Subsampled Renyi Differential Privacy and Analytical Moments Accountant

Yu-Xiang Wang

UC Santa Barbara

Joint work with Borja Balle and Shiva Kasiviswanathan

Outline

- Preliminary:
  - Algorithm-specific privacy analysis and Renyi DP
  - Privacy amplification by subsampling

- Renyi DP of Subsampled Algorithms

- Composition and Analytical moments accountant

Renyi DP and algorithm-specific DP analysis

- Ɛ-DP is a crude summary of the privacy guarantee

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

- RDP (Mironov, 2017) and 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)

Subsampled Randomized Algorithm

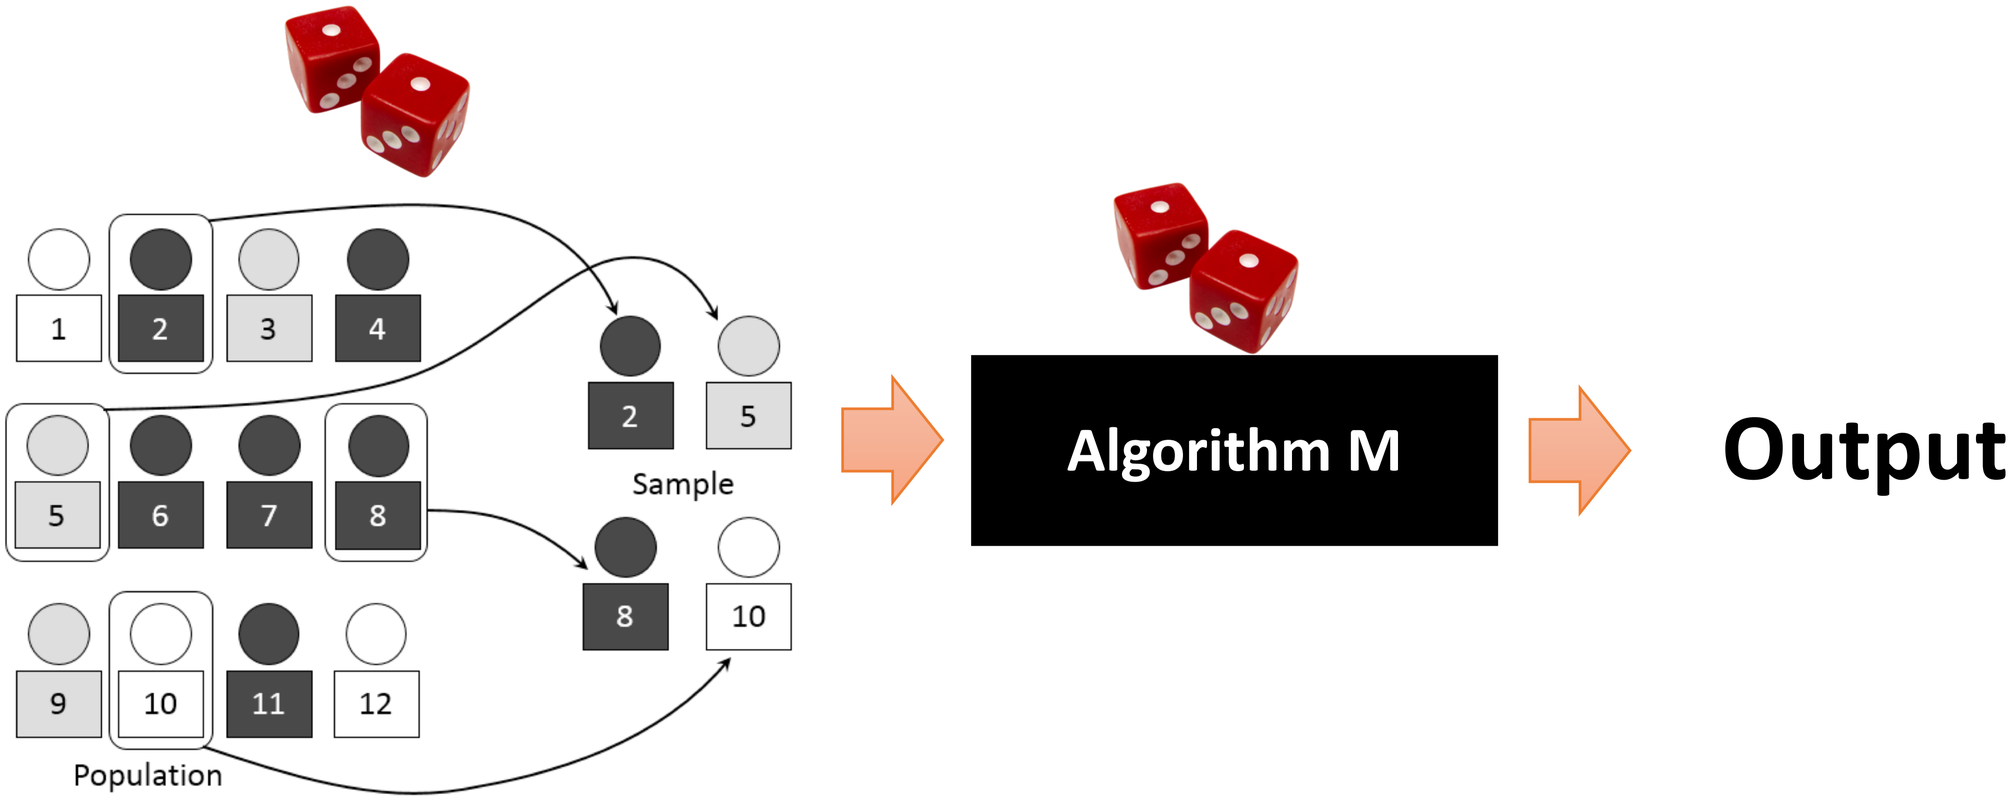

Example: The Noisy SGD algorithm (Song et al. 2013; Bassily et. al. 2014)

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

- 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 (Breiman)
- Bootstraps, Jackknife, subsampling bootstrap (Efron; Stein; Politis and Romano)
- Sublinear algorithms in exploratory data analysis
  - Sketching
  - Property testing

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 and , 2015) (Wang et al., 2016)
- Most recent "tightened" revision above in:
  - Borja Balle, Gilles Barthe, Marco Gaboardi (NeurIPS'18)

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.
  - Tight constant.

