Benefits and Pitfalls of the Exponential Mechanism

with Applications to Hilbert Spaces and Functional PCA

Jordan Awan
Ana Kenney, Matthew Reimherr, Aleksandra Slavković

Department of Statistics, Penn State University

Thirty-sixth International Conference on Machine Learning
Long Beach, CA

Differential Privacy

awan@psu.edu

Definition (DMNS06)

A privacy mechanism $\{\mu_X : X \in \mathcal{X}^n\}$ satisfies ϵ-Differential Privacy (ϵ-DP) if for all measurable B and adjacent $X, X' \in \mathcal{X}^n$,

$$\mu_X(B) \leq \mu_{X'}(B) \exp(\epsilon).$$

- Distribution of outputs does not change much if the input changes in one entry

Exponential Mechanism [MT07]

- Given an objective function $\xi_X : \mathcal{Y} \to \mathbb{R}$ for any $X \in \mathcal{X}^n$
- The Exponential Mechanism samples $\widetilde{b}$ from the density

$$f_X(b) \propto \exp\left[\left(\frac{\epsilon}{2\Delta}\right)\xi_X(b)\right]$$

and satisfies ϵ-DP.

Utility of Exponential Mechanism

awan@psu.edu

Background
Utility
Extensions
References

Theorem

Let $(X_i)_{i=1}^{\infty}$ *such that* $X_i \in \mathcal{X}$. *Define* $\xi_n(b) := \xi_{X_1,\ldots,X_n}(b)$ *for any* $b \in \mathbb{R}^p$. *Assume that*

- $-\frac{1}{n}\xi_n$ *is twice differentiable,* α*-strongly convex, and has constant sensitivity* Δ
- *the minimizers* $\hat{b}$ *converge to some* b^*
- $-\frac{1}{n}\xi''(\hat{b}) \to \Sigma$, *a positive definite matrix*

Then,

$$\sqrt{n}(\widetilde{b} - \hat{b}) \xrightarrow{d} N_p\left(0, \left(\frac{2\Delta}{\epsilon}\right)\Sigma\right)$$

Consequences

- Large class of objective functions
- Noise introduced by Exp Mech is asymptotically normal and $O(1/\sqrt{n})$.
- Same order as statistical estimation error
- Results in increased asymptotic variance compared to non-private estimator
- Unifies the results of [WZ10, WFS15, FGWC16]

Extensions to Hilbert Spaces

- Require non-trivial base measure. Propose Gaussian process
- Give analogous utility result in infinite-dimensional spaces. GP must be chosen carefully.
- Apply Exp Mech to release DP functional principal components, extending [CSS13]

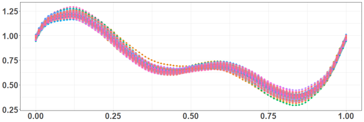

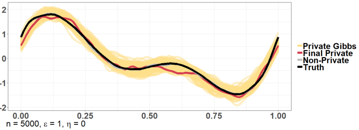

Thank You!

- NSF Grant SES-1534433
- NSF Grant DMS-1712826
- NIH Grant UL1 TR002014
- NIH Grant 5T32LM012415-03

Awan, J., Kenney, A., Reimherr, M., Slavković A. (2019) "**Benefits and Pitfalls of the Exponential Mechanism with Applications to Hilbert Spaces and Functional PCA.**" Proceedings of the 36th International Conference on International Conference on Machine Learning. arXiv:1901.10864.

Poster #179

References

[CSS13] Kamalika Chaudhuri, Anand D. Sarwate, and Kaushik Sinha. A near-optimal algorithm for differentially-private principal components. *Journal of Machine Learning Research*, 14(1):2905–2943, January 2013.

[DMNS06] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. *Calibrating Noise to Sensitivity in Private Data Analysis*, pages 265–284. Springer Berlin Heidelberg, Berlin, Heidelberg, 2006.

[FGWC16] James Foulds, Joseph Geumlek, Max Welling, and Kamalika Chaudhuri. On the theory and practice of privacy-preserving bayesian data analysis. *arXiv preprint arXiv:1603.07294*, 2016.

[MT07] Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '07, pages 94–103, Washington, DC, USA, 2007. IEEE Computer Society.

[WFS15] Yu-Xiang Wang, Stephen E. Fienberg, and Alexander J. Smola. Privacy for free: Posterior sampling and stochastic gradient monte carlo. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37*, ICML'15, pages 2493–2502. JMLR.org, 2015.

[WZ10] Larry Wasserman and Shuheng Zhou. A statistical framework for differential privacy. *JASA*, 105:489:375–389, 2010.