UC SANTA BARBARA

Poisson Subsampled Rényi Differential Privacy

Yuqing Zhu

joint work with Yu-Xiang Wang

Privacy Amplification by Sampling

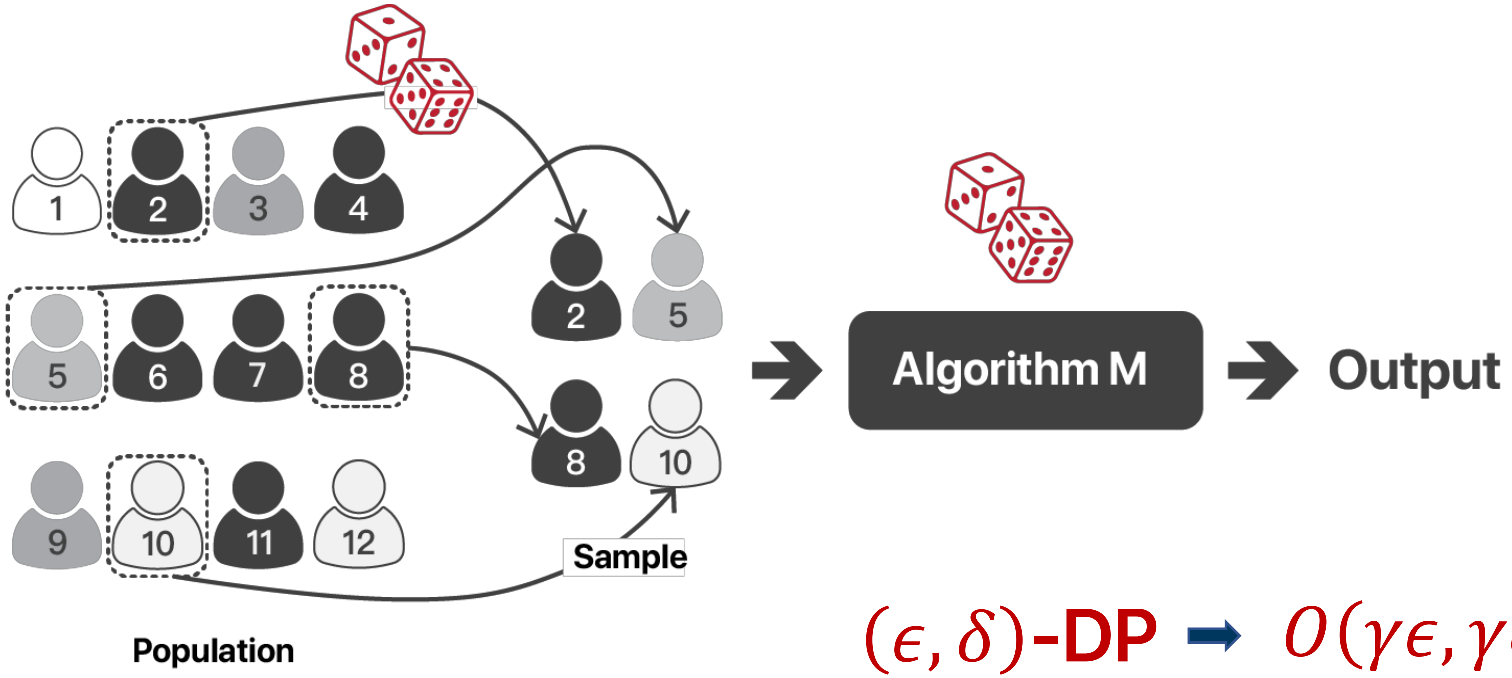

Sampling probability γ

(ϵ, δ)-DP ➡ $O(\gamma\epsilon, \gamma\delta)$-DP

[KLNRS'08], [Li et al., 2011]

(ϵ, α)-Rényi DP ➡ What's the optimal bound ?

Strong composition tool

Example: The Noisy SGD Algorithm

Song et al. 2013; Bassily et al. 2014

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

1. Randomly chosen minibatch (Poisson subsampling)
2. 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

Exact RDP of Subsampled Mechanism

Let M be any randomized algorithm that obeys $(\alpha,\epsilon(\alpha))$–RDP
γ be the subsampling probability and for integer $\alpha \geq 2$

Asymptotic rate

$$\epsilon_{M\circ \boldsymbol{sample}}(\alpha) \leq \ O(\alpha\gamma^2\epsilon(2))$$

This asymptotic rate holds for any mechanism M !