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

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

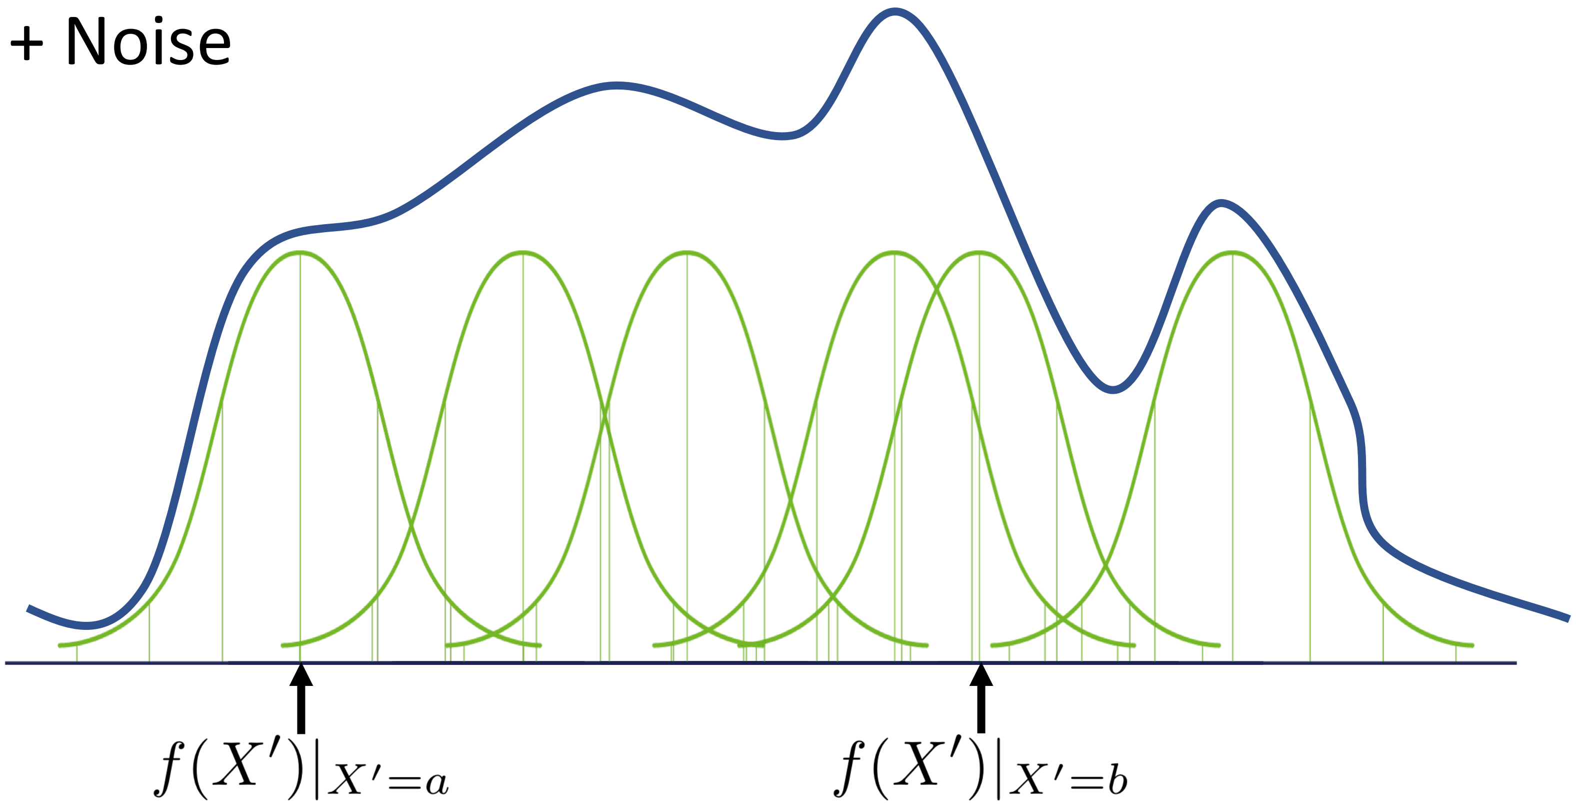

Changing to an adjacent data set

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

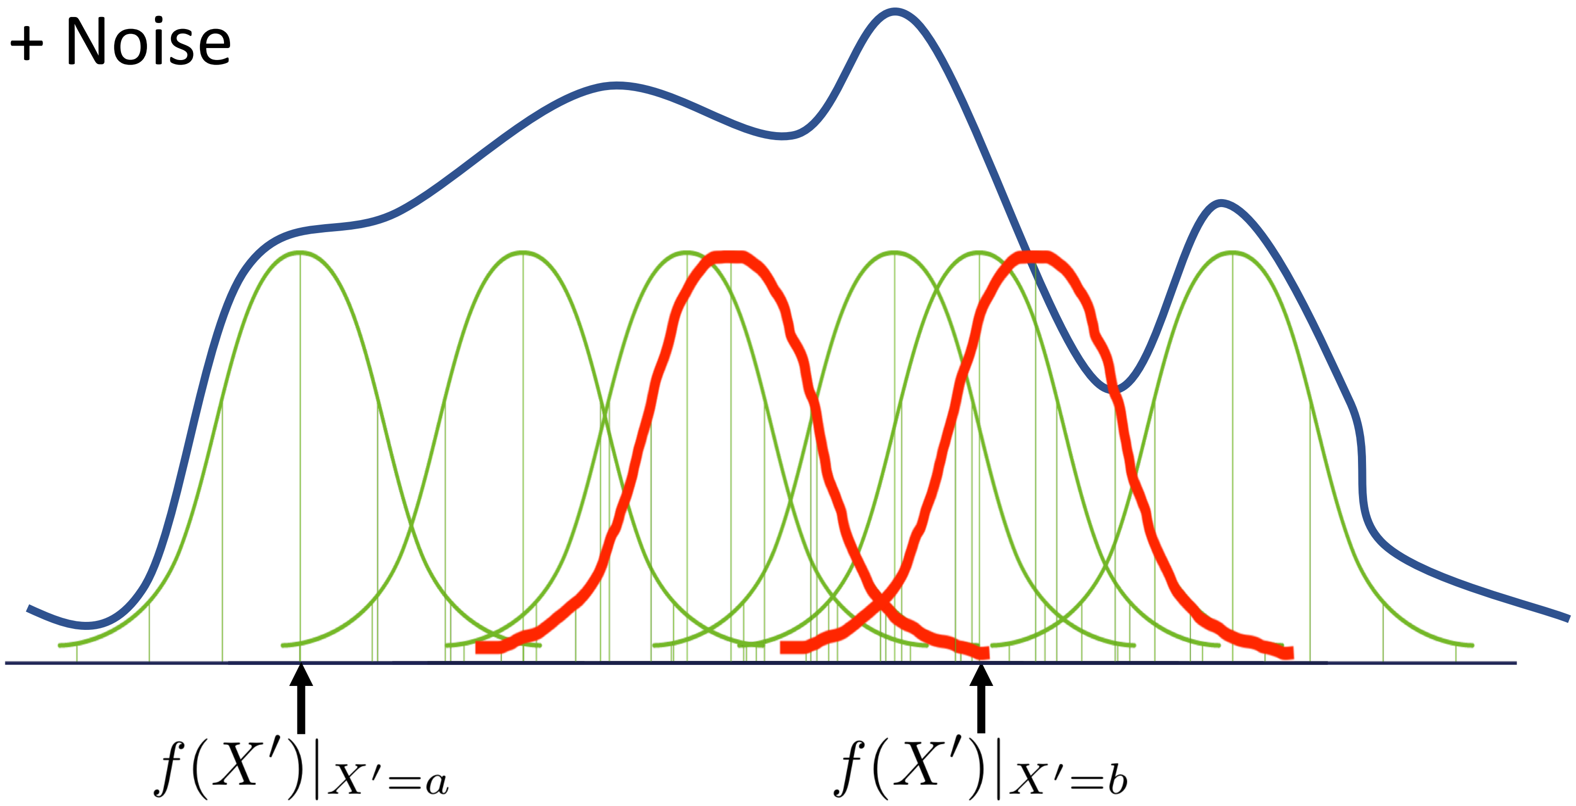

Main technical results

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

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

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

$$\epsilon'(\alpha) \geq \frac{\alpha}{\alpha-1}\log(1-\gamma)+\frac{1}{\alpha-1}\log\left(1+\alpha\frac{\gamma}{1-\gamma}+\sum_{j=2}^{\alpha}\binom{\alpha}{j}(\frac{\gamma}{1-\gamma})^j e^{(j-1)\epsilon(j)}\right).$$

