Flipped Huber: A new additive noise mechanism for differential privacy

Sheetal Kalyani

Joint work with Gokularam M

Department of Electrical Engineering

Prologue

Why does data privacy matter and what can one do?

- What one wants to reveal should be their choice
- However each and every mobile app collects your data
- From all the collected data one can track any individual
- Every survey, every test, every hypothesis needs data
- Data collection or release has to be privatized
- My participation in a survey should not reveal my identity

With great accuracy comes great loss of privacy

- As an algorithm becomes more and more accurate we compromise data privacy
- Data privacy is increasingly hard
- Information about an individual is available from multiple sources
- Example - apps which reveal caller ids
- Plain anonymization doesn't ensure complete protection

Algorithm output can reveal information about individual data

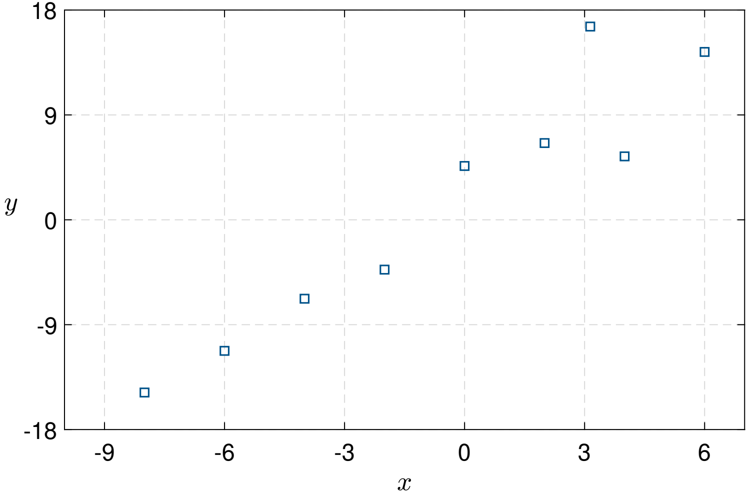

Algorithm output can reveal information about individual data

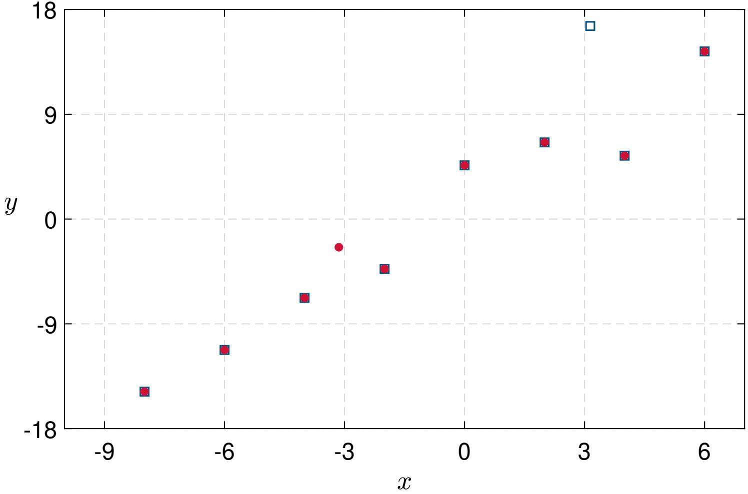

Algorithm output can reveal information about individual data

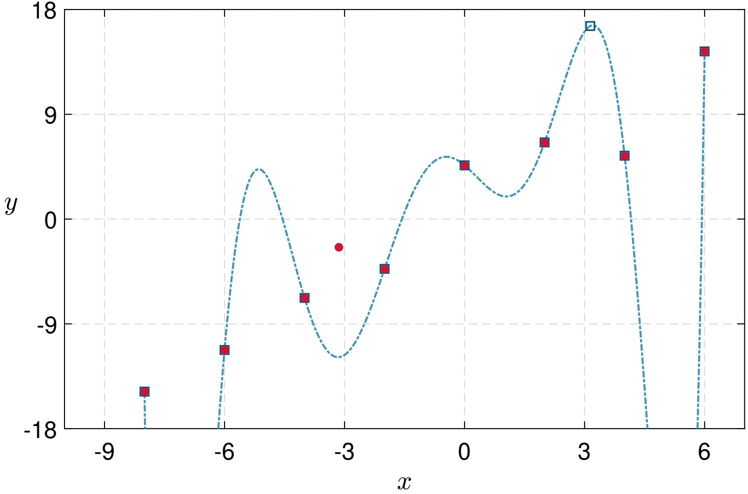

Algorithm output can reveal information about individual data

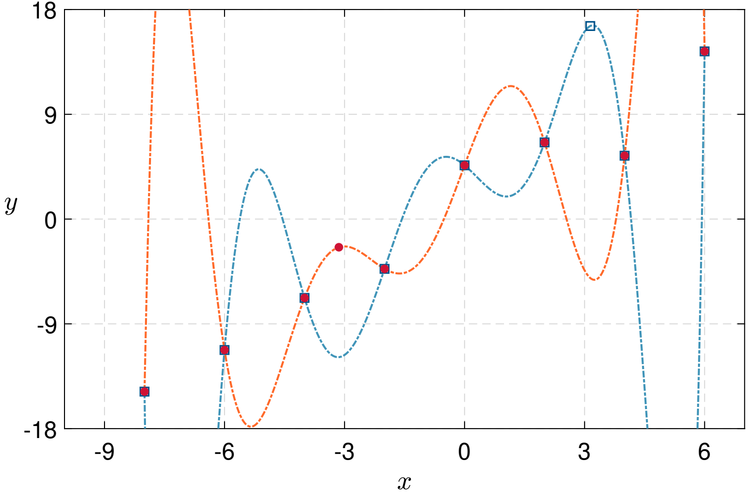

Algorithm output can reveal information about individual data

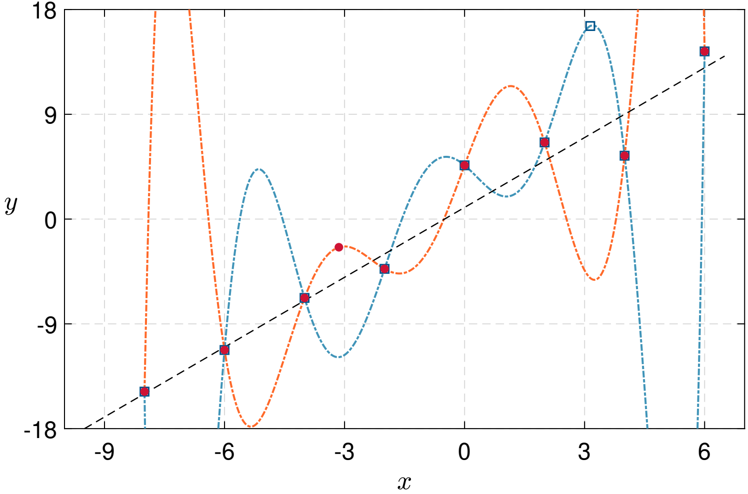

Algorithm output can reveal information about individual data

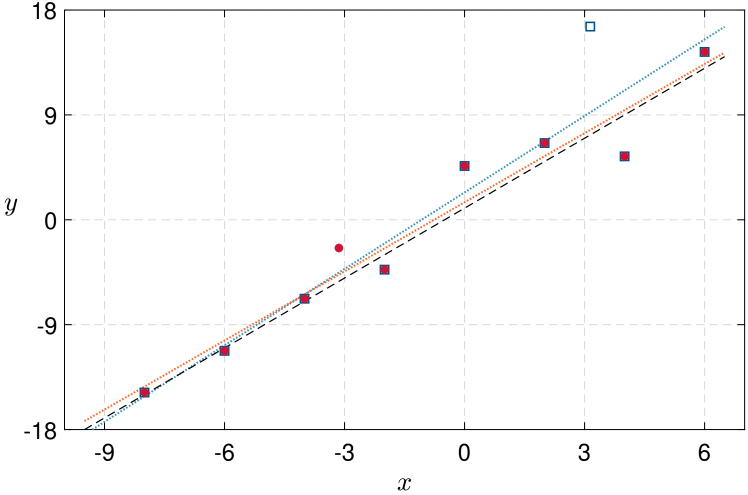

Differential privacy: Ensuring the privacy of individuals in datasets

Differential privacy (DP)

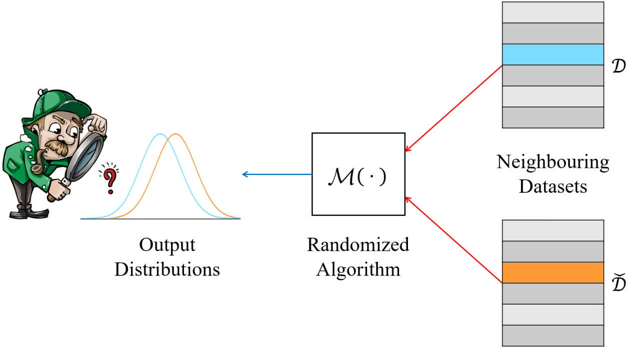

Differential privacy (DP)

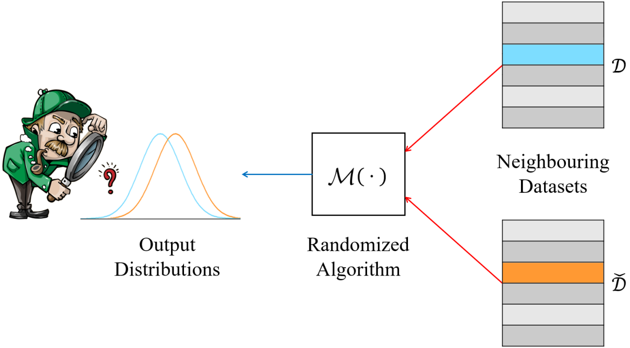

Definition: (ϵ, δ)-DP

$\mathcal{M}$ is (ϵ, δ)-DP if for every measurable $\mathcal{S} \subseteq \mathcal{Y}$ and $\mathcal{D} \asymp_{\mathcal{X}} \check{\mathcal{D}}$,

$$\mathbb{P}\{\mathcal{M}(\mathcal{D}) \in \mathcal{S}\} \leq e^{\epsilon}\, \mathbb{P}\{\mathcal{M}(\check{\mathcal{D}}) \in \mathcal{S}\} + \delta\,.$$

Additive noise mechanism

- Consider numeric vector query $f : \mathcal{X} \rightarrow \mathbb{R}^K$
  - f acts on $\mathcal{D} \in \mathcal{X}$ and provides the response $f(\mathcal{D})$
- Additive noise mechanism imparts DP by perturbing $f(\mathcal{D})$ as $\mathcal{M}(\mathcal{D}) = f(\mathcal{D}) + \mathbf{t}$
- $\mathbf{t} = [t_1 \ t_2 \ \cdots \ t_K]^\top \in \mathbb{R}^K$: noise (typically i.i.d.) sampled from known distribution
- Popular choices are Laplace[1] and Gaussian[2]

[1] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Proc. Theory Cryptogr. Conf.* Springer, 2006, pp. 265–284

[2] B. Balle and Y.-X. Wang, "Improving the Gaussian mechanism for differential privacy: Analytical calibration and optimal denoising," in *Proc. Int. Conf. Mach. Learn.* PMLR, 2018, pp. 394–403

Additive noise mechanism (cont.,)

- Sensitivity determines amount of noise $\Delta_p = \sup_{\mathcal{D} \asymp_{\mathcal{X}} \check{\mathcal{D}}} \left\| f(\mathcal{D}) - f(\check{\mathcal{D}}) \right\|_p$
- Deviation in query result: $\mathbf{d} = f(\mathcal{D}) - f(\check{\mathcal{D}})$

[3] C. Dwork, G. N. Rothblum, and S. Vadhan, "Boosting and differential privacy," in *Proc. IEEE Annu. Symp. Found. Comput. Sci.* IEEE, 2010, pp. 51–60

[4] B. Balle and Y.-X. Wang, "Improving the Gaussian mechanism for differential privacy: Analytical calibration and optimal denoising," in *Proc. Int. Conf. Mach. Learn.* PMLR, 2018, pp. 394–403

Additive noise mechanism (cont.,)

- Sensitivity determines amount of noise $\Delta_p = \sup_{\mathcal{D} \asymp_{\mathcal{X}} \check{\mathcal{D}}} \left\| f(\mathcal{D}) - f(\check{\mathcal{D}}) \right\|_p$
- Deviation in query result: $\mathbf{d} = f(\mathcal{D}) - f(\check{\mathcal{D}})$
- Privacy loss random variable[3]: $\zeta_{\mathbf{d}}(\mathbf{T}) = \log \frac{g_{\mathbf{T}}(\mathbf{t})}{g_{\mathbf{T}}(\mathbf{t}+\mathbf{d})}$
  - Additive under i.i.d. noise: $\zeta_{\mathbf{d}}(\mathbf{t}) = \sum_{i=1}^{K} \zeta_{d_i}(t_i)$
  - Centered privacy loss $\widetilde{\zeta}_d(t) = \zeta_d\big(t - \frac{d}{2}\big)$

