Fighting Opinion Control in Social Networks via Link Recommendation

Victor Amelkin (University of Pennsylvania) Ambuj K. Singh (University of California, Santa Barbara)

Introduction

- **Goal:** Strategically recommend links to recover weighted average user opinion from exogenous node-level attacks.
- Large directed strongly connected social network of n users
- W – interpersonal influence adjacency matrix ($W\mathbb{1} = \mathbb{1}$)
- $x \in [0,1]^n$, $(\widetilde{x})$ – user opinions before (after) the attack
- $\pi \in \mathbb{R}^n$ – network nodes' eigenvector centralities
- $\langle \pi, x \rangle$ – **(weighted) average opinion**

Problem Statement by Example

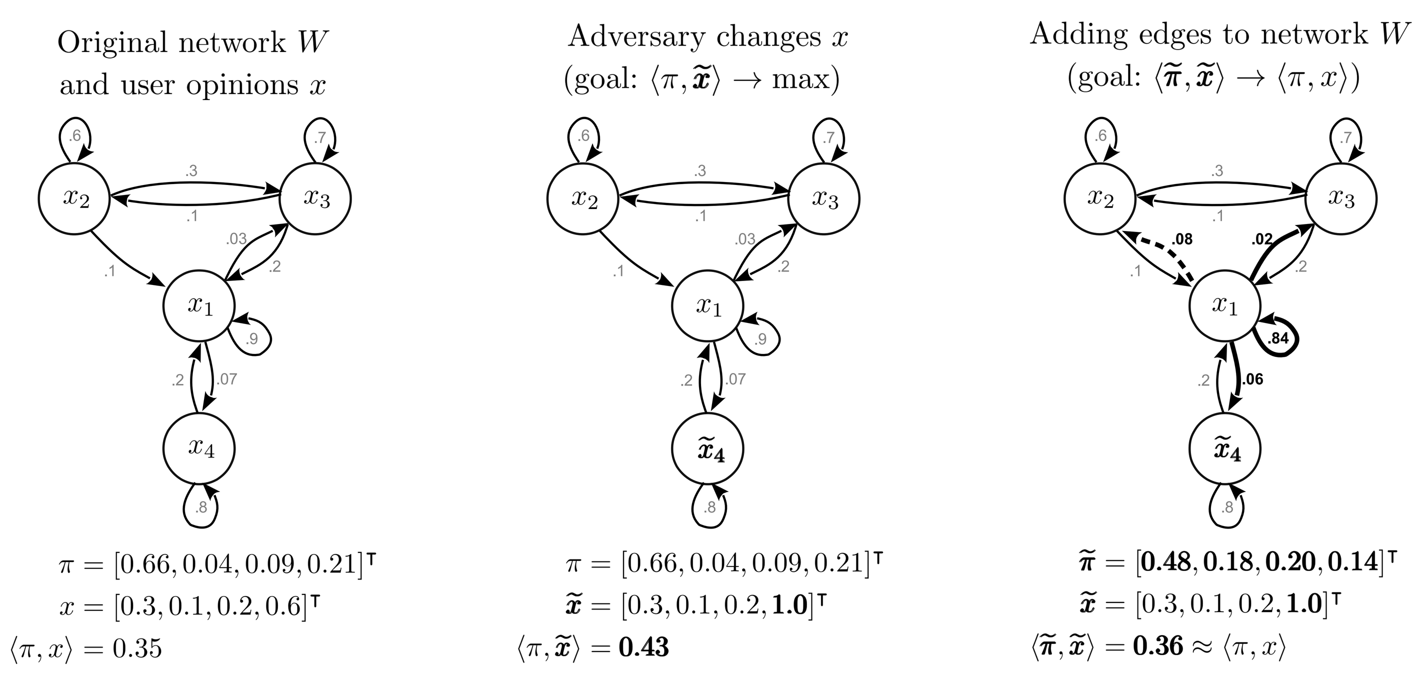

Problem and Its Hardness

- **Adversary's Goal:** Maximize $\langle \pi, \widetilde{x} \rangle$ via altering $x \to \widetilde{x}$
- **Our Goal:** Return $\langle \widetilde{\pi}, \widetilde{x} \rangle$ back to $\langle \pi, x \rangle$ via altering $\pi \to \widetilde{\pi}$ through edge addition. Single-edge (r, c) perturbation:

