Privacy Amplification by Subsampling and Renyi Differential Privacy

Yu-Xiang Wang
UC Santa Barbara

Joint work with Borja Balle and Shiva Kasiviswanathan

Outline

- Preliminary:
  - From DP to Renyi DP
  - Subsampled mechanisms and Privacy amplification
- Renyi DP of Subsampled Algorithms
- Composition and Analytical moments accountant
- Proof ideas
- Open problems

Renyi DP and algorithm-specific DP analysis

- Ɛ-DP is a one number summary of the privacy guarantee

$$\log \frac{p_{\mathcal{M}}(X)(h)}{p_{\mathcal{M}}(X')(h)} \leq \epsilon$$

- RDP (Mironov, 2017) characterizes the full-distribution of the privacy R.V. induced by a specific algorithm

$$D_\alpha(\mathcal{M}(X)\|\mathcal{M}(X')) = \frac{1}{\alpha - 1}\log(\mathrm{MGF}_\epsilon(\alpha - 1)) \leq \epsilon(\alpha)$$

  - Also closely related to CDP (Dwork & Rothblum,2016) and zCDP (Bun & Steinke,2016)

Renyi DP is natural for composition

- Compose linearly $$\epsilon_{\mathcal{M}_1 \times \mathcal{M}_2}(\alpha) = \epsilon_{\mathcal{M}_1}(\alpha) + \epsilon_{\mathcal{M}_2}(\alpha)$$
- RDP => (Ɛ, δ)-DP

$$\begin{aligned} \delta \Rightarrow \epsilon: \qquad & \epsilon(\delta) = \min_{\alpha > 1} \frac{\log(1/\delta)}{\alpha - 1} + \epsilon_{\mathcal{M}}(\alpha - 1), \\ \epsilon \Rightarrow \delta: \qquad & \delta(\epsilon) = \min_{\alpha > 1} e^{(\alpha-1)(\epsilon_{\mathcal{M}}(\alpha-1)-\epsilon)}. \end{aligned}$$

- Comparing to the composition theorems for (Ɛ, δ)-DP
  - Cleaner, no need to choose individual ($\mathcal{E}_i$, δ_i)
  - Elegantly handle the advanced composition of heterogenous mechanisms.
  - Efficiently computable, nothing #P-complete. (Murtagh&Vadhan, 2017)
  - Often better than the optimal composition with just ($\mathcal{E}_i$, δ_i)-DP.

Increasing list of mechanisms where we know how to precisely calculate their RDP

$$\epsilon_{\mathrm{Gaussian}(\alpha)} = \frac{\alpha}{2\sigma^2},$$

$$\epsilon_{\mathrm{Laplace}(\alpha)} = \frac{1}{\alpha-1}\log\left(\left(\frac{\alpha}{2\alpha-1}\right)e^{(\alpha-1)/\lambda} + \left(\frac{\alpha-1}{2\alpha-1}\right)e^{-\alpha/\lambda}\right) \text{ for } \alpha > 1,$$

$$\epsilon_{\mathrm{RandResp}(\alpha)} = \frac{1}{\alpha-1}\log\left(p^{\alpha}(1-p)^{1-\alpha} + (1-p)^{\alpha}p^{1-\alpha}\right) \text{ for } \alpha > 1.$$

Many DP mechanisms that samples from an exponential family distribution have their RDP readily available in closed-form. (Geumlek, Song, Chaudhuri, 2017)

Subsampled Randomized Algorithm

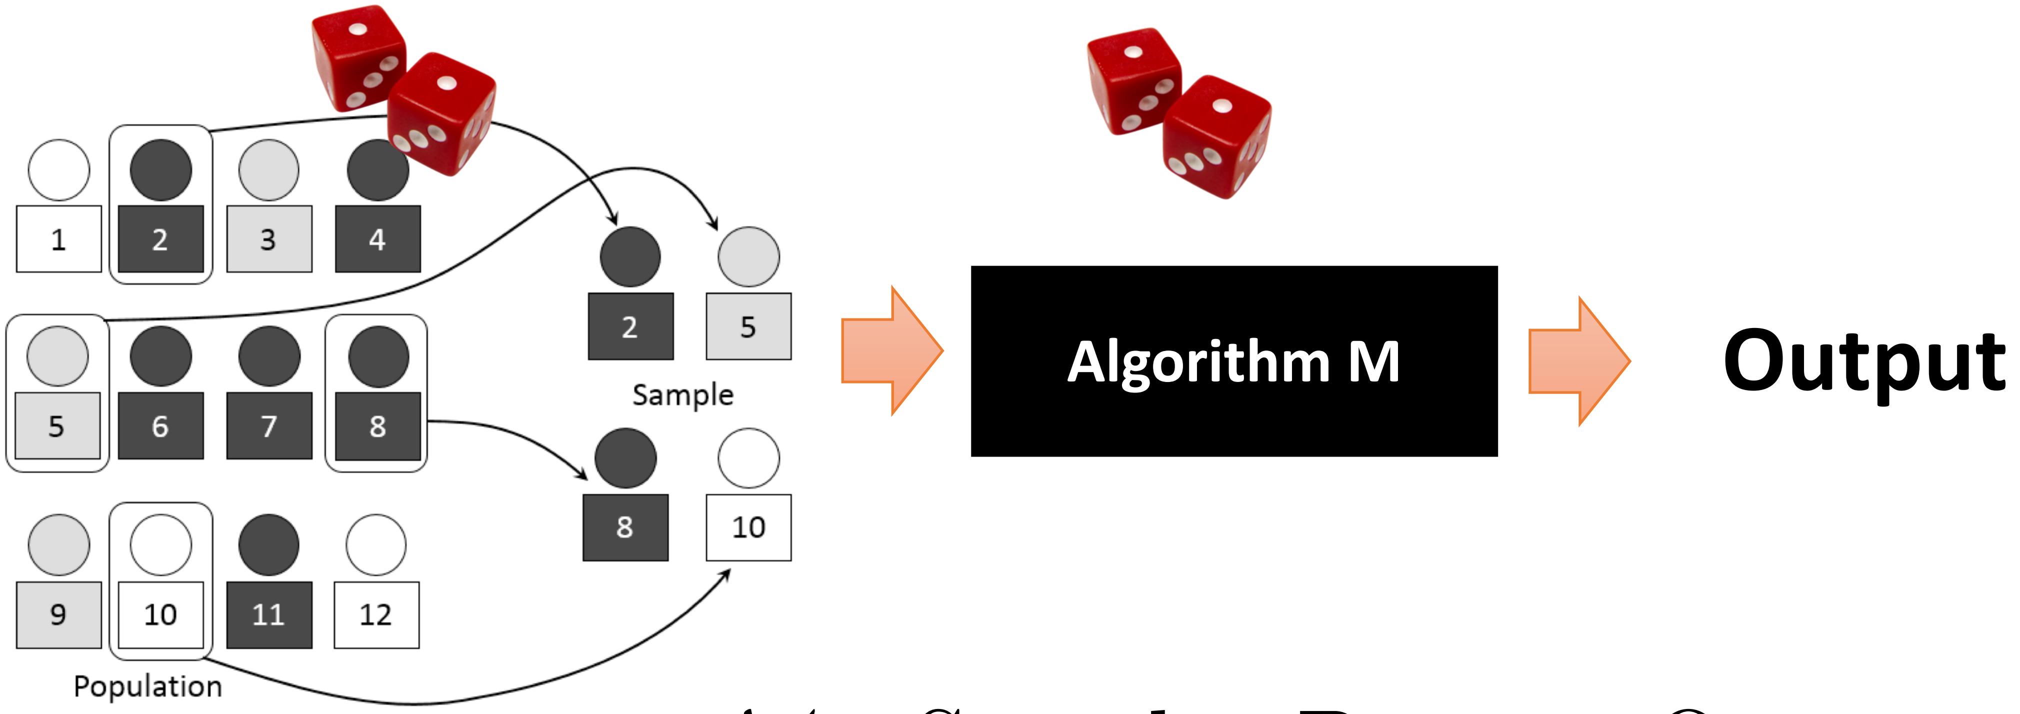

$$\mathcal{M} \circ \mathrm{Sample} : \mathrm{Data} \to \mathrm{Output}$$

Example: The Noisy SGD algorithm (Song et al. 2013; Bassily et. al. 2014)

$$\theta_{t+1} \leftarrow \theta_t - \eta_t \left(\frac{1}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} \nabla f_i(\theta_t) + Z_t \right)$$

- Randomly chosen minibatch (Subsampling)
- Then add gaussian noise (Gaussian mechanism)
- RDP analysis for subsampled Gaussian mechanism (Abadi et al., 2016)
  - Really what makes Deep Learning with Differential Privacy practical.

More general use of subsampling in algorithm designs

- Ensemble learning with Bagging / Random Forest / Boosting (Breiman)

- Bootstraps, Jackknife, subsampling bootstrap (Efron; Stein; Politis and Romano)

- Sublinear time algorithms in exploratory data analysis
  - Sketching, mean, quantiles, data cleaning.

Do we have to do these on a case-by-case basis?

