Private Online Community Detection for Censored Block Models

Mohamed Seif*, **Liyan Xie**†, Andrea J. Goldsmith*, H. Vincent Poor*

*Department of Electrical and Computer Engineering, Princeton University

†Department of Industrial and Systems Engineering, University of Minnesota

Department of Statistics, Iowa State University, Feb 2025

Outline

1. Motivation and Problem Setup: Censored Block Models
2. Private Online Detection Procedures
3. Numerical Examples
4. Open problems & Challenges…

Biological Networks

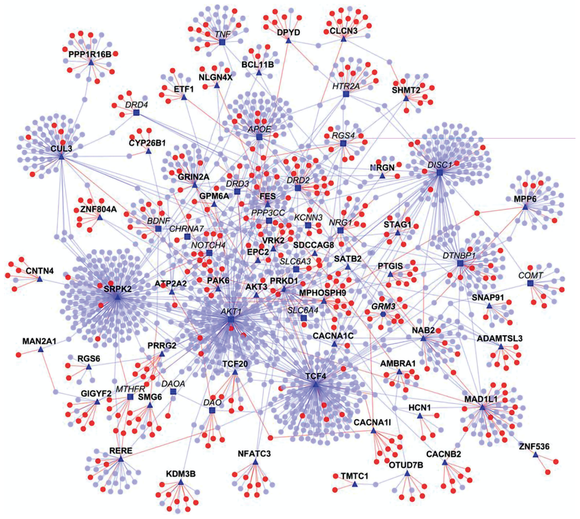

Contact-tracing Networks

Collaboration Networks

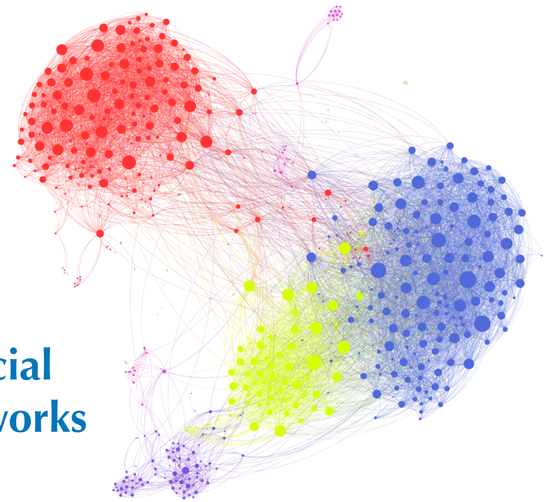

Social Networks

Data Analysis for Graphs

- Much of the data of scientific interest comes in the form of **graphs**

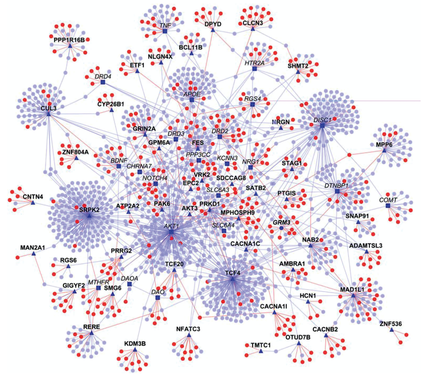

Biological

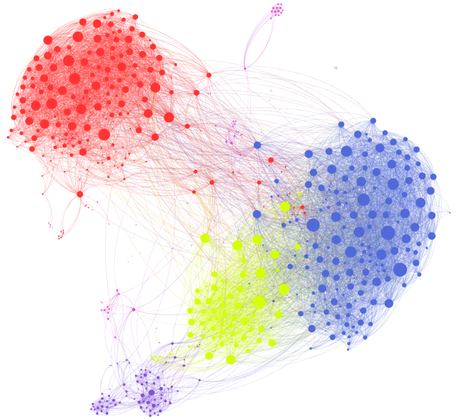

Social Networks

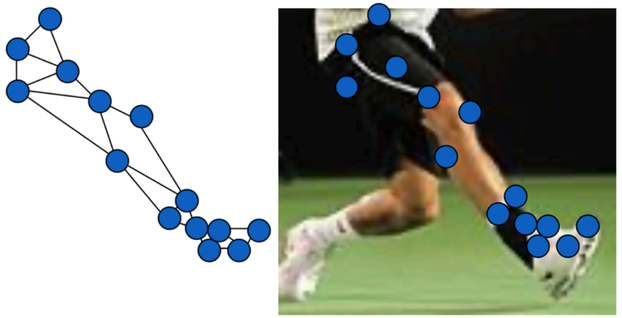

Image Segmentation

Data Analysis for Graphs

- Much of the data of scientific interest comes in the form of **graphs**

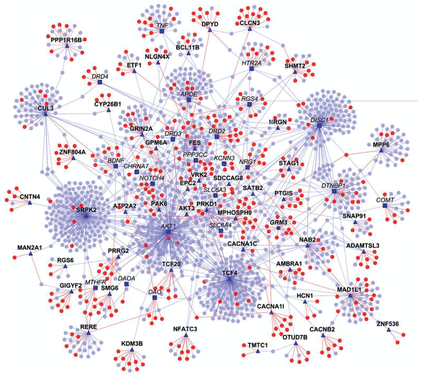

Biological

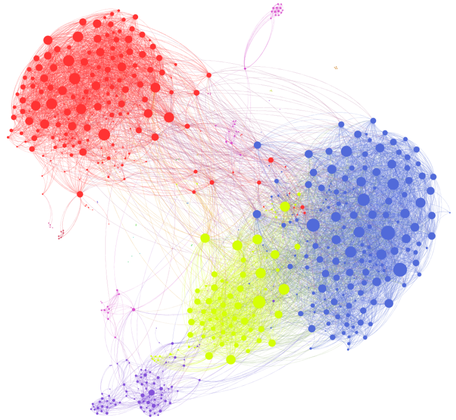

Social Networks

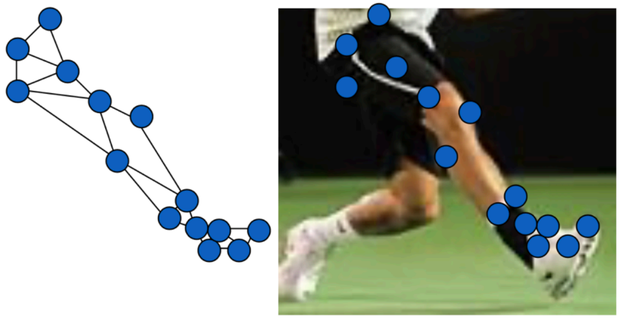

Image Segmentation

- **Dynamic** data: Change in community memberships

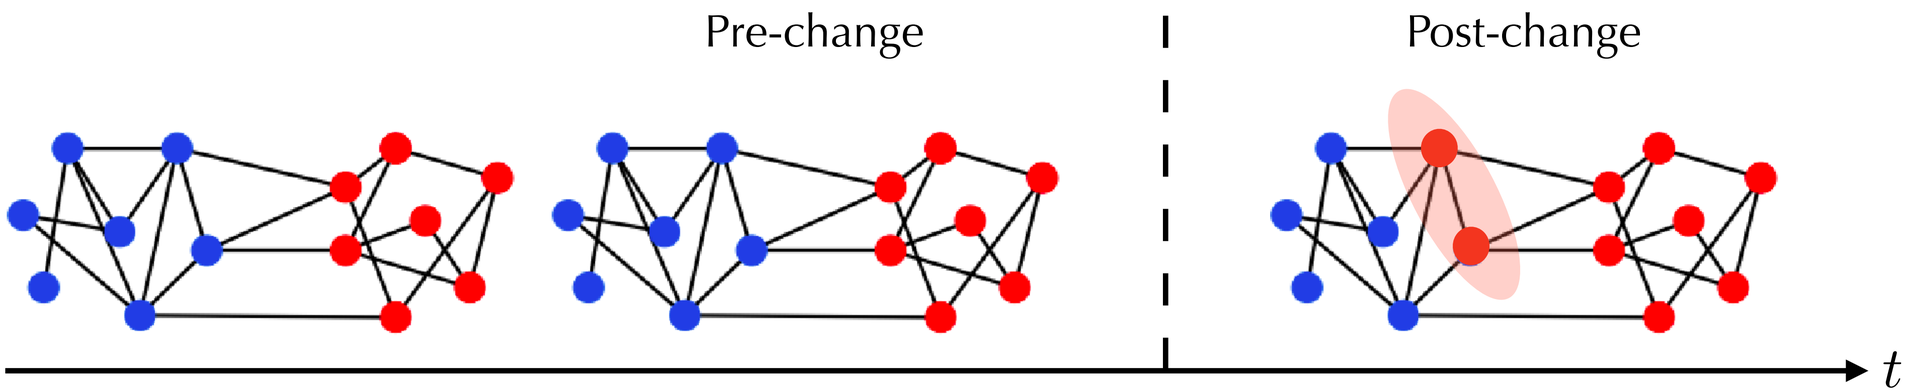

Dynamics in Graphs

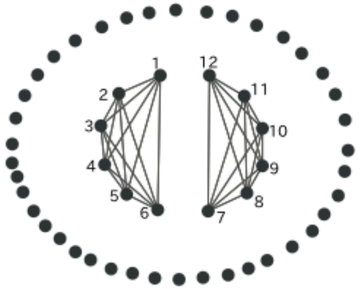

(a) Original Graph

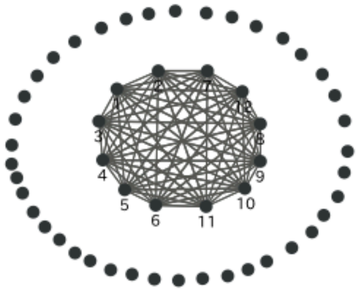

(b) Merge

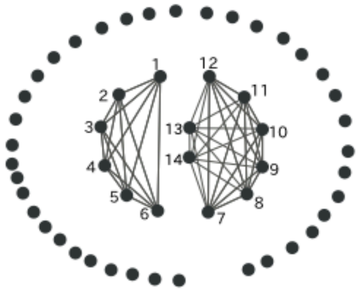

(d) Growth

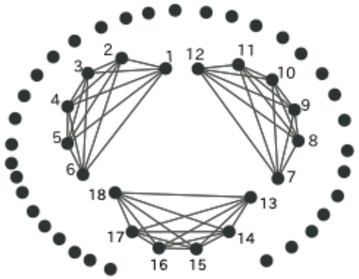

(f) Birth

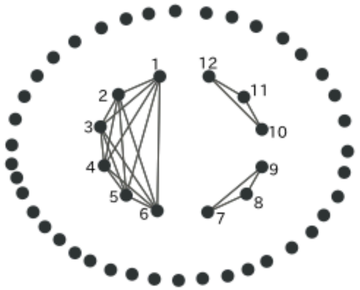

(c) Split

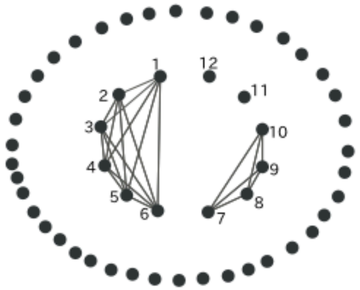

(e) Shrink

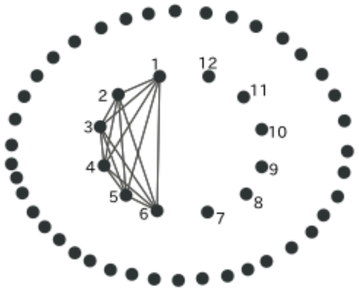

(g) Death

Privacy Threats in Information Networks

- Growing aggregation & use of personal data

