

Privacy in Statistics and Machine Learning
In-class Exercises for Lecture 14 (Projection Mechanism)
March 15, 2021

Spring 2021

Adam Smith and Jonathan Ullman

Problems with marked with an asterisk (*) are more challenging or open-ended.

- (Exploring Projection) If you're like us, you find it pretty magical why the projection mechanism is able to take answers that are so noisy as to seemingly contain no information about the true answers, and somehow produce very accurate answers. Here is an example that I think demystifies it a bit. Suppose we have k copies of the exact same count query f_1, \dots, f_k , so that $\mathbf{F} \in \{0, 1\}^{k \times m}$ is a matrix whose rows are all the same. Suppose you compute noisy answers $\hat{\mathbf{a}} \in \mathbb{R}^k$ with the Gaussian mechanism and project into the set of feasible answers C .
 - What is the set of consistent answers C ?
 - Given noisy answers $\hat{\mathbf{a}}$, what is the projection $\tilde{\mathbf{a}} = \Pi_C(\hat{\mathbf{a}})$?
 - What is the error of the noisy answers $\hat{\mathbf{a}}$, and the projected answers $\tilde{\mathbf{a}}$? Does it increase with k ?

- (Projection for Monotone Queries) This question explores the special case of the projection mechanism that arises from threshold queries over the domain $\mathcal{U} = \{1, \dots, D\}$. Here there are $D - 1$ non-trivial queries¹

- When $D = 3$, there are 2 non-trivial queries, and the feasible set is

$$C = \{(a_1, a_2) \in \mathbb{R}^2 : 0 \leq a_1 \leq a_2 \leq 1\}$$

Write (or illustrate) the projection operation $\Pi_C(a_1, a_2)$ for this specific set C .

- (*) For the general case, the feasible set has the form

$$C = \{(a_1, \dots, a_{D-1}) : 0 \leq a_1 \leq \dots \leq a_{D-1} \leq 1\} \tag{1}$$

Design a polynomial time algorithm for computing projection into the set C .

[**Hints:** (1) Dynamic programming! (2) It's a bit easier if we further require that all the projected answers are integer multiples of some discretization parameter $\gamma > 0$.]

- (*) Suppose the data happens to be such that all the entries of the true answer a_1, \dots, a_{D-1} are the same. Projecting a noisy version of this vector onto the set C in (1) will eliminate much of noise. How good a bound can you give on the ℓ_2 error after projection?
- (Bias and Variance) Unlike the Gaussian mechanism which is *unbiased*, meaning $\mathbb{E}(\hat{\mathbf{a}}) = \mathbf{a}$, the projection mechanism is not necessarily unbiased.

¹The query $f_D(\mathbf{x}) = \#\{i : x_i \leq D\}$ evaluates to 1 for every dataset.

- (a) Consider the case of a single count query, for which the feasible set is just $C = \{a : 0 \leq a \leq 1\}$. Compute an upper bound on the bias $\mathbb{E}(\hat{a} - a)$ of the projection mechanism on any dataset. How does the bias compare to the standard deviation of the noise?
- (b) (*) For a single count query, show that *every* (ϵ, δ) -differentially private mechanism whose output is always in the range $[0, 1]$ must have bias $\Omega(1/\epsilon n)$, at least when δ is sufficiently small.
4. (Projection as Maximum-Likelihood Estimation) A common paradigm in statistical estimation is *maximum-likelihood*. Here we are given a family of distributions $\{P_\theta\}_{\theta \in \Theta}$ and an observation y that is drawn from some distribution P_θ , and our goal is to estimate the hidden parameter θ . To do so, we're given an observation $y \sim P_\theta$ and we estimate

$$\tilde{\theta} = \arg \max_{\theta' \in \Theta} P_{\theta'}(y)$$

Show that the projection mechanism is an instance of maximum likelihood estimation. What is the family of distribution P_θ and what is $P_\theta(y)$?