Privacy “amplification” by subsampling

Subsampling Lemma: If M obeys (Ɛ,δ)-DP, then M ∘ Subsample obeys that (Ɛ’,δ’)-DP with $\delta' = \gamma\delta$

$$\epsilon' = \log(1 + \gamma(e^{\epsilon} - 1)) = O(\gamma\epsilon)$$

- First seen in “What can we learn privately?” (Kasiviswanathan et al., 2008)
- Subsequently used as a fundamental technical tool for learning theory with DP:
  - (Beimel et al., 2013) (Bun et al, 2015) (Wang et al., 2016)
- Most recent “tightened” revision above in:
  - Borja Balle, Gilles Barthe, Marco Gaboardi (2018)

This work: Privacy amplification by subsampling using Renyi Differential Privacy

- Can we prove a similar theorem for RDP?
  - Laplace mech., Randomized responses, posterior sampling and etc.
  - New tool in DP algorithm design.
  - Explicit constant.

Two different types of subsampling

- Sampling without replacement
  - Random subset of size m from a data set of size n
  - Replace-one version of DP

- Poisson sampling
  - Each data point is included independently with probability
  - Equivalent to $m \sim \text{Binomial}(\gamma, n)$, then sample without replacement.
  - Add-remove version of DP
  - The mechanism M needs to be well-defined for all data size

A subsampled mechanism samples from a mixture distribution with many mixture components!

