The Algorithmic Foundations of Adaptive Data Analysis November 21, 2017

Homework 1

Instructors: Aaron Roth and Adam Smith

1. **(Analyst strategies for simple mechanisms)**

   In this problem, we will revisit and optimize the analyst strategies from Lecture 3.

   - the rounded empirical mechanism $OT_{1/n}$, which returns the answer to a statistical query rounded to the nearest multiple of $1/n$.
   - The Guassian mechanism with parameter σ.

   (a) For each the two mechanisms above, implement the mechanism as a function that takes a data set and a query function and returns the appropriate answer (passing a query is easiest in a language with first-order functions; Python has these).

   (b) Implement and test the attack from Lecture 3 using data drawn uniformly from $\{0,1\}^{k+1}$ for each of the two mechanisms.

   Do this for $k \in \{10, 100, 1000\}$ and $n \in \left\{\frac{k}{10}, \frac{k}{2}, k, 2k, 10k\right\}$. When you run the experiment, record both the empirical and population error of the mechanism. Run each setting several times (at least 10) to get a sense of the expected performance and its variability.

   How does the value of σ affect the accuracy of the Gaussian mechanism? Is there an optimal setting of σ?

2. (**KL stability and mean-squared error**) Let $M : \mathcal{X}^n \to (\mathcal{X} \to [0,1])$ be a randomized algorithm that outputs a statistical query $\phi : \mathcal{X} \to [0,1]$. In this exercise you will show if M is τ-KL-stable, then its expected squared bias is bounded, namely

$$\mathop{\mathbb{E}}_{\substack{\mathbf{S}\sim\mathcal{D}^n \\ \phi_{\mathbf{S}}\sim M(\mathbf{S})}} \left((\phi_{\mathbf{S}}(\mathbf{S}) - \phi_{\mathbf{S}}(\mathcal{D}))^2 \right) \leq O(\frac{1}{n} + \tau)\,. \tag{1}$$

   (a) Consider the distributions $(\mathbf{S}, M(\mathbf{S}))$ and $(\mathbf{S}', M(\mathbf{S}))$ where $\mathbf{S}'$ is a fresh sample of size n drawn i.i.d. from $\mathcal{D}$.

   Show that $D_{KL}((\mathbf{S}, M(\mathbf{S}))\|(\mathbf{S}', M(\mathbf{S})) \leq n\tau$. The divergence $D_{KL}((\mathbf{S}, M(\mathbf{S}))\|(\mathbf{S}', M(\mathbf{S}))$ is called the *mutual information* between $\mathbf{S}$ and $M(\mathbf{S})$.

   (b) Let A, B be random variables taking values in the same set, with distributions P, Q respectively, such that $D_{KL}(A\|B)$ is well-defined and finite. Show that for every real-valued function f, we have

$$\mathbb{E}(f(A)) \leq D_{KL}(A\|B) + \ln\left(\mathbb{E}\left(e^{f(B)}\right)\right).$$

   [*Hint:* Use Jensen's inequality. For your proof, it's ok to assume that the set in which A, B take values is finite.]

   (c) Show that, for every $\lambda \in (0,1)$, for every statistical query ϕ, and for every distribution $\mathcal{D}$ on $\mathcal{X}$, we have

$$\mathbb{E}(\exp(\frac{\lambda}{2\sigma^2}(\phi(\mathbf{S}) - \phi(\mathcal{D}))^2)) \leq \frac{1}{\sqrt{1-\lambda}}\,,$$

   where $\sigma^2 = 1/n$. (Note that here ϕ is fixed and independent of $\mathbf{S}$.)

   To do this, first use the Chernoff bound to show that for every value $t > 0$, we have $\Pr(|\phi(\mathbf{S}) - \phi(\mathcal{D})| > t) \leq \Pr(|Z| > t)$ where Z is an appropriate Gaussian distribution. Then calculate $\mathbb{E}(\exp(\frac{\lambda}{2\sigma^2}Z^2))$.

(d) Prove the bound in (1) above by using part (b) to bound the expectation of $f(\mathbf{S}) = \frac{\lambda}{2\sigma^2}(\phi_{\mathbf{S}}(\mathbf{S}) - \phi_{\mathbf{S}}(\mathcal{D}))^2$. You will have to choose λ (but many different choices will get the right asymptotics).

3. **(Differentially Private Algorithms)**

   (a) Show that the exponentnial mechanism seen in class is equivalent to the report noisy max mechanism. [*Hint:* Consider two outputs a, b. For a fixed input $\mathbf{s}$, what is $\frac{P(a|\mathbf{s})}{P(b|\mathbf{s})}$?]

   (b) Show that if we did not add noise to T (that is, we set $\tilde{T} = T$), then the Sparse Vector Mechanism would not be $(\epsilon, 0)$-differentially private for any finite value of ϵ.

   (c) Show that the following algorithm is (ϵ, δ)-differentially private over any domain $\mathcal{X}$.

```
Algorithm 1: Stable Histogam(s; ε, δ)
for every x ∈ 𝒳 that appears in s do
    c̃_x = #{i : x_i = x} + Lap(2/ε);
Release the set of pairs {(x, c̃_x) : c̃_x > τ} where τ = 1 + 2 ln(1/δ)/ε.
```

   (d) Suppse we modify the stable histogram to only output the set $\{x : \#\{i : x_i = x\} > \tau\}$. How compressible is the stable histogram algorithm as a function of n, $|\mathcal{X}|$ and the cutoff τ? What happens when $\mathcal{X}$ is infinite?

4. **(Differentially Private Ladder)** Complete Exercise 2 from Lecture 12. It is ok to only do the analysis for the Ladder algorithm.