Exact RDP of Subsampled Mechanism

Let M be any randomized algorithm that obeys $(\alpha,\epsilon(\alpha))$–RDP
γ be the subsampling probability and for integer $\alpha \geq 2$

$$\epsilon_{\mathrm{M}\circ Sample}(\alpha) \leq \frac{1}{\alpha}\log\{ (1-\gamma)^{\alpha-1}(\alpha\gamma - \gamma + 1) + \binom{\alpha}{2}\gamma^2(1-\gamma)^{\alpha-2}e^{\epsilon(2)} + 3\sum_{\ell=3}^{\alpha}\binom{\alpha}{\ell}(1-\gamma)^{\alpha-\ell}\gamma^{\ell}e^{(\ell-1)\epsilon(\ell)}\}$$

This bound is optimal, up to a factor of 3 on a low order term

Exact Amplification Bound for RDP

Let M be any randomized algorithm that obeys $(\alpha,\epsilon(\alpha))$–RDP
γ be the subsampling probability and for integer $\alpha \geq 2$

$$\epsilon_{\mathrm{M}\circ Sample}(\alpha) \leq \frac{1}{\alpha}\log\{ (1-\gamma)^{\alpha-1}(\alpha\gamma - \gamma + 1) + \binom{\alpha}{2}\gamma^2(1-\gamma)^{\alpha-2}e^{\epsilon(2)}$$
$$+\boxed{3}\sum_{\ell=3}^{\alpha}\binom{\alpha}{\ell}(1-\gamma)^{\alpha-\ell}\gamma^{\ell}e^{(\ell-1)\epsilon(\ell)}\}$$