- X’ <- Subsample(X)
- h <- f(X’) + Noise

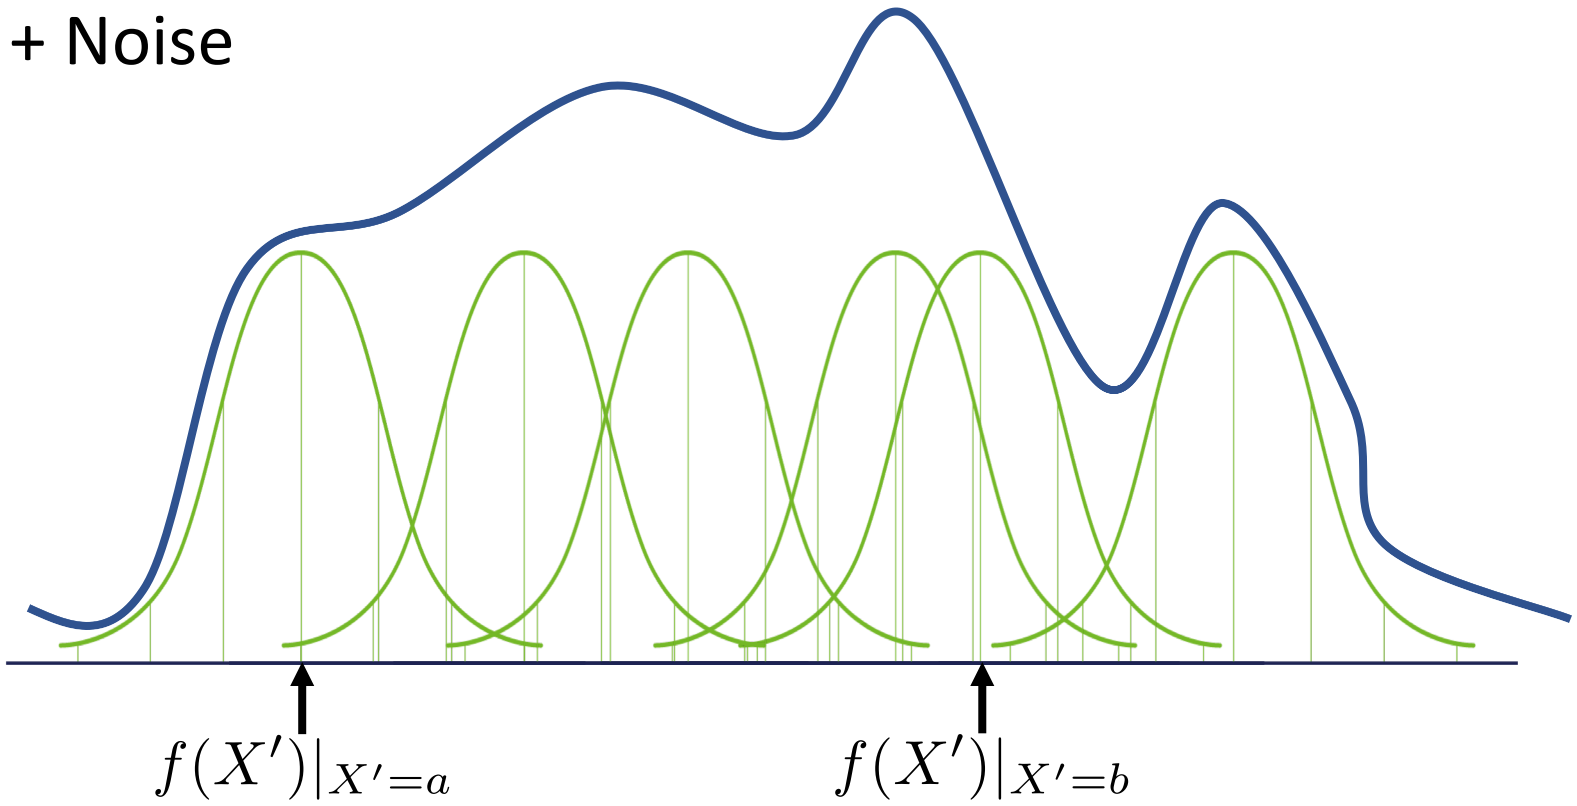

Changing to an adjacent data set

- X’ <- Subsample(X)
- h <- f(X’) + Noise

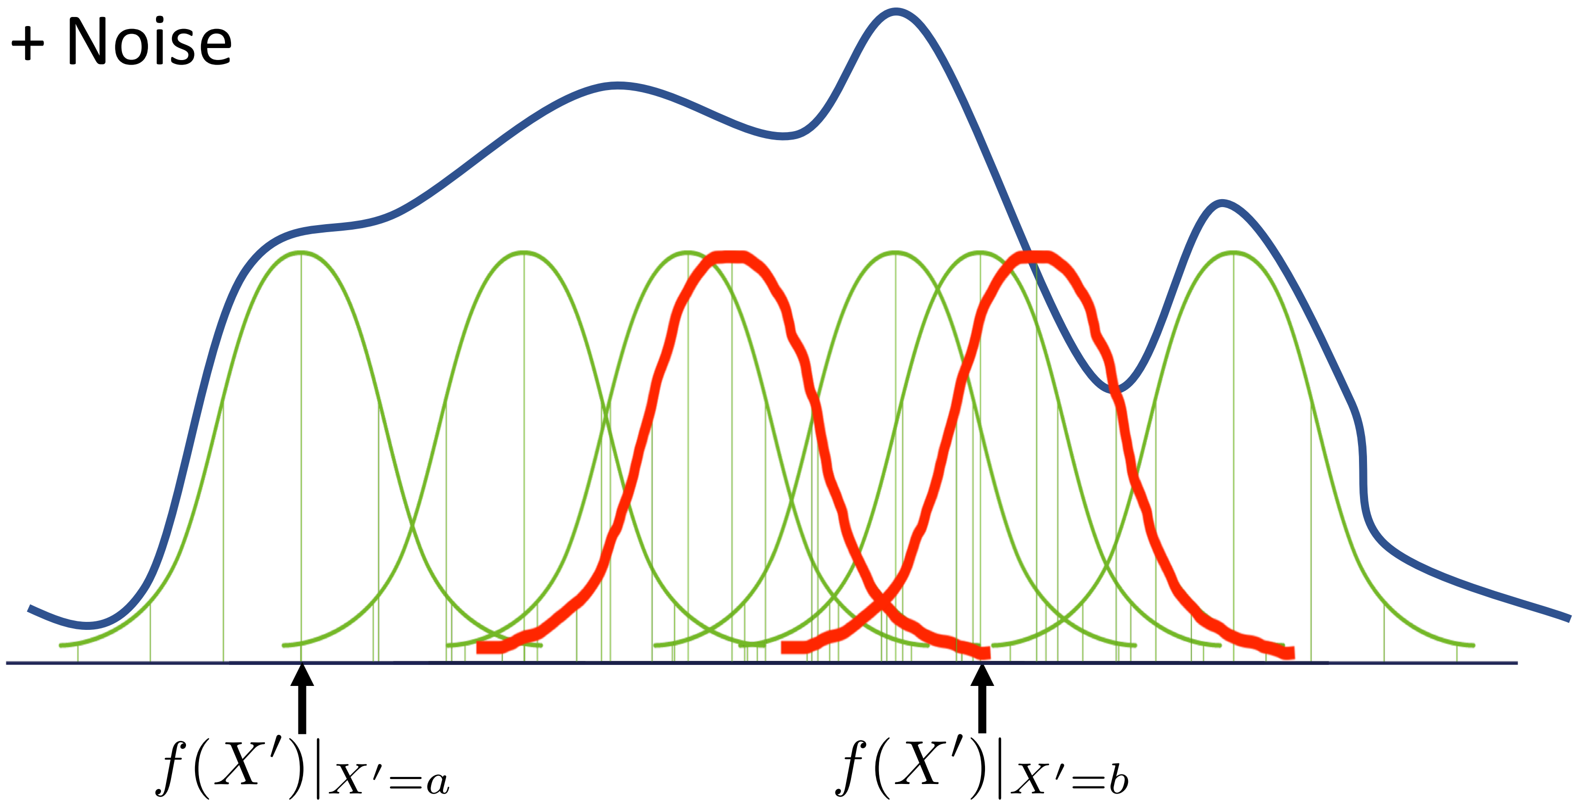

Changing to an adjacent data set

- X’ <- Subsample(X)
- h <- f(X’) + Noise

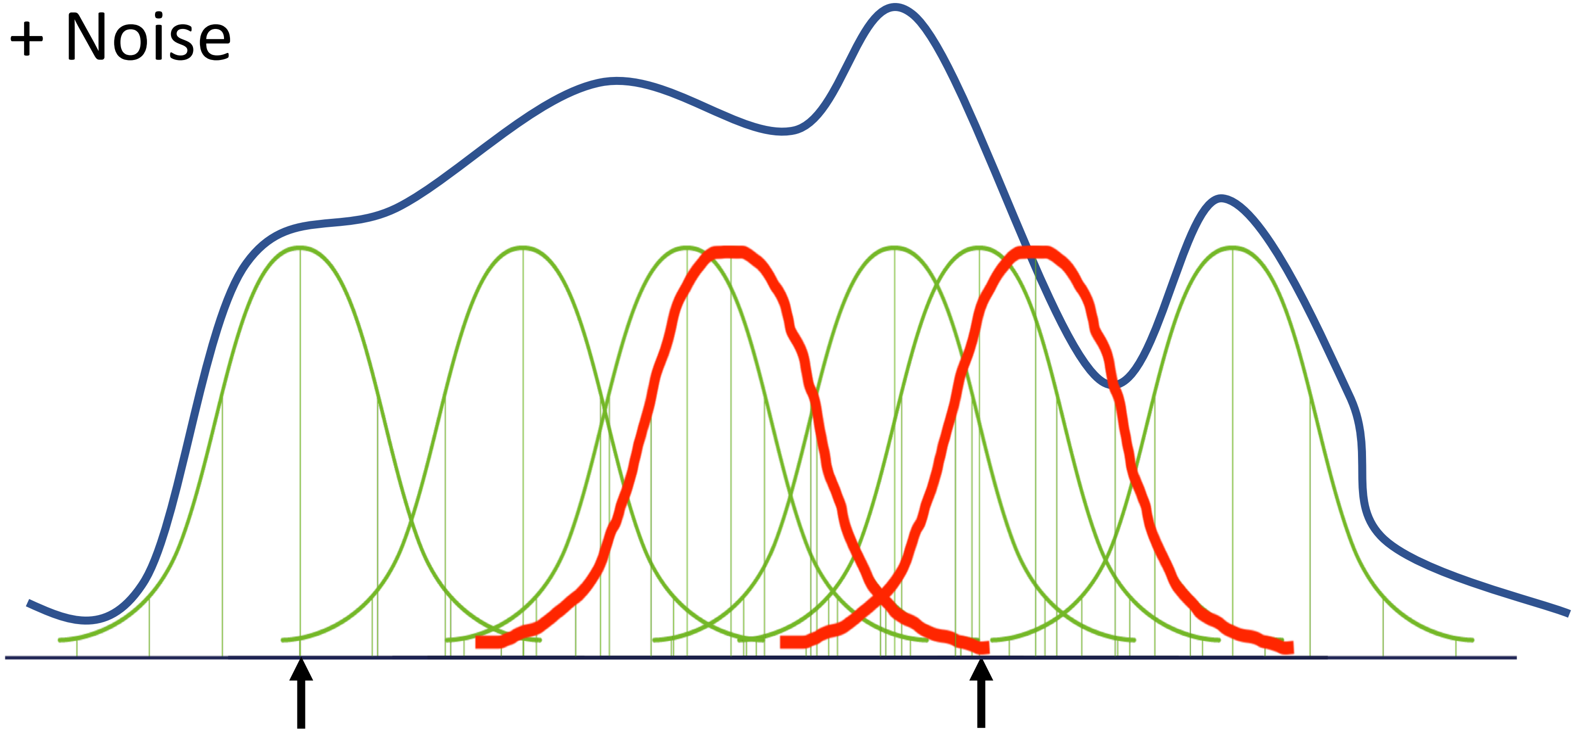

$\mathbb{E}[\mathcal{M}(X[:, \mathrm{idx}])|\mathrm{idx} = a]$ $\mathbb{E}[\mathcal{M}(X[:, \mathrm{idx}])|\mathrm{idx} = b]$

Main technical results

Theorem (Upper bound): Let M obeys (α ,Ɛ(α))-RDP for all α. Then M(subsample(DATA)) obeys

$$\epsilon'(\alpha) \le \frac{1}{\alpha-1}\log\Bigg(1+\gamma^2\binom{\alpha}{2}\min\Big\{4(e^{\epsilon(2)}-1), e^{\epsilon(2)}\min\{2,(e^{\epsilon(\infty)}-1)^2\}\Big\} + \sum_{j=3}^{\alpha}\gamma^j\binom{\alpha}{j}e^{(j-1)\epsilon(j)}\min\{2,(e^{\epsilon(\infty)}-1)^j\}\Bigg).$$

Theorem (lower bound): Let M satisfies some mild conditions

$$\epsilon'(\alpha) \ge \frac{\alpha}{\alpha-1}\log(1-\gamma) + \frac{1}{\alpha-1}\log\Big(1+\alpha\frac{\gamma}{1-\gamma}+\sum_{j=2}^{\alpha}\binom{\alpha}{j}\big(\frac{\gamma}{1-\gamma}\big)^j e^{(j-1)\epsilon(j)}\Big).$$