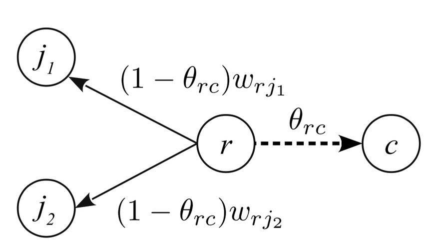

- **(NP-hard) Problem:**

$$\boxed{\mathrm{DIVER}(W, k, x, \widetilde{x}) = \arg\min_{\widetilde{W}} |\langle \widetilde{\pi}(\widetilde{W}), \widetilde{x} \rangle - \langle \pi, x \rangle|,}$$

where the perturbed $\widetilde{W}$ differs from W by k new edges, we cannot choose weight θ_{ij} of an added edge (i, j).

General Solution for DIVER

$$\mathrm{DIVER}(W, k, x, \widetilde{x}) = \arg\min_{\widetilde{W}} |\langle \widetilde{\pi}, \widetilde{x} \rangle - \langle \pi, x \rangle|$$

- **Method:** reduce $\langle \widetilde{\pi}, \widetilde{x} \rangle$ through iterative edge addition until it gets close enough to $\langle \pi, x \rangle$

$$\langle \pi, \widetilde{x} \rangle - \langle \widetilde{\pi}, \widetilde{x} \rangle \to \max$$

- **Central Question:** How does $\langle \widetilde{\pi}, \widetilde{x} \rangle$ change when a single edge (r, c) with weight θ_{rc} is added to network W?

Network Perturbation Analysis

- Adding a single edge to the network:

$$\widetilde{W} = W - \theta_{rc}\,\mathrm{diag}(e_r) W + \theta_{rc} e_r e_c^\intercal$$

Theorem 1. *Under single-edge perturbation of W with edge (r, c) having weight θ_{rc}, the eigenvector centrality changes as follows:*

$$\widetilde{\pi}_j = \pi_j \left[1 - \frac{\theta_{rc}(m_{cj} \cdot (1 - \delta\{j, c\}) - m_{rj} + 1)}{m_{rr} + \theta_{rc}(m_{cr} - m_{rr} + 1)}\right],$$

where m_{ij} is the mean first passage time (MFPT) from state i to state j of Markov chain W, and $\delta\{\cdot,\cdot\}$ is Kronecker's delta. In particular,

$$\widetilde{\pi}_r = \frac{1}{m_{rr} + \theta_{rc}(m_{cr} - m_{rr} + 1)},$$

$$\widetilde{\pi}_c = 1 + \theta_{rc} \cdot \frac{m_{rc} - 1}{m_{rr} + \theta_{rc}(m_{cr} - m_{rr} + 1)}.$$

Theorem 2. *Under single-edge perturbation of W with edge (r, c) having weight θ_{rc}, the weighted average opinion changes as follows:*

$$f_\pi(r, c) = \langle \pi, \widetilde{x} \rangle - \langle \widetilde{\pi}, \widetilde{x} \rangle$$

$$= \theta_{rc} \frac{\sum_{j=1}^{n} \pi_j (m_{cj} \cdot (1 - \delta\{j, c\}) - m_{rj} + 1)\widetilde{x}_j}{m_{rr} + \theta_{rc}(m_{rc} - m_{rr} + 1)}.$$

How to solve DIVER in very large networks?

- Approach to DIVER: iteratively adding edges (r, c) with top values $f_\pi(r, c)$ until satisfied with the value of $\langle \widetilde{\pi}, \widetilde{x} \rangle$.
- **Issue 1:** There are $\mathcal{O}(n^2)$ candidate edges in a sparse network.
- **Issue 2:** How to efficiently compute $f_\pi(r, c)$?
  - Evaluation of a single $f_\pi(r, c)$ involves summation over $\mathcal{O}(n)$ terms.
  - Direct computation of MFPTs m_{ij} would cost at least $\mathcal{O}(n^3)$.