[3] C. Dwork, G. N. Rothblum, and S. Vadhan, "Boosting and differential privacy," in *Proc. IEEE Annu. Symp. Found. Comput. Sci.* IEEE, 2010, pp. 51–60

[4] B. Balle and Y.-X. Wang, "Improving the Gaussian mechanism for differential privacy: Analytical calibration and optimal denoising," in *Proc. Int. Conf. Mach. Learn.* PMLR, 2018, pp. 394–403

Additive noise mechanism (cont.,)

- Sensitivity determines amount of noise $\Delta_p = \sup_{\mathcal{D} \asymp_{\mathcal{X}} \check{\mathcal{D}}} \left\| f(\mathcal{D}) - f(\check{\mathcal{D}}) \right\|_p$
- Deviation in query result: $\mathbf{d} = f(\mathcal{D}) - f(\check{\mathcal{D}})$
- Privacy loss random variable[3]: $\zeta_{\mathbf{d}}(\mathbf{T}) = \log \frac{g_{\mathbf{T}}(\mathbf{t})}{g_{\mathbf{T}}(\mathbf{t}+\mathbf{d})}$
  - Additive under i.i.d. noise: $\zeta_{\mathbf{d}}(\mathbf{t}) = \sum_{i=1}^{K} \zeta_{d_i}(t_i)$
  - Centered privacy loss $\widetilde{\zeta}_d(t) = \zeta_d(t - \frac{d}{2})$
- Equivalent condition for (ϵ, δ)-DP[4]: $\sup_{\mathcal{D} \asymp_{\mathcal{X}} \check{\mathcal{D}}} \mathbb{P}\{\zeta_{\mathbf{d}}(\mathbf{T}) \geq \epsilon\} - e^{\epsilon}\, \mathbb{P}\{\zeta_{-\mathbf{d}}(\mathbf{T}) \leq -\epsilon\} \leq \delta$

---

[3] C. Dwork, G. N. Rothblum, and S. Vadhan, "Boosting and differential privacy," in *Proc. IEEE Annu. Symp. Found. Comput. Sci.* IEEE, 2010, pp. 51–60

[4] B. Balle and Y.-X. Wang, "Improving the Gaussian mechanism for differential privacy: Analytical calibration and optimal denoising," in *Proc. Int. Conf. Mach. Learn.* PMLR, 2018, pp. 394–403

Proposed mechanism : Flipped Huber

What have we lost by introducing DP?

- By adding noise to the output query or by randomizing it I have lost information
- No free lunch: To get more privacy you have to sacrifice more utility or accuracy
- What is the minimum noise I should add to get desired DP with least loss of utility
- The noise added is a function of ϵ, δ and Δ_p

Is there an optimal additive mechanism?

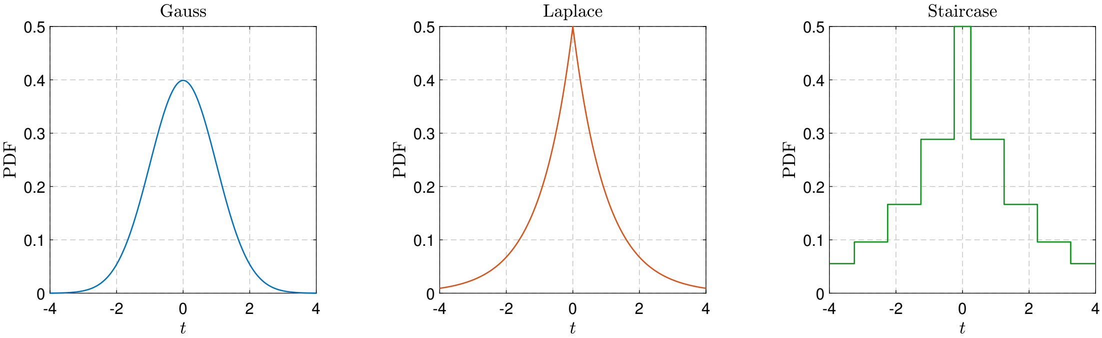

Two tale(il)s, one story

- Heavy-tailed noise is undesirable
  - Tail of $T \rightarrow$ affects accuracy
- $\mathbb{P}\{\zeta_d(T) \geq \epsilon\} \leq \delta \;\Rightarrow\; (\epsilon, \delta)$-DP
  - Tail of $\zeta_d(T) \rightarrow$ affects privacy

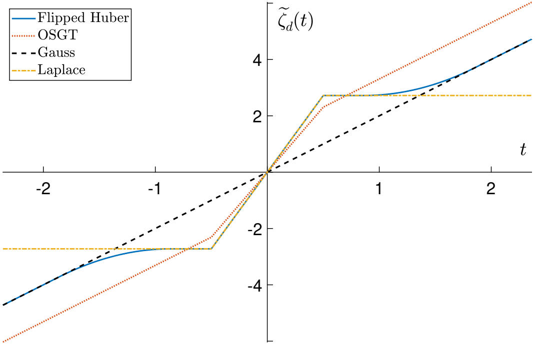

Characteristics of tails of both T and $\zeta_d(t)$ determine privacy-accuracy trade-off

What about popular existing distributions?

- Laplace noise:
  - Optimal ϵ-DP mechanism in high privacy regime
    - Bounded $\zeta_d(T)$
  - Outputs are *more informative* of the true response
  - Excessive noise for large K and results in outliers

What about popular existing distributions?

- Laplace noise:
  - Optimal ϵ-DP mechanism in high privacy regime
    - Bounded $\zeta_d(T)$
  - Outputs are *more informative* of the true response
  - Excessive noise for large K and results in outliers
- Gaussian noise:
  - Light tailed noise
  - Privacy loss is also Gaussian $\rightarrow$ light tailed $\zeta_d(T)$
    - Composes well
  - Outputs are *less informative* of the true response

Can grafting Laplace and Gaussian help?

- Noise density design to have the best of both Laplace and Gaussian

Can grafting Laplace and Gaussian help?

- Noise density design to have the best of both Laplace and Gaussian

Can grafting Laplace and Gaussian help?

- Noise density design to have the best of both Laplace and Gaussian
- *Hybridize* the densities $\rightarrow$ splice Laplace centre and Gaussian tails

Are there other grafted densities?

- Huber's distribution : Gaussian in the centre and Laplacian in the tails

$$g_{\mathcal{H}}(t) = (1-\tau) \times \begin{cases} \phi(t)\,, & |t| \leq \alpha \\ e^{-\alpha\left(|t|-\frac{\alpha}{2}\right)}\,, & |t| > \alpha \end{cases}, \qquad \tau = \left(1 + \frac{\alpha}{2(\phi(\alpha)-\alpha Q(\alpha))}\right)^{-1}$$

- Least Fisher information among symmetric distributions[5] of the form $(1-\tau)\phi(t) + \tau h(t)$

- We want actually the most favorable distribution which can satisfy the DP requirement

[5] P. J. Huber and E. M. Ronchetti, *Robust statistics*. John Wiley & Sons, 2009

Are there other grafted densities? (cont.,)

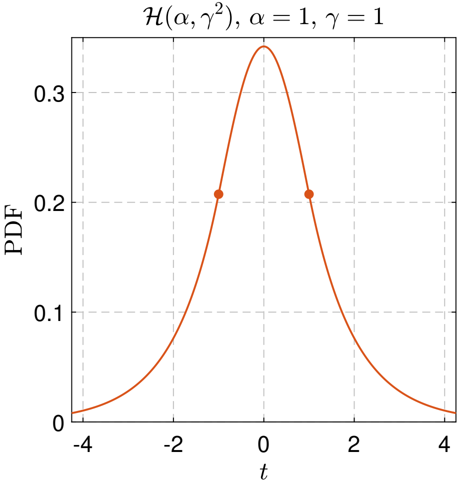

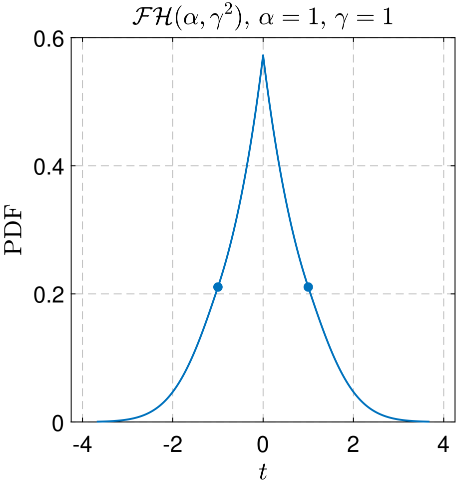

Flipped Huber distribution

- Flipped Huber loss function: $\rho_\alpha(t) = \begin{cases} \alpha|t|\,, & |t| \leq \alpha \\ (t^2 + \alpha^2)/2\,, & |t| > \alpha \end{cases}$
  - Symmetric and convex

Flipped Huber distribution

- Flipped Huber loss function: $\rho_\alpha(t) = \begin{cases} \alpha|t|\,, & |t| \leq \alpha \\ (t^2+\alpha^2)/2\,, & |t| > \alpha \end{cases}$
  - Symmetric and convex

Definition: Flipped Huber Distribution

The flipped Huber distribution $\mathcal{FH}(\alpha, \gamma^2)$ is specified by the density function,

$$g_{\mathcal{FH}}(t; \alpha, \gamma^2) = \frac{1}{\kappa} \exp\left(-\frac{\rho_\alpha(t)}{\gamma^2}\right),$$

where $\kappa = \gamma\omega e^{-\alpha^2/2\gamma^2}$ and $\omega = 2\left[\sqrt{2\pi}Q\left(\frac{\alpha}{\gamma}\right) + \frac{2\gamma}{\alpha}\sinh\left(\frac{\alpha^2}{2\gamma^2}\right)\right]$.

Flipped Huber distribution (cont.,)

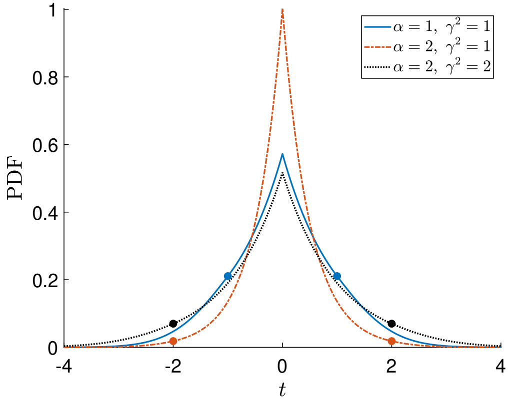

$\mathcal{FH}(\alpha, \gamma^2)$ for various choices of α and γ

Properties of flipped Huber

- Variance: $\sigma^2_{\mathcal{FH}} = \gamma^2\left[1 - \frac{1}{\omega}\left(\frac{2\gamma}{\alpha}\right)^3\left(\frac{\alpha^2}{2\gamma^2}\cosh\left(\frac{\alpha^2}{2\gamma^2}\right) - \sinh\left(\frac{\alpha^2}{2\gamma^2}\right)\right)\right]$
- Fisher Information: $\mathcal{I}_{\mathcal{FH}} = \frac{1}{\gamma^2}\left[1 + \frac{4\gamma}{\alpha\omega}\left(\frac{\alpha^2}{2\gamma^2}e^{\alpha^2/2\gamma^2} - \sinh\left(\frac{\alpha^2}{2\gamma^2}\right)\right)\right]$
- $\sigma^2_{\mathcal{FH}} \leq \gamma^2$ and $\mathcal{I}_{\mathcal{FH}} \geq \frac{1}{\gamma^2}$

Properties of flipped Huber (cont.,)

- Normalized Fisher information: $\widetilde{\mathcal{I}}_T = \mathcal{I}_T \times \sigma_T^2$

Properties of flipped Huber (cont.,)

- Normalized Fisher information: $\widetilde{\mathcal{I}}_T = \mathcal{I}_T \times \sigma_T^2$

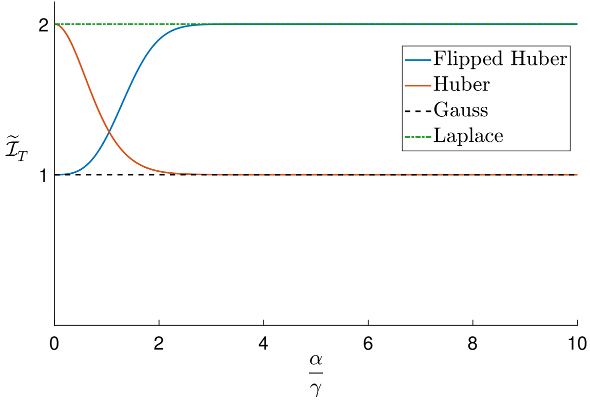

$\widetilde{\mathcal{I}}_T$ for $T \sim \mathcal{FH}(\alpha, \gamma^2)$ compared with that of $T \sim \mathcal{H}(\alpha, \gamma^2)$, $T \sim \mathcal{N}(0, \sigma^2)$ and $T \sim \mathcal{L}(0, \beta)$

Properties of flipped Huber (cont.,)

Lemma: Sub-Gaussianity of Flipped Huber

$\mathcal{FH}(\alpha, \gamma^2)$ is sub-Gaussian with proxy variance γ^2, i.e., $\mathcal{FH}(\alpha, \gamma^2) \in \mathcal{SG}(\gamma^2)$.