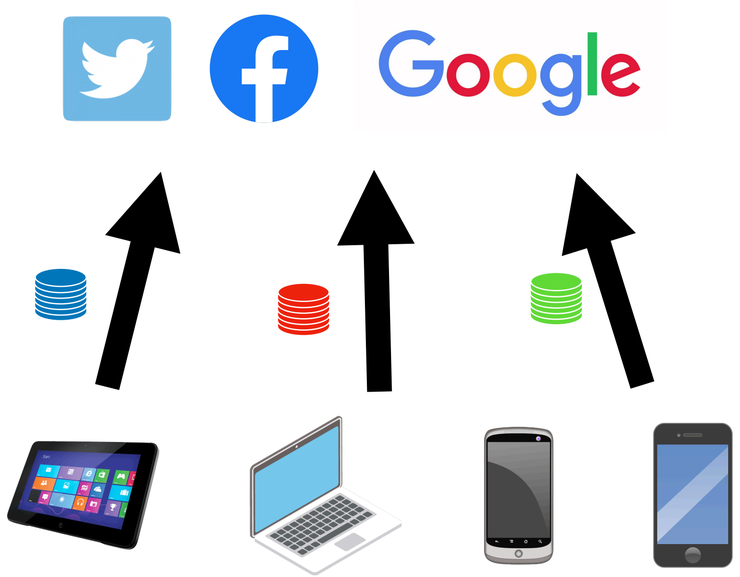

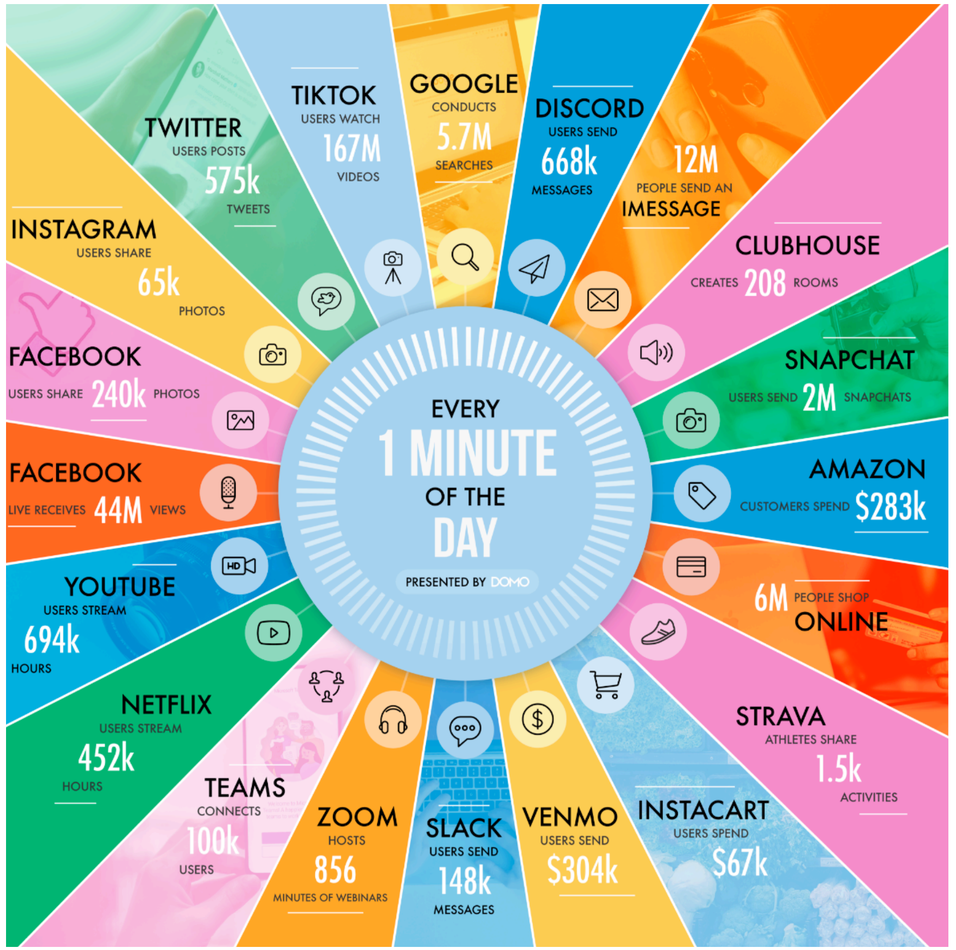

Privacy Threats in Information Networks

- Growing aggregation & use of personal data

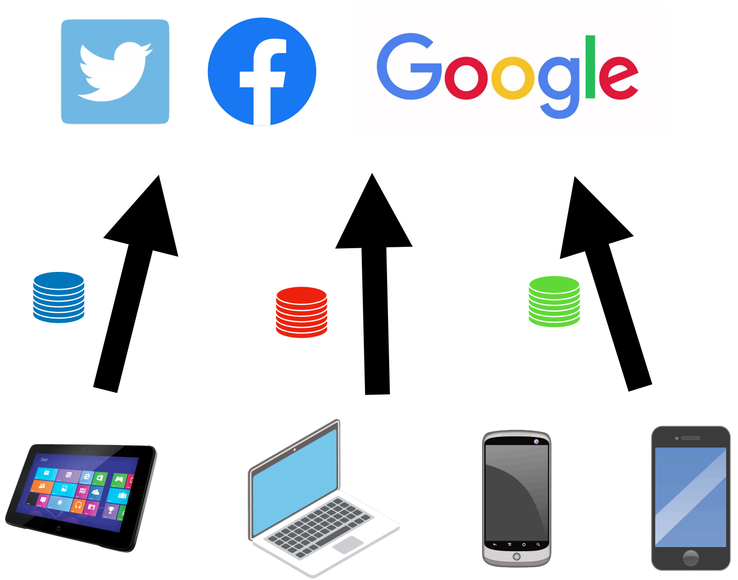

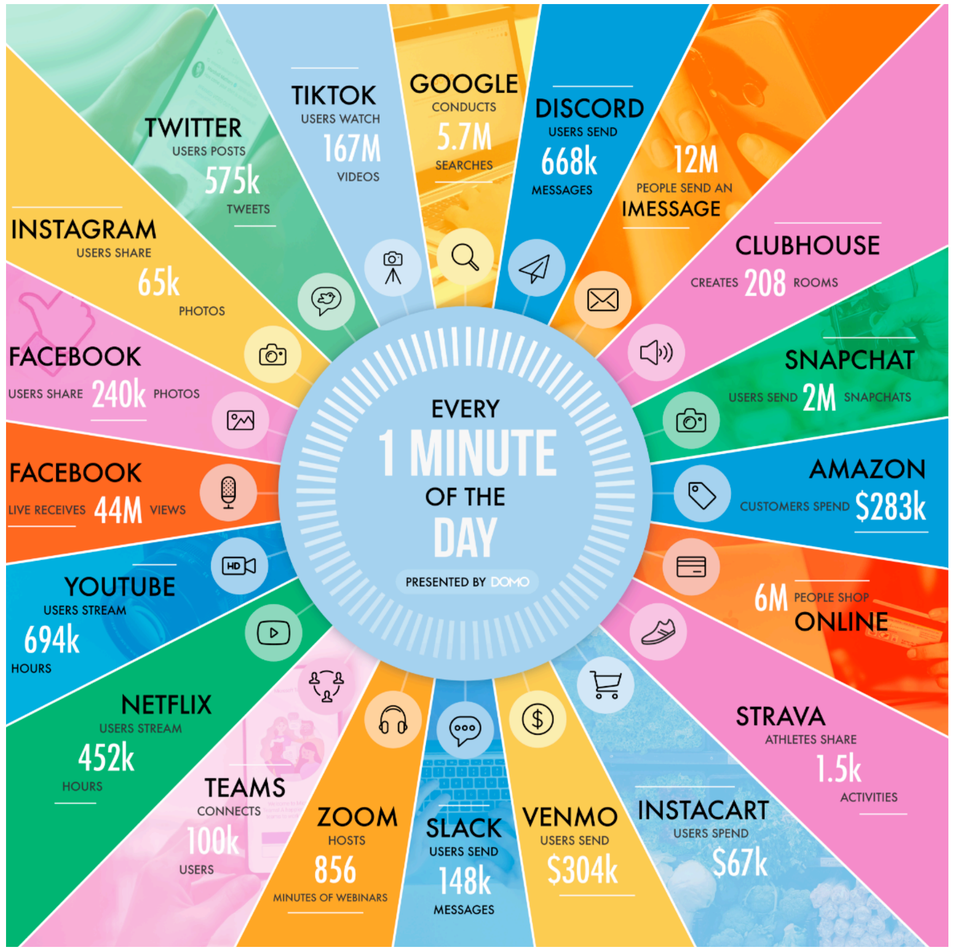

- Data Release Applications
  - Google Maps: traffic monitoring, location tracking
  - YouTube: recommendations
  - Apple: learns frequent emojis

Privacy Threats in Information Networks

CNN BUSINESS Markets **Tech** Media Success Perspectives Videos LIVE TV Ed

Half a billion Facebook users' information posted on hacking website, cyber experts say

By Donie O'Sullivan, CNN Business

Updated 7:01 AM ET, Mon April 5, 2021

Forbes

Sep 23, 2019, 05:00am EDT | 154,130 views

Security Warning For 23 Million YouTube Creators Following 'Massive' Hack Attack

Davey Winder Senior Contributor

Cybersecurity

I report and analyse breaking cybersecurity and privacy stories

DIGIT NEWS News Eve

Share

Scottish Power Parent Company Iberdrola Hit by Cyber-attack

Michael Behr

04 April 2022, 01.15pm

Privacy Threats in Information Networks

CNN BUSINESS Markets **Tech** Media Success Perspectives Videos LIVE TV Ed

Half a billion Facebook users' information posted on hacking website, cyber experts say

By Donie O'Sullivan, CNN Business

Updated 7:01 AM ET, Mon April 5, 2021

Forbes

Sep 23, 2019, 05:00am EDT | 154,130 views

Security Warning For 23 Million YouTube Creators Following 'Massive' Hack Attack

Davey Winder Senior Contributor

Cybersecurity

I report and analyse breaking cybersecurity and privacy stories

DIGIT NEWS News Eve

Share

Scottish Power Parent Company Iberdrola Hit by Cyber-attack

Michael Behr

04 April 2022, 01.15pm

Research Question: How to detect community change in an online and private way?

