Privacy Odometers and Filters: Pay-as-you-Go Composition

Ryan Rogers, Aaron Roth, Jonathan Ullman, and Salil Vadhan

Penn University of Pennsylvania · Northeastern University · Harvard

Differential Privacy [DMNS]

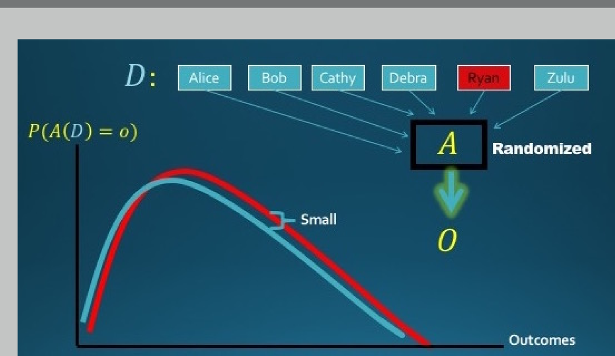

- Outcome of algorithm $A : \mathcal{X}^n \to \mathcal{O}$ should *roughly* stay the same if one person's data changes.

For all neighboring x, x' and outcomes $O \subseteq \mathcal{O}$,

$$\mathbb{P}[A(x) \in O] \leq e^{\epsilon}\mathbb{P}[A(x') \in O] + \delta$$

- Parameter $\epsilon > 0$ measures the *privacy loss*.
- Parameter $\delta > 0$ is the *failure probability* where the privacy loss can be much larger than ϵ.

Composition Theorems - Prior Work

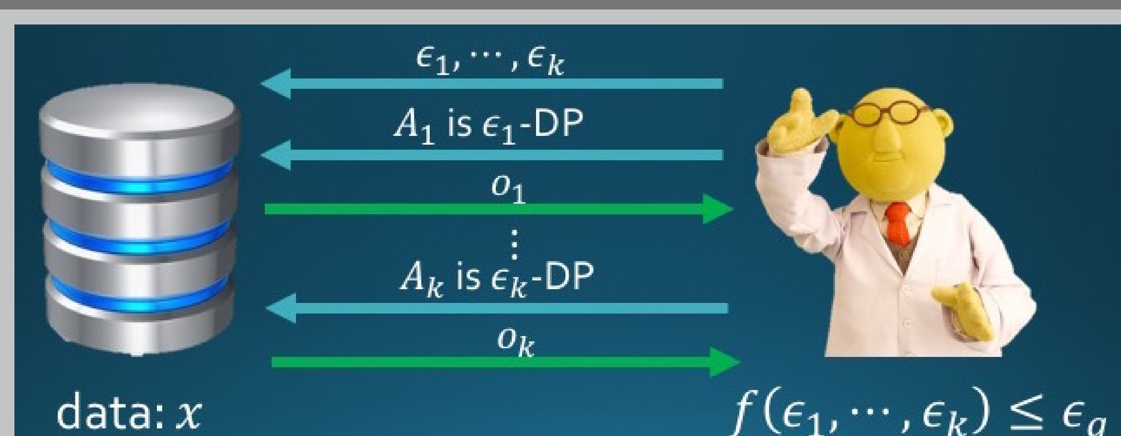

- Run algorithm A_i which is ϵ_i-DP on the data $\mathbf{x}$ as a function of the outcomes of previous algorithms $A_1, \cdots, A_{i-1}$. Then A is (ϵ_g, δ_g)-DP where $f(\epsilon_1, \cdots, \epsilon_k) \leq \epsilon_g$ and
$$A(x) = A_k \circ \cdots \circ A_1(x).$$
- **Basic Composition** [DMNS]: $f(\epsilon_1, \cdots, \epsilon_k) = \sum \epsilon_i$.
- **Advanced Composition** [DRV]: quadratic improvement for $\delta_g > 0$,
$$f(\epsilon_1, \cdots, \epsilon_k) = \tilde{O}\left(\sqrt{\sum \epsilon_i^2}\right).$$
- **Optimal Composition** [KOV, MV]: complex form.

