CSC2412:Algorithms for Private Data Analysis Updated Fall 2020

Projection Mechanism

Aleksandar Nikolov *Scribe: Tyrone Strangway, Sepehr Abbasi Zadeh*

1 Review: the Query Release Problem

Let us recall the query release problem. An instance of this problem is defined by a set of k linear queries, $Q = \{q_1, \ldots, q_k\}$ and a dataset X, which is a multiset of n elements from the universe $\mathcal{X}$. For a single point x, $q_i(x) \in [0,1]\}$ and $q_i(X) = \frac{1}{n}\sum_{x \in X} q_i(x)$. When $q_i(x) \in \{0,1\}$ for all $x \in \mathcal{X}$, we can think of $q_i(x)$ as a property, and of $q_i(X)$ as the fraction of points in X with that property. We want to accurately estimate the query answers $Q(X) = (q_1(X), \ldots, q_k(X))$ on X using an (ε, δ)-differentially private mechanism $\mathcal{M}$. I.e., for all i, we want to have $|\mathcal{M}(x)_i - q_i(x)| \leq \alpha$. The main question we ask is, given Q, what is the smallest dataset size n for which this task is solvable for all X.

In this course, we have seen a couple of mechanisms to solve this problem:

1. **Gaussian Mechanism:** For each $i \in [k]$ output $q_i(X) + Z_i$, where $Z_i \sim \mathcal{N}(0, \frac{k}{n^2\rho})$. That is we add independent Gaussian noise to each query answer. We showed that for some $\rho = \rho_{\varepsilon,\delta} \approx \frac{\varepsilon^2}{\log \frac{1}{\delta}}$, the Gaussian Mechanism is (ε, δ)-differentially private. It is easy to show that to get accuracy α for every query, we need to set $n \gg \frac{\sqrt{k \log k}\sqrt{\log \frac{1}{\delta}}}{\alpha\varepsilon}$. The running time of the Gaussian Mechanism is $O(kn)$, assuming each query can be evaluated in constant time on a dataset point.
2. **Private Multiplicative Weights:** We try to learn a distribution that approximates the empirical distribution of X, by starting with the uniform distribution, and iteratively refining it. The refinement is based on the multiplicative weights method, hence the name. This mechanism is (ε, δ)-differentially private with an accuracy of α when $n \gg \frac{\log k \sqrt{\log |\mathcal{X}|}\sqrt{\log \frac{1}{\delta}}}{\alpha^2\varepsilon}$. This mechanism requires running time $O(knL|\mathcal{X}|)$, where $L = O(\frac{\log |\mathcal{X}|}{\alpha^2})$ is the number of refinements we execute.

The Gaussian mechanism is conceptually simple, but does not get the same accuracy performance as the more complex multiplicative weights mechanism. On the flip side the Gaussian mechanism is far faster, and does not run in time relative to $|\mathcal{X}|$ which can often be exponential in the size of a natural dataset representation.

2 The Projection Mechanism

Here we study the Projection Mechanism. This is a conceptually very simple mechanism, as it is a slight modification of the Gaussian Mechanism, and has accuracy comparable to that of the Multiplicative Weights Mechanism. But there is a caveat: the error guarantee we give will no

longer be in the worst case but instead on average. This turns out not to be a huge burden, since there are methods that can turn average error guarantees into worst case guarantees, by utilizing the technique of boosting from machine learning.

The Projection Mechanism is also interesting because it and its analysis make use of the geometry of the query space. What we mean by this is that we view $\{Q(X) : X \in \mathcal{X}^n\} \subseteq \mathbb{R}^k$ as a set of points in k dimensional space, and the vector of query answers on some dataset is just a point in this k-dimensional set. We utilize the geometric properties of this set of feasible query answers in the mechanism and its analysis. This geometric analysis also allows us to get improved error guarantees for "nice" query sets Q.

Finally, the Projection Mechanism is appealing because it outputs query answers that are consistent with some dataset, although usually not with the actual private dataset. This makes the answers easier to interpret.