Numerical evaluation of the bounds

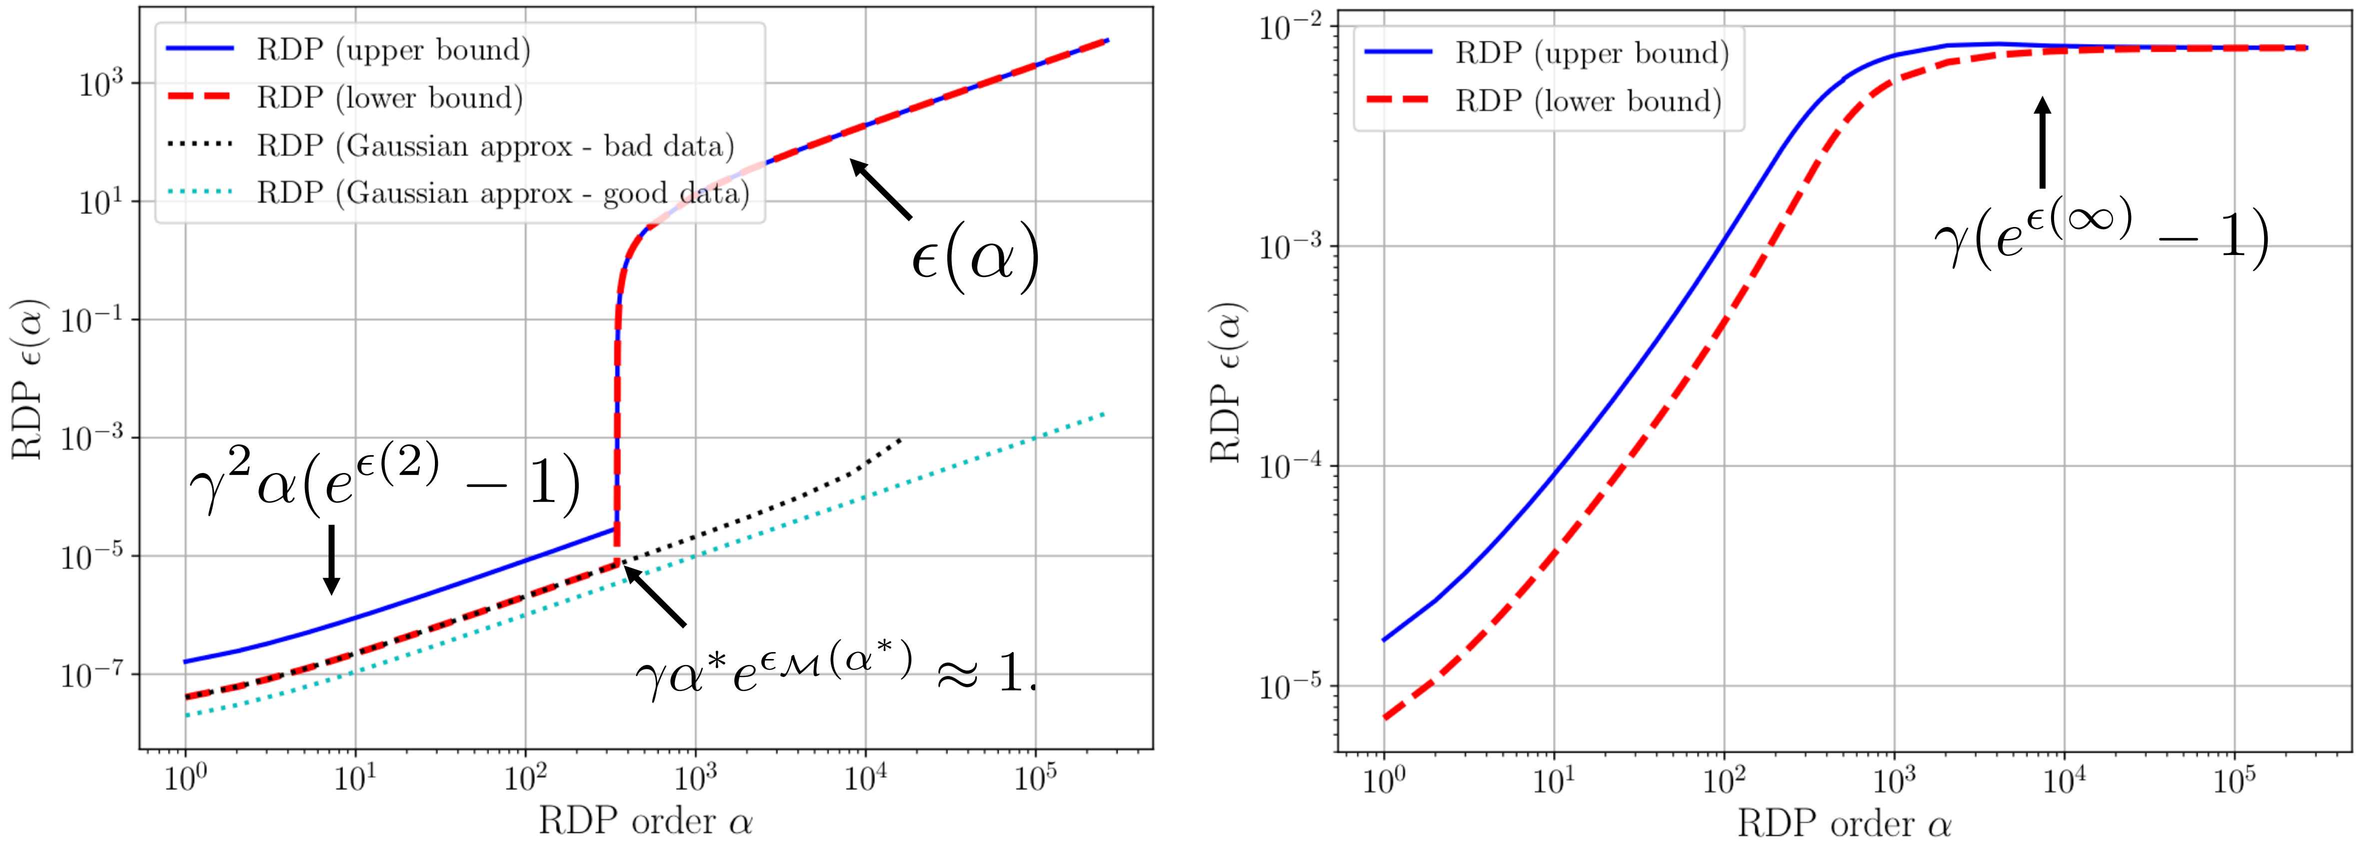

(a) RDP of Subsampled Gaussian with σ = 5

(b) RDP of Subsampled Laplace with b = 0.5

Comparing to zCDP and tCDP

- zCDP: linear upper bound of the entire RDP function
  - Doesn’t get amplified by subsampling

- tCDP: linear upper bound of the RDP up to a fixed threshold
  - Does get amplified by subsampling

- Not able to capture the fine-grained shape

Analytical moments accountant

- Tracking RDP for all order as a symbolic functions.
- Numerical calculations for (Ɛ, δ)-DP guarantees.
- Automatically DP calculations for complex algorithms.
- Enable state-of-the-art DP for non-experts.