Numerical evaluation of the bounds

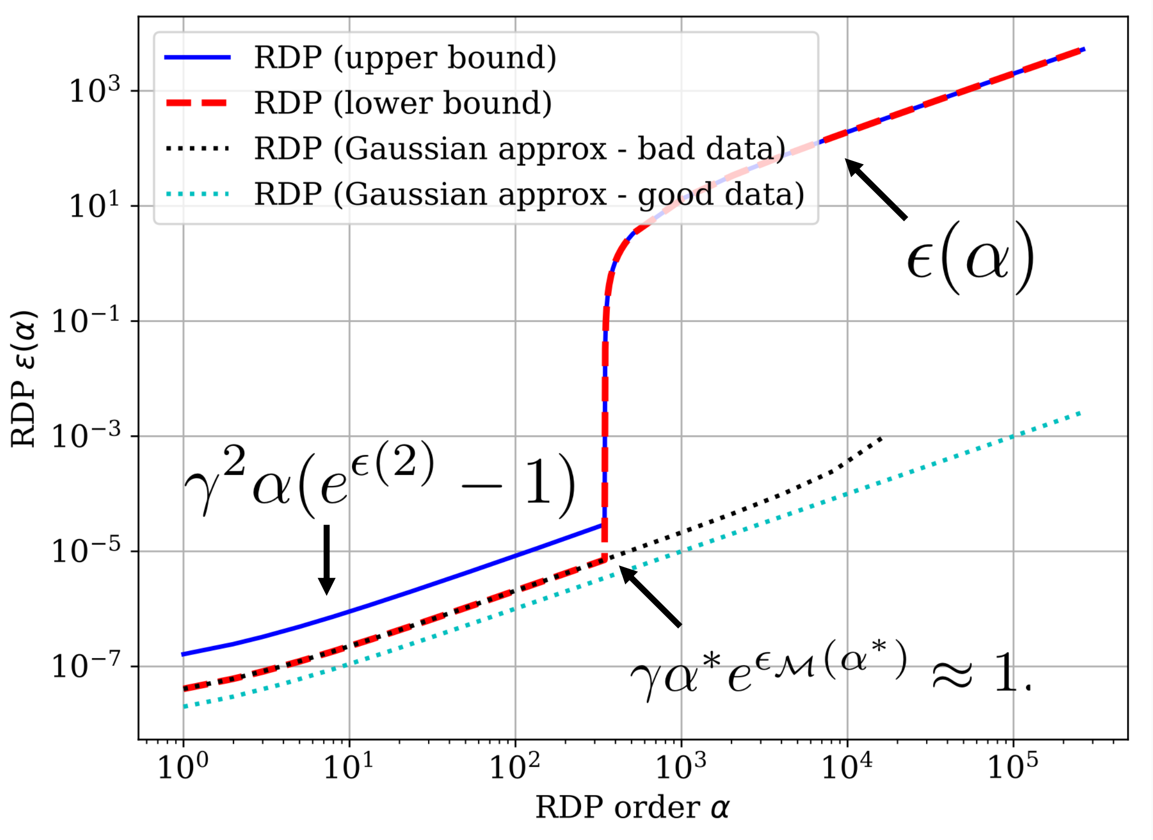

(a) Subsampled Gaussian with $\sigma = 5$.