2.1 Specifying the Mechanism

Before we introduce the Projection Mechanism we need the following definitions:

Definition 1. *The convex hull of a set $S \subseteq \mathbb{R}^k$ is the minimal (with respect to inclusion) convex set which contains S. A convex set is a set where the line segment connecting any two points in the set is itself entirely contained in the set.*

For a finite set of points S, the convex hull can also be equivalently defined as $\mathrm{conv}(S) = \{\sum_{x\in S} \alpha_x x : \forall x\ \alpha_x \geq 0, \sum_{x\in S} \alpha_x = 1\}$. *For an infinite set S, we can take* $\mathrm{conv}(S)$ *to be the closure of* $\{\sum_{x\in S} \alpha_x x\}$ *over all* $(\alpha_x)_{x\in S}$ *with finite support that satisfy* $\alpha_x \geq 0$ *for all* x *and* $\sum_{x\in S} \alpha_x = 1$, *as above.*

Let $S_Q = \{Q(\{x\}) : x \in \mathcal{X}\}$, i.e., all possible query answers on single point datasets. For any size n dataset X, we get

$$Q(X) = \frac{1}{n} \sum_{x\in X} Q(\{x\}) \in \mathrm{conv}(S_Q).$$

Moreover, it is not hard to see that $\mathrm{conv}(S_Q)$ is the closure of the set of all possible answer vectors $Q(X)$ on all possible data sets X of all possible sizes n. Let us use the notation $K_Q = \mathrm{conv}(S_Q)$ from now on.

We can now define the Projection Mechanism.

Algorithm 1 Projection Mechanism, with D.B. x and parameters $\varepsilon, \delta, \alpha$:

1: Let $\widetilde{Y} = Q(X) + Z$, where $Z \sim \mathcal{N}(0, \frac{k}{n^2\rho_{\varepsilon,\delta}} \cdot I)$ (I is the identity matrix and $\rho_{\varepsilon,\delta} \approx \frac{\varepsilon^2}{\log\frac{1}{\delta}}$)

2: Return $\widehat{Y} = \arg\min\{\|\widetilde{Y} - y\|_2 : y \in K_Q\}$

The first line is the (ε, δ)-differentially private Gaussian Mechanism. This mechanism may push the answer vector far away from $Q(X)$, and it may end up outside the set of feasible answers K_Q. The second line simply projects the answer back onto the closest point in K_Q. If the first line did not move the answer out of K_Q, then the second line does nothing. We will see that, in a certain sense, this second step can only improve the accuracy, and often does so significantly.

2.2 Privacy of the Mechanism

Theorem 2. *The projection mechanism is (ε, δ)-differentially private.*

The proof of this is straightforward. As shown in a previous lecture, the Gaussian Mechanism with the above parameters is (ε, δ)-differentially private. The actual projection step is simply post processing: it only relies on publicly available information, and thus it is $(0, 0)$-differentially private. By composition, the Projection Mechanism is (ε, δ)-differentially private.

2.3 Accuracy of the Mechanism

We first formalize the notion of average error.

Definition 3. *A mechanism $\mathcal{M}$ has average error (equivalently, root mean squared error) α if for all datasets X of size n we have*

$$\sqrt{\mathbb{E}\frac{1}{k}\sum_{i=1}^{k}(\mathcal{M}(X)_i - q_i(X))^2} \leq \alpha.$$

This can be equivalently formulated as

$$\sqrt{\mathbb{E}\frac{1}{k}\|\mathcal{M}(X) - Q(X)\|_2^2} \leq \alpha.$$

Above, the expectation is taken over the randomness of the mechanism.

If the average error is small, it is still possible that for some queries the error is quite high, while on others it is very low. In contrast, for worst case error error analysis we required each of the k queries be close to what the mechanism returned. Rather than measuring the error in expectation, we can ask for a high probability guarantee, and introduce another parameter β for the probability that the average error is larger than α. For the projection mechanism this would increase the lower bound on n by a $\log \frac{1}{\beta}$ term. We do not pursue this further, and instead stick with expectation to avoid dealing with yet another parameter.

We now state our first result regarding the accuracy of the Projection Mechanism.

Theorem 4. *The Projection Mechanism has average error at most α as long as $n \gg \frac{\sqrt{\log|\mathcal{X}|}\cdot\sqrt{\log\frac{1}{\delta}}}{\alpha^2\cdot\varepsilon}$.*

To show this, we will prove a refined average error guarantee that relies on the geometry of K_Q, and in particular on its size. The less precise result in Theorem 4 is useful for comparison with other mechanisms, like the Multiplicative Weights Mechanism. Indeed, the bound on n is better than the one for the Multiplicative Weights by a $\log k$ factor, albeit holding only for average error.

