Privacy-Preserving Information Release from Social Networks

Avrim Blum
Anupam Datta
Jeremiah Blocki
Or Sheffet

The problem

- Given a sensitive social network:

Friendship network

Email or phone-call graph

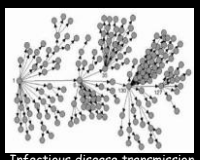
Infectious disease transmission network

- Can we release a sanitized version or (noisy) information about it an a way that preserves privacy and is still useful?

How can one formally talk about preserving privacy?

The science of privacy

- Fundamental breakthrough came from MSR in work of [Dwork-McSherry-Nissim-Smith] building on earlier work of Dwork et al, in definition of differential privacy.

A semantically-meaningful,
composable,
plausibly achievable,
notion of privacy

The science of privacy

- Fundamental breakthrough came from MSR in work of [Dwork-McSherry-Nissim-Smith] building on earlier work of Dwork et al, in definition of differential privacy.

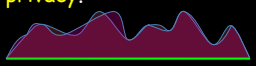

Any participant should be able to plausibly deny any fact claimed about them (the probability of any given output of mechanism would change by only $1 \pm \epsilon$)

The science of privacy

- Fundamental breakthrough came from MSR in work of [Dwork-McSherry-Nissim-Smith] building on earlier work of Dwork et al, in definition of differential privacy.
- Powerful tools, including
  - Smooth sensitivity, sample-and-aggregate of [Nissim-Raskhodnikova-Smith]
  - Exponential mechanism of [McSherry-Talwar] and use for outputting sanitized databases of [B-Ligett-Roth]
  - Composition theorems of [Dwork-Rothblum-Vadhan]
  - Multiplicative weights of [Hardt-Talwar] and Iterative Database Construction of [Gupta-Roth-Ullman]

What can be done when your "database" is a social network?

Thrust #1: output sanitized network approximately preserving cut-values.

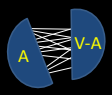

- Natural for understanding high-level connectivity from various sets A to the rest of the network.

What can be done when your "database" is a social network?

Thrust #1: output sanitized network approximately preserving cut-values.

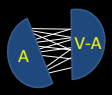

- Recent work of [Gupta-Roth-Ullman] gave interactive query-answering protocol with good properties.

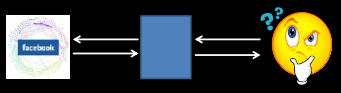

What can be done when your "database" is a social network?

Thrust #1: output sanitized network approximately preserving cut-values.

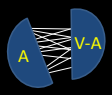

- Our work: produce output s.t. for (almost) all sets A, error is only $O(|A|)$ [+ multiplicative $(1 \pm \epsilon)$]

What can be done when your "database" is a social network?

Thrust #1: output sanitized network approximately preserving cut-values.

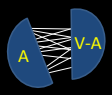

How it works:

1. Each edge (i,j) picks random $w_{ij} \sim N(0,1)$. Sends w_{ij} to i, $-w_{ij}$ to j. [Technically, non-edges do this too, scaled by $O(1/n)$]
2. Vertices add up and publish their sums.

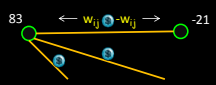

So, what happens if add these up over all $i \in A$?

What can be done when your "database" is a social network?

Thrust #1: output sanitized network approximately preserving cut-values.

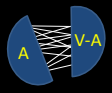

How it works:

- Repeat several times to get good estimate of $|E(A,V\text{-}A)|$.
- Equivalent to random (JL) projection of edge-adjacency matrix.
- Privacy is wrt individual edge changes. Doing anything wrt arbitrary node changes is much harder….

What can be done when your "database" is a social network?

Thrust #2: answering queries while preserving privacy wrt arbitrary node changes (node privacy).

- Considering queries of form: how many nodes participate in a given local pattern?

- Note: could be very sensitive to single node change (adding many nbrs, changing profession)…

What can be done when your "database" is a social network?

Thrust #2: answering queries while preserving privacy wrt arbitrary node changes (node privacy).

- But would be OK if, e.g., all nodes had low degree.

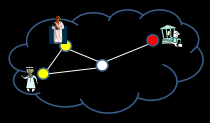

- Note: could be very sensitive to single node change (adding many nbrs, changing profession)…

What can be done when your "database" is a social network?

Thrust #2: answering queries while preserving privacy wrt arbitrary node changes (node privacy).

Approach:

- User specifies property $\mathcal{P}$ of graphs (e.g., degree $\leq$ k) s.t. query has low sensitivity over $\mathcal{P}$, and user believes $G \in \mathcal{P}$.

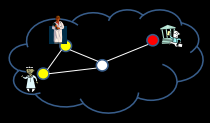

- Want to preserve privacy no matter what.
- And give accurate answer if $G \in \mathcal{P}$.

What can be done when your "database" is a social network?

Thrust #2: answering queries while preserving privacy wrt arbitrary node changes (node privacy).

Approach:

- Method #1: construct modified query q' that is smooth over all graphs, and equiv to actual query q over the nice graphs.

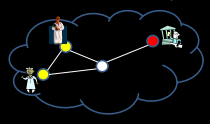

- Method #2: map graphs into the nice set in a smooth manner.
- See poster (manned by Or and Jeremiah) for details.

Additional notes

- In discussion for visits to MSR-SVC.
- 2010-2011 support in area of algorithmic pricing led to results enabling successful NSF ICES grant. (interface b/w CS, Economics, and Social Sciences)
- Current support has enabled collaboration that is very helpful in connection to current NSF center proposal.

Thank you for your support!