Properties of flipped Huber (cont.,)

Lemma: Sub-Gaussianity of Flipped Huber

$\mathcal{FH}(\alpha, \gamma^2)$ is sub-Gaussian with proxy variance γ^2, i.e., $\mathcal{FH}(\alpha, \gamma^2) \in \mathcal{SG}(\gamma^2)$.

Proof: Prove by showing Orlicz condition, $\mathbb{E}\left[\exp\left(\frac{sT^2}{2\gamma^2}\right)\right] \leq \frac{1}{\sqrt{1-s}} \;\; \forall\, s \in [0,1)$.

When $T \sim \mathcal{FH}(\alpha, \gamma^2)$, $\mathbb{E}\left[\exp\left(\frac{sT^2}{2\gamma^2}\right)\right] = \frac{1}{\sqrt{1-s}} \mathcal{C}_{\alpha,\gamma}(s)$, where

$$\mathcal{C}_{\alpha,\gamma}(s) = \frac{\sqrt{2\pi}}{\omega}\left[\sqrt{\frac{1}{s}-1}\ \exp\left(-\left(\frac{1}{s}-1\right)\frac{\alpha^2}{2\gamma^2}\right)\left(\mathrm{erfi}\left(\frac{1}{\sqrt{2s}}\frac{\alpha}{\gamma}\right) - \mathrm{erfi}\left(\frac{(1-s)}{\sqrt{2s}}\frac{\alpha}{\gamma}\right)\right) + 2Q\left(\sqrt{1-s}\,\frac{\alpha}{\gamma}\right)\right].$$

$\mathcal{C}_{\alpha,\gamma}(s)$ is a decreasing function in $s \in [0,1)$ and $\lim_{s \to 0^+} \mathcal{C}_{\alpha,\gamma}(s) = 1 \Rightarrow \mathcal{C}_{\alpha,\gamma}(s) \leq 1$.

Analysis for one dimension

Privacy guarantee in one dimension

The one-dimensional flipped Huber mechanism guarantees (ϵ,δ)-DP $\iff$ $\delta^{(1)}_{\mathcal{FH}}(\epsilon)\le\delta$, where

$$
\delta^{(1)}_{\mathcal{FH}}(\epsilon)=\begin{cases}
\left(1-\frac{\sqrt{2\pi}}{\omega}\right)+\frac{\sqrt{2\pi}}{\omega}\left[Q\left(\frac{\gamma\epsilon}{\Delta}-\frac{\Delta}{2\gamma}\right)-e^{\epsilon}Q\left(\frac{\gamma\epsilon}{\Delta}+\frac{\Delta}{2\gamma}\right)\right], & 0\le\epsilon<\frac{(\Delta-2\alpha)\Delta}{2\gamma^2}\\
\frac{1}{2}(1-e^{\epsilon})+\frac{\gamma}{\alpha\omega}e^{\alpha^2/2\gamma^2}\left(1+e^{\epsilon}-2\exp\left(\frac{\epsilon}{2}-\frac{\alpha\Delta}{2\gamma^2}\right)\right), & 0\le\epsilon<\frac{((2\alpha-\Delta)\wedge\Delta)\alpha}{\gamma^2}\\
\frac{1}{2}+\frac{\gamma}{\alpha\omega}e^{\alpha^2/2\gamma^2}\left[1-\exp\left(\frac{\alpha}{\gamma^2}\left(-\alpha+\sqrt{2(\gamma^2\epsilon+\alpha\Delta)}-\Delta\right)\right)\right]-e^{\epsilon}\frac{\sqrt{2\pi}}{\omega}Q\left(\frac{\sqrt{2(\gamma^2\epsilon+\alpha\Delta)}-\alpha}{\gamma}\right), & \\
 & \frac{(2\alpha\vee\Delta)^2-2\alpha\Delta}{2\gamma^2}\le\epsilon<\frac{([\Delta-\alpha]_+)^2+2\alpha\Delta}{2\gamma^2}\\
\frac{1}{2}-\frac{\gamma}{\alpha\omega}e^{\alpha^2/2\gamma^2}\left[1-\exp\left(\frac{\alpha}{\gamma^2}\left(-\alpha-\sqrt{2(\gamma^2\epsilon-\alpha\Delta)}+\Delta\right)\right)\right]-e^{\epsilon}\frac{\sqrt{2\pi}}{\omega}Q\left(\frac{\sqrt{2(\gamma^2\epsilon-\alpha\Delta)}+\alpha}{\gamma}\right), & \\
 & \frac{([\Delta-\alpha]_+)^2+2\alpha\Delta}{2\gamma^2}\le\epsilon<\frac{(\Delta+2\alpha)\Delta}{2\gamma^2}\\
\frac{\sqrt{2\pi}}{\omega}\left[Q\left(\frac{\gamma\epsilon}{\Delta}-\frac{\Delta}{2\gamma}\right)-e^{\epsilon}Q\left(\frac{\gamma\epsilon}{\Delta}+\frac{\Delta}{2\gamma}\right)\right], & \epsilon\ge\frac{(\Delta+2\alpha)\Delta}{2\gamma^2}
\end{cases}
$$

Proof sketch

$$\delta^{(1)}_{\mathcal{F}\mathcal{H}}(\epsilon) = \int_{\mathbb{R}} \left[g_{\mathcal{F}\mathcal{H}}(t) - e^{\epsilon} g_{\mathcal{F}\mathcal{H}}(t+d)\right]_+ \mathrm{d}t \leq \delta$$

$$d = f(\mathcal{D}) - f(\check{\mathcal{D}}) \qquad \Delta = \sup_{\mathcal{D} \asymp_{\mathcal{X}} \check{\mathcal{D}}} |d|$$

$$\delta^{(1)}_{\mathcal{F}\mathcal{H}}(\epsilon) = \overline{G}_{\mathcal{F}\mathcal{H}}\Big(\zeta_{\Delta}^{-1}(\epsilon)\Big) - e^{\epsilon}\,\overline{G}_{\mathcal{F}\mathcal{H}}\Big(\zeta_{\Delta}^{-1}(\epsilon) + \Delta\Big)$$

Proof sketch

$$\delta^{(1)}_{\mathcal{FH}}(\epsilon) = \int_{\mathbb{R}} \left[g_{\mathcal{FH}}(t) - e^{\epsilon} g_{\mathcal{FH}}(t+d)\right]_+ \mathrm{d}t \leq \delta$$

$$d = f(\mathcal{D}) - f(\check{\mathcal{D}}) \qquad \Delta = \sup_{\mathcal{D} \asymp_{\mathcal{X}} \check{\mathcal{D}}} |d|$$

$$\delta^{(1)}_{\mathcal{FH}}(\epsilon) = \overline{G}_{\mathcal{FH}}\left(\zeta_{\Delta}^{-1}(\epsilon)\right) - e^{\epsilon}\,\overline{G}_{\mathcal{FH}}\left(\zeta_{\Delta}^{-1}(\epsilon) + \Delta\right)$$

- Piecewise density $\rightarrow$ piecewise $\zeta_d(t)$
- 3 different functional forms of $\zeta_d(t)$

Proof sketch

$$\delta^{(1)}_{\mathcal{FH}}(\epsilon) = \int_{\mathbb{R}} \left[g_{\mathcal{FH}}(t) - e^{\epsilon} g_{\mathcal{FH}}(t+d)\right]_+ \mathrm{d}t \leq \delta$$

$$d = f(\mathcal{D}) - f(\check{\mathcal{D}}) \qquad \Delta = \sup_{\mathcal{D} \asymp_{\mathcal{X}} \check{\mathcal{D}}} |d|$$

$$\delta^{(1)}_{\mathcal{FH}}(\epsilon) = \overline{G}_{\mathcal{FH}}\Big(\zeta_{\Delta}^{-1}(\epsilon)\Big) - e^{\epsilon}\,\overline{G}_{\mathcal{FH}}\Big(\zeta_{\Delta}^{-1}(\epsilon) + \Delta\Big)$$

- Piecewise density $\rightarrow$ piecewise $\zeta_d(t)$
- 3 different functional forms of $\zeta_d(t)$

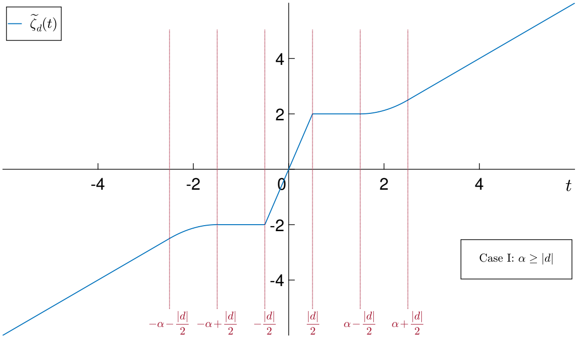

Proof sketch

$$\delta^{(1)}_{\mathcal{FH}}(\epsilon) = \int_{\mathbb{R}} \left[g_{\mathcal{FH}}(t) - e^{\epsilon} g_{\mathcal{FH}}(t+d)\right]_{+} \mathrm{d}t \le \delta$$

$$d = f(\mathcal{D}) - f(\check{\mathcal{D}}) \qquad \Delta = \sup_{\mathcal{D} \asymp_{\mathcal{X}} \check{\mathcal{D}}} |d|$$

$$\delta^{(1)}_{\mathcal{FH}}(\epsilon) = \overline{G}_{\mathcal{FH}}\left(\zeta_{\Delta}^{-1}(\epsilon)\right) - e^{\epsilon}\,\overline{G}_{\mathcal{FH}}\left(\zeta_{\Delta}^{-1}(\epsilon) + \Delta\right)$$

- Piecewise density $\rightarrow$ piecewise $\zeta_d(t)$
- 3 different functional forms of $\zeta_d(t)$

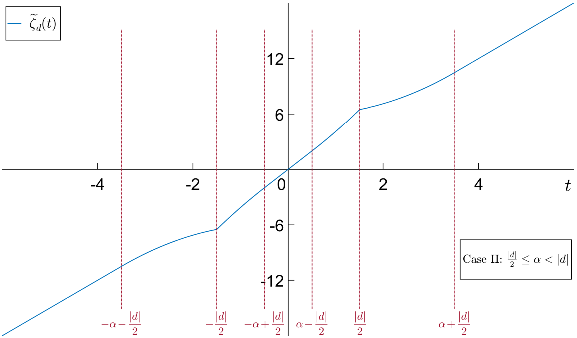

Proof sketch

$$\delta^{(1)}_{\mathcal{FH}}(\epsilon) = \int_{\mathbb{R}} \left[g_{\mathcal{FH}}(t) - e^{\epsilon} g_{\mathcal{FH}}(t+d)\right]_+ \mathrm{d}t \leq \delta$$

$$d = f(\mathcal{D}) - f(\check{\mathcal{D}}) \qquad \Delta = \sup_{\mathcal{D} \asymp_{\mathcal{X}} \check{\mathcal{D}}} |d|$$

$$\delta^{(1)}_{\mathcal{FH}}(\epsilon) = \overline{G}_{\mathcal{FH}}\left(\zeta_{\Delta}^{-1}(\epsilon)\right) - e^{\epsilon}\,\overline{G}_{\mathcal{FH}}\left(\zeta_{\Delta}^{-1}(\epsilon) + \Delta\right)$$

- Piecewise density $\rightarrow$ piecewise $\zeta_d(t)$
- 3 different functional forms of $\zeta_d(t)$

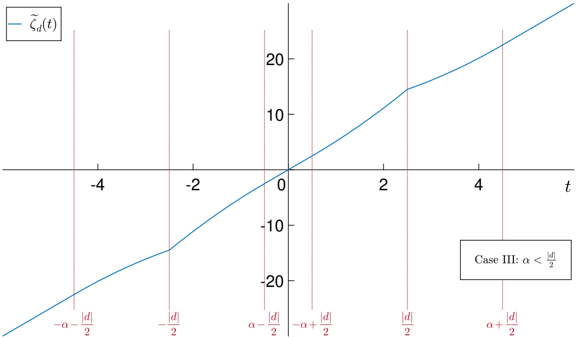

Proof sketch

$$\delta^{(1)}_{\mathcal{FH}}(\epsilon) = \int_{\mathbb{R}} \left[g_{\mathcal{FH}}(t) - e^{\epsilon} g_{\mathcal{FH}}(t+d)\right]_+ \mathrm{d}t \leq \delta$$

$$d = f(\mathcal{D}) - f(\check{\mathcal{D}}) \qquad \Delta = \sup_{\mathcal{D} \asymp_{\mathcal{X}} \check{\mathcal{D}}} |d|$$

$$\delta^{(1)}_{\mathcal{FH}}(\epsilon) = \overline{G}_{\mathcal{FH}}\left(\zeta_{\Delta}^{-1}(\epsilon)\right) - e^{\epsilon}\,\overline{G}_{\mathcal{FH}}\left(\zeta_{\Delta}^{-1}(\epsilon) + \Delta\right)$$