Get rid of it

Matches the lower bound when M is Gaussian or Laplace mechanism

Overall (ϵ, δ)-DP over composition

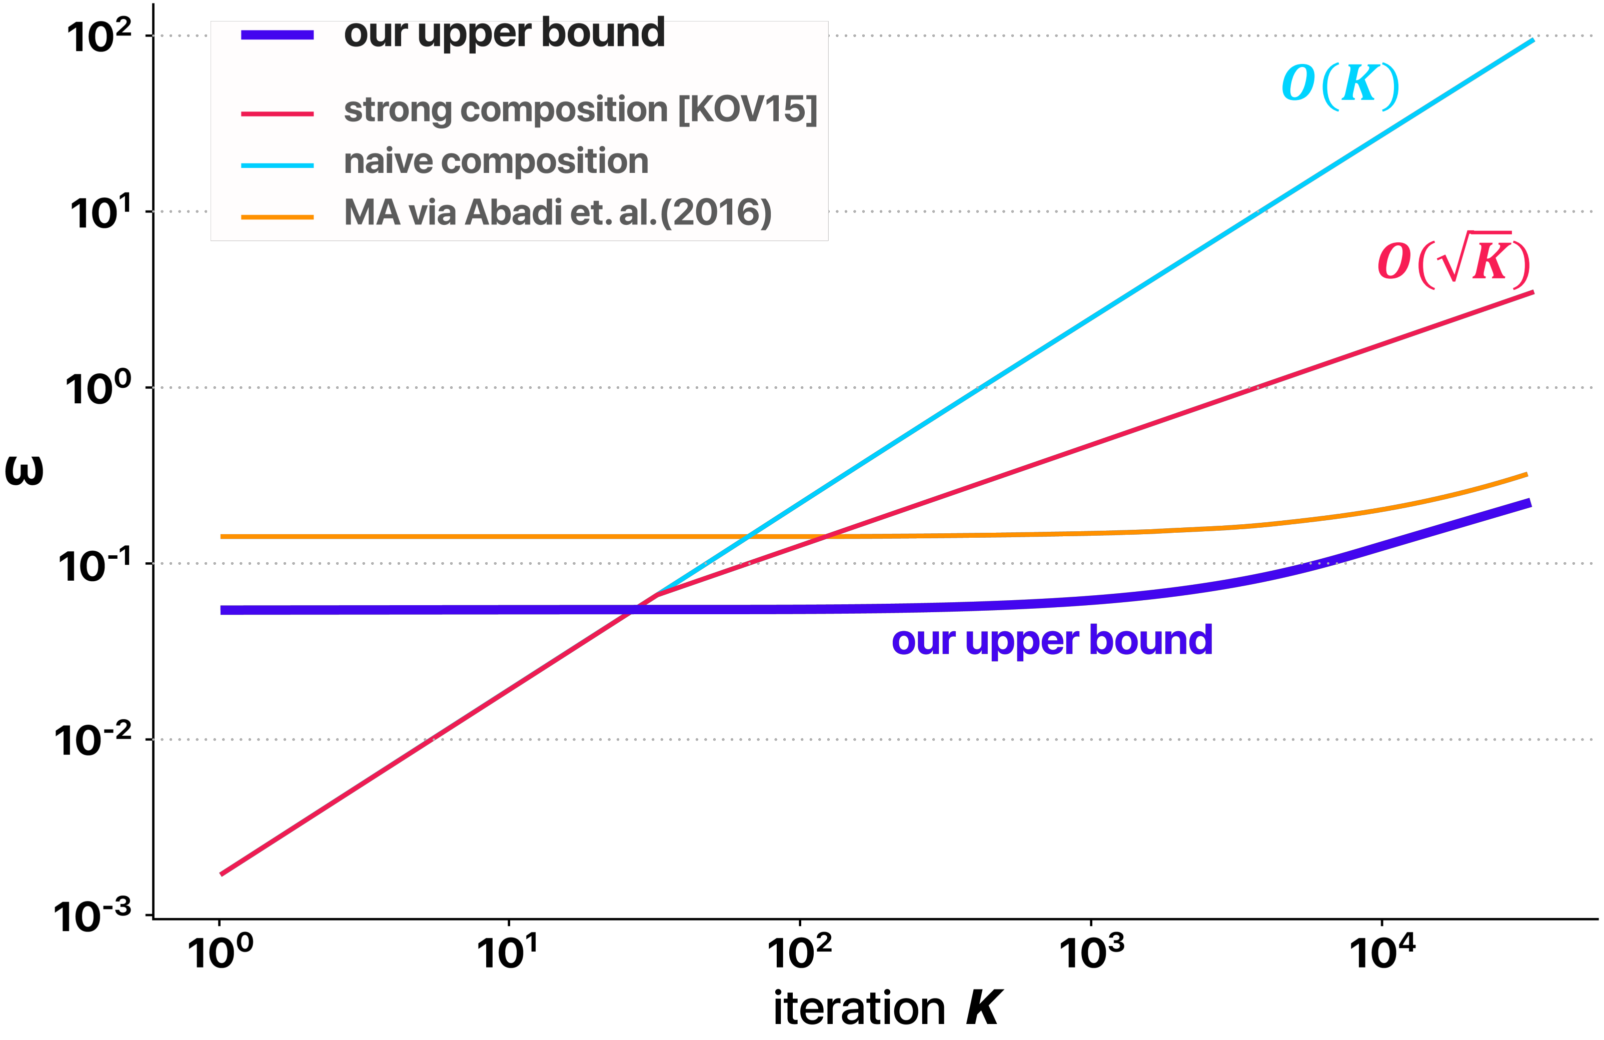

Subsampled Gaussian Mechanism

$\sigma = 5, \gamma = 1e - 3$

Low Privacy Regime