Application: Adaptive Data Analysis

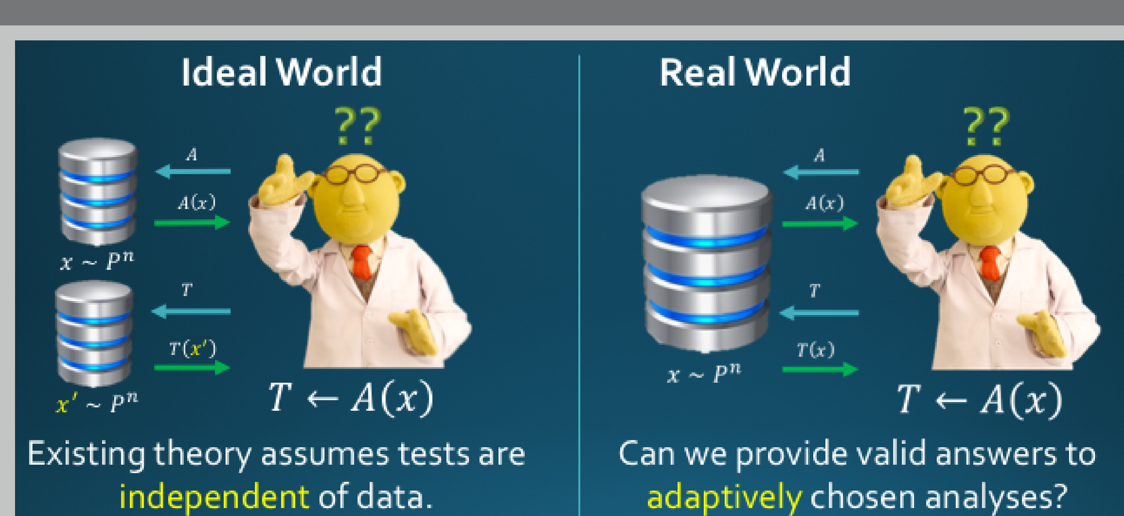

Proposed Solution [DFH+]:

- Limit *information* learned from x through $A(x)$ $\implies$ $A(x)$ and x are "close" to independent.
- One way to limit info is to have analysis be *DP*

Our Focus: Adaptive Privacy Parameters

- As the analyst determines what (and how many) analyses to run he will want to allocate his privacy budget adaptively.
- These composition theorems crucially rely on the choice of parameters ϵ_i and the number of algorithms k to be fixed up front.

- Which composition theorems still apply when we can select the parameters *adaptively*?
- How can we even *define* differential privacy in this adaptively parameter setting?

Privacy Loss and the Analyst

- The privacy loss for algorithm $A : \mathcal{X}^n \to \mathcal{O}$ on neighboring x, x' for outcome $o \in \mathcal{O}$

$$L(o) = \log\left(\frac{\mathbb{P}[A(x) = o]}{\mathbb{P}[A(x') = o]}\right)$$

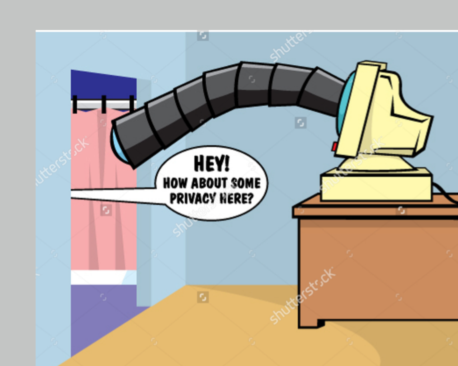

- Privacy loss random variable $L(o)$ where $o \sim A(x)$.
- The *analyst* $\mathcal{A}$ fixes a prob of failure δ_g beforehand.
- $\mathcal{A}$ selects $\epsilon_i \geq 0$ and A_i, which is ϵ_i-DP, as a function of previous outcomes in an adversarial way to try to make the privacy loss large.