- Piecewise density $\rightarrow$ piecewise $\zeta_d(t)$
- 3 different functional forms of $\zeta_d(t)$

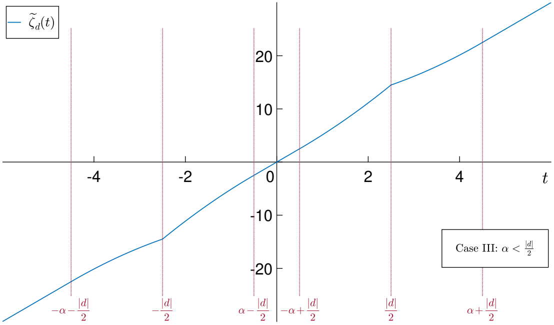

Depending on value of ϵ, we may get different values of $\zeta_{\Delta}^{-1}(\epsilon)$ for each of the 3 cases

Proof sketch (cont.,)

	Range of ϵ	Non-empty only when	$t_1 = \zeta_\Delta^{-1}(\epsilon)$	$t_2 = \zeta_\Delta^{-1}(\epsilon) + \Delta$	Intervals t_1	Intervals t_2
(i)	$\left[0, \frac{(\Delta - 2\alpha)\Delta}{2\gamma^2}\right)$	$\alpha < \frac{\Delta}{2}$	$\frac{\gamma^2\epsilon}{\Delta} - \frac{\Delta}{2}$	$\frac{\gamma^2\epsilon}{\Delta} + \frac{\Delta}{2}$	$(-\infty, -\alpha)$	(α, ∞)
(ii)	$\left[0, \frac{((2\alpha - \Delta) \wedge \Delta)\alpha}{\gamma^2}\right)$	$\alpha > \frac{\Delta}{2}$	$\frac{\gamma^2\epsilon}{2\alpha} - \frac{\Delta}{2}$	$\frac{\gamma^2\epsilon}{2\alpha} + \frac{\Delta}{2}$	$(-\alpha, 0)$	$(0, \alpha)$
(iii)	$\left[\frac{(2\alpha \vee \Delta)^2 - 2\alpha\Delta}{2\gamma^2}, \frac{([\Delta - \alpha]_+)^2 + 2\alpha\Delta}{2\gamma^2}\right)$	$\alpha < \Delta$	$\sqrt{2(\gamma^2\epsilon + \alpha\Delta)} - \alpha - \Delta$	$\sqrt{2(\gamma^2\epsilon + \alpha\Delta)} - \alpha$	$[-\alpha, 0)$	$[\alpha, \infty)$
(iv)	$\left[\frac{([\Delta - \alpha]_+)^2 + 2\alpha\Delta}{2\gamma^2}, \frac{(\Delta + 2\alpha)\Delta}{2\gamma^2}\right)$	–	$\sqrt{2(\gamma^2\epsilon - \alpha\Delta)} + \alpha + \Delta$	$\sqrt{2(\gamma^2\epsilon - \alpha\Delta)} + \alpha$	$[0, \alpha)$	$[\alpha, \infty)$
(v)	$\left[\frac{(\Delta + 2\alpha)\Delta}{2\gamma^2}, \infty\right)$	–	$\frac{\gamma^2\epsilon}{\Delta} - \frac{\Delta}{2}$	$\frac{\gamma^2\epsilon}{\Delta} + \frac{\Delta}{2}$	$[\alpha, \infty)$	$[\alpha, \infty)$

- $\epsilon \geq \frac{([\Delta-\alpha]_+)^2+2\alpha\Delta}{2\gamma^2} \to$ no ambiguity, otherwise determine $\zeta_\Delta^{-1}(\epsilon)$ based on α and Δ

Empirical results for single dimension

Performance in single dimension

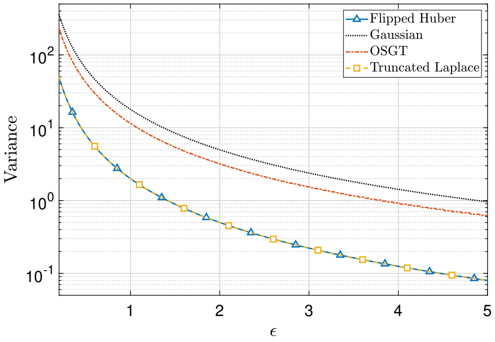

Variances of flipped Huber, Gaussian, truncated Laplace[6] and OSGT[7] noises when $\delta = 10^{-6}$ and $\Delta = 1$

[6] Q. Geng, W. Ding, R. Guo, and S. Kumar, "Tight analysis of privacy and utility tradeoff in approximate differential privacy," in *Proc. Int. Conf. Artif. Intell. Statist.* PMLR, 2020, pp. 89–99

[7] P. Sadeghi and M. Korki, "Offset-symmetric Gaussians for differential privacy," *IEEE Trans. Inf. Forensics Security*, vol. 17, pp. 2394–2409, 2022

How to choose the parameters of flipped Huber?

- (α, γ) that results in lowest variance while satisfying DP can be selected through grid search

- Illustrative values ($\delta = 10^{-6}$):

ϵ	α	γ
0.5	20.48	6.4
2	6.48	1.8
4	4	1

Performance in single dimension (cont.,)

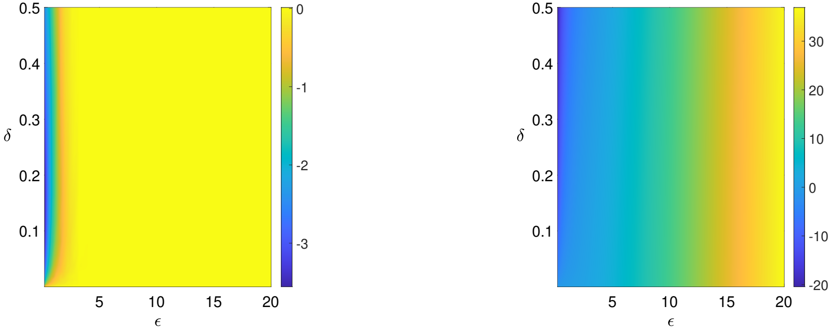

(a) (b)

Ratio of the variance (in dB) of flipped Huber noise to that of (a) Laplace and (b) staircase[8] noises

[8] Q. Geng and P. Viswanath, "The optimal noise-adding mechanism in differential privacy," *IEEE Trans. Inf. Theory*, vol. 62, no. 2, pp. 925–951, 2016

Noise densities for $\epsilon = 2$, $\delta = 10^{-6}$

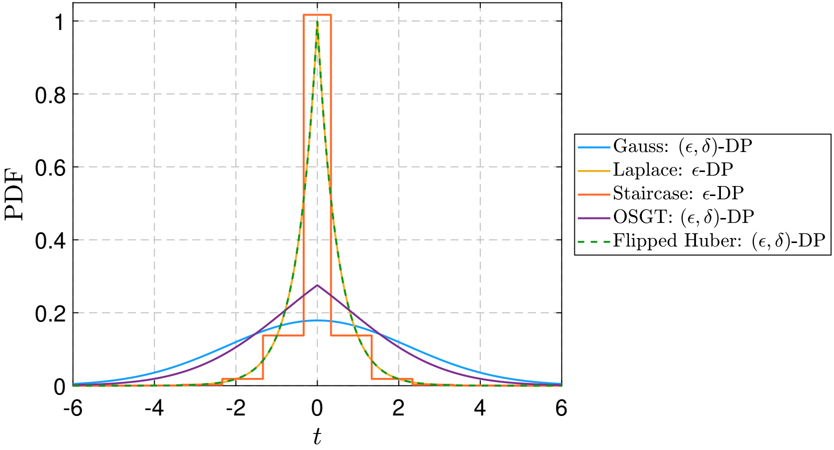

The story so far

- Approximate DP guaranteed by the flipped Huber mechanism
- Outperforms both Gaussian and OSGT mechanism by a significant margin
- However Laplace is clearly superior and gives pure DP
- Staircase is the optimal pure DP noise mechanism in single and two dimensions
- So why flipped Huber?

Flipped Huber strikes back in higher dimensions

DP in higher dimensions

- Machine learning applications are typically high-dimensional
  - Linear regression[9]: few 10's
  - Principal component analysis[10]: few 100's or 1000's
  - Deep learning[11]: several millions
- Need for efficient DP mechanisms without killing the utility

[9] Y.-X. Wang, "Revisiting differentially private linear regression: optimal and adaptive prediction & estimation in unbounded domain," in *Uncertainty in Artif. Intell.*, 2018

[10] C. Dwork, K. Talwar, A. Thakurta, and L. Zhang, "Analyze Gauss: optimal bounds for privacy-preserving principal component analysis," in *Proc. Annu. ACM Symp. Theory of Comput.*, 2014, pp. 11–20

[11] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, "Deep learning with differential privacy," in *Proc. ACM SIGSAC Conf. Computer and Communications security*, 2016, pp. 308–318

Performance in $K = 20$ dimensions for $\delta = 10^{-6}$

Variance of noise added by various mechanisms

ϵ	0.2	0.4	1	2.2	5
Flipped Huber (sufficient)	7237.09	1971.36	359.57	87.09	19.49
Gaussian	11209.84	2979.23	520.26	117.77	25.95
Laplace	$2 \cdot 10^4$	5000	800	165.29	32
Staircase (independent)	19999.92	4999.92	799.92	165.21	31.92

Noise densities in $K = 20$ dimensions for $\epsilon = 5$, $\delta = 10^{-6}$

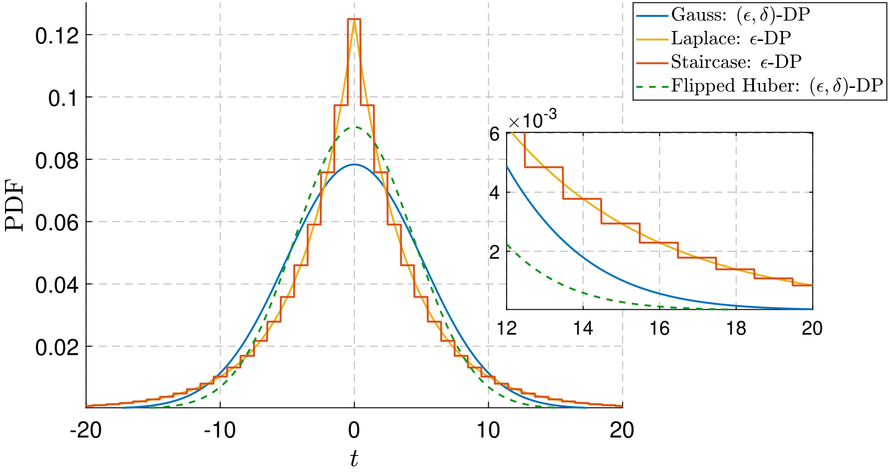

The problem with DP in higher dimensions

- Privacy loss distribution is difficult to characterize (except for Gaussian)
  - Difficult even for i.i.d. noise (convolution)

[12] Q. Geng, P. Kairouz, S. Oh, and P. Viswanath, “The staircase mechanism in differential privacy,” *IEEE J. Sel. Topics Signal Process.*, vol. 9, no. 7, pp. 1176–1184, 2015

The problem with DP in higher dimensions

- Privacy loss distribution is difficult to characterize (except for Gaussian)
  - Difficult even for i.i.d. noise (convolution)
- Numerical integration - computationally prohibitive even for small dimensions

[12] Q. Geng, P. Kairouz, S. Oh, and P. Viswanath, "The staircase mechanism in differential privacy," *IEEE J. Sel. Topics Signal Process.*, vol. 9, no. 7, pp. 1176–1184, 2015

The problem with DP in higher dimensions

- Privacy loss distribution is difficult to characterize (except for Gaussian)
    - Difficult even for i.i.d. noise (convolution)
- Numerical integration - computationally prohibitive even for small dimensions
- Composition approach - not tight

[12] Q. Geng, P. Kairouz, S. Oh, and P. Viswanath, “The staircase mechanism in differential privacy,” *IEEE J. Sel. Topics Signal Process.*, vol. 9, no. 7, pp. 1176–1184, 2015

The problem with DP in higher dimensions

- Privacy loss distribution is difficult to characterize (except for Gaussian)
  - Difficult even for i.i.d. noise (convolution)
- Numerical integration - computationally prohibitive even for small dimensions
- Composition approach - not tight
- Optimal noise distribution for arbitrary dimension is not known yet
  - All optimal distributions in literature are for single-dimensional queries
  - High dimensional functional optimization - difficult

[12] Q. Geng, P. Kairouz, S. Oh, and P. Viswanath, "The staircase mechanism in differential privacy," *IEEE J. Sel. Topics Signal Process.*, vol. 9, no. 7, pp. 1176–1184, 2015