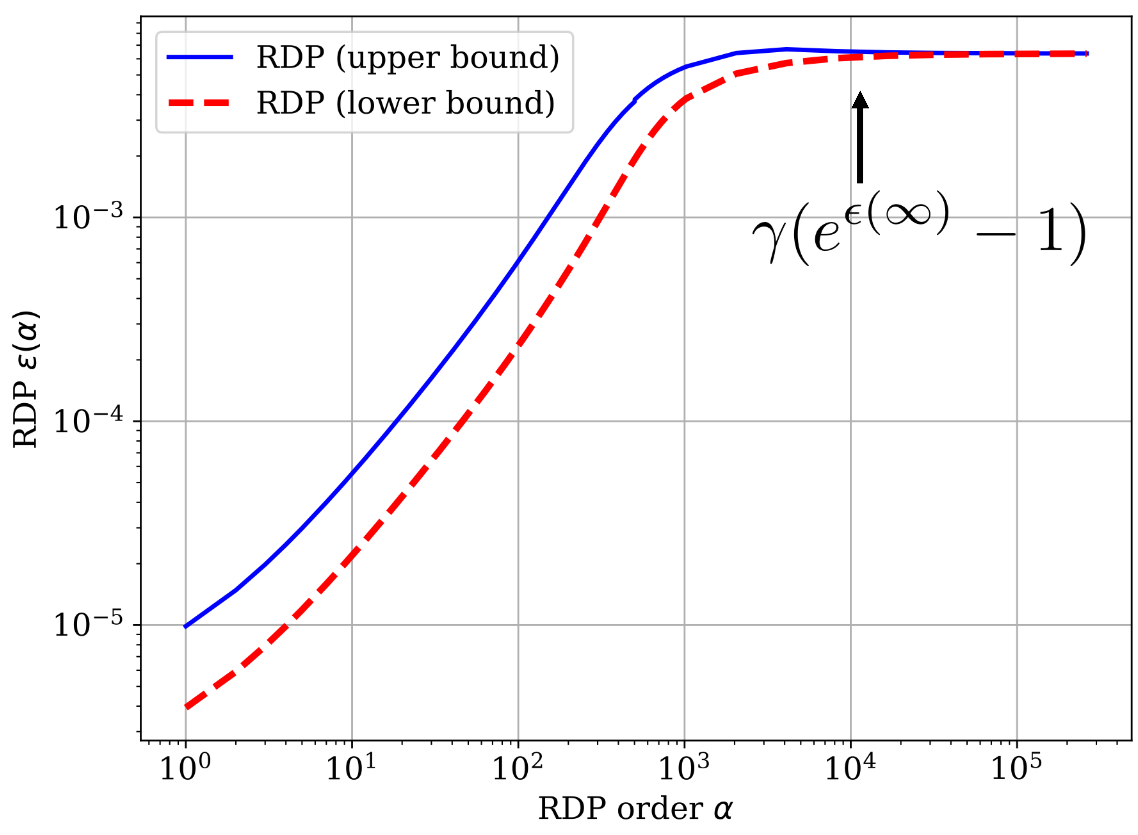

(e) Subsampled Laplace with $b = 0.5$.

New techniques in the proof

- Moments of Linearized Privacy loss R.V.
  - discrete difference operators ---- continuous derivative operators
  - Newton series expansions ----- Taylor series