Before we state the refined geometric bound, we need to define some measure of the "size" of K_Q. To do so we first introduce the notion of a support function:

Definition 5. *The support function $h_K : \mathbb{R}^k \to \mathbb{R}$ of a set $K \subseteq \mathbb{R}^k$ is defined by $h_K(y) = \sup\{\langle x, y\rangle : x \in K\}$.*

For some intuition, we mention that when y is of unit Euclidean norm, i.e. $\|y\|_2 = 1$, then $h_K(y) + h_K(-y)$ is its *width* in the direction of y. I.e. it is the smallest w so that we can sandwich K between two parallel hyperplanes, both orthogonal to y, and distance w apart.

The support function satisfies the following properties for all k-dimensional vectors x and y:

1. $\forall t \geq 0$, $h_K(t \cdot y) = t \cdot h_K(y)$
2. $h_K(x + y) \leq h_K(x) + h_K(y)$
3. if $K \subseteq L$, then for any y we have $h_K(y) \leq h_L(y)$.

Since $K_Q = \text{conv}(S_Q)$, we can write $h_{K_Q}(y)$ as $\max\{\langle x, y\rangle : x \in S_Q\}$, where $y \in \mathbb{R}^k$.

Now we can introduce a way to measure the average width of K_Q . We first define the mean width:

Definition 6. *The mean width of a convex set K is $M^*(K) = \mathbb{E}h_K(Y)$, where Y is chosen uniformly at random from the set of vectors where y with $\|y\|_2 = 1$. I.e., y is chosen according to the unique rotationally invariant probability measure on the unit sphere centered about the origin.*

We also define the Gaussian width as:

Definition 7. *The Gaussian width of a convex set K is $\ell^*(K) = \mathbb{E}h_K(G)$, where $G \sim \mathcal{N}(0, I)$ and I is the k-dimensional identity matrix. That is G is chosen according to the standard k dimensional Gaussian centered about the origin.*

It turns out the Gaussian width and the mean width are closely related: we have $\ell^*(K) = (\mathbb{E}\|G\|_2) \cdot M^*(K) = c_k \cdot M^*(K)$ where $c_k = \frac{\sqrt{2}\Gamma\left(\frac{k+1}{2}\right)}{\Gamma\left(\frac{k}{2}\right)} = \Theta(\sqrt{k})$.

We can now state our refined accuracy measure, which is in terms of the Gaussian width, and thus also the mean width.

Theorem 8. *The Projection Mechanism has average error α if $n \gg \frac{\ell^*(K_Q)\sqrt{\log \frac{1}{\delta}}}{\sqrt{k}\alpha^2\varepsilon}$, or, equivalently, $n \gg \frac{M^*(K_Q)\sqrt{\log \frac{1}{\delta}}}{\alpha^2\varepsilon}$.*

Now that we have the geometric bound, we can prove Theorem 4. To do so we will show that $\ell^*(K_Q) \lesssim \sqrt{k \log |\mathcal{X}|}$. Plugging this into Theorem 8 proves Theorem 4.

We need the following three facts:

1. For any query q_i and data point x, $q_i(x) \in [0, 1]$, so $S_Q \subseteq [0, 1]^k$. Thus, we get that $\forall y \in S_Q, \|y\|_2 \leq \sqrt{k}$.
2. The Gaussian moment generating function: for any $y \in \mathbb{R}^k$, and a standard Gaussian $G \sim N(0, I)$, $\mathbb{E}[e^{\langle y, G\rangle}] = e^{\|y\|_2^2/2}$.
3. Jensen's inequality: for any random variable $A \in \mathbb{R}$, and any concave function $f : \mathbb{R} \to \mathbb{R}$, $\mathbb{E}[f(A)] \leq f(\mathbb{E}[A])$.

We have, for any $\lambda > 0$,

$$\begin{aligned}
\ell^*(K_Q) = \mathbb{E}\, h_{K_Q}(G) &= \mathbb{E} \max\{\langle y, G\rangle : y \in K_Q\} \\
&= \mathbb{E} \max\{\langle y, G\rangle : y \in S_Q\} \\
&= \mathbb{E}\, \frac{1}{\lambda} \ln \max\{e^{\lambda\langle y,G\rangle} : y \in S_Q\} \\
&\le \frac{1}{\lambda} \ln \mathbb{E} \max\{e^{\lambda\langle y,G\rangle} : y \in S_Q\} \\
&\le \frac{1}{\lambda} \ln \sum_{y\in S_Q} \mathbb{E}\, e^{\lambda\langle y,G\rangle} \\
&= \frac{1}{\lambda} \ln \sum_{y\in S_Q} e^{\lambda^2\|y\|_2^2/2} \\
&\le \frac{1}{\lambda} \ln(|\mathcal{X}| e^{\lambda^2 k/2}) = \frac{1}{\lambda} \ln|\mathcal{X}| + \frac{\lambda k}{2}.
\end{aligned}$$

The second line follows because K_Q is the convex hull of the points in S_Q. The third line uses the monotonicity of the logarithm to exchange max and ln. The fourth line follows from the concavity of ln and from Jensen's inequality. The fifth line follows since the maximum of any non-negative numbers is bounded by their sum, and also by linearity of expectation. The sixth line follows from the formula for the moment generating function of the Gaussian, and the seventh from the bound on the Euclidean norm of any element of S_Q.

