Lecture 15 Propose-Test-Release

Yu-Xiang Wang

COMPUTER SCIENCE
UC SANTA BARBARA

Computing. ReInvented.

Logistics

- Please submit your HW2.

- The coding part should be pretty easy given my template.
  - Let me know if you run into troubles.

- HW3 will be light-weighted so you have time to work on your project.

Recap: Beyond worst-case noise in DP query release

- Global sensitivity
- Local sensitivity
$$\mathrm{LS}_q(x) = \max\left\{q(x) - q(x')) | : x' \sim x\right\}.$$
- Smooth sensitivity

Recap: Admissible noise

Notation. For a subset $\mathcal{S}$ of $\mathbb{R}^d$, we write $\mathcal{S} + \Delta$ for the set $\{z + \Delta \mid z \in \mathcal{S}\}$, and $e^\lambda \cdot \mathcal{S}$ for the set $\{e^\lambda \cdot z \mid z \in \mathcal{S}\}$. We also write $a \pm b$ for the interval $[a-b, a+b]$.

Definition 2.5 (Admissible Noise Distribution). *A probability distribution on $\mathbb{R}^d$, given by a density function h, is (α, β)-*admissible *(with respect to ℓ_1) if, for $\alpha = \alpha(\epsilon, \delta), \beta = \beta(\epsilon, \delta)$, the following two conditions hold for all $\Delta \in \mathbb{R}^d$ and $\lambda \in \mathbb{R}$ satisfying $\|\Delta\|_1 \leq \alpha$ and $|\lambda| \leq \beta$, and for all measurable subsets $\mathcal{S} \subseteq \mathbb{R}^d$:*

Sliding Property: $$\Pr_{Z\sim h}\Big[Z \in \mathcal{S}\Big] \leq e^{\frac{\epsilon}{2}} \cdot \Pr_{Z\sim h}\Big[Z \in \mathcal{S} + \Delta\Big] + \frac{\delta}{2}.$$

Dilation Property: $$\Pr_{Z\sim h}\Big[Z \in \mathcal{S}\Big] \leq e^{\frac{\epsilon}{2}} \cdot \Pr_{Z\sim h}\Big[Z \in e^\lambda \cdot \mathcal{S}\Big] + \frac{\delta}{2}.$$

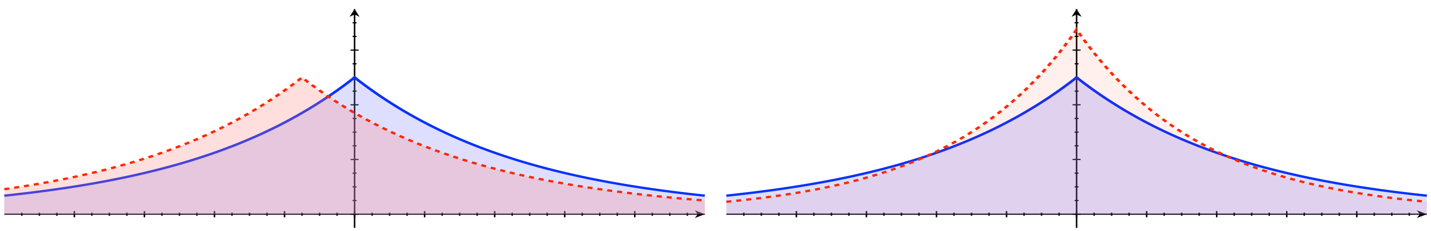

Figure 1: Sliding and dilation for the Laplace distribution with p.d.f. $h(z) = \frac{1}{2}e^{-|z|}$, plotted as a solid line. The dotted lines plot the densities $h(z + 0.3)$ (left) and $e^{0.3}h(e^{0.3}z)$ (right).

- Then $\mathcal{A}(x) = f(x) + \frac{S(x)}{\alpha} \cdot Z$ satisfies (ε,δ)-DP.