Problem Setup: Censored Block Models (CBMs)

- A random graph model with *two* communities

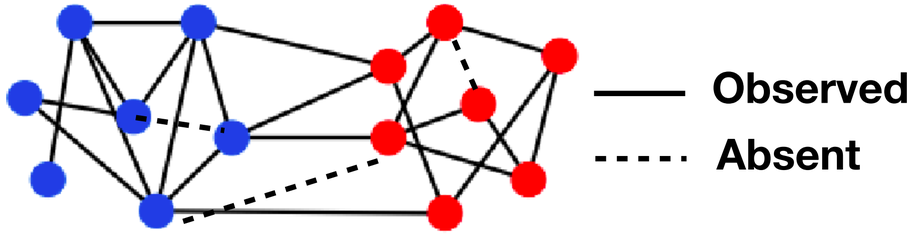

$\mathrm{CBM}(n, p, \zeta)$

Nodes **within** a community tend to connect **more** compared to nodes across different communities

Problem Setup: Censored Block Models (CBMs)

- A random graph model with *two* communities
  - n nodes

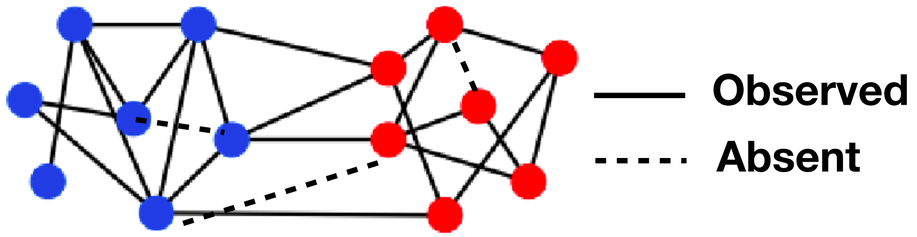

$\mathrm{CBM}(n, p, \zeta)$

Nodes **within** a community tend to connect **more** compared to nodes across different communities

Problem Setup: Censored Block Models (CBMs)

- A random graph model with *two* communities
  - n nodes
  - Community **labels**: $\sigma_i \in \{-1, +1\}, \forall i \in [n]$

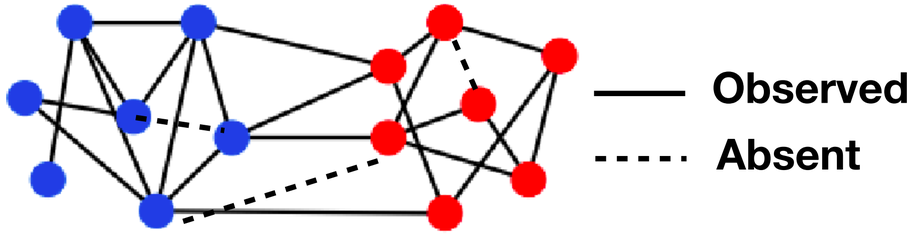

$\mathrm{CBM}(n, p, \zeta)$

Nodes **within** a community tend to connect **more** compared to nodes across different communities

Problem Setup: Censored Block Models (CBMs)

- A random graph model with *two* communities
  - n nodes
  - Community **labels**: $\sigma_i \in \{-1, +1\}, \forall i \in [n]$
  - Given the labels, edges are drawn independently

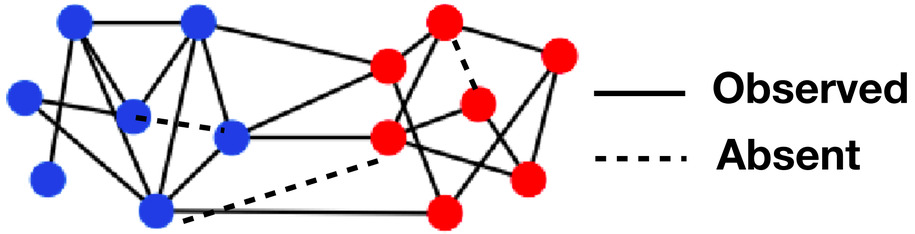

$\mathrm{CBM}(n, p, \zeta)$

Nodes **within** a community tend to connect **more** compared to nodes across different communities

Problem Setup: Censored Block Models (CBMs)

- A random graph model with *two* communities
  - n nodes
  - Community **labels**: $\sigma_i \in \{-1, +1\}, \forall i \in [n]$
  - Given the labels, edges are drawn independently
  - The adjacency matrix $\mathbf{A}$, with $A_{i,j} \in \{+1, -1, 0\}$

$$\mathbb{P}(A_{i,j} = a) = \begin{cases} p(1-\zeta) & a = \sigma_i \cdot \sigma_j \\ p\zeta & a = -\sigma_i \cdot \sigma_j \\ 1-p & a = 0 \end{cases}$$

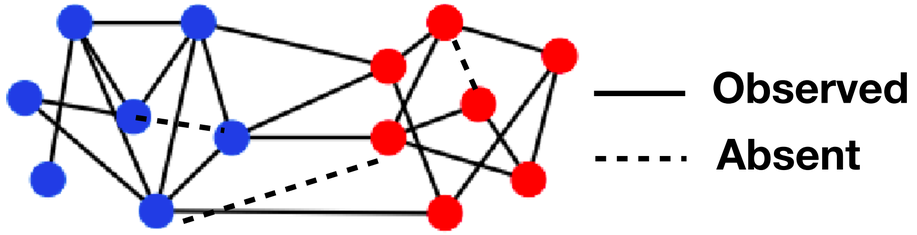

$\mathrm{CBM}(n, p, \zeta)$

Nodes **within** a community tend to connect **more** compared to nodes across different communities

Problem Setup: Censored Block Models (CBMs)

- A random graph model with *two* communities
  - n nodes
  - Community **labels**: $\sigma_i \in \{-1, +1\}, \forall i \in [n]$
  - Given the labels, edges are drawn independently
  - The adjacency matrix $\mathbf{A}$, with $A_{i,j} \in \{+1, -1, 0\}$

$$\mathbb{P}(A_{i,j} = a) = \begin{cases} p(1-\zeta) & a = \sigma_i \cdot \sigma_j \\ p\zeta & a = -\sigma_i \cdot \sigma_j \\ 1-p & a = 0 \end{cases}$$

  - Some edges might be **censored** (due to some errors)

    p : probability of an edge is being observed

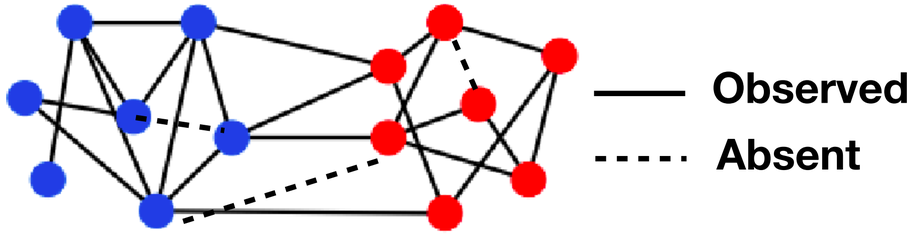

$\mathrm{CBM}(n, p, \zeta)$

Nodes **within** a community tend to connect **more** compared to nodes across different communities

Problem Setup: Change-point in CBMs

- Censored stochastic block model with two communities
- **Change point** at time v

$$\sigma_t^* = \begin{cases} \sigma^{\text{pre}}, & t = 1, \dots, \nu - 1, \\ \sigma^{\text{post}}, & t = \nu, \tau + 1, \dots \end{cases}$$

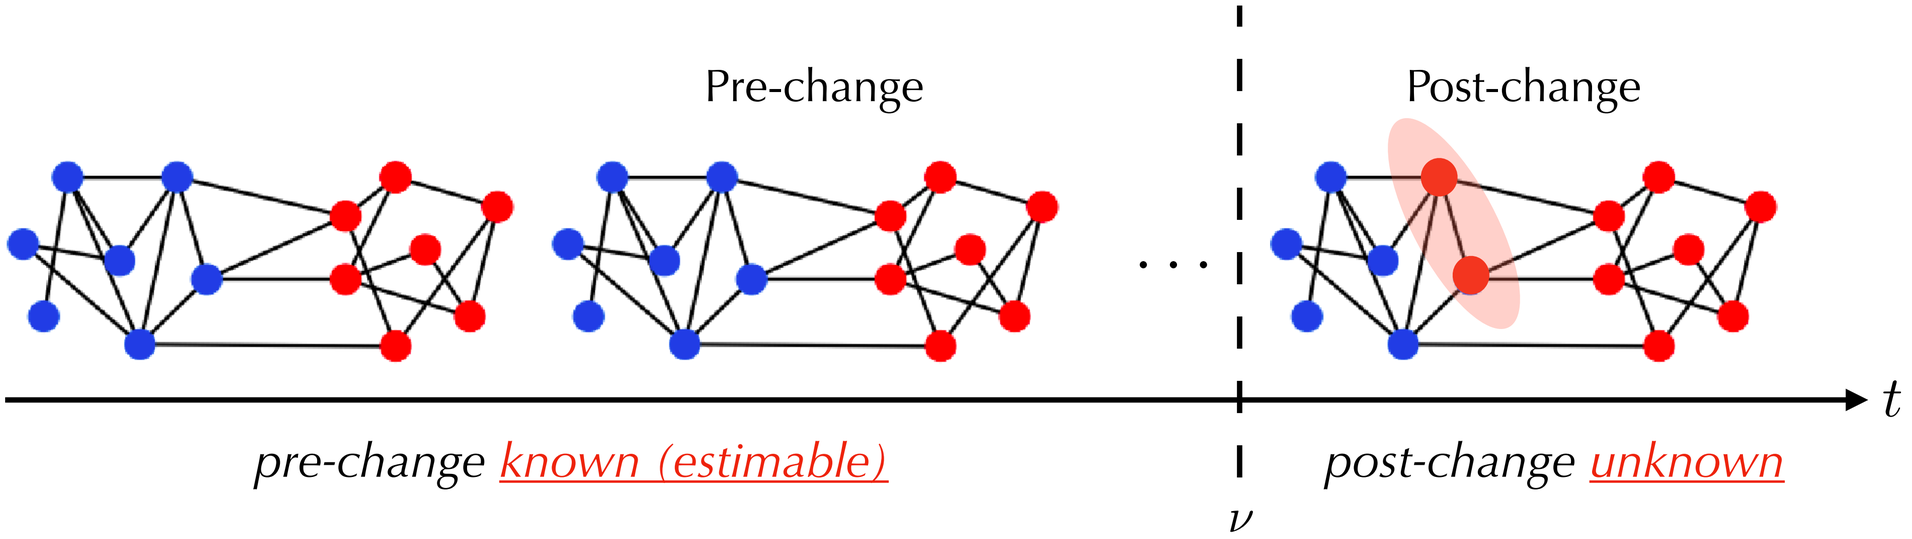

Problem Setup: Change-point in CBMs

- **Goal:** Detect the unknown change-point as quickly as possible, while controlling the false alarm rate and subject to privacy requirements

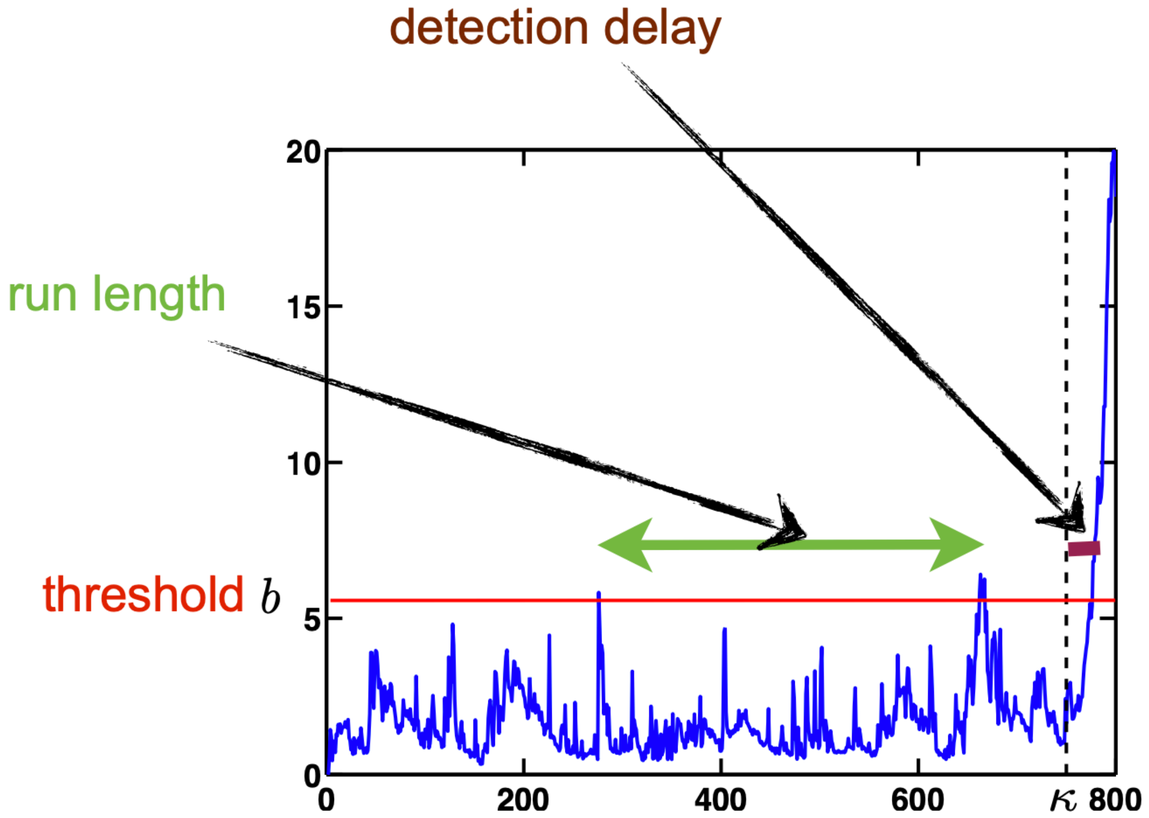

Problem Setup: *Graph* Differential Privacy (DP)