The problem with DP in higher dimensions

- Privacy loss distribution is difficult to characterize (except for Gaussian)
  - Difficult even for i.i.d. noise (convolution)
- Numerical integration - computationally prohibitive even for small dimensions
- Composition approach - not tight
- Optimal noise distribution for arbitrary dimension is not known yet
  - All optimal distributions in literature are for single-dimensional queries
  - High dimensional functional optimization - difficult
- Staircase is the optimal noise for ϵ-DP (under ℓ_1-error) in two dimensions[12]
  - PDF is not characterized in K dimensions
  - Use i.i.d. samples from one-dimensional staircase distribution

[12] Q. Geng, P. Kairouz, S. Oh, and P. Viswanath, "The staircase mechanism in differential privacy," *IEEE J. Sel. Topics Signal Process.*, vol. 9, no. 7, pp. 1176–1184, 2015

Privacy guarantee in K dimensions

- K-dimensional query: Add i.i.d. $\mathcal{FH}(\alpha, \gamma^2)$ to each coordinate
- Necessary and sufficient condition $\rightarrow$ intractable
  - Hybrid, piecewise nature of $\mathcal{FH}(\alpha, \gamma^2)$
  - Complex expression for $\zeta_{\mathbf{d}}(\mathbf{T})$

Privacy guarantee in K dimensions

- K-dimensional query: Add i.i.d. $\mathcal{FH}(\alpha, \gamma^2)$ to each coordinate
- Necessary and sufficient condition $\rightarrow$ intractable
  - Hybrid, piecewise nature of $\mathcal{FH}(\alpha, \gamma^2)$
  - Complex expression for $\zeta_{\mathbf{d}}(\mathbf{T})$

Theorem: Sufficient Condition for (ϵ, δ)-DP in K Dimension

The K-dimensional flipped Huber mechanism guarantees (ϵ, δ)-DP if $\mathcal{R}_\Delta(\alpha) = \alpha^2 - ([\alpha - \Delta]_+)^2 \leq (2\gamma^2\epsilon - \Delta_2^2)/K$ and

$$Q\left(\frac{\gamma\epsilon}{\Delta_2} - \frac{\Delta_2}{2\gamma} - \frac{K\mathcal{R}_\Delta(\alpha)}{2\gamma\Delta_2}\right) - e^\epsilon Q\left(\frac{\gamma\epsilon}{\Delta_2} + \frac{\Delta_2}{2\gamma} + \frac{K\mathcal{R}_\Delta(\alpha)}{2\gamma\Delta_2} + \frac{\theta\Delta_1}{\gamma\Delta_2}\right) \leq \delta,$$

where $\theta = \gamma Q^{-1}\left(\frac{1}{\omega}\sqrt{\frac{\pi}{2}}\right)$.

Proof sketch

- $\mathbf{T} \overset{\mathrm{d}}{=} -\mathbf{T}$ and $\zeta_{-\mathbf{d}}(\mathbf{T}) = \zeta_{\mathbf{d}}(-\mathbf{T}) \overset{\mathrm{d}}{=} \zeta_{\mathbf{d}}(\mathbf{T})$

$$\mathbb{P}\{\zeta_{\mathbf{d}}(\mathbf{T}) \geq \epsilon\} - e^{\epsilon}\,\mathbb{P}\{\zeta_{\mathbf{d}}(\mathbf{T}) \leq -\epsilon\} \leq \delta$$

Proof sketch

- $\mathbf{T} \overset{\mathrm{d}}{=} -\mathbf{T}$ and $\zeta_{-\mathbf{d}}(\mathbf{T}) = \zeta_{\mathbf{d}}(-\mathbf{T}) \overset{\mathrm{d}}{=} \zeta_{\mathbf{d}}(\mathbf{T})$

$$\mathbb{P}\{\zeta_{\mathbf{d}}(\mathbf{T}) \geq \epsilon\} - e^{\epsilon}\,\mathbb{P}\{\zeta_{\mathbf{d}}(\mathbf{T}) \leq -\epsilon\} \leq \delta$$

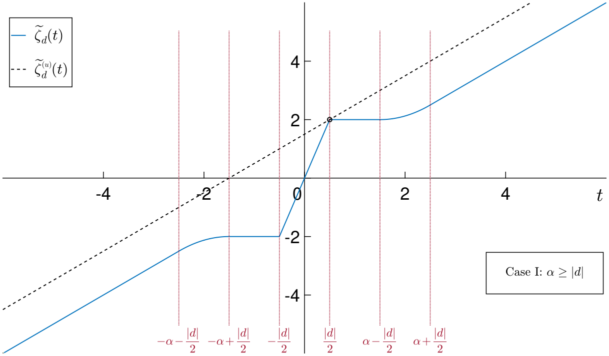

Proof sketch

- $\mathbf{T} \overset{\mathrm{d}}{=} -\mathbf{T}$ and $\zeta_{-\mathbf{d}}(\mathbf{T}) = \zeta_{\mathbf{d}}(-\mathbf{T}) \overset{\mathrm{d}}{=} \zeta_{\mathbf{d}}(\mathbf{T})$

$$\mathbb{P}\{\zeta_{\mathbf{d}}(\mathbf{T}) \geq \epsilon\} - e^{\epsilon}\,\mathbb{P}\{\zeta_{\mathbf{d}}(\mathbf{T}) \leq -\epsilon\} \leq \delta$$

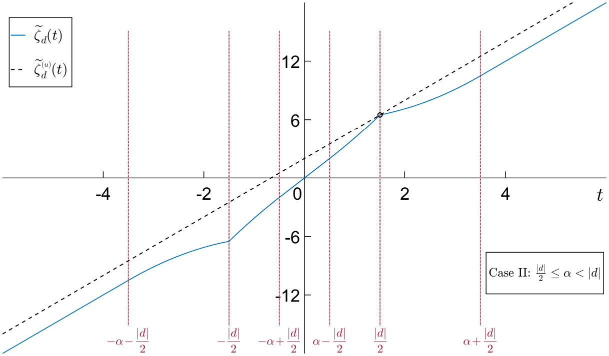

Proof sketch

- $\mathbf{T} \overset{\mathrm{d}}{=} -\mathbf{T}$ and $\zeta_{-\mathbf{d}}(\mathbf{T}) = \zeta_{\mathbf{d}}(-\mathbf{T}) \overset{\mathrm{d}}{=} \zeta_{\mathbf{d}}(\mathbf{T})$

$$\mathbb{P}\{\zeta_{\mathbf{d}}(\mathbf{T}) \geq \epsilon\} - e^{\epsilon}\,\mathbb{P}\{\zeta_{\mathbf{d}}(\mathbf{T}) \leq -\epsilon\} \leq \delta$$

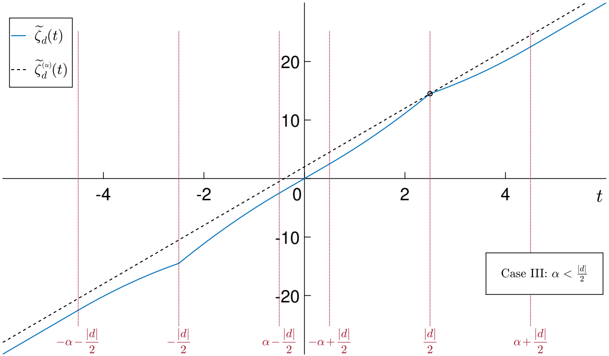

Proof Sketch

- $\mathbf{T} \overset{\mathrm{d}}{=} -\mathbf{T}$ and $\zeta_{-\mathbf{d}}(\mathbf{T}) = \zeta_{\mathbf{d}}(-\mathbf{T}) \overset{\mathrm{d}}{=} \zeta_{\mathbf{d}}(\mathbf{T})$

  $$\mathbb{P}\{\zeta_{\mathbf{d}}(\mathbf{T}) \geq \epsilon\} - e^{\epsilon}\,\mathbb{P}\{\zeta_{\mathbf{d}}(\mathbf{T}) \leq -\epsilon\} \leq \delta$$

- $\zeta_{\mathbf{d}}(\mathbf{T}) \leq \zeta_{\mathbf{d}}^{(u)}(\mathbf{t}) = \dfrac{\mathbf{t}^\top \mathbf{d}}{\gamma^2} + \dfrac{\|\mathbf{d}\|_2^2}{2\gamma^2} + \dfrac{K\mathcal{R}_\Delta(\alpha)}{2\gamma^2}$

  $$\mathbb{P}\{\zeta_{\mathbf{d}}^{(u)}(\mathbf{T}) \geq \epsilon\} - e^{\epsilon}\,\mathbb{P}\{\zeta_{\mathbf{d}}^{(u)}(\mathbf{T}) \leq -\epsilon\} \leq \delta$$

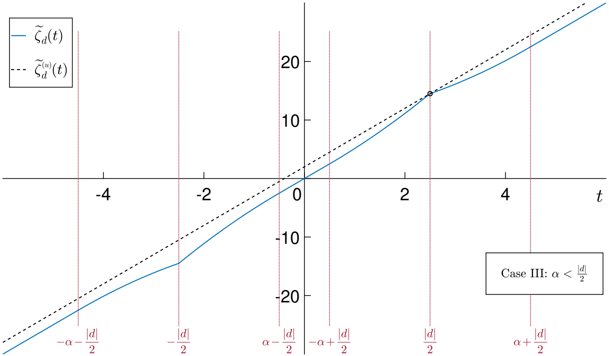

Proof sketch (cont.,)

- Upper bound the upper tail probability of privacy loss: Sub-Gaussianity

Proof sketch (cont.,)

- Upper bound the upper tail probability of privacy loss: Sub-Gaussianity

Lemma: Upper Bound on the Upper Tail Probability

Let $T_1, T_2, \ldots, T_K$ be i.i.d. flipped Huber RVs, $T_i \sim \mathcal{FH}(\alpha, \gamma^2)$ and $\mathbf{T} = [\, T_1 \; T_2 \; \cdots \; T_K \,]^\top$. If $\mathcal{R}_\Delta(\alpha) = \alpha^2 - ([\alpha - \Delta]_+)^2 \leq (2\gamma^2\epsilon - \Delta_2^2)/K$, then

$$\mathbb{P}\{\zeta_{\mathbf{d}}^{(u)}(\mathbf{T}) \geq \epsilon\} \leq Q\left(\frac{\gamma\epsilon}{\Delta_2} - \frac{\Delta_2}{2\gamma} - \frac{K\mathcal{R}_\Delta(\alpha)}{2\gamma\Delta_2}\right).$$

Proof sketch (cont.,)

- Lower bound the lower tail probability of privacy loss: Stochastic ordering
  - $X \leq_{\mathtt{st}} Y$ if $\overline{G}_X(a) \leq \overline{G}_Y(a)\ \forall\, a \in \mathbb{R}$

Proof sketch (cont.,)

- Lower bound the lower tail probability of privacy loss: Stochastic ordering
  - $X \leq_{\mathsf{st}} Y$ if $\overline{G}_X(a) \leq \overline{G}_Y(a)\ \forall\, a \in \mathbb{R}$

Lemma: Stochastic Upper Bound for $\mathcal{FH}(\alpha, \gamma^2)$

$\mathcal{FH}(\alpha, \gamma^2) \leq_{\mathsf{st}} \mathcal{N}(\theta, \gamma^2)$, where $\theta = \gamma Q^{-1}(\frac{1}{\omega}\sqrt{\frac{\pi}{2}})$.

Proof sketch (cont.,)

- Lower bound the lower tail probability of privacy loss: Stochastic ordering
  - $X \leq_{\mathsf{st}} Y$ if $\overline{G}_X(a) \leq \overline{G}_Y(a)\ \forall\, a \in \mathbb{R}$

Lemma: Stochastic Upper Bound for $\mathcal{FH}(\alpha, \gamma^2)$

$\mathcal{FH}(\alpha, \gamma^2) \leq_{\mathsf{st}} \mathcal{N}(\theta, \gamma^2)$, where $\theta = \gamma Q^{-1}(\frac{1}{\omega}\sqrt{\frac{\pi}{2}})$.

Lemma: Lower Bound on the Lower Tail Probability

Let $T_1, T_2, \ldots, T_K$ be i.i.d. flipped Huber random variables, $T_i \sim \mathcal{FH}(\alpha, \gamma^2)$ and $\mathbf{T} = [\,T_1\ T_2\ \cdots\ T_K\,]^\top$. We have

$$\mathbb{P}\{\zeta_{\mathbf{d}}^{(u)}(\mathbf{T}) \leq -\epsilon\} \geq Q\left(\frac{\gamma\epsilon}{\Delta_2} + \frac{\Delta_2}{2\gamma} + \frac{K\mathcal{R}_\Delta(\alpha)}{2\gamma\Delta_2} + \frac{\theta\Delta_1}{\gamma\Delta_2}\right),$$

where $\theta = \gamma Q^{-1}(\frac{1}{\omega}\sqrt{\frac{\pi}{2}})$.

