CSCI4230 Computational Learning Theory Spring 2021

Lecturer: Siu On Chan *Based on Maria-Florina Balcan's notes and Graham Cormode's slides*

Notes 21: Differential privacy and Statistical Query model

1. Differential privacy from statistical query algorithms

If $\mathcal{C}$ is efficiently learnable from SQ's, then $\mathcal{C}$ is efficiently PAC-learnable, differential-privately

Theorem 1. *Suppose some algorithm A efficiently learns $\mathcal{C}$ to error ε from M statistical queries of tolerance τ. Then some algorithm B efficiently PAC-learns $\mathcal{C}$ to error ε with probability $\geqslant 1-\delta$ while satisfying α-differential privacy, using*

$$O\left(\left(\frac{M}{\alpha\tau}+\frac{M}{\tau^2}\right)\ln\frac{2M}{\delta}\right) \quad \textit{samples}$$

Proof. Algorithm B draws $O\left(\left(\frac{M}{\alpha\tau}+\frac{M}{\tau^2}\right)\ln\frac{2M}{\delta}\right)$ random samples (call them S)

Break S into M disjoint chunks $S_1,\dots,S_M$, each of size $O\left(\left(\frac{1}{\alpha\tau}+\frac{1}{\tau^2}\right)\ln\frac{2M}{\delta}\right)$

Answer i-th statistical query (φ_i,τ) of A using S_i (taking average of φ_i over S_i)

To each response, add Laplacian noise of scale $M/|S_i|\alpha$

Finally return A's hypothesis h

Privacy: Each query is the average of $|S_i|$ values, each between 0 and 1

By Theorem in Notes20, each response satisfies α/M-differential privacy

By Composition property, the collection of all M responses satisfies α-differential privacy

Error: By Hoeffding, with prob $\geqslant 1-\delta/(2M)$,

empirical average of φ_i over S_i (before adding noise) is within $\tau/2$ of the true expectation

With prob $\geqslant 1-\delta/(2M)$, the Laplace noise has magnitude

$$O\left(\frac{1}{\alpha|S_i|}\ln\frac{2M}{\delta}\right)\leqslant\frac{\tau}{2} \qquad \text{since } |S_i|\geqslant\frac{C}{\alpha\tau}\ln\frac{2M}{\delta} \text{ for some large } C$$

Hence with prob $\geqslant 1-\delta/M$, the i-th response $\hat{P}_{\varphi_i}$ is within τ of the true expectation P_{φ_i}

By union bound over all M queries, with prob $\geqslant 1-\delta$

all responses are within τ of their true averages, and algorithm A succeeds □

---

2. Geometric mechanism

When response of $\mathrm{STAT}(c,\mathcal{D})$ is integer-valued, geometric mechanism may be used

Geometric mechanism adds noise that is a (symmetric) geometric random variable

(Symmetric) **geometric distribution** with parameter $\alpha>1$ has pmf $f(k)=\alpha^{-|k|}(\alpha-1)/(\alpha+1)$

Like the Laplace distribution, symmetric geometric distribution changes by (at most) the same multiplicative factor when shifted, i.e.

$$f(k+j)/f(k)=\alpha^{-|k-j|}/\alpha^{-|k|}\leqslant\alpha^{|j|} \qquad \text{for any } j,k\in\mathbb{Z}$$

In fact the distribution is defined so that this inequality is achieved as an equality for certain j,k

If symmetric geometric noise with parameter α is added to the output of an integer-valued function g

Then the mechanism satisfies ε-differential privacy where $e^\varepsilon=\alpha^{\Delta g}$ (exercise)

Again Δg = maximum change to g's output when just one data point changes

By the same calculations as the Laplace mechanism

In practice a response of $\mathrm{STAT}(c,\mathcal{D})$ may be required to bounded, say between 0 and b

Can **truncate** the response to force it to lie in the desired range, without hurting privacy (exercise)