- Edge DP:
  - Hide the presence or absence of an edge
  - **Notion:** (ϵ, δ)- edge DP

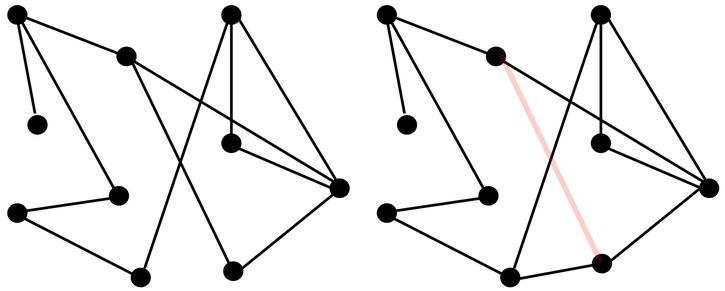

$$\Pr(\hat{\sigma}(\mathbf{A}) = \sigma) \leq e^{\epsilon} \Pr(\hat{\sigma}(\mathbf{A}') = \sigma) + \delta$$

$\forall \mathbf{A}, \mathbf{A}'$ that differ in one edge

- Romantic relationships
- "Friendships"
- Financial transactions

Problem Setup: *Graph* Differential Privacy (DP)

- Edge DP:
  - Hide the presence or absence of an edge
  - **Notion:** (ϵ, δ)- edge DP

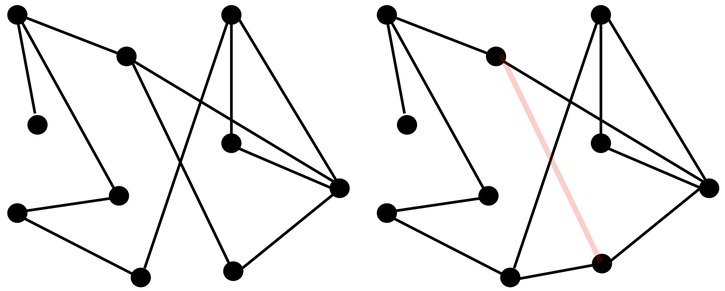

$$\Pr(\hat{\sigma}(\mathbf{A}) = \sigma) \leq e^{\epsilon} \Pr(\hat{\sigma}(\mathbf{A}') = \sigma) + \delta$$

$\forall \mathbf{A}, \mathbf{A}'$ that differ in one edge

- Romantic relationships
- "Friendships"
- Financial transactions

- The Privacy Toolkit: Techniques to achieve *Graph* DP

Problem Setup: *Graph* Differential Privacy (DP)

- Edge DP:
  - Hide the presence or absence of an edge
  - **Notion:** (ϵ, δ)- edge DP

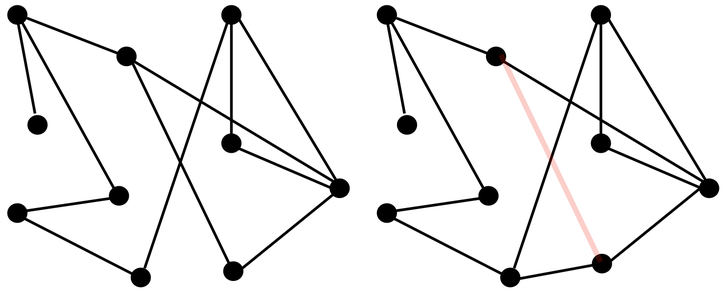

$$\Pr(\hat{\sigma}(\mathbf{A}) = \sigma) \leq e^{\epsilon} \Pr(\hat{\sigma}(\mathbf{A}') = \sigma) + \delta$$

$\forall \mathbf{A}, \mathbf{A}'$ that differ in one edge

  - Romantic relationships
  - "Friendships"
  - Financial transactions

- The Privacy Toolkit: Techniques to achieve *Graph* DP

Data (Adjacency matrix $\mathbf{A}$)

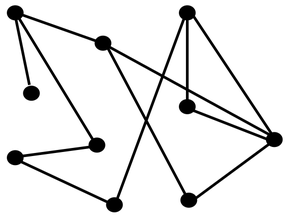

Goal: Compute the community estimate $\hat{\sigma}(\mathbf{A})$ subject to DP

Problem Setup: *Graph* Differential Privacy (DP)

- Edge DP:
  - Hide the presence or absence of an edge
  - **Notion:** (ϵ, δ)- edge DP

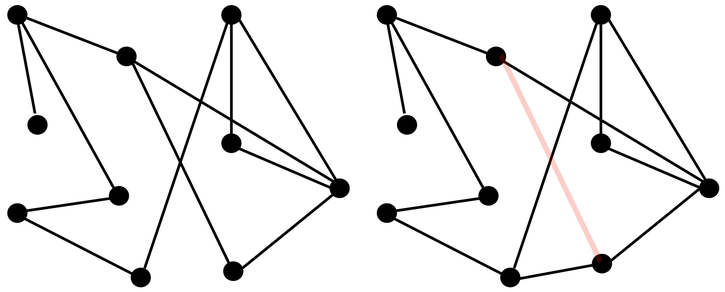

$$\Pr(\hat{\sigma}(\mathbf{A}) = \sigma) \leq e^{\epsilon} \Pr(\hat{\sigma}(\mathbf{A}') = \sigma) + \delta$$

$$\forall \mathbf{A}, \mathbf{A}' \text{ that differ in one edge}$$

- Romantic relationships
- "Friendships"
- Financial transactions

- The Privacy Toolkit: Techniques to achieve *Graph* DP

Data (Adjacency matrix $\mathbf{A}$)

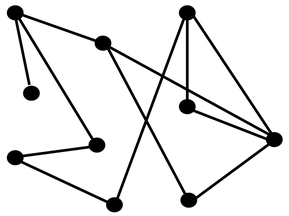

Goal: Compute the community estimate $\hat{\sigma}(\mathbf{A})$ subject to DP

Input Perturbation (Local-DP)

- "Perturb" data: $\mathbf{A} \rightarrow \tilde{\mathbf{A}}$
- Compute $\hat{\sigma}(\tilde{\mathbf{A}})$

Problem Setup: *Graph* Differential Privacy (DP)

- Edge DP:
  - Hide the presence or absence of an edge
  - **Notion:** (ϵ, δ)- edge DP

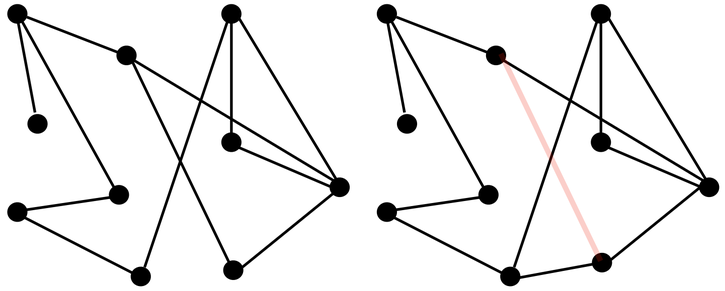

$$\Pr(\hat{\sigma}(\mathbf{A}) = \sigma) \leq e^{\epsilon} \Pr(\hat{\sigma}(\mathbf{A}') = \sigma) + \delta$$

$\forall \mathbf{A}, \mathbf{A}'$ that differ in one edge

- Romantic relationships
- "Friendships"
- Financial transactions

- The Privacy Toolkit: Techniques to achieve *Graph* DP

Data (Adjacency matrix $\mathbf{A}$)

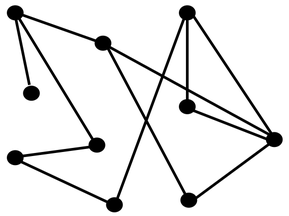

Goal: Compute the community estimate $\hat{\sigma}(\mathbf{A})$ subject to DP

Input Perturbation (Local-DP)
- "Perturb" data: $\mathbf{A} \rightarrow \tilde{\mathbf{A}}$
- Compute $\hat{\sigma}(\tilde{\mathbf{A}})$

Output Perturbation (Central-DP)
- Compute $\hat{\sigma}(\mathbf{A})$ on the original data
- "Perturb" $\hat{\sigma}(\mathbf{A})$ and then release

Related work

Related work

- Private change detection
  - Known pre- and post-change [Cummings et al, 2018]
  - Unknown but 1-d distributions [Cummings et al, 2020]
  - Multivariate nonparametric regression under local DP [Berrett & Yu, 2021]
  - Offline change-point detection (localization) [Li et al, 2021]

Related work

- Private change detection
  - Known pre- and post-change [Cummings et al, 2018]
  - Unknown but 1-d distributions [Cummings et al, 2020]
  - Multivariate nonparametric regression under local DP [Berrett & Yu, 2021]
  - Offline change-point detection (localization) [Li et al, 2021]
- Community recovery (estimation)
  - Recovery conditions for SBM and CBM under dense graphs [Hajek et al., 2016)] [Dhara et al. 2021] and sparse graphs [Saade et al., 2015], [Lelarge et al., 2013]
  - Exact recovery conditions under DP constraints [Seif et al, 2022], [Seif et al, 2024].

Related work

- Private change detection
  - Known pre- and post-change [Cummings et al, 2018]
  - Unknown but 1-d distributions [Cummings et al, 2020]
  - Multivariate nonparametric regression under local DP [Berrett & Yu, 2021]
  - Offline change-point detection (localization) [Li et al, 2021]
- Community recovery (estimation)
  - Recovery conditions for SBM and CBM under dense graphs [Hajek et al., 2016)] [Dhara et al. 2021] and sparse graphs [Saade et al., 2015], [Lelarge et al., 2013]
  - Exact recovery conditions under DP constraints [Seif et al, 2022], [Seif et al, 2024].

