KNG: The K-Norm Gradient Mechanism

Matthew Reimherr and **Jordan Awan**
Department of Statistics, Pennsylvania State University, University Park, PA

Overview

Objective: New privacy mechanism for producing efficient point estimates for empirical risk problems

Contributions:

- Propose the K-Norm Gradient Mechanism (KNG)
- Prove KNG introduces $O_p(1/n)$ noise, asymptotically negligible compared to estimation error
- Demonstrate that KNG unifies popular mechanisms such as Laplace, K-norm, and PrivateQuantile
- Develop the first DP mechanism for quantile regression
- Illustrate through simulations the performance gain of KNG over exponential mechanism for linear and quantile regression

Differential Privacy: [DMNS06]

- Let $\mathcal{X}$ be a set, $\mathcal{P} = \{P_{\underline{x}} \mid \underline{x} \in \mathcal{X}^n\}$ be a set of probability measures
- $\mathcal{P}$ satisfies **ϵ-Differential Privacy** (ϵ - DP) if for all measurable setws B and all $\underline{x}, \underline{x}'$ differing in one entry, we have $P_{\underline{x}}(B) \leqslant e^{\epsilon} P_{\underline{x}'}(B)$

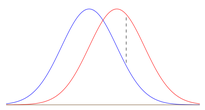

KNG, Exponential, Objective Perturbation

Suppose that $\ell_n(\theta; D)$ is an objective function.

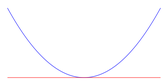

Exponential mechanism [MT07] draws from

$$f_E(\theta) \propto \exp\{-c_1 \ell_n(\theta; D)\}.$$

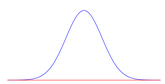

Our proposed KNG mechanism draws from

$$f_n(\theta) \propto \exp\{-c_2 \|\nabla \ell_n(\theta; D)\|_K\}.$$

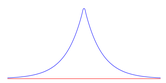

Similar to *objective perturbation* [CMS11]:

$$\tilde{\theta}_O = \text{argmin}_{\theta \in \Theta} \left(\ell_n(\theta; D) + c_3^{-1} \theta^\top b \right),$$

or equivalently,

$$f_O(\theta) \propto \exp\{-c_3 \|\nabla \ell_n(\theta)\|_K\} |\det(\nabla^2 \ell_n(\theta))|.$$

Limitations of Previous Mechanisms

- Exponential mechanism is often asymptotically inefficient[AKRS19]
- In exponential mechanism, difficult to bound sensitivity
- Objective perturbation has many assumptions on objective function

Theorem (KNG)

Let $\{\ell_n(\theta; D) : \Theta \to \mathbb{R} \mid D \in \mathcal{D}^n\}$ *be differentiable a.e. If* $\exists \Delta : \Theta \to \mathbb{R}$ *s.t.*

$$\|\nabla \ell_n(\theta; D) - \nabla \ell_n(\theta; D')\|_K \leqslant \Delta(\theta) < \infty,$$

then the mechanism $\{\mu_D \mid D \in \mathcal{D}\}$*, with densities*

$$f_D(\theta) \propto \exp\left[\left(\frac{-\epsilon}{2\Delta(\theta)} \right) \|\nabla \ell_n(\theta; D)\|_K \right] \quad \textit{satisfies } \epsilon\textit{-DP.}$$

Intuition

- Use the normed gradient to obtain a distribution similar to Laplace or K-norm
- Exponential mechanism results in distributions more like Guassian
- Similar to objective perturbation, but with fewer assumptions. Applicable to more problems

Theorem (Utility of KNG)

Let $\ell_n(\theta) := \ell_n(\theta; D)$ *have sensitivity* $\Delta(\theta)$*. Assume*

1. $n^{-1}\ell_n(\theta)$ *are twice differentiable (a.e.)* α*-strongly convex functions*
2. *the minimizers satisfy* $\hat{\theta} \to \theta^\star \in \mathbb{R}^d$ *and* $n^{-1}\mathbf{H}_n(\hat{\theta}) \to \Sigma^{-1}$ *where* Σ *is p.d.*
3. $\Delta(\theta)$ *is continuous in* θ *and constant in* n

Then the density of $Z = n(\tilde{\theta} - \hat{\theta})$ *converges to a K-norm distribution, with density proportional to* $f(z) \propto \exp\left(\frac{-\epsilon}{2\Delta(\theta^*)} \|\Sigma^{-1} z\|_K \right)$.