Open source project:
https://github.com/yuxiangw/autodp

pip install autodp

Using our bounds for advanced composition

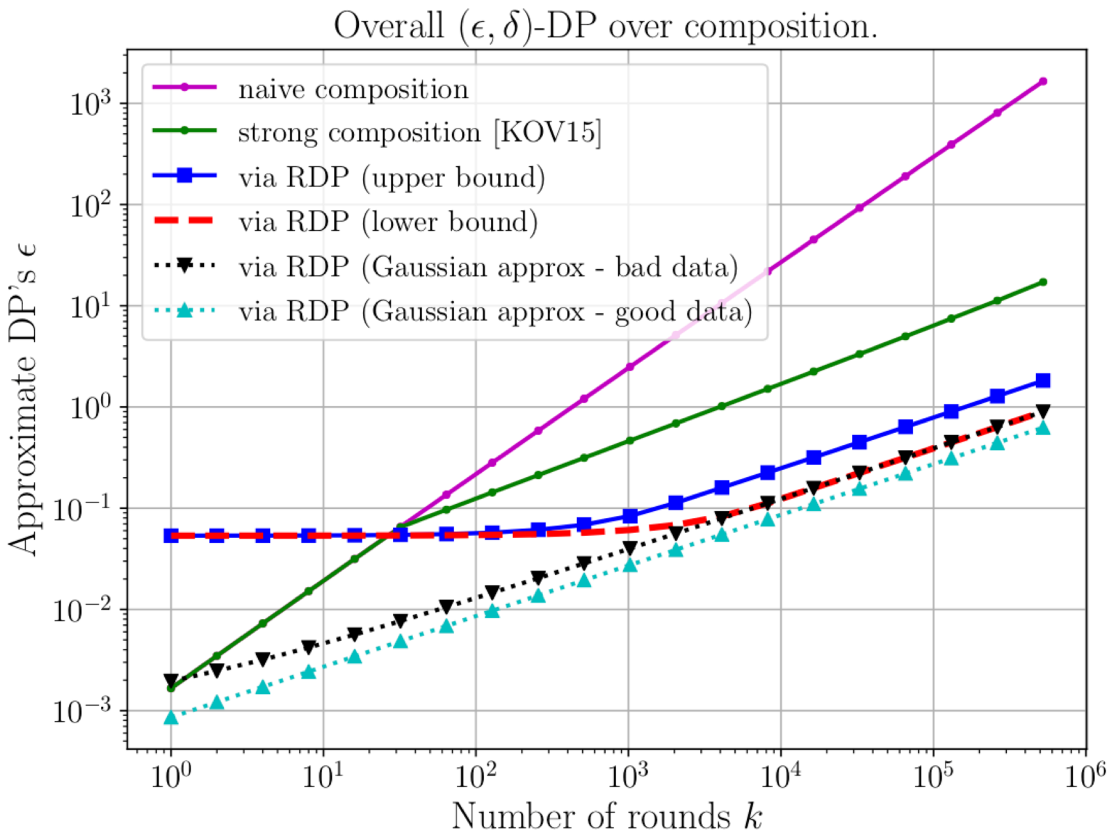

(a) Subsampled Gaussian with σ = 5

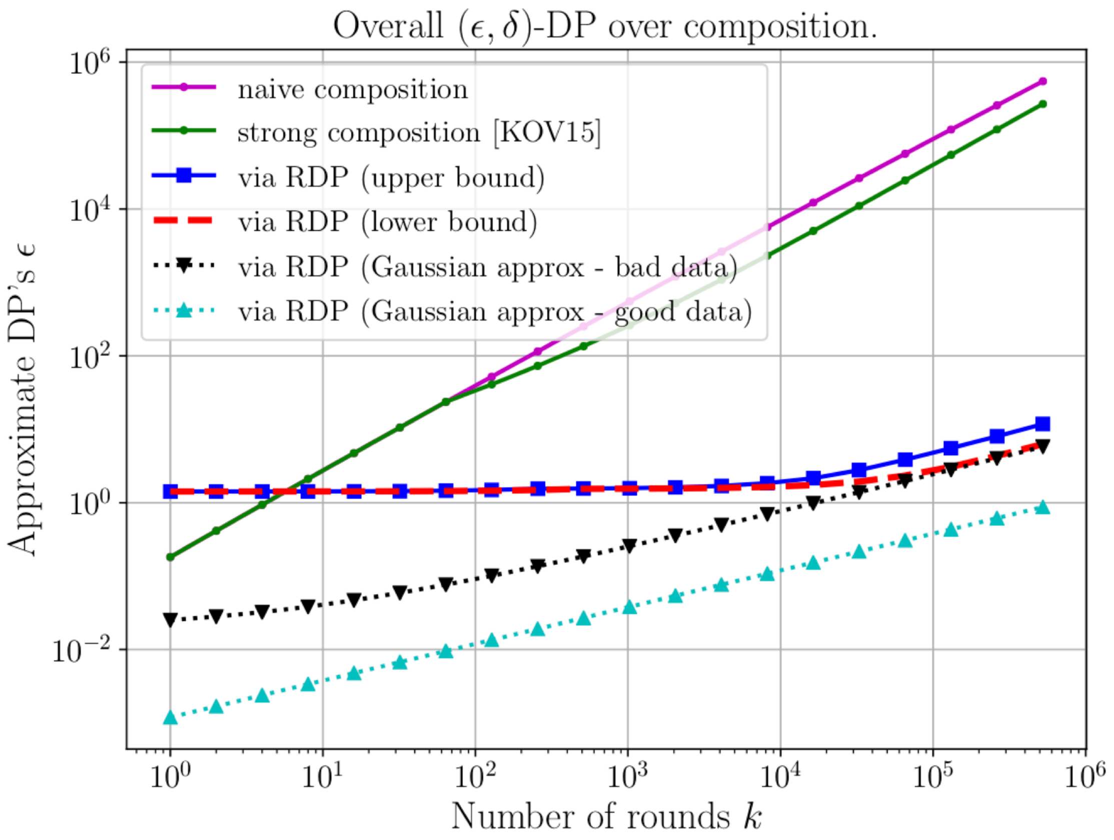

(a) Subsampled Gaussian with σ = 0.5

Using our bounds for advanced composition

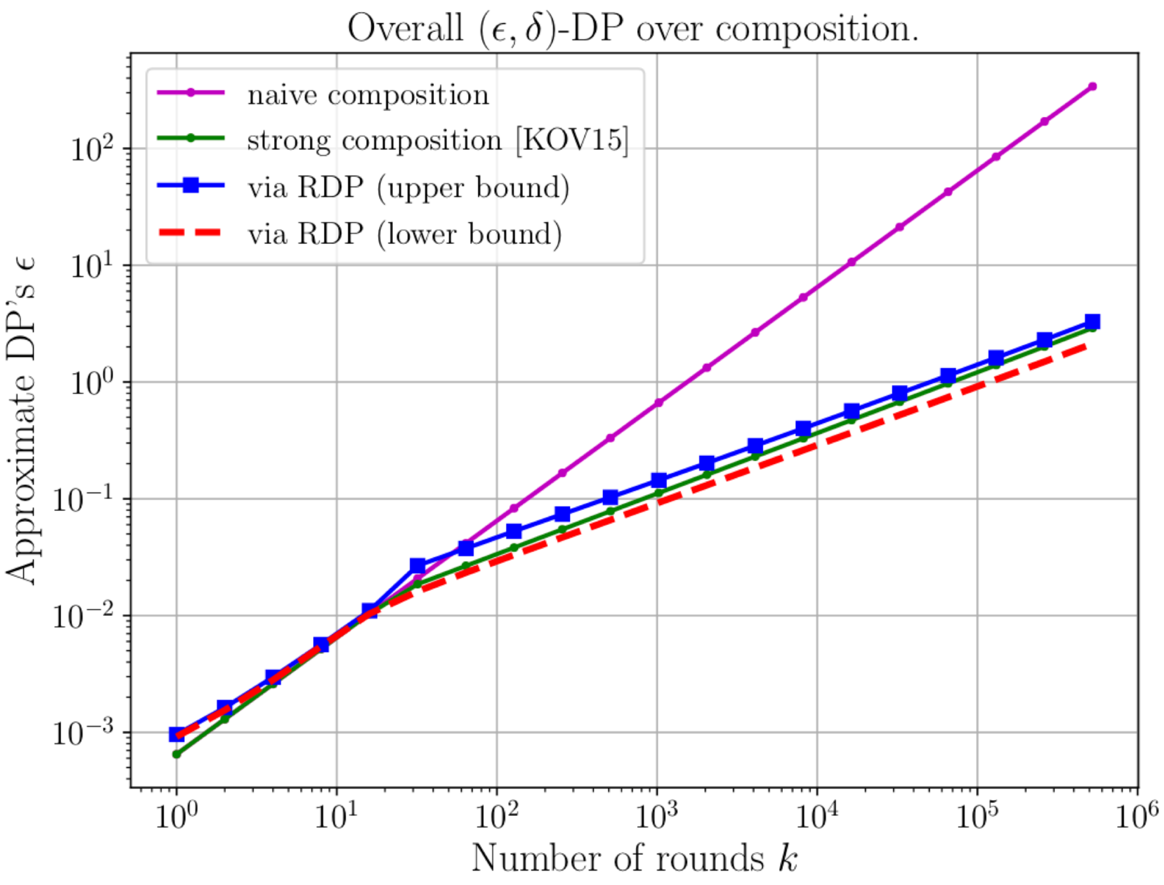

(c) Subsampled Laplace with b = 2

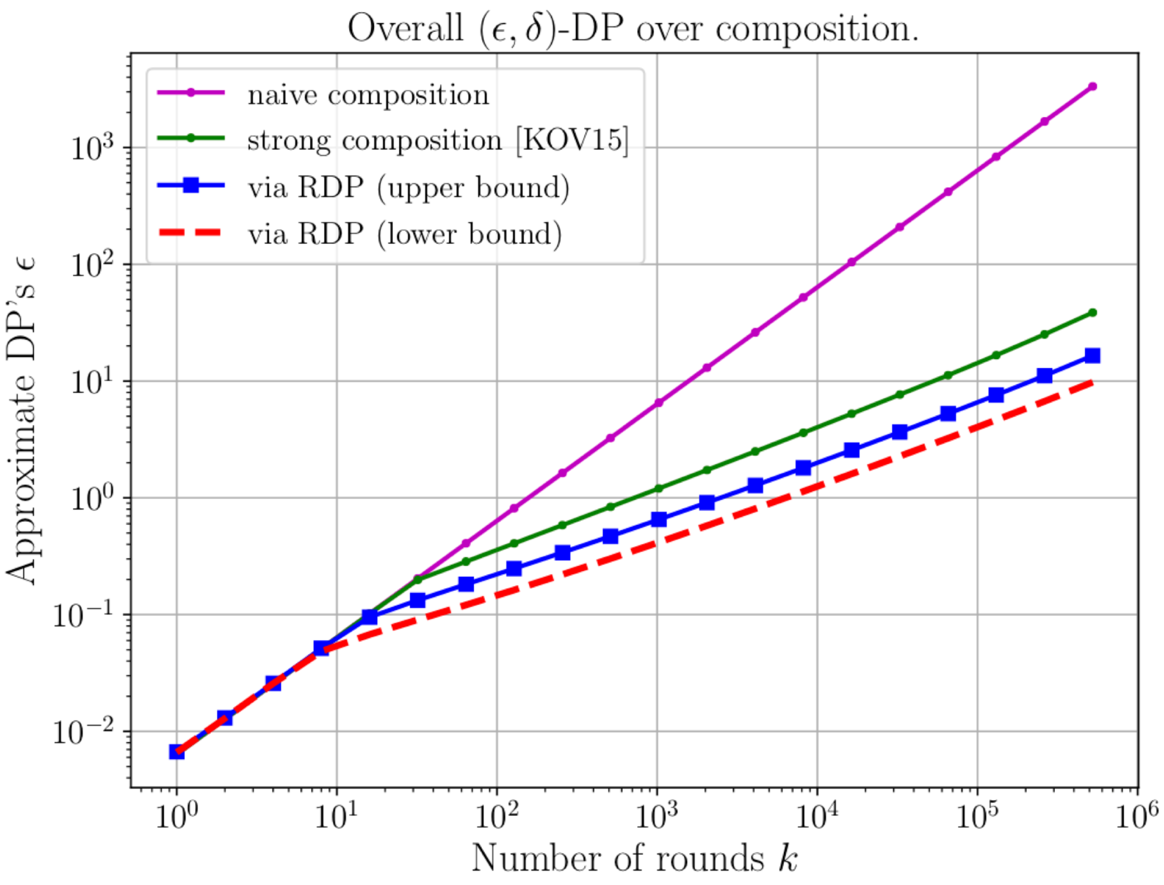

(d) Subsampled Laplace with b = 0.5

Proof idea (Upper bound)

