## Practical Differentially Private Top-k Selection with Pay-what-you-get Composition

in
David Durfee and Ryan Rogers Spotlight Presentation

NeurIPS 2019

in

## Differential Privacy [Dwork, McSherry, Nissim, Smith ‘06]

A randomized algorithm $A: \mathcal{D} \rightarrow \mathcal{Y}$ is $(\varepsilon, \delta)$-DP if for any neighboring data sets $x, x^{\prime} \in \mathcal{D}$ and any outcome $S \subseteq \mathcal{Y}$ we have:

$$
P(A(x) \in S) \leq e^{\varepsilon} P\left(A\left(x^{\prime}\right) \in S\right)+\delta
$$

Privacy loss

## Privatizing Top-k Queries

## Query: Top-10 articles with distinct user engagement?



## Privatizing Top-k Queries

## Query: Top-10 articles with distinct user engagement?



## Privatizing Top-k Queries

## Query: Top-10 articles with distinct user engagement?



## Unknown Domain Setting

- Previous algorithms require knowing the full data domain
- They require adding noise to counts even when the true count is zero
- Typically, the domain is unknown or very large (e.g. all possible articles)
- Lots of prior work for Frequent Itemsets, but requires knowing structure of the data domain universe.

Existing Systems for Data Analytics


First Attempt


## First Attempt



## Firstricempt



## Second Attempt - Include Threshold



## Second Attempt - Include Threshold



## BR Composition versus Optimal DP Composition



## Linked in

Come to Poster \#161 for more details