- Ternary Pearson-Vajda divergences.
  - Natural for handling subsampling.

Analytical moments accountant

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

Using our bounds for advanced composition

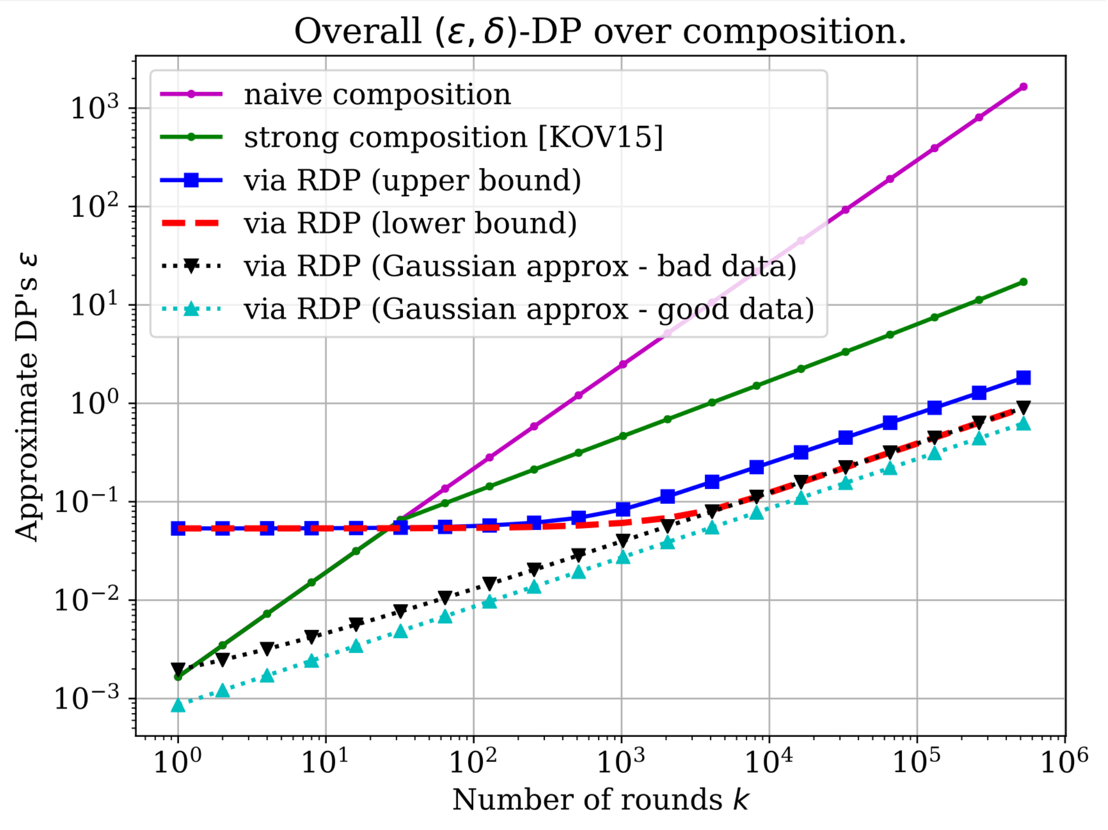

(a) Subsampled Gaussian with $\sigma = 5$.