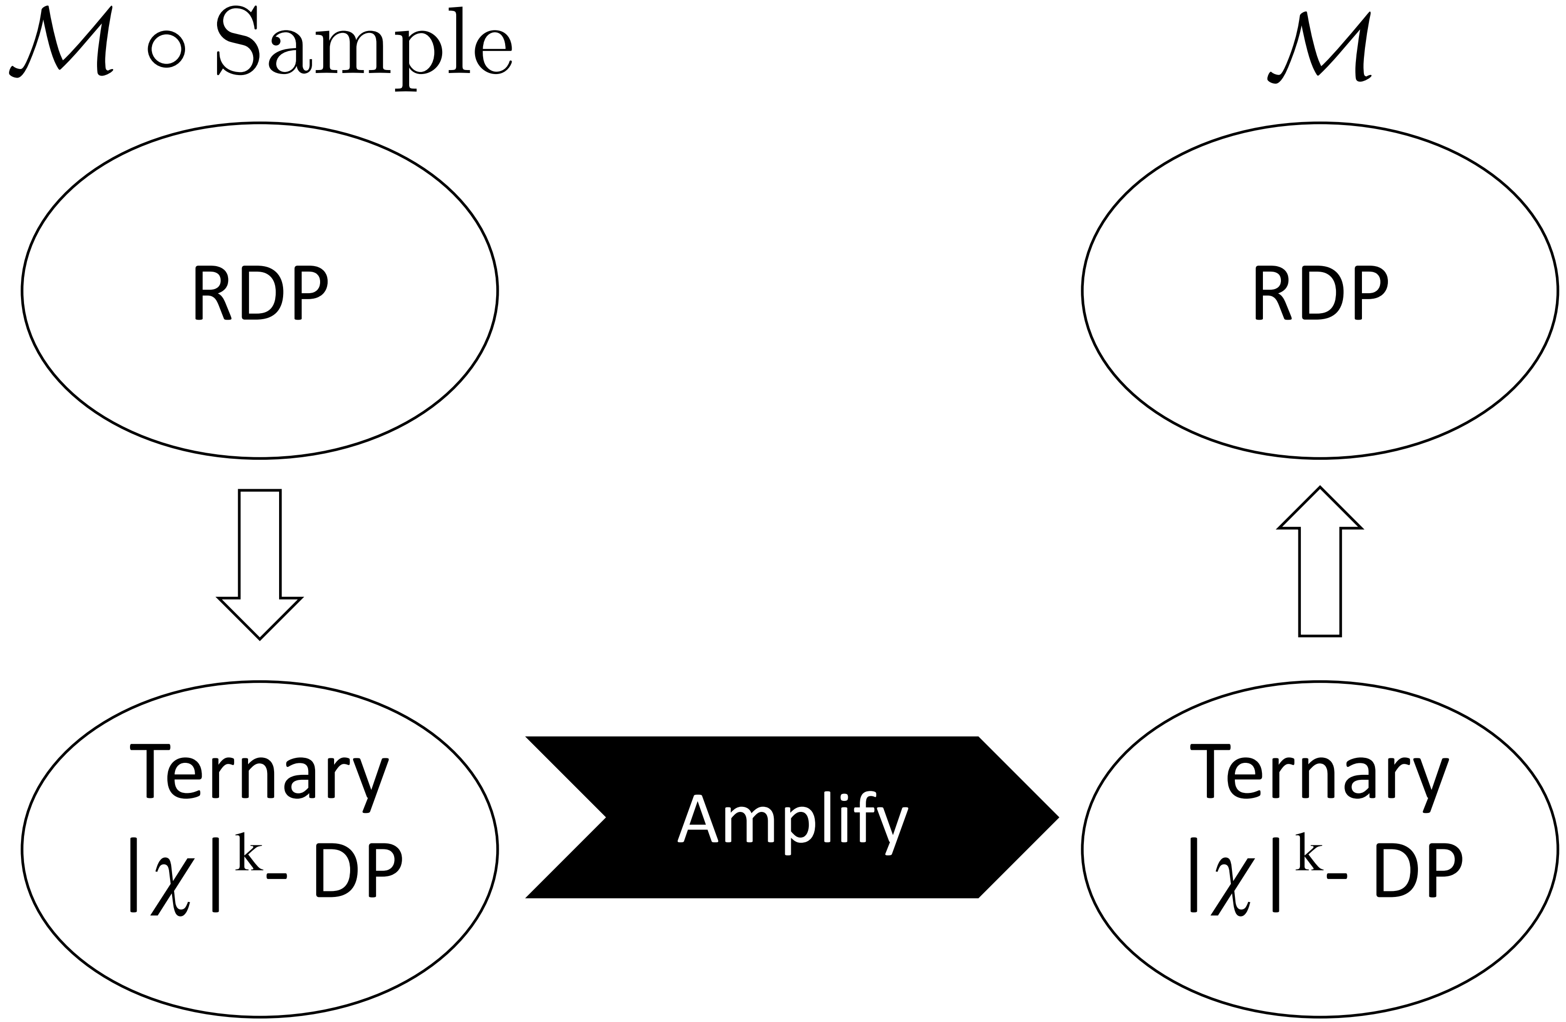

A short detour to divergences

- Renyi divergence $D_\alpha(p\|q) := \frac{1}{\alpha - 1}\log \mathbb{E}_q[e^{\alpha \log(p/q)}]$
  - $D_{1/2}(p\|q) = -2\log(1 - \frac{\mathrm{Hel}(p\|q)}{2})$
  - $\lim_{\alpha\to 1} D_\alpha(p\|q) = \mathrm{KL}(p\|q)$
  - $D_2(p\|q) = \log(1 + \chi^2(p\|q))$
- Pearson-Vajda Divergences

$$\chi^\ell(p\|q) := \mathbb{E}_q[(p/q - 1)^\ell]$$

$$|\chi|^\ell(p\|q) := \mathbb{E}_q[|p/q - 1|^\ell]$$

f-divergence

$$D_f(p\|q) := \mathbb{E}_q[f(p/q)]$$

Pearson-Vajda divergences are moments of the linearized privacy loss

$$\mathbb{E}[\log(p/q)^\alpha] = \frac{\partial^\alpha}{\partial t^\alpha}[e^{K_{\mathcal{M}}(t)}](0),$$

$$\mathbb{E}[(p/q - 1)^\alpha] = \Delta^{(\alpha)}[e^{K_{\mathcal{M}}(\cdot)}](0).$$

Discrete Derivative

Ternary $|\chi|^\alpha$-divergences and $|\chi|^\alpha$-DP

$$D_{|\chi|^\alpha}(p, q \| r) := \mathbb{E}_r \left[\left| \frac{p - q}{r} \right|^\alpha \right].$$

- Take supremum over three data sets that are **mutually adjacent**

p

q

r

$$\sup_{X, X', X'' \text{ mutually adjacent}} \left(D_{|\chi|^\alpha}(\mathcal{M}(X), \mathcal{M}(X') \| \mathcal{M}(X'')) \right)^{1/\alpha} \leq \zeta(\alpha).$$

Ternary $|\chi|^{\alpha}$ –DP and Binary $|\chi|^{\alpha}$ –DP are roughly the same

$$\sup_{X,X',X''\,\text{mutually adjacent}} \left(D_{|\chi|^{\alpha}}(\mathcal{M}(X),\mathcal{M}(X')\|\mathcal{M}(X''))\right)^{1/\alpha} \leq \zeta(\alpha).$$

$$\sup_{X,X':d(X,X')\leq 1} \left(D_{|\chi|^{\alpha}}(\mathcal{M}(X)\|\mathcal{M}(X'))\right)^{1/\alpha} \leq \xi(\alpha).$$

Lemma: Ternary $|\chi|^{\alpha}$ -DP ≈ Binary $|\chi|^{\alpha}$ -DP.

$$\xi(\alpha)^{\alpha} \leq \zeta(\alpha)^{\alpha} \leq 4\xi(\alpha)^{\alpha}$$

Step 1. Ternary $|\chi|^k$-DP is natural for subsampling

Proposition (Privacy amplification for Ternary $|\chi|^k$-DP)

Let a mechanism $\mathcal{M}$ obey ζ-ternary-$|\chi|^\alpha$-DP, then the algorithm $\mathcal{M} \circ \text{sample}$ obeys $\gamma\zeta$-ternary-$|\chi|^\alpha$-DP.

$$p = \gamma p(\cdot|E) + (1-\gamma)p(\cdot|E^c)$$
$$q = \gamma q(\cdot|E) + (1-\gamma)q(\cdot|E^c).$$

Still mixture distributions!

$$D_{|\chi|^j}(p,q\|r) = \mathbb{E}_r\left[\left(\frac{|p-q|}{r}\right)^j\right] = \gamma^j \mathbb{E}_r\left[\left(\boxed{\frac{|p(\cdot|E)-q(\cdot|E)|}{r}}\right)^j\right]$$
$$= \gamma^j D_{|\chi|^j}(p(\cdot|E), q(\cdot|E)\|r).$$

Step 2. Bounding RDP with Ternary $|\chi|^k$-DP

$$\mathbb{E}_q\left[\left(\frac{p}{q}\right)^{\alpha}\right] = 1 + \binom{\alpha}{1}\mathbb{E}_q\left[\frac{p}{q} - 1\right] + \sum_{j=2}^{\alpha}\binom{\alpha}{j}\mathbb{E}_q\left[\left(\frac{p}{q} - 1\right)^{j}\right].$$

Bound binary with ternary:

$$\max_{p,q}\ \mathbb{E}_q\left[\left(\frac{p-q}{q}\right)^{j}\right] \le \max_{p,q,r}\ \mathbb{E}_r\left[\left(\frac{p-q}{r}\right)^{j}\right]$$

Apply Natural Subsampling:

$$\mathbb{E}_q\left[\left(\frac{p}{q}\right)^{\alpha}\right] \le 1 + \sum_{j=2}^{\alpha}\binom{\alpha}{j}\gamma^{j}\zeta(j)^{j},$$

Step 3. Bounding Ternary $|\chi|^k$-DP with RDP

- From Ternary to Binary $|\chi|^k$-DP, we lose a factor of 4, then

$$D_2(p\|q) = \log(1 + \chi^2(p\|q))$$

Lemma 16. Let X, Y be nonnegative random variables, for any $j \geq 1$,

$$\mathbb{E}[|X - Y|^j] \leq \mathbb{E}[X^j] + \mathbb{E}[Y^j].$$

Lemma 17. Let X, Y be nonnegative random variables and with probability 1, $e^{-\varepsilon}Y \leq X \leq e^{\varepsilon}Y$. Then for any $j \geq 1$,

$$\mathbb{E}[|X - Y|^j] \leq \mathbb{E}[Y^j](e^{\varepsilon} - 1)^j.$$

Step 3. Bounding Ternary $|\chi|^k$-DP with RDP

Theorem (Upper bound): Let M obeys (α ,Ɛ(α))-RDP for all α. Then M(subsample(DATA)) obeys

$$\epsilon'(\alpha) \leq \frac{1}{\alpha-1}\log\left(1+\gamma^2\binom{\alpha}{2}\min\left\{4(e^{\epsilon(2)}-1), e^{\epsilon(2)}\min\{2,(e^{\epsilon(\infty)}-1)^2\}\right\}\right.$$
$$\left.+\sum_{j=3}^{\alpha}\gamma^j\binom{\alpha}{j}e^{(j-1)\epsilon(j)}\min\{2,(e^{\epsilon(\infty)}-1)^j\}\right).$$

Lower bound by constructing a data sets pair

- Construct a specific pair of data set
  - X = [0,0,0,0,…,0,1]
  - X’ = [0,0,0,0,…,0,0]
- All subsamples from X’ are identical! If the last data point is not chosen, so are the subsample from X

$$\mathbb{E}_q\left[\left(\frac{(1-\gamma)q+\gamma p}{q}\right)^{\alpha}\right] = \mathbb{E}_q\left[\left(1-\gamma+\gamma\frac{p}{q}\right)^{\alpha}\right] = (1-\gamma)^{\alpha}\mathbb{E}_q\left[\left(1+\frac{\gamma}{1-\gamma}\frac{p}{q}\right)^{\alpha}\right]$$
$$= (1-\gamma)^{\alpha}\left(1+\alpha\frac{\gamma}{1-\gamma}+\sum_{j=2}^{\alpha}\binom{\alpha}{j}\left(\frac{\gamma}{1-\gamma}\right)^{j}\mathbb{E}_q\left[\left(\frac{p}{q}\right)^{j}\right]\right).$$

Constants matter in Differential Privacy. Can we close the constant gap?

