_{-i}^2 \\
&= \frac{\psi(s_i)^2 - \mu^2 - \sigma^2}{n-1} - (\mu^2 - \mu_{-i}^2) \\
&= \frac{(\psi(s_i) - \mu)(\psi(s_i) + \mu) - \sigma^2}{n-1} - (\mu - \mu_{-i})(\mu + \mu_{-i}) \\
&= \frac{(\psi(s_i) - \mu)(\psi(s_i) + \mu) - \sigma^2}{n-1} - \frac{\psi(s_i) - \mu}{n-1} (\mu + \mu_{-i}) \quad (\text{by (16)}) \\
&= \frac{\psi(s_i) - \mu}{n-1} \left((\psi(s_i) - \mu) + (\mu - \mu_{-i}) \right) - \frac{\sigma^2}{n-1} \\
&= \frac{\psi(s_i) - \mu}{n-1} \left((\psi(s_i) - \mu) + \frac{\psi(s_i) - \mu}{n-1} \right) - \frac{\sigma^2}{n-1} \quad (\text{by (16)}) \\
&= \frac{\frac{n}{n-1} (\psi(s_i) - \mu)^2 - \sigma^2}{n-1} \tag{18}
\end{aligned}$$

and

$$\left| \sigma^2 - \sigma_{-i}^2 \right| \leq \frac{\max\left\{ \frac{n}{n-1} (\psi(s_i) - \mu)^2, \sigma^2 \right\}}{n-1} \leq \frac{n}{(n-1)^2}. \tag{19}$$

By (18),

$$\begin{aligned}
\frac{1}{n} \sum_{i \in [n]} (\sigma^2 - \sigma_{-i}^2)^2 &= \frac{1}{n(n-1)^2} \sum_{i \in [n]} \left(\frac{n}{n-1} (\psi(s_i) - \mu)^2 - \sigma^2 \right)^2 \\
&= \frac{1}{n(n-1)^2} \sum_{i \in [n]} \frac{n^2}{(n-1)^2} (\psi(s_i) - \mu)^4 - \frac{2n}{n-1} (\psi(s_i) - \mu)^2 \sigma^2 + \sigma^4 \\
&\leq \frac{1}{n(n-1)^2} \sum_{i \in [n]} \frac{n^2}{(n-1)^2} (\psi(s_i) - \mu)^2 - \frac{2n}{n-1} (\psi(s_i) - \mu)^2 \sigma^2 + \sigma^4 \\
&\quad (\text{Since } 0 \leq (\psi(s_i) - \mu)^2 \leq 1.) \\
&= \frac{1}{(n-1)^2} \left(\frac{n^2}{(n-1)^2} \sigma^2 - \frac{2n}{n-1} \sigma^2 \sigma^2 + \sigma^4 \right) \\
&= \frac{\sigma^2}{(n-1)^2} \left(\frac{n^2}{(n-1)^2} - \sigma^2 \frac{n+1}{n-1} \right) \\
&\leq \frac{\sigma^2}{(n-1)^2} \frac{n^2}{(n-1)^2}. \tag{20}
\end{aligned}$$

Let

$$\gamma = \max_i \left| \frac{\max \left\{ \frac{\sigma^2}{t}, \frac{1}{T} \right\}}{\max \left\{ \frac{\sigma_{-i}^2}{t}, \frac{1}{T} \right\}} - 1 \right|.$$

By (19), for some i ,

$$\gamma = \frac{\left| \max \left\{ \frac{\sigma_{-i}^2}{t}, \frac{1}{T} \right\} - \max \left\{ \frac{\sigma^2}{t}, \frac{1}{T} \right\} \right|}{\max \left\{ \frac{\sigma_{-i}^2}{t}, \frac{1}{T} \right\}} \leq \frac{\left| \frac{\sigma_{-i}^2}{t} - \frac{\sigma^2}{t} \right|}{\frac{1}{T}} = \frac{T}{t} |\sigma^2 - \sigma_{-i}^2| \leq \frac{Tn}{t(n-1)^2}. \quad (21)$$

