Algorithms for Private Data Analysis — Spring 2021
Lecture 7: Solving Selection Problems Privately

Steven Wu

1 Report Noisy Max

When the domain is finite, it is often more convenient to work with a another algorithm which behaves very similarly to the exponential mechanism. The setup is the same—we have a set of outcomes $\mathcal{Y}$ (now required to be finite) and a score function with sensitivty at most δ for each outcome. The idea is to add noise with expected magnitude Δ/ε to each item's score, independent of the number of possible outputs. The algorithm returns the *output with the highest noisy score*:

Algorithm 1: Report-Noisy-Max $A_{RNM}(\mathbf{x}, q(\cdot;\cdot), \Delta, \varepsilon)$

Input: Assume that $q(y;\cdot)$ is Δ-sensitive for every $y \in \mathcal{Y}$, and $\mathcal{Y} = \{1, ..., d\}$ is finite

Select $Z_1, ..., Z_d \sim \mathsf{Exp}(2\Delta/\varepsilon)$ i.i.d. ;

return $\arg\max_{y\in\{1,...,d\}} \left(q(y;\mathbf{x}) + Z_y\right)$;

The distribution being used to generate noise is the *exponential distribution* $\mathsf{Exp}(\lambda)$, a distribution over the nonegative real numbers $[0, +\infty)$ with density $h_\lambda(y) = \frac{1}{\lambda}\exp(-y/\lambda)$.

This algorithm is generally much easier to implement than the exponential mechanism, since it does not require explicitly computing any probabilities and can make use of standard libraries for sampling from the exponential distribution. It satisfies a very similar guarantee to the exponential mechanism.

Exercise 1.1. Show that report noisy max is ε-differentially private. [*Hint:* Consider two outputs a, b. For a fixed input $\mathbf{x}$, what is $\frac{P(a|\mathbf{x})}{P(b|\mathbf{x})}$?]

In addition to making implementation easier, the utility analysis of report-noisy-max is more intuitive. We just need to bound the probability that all d noise random variables are small:

Lemma 1.2 (Tail Bounds for Exponential Distributions)**.**

1. *If* $Z \sim \mathsf{Exp}(\lambda)$*, then* $\Pr(Z \geq t\lambda) = e^{-t}$ *for all* $t \geq 0$.
2. *If* $Z_1, ..., Z_d \sim \mathsf{Exp}(\lambda)$ *i.i.d., and* $Z_{\max} = \max_{i=1}^d Z_i$ *then* $\Pr(Z_{\max} > \lambda(\ln(d) + t)) = e^{-t}$ *for all* $t \geq 0$*, and* $\mathbb{E}(Z_{\max}) \leq \lambda(\ln(d) + 1)$.

Note that if Y_i are independent Laplace random variables with $Y_i \sim \mathsf{Lap}(\mu_i, \lambda_i)$ and $Z_i = |Y_i - \mu_i|$, then the Z_i's will be exponentially distributed with parameter λ and so Lemma 1.2 above applies.

Proof. The first part follows from a direct computation of the CDF:

$$\Pr(Z > \lambda t) = \int_{y\geq\lambda t} \tfrac{1}{\lambda} e^{-y/\lambda} dy = \tfrac{1}{\lambda}\left[-\lambda e^{-y/\lambda}\right]_{y=\lambda t}^{\infty} = e^{-t}.$$

The second part follows by a union bound: the probability that any particular Z_i exceeds $\lambda(\ln(d) + t)$ is $\frac{e^{-t}}{d}$ by part 1, so the probability that any of the Z_i's exceeds the bound is at most e^{-t}. The expectation calculation is essentially the same as in the proof the exponential mechanism's utility (Proposition **??**). □

We can also use Lemma 1.2 to prove the following, which is essentially identical to what we proved about the exponential mechanism.

Theorem 1.3. *If $q(y;\cdot)$ is Δ-sensitive for every $y \in \{1,\ldots,d\}$, then for every data set $\mathbf{x}$ in $\mathcal{X}^n$ and every $t > 0$, the output of report-noisy-max $Y \leftarrow A_{RNM}(\mathbf{x}, score, \Delta, \varepsilon)$ satisfies*

$$\Pr\left(q_{max}(\mathbf{x}) - q(Y,\mathbf{x}) \geq \frac{2\Delta(\ln(d)+t)}{\varepsilon}\right) \leq e^{-t}\,, \text{ where } q_{max}(\mathbf{x}) = \max_{y=1}^{d} q(y;\mathbf{x})\,,$$

and

$$\mathbb{E}\left(q_{max}(\mathbf{x}) - q(Y,\mathbf{x})\right) \leq \frac{2\Delta(\ln(d)+1)}{\varepsilon}\,.$$

Exercise 1.4. Prove Theorem 1.3.

Additional Reading

- McSherry and Talwar's paper that defined the exponential mechanism [MT07]
- The "Permute-and-Flip" mechanism [MS20] is an equivalent algorithm to Report-Noisy-Max [Ste20]. McKenna and Sheldon [MS20] argue that the algorithm is optimal among a natural class of selection algorithms.
- Further optimizations on the exact privacy cost of the exponential mechanism can be found in [DDR20, DWX+20].

Acknowledge This lecture note is built on the note written by Adam Smith and Jonathan Ullman.

References

[DDR20] Jinshuo Dong, David Durfee, and Ryan Rogers. Optimal differential privacy composition for exponential mechanisms. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*. PMLR, 2020.

[DWX+20] Zeyu Ding, Yuxin Wang, Yingtai Xiao, Guanhong Wang, Danfeng Zhang, and Daniel Kifer. Free gap estimates from the exponential mechanism, sparse vector, noisy max and related algorithms. *arxiv [CoRR]*, abs/2012.01592, 2020.

[MS20] Ryan McKenna and Daniel R. Sheldon. Permute-and-flip: A new mechanism for differentially private selection. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.

[MT07] Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *IEEE Symposium on Foundations of Computer Science*, FOCS '07, 2007.

[Ste20] Thomas Steinke. Personal communication, November 2020.