To finish the proof, we optimize over λ. The right hand side is minimized for $\lambda = \sqrt{\frac{2\ln|\mathcal{X}|}{k}}$, and gives

$$\ell^*(K_Q) \le \sqrt{2k \ln|\mathcal{X}|}.$$

3 Geometric Analysis

In this section we prove Theorem 8, i.e., the geometric guarantee on the accuracy of the projection mechanism.

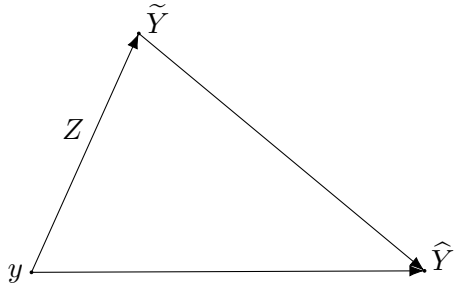

Figure 1: Vector view of query answers.

We want to show $\mathbb{E}\,\frac{1}{k}\left\|\widehat{Y} - y\right\|_2^2 \le \alpha^2$ when $n \gg \frac{\ell^*(K_Q)}{\sqrt{k}\alpha^2\rho_{\epsilon,\delta}}$. We can rewrite the distance inside the expectation as (see Figure 1):

$$\left\|\widehat{Y} - y\right\|_2^2 = \langle \widehat{Y} - y, \widetilde{Y} - y\rangle + \langle \widehat{Y} - y, \widehat{Y} - \widetilde{Y}\rangle. \tag{1}$$

We claim that

$$\langle \widehat{Y} - y, \widehat{Y} - \widetilde{Y} \rangle \leq 0. \tag{2}$$

Geometrically, this means that the triangle formed by the points y, $\widehat{Y}$, and $\widetilde{Y}$ is obtuse, and the obtuse angle is at $\widehat{Y}$ (see Figure 2).

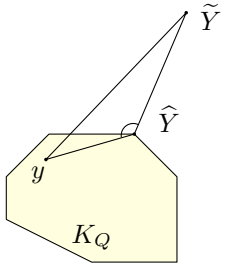

Figure 2: $\widehat{Y}$ is the projection of $\widetilde{Y}$ to the convex hull K_Q.

Now let us prove this claim, which follows from the optimality of $\widehat{Y}$ as the closest point to $\widetilde{Y}$. One way to see this is to use a bit of calculus. Let us define the function $f(\lambda) = \left\| \widetilde{Y} - (1-\lambda)\widehat{Y} - \lambda y \right\|_2^2$. For any $\lambda \in [0,1]$, the point $\widehat{Y}_\lambda = (1-\lambda)\widehat{Y} + \lambda y$ lies on the line segment between $\widehat{Y}$ and y, and, since K_Q is convex, $\widehat{Y}_\lambda$ is also in K_Q. The function $f(\lambda)$ is simply $\left\| \widetilde{Y} - \widehat{Y}_\lambda \right\|_2^2$. For $\widehat{Y}$ to be optimal, it must be the case that the minimum of $f(\lambda)$ over $\lambda \in [0,1]$ is achieved at 0, which implies

$$f'(0) = 2\langle \widehat{Y} - y, \widetilde{Y} - \widehat{Y} \rangle \geq 0 \iff \langle \widehat{Y} - y, \widehat{Y} - \widetilde{Y} \rangle \leq 0,$$

as we wanted.