Performance in K dimensions

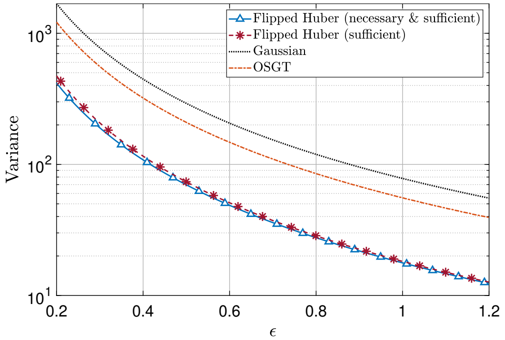

(a) $K = 3$

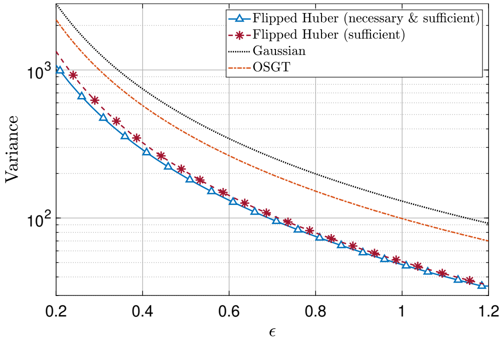

(b) $K = 5$

$$\delta = 10^{-8},\ \ \Delta = 1,\ \ \Delta_2 = \sqrt{K}\ \text{ and }\ \Delta_1 = K$$

Applications

- Any estimation/detection/ML can be privatized
- What about iterative algorithms?
  - Composition and guarantees for the same required
- What about neural networks?
  - Gradient clipping typically required

Composition

- Consider a set of DP mechanisms $\mathcal{M}_l(\cdot)$, $l = 1, 2, \ldots, L$

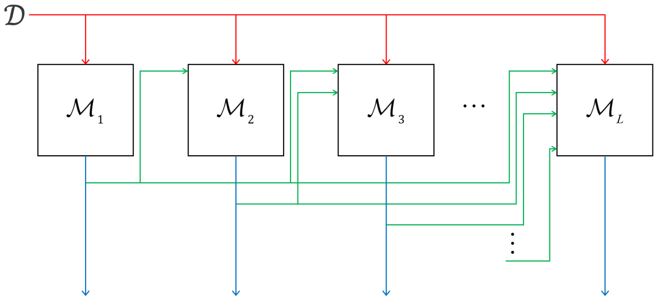

Composition

- Consider a set of DP mechanisms $\mathcal{M}_l(\cdot)$, $l = 1,\ 2,\ \ldots,\ L$

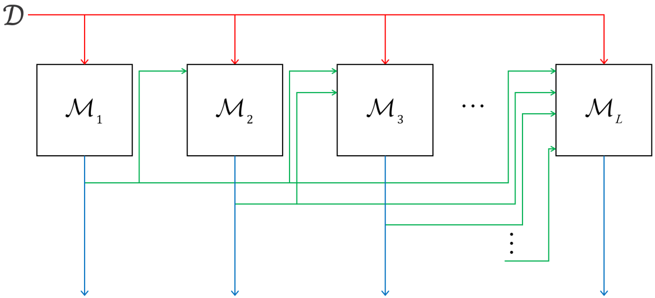

- $\mathcal{D} \mapsto \big(\mathcal{M}_1(\mathcal{D}), \ldots, \mathcal{M}_L(\mathcal{D})\big)$ is also DP (with graceful degradation)

Composition

- Consider a set of DP mechanisms $\mathcal{M}_l(\cdot)$, $l = 1, 2, \ldots, L$

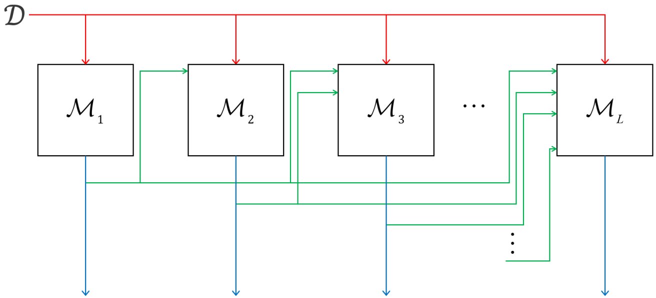

- $\mathcal{D} \mapsto \big(\mathcal{M}_1(\mathcal{D}), \ldots, \mathcal{M}_L(\mathcal{D})\big)$ is also DP (with graceful degradation)
- Non-adaptive composition: without *side links* $\rightarrow$ Multi-dimensional query

Zero concentrated differential privacy (zCDP)

Definition: (ξ, η)-zCDP[13]

The randomized mechanism $\mathcal{M} : \mathcal{X} \to \mathcal{Y}$ is said to satisfy (ξ, η)-zCDP if

$$\mathfrak{D}_{\Lambda}^{(\mathsf{R})}(\mu \,\|\, \breve{\mu}) \leq \xi + \Lambda\eta \quad \forall\, \Lambda \in (1, \infty) \;\text{ and }\; \mathcal{D} \asymp_{\mathcal{X}} \breve{\mathcal{D}},$$

where $\mathfrak{D}_{\Lambda}^{(\mathsf{R})}(\mu \,\|\, \breve{\mu})$ is the Λ-Rényi divergence between the distributions of $\mathcal{M}(\mathcal{D})$ and $\mathcal{M}(\breve{\mathcal{D}})$.

[13] M. Bun and T. Steinke, “Concentrated differential privacy: Simplifications, extensions, and lower bounds,” in *Proc. Int. Conf. Theory of Cryptogr. Part I.* Springer, 2016, pp. 635–658

Zero concentrated differential privacy (zCDP)

Definition: (ξ, η)-zCDP[13]

The randomized mechanism $\mathcal{M} : \mathcal{X} \to \mathcal{Y}$ is said to satisfy (ξ, η)-zCDP if

$$\mathfrak{D}_{\Lambda}^{(\mathsf{R})}(\mu \,\|\, \check{\mu}) \leq \xi + \Lambda\eta \quad \forall\, \Lambda \in (1, \infty) \;\; \text{and} \;\; \mathcal{D} \asymp_{\mathcal{X}} \check{\mathcal{D}},$$

where $\mathfrak{D}_{\Lambda}^{(\mathsf{R})}(\mu \,\|\, \check{\mu})$ is the Λ-Rényi divergence between the distributions of $\mathcal{M}(\mathcal{D})$ and $\mathcal{M}(\check{\mathcal{D}})$.

- zCDP $\to$ bound on the MGF of $\zeta_d(T)$: $\mathbb{E}\left[\exp\left(s\,\zeta_d(T)\right)\right] \leq e^{s(\xi+(s+1)\eta)}$ $\forall s > 0$

[13] M. Bun and T. Steinke, “Concentrated differential privacy: Simplifications, extensions, and lower bounds,” in *Proc. Int. Conf. Theory of Cryptogr. Part I.* Springer, 2016, pp. 635–658

Zero concentrated differential privacy (zCDP)

Definition: (ξ,η)-zCDP[13]

The randomized mechanism $\mathcal{M}:\mathcal{X}\to\mathcal{Y}$ is said to satisfy (ξ,η)-zCDP if

$$\mathfrak{D}_{\Lambda}^{(\mathrm{R})}(\mu\,\|\,\check{\mu}) \leq \xi + \Lambda\eta \quad \forall\,\Lambda \in (1,\infty) \ \text{ and } \ \mathcal{D} \asymp_{\mathcal{X}} \check{\mathcal{D}},$$

where $\mathfrak{D}_{\Lambda}^{(\mathrm{R})}(\mu\,\|\,\check{\mu})$ is the Λ-Rényi divergence between the distributions of $\mathcal{M}(\mathcal{D})$ and $\mathcal{M}(\check{\mathcal{D}})$.

- zCDP $\rightarrow$ bound on the MGF of $\zeta_d(T)$: $\mathbb{E}\Big[\exp\Big(s\,\zeta_d(T)\Big)\Big] \leq e^{s(\xi+(s+1)\eta)}$ $\forall s > 0$
- L-fold (adaptive) composition of (ξ_l,η_l)-zCDP mechanisms $\rightarrow$ $\left(\sum_{l=1}^{L}\xi_l, \sum_{l=1}^{L}\eta_l\right)$-zCDP

[13] M. Bun and T. Steinke, "Concentrated differential privacy: Simplifications, extensions, and lower bounds," in *Proc. Int. Conf. Theory of Cryptogr. Part I.* Springer, 2016, pp. 635–658

Zero concentrated differential privacy (zCDP)

Definition: (ξ,η)-zCDP[13]

The randomized mechanism $\mathcal{M}:\mathcal{X}\to\mathcal{Y}$ is said to satisfy (ξ,η)-zCDP if

$$\mathfrak{D}_{\Lambda}^{(\mathsf{R})}(\mu\,\|\,\check{\mu}) \le \xi + \Lambda\eta \quad \forall\,\Lambda\in(1,\infty) \ \text{ and } \ \mathcal{D} \asymp_{\mathcal{X}} \check{\mathcal{D}},$$

where $\mathfrak{D}_{\Lambda}^{(\mathsf{R})}(\mu\,\|\,\check{\mu})$ is the Λ-Rényi divergence between the distributions of $\mathcal{M}(\mathcal{D})$ and $\mathcal{M}(\check{\mathcal{D}})$.

- zCDP $\to$ bound on the MGF of $\zeta_d(T)$: $\mathbb{E}\Big[\exp\Big(s\,\zeta_d(T)\Big)\Big] \le e^{s(\xi+(s+1)\eta)}\ \forall s>0$
- L-fold (adaptive) composition of (ξ_l,η_l)-zCDP mechanisms $\to$ $\left(\sum_{l=1}^{L}\xi_l,\sum_{l=1}^{L}\eta_l\right)$-zCDP
- zCDP offers tightest characterization for Gaussian

[13] M. Bun and T. Steinke, "Concentrated differential privacy: Simplifications, extensions, and lower bounds," in *Proc. Int. Conf. Theory of Cryptogr. Part I.* Springer, 2016, pp. 635–658

zCDP of flipped Huber

Theorem: zCDP of $\mathcal{FH}(\alpha, \gamma^2)$

The one-dimensional flipped Huber mechanism guarantees $\left(\frac{\mathcal{R}_\Delta(\alpha)}{2\gamma^2}, \frac{\Delta^2}{2\gamma^2}\right)$-zCDP, where $\mathcal{R}_\Delta(\alpha) = \alpha^2 - ([\alpha - \Delta]_+)^2$.

zCDP of flipped Huber

Theorem: zCDP of $\mathcal{FH}(\alpha, \gamma^2)$

The one-dimensional flipped Huber mechanism guarantees $\left(\frac{\mathcal{R}_\Delta(\alpha)}{2\gamma^2}, \frac{\Delta^2}{2\gamma^2}\right)$-zCDP, where $\mathcal{R}_\Delta(\alpha) = \alpha^2 - ([\alpha - \Delta]_+)^2$.

- K-dimensional flipped Huber mechanism $\rightarrow$ $\left(\frac{K\mathcal{R}_\Delta(\alpha)}{2\gamma^2}, \frac{\Delta_2^2}{2\gamma^2}\right)$-zCDP

zCDP of flipped Huber

Theorem: zCDP of $\mathcal{FH}(\alpha, \gamma^2)$

The one-dimensional flipped Huber mechanism guarantees $\left(\frac{\mathcal{R}_\Delta(\alpha)}{2\gamma^2}, \frac{\Delta^2}{2\gamma^2}\right)$-zCDP, where $\mathcal{R}_\Delta(\alpha) = \alpha^2 - ([\alpha - \Delta]_+)^2$.