No comprehensive work analyzing the impact of privacy, efficiency & change detection in CBMs

Outline

1. Motivation and Problem Setup: Censored Block Models

2. **Private Online Detection Procedures**

3. Numerical Examples

4. Open problems & Challenges…

Biological Networks

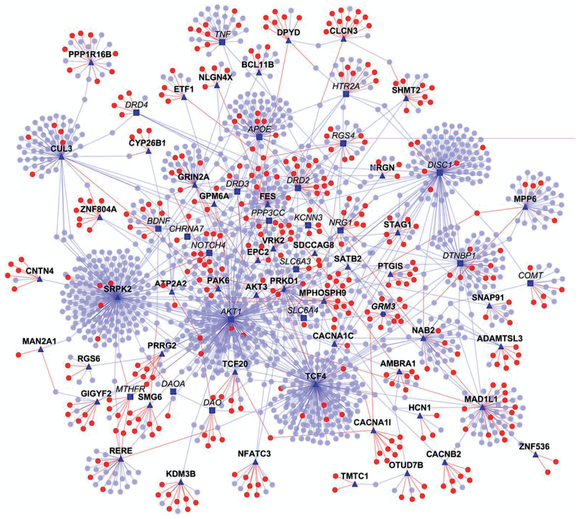

Contact-tracing Networks

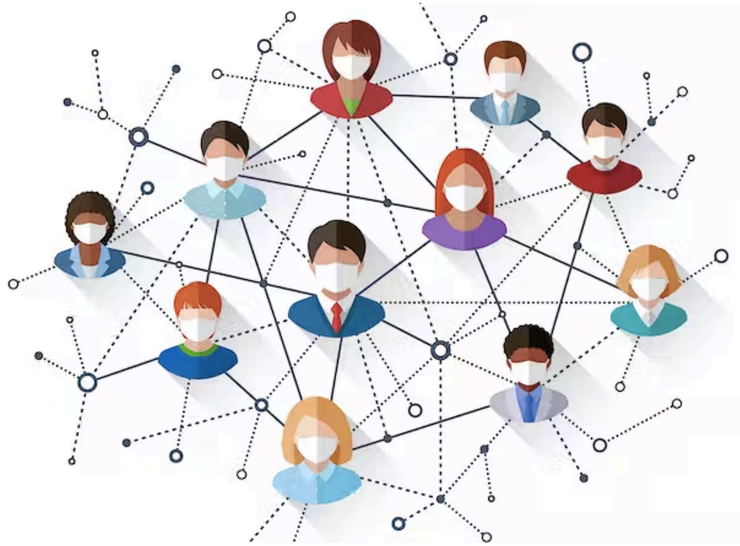

Collaboration Networks

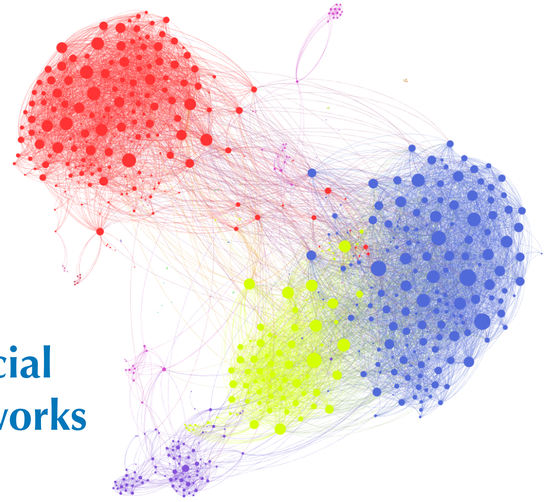

Social Networks

Overview: Proposed two-stage private detection

- Proposed Algorithm: Private adaptive cumulative sum (CUSUM) test

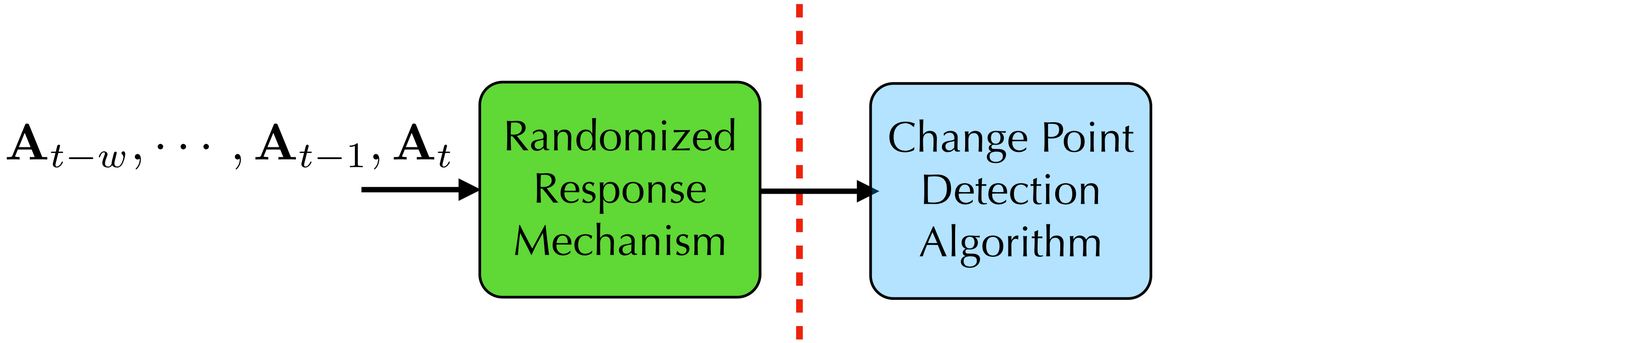

Adaptive CUSUM type detector

Overview: Proposed two-stage private detection

- Proposed Algorithm: Private adaptive cumulative sum (CUSUM) test

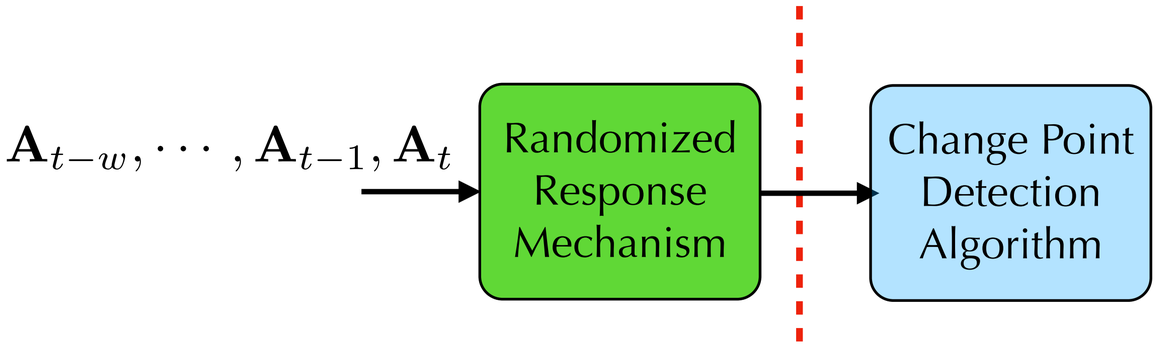

Mechanism (1): Input Perturbation (Local)

- "Perturb" data: $\mathbf{A} \to \tilde{\mathbf{A}}$
- Compute $\hat{\sigma}(\tilde{\mathbf{A}})$

Adaptive CUSUM type detector

$$S_t = (S_{t-1})^+ + \log \frac{\Pr(\tilde{\mathbf{A}}_t; \hat{\sigma}_{t-1})}{\Pr(\tilde{\mathbf{A}}_t; \sigma^{\text{pre}})}$$

Overview: Proposed two-stage private detection

- Proposed Algorithm: Private adaptive cumulative sum (CUSUM) test

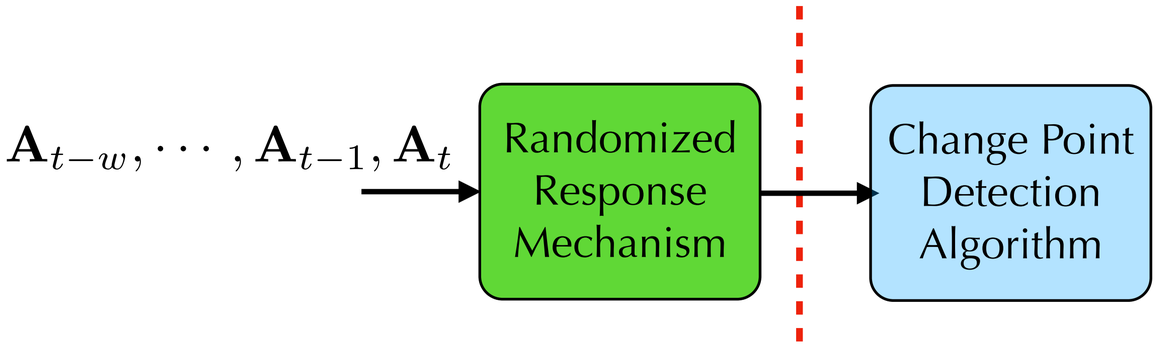

Adaptive CUSUM type detector

Mechanism (1): Input Perturbation (Local)

- "Perturb" data: $\mathbf{A} \to \tilde{\mathbf{A}}$
- Compute $\hat{\sigma}(\tilde{\mathbf{A}})$

$$\longrightarrow \quad S_t = (S_{t-1})^+ + \log \frac{\Pr(\tilde{\mathbf{A}}_t; \hat{\sigma}_{t-1})}{\Pr(\tilde{\mathbf{A}}_t; \sigma^{\text{pre}})}$$

Mechanism (2): Output Perturbation (Central)

- Compute $\hat{\sigma}(\mathbf{A})$ on the original data
- "Perturb" $\hat{\sigma}(\mathbf{A})$ and then release

$$\longrightarrow \quad S_t = (S_{t-1})^+ + \log \frac{\Pr(\mathbf{A}_t; \hat{\sigma}_{t-1})}{\Pr(\mathbf{A}_t; \sigma^{\text{pre}})} \qquad \tilde{S}_t = S_t + \text{Lap}\left(\frac{4\text{C}}{\epsilon}\right)$$

Mechanism (1): Graph Perturbation

$$\mathbf{A} = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}$$

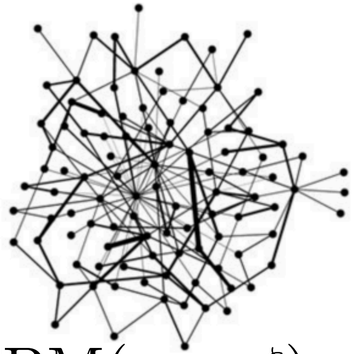

$\mathrm{CBM}(n, p, \zeta)$

Mechanism (1): Graph Perturbation

$$\mathbf{A} = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}$$

$\mathrm{CBM}(n, p, \zeta)$

Randomized Response Mechanism

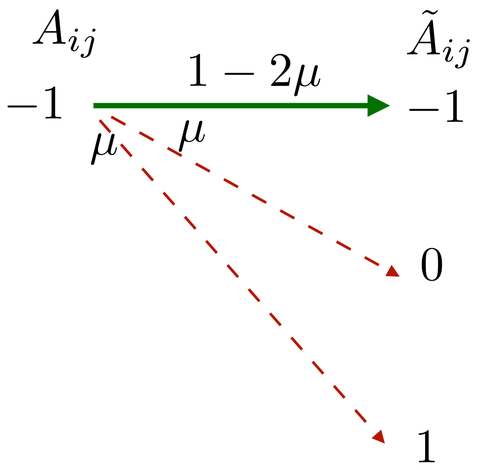

Mechanism (1): Graph Perturbation

$$\mathbf{A} = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}$$

$\mathrm{CBM}(n, p, \zeta)$

Randomized Response Mechanism

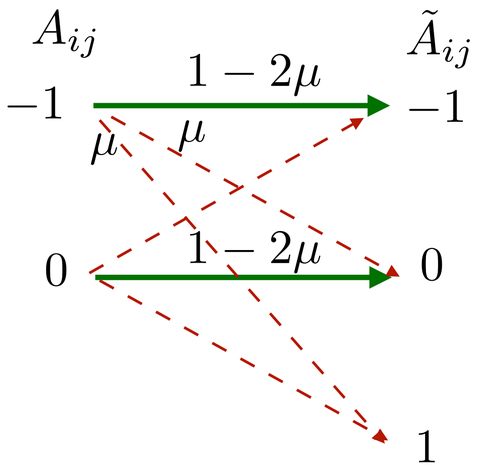

Mechanism (1): Graph Perturbation

$$\mathbf{A} = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}$$

$\mathrm{CBM}(n, p, \zeta)$

Randomized Response Mechanism

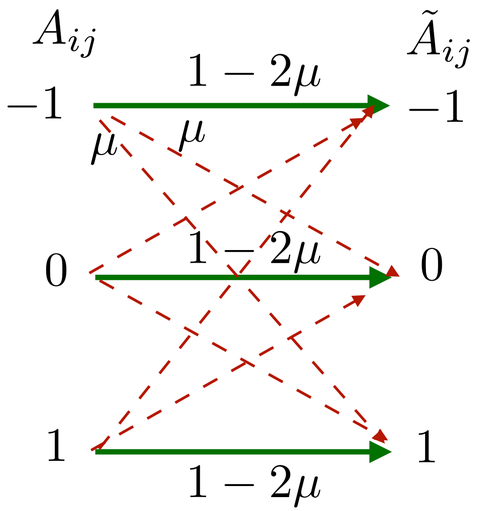

Mechanism (1): Graph Perturbation

$$\mathbf{A} = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}$$

$\mathrm{CBM}(n, p, \zeta)$

Randomized Response Mechanism

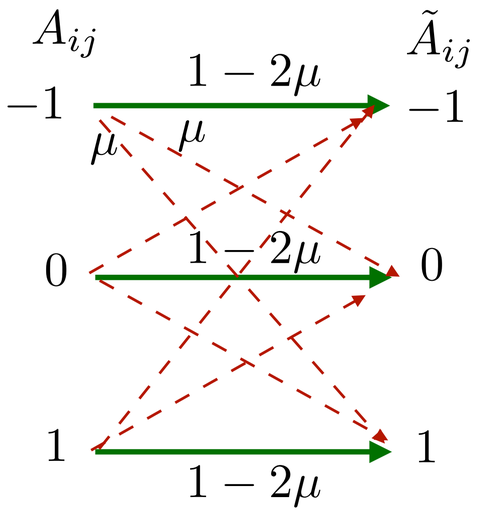

$$\mu = \frac{1}{e^{\epsilon} + 2}$$

small $\epsilon \leftrightarrow$ higher privacy (less leakage)