Efficient candidate edge selection

- Focus on $\mathcal{O}(n)$ candidate edges, outgoing from $n_{src} \ll n$ nodes.
- Most good candidate edges emanate from a small number of nodes.
- In hierarchical networks, these edge sources are top-centrality nodes.

Efficient computation of $f_\pi(r, c)$

- In hierarchical networks, f_π is largely determined by a small (n_{src}) number of top-centrality nodes.
- We can estimate MFPTs via finite-time random walks; all the MFPTs to and from n_{src} top-centrality nodes converge in $\mathcal{O}(n)$ time in practice.

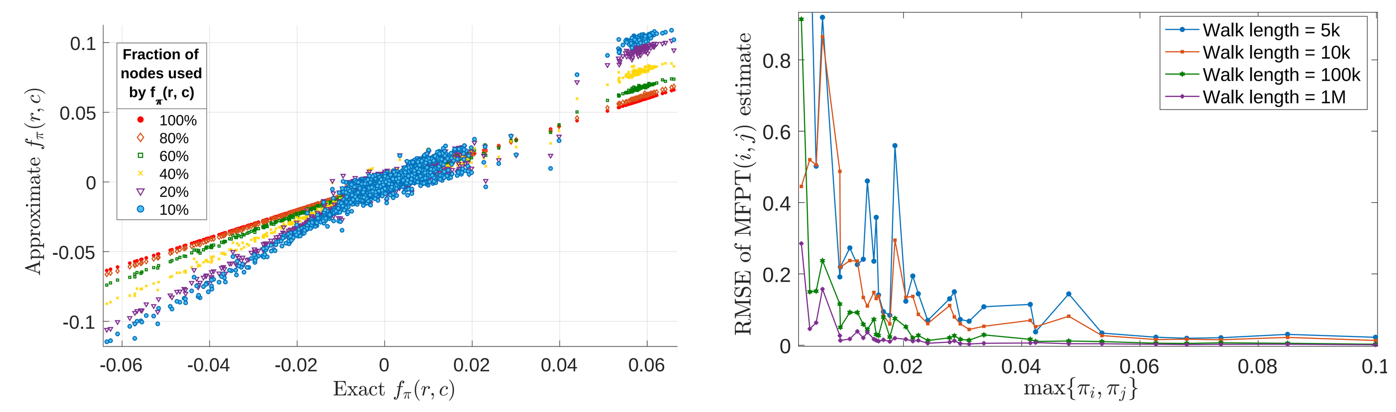

- **Outcome:** $\mathcal{O}(n)$-time heuristic for DIVER for hierarchical (scale-free-like) networks.

Solving DIVER

- Driving $\langle \widetilde{\pi}, \widetilde{x} \rangle - \langle \pi, x \rangle$ to 0 by adding new edges:

