

POTION : Optimizing Graph Structure for Targeted Diffusion

Sixie Yu * ^{†1}, Leo Torres^{‡2}, Scott Alfeld^{§3}, Tina Eliassi-Rad ^{¶2}, and Yevgeniy Vorobeychik ^{||1}

¹Washington University in St. Louis

²Northeastern University

³Amherst College

Abstract

The problem of diffusion control on networks has been extensively studied, with applications ranging from marketing to controlling infectious disease. However, in many applications, such as cybersecurity, an attacker may want to attack a *targeted* subgraph of a network, while limiting the impact on the rest of the network in order to remain undetected. We present a model POTION in which the principal aim is to optimize graph structure to achieve such targeted attacks. We propose an algorithm POTION-ALG for solving the model at scale, using a gradient-based approach that leverages Rayleigh quotients and pseudospectrum theory. In addition, we present a condition for certifying that a targeted subgraph is immune to such attacks. Finally, we demonstrate the effectiveness of our approach through experiments on real and synthetic networks.

1 Introduction

Many diverse phenomena that propagate through a network, such as epidemic spread, cascading failures, and chemical reactions, can be modeled by network diffusion models [3, 5, 23, 41, 38]. The problem of controlling diffusion has, as a result, received much attention in the literature, with primary focus on two mechanisms for control: the choice of initial nodes to start the spread [14, 7, 43], and the modification of network structure [16, 33, 42, 34]. To date, most work on diffusion control (either promotion or inhibition) has considered diffusion over the entire network. However, in many problems, the focus is instead on diffusion that is *targeted* to a particular subgraph of the network. For example, in cybersecurity, diffusion commonly represents malware spread, but malware attacks are often targeted

at particular subsets of critical devices [12], which should be accounted for when modeling attacking behavior. Congestion cascades of ground traffic or flight networks are other examples, where the goal of resilience may be to ensure that cascades concentrate on a subset of high-capacity nodes that can handle them, limiting the impact on the rest of the network [10, 9]. In another domain, medical treatments for certain diseases such as cancer may leverage a molecular signaling network, with the goal of targeting just the pathogenic portion of it, while limiting the deleterious effects on the rest [39].

We study the problem of targeted diffusion in which an attacker¹ can modify the graph structure $G = (\mathcal{V}, \mathcal{E})$ to achieve two goals: 1) maximize the diffusion spread to a target subgraph G_S , and 2) minimize the impact on the remaining graph $G \setminus G_S$. We capture the first goal by maximizing a utility function that incorporates spectral information of the adjacency matrix of G , specifically its largest (in magnitude) eigenvalue, eigenvector centrality, and the normalized cut of the target subgraph. The second goal is achieved by limiting the modifications made outside of the target subgraph. We present a scalable algorithmic framework for solving this problem. Our framework leverages a combination of gradient ascent with the use of Rayleigh quotients and pseudospectrum theory, which yields differentiable approximations of our objective and allows us to avoid projection steps that would otherwise be costly and imprecise. Moreover, we derive a condition that enables us to certify if a network is robust against a broad class of targeted diffusion attacks. Finally, we demonstrate the effectiveness of our approach through extensive experiments.

In summary, our contributions are:

1. We propose POTION (oPtimizing graph structures fOr Targeted diffusION): a model for targeted diffusion attack by optimizing graph structures.

*The first two authors contributed equally to the paper.

[†]sixie.yu@wustl.edu.

[‡]leo@leotrs.com

[§]salfeld@amherst.edu

[¶]t.eliassirad@northeastern.edu

^{||}yvorobeychik@wustl.edu

¹The attacker is the agent who initiates diffusion.

2. We present POTION-ALG : an efficient algorithm to optimize POTION by leveraging Rayleigh quotient and pseudospectrum theory.
3. We describe a condition for certifying that a targeted subgraph is immune to such attacks.
4. We demonstrate the effectiveness and efficiency of POTION and POTION-ALG on synthetic and real-world networks; and against baseline and competing methods.²

2 Related Work

Various dynamical processes can be modeled as diffusion dynamics on networks, including the spread of infectious diseases [3, 5], cascading failures in infrastructure networks [23, 41], and information spread (e.g., rumors, fake news) on social networks [19, 20]. One line of research assesses the impact of cascading failures. Yang et al. [41] simulated cascading failures to quantify the vulnerability of the power grid in North America. Fleurquin et al. [10] studied the impact of flight delays as a cascading failure diffusing through the network. Motter and Lai [23] investigated the cascading failures on a network due to the malfunction of a single node. Another line of research concerns diffusion control, for example, selecting a set of nodes such that if the diffusion originated from them, it reaches as many nodes as possible [14, 7, 43]; or modifying network structures to increase or limit some diffusion [33, 28]. However, these lines of research do not differentiate between targeted and non-targeted nodes. Ho et al. [13] studied targeted diffusion controlled by changing nodal status. We focus on the problem where an attacker manipulates underlying network structures in order to achieve targeted diffusion.

Another relevant research thread is network design, which is the problem of modifying network structure to induce certain desirable outcomes. Some prior work [33, 42, 31] considered the containment of spreading dynamics by adding or removing nodes or edges from the network, while others [37, 29, 16, 6, 34] considered limiting the spread of infectious disease by minimizing the largest eigenvalue of the network. Kempe et al. [15] studied modifying network structure to induce certain outcomes from a game-theoretic perspective, but they did not consider diffusion dynamics. Others have studied the problem of manipulating node centrality measures (e.g., eigenvector or PageRank centrality) [2, 1] or node similarity measures (e.g., Katz similarity) [44] through

edge perturbation. All of these prior efforts focus on the impact either at the network level or at the node-level properties, while our focus is on the impact of diffusion dynamics on a targeted subgraph of the network.