Mechanism (1): Graph Perturbation

$$\mathbf{A} = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix} \qquad \tilde{\mathbf{A}} = \begin{bmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{bmatrix}$$

$\mathrm{CBM}(n, p, \zeta)$

Randomized Response Mechanism

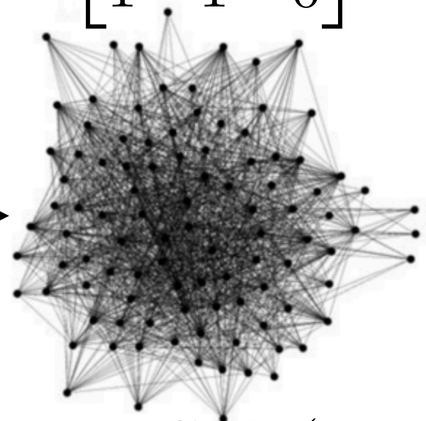

$\mathrm{CBM}(n, \tilde{p}, \tilde{\zeta})$

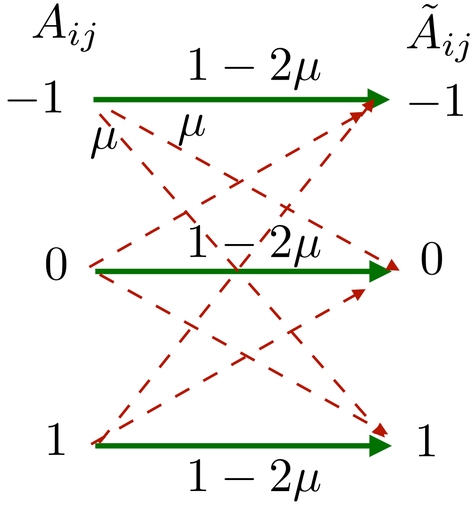

$$\mu = \frac{1}{e^{\epsilon} + 2}$$

small $\epsilon \leftrightarrow$ higher privacy (less leakage)

Mechanism (1): Graph Perturbation

$$\mathbf{A} = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix} \qquad \tilde{\mathbf{A}} = \begin{bmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{bmatrix}$$

$\mathrm{CBM}(n, p, \zeta)$

Randomized Response Mechanism

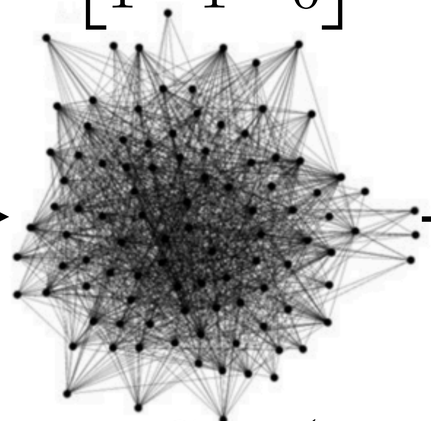

$\mathrm{CBM}(n, \tilde{p}, \tilde{\zeta})$

Post-processing

community label estimator

$\hat{\sigma}(\tilde{\mathbf{A}})$

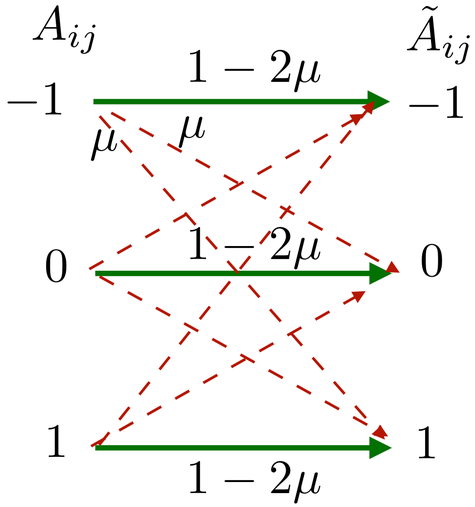

$$\mu = \frac{1}{e^{\epsilon} + 2}$$

small $\epsilon \leftrightarrow$ higher privacy (less leakage)

Detection Procedure

We apply the adaptive CUSUM on the **perturbed data**

- The detection statistic:

$$S_t = (S_{t-1})^+ + \log \frac{\Pr(\tilde{\mathbf{A}}_t; \hat{\sigma}_{t-1})}{\Pr(\tilde{\mathbf{A}}_t; \sigma^{\text{pre}})} \qquad t \geq 1, S_0 = 0$$

Detection Procedure

We apply the adaptive CUSUM on the **perturbed data**

- The detection statistic: *recursively accumulate the* ***log-likelihood ratio***

$$S_t = (S_{t-1})^+ + \log \frac{\Pr(\tilde{\mathbf{A}}_t; \hat{\sigma}_{t-1})}{\Pr(\tilde{\mathbf{A}}_t; \sigma^{\text{pre}})} \quad t \geq 1, S_0 = 0$$

Detection Procedure

We apply the adaptive CUSUM on the **perturbed data**

- The detection statistic: *recursively accumulate the* ***log-likelihood ratio***

$$S_t = (S_{t-1})^+ + \log \frac{\Pr(\tilde{\mathbf{A}}_t; \hat{\sigma}_{t-1})}{\Pr(\tilde{\mathbf{A}}_t; \sigma^{\text{pre}})} \quad t \geq 1, S_0 = 0$$

- Stopping time: $T = \inf\{t : S_t \geq b\}$

Detection Procedure

We apply the adaptive CUSUM on the **perturbed data**

- The detection statistic: *recursively accumulate the* ***log-likelihood ratio***

$$S_t = (S_{t-1})^+ + \log \frac{\Pr(\tilde{\mathbf{A}}_t; \hat{\sigma}_{t-1})}{\Pr(\tilde{\mathbf{A}}_t; \sigma^{\text{pre}})} \quad t \geq 1, S_0 = 0$$

- Stopping time: $T = \inf\{t : S_t \geq b\}$

- Differentially Private Exact Recovery Condition ($\mathrm{P}(\hat{\sigma} = \sigma^*) = 1 - o(1)$):

$$a(\sqrt{1-\zeta} - \sqrt{\zeta})^2 > \left(\frac{\sqrt{n}}{\sqrt{n}-1}\right) \times \left(\frac{e^\epsilon + 1}{e^\epsilon - 1}\right) \qquad \epsilon = \Omega(\log(n)) \quad p = a \times \frac{\log(n)}{n}$$

Detection Procedure

We apply the adaptive CUSUM on the **perturbed data**

- The detection statistic: *recursively accumulate the **log-likelihood ratio***

$$S_t = (S_{t-1})^+ + \log \frac{\Pr(\tilde{\mathbf{A}}_t; \hat{\sigma}_{t-1})}{\Pr(\tilde{\mathbf{A}}_t; \sigma^{\text{pre}})} \quad t \geq 1, S_0 = 0$$

- Stopping time: $T = \inf\{t : S_t \geq b\}$

- Differentially Private Exact Recovery Condition ($\mathrm{P}(\hat{\sigma} = \sigma^*) = 1 - o(1)$):

$$a(\sqrt{1-\zeta} - \sqrt{\zeta})^2 > \left(\frac{\sqrt{n}}{\sqrt{n}-1}\right) \times \left(\frac{e^\epsilon + 1}{e^\epsilon - 1}\right) \qquad \epsilon = \Omega(\log(n)) \quad p = a \times \frac{\log(n)}{n}$$

- Detection Delay: $\mathrm{WADD}(T) = \frac{\log \gamma}{\tilde{I}_0}(1 + o(1))$ $\qquad \tilde{I}_0 = \frac{1}{2}(\log \frac{1-\tilde{\zeta}}{\tilde{\zeta}})\tilde{p}(1-2\tilde{\zeta})(\binom{n}{2} - \sum_{i<j} \sigma_i^{\text{pre}} \sigma_j^{\text{pre}} \sigma_i^{\text{post}} \sigma_j^{\text{post}})$

Mechanism (2): Perturbation Stability-based Mechanism

$$\mathbf{A} = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}$$

Non-private estimator

$\hat{\sigma}(G)$

[Seif et al. ICML 2022]

Mechanism (2): Perturbation Stability-based Mechanism

$$\mathbf{A} = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}$$

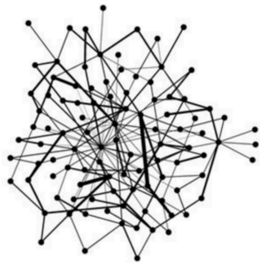

Non-private estimator

$\hat{\sigma}(G)$ → $d_{\hat{\sigma}}$

(compute stability of $\hat{\sigma}(G)$)

More stable → Less influence on the estimator → More Private

Stability of an estimator given *G*:

$$d_{\hat{\boldsymbol{\sigma}}}(G) = \{\min k : \exists G', dist(G, G') \leq k, \hat{\boldsymbol{\sigma}}(G) \neq \hat{\boldsymbol{\sigma}}(G')\}$$

[Seif et al. ICML 2022]

Mechanism (2): Perturbation Stability-based Mechanism

$$\mathbf{A} = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}$$

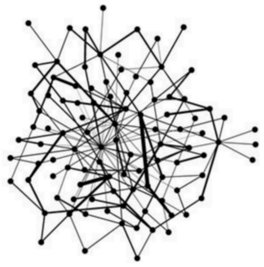

Non-private estimator

$\hat{\sigma}(G)$ → $d_{\hat{\sigma}}$

(compute stability of $\hat{\sigma}(G)$)

More stable → Less influence on the estimator → More Private

$\hat{\sigma}(G)$

Stability of an estimator given *G*:

$$d_{\hat{\boldsymbol{\sigma}}}(G) = \{\min k : \exists G', dist(G, G') \leq k, \hat{\boldsymbol{\sigma}}(G) \neq \hat{\boldsymbol{\sigma}}(G')\}$$

[Seif et al. ICML 2022]

Mechanism (2): Perturbation Stability-based Mechanism

$$\mathbf{A} = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}$$

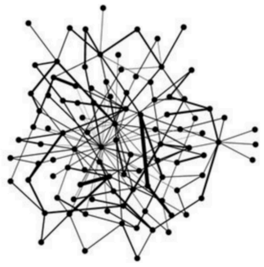

Non-private estimator

$\hat{\sigma}(G)$ → $d_{\hat{\sigma}}$

(compute stability of $\hat{\sigma}(G)$)

More stable → Less influence on the estimator → More Private

$\hat{\sigma}(G)$

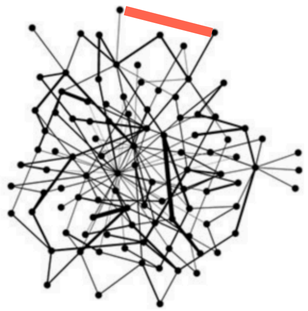

$\hat{\sigma}(G') = \hat{\sigma}(G)$

Stability of an estimator given *G*:

$$d_{\hat{\boldsymbol{\sigma}}}(G) = \{\min k : \exists G', dist(G, G') \leq k, \hat{\boldsymbol{\sigma}}(G) \neq \hat{\boldsymbol{\sigma}}(G')\}$$