Interpretation

- KNG introduces only $O_p(1/n)$ noise, asymptotically negligible
- Exponential Mechanism introduces $O_p(1/\sqrt{n})$
- Asymptotically, number of samples for KNG same as for MLE
- Theorem assumptions can likely be weakened

Connection to other mechanisms

- Mean Estimation: $\ell_n(\theta; D) = \sum_{i=1}^n \|x_i - \theta\|_2^2$. KNG recovers the K-norm mechanism [HT10], a generalization of Laplace
- Quantile Estimation: $\ell_n(\theta; D) = \sum_{i=1}^n \rho_\tau(y_i - \theta)$, where $\rho_\tau(y) = y(\tau - I_{(y<0)})$. KNG recovers PrivateQuantile [Smi11].

Simulations

- Linear Regression: $\ell(\theta; D) = \sum_{i=1}^n (y_i - x_i^\top \theta)^2$.
- Quantile Regression: $\ell_n(\theta; D) = \sum_{i=1}^n \rho_\tau(y_i - x_i^\top \theta)$

For error $= Cn^{-1/2}$, $\quad \log(\text{error}) = (-1/2)\log(n) + \log(C)$,

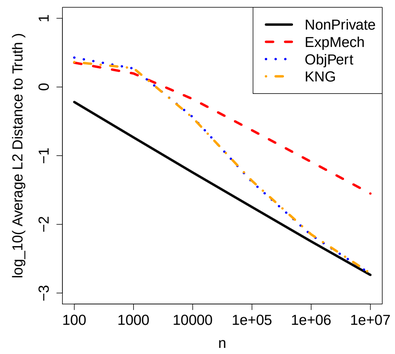

Figure: Simulation comparing the non-private MLE, exponential mechanism, objective perturbation, and KNG for linear regression.

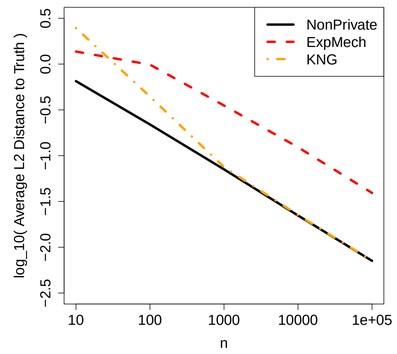

Figure: Simulation comparing the non-private, exponential mechanism, and KNG for quantile regression.

Discussion

- Principled approach to develop practical and efficient mechanisms
- First efficient mechanism for quantile regression
- Sampling is non-trivial, often resort to MCMC

References

[AKRS19] Jordan Awan, Ana Kenney, Matthew Reimherr, and Aleksandra Slavković. Benefits and pitfalls of the exponential mechanismwith applications to hilbert spaces and functional pca. In *Proceedings of the 36th International Conference on International Conference on Machine Learning*, ICML'19, pages 374–384. JMLR.org, 2019.

[CMS11] Kamalika Chaudhuri, Claire Monteleoni, and D. Sarwate. Differentially private empirical risk minimization. In *Journal of Machine Learning Research*, volume 12, pages 1069–1109, 2011.

[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.

[HT10] Moritz Hardt and Kunal Talwar. On the geometry of differential privacy. In *Proceedings of the Forty-second ACM Symposium on Theory of Computing*, STOC '10, pages 705–714, New York, NY, USA, 2010. ACM.

[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.

[Smi11] Adam Smith. Privacy-preserving statistical estimation with optimal convergence rates. In *Proceedings of the Forty-third Annual ACM Symposium on Theory of Computing*, STOC '11, pages 813–822, New York, NY, USA, 2011. ACM.

NSF Awards SES-1534433, SES-1853209, and DMS-1712826 to The Pennsylvania State University

awan@psu.edu