Privacy Odometer

- Privacy odometer provides a running upper bound on privacy loss.

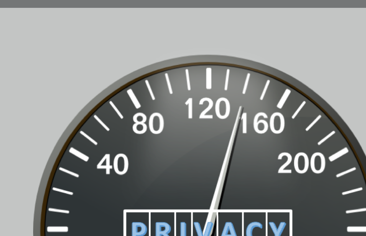

- A *valid privacy odometer* is a function $\texttt{COMP}_{\delta_g} : \mathbb{R}^* \to \mathbb{R}$ where for any analyst $\mathcal{A}$ who selects $\epsilon_1, \cdots, \epsilon_k$ adaptively, then w.p. $\geq 1 - \delta_g$,
$$L(o_1, \cdots, o_k) < \texttt{COMP}_{\delta_g}(\epsilon_1, \cdots, \epsilon_k)$$

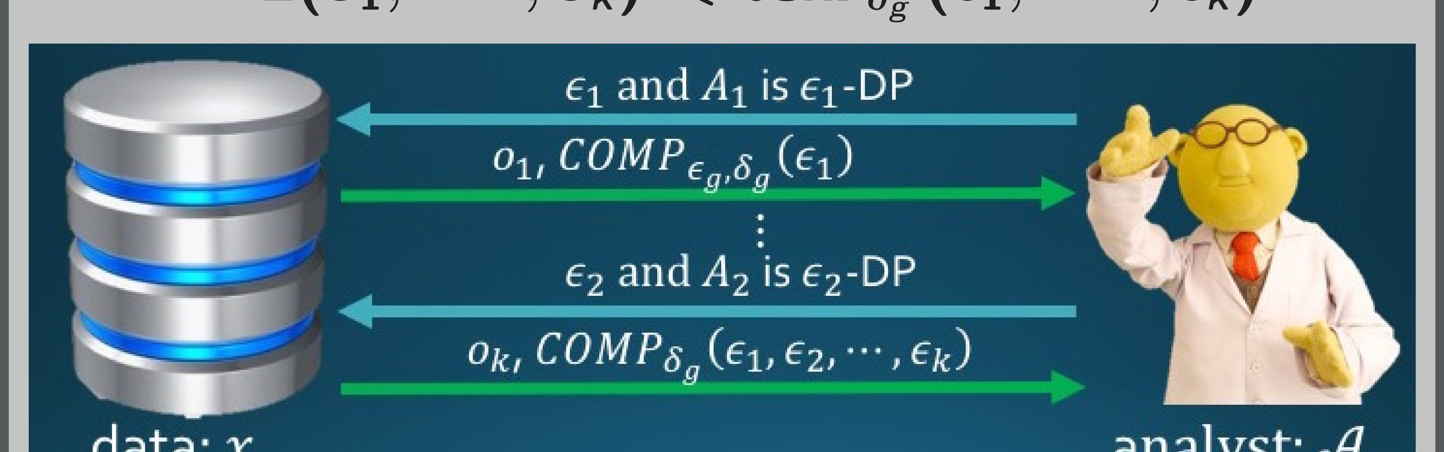

Privacy Filter

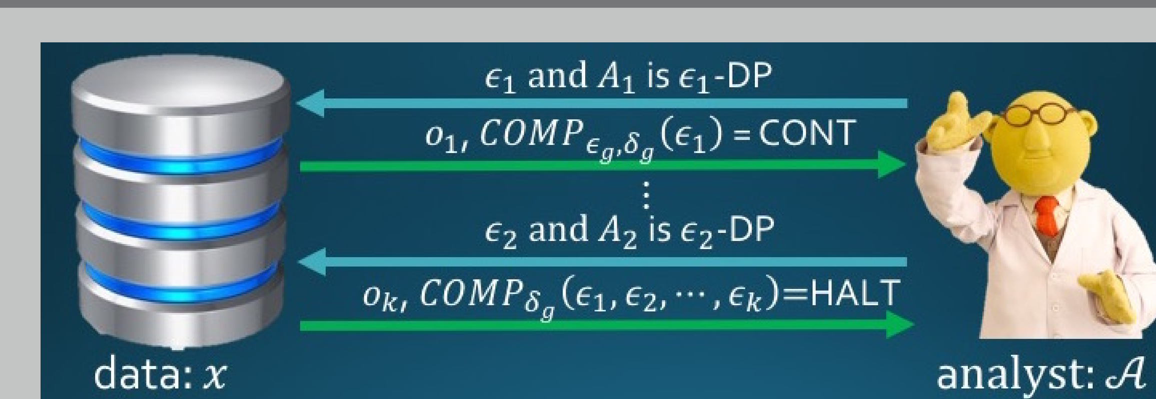

- Privacy filter is a *stopping rule*, so w.h.p. a given privacy budget ϵ_g will not be exceeded.

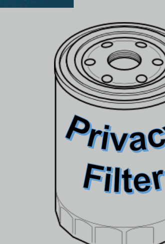

- A *valid privacy filter* $\texttt{COMP}_{\epsilon_g,\delta_g} : \mathbb{R}^k \to \{\texttt{HALT}, \texttt{CONT}\}$ where for any analyst $\mathcal{A}$ who selects $\epsilon_1, \cdots, \epsilon_k$ adaptively, then w.p. $> 1 - \delta_g$, we have $L < \epsilon_g$ and
$$\texttt{COMP}_{\epsilon_g,\delta_g}(\epsilon_1, \cdots, \epsilon_k) = \texttt{HALT}$$
- The mechanism that stops before $\texttt{HALT}$ is (ϵ_g, δ_g)-DP

Main Results

- Basic composition still applies in adaptive setting.
- A valid privacy filter is the following $\texttt{COMP}_{\epsilon_g,\delta_g}(\epsilon_1, \cdots, \epsilon_k) = \texttt{CONT}$ if