- K-dimensional flipped Huber mechanism $\rightarrow$ $\left(\frac{K\mathcal{R}_\Delta(\alpha)}{2\gamma^2}, \frac{\Delta_2^2}{2\gamma^2}\right)$-zCDP

Proof:

$$\zeta_d(T) \leq_{\mathtt{st}} \zeta_d^{(u)}(T)$$

$$\begin{aligned}
\mathbb{M}_{\zeta_d(T)}(s) &= \mathbb{E}\left[e^{s\zeta_d(T)}\right] \leq \mathbb{E}\left[e^{s\zeta_d^{(u)}(T)}\right] \\
&= \exp\left(s\frac{\mathcal{R}_{|d|}(\alpha)}{2\gamma^2} + s\frac{d^2}{2\gamma^2}\right) \times \mathbb{E}\left[\exp\left(\frac{s\,d}{\gamma^2}T\right)\right] \\
&\leq \exp\left(s\frac{\mathcal{R}_{|d|}(\alpha)}{2\gamma^2} + s(s+1)\frac{d^2}{2\gamma^2}\right) \leq \exp\left(s\frac{\mathcal{R}_\Delta(\alpha)}{2\gamma^2} + s(s+1)\frac{\Delta^2}{2\gamma^2}\right)
\end{aligned}$$

Coordinate descent (CD)

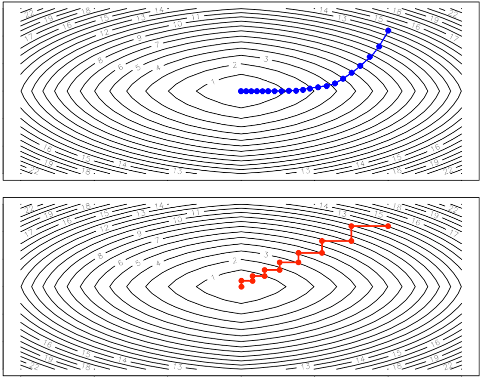

Differentially private coordinate descent (DP-CD)[14]

- Perturb gradient updates of CD
- Empirical Risk Minimization (ERM):

$$\min_{\boldsymbol{\theta} \in \mathbb{R}^K} \frac{1}{n} \sum_{n=1}^{N} J(\boldsymbol{\theta}; \mathcal{D}_n) + \psi(\boldsymbol{\theta}),$$

- $\boldsymbol{\theta} \in \mathbb{R}^K$ – model parameter
- $\mathcal{D} = (\mathcal{D}_1, \mathcal{D}_2, \ldots, \mathcal{D}_N) \in \mathcal{X}$ – dataset of N samples; $\mathcal{D}_n = (\mathbf{x}_n, y_n)$
- $J : \mathbb{R}^K \times \mathcal{X} \to \mathbb{R}$ – convex and smooth loss function
- $\psi : \mathbb{R}^K \to \mathbb{R}$ – convex and separable regularizing function, $\psi(\boldsymbol{\theta}) = \sum_{i=1}^{K} \psi_i(\theta_i)$

[14] P. Mangold, A. Bellet, J. Salmon, and M. Tommasi, “Differentially private coordinate descent for composite empirical risk minimization,” in *Proc. Int. Conf. Mach. Learn.* PMLR, 2022, pp. 14 948–14 978

DP-CD (cont.,)

Algorithm 3 Differentially Private Coordinate Descent (DP-CD).

Input: Dataset $\mathcal{D}$, privacy parameters $\epsilon \in \mathbb{R}_{++}$ and $\delta \in (0,1)$, iteration budget $L \in \mathbb{N}$, initial point $\boldsymbol{\theta}^{(0)} \in \mathbb{R}^K$, Clipping constants $\{C_i\}_{i=1}^K$, and step sizes $\{\tau_i\}_{i=1}^K$

1: Determine noise parameters $\{\sigma_i\}_{i=1}^K$ for Gaussian and $\{\alpha_i, \gamma_i\}_{i=1}^K$ for flipped Huber from the privacy parameters.
2: **for** $l = 1, 2, \ldots, L$ **do**
3: $\quad \boldsymbol{\theta}^{(l)} \leftarrow \boldsymbol{\theta}^{(l-1)}$.
4: $\quad$ **for** $i = 1, 2, \ldots, K$ **do**
5: $\quad\quad$ Sample $t_i^{(l)} \sim \mathcal{N}(0, \sigma_j^2)$ or $t \sim \mathcal{FH}(\alpha_i, \gamma_i^2)$
6: $\quad\quad$ $\nu_i^{(l)} = \frac{1}{n} \sum_{n=1}^{N} \text{clip}\left(\nabla_i J(\boldsymbol{\theta}^{(l)}; \mathcal{D}_n); C_i\right) + t_i^{(l)}$
7: $\quad\quad$ $\theta_i^{(l)} \leftarrow \text{prox}_{\tau_i \psi_i}\left(\theta_i^{(l)} - \tau_i \nu_i^{(l)}\right)$
8: $\quad$ **end for**
9: **end for**

Output: $\widehat{\boldsymbol{\theta}} = \boldsymbol{\theta}^{(L)}$

DP-CD results

$\epsilon = 1 \qquad \delta = \frac{1}{N^2}$

	Dataset	Regularization and parameter	Gaussian NMSE	Gaussian Test error	Flipped Huber NMSE	Flipped Huber Test error
Logistic regression	Houses	$(\ell_2, 0.1)$	$0.6371 \cdot 10^{-3}$	0.0391	$0.6165 \cdot 10^{-3}$	0.0389
	Wine quality	$(\ell_2, 2 \cdot 10^{-4})$	0.2250	0.0614	0.1421	0.0513
	Pumpkin seeds	$(\ell_2, 0.1)$	0.0255	0.1224	0.0189	0.1152
	Heart	$(\ell_2, 0.1)$	0.2384	0.1741	0.1989	0.1556
Linear regression	California	$(\ell_1, 0.01)$	0.0479	0.4532	0.0465	0.4298
	Boston housing	$(\ell_1, 0.01)$	0.3579	0.3743	0.3406	0.3253
	Airfoil	$(\ell_1, 0.01)$	0.0206	0.5161	0.0190	0.4558
	Diabetes	$(\ell_1, 0.1)$	0.2515	0.5741	0.1489	0.4384

DP-CD results (cont.,)

Logistic Regression

Dataset	N	K
Houses	16512	8
Wine quality	5198	11
Pumpkin seeds	2000	12
Heart	216	13

Linear Regression

Dataset	N	K
California	16512	8
Boston housing	405	13
Airfoil	1202	5
Diabetes	354	10

Summary

The pros of flipped Huber

- Laplace: can lead to large amount of noise for large K and results in outliers
- Gaussian: light tailed, but renders least Fisher information
- Flipped Huber: Hybrid noise mechanism with density having lighter tails and sharper center
- More accurate for given privacy constraints compared to other mechanisms
  - Seems to significantly outperform in higher dimensions
  - Shows good results in real datasets e.g. private ERM
- Theoretically characterized
  - Necessary and sufficient conditions in one dimension
  - a sufficient condition in K dimension for (ϵ, δ)-DP
  - Composition using zCDP with application to CD

The cons of flipped Huber

- Requires several measures of sensitivities
  - Unknown Sensitivities can be loosely bounded
  - Cleverly handled by smart clipping in DP-CD
- In very high levels of composition, performs similar to Gaussian

Could flipped Huber be even better than stated?

- The sufficient condition in K dimension involves several bounds
  - Bounds loose for small ϵ
  - Bounds loose with increasing K
- zCDP is tight for Gaussian
  - Our composition results may be loose compared to composition results for Gaussian
  - We may be adding more noise than required

Some applications of DP in wireless systems

- Uplink channel estimation in cell-free MIMO[15]
  - Matrix completion for estimating channel with lesser number of pilots
  - Use DP Low rank matrix completion to protect user locations
- Wireless federated learning local DP (curator-free model)[16]
  - Superposition of gradients over non-orthogonal channel $\rightarrow$ more privacy

[15] J. Xu, X. Wang, P. Zhu, and X. You, "Privacy-preserving channel estimation in cell-free hybrid massive MIMO systems," *IEEE Trans. Wireless Commun.*, vol. 20, no. 6, pp. 3815–3830, 2021

[16] M. Seif, R. Tandon, and M. Li, "Wireless federated learning with local differential privacy," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, 2020, pp. 2604–2609

Some applications of DP in wireless systems (cont.,)

- Radio positioning and sensing[17]
  - DP through channel randomization and beam steering
- Energy harvesting through IRS[18]
  - Exponential mechanism for preserving location
- Edge computing over wireless big data[19]
  - Output perturbation and objective perturbation with Laplace noise

[17] V.-L. Nguyen, R.-H. Hwang, B.-C. Cheng, Y.-D. Lin, and T. Q. Duong, "Understanding privacy risks of high-accuracy radio positioning and sensing in wireless networks," *IEEE Commun. Mag.*, 2023

[18] Q. Pan, J. Wu, X. Zheng, W. Yang, and J. Li, "Differential privacy and irs empowered intelligent energy harvesting for 6g internet of things," *IEEE Internet Things J.*, vol. 9, no. 22, pp. 22 109–22 122, 2021

[19] M. Du, K. Wang, Z. Xia, and Y. Zhang, "Differential privacy preserving of training model in wireless big data with edge computing," *IEEE Trans. Big Data*, vol. 6, no. 2, pp. 283–295, 2018

Some applications of DP in wireless systems (cont.,)

- Split learning for integrated terrestrial and non-terrestrial networks[20]
  - Data owner and label owner train different parts of the deep learning model
- Cognitive radio networks[21]
  - DP in spectrum sensing, spectrum analysis, spectrum sharing
- Cyber physical systems[22] - time-series and statistical data
  - DP in smart grid, transportation, healthcare and IIoT

[20] M. Wu, G. Cheng, P. Li, R. Yu, Y. Wu, M. Pan, and R. Lu, "Split learning with differential privacy for integrated terrestrial and non-terrestrial networks," *IEEE Wireless Commun.*, 2023

[21] M. U. Hassan, M. H. Rehmani, M. Rehan, and J. Chen, "Differential privacy in cognitive radio networks: a comprehensive survey," *Cogn. Comput.*, vol. 14, no. 2, pp. 475–510, 2022

[22] M. U. Hassan, M. H. Rehmani, and J. Chen, "Differential privacy techniques for cyber physical systems: A survey," *IEEE Commun. Surveys Tuts.*, vol. 22, no. 1, pp. 746–789, 2020

Epilogue

Improving accuracy with i.n.i.d. noise

- i.i.d. noise: noise parameters depend on overall sensitivity measure

[23] G. Muthukrishnan and S. Kalyani, “Differential privacy with higher utility by exploiting coordinate-wise disparity: Laplace mechanism can beat Gaussian in high dimensions,” *arXiv:2302.03511*, 2024

Improving accuracy with i.n.i.d. noise

- i.i.d. noise: noise parameters depend on overall sensitivity measure
- Sensitivity of i-th coordinate of query response – λ_i

[23] G. Muthukrishnan and S. Kalyani, "Differential privacy with higher utility by exploiting coordinate-wise disparity: Laplace mechanism can beat Gaussian in high dimensions," *arXiv:2302.03511*, 2024

Improving accuracy with i.n.i.d. noise

- i.i.d. noise: noise parameters depend on overall sensitivity measure
- Sensitivity of i-th coordinate of query response – λ_i
- Whenever there is disparity in $\{\lambda_i\}_{i=1}^{K}$, performance can be improved

[23] G. Muthukrishnan and S. Kalyani, "Differential privacy with higher utility by exploiting coordinate-wise disparity: Laplace mechanism can beat Gaussian in high dimensions," *arXiv:2302.03511*, 2024

Improving accuracy with i.n.i.d. noise

- i.i.d. noise: noise parameters depend on overall sensitivity measure
- Sensitivity of i-th coordinate of query response – λ_i
- Whenever there is disparity in $\{\lambda_i\}_{i=1}^{K}$, performance can be improved
- Add non-identical (but still independent) noise samples[23] across coordinates

[23] G. Muthukrishnan and S. Kalyani, "Differential privacy with higher utility by exploiting coordinate-wise disparity: Laplace mechanism can beat Gaussian in high dimensions," *arXiv:2302.03511*, 2024

Improving accuracy with i.n.i.d. noise

- i.i.d. noise: noise parameters depend on overall sensitivity measure
- Sensitivity of i-th coordinate of query response $-$ λ_i
- Whenever there is disparity in $\{\lambda_i\}_{i=1}^{K}$, performance can be improved
- Add non-identical (but still independent) noise samples[23] across coordinates
- Gaussian and Laplace - lesser noise for more dispersed $\{\lambda_i\}_{i=1}^{K}$

[23] G. Muthukrishnan and S. Kalyani, "Differential privacy with higher utility by exploiting coordinate-wise disparity: Laplace mechanism can beat Gaussian in high dimensions," *arXiv:2302.03511*, 2024

Results

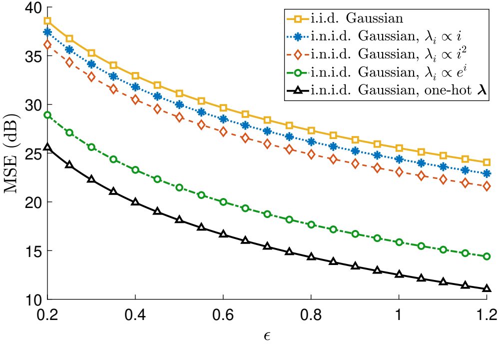

Gaussian: $K = 20$ and $\delta = 10^{-6}$

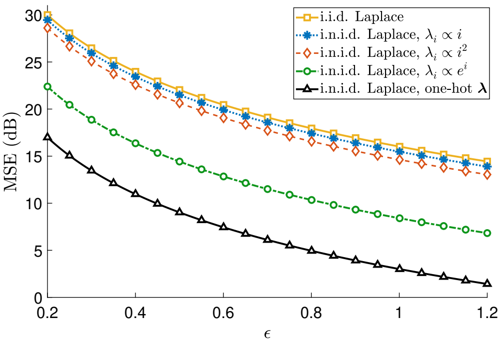

Laplace: $K = 20$ and $\delta = 0$

Results (cont.,)

