Leveraging Privacy in Data Analysis and Mechanism Design

Ryan Rogers - Applied Mathematics and Computational Sciences

Penn — University of Pennsylvania

Differential Privacy [Dwork,McSherry,Nissim,Smith'06]

A is (ϵ, δ)-**DP** if for any neighboring $\boldsymbol{x}, \boldsymbol{x}' \in D^n$ and outcome S:

$$\mathbf{Pr}\left[A(\boldsymbol{x}) \in S\right] \leq e^{\epsilon}\,\mathbf{Pr}\left[A(\boldsymbol{x}') \in S\right] + \delta.$$

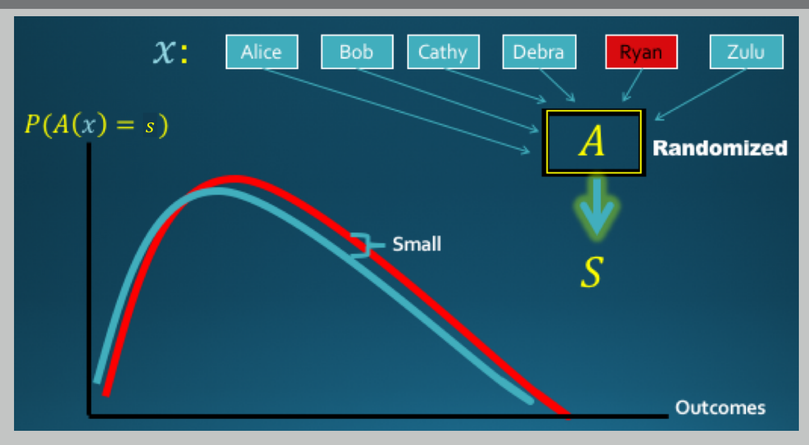

Applications of DP:

- Add privacy as a constraint into an otherwise classical problem.
- Leverage stability guarantees of DP to answer new questions.

Adaptive Data Analysis

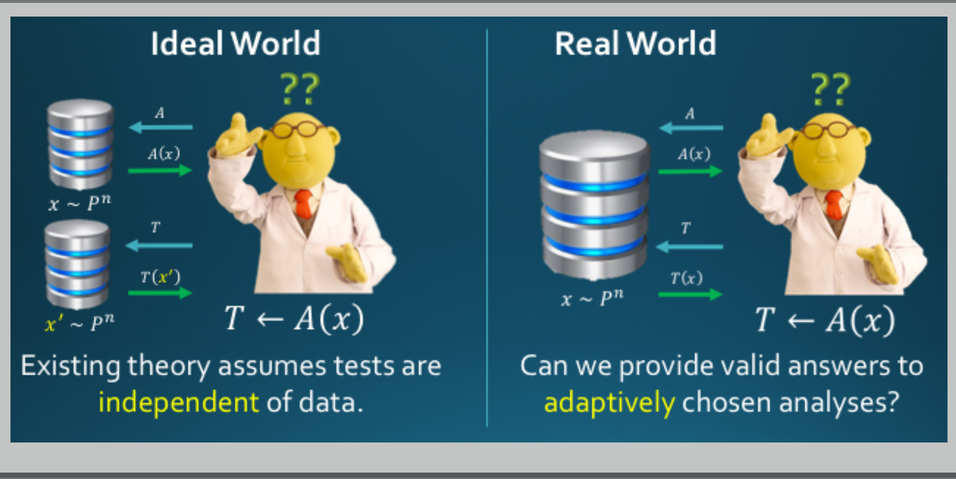

Proposed Solution [Dwork,Feldman,Hardt,Pitassi,Reingold,Roth'15]

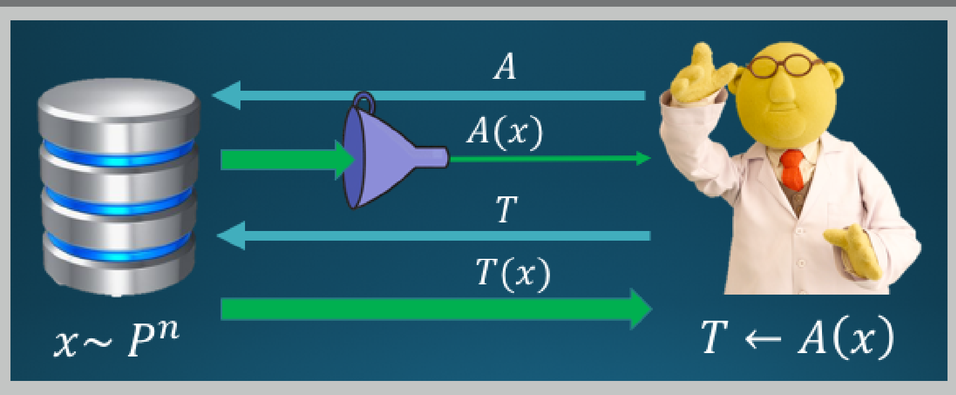

- Limit *information* (max-info) learned from $\boldsymbol{x}$ through $A(x)$ $\implies$ $A(\boldsymbol{x})$ and $\boldsymbol{x}$ are "close" to independent.
- If A is $(\epsilon, 0)$-DP, then it has bounded max-info.