Recap: Summary of the noises that are known to work

- Cauchy distribution
- Student t-distribution

- Laplace-log-normal
- Uniform-log-normal
- Arcsinh-normal

- Gaussian
- Laplace

Recap: Laplace-log-normal noise and CDP

- Adding log-normal noise

$$Z = X \cdot e^{\sigma Y}$$

- X drawn from Laplace and Y from a standard Normal.

Proposition 3. *Let $f : \mathcal{X}^n \to \mathbb{R}$ and let $Z \leftarrow \mathsf{LLN}(\sigma)$ for some $\sigma > 0$. Then, for all $s, t > 0$, the algorithm $M(x) = f(x) + \frac{1}{s} \cdot \mathsf{S}_f^t(x) \cdot Z$ guarantees $\frac{1}{2}\varepsilon^2$-CDP for $\varepsilon = t/\sigma + e^{3\sigma^2/2} s$.*

Bun and Steinke (2019): “Average case averages”: https://arxiv.org/pdf/1906.02830.pdf

This lecture

- Finish smooth sensitivity
  - Sketching the idea of the zCDP proof for Laplace log-normal.
  - Empirical results on truncated mean.
- Propose-Test-Release
- Easy-to-use recipes for PTR and examples

Reading materials

- Vadhan book Section 3.2 – 3.4
- Dwork and Lei “Differential Privacy and Robust Statistics”
  - Original paper for PTR.
- W. (2018) “Revisiting Differentially Private Linear Regression” https://arxiv.org/abs/1803.02596
  - A good example for deriving data—dependent DP algorithm

Concentrated DP analysis of Smoothed Sensitivity

- Adding log-normal noise

$$Z = X \cdot e^{\sigma Y}$$

- X drawn from Laplace and Y from a standard Normal.

Proposition 3. *Let* $f : \mathcal{X}^n \to \mathbb{R}$ *and let* $Z \leftarrow \mathsf{LLN}(\sigma)$ *for some* $\sigma > 0$*. Then, for all* $s, t > 0$*, the algorithm* $M(x) = f(x) + \frac{1}{s} \cdot \mathsf{S}_f^t(x) \cdot Z$ *guarantees* $\frac{1}{2}\varepsilon^2$*-CDP for* $\varepsilon = t/\sigma + e^{3\sigma^2/2}s$*.*

Bun and Steinke (2019): "Average case averages": https://arxiv.org/pdf/1906.02830.pdf

Summary of the noises that are known to work

- Cauchy distribution
- Student t-distribution

- Laplace-log-normal
- Uniform-log-normal
- Arcsinh-normal

- Gaussian
- Laplace

Sketch of the proof for the Laplace-Log-Normal

- Let's say for all neighboring datasets

$$|f(x) - f(x')| \leq g(x) \qquad \text{and} \qquad e^{-t} g(x) \leq g(x') \leq e^t g(x).$$

- Algorithm: $M(x) = f(x) + \frac{g(x)}{s} \cdot Z \qquad \text{for} \qquad Z \leftarrow \mathsf{LLN}(\sigma).$

- We have that $\mathrm{D}_\alpha\left(M(x) \| M(x')\right) = \mathrm{D}_\alpha\left(Z \middle\| \frac{f(x') - f(x)}{g(x)} \cdot s + \frac{g(x')}{g(x)} \cdot Z\right).$

Technical tools

- Group privacy for CDP:

Lemma 11. *Let P, Q, R be probability distributions. Suppose* $\mathrm{D}_\alpha(P\|R) \le a\cdot\alpha$ *and* $\mathrm{D}_\alpha(R\|Q) \le b\cdot\alpha$ *for all* $\alpha \in (1,\infty)$*. Then, for all* $\alpha \in (1,\infty)$,

$$\mathrm{D}_\alpha(P\|Q) \le \alpha \cdot (\sqrt{a}+\sqrt{b})^2 \le 2\alpha\cdot(a+b).$$