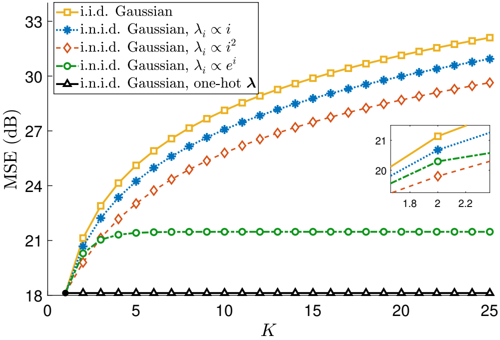

Gaussian: $\epsilon = 0.5$ and $\delta = 10^{-6}$

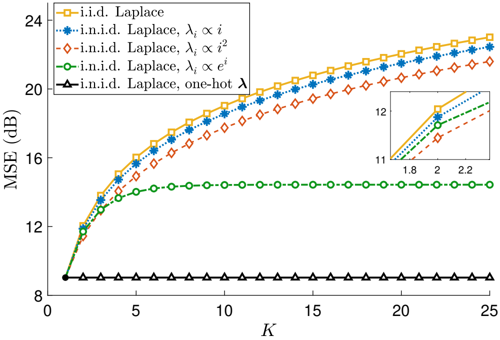

Laplace: $\epsilon = 0.5$ and $\delta = 0$

Results (cont.,)

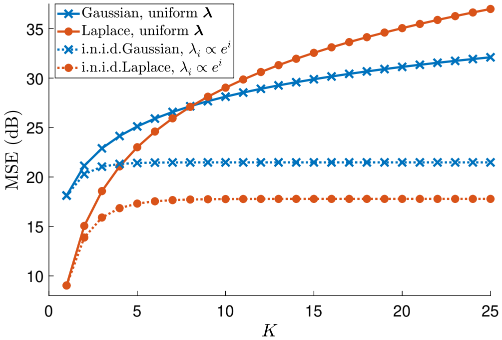

Comparison of $(0.5, 10^{-6})$-DP Gaussian and $(0.5, 0)$-DP Laplace

Improved DP-CD

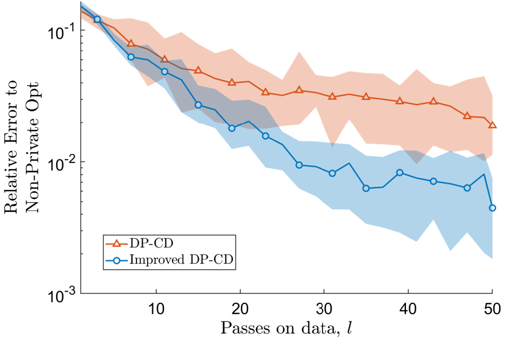

California dataset:
ℓ_1-regularized linear regression

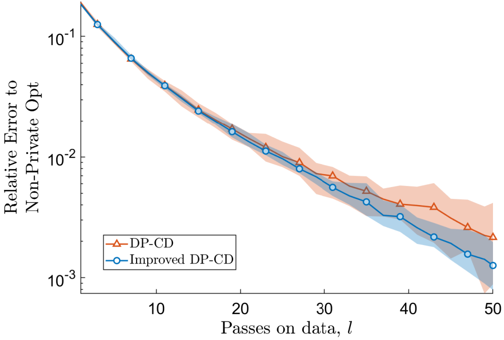

Electricity dataset:
ℓ_2-regularized logistic regression

References

[1] C. Dwork, F. McSherry, K. Nissim, and A. Smith, “Calibrating noise to sensitivity in private data analysis,” in *Proc. Theory Cryptogr. Conf.* Springer, 2006, pp. 265–284.

[2] B. Balle and Y.-X. Wang, “Improving the Gaussian mechanism for differential privacy: Analytical calibration and optimal denoising,” in *Proc. Int. Conf. Mach. Learn.* PMLR, 2018, pp. 394–403.

[3] C. Dwork, G. N. Rothblum, and S. Vadhan, “Boosting and differential privacy,” in *Proc. IEEE Annu. Symp. Found. Comput. Sci.* IEEE, 2010, pp. 51–60.

[4] P. J. Huber and E. M. Ronchetti, *Robust statistics*. John Wiley & Sons, 2009.

[5] Q. Geng, W. Ding, R. Guo, and S. Kumar, “Tight analysis of privacy and utility tradeoff in approximate differential privacy,” in *Proc. Int. Conf. Artif. Intell. Statist.* PMLR, 2020, pp. 89–99.

[6] P. Sadeghi and M. Korki, "Offset-symmetric Gaussians for differential privacy," *IEEE Trans. Inf. Forensics Security*, vol. 17, pp. 2394–2409, 2022.

[7] Q. Geng and P. Viswanath, "The optimal noise-adding mechanism in differential privacy," *IEEE Trans. Inf. Theory*, vol. 62, no. 2, pp. 925–951, 2016.

[8] Y.-X. Wang, "Revisiting differentially private linear regression: optimal and adaptive prediction & estimation in unbounded domain," in *Uncertainty in Artif. Intell.*, 2018.

[9] C. Dwork, K. Talwar, A. Thakurta, and L. Zhang, "Analyze Gauss: optimal bounds for privacy-preserving principal component analysis," in *Proc. Annu. ACM Symp. Theory of Comput.*, 2014, pp. 11–20.

[10] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, "Deep learning with differential privacy," in *Proc. ACM SIGSAC Conf. Computer and Communications security*, 2016, pp. 308–318.

[11] Q. Geng, P. Kairouz, S. Oh, and P. Viswanath, “The staircase mechanism in differential privacy,” *IEEE J. Sel. Topics Signal Process.*, vol. 9, no. 7, pp. 1176–1184, 2015.

[12] M. Bun and T. Steinke, “Concentrated differential privacy: Simplifications, extensions, and lower bounds,” in *Proc. Int. Conf. Theory of Cryptogr. Part I.* Springer, 2016, pp. 635–658.

[13] P. Mangold, A. Bellet, J. Salmon, and M. Tommasi, “Differentially private coordinate descent for composite empirical risk minimization,” in *Proc. Int. Conf. Mach. Learn.* PMLR, 2022, pp. 14 948–14 978.

[14] J. Xu, X. Wang, P. Zhu, and X. You, “Privacy-preserving channel estimation in cell-free hybrid massive MIMO systems,” *IEEE Trans. Wireless Commun.*, vol. 20, no. 6, pp. 3815–3830, 2021.

[15] M. Seif, R. Tandon, and M. Li, “Wireless federated learning with local differential privacy,” in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, 2020, pp. 2604–2609.

[16] V.-L. Nguyen, R.-H. Hwang, B.-C. Cheng, Y.-D. Lin, and T. Q. Duong, "Understanding privacy risks of high-accuracy radio positioning and sensing in wireless networks," *IEEE Commun. Mag.*, 2023.

[17] Q. Pan, J. Wu, X. Zheng, W. Yang, and J. Li, "Differential privacy and irs empowered intelligent energy harvesting for 6g internet of things," *IEEE Internet Things J.*, vol. 9, no. 22, pp. 22 109–22 122, 2021.

[18] M. Du, K. Wang, Z. Xia, and Y. Zhang, "Differential privacy preserving of training model in wireless big data with edge computing," *IEEE Trans. Big Data*, vol. 6, no. 2, pp. 283–295, 2018.

[19] M. Wu, G. Cheng, P. Li, R. Yu, Y. Wu, M. Pan, and R. Lu, "Split learning with differential privacy for integrated terrestrial and non-terrestrial networks," *IEEE Wireless Commun.*, 2023.

[20] M. U. Hassan, M. H. Rehmani, M. Rehan, and J. Chen, "Differential privacy in cognitive radio networks: a comprehensive survey," *Cogn. Comput.*, vol. 14, no. 2, pp. 475–510, 2022.

[21] M. U. Hassan, M. H. Rehmani, and J. Chen, “Differential privacy techniques for cyber physical systems: A survey,” *IEEE Commun. Surveys Tuts.*, vol. 22, no. 1, pp. 746–789, 2020.

[22] G. Muthukrishnan and S. Kalyani, “Differential privacy with higher utility by exploiting coordinate-wise disparity: Laplace mechanism can beat Gaussian in high dimensions,” *arXiv:2302.03511*, 2024.

[23] F. Liu, “Generalized Gaussian mechanism for differential privacy,” *IEEE Trans. Knowledge and Data Engg.*, vol. 31, no. 4, pp. 747–756, 2018.

[24] C. L. Canonne, G. Kamath, and T. Steinke, “The discrete Gaussian for differential privacy,” in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 33. PMLR, 2020, pp. 15 676–15 688.

[25] J. Awan and A. Slavković, “Structure and sensitivity in differential privacy: Comparing K-norm mechanisms,” *J. Amer. Stat. Assoc.*, vol. 116, no. 534, pp. 935–954, 2021.

Thank You!

Noise mechanisms in literature

- Laplace mechanism[24]: noise sampled from density $\frac{1}{2\beta}\exp\left(-\frac{|x|}{\beta}\right)$
  - ϵ-DP for $\beta \geq \frac{\Delta_1}{\epsilon}$
- Gaussian mechanism[25]: noise sampled from density $\frac{1}{\sqrt{2\pi\sigma^2}}\exp\left(-\frac{x^2}{2\sigma^2}\right)$
  - (ϵ, δ)-DP for $\sigma \geq \sigma_0$, where $Q\left(\frac{\sigma_0\epsilon}{\Delta_2} - \frac{\Delta_2}{2\sigma_0}\right) - e^{\epsilon}Q\left(\frac{\sigma_0\epsilon}{\Delta_2} + \frac{\Delta_2}{2\sigma_0}\right) = \delta$
- OSGT mechanism[26]: noise sampled from density $\frac{1}{2Q(\frac{\vartheta}{\varrho})}\phi(|t|\,; -\vartheta,\ \varrho^2)$

[24] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Proc. Theory Cryptogr. Conf.* Springer, 2006, pp. 265–284

[25] B. Balle and Y.-X. Wang, "Improving the Gaussian mechanism for differential privacy: Analytical calibration and optimal denoising," in *Proc. Int. Conf. Mach. Learn.* PMLR, 2018, pp. 394–403

[26] P. Sadeghi and M. Korki, "Offset-symmetric Gaussians for differential privacy," *IEEE Trans. Inf. Forensics Security*, vol. 17, pp. 2394–2409, 2022

Noise mechanisms in literature (cont.,)

- Subbotin or generalized Gaussian mechanism[27]: noise density $\frac{p^{1-\frac{1}{p}}}{2\,\xi\,\Gamma\left(\frac{1}{p}\right)} \exp\left(-\frac{|x|^p}{p\xi^p}\right)$
- Discrete Gaussian mechanism[28]
- K-norm mechanism[29]: noise density for ϵ-DP $\rightarrow \frac{1}{\Gamma(K+1)\lambda\left(\frac{\Delta}{\epsilon}\mathcal{K}\right)} \exp\left(-\frac{\epsilon}{\Delta}\|\mathbf{x}\|_{\mathcal{K}}\right)$
  - Difficult to characterize sensitivity space and construct $\mathcal{K}$

[27] F. Liu, "Generalized Gaussian mechanism for differential privacy," *IEEE Trans. Knowledge and Data Engg.*, vol. 31, no. 4, pp. 747–756, 2018

[28] C. L. Canonne, G. Kamath, and T. Steinke, "The discrete Gaussian for differential privacy," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 33. PMLR, 2020, pp. 15 676–15 688

[29] J. Awan and A. Slavković, "Structure and sensitivity in differential privacy: Comparing K-norm mechanisms," *J. Amer. Stat. Assoc.*, vol. 116, no. 534, pp. 935–954, 2021

Optimal DP noise mechanisms

- Staircase mechanism: optimal ϵ-DP mechanism for one-dimensional queries[30]
  - Laplace is optimal ϵ-DP mechanism for small ϵ
- Staircase is the optimal noise for ϵ-DP (under ℓ_1-error) in two dimensions[31]
- Truncated Laplace: optimal (ϵ, δ)-DP mechanism for one-dimensional queries[32]
  - Optimal in high privacy regime $(\epsilon, \delta) \to (0, 0)$
  - Bounded support $\to$ $\text{supp}\left(\mathcal{M}(\mathcal{D})\right) \setminus \text{supp}\left(\mathcal{M}(\check{\mathcal{D}})\right)$ is non empty
    - Can perfectly distinguish $\mathcal{D}$ and $\check{\mathcal{D}}$ with probability up to δ

[30] Q. Geng and P. Viswanath, "The optimal noise-adding mechanism in differential privacy," *IEEE Trans. Inf. Theory*, vol. 62, no. 2, pp. 925–951, 2016

[31] Q. Geng, P. Kairouz, S. Oh, and P. Viswanath, "The staircase mechanism in differential privacy," *IEEE J. Sel. Topics Signal Process.*, vol. 9, no. 7, pp. 1176–1184, 2015

[32] Q. Geng, W. Ding, R. Guo, and S. Kumar, "Tight analysis of privacy and utility tradeoff in approximate differential privacy," in *Proc. Int. Conf. Artif. Intell. Statist.* PMLR, 2020, pp. 89–99