$$\tilde{O}\left(\sqrt{\left(\epsilon_g^2 + \sum \epsilon_i^2\right)}\right) < \epsilon_g$$
and otherwise $\texttt{COMP}_{\epsilon_g,\delta_g}(\epsilon_1, \cdots, \epsilon_k) = \texttt{HALT}$.
- A valid privacy odometer is
$$\texttt{COMP}_{\delta_g}(\epsilon_1, \cdots, \epsilon_k) = \tilde{O}\left(\sqrt{\sum \epsilon_i^2 \log\log(n)}\right)$$
as long as $\sum \epsilon_i^2 > 1/n^2$.
- There is a provable gap between privacy filters and odometers – there is *no* valid privacy odometer where $\texttt{COMP}_{\delta_g}(\epsilon_1, \cdots, \epsilon_k)$ is $\tilde{o}\left(\sqrt{\sum \epsilon_i^2 \log\log(n)}\right)$.

Key to Proofs

- The advanced composition theorem used *martingale* concentration inequalities, like Azuma's inequality, but they no longer apply when the bounds are random.
- We then apply concentration bounds from *self normalizing processes* [PnKLL].

References

[DFH+] C. Dwork, V. Feldman, M. Hardt, T. Pitassi, O. Reingold, and A. Roth. In *NIPS'15*.
[DMNS] C. Dwork, F. McSherry, K. Nissim, and A. Smith. In *TCC '06*.
[DRV] C. Dwork, G. Rothblum, and S. Vadhan. In *FOCS '10*.
[KOV] P. Kairouz, S. Oh, and P. Viswanath. In *ICML '15*.
[MV] J. Murtagh and S. Vadhan. In *TCC '16*.
[PnKLL] V. Peña, M. Klass, and T. Leung Lai. *The Annals of Probability '04*.