[Seif et al. ICML 2022]

Mechanism (2): Perturbation Stability-based Mechanism

$$\mathbf{A} = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}$$

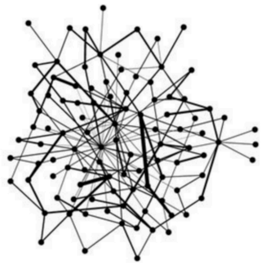

Non-private estimator

$\hat{\sigma}(G)$ → $d_{\hat{\sigma}}$

(compute stability of $\hat{\sigma}(G)$)

More stable → Less influence on the estimator → More Private

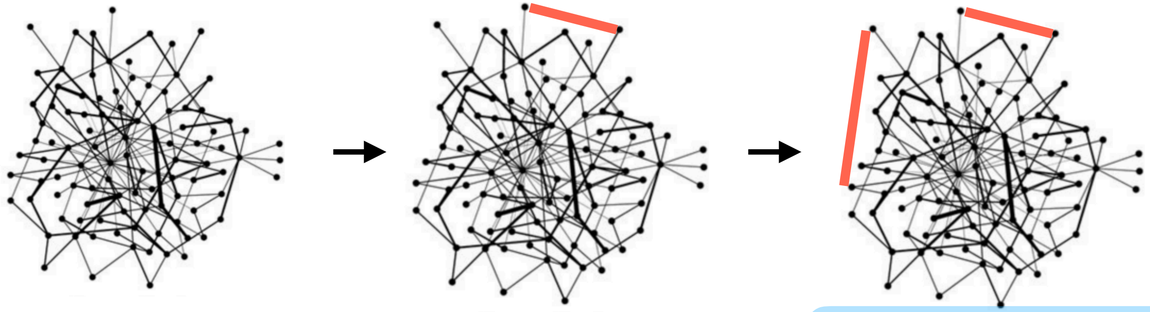

$\hat{\sigma}(G)$ $\quad$ $\hat{\sigma}(G') = \hat{\sigma}(G)$ $\quad$ $\hat{\sigma}(G'') \neq \hat{\sigma}(G)$

Stability of an estimator given *G*:

$$d_{\hat{\boldsymbol{\sigma}}}(G) = \{\min k : \exists G', dist(G, G') \leq k, \hat{\boldsymbol{\sigma}}(G) \neq \hat{\boldsymbol{\sigma}}(G')\}$$

[Seif et al. ICML 2022]

Mechanism (2): Perturbation Stability-based Mechanism

$$\mathbf{A} = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}$$

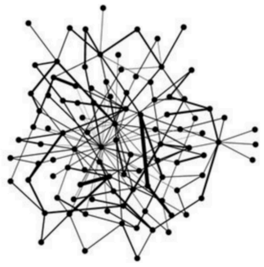

Non-private estimator

$\hat{\sigma}(G)$

(compute stability of $\hat{\sigma}(G)$)

$d_{\hat{\sigma}}$

Laplacian Mechanism

$\tilde{d}$

Private estimate of Stability

More stable → Less influence on the estimator → More Private

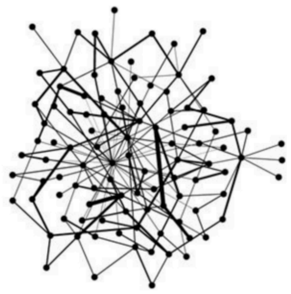

$\hat{\sigma}(G)$

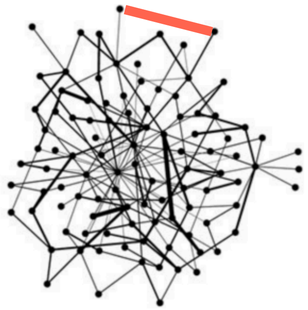

$\hat{\sigma}(G') = \hat{\sigma}(G)$

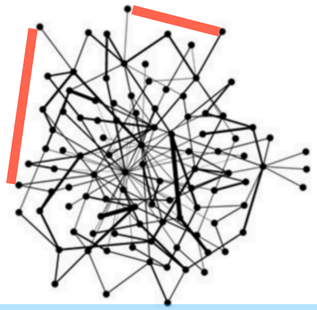

$\hat{\sigma}(G'') \neq \hat{\sigma}(G)$

Stability of an estimator given *G*:

$$d_{\hat{\boldsymbol{\sigma}}}(G) = \{\min k : \exists G', dist(G, G') \leq k, \hat{\boldsymbol{\sigma}}(G) \neq \hat{\boldsymbol{\sigma}}(G')\}$$

[Seif et al. ICML 2022]

Mechanism (2): Perturbation Stability-based Mechanism

$$\mathbf{A} = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}$$

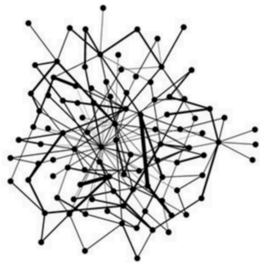

Non-private estimator

$\hat{\sigma}(G)$

(compute stability of $\hat{\sigma}(G)$)

$d_{\hat{\sigma}}$

Laplacian Mechanism

$\tilde{d}$

Private estimate of Stability

$\tilde{d} \overset{?}{>} \log(1/\delta)/\epsilon$

Yes → $\hat{\sigma}(G)$

No → $\perp$

More stable → Less influence on the estimator → More Private

$\hat{\sigma}(G)$

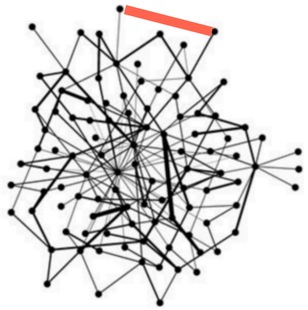

$\hat{\sigma}(G') = \hat{\sigma}(G)$

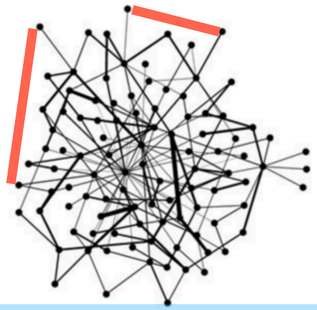

$\hat{\sigma}(G'') \neq \hat{\sigma}(G)$

Stability of an estimator given *G*:

$$d_{\hat{\boldsymbol{\sigma}}}(G) = \{\min k : \exists G', dist(G, G') \leq k, \hat{\boldsymbol{\sigma}}(G) \neq \hat{\boldsymbol{\sigma}}(G')\}$$

[Seif et al. ICML 2022]

Detection procedure

- The adaptive CUSUM statistic:

$$S_t = (S_{t-1})^+ + \log \frac{\Pr(\mathbf{A}_t; \hat{\sigma}_{t-1})}{\Pr(\mathbf{A}_t; \sigma^{\text{pre}})}$$

Add Laplacian noice to protect individual data

$$\tilde{S}_t = S_t + \text{Lap}\left(\frac{4\text{C}}{\epsilon}\right) \quad C = 2\log\frac{1-\zeta}{\zeta}$$

Detection procedure

- The adaptive CUSUM statistic:

$$S_t = (S_{t-1})^+ + \log \frac{\Pr(\mathbf{A}_t; \hat{\sigma}_{t-1})}{\Pr(\mathbf{A}_t; \sigma^{\text{pre}})}$$

Add Laplacian noice to protect individual data

$$\tilde{S}_t = S_t + \text{Lap}\left(\frac{4\text{C}}{\epsilon}\right) \quad C = 2\log\frac{1-\zeta}{\zeta}$$

- Stopping time: $T = \inf\{t : \tilde{S}_t \geq \tilde{b}\},$

Detection procedure

- The adaptive CUSUM statistic:

$$S_t = (S_{t-1})^+ + \log \frac{\Pr(\mathbf{A}_t; \hat{\sigma}_{t-1})}{\Pr(\mathbf{A}_t; \sigma^{\text{pre}})}$$

Add Laplacian noice to protect individual data

$$\tilde{S}_t = S_t + \text{Lap}\left(\frac{4\text{C}}{\epsilon}\right) \quad C = 2\log\frac{1-\zeta}{\zeta}$$

- Stopping time: $T = \inf\{t : \tilde{S}_t \geq \tilde{b}\},$

- Differentially Private Exact Recovery Condition ($\mathbf{P}(\hat{\sigma} = \sigma^*) = 1 - o(1)$):

$$a\left(\sqrt{1-\zeta} - \sqrt{\zeta}\right)^2 > 1 \qquad \epsilon = O(1)$$

Detection procedure

- The adaptive CUSUM statistic:

$$S_t = (S_{t-1})^+ + \log \frac{\Pr(\mathbf{A}_t; \hat{\sigma}_{t-1})}{\Pr(\mathbf{A}_t; \sigma^{\text{pre}})}$$

Add Laplacian noice to protect individual data

$$\tilde{S}_t = S_t + \text{Lap}\left(\frac{4\text{C}}{\epsilon}\right) \quad C = 2\log\frac{1-\zeta}{\zeta}$$

- Stopping time: $T = \inf\{t : \tilde{S}_t \geq \tilde{b}\},$

- Differentially Private Exact Recovery Condition ($\mathrm{P}(\hat{\sigma} = \sigma^*) = 1 - o(1)$):

$$a\left(\sqrt{1-\zeta} - \sqrt{\zeta}\right)^2 > 1 \qquad \epsilon = O(1)$$

- Detection Delay:

$$\text{WADD}(T) = \frac{\log\gamma}{I_0}(1 + o(1)) \qquad I_0 = \frac{1}{2}(\log\frac{1-\zeta}{\zeta})\tilde{p}(1-2\zeta)\left(\binom{n}{2} - \sum_{i<j}\sigma_i^{\text{pre}}\sigma_j^{\text{pre}}\sigma_i^{\text{post}}\sigma_j^{\text{post}}\right)$$

Outline

1. Motivation and Problem Setup: Censored Block Models
2. Private Online Detection Procedures
3. **Numerical Examples**
4. Open problems & Challenges…