- Decompose what we want to bound

$$\mathrm{D}_\alpha(Z\|e^t Z + s)$$

$$\mathrm{D}_\alpha(e^t Z + s\|Z)$$

Bounding the two parts separately

Lemma 19. *Let* $Z \leftarrow \mathsf{LLN}(\sigma)$ *for* $\sigma > 0$. *Let* $t \in \mathbb{R}$ *and* $\alpha \in (1, \infty)$. *Then*

$$\mathrm{D}_\alpha\left(Z \middle\| e^t Z\right) \le \frac{\alpha t^2}{2\sigma^2}.$$

- Proof:

$$\mathrm{D}_\alpha\left(Z \middle\| e^t Z\right) = \mathrm{D}_\alpha\left(Xe^{\sigma Y} \middle\| Xe^{\sigma Y + t}\right) \le \sup_x \mathrm{D}_\alpha\left(xe^{\sigma Y} \middle\| xe^{\sigma Y + t}\right) \le \mathrm{D}_\alpha\left(\sigma Y \| \sigma Y + t\right).$$

Lemma 20. *Let* $Z \leftarrow \mathsf{LLN}(\sigma)$ *for* $\sigma > 0$. *Let* $s \in \mathbb{R}$ *and* $\alpha \in (1, \infty)$. *Then*

$$\mathrm{D}_\alpha\left(Z \| Z + s\right) \le \min\left\{\frac{1}{2} e^{3\sigma^2} s^2 \alpha, e^{\frac{3}{2}\sigma^2} s\right\}.$$

- Proof:

Improvement from running smoothed sensitivity is substantial!

$$\mathrm{trim}_m(x) = \frac{x_{(m+1)} + x_{(m+2)} + \cdots + x_{(n-m)}}{n - 2m},$$