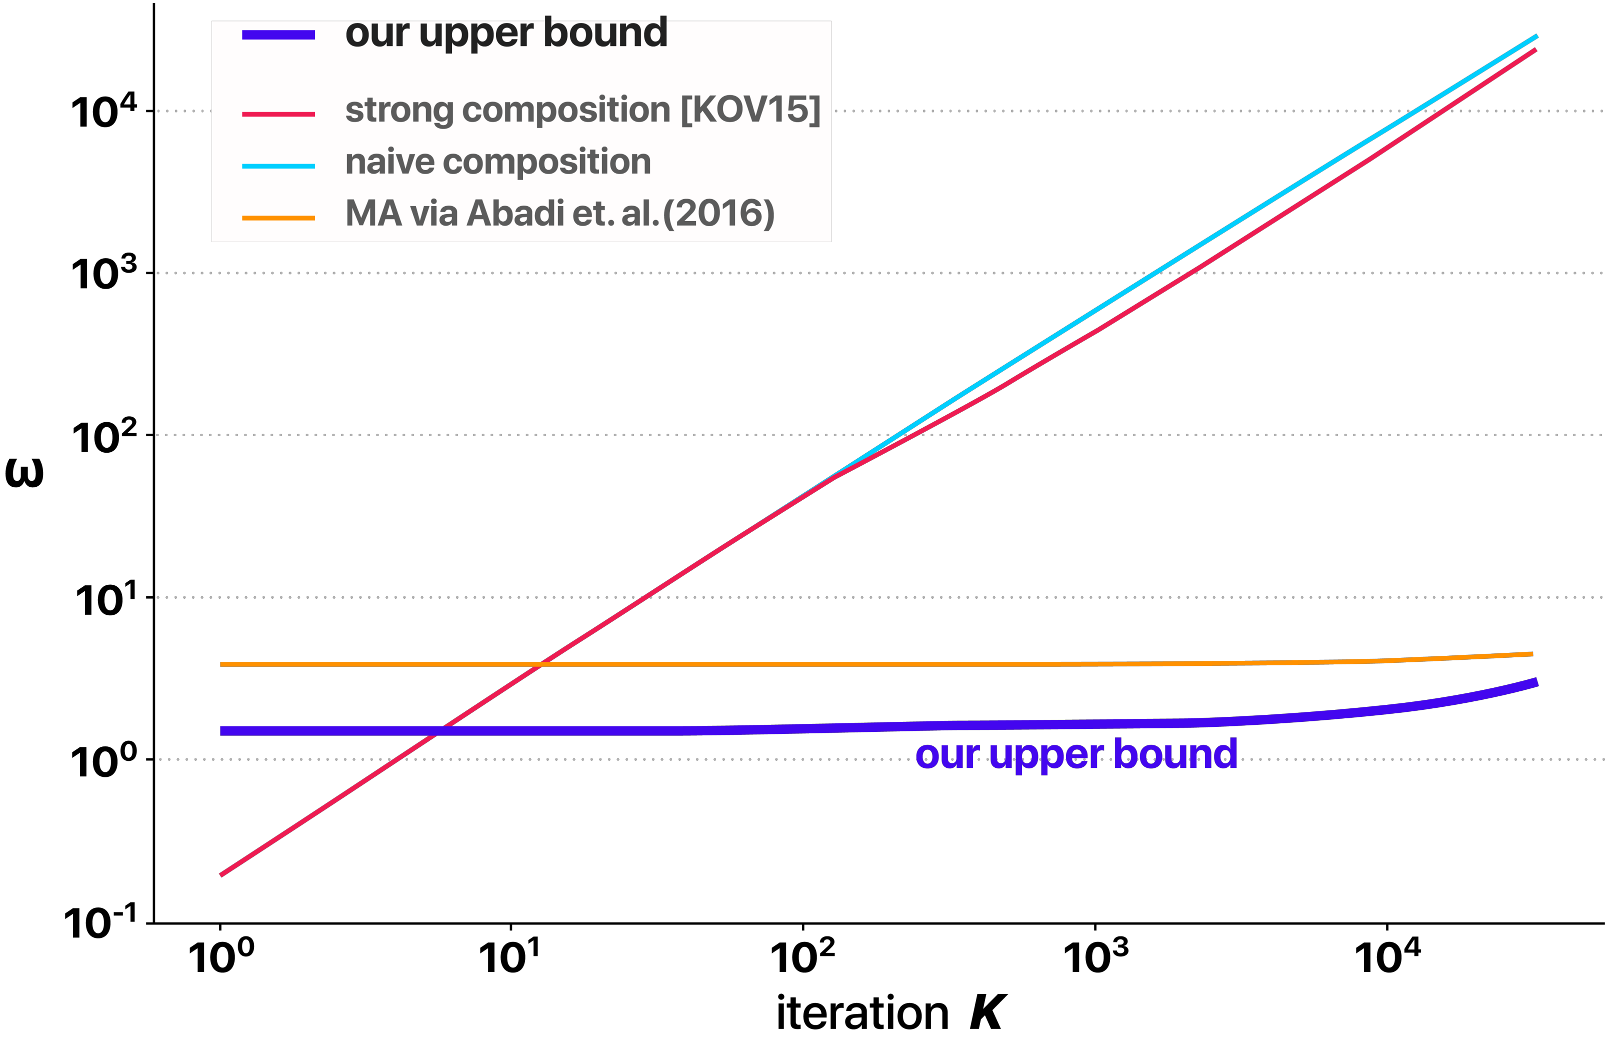

Subsampled Gaussian Mechanism

$$\sigma = 1, \gamma = 1e - 3$$

Thank you!

Poster Number Pacific Ballroom **#178**

Code available:
https://github.com/yuxiangw/autodp

Or just use:
pip install autodp

Get Paper