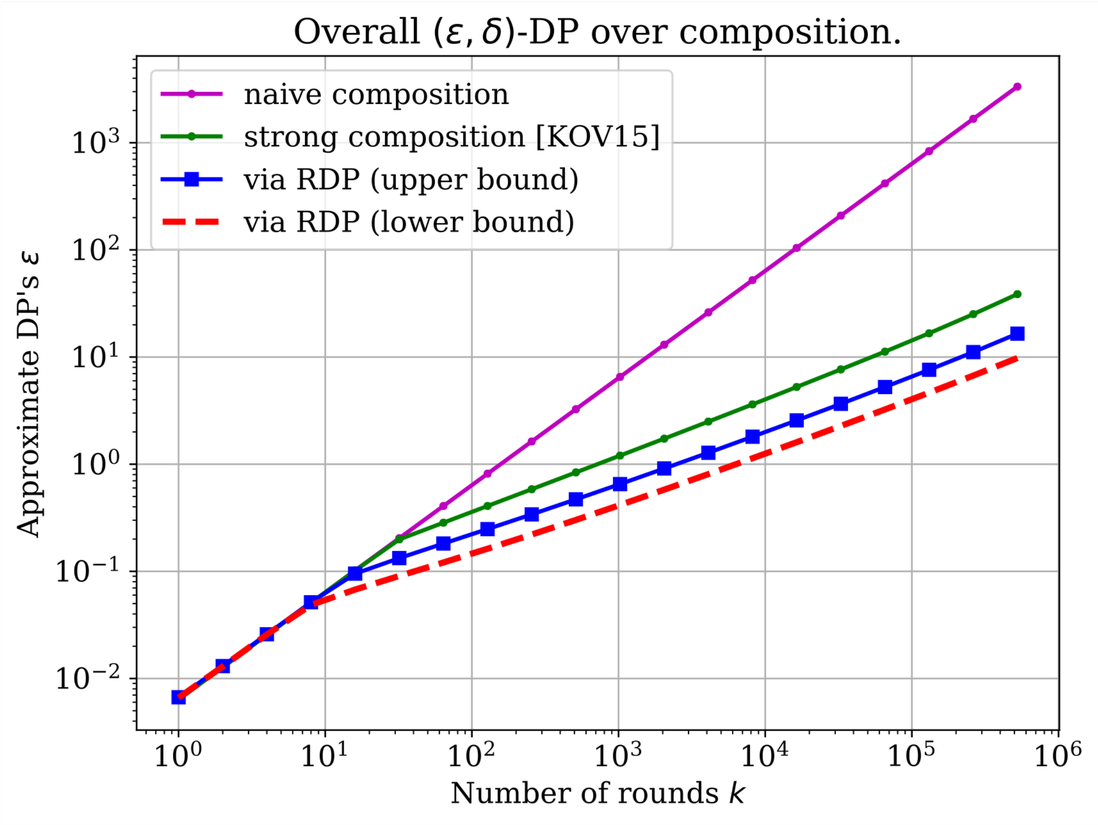

(e) Subsampled Laplace with $b = 0.5$.

Take-home messages and future work

1. The first generic subsampling lemma for RDP mechanism.
2. Stronger composition than advanced composition

- Future work:
  - Closing the constant gap in the upper/lower bounds
  - Other types of subsampling (e.g., Poisson subsampling)
  - Other types of privacy amplification in RDP

Wang, Y. X., Balle, B., & Kasiviswanathan, S. (2018). Subsampled R\'enyi Differential Privacy and Analytical Moments Accountant. *arXiv preprint arXiv:1808.00087*.

Open source software will be released soon! Stay tuned.

Thank you for your attention!