Plugging (2) into (1), we get

$$\left\| \widehat{Y} - y \right\|_2^2 \leq \langle \widehat{Y} - y, \widetilde{Y} - y \rangle = \langle \widehat{Y} - y, Z \rangle.$$

We can write $Z = \frac{\sqrt{k}}{n\sqrt{\rho_{\varepsilon,\delta}}} G$, where G is a standard Gaussian (i.e., $G \sim \mathcal{N}(0, I)$ and I is the identity matrix). Using the inequality above, we have

$$\begin{aligned}
\mathbb{E}\left(\frac{1}{k} \left\| \widehat{Y} - y \right\|_2^2 \right) &\leq \mathbb{E}\left(\frac{1}{k} \langle \widehat{Y} - y, Z \rangle \right) \\
&= \frac{1}{k}\, \mathbb{E}(\langle \widehat{Y}, Z \rangle) - \frac{1}{k}\, \mathbb{E}(\langle y, Z \rangle) \\
&= \frac{1}{k}\, \mathbb{E}(\langle \widehat{Y}, Z \rangle) = \frac{1}{\sqrt{k} n \sqrt{\rho_{\varepsilon,\delta}}}\, \mathbb{E}(\langle \widehat{Y}, G \rangle) \\
&\leq \frac{1}{\sqrt{k} n \sqrt{\rho_{\varepsilon,\delta}}}\, \mathbb{E}\, h_{K_Q}(G) \\
&= \frac{1}{\sqrt{k} n \sqrt{\rho_{\varepsilon,\delta}}} \ell^*(K_Q).
\end{aligned}$$

Above, we used that $\mathbb{E}(\langle y, Z \rangle) = \langle y, \mathbb{E}(Z) \rangle = 0$, and that $\langle \widehat{Y}, G \rangle \leq \max_{x \in K_Q} \langle x, G \rangle = h_{K_Q}(G)$.

Finally, plugging in $n \gg \frac{\ell^*(K_Q)}{\sqrt{k}\alpha^2 \rho_{\epsilon,\delta}}$ gives $\mathbb{E}(\frac{1}{k} \left\| \widehat{Y} - y \right\|_2^2) \leq \alpha^2$.

Marginal Queries

Consider a database $\mathcal{X}^n$ where $\mathcal{X} \in \{0,1\}^d$ (i.e., a database of n rows of records with d binary attributes).

Definition 9. *A k-way marginal query is query over a conjunction of a subset of size k of attributes or their negations. It evaluates to one for a row, if this row satisfies this query and to zero otherwise. The answer to a marginal query on $\mathcal{X}^n$ is the fraction of rows which evaluate to one.*

For example, $c_1 \wedge \overline{c_2} \wedge c_3$ a 3-way conjunction query which evaluates to one for querying a row x if and only if $x_1 = 1, x_2 = 0$, and $x_3 = 1$ for this specific row. For example, c_1, c_2, and c_3 can respectively encode the gender, whether this person is smoking or not, and does he/she have lung cancer in each row. This specific query asks for the fraction of male people in this database who do not smoke and have lung cancer.

It can be seen that we have $\binom{d}{k}2^k$ different k-way marginal queries.

Like before, we can compute lower bounds for n when we want to have less than α error on running distinct mechanisms such as Private Multiplicative Weights (PMW) and Gaussian Noise with marginal queries:

$$n \geq \begin{cases} \frac{Cd^{k/2}\sqrt{\log 1/\delta}}{\alpha\epsilon} & Gaussian \\ \frac{C\sqrt{d}\log d\sqrt{\log 1/\delta}}{\alpha^2\epsilon} & PMW \end{cases}$$

In the above equations, C denotes a large enough *constant* which is independent from ϵ, δ and α.

We can show that the Gaussian noise mechanism's running time is $\text{poly}(n, d^k)$ and PMW's running time is $\text{poly}(n, 2^d)$.

When we are using the projection mechanism, we want to solve the following minimization problem:

$$\underset{z \in K_Q \subseteq \mathbb{R}^m}{\text{argmin}} \{\left\|\widetilde{Y} - z\right\|_2^2\}$$

Basically, it means that we are minimizing a convex function subject to a convex constraint. To do so, we need a **separation oracle**. A separation oracle for a convex set $K_Q \subseteq \mathbb{R}^m$ is an algorithm which takes a point $z \in \mathbb{R}^m$ as the input and returns *"inside"* if it is inside K_Q or otherwise it returns a hyperplane that separates K_Q and the queried point.

By having a polynomial time separation oracle for the K_Q we can solve the minimization problem of the projection mechanism. However, the structure of the K_Q polytope is too complicated when we have k-way marginal queries for $k \geq 2$. It turns out, nevertheless, that there is another convex body L such that $\frac{1}{C}L \subseteq K \subseteq L$ for an absolute constant $C > 0$, and, moreover, there is an efficient separation oracle for L. This meansthat we can run the projection mechanism so that we project on L rather than K, and we can do so efficiently. Moreover, since $\ell^*(L) \leq C\ell^*(K)$, this does not cost us more than a constant factor in terms of the error bound we can guarantee.