Post-Selection Hypothesis Tests [R,Roth,Smith,Thakkar'16]

- Hypothesis Test: model $H_0 \subseteq \boldsymbol{\Delta}(D^n)$ and statistic $T : D^n \to \{\text{Inconclusive,Reject}\}$
- `False Discovery` if $\boldsymbol{x} \sim P^n \in H_0$ but $T(\boldsymbol{x}) =$Reject.
- **Result**: If A has small max-info then we can bound $\mathbf{Pr}\left[\texttt{False Discovery}\right]$ when $T \leftarrow A(\boldsymbol{x})$.
- **Result**: If A is (ϵ, δ)-DP, then it has bounded max-info.
- k adaptive rounds: max-info scales with k rather than k^2.

DP Hypothesis Tests [Gaboardi,Lim,R,Vadhan'16]

- Want private test s.t. $\mathbf{Pr}\left[\texttt{False Discovery}\right] \leq 0.05$
- Categorical data: $\boldsymbol{x} \sim$Multinomial$(n, \boldsymbol{p})$.
- New private χ^2-tests: independence, GOF ($H_0 : \boldsymbol{p} = \boldsymbol{p}^0$).

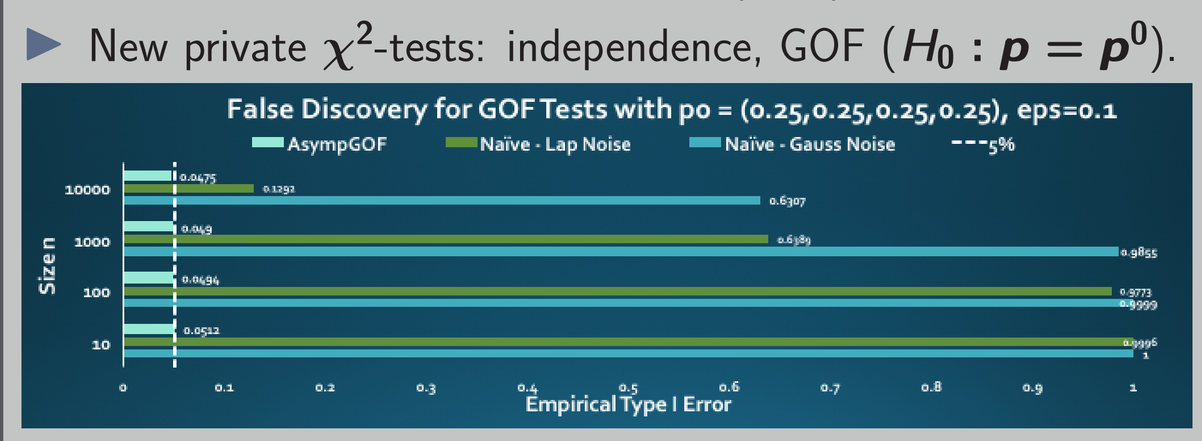

Privacy Odometers and Filters [R,Roth,Ullman,Vadhan'16]

- Privacy definition where analyst may adaptively select privacy parameters and number of analyses.
- **Result**: asymptotically tight composition bounds.

Mechanism Design with Joint DP [Kearns,Pai,R,Roth,Ullman'15]

- DP too strong for applications in mechanism design, so we use a relaxation.
- Mechanism $A : D^n \to Y^n$ is (ϵ, δ)-JDP if for any $\boldsymbol{x}_{-i}, x_i, x_i'$ and $S \subseteq Y^{n-1}$:

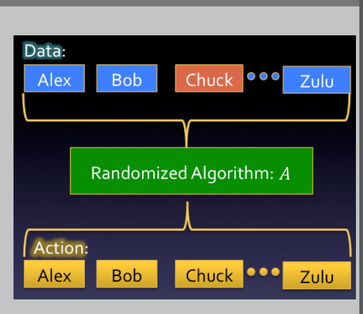

$$\mathbf{Pr}\left[A(\boldsymbol{x}_{-i}, x_i)_{-i} \in S\right] \leq e^{\epsilon}\,\mathbf{Pr}\left[A(\boldsymbol{x}_{-i}, x_i')_{-i} \in S\right] + \delta.$$

Mediators in Large Games

- In *incomplete information* games, we want to coordinate players to play a strategy as if they knew all the information.

- A *mediator* cannot enforce participation and only suggests actions to participants.
- [KP**R**RU'15]: If mediator A is JDP and computes an equilibrium for each input, then using and following the suggestion of A is an *ex-post* Nash equilibrium.
- [KP**R**RU'15] For any large game, there exists such a mediator so that players coordinate to a *correlated equilibrium*.
- [KP**R**RU'15] For large congestion games, there exists such a mediator so that players coordinate to a *Nash equilibrium*.
- [**R**,Roth,Ullman,Wu'15] Further, we can coordinate players to the *socially optimal strategy* if mediator can charge *tolls*.

Private Pareto Optimal Exchange [Kannan,Morgenstern,R,Roth'15]

- In a barter exchange market, we want to privately obtain an *individually rational* and *Pareto optimal* allocation.
- Impossibility result if JDP allocation.
- Positive result for relaxed *Marginal*-DP.
- MDP *Top Trading Cycles Algorithm*.