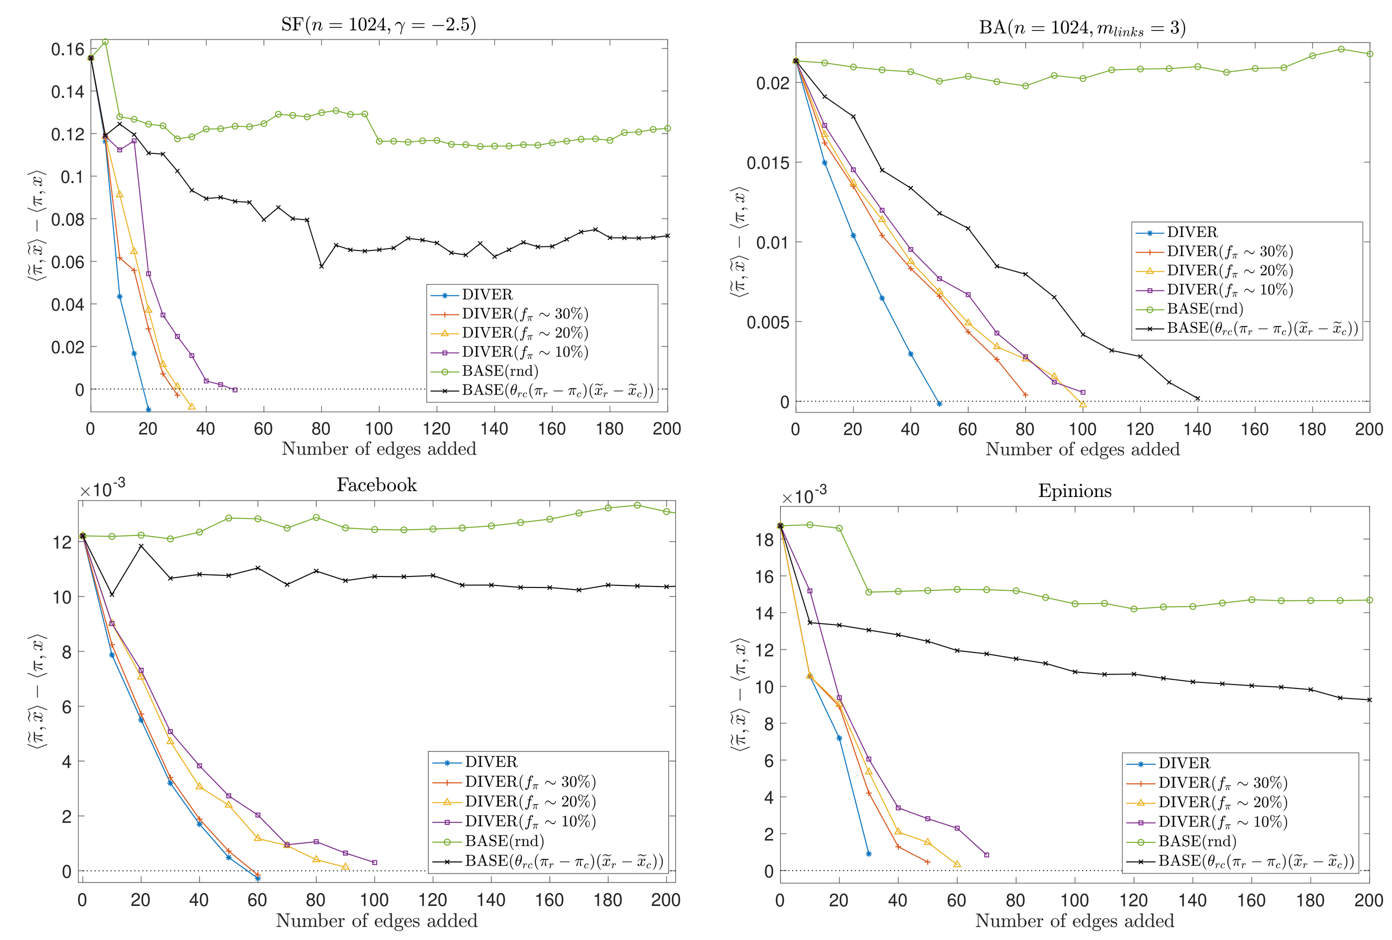

- DIVER: optimistic DIVER estimates f_π using *all* (not $\mathcal{O}(1)$) top-centrality nodes;
- DIVER($f_\pi \sim X\%$): estimates f_π using only fraction X of top-centrality nodes;
- BASE(rnd): worst-case baseline, choosing edges uniformly randomly;
- BASE($\theta_{rc}(\pi_r - \pi_c)(\widetilde{x}_r - \widetilde{x}_c)$) – smarter baseline that relies only on absolute centrality π_i, as well as edge weights θ_{rc}, and user opinions $\widetilde{x}_i$.

References

- Amelkin V., Singh A.K., "Fighting Opinion Control in Social Networks via Link Recommendation", in Proc. of ACM SIGKDD, 2019

The work was supported by the U.S. Army Research Laboratory and the U.S. Army Research Office under grant W911NF-15-1-0577, by the National Science Foundation under grant IIS-1817046, and the Rockefeller Foundation under grant 2017 PRE 301.

https://victoramelkin.com/pub/diver/
vctr@seas.upenn.edu; ambuj@cs.ucsb.edu