Improved Privacy Filters and Odometers: Time-Uniform Bounds in Privacy Composition

Justin Whitehouse*, Aaditya Ramdas*, Steven Wu*, and Ryan Rogers^
*Carnegie Mellon University, ^LinkedIn

Privacy Composition: Limitations and a Solution

- Advanced composition results allows analysts to adaptively select private algorithms to run while maintaining privacy [1, 2, 5].
- However, analysts cannot adaptively pick privacy parameters.
- *Privacy Filters/Odometers* allow analysts to adaptively choose privacy parameters to meet a desired privacy level.
- Existing filters only apply for special cases (e.g., probabilistic DP [3], Rényi DP [4]), and are looser than advanced composition.
- We eliminate the limitations by designing essentially tight filters and flexible odometers.

Conditional DP and Probabilistic DP

- The *privacy loss* of an algorithm A acting on input x vs. input x' is the random variable $L(x,x') := \log\left(p_x(A(x))/p_{x'}(A(x))\right)$ where p_x and $p_{x'}$ are the densities for $A(x)$ and $A(x')$ respectively.
- An algorithm is *probabilistically differentially private* (pDP) if, for any neighboring datasets, the privacy loss is small with high probability.

$$\sup_{x\sim x'} \mathbb{P}\left(L(x,x') > \epsilon\right) \leq \delta$$

- The nth algorithm in a sequence is *conditionally pDP* if, conditioned on the outputs of the previous $n-1$ algorithms, the nth algorithm is pDP, i.e.

$$\sup_{x\sim x'} \mathbb{P}\left(L_n(x,x') > \epsilon_n \mid A_{1:n-1}(x)\right) \leq \delta$$

- Likewise, the nth algorithm in a sequence is *conditionally DP* if,

$$\mathbb{P}\left(A_n(x) \in B \mid A_{1:n-1}(x)\right) \leq e^{\epsilon}\mathbb{P}\left(A_n(x') \in B \mid A_{1:n-1}(x)\right) + \delta, \quad \forall x \sim x', \forall G$$

Privacy Filters and Odometers

- **Filters:** An (ϵ,δ)-privacy filter is a data-dependent stopping rule τ such that the mechanism which releases the outputs of the first τ algorithms is (ϵ,δ)-DP, i.e.

$$\mathbb{P}\left(A_{1:\tau}(x) \in B\right) \leq e^{\epsilon}\mathbb{P}\left(A_{1:\tau}(x') \in B\right) + \delta, \quad \forall x \sim x', \forall G$$

- **Odometers:** A $\boldsymbol{\delta}$-privacy odometer is a sequence of upper bounds (U_n) satisfying:

$$\sup_{x\sim x'} \mathbb{P}\left(\exists n \in \mathbb{N} : L_n(x,x') > U_n\right) \leq \delta$$

Fully Adaptive Composition via Privacy Filters

Filter Theorem (Informal): Suppose we have a sequence of algorithms which are *conditionally* (ϵ_n,δ_n)-pDP, where (ϵ_n) and $(\delta_n)$are adaptively chosen. Fix $\epsilon > 0, \delta = \delta' + \delta'' > 0$. Define the times T_1 and T_2 by

$$T_1 := \inf\left\{ n \in \mathbb{N} : \epsilon \leq \sqrt{2\log\left(\frac{1}{\delta'}\right)\sum_{m\leq n+1}\epsilon_m^2} + \frac{1}{2}\sum_{m\leq n+1}\epsilon_m^2 \right\},$$

$$T_2 := \inf\left\{ n \in \mathbb{N} : \delta'' \leq \sum_{m\leq n+1}\delta_m \right\}.$$

Then, the time $\tau := T_1 \wedge T_2$ is an (ϵ,δ)-privacy filter.

Fully Adaptive Theorem (informal): We can get the same guarantee as the **Filter Theorem,** replacing the assumption of *conditional* (ϵ_n,δ_n)-pDP with the weaker assumption of conditional (ϵ_n,δ_n)-DP.

References

[1] Cynthia Dwork, Guy N Rothblum, and Salil Vadhan. Boosting and differential privacy. In IEEE 51st Annual Symposium on Foundations of Computer Science, 2010.
[2] Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. Foundations and Trends in Theoretical Computer Science, 2014.
[3] Ryan M Rogers, Aaron Roth, Jonathan Ullman, and Salil Vadhan. Privacy odometers and filters: Pay-as-you-go composition. NeurIPS, 2016.
[4] Vitaly Feldman and Tijana Zrnic. Individual privacy accounting via a Renyi filter. arXiv preprint, arXiv:2008.11193, 2020.
[5] Peter Kairouz, Sewoong Oh, and Pramod Viswanath. The composition theorem for differential privacy. ICML, 2015.

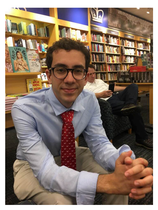


General Odometers via Stitching

Stitched Odometer (Informal): Given a sequence of algorithms which are conditionally (ϵ_n,δ_n)-pDP and $\delta = \delta' + \delta'' > 0$, the sequence (U_n) given by

$$U_n = 1.7\sqrt{\left(\sum_{i\leq n}\epsilon_i^2\right)\left(\log\log\left(\frac{2\sum_{i\leq n}\epsilon_i^2}{\epsilon_1^2}\right) + 0.72\log\left(\frac{5.2}{\delta'}\right)\right)}\mathbf{1}_{\sum_{m\leq n}\delta_m \leq \delta''} + \infty\mathbf{1}_{\sum_{m\leq n}\delta_m > \delta''}$$

is a δ-privacy odometer.

Conjugate Mixture Odometer (Informal): Given the same setup as above, we have that (U_n) given by

$$U_n = \sqrt{2\log\left(\frac{1}{2\delta}\sqrt{\frac{\sum_{n\leq m}\epsilon_m^2 + \rho}{\rho}} + 1\right)\left(\rho + \sum_{m\leq n}\epsilon_m^2\right)}\mathbf{1}_{\sum_{m\leq n}\delta_m \leq \delta''} + \infty\mathbf{1}_{\sum_{m\leq n}\delta_m > \delta''}$$

is a δ-privacy odometer for any $\rho > 0$.

Odometer Comparison

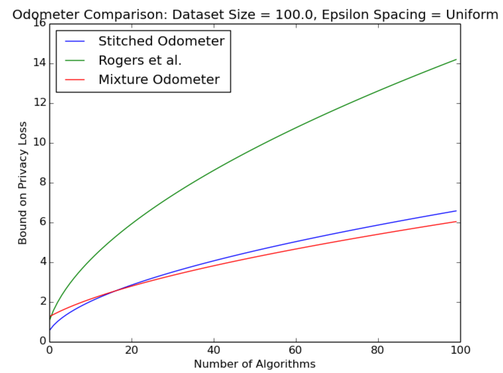

A comparison of the original odometer found in [3] with our odometers. Here, we compose 100 algorithms, all with the same privacy parameters of (0.1, 0)-DP. Both the mixture odometer and the stitched odometer significantly improve over the original odometer, optimizing tightness at different points in time.

jwhiteho@andrew.cmu.edu