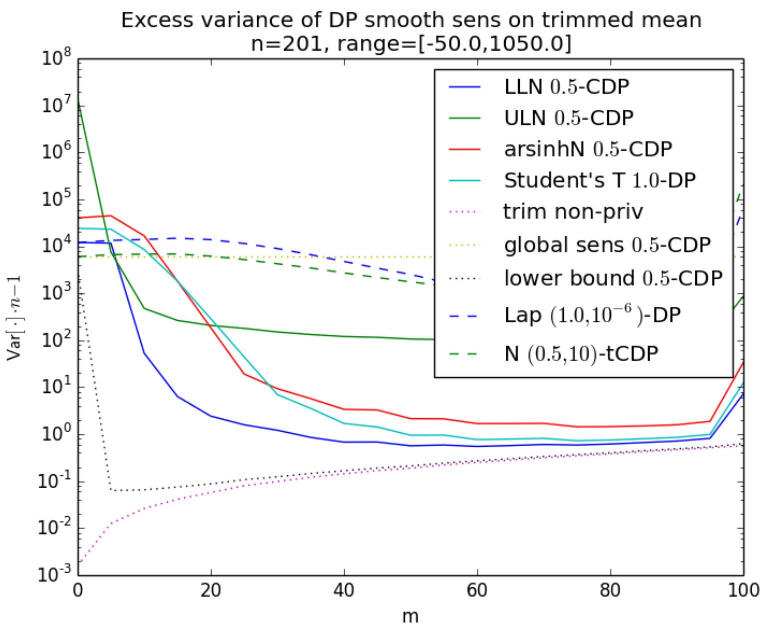

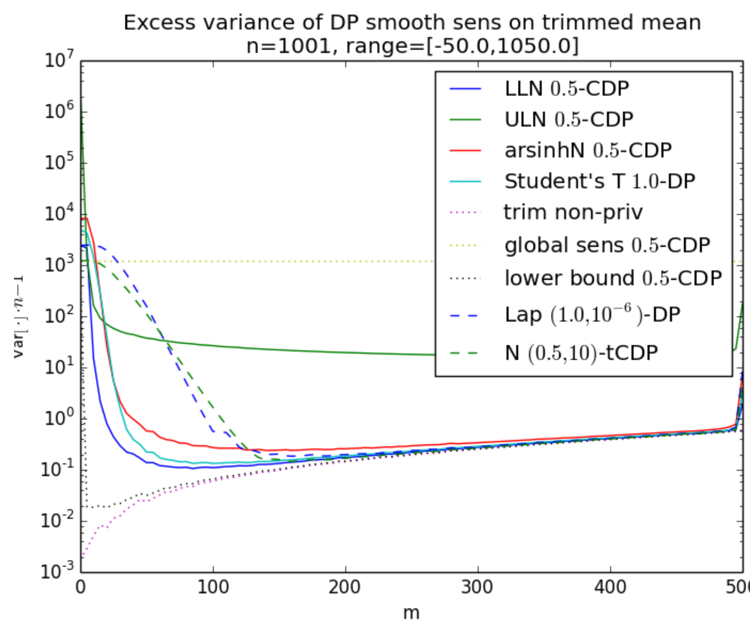

Bun and Steinke (2019): “Average case averages”: https://arxiv.org/pdf/1906.02830.pdf

Drawbacks of Smooth Sensitivity

- Restricted to numerical valued outputs.

- Requires elaborate design of the noise, generally with a much heavier tail

- Does not generalize well to high-dimension

- Are there more flexible recipes for deriving data-dependent DP algorithms?

Example: Releasing reciprocal

- Let f(D) be a counting query, define g(D) = 1/f(D)
  - What is the global and local sensitivity of g(D)?

  - What is the smooth sensitivity of g(D)?

- Example: the prediction variance of linear regression on a new dataset
  - Useful for statistical inference / uncertainty quantification

Examples: Private Argmax

- Voting: Who won the election?
- Model selection: Which is the best performing model when evaluating on a private dataset?
- Netflix: What is the Top-k most-popular movie last week?

Release Stable Values without adding noise.

Define "Dist2Instability" function:

"Dist2Instability":

1.

2.

The privacy analysis of "Dist2Instability"

- Case A:

- Case B:

Utility of “Dist2Instability”

- Perfect utility with high probability when “margin is large”

- No utility at all when the margin is small.

- Comparing to exponential mechanism
  - Homework 3 question.

Propose-Test-Release

1. Propose a bound on local-sensitivity

2. Test the validity of this bound

$$\hat{d} = d(x, \{x' : \mathrm{LS}_q(x') > \beta\}) + \mathrm{Lap}(1/\varepsilon),$$

3. Release:

Proposition 3.2 (propose-test-release [33]). *For every query $q : \mathcal{X}^n \to \mathbb{R}$ and $\varepsilon, \delta, \beta \geq 0$, the above algorithm is $(2\varepsilon, \delta)$-differentially private.*

The privacy analysis of PTR

- Case 1:

- Case 2:

Two remaining issues with PTR

1. How do I know what bound to propose?

2. Isn't it still relying on local sensitivity and noise-adding? How does it help to go beyond releasing numerical queries?

How do I know what bound to propose? Privately releasing “a high probability bound” of local sensitivity.

- Example: Estimating the number of triangles in a graph under Edge Differential Privacy.

- Global sensitivity: n-2
- Local sensitivity: the max degree of G

- Private releasing local sensitivity?

Privacy analysis of the approach to release local sensitivity privately.

Lemma: Let $\tilde{\Delta}_f(D)$ satisfies ε-DP and

$$\mathbb{P}\left[\Delta_f(D) \geq \tilde{\Delta}_f(D)\right] \leq \delta$$

Then $f(D) + \mathrm{Lap}(\tilde{\Delta}_f(D)/\epsilon)$ satisfies $(2\varepsilon,\delta)$-DP.

- Proof:

See a more general statement and proof in
Appendix G.6 of this paper: https://sites.cs.ucsb.edu/~yuxiangw/docs/spectral_privatelda.pdf

Beyond local sensitivity / noise-adding approaches

- What happens when the output space is not numerical?

- How to design data-adaptive versions of posterior-sampling, or objective-perturbation, or NoisySGD rather than just noise adding?

Topic of the next (and final) lecture

- Beyond local sensitivity
  - Per-instance differential privacy
  - pDP to DP conversion

- Data-dependent algorithms in differentially private machine learning