By Corollary 2.4, we have

$$\begin{aligned} & D \left(\mathcal{N} \left(\mu, \max \left\{ \frac{\sigma^2}{t}, \frac{1}{T} \right\} \right) \middle\| \mathcal{N} \left(\mu_{-i}, \max \left\{ \frac{\sigma_{-i}^2}{t}, \frac{1}{T} \right\} \right) \right) \\ & \leq \frac{1}{2} \left(\frac{(\mu - \mu_{-i})^2}{\max \left\{ \frac{\sigma^2}{t}, \frac{1}{T} \right\}} + \left(\frac{\max \left\{ \frac{\sigma_{-i}^2}{t}, \frac{1}{T} \right\}}{\max \left\{ \frac{\sigma^2}{t}, \frac{1}{T} \right\}} - 1 \right)^2 \cdot \min \left\{ 1, \frac{2 + (1 + \gamma)}{6} \right\} \right) \cdot (1 + \gamma) \\ & \leq \frac{1}{2} \left(\frac{(\mu - \mu_{-i})^2}{\max \left\{ \frac{\sigma^2}{t}, \frac{1}{T} \right\}} + \frac{\left(\max \left\{ \frac{\sigma_{-i}^2}{t}, \frac{1}{T} \right\} - \max \left\{ \frac{\sigma^2}{t}, \frac{1}{T} \right\} \right)^2}{\max \left\{ \frac{\sigma^2}{t}, \frac{1}{T} \right\}^2} \cdot \frac{1}{2} \left(1 + \frac{\gamma}{3} \right) \right) \cdot (1 + \gamma) \\ & \leq \frac{1}{2} \left(\frac{(\mu - \mu_{-i})^2}{\frac{\sigma^2}{t}} + \frac{\left(\frac{\sigma_{-i}^2}{t} - \frac{\sigma^2}{t} \right)^2}{\frac{\sigma^2}{t} \cdot \frac{1}{T}} \cdot \frac{1 + \gamma/3}{2} \right) \cdot (1 + \gamma). \end{aligned} \quad (22)$$

Combining (22), (17), (20), and (21), we have

$$\begin{aligned}
& \frac{1}{n} \sum_{i \in [n]} D \left(\mathcal{N} \left(\mu, \max \left\{ \frac{\sigma^2}{t}, \frac{1}{T} \right\} \right) \left\| \mathcal{N} \left(\mu_{-i}, \max \left\{ \frac{\sigma_{-i}^2}{t}, \frac{1}{T} \right\} \right) \right) \\
& \leq \frac{1}{n} \sum_{i \in [n]} \frac{1}{2} \left(\frac{(\mu - \mu_{-i})^2}{\frac{\sigma^2}{t}} + \frac{(\sigma^2 - \sigma_{-i}^2)^2}{2 \cdot t^2 \cdot \frac{\sigma^2}{t} \cdot \frac{1}{T}} (1 + \gamma/3) \right) \cdot (1 + \gamma) \\
& \leq \frac{1}{2} \left(\frac{\frac{\sigma^2}{(n-1)^2}}{\frac{\sigma^2}{t}} + \frac{\frac{\sigma^2}{(n-1)^2} \frac{n^2}{(n-1)^2}}{2 \cdot t^2 \cdot \frac{\sigma^2}{t} \cdot \frac{1}{T}} (1 + \gamma/3) \right) \cdot (1 + \gamma) \\
& = \frac{1}{4(n-1)^2} \left(2t + \frac{T}{t} \frac{n^2}{(n-1)^2} (1 + \gamma/3) \right) \cdot (1 + \gamma) \\
& \leq \frac{1}{4(n-1)^2} \left(2t + \frac{T}{t} \frac{n^2}{(n-1)^2} \left(1 + \frac{Tn}{3t(n-1)^2} \right) \right) \cdot \left(1 + \frac{Tn}{t(n-1)^2} \right) \\
& = \frac{1}{4n^2} \left(2t + \frac{T}{t} \cdot \left(1 + \frac{1}{n-1} \right)^2 \left(1 + \frac{Tn}{3t(n-1)^2} \right) \right) \cdot \left(1 + \frac{1}{n-1} \right)^2 \left(1 + \frac{Tn}{t(n-1)^2} \right) \\
& \leq \frac{1}{4n^2} \left(2t + \frac{T}{t} \cdot (1 + \zeta) \right) \cdot (1 + \zeta). \tag{23}
\end{aligned}$$

□