Biological Networks

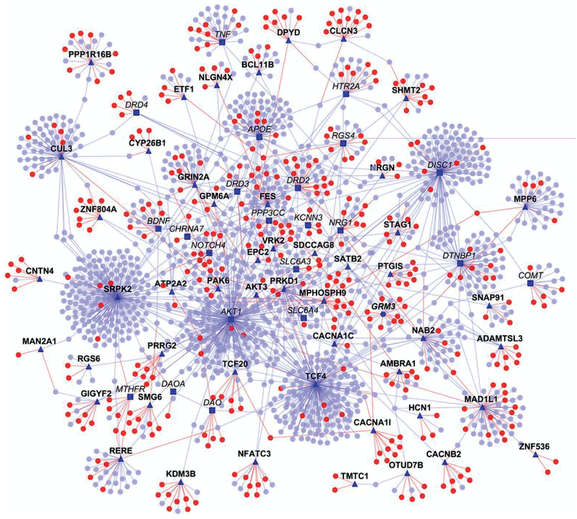

Contact-tracing Networks

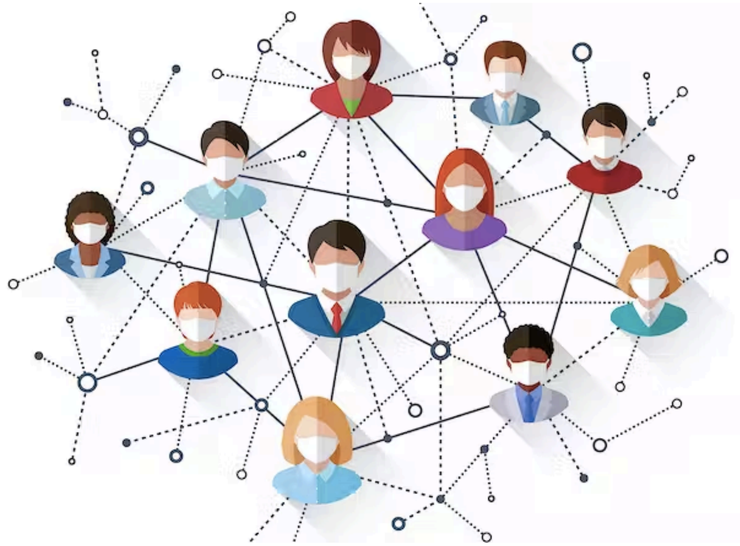

Collaboration Networks

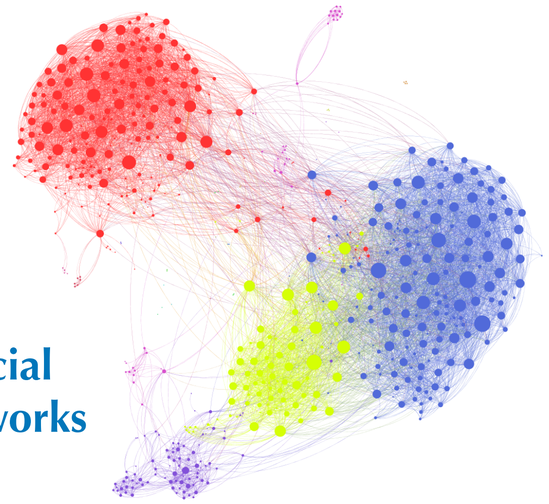

Social Networks

Real-world Datasets

- **Dataset:** Worldwide Agricultural Trade Dataset

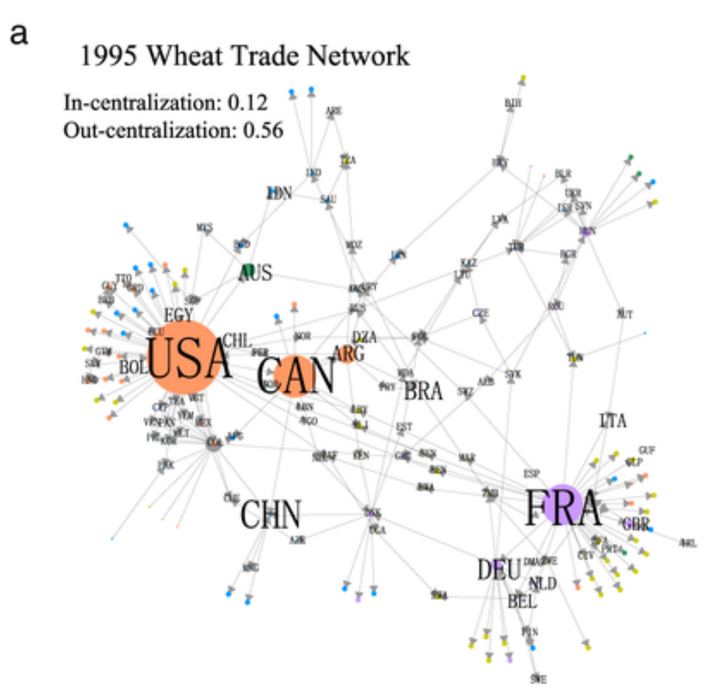

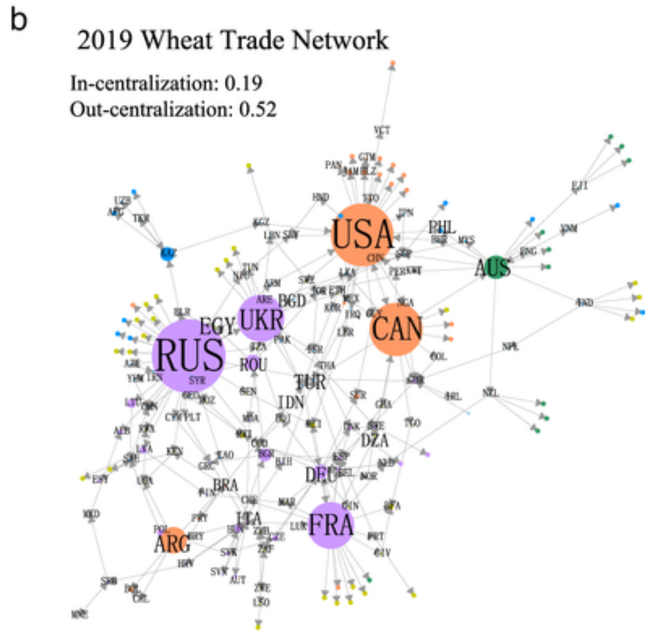

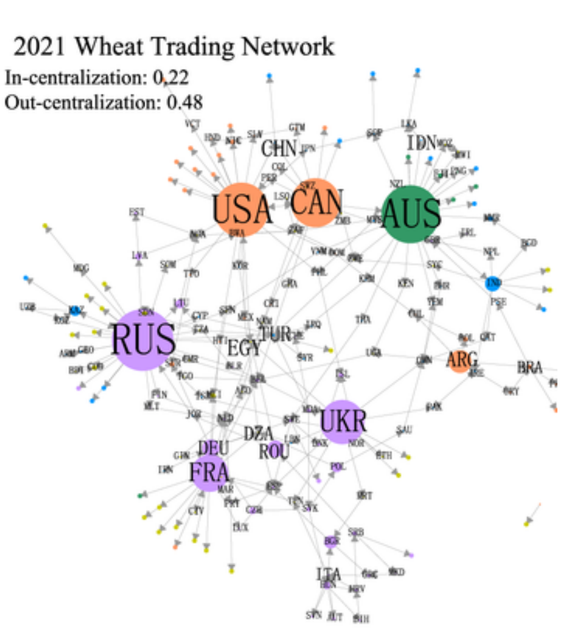

$n = 50$

$A_{ij} = +1$ (trade between two countries exceeds 4 products)

$A_{ij} = -1$ (trade between two countries is less than 4 products)

$A_{ij} = 0$ (no trade at all)

Real-world Datasets

Impact of Graph Perturbation

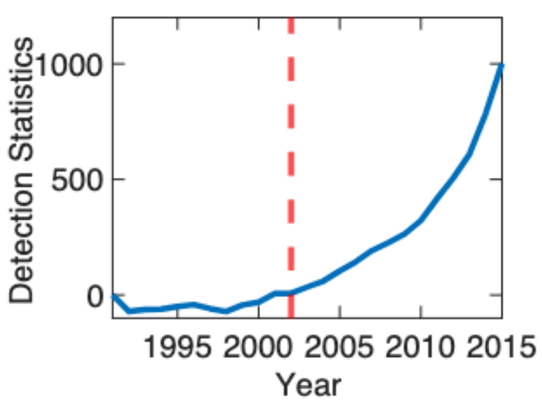

Raw Network ($\epsilon = +\infty$)

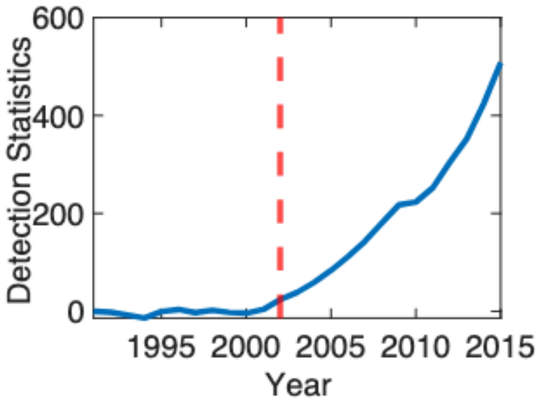

$\epsilon = 1.5$

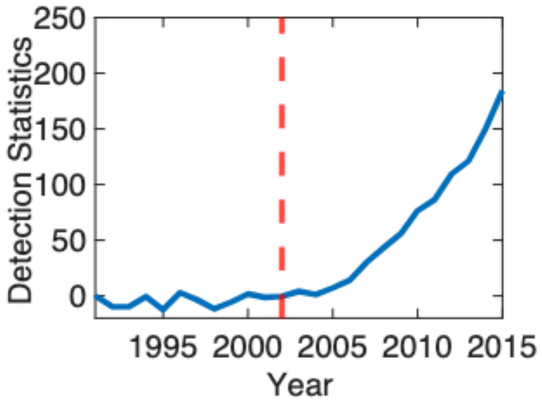

$\epsilon = 1$

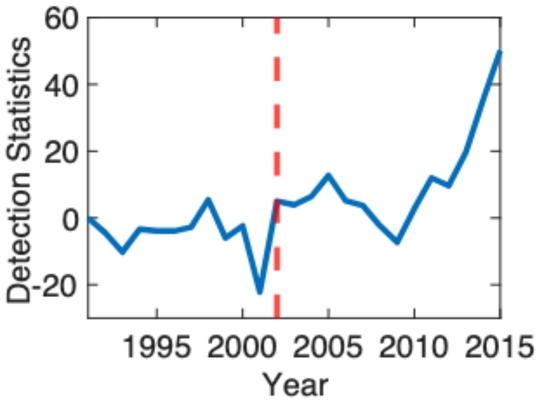

$\epsilon = 0.8$

We treat the year **2002 as the true change-point** location, as commonly found in previous studies.

- 2002 marked a significant turning point in global commodity prices, which likely influenced trade patterns of various countries

Conclusions & Open Problems

Conclusions & Open Problems

- Studied the problem of community change detection subject to *edge* differential privacy

Conclusions & Open Problems

- Studied the problem of community change detection subject to *edge* differential privacy
- Interesting interplay between **privacy**, **recovery** & **detection** performance

Conclusions & Open Problems

- Studied the problem of community change detection subject to *edge* differential privacy
- Interesting interplay between **privacy**, **recovery** & **detection** performance
- Several Open Problems

Conclusions & Open Problems

- Studied the problem of community change detection subject to *edge* differential privacy
- Interesting interplay between **privacy**, **recovery** & **detection** performance
- Several Open Problems
  - Studying the impact of privacy on other recovery notions (e.g., weak recovery) and the impact on detection efficiency.

Conclusions & Open Problems

- Studied the problem of community change detection subject to *edge* differential privacy
- Interesting interplay between **privacy**, **recovery** & **detection** performance
- Several Open Problems
  - Studying the impact of privacy on other recovery notions (e.g., weak recovery) and the impact on detection efficiency.
  - Studying other realistic random graph models (e.g., degree corrected SBMs, overlapping communities…)

Conclusions & Open Problems

- Studied the problem of community change detection subject to *edge* differential privacy
- Interesting interplay between **privacy**, **recovery** & **detection** performance
- Several Open Problems
  - Studying the impact of privacy on other recovery notions (e.g., weak recovery) and the impact on detection efficiency.
  - Studying other realistic random graph models (e.g., degree corrected SBMs, overlapping communities…)

Seif, Mohamed, Liyan Xie, Andrea J. Goldsmith, and H. Vincent Poor. "Private Online Community Detection for Censored Block Models." 2024. https://arxiv.org/pdf/2405.05724

Conclusions & Open Problems

- Studied the problem of community change detection subject to *edge* differential privacy
- Interesting interplay between **privacy**, **recovery** & **detection** performance
- Several Open Problems
  - Studying the impact of privacy on other recovery notions (e.g., weak recovery) and the impact on detection efficiency.
  - Studying other realistic random graph models (e.g., degree corrected SBMs, overlapping communities…)

Seif, Mohamed, Liyan Xie, Andrea J. Goldsmith, and H. Vincent Poor. "Private Online Community Detection for Censored Block Models." 2024. https://arxiv.org/pdf/2405.05724

Thank you!

Email: liyanxie@umn.edu