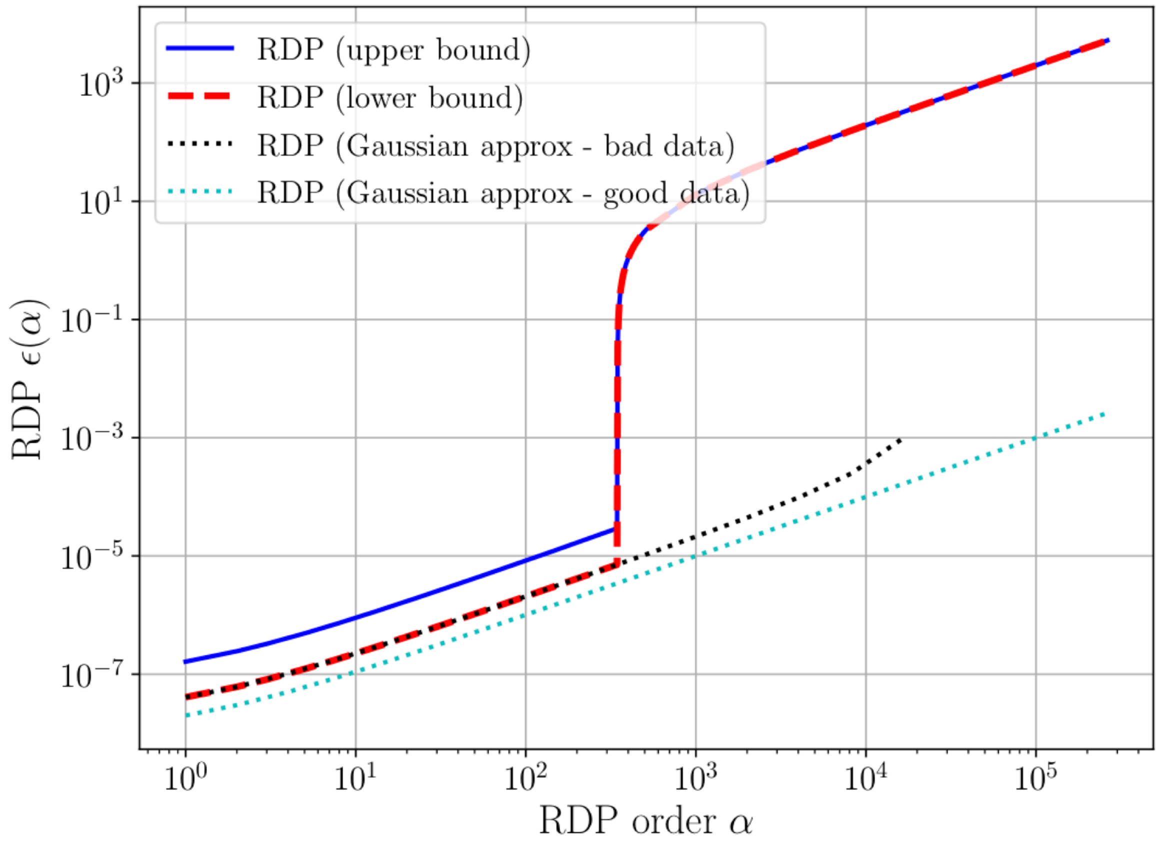

(a) RDP of Subsampled Gaussian with σ = 5

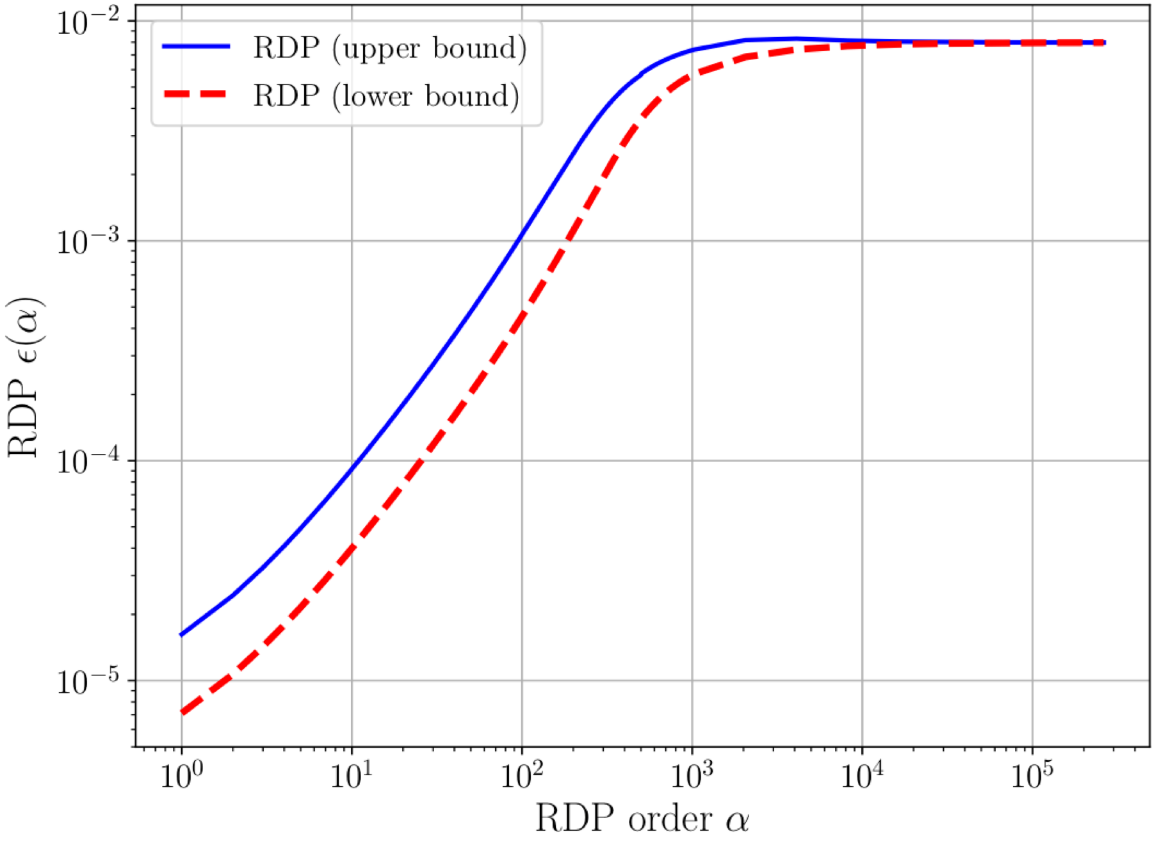

(b) RDP of Subsampled Laplace with b = 0.5

Sometimes we can improve it somewhat.

- If there is a pair of worst case data sets that attains the RDP bound for all α.
- If the same pair of data sets also attains the Binary $|\chi|^k$-DP bounds.
- Then we have an improved bound.
- This is true for Gaussian mechanism.

(New Results) RDP Amplification Under Poisson sampling

Work with my student
Yuqing Zhu

Theorem (Poisson Sampling):

$$\epsilon_{\mathcal{M}\circ\mathsf{PoissonSample}}(\alpha) \leq \frac{1}{\alpha-1}\log\Big\{(1-\gamma)^{\alpha-1}(\alpha\gamma-\gamma+1) + \binom{\alpha}{2}\gamma^2(1-\gamma)^{\alpha-2}e^{\epsilon(2)} + \boxed{3}\sum_{\ell=3}^{\alpha}\binom{\alpha}{\ell}(1-\gamma)^{\alpha-\ell}\gamma^{\ell}e^{(\ell-1)\epsilon(\ell)}\Big\}.$$

Remark:

- Multiplicative error O(1+γ) for small α, additive error log(3)/(α-1) for large α.
- The factor of 3 in the lower order term can be removed if **odd-order Pearson-Vajda divergences** > 0
- Allows us to prove **exact bound for Gaussian** mechanism and **Laplace** mechanism.

Is the lower bound always achievable by all M? Counterexample from: (Nielsen and Nock, 2014)

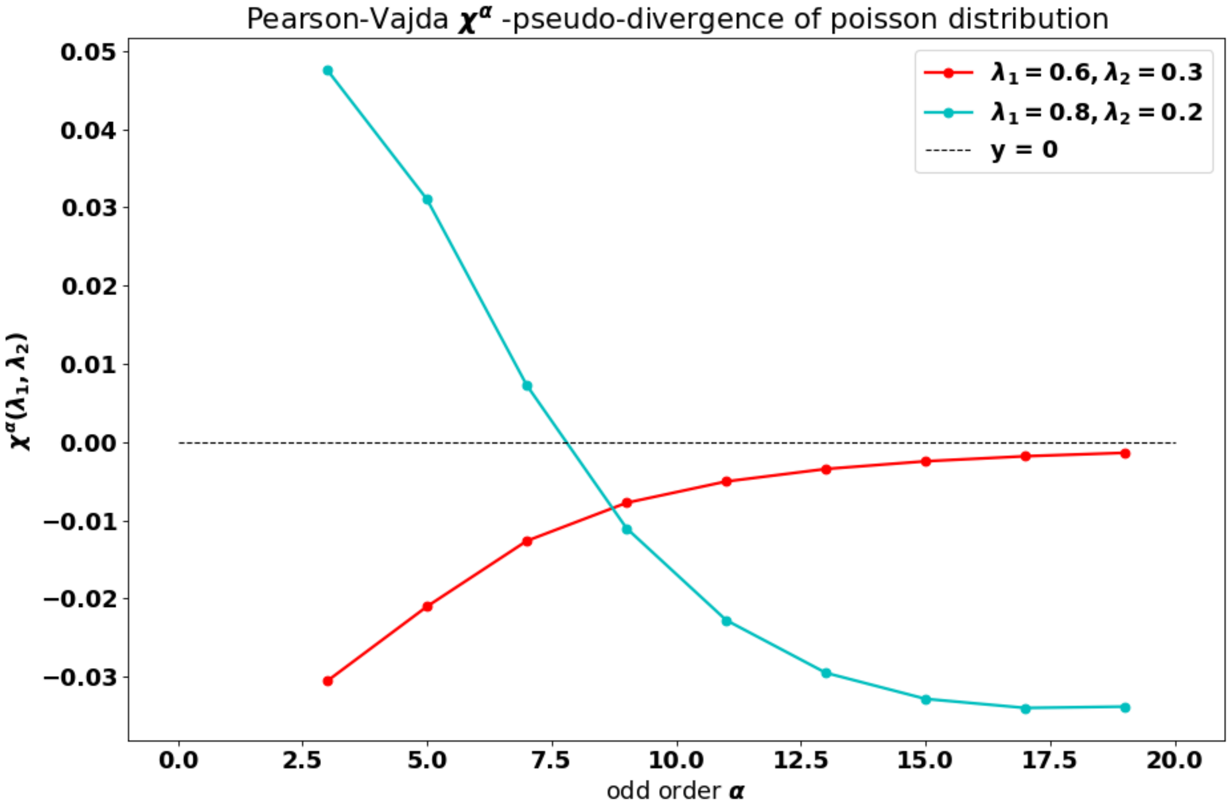

Figure 2. Negative χ^{α} divergence in Poisson distribution.

Main challenge in the Poisson Case

- Asymmetry: X has n data points,   X' has n+1 data points.
- Need to bound not just E[(p/q)^k] but also E[(q/p)^k].
- E[(q/p)^k] is easy, E[(q/p)^k] is challenging
  - Requires an explicit knowledge on the worst pair of data sets.

Take-home messages and open problems

1. The first generic subsampling lemma for RDP mechanism.
2. Exact formula under Poisson sampling for some mechanisms.
3. Stronger composition than advanced composition.

- Open problems / interesting directions:
  - Closing the constant gap in the upper/lower bounds
  - Exploiting randomness from the data

W., Balle & Kasiviswanathan (2018). Subsampled Renyi Differential Privacy and Analytical Moments Accountant. *AISTATS'2019*
Zhu & W. (2019) Poisson Subsampled Renyi Differential Privacy. Upcoming.

Open problem: Exploit the noise from the data in a valid way?

- Subsample with too small a noise added does not amplify privacy.
- Subsample with slightly larger noise smooth things out.
- Your peers may be hiding you underneath a **privacy blanket**!

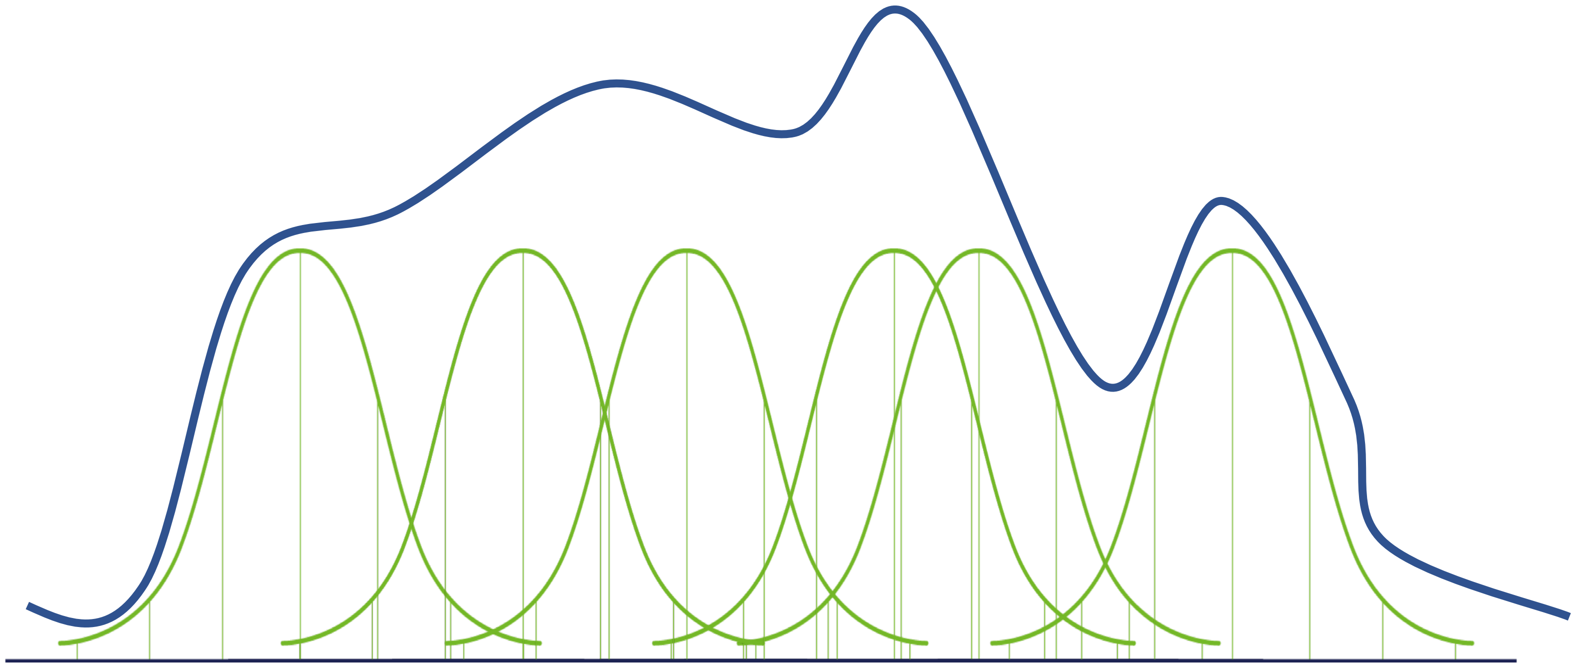

References

- Privacy amplification by subsampling in DP
  - Kasiviswanathan, S. P., Lee, H. K., Nissim, K., Raskhodnikova, S., & Smith, A. (2011). What can we learn privately?. SIAM Journal on Computing, 40(3), 793-826.
- Tight bounds for (Ɛ, δ)-DP
  - Li, N., Qardaji, W., & Su, D. (2012). On sampling, anonymization, and differential privacy or, k-anonymization meets differential privacy. In CCS'12. ACM.
  - Balle, B., Barthe, G., & Gaboardi, M. (2018). Privacy amplification by subsampling: Tight analyses via couplings and divergences. In NeurIPS'18.
- RDP / Moments Accountant:
  - Mironov, I. (2017, August). Rényi differential privacy. In *IEEE CSF'17*. IEEE.
  - Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. "Deep learning with differential privacy." In CCS'16. ACM, 2016.
- Privacy amplification under RDP
  - W., Balle, Kasiviswanathan (2019) Subsampled RDP and Analytical Moments Accountant. In AISTATS'19.
  - Zhu and W. (2019) Poisson Subsampled RDP. Available soon.

By the joint convexity argument, we get:

$$\mathbb{E}_p[(q/p)^\alpha] \leq \sum_J \mathbb{P}(J)\mathbb{E}_{\mu_0(J)} \left(\frac{(1-\gamma)\mu_0(J) + \gamma\mu_1(J)}{\mu_0(J)} \right)^\alpha$$

$$\mathbb{E}_q[(p/q)^\alpha] \leq \sum_J \mathbb{P}(J)\mathbb{E}_{(1-\gamma)\mu_0(J)+\gamma\mu_1(J)} \left(\frac{\mu_0(J)}{(1-\gamma)\mu_0(J) + \gamma\mu_1(J)} \right)^\alpha .$$

- But the latter is really hard to work with given only RDP upper bounds.
- Finding the pair of data sets that maximizes the latter is where things get a bit challenging.
- Our proof involves proposing an alternative decomposition to replace the second inequality.