3 POTION : Proposed Model

We present a model for targeted diffusion through graph structure optimization. We refer to the agent who initiates diffusion as *the attacker*. We use cybersecurity as a running example. Here the attacker initiates the diffusion (e.g., the spread of malware) on a network of computers. We define the impact of the diffusion as the number of infected nodes (e.g., compromised with malware). The attacker has two objectives: 1) she wishes to maximize the impact of the diffusion on a targeted set of nodes (e.g., computing nodes with access to critical assets), and 2) to limit the impact on non-targeted nodes to ensure stealth [12].

Let $G = (\mathcal{V}, \mathcal{E})$ be a connected, weighted or unweighted, undirected graph with no self-loops. Let $n = |\mathcal{V}|$ be the number of nodes in G and $\mathbf{A}$ be its adjacency matrix. Throughout this paper, the eigenvalues of $\mathbf{A}$ are ranked in descending order $\lambda_1(\mathbf{A}) \geq \dots \geq \lambda_n(\mathbf{A})$. Suppose the attacker targets a subgraph G_S where $\mathcal{S} \subseteq \mathcal{V}$ is the node set of G_S . Let $\mathcal{S}' = \mathcal{V} \setminus \mathcal{S}$, and its induced subgraph $G_{S'}$. Throughout the paper we assume G_S is connected, and denote its adjacency matrix by $\mathbf{A}_S$. To achieve her objectives, the attacker modifies the structure of G . The modified graph and targeted subgraph are represented by $\tilde{G}$ and $\tilde{G}_S$, respectively. Formally, the attacker's action is to add a perturbation $\Delta \in \mathbb{R}^{n \times n}$ to $\mathbf{A}$, which results in the perturbed adjacency matrix $\tilde{\mathbf{A}} = \mathbf{A} + \Delta$. The adjacency matrix of $\tilde{G}_S$ is denoted by $\tilde{\mathbf{A}}_S$.

3.1 Diffusion Dynamics The status of a node is modeled by the well-known SIS (Susceptible-Infected-Susceptible) diffusion dynamics, where it alternates between “infected” and “susceptible”.³ Due to the malware spread by infected neighbors, a susceptible node becomes infected with probability β . An infected node becomes susceptible again (e.g., malware is removed) with probability δ . Following Chakrabarti et al. [5], this process is modeled by a nonlinear dynamical system. Let π_i be the probability of node i becoming infected (e.g., compromised with malware) in the steady state of this dynamical system, with $\boldsymbol{\pi}$ the vector of these probabilities. A key result in [5] is that when $\lambda_1(\mathbf{A}) < \delta/\beta$ the system converges to the steady state $\boldsymbol{\pi} = \mathbf{0}$, which implies that the diffusion process quickly dies out.

²The code to replicate the experimental results is at <https://github.com/marsplus/POTION>.

³Due to brevity, a discussion on generalization of our approach to other diffusion dynamics is at <https://arxiv.org/abs/2008.05589>.

However, when $\lambda_1(\mathbf{A}) \geq \delta/\beta$ the system converges to another steady state $\boldsymbol{\pi} \neq \mathbf{0}$. We leverage this connection between graph structure, dynamical model of epidemic spread, and the epidemic threshold, in constructing our *threat model*, as discussed next.

3.2 Threat Model Maximizing the Impact on G_S : To maximize the impact of diffusion on G_S , the attacker has two goals: 1) ensure that epidemics starting in G_S spread rather than die out, and 2) ensure that epidemics starting outside G_S are likely to reach it. We capture the first goal by maximizing the largest (in modulus) eigenvalue of G_S , $\lambda_1(\tilde{\mathbf{A}}_S)$, which corresponds to the epidemic threshold of the targeted subgraph.⁴ The second goal is captured by maximizing the *normalized cut* of G_S , $\phi(\mathcal{S})$, where $\mathcal{S}$ is the set of nodes in G_S and $\mathcal{S}'$ are the nodes in the remaining graph. The normalized cut is formally defined as follows:

$$(3.1) \quad \phi(\mathcal{S}) = \text{cut}(\mathcal{S}, \mathcal{S}') \left(\frac{1}{\text{vol}(\mathcal{S})} + \frac{1}{\text{vol}(\mathcal{S}')} \right),$$

where $\text{cut}(\mathcal{S}, \mathcal{S}')$ is the sum of the weights on the edges across $\mathcal{S}$ and $\mathcal{S}'$ (unit weights for unweighted graphs), and $\text{vol}(\mathcal{S})$ (resp. $\text{vol}(\mathcal{S}')$) is the sum of degrees of the nodes in $\mathcal{S}$ (resp. $\mathcal{S}'$). The formal rationale for using the normalized cut is based on Meila et al. [21], which showed that increasing $\phi(\mathcal{S})$ increases the probability that a random walker transitions from $\mathcal{S}'$ to $\mathcal{S}$, if we assume that G_S is smaller than $G_{S'}$.

Limiting the Impact on $G_{S'}$: Another important objective of the attacker is to limit the impact on $G_{S'}$, the non-targeted part of the graph. We capture this goal in two different ways. First, by limiting the likelihood of the epidemic spreading to $G_{S'}$, which we define as minimizing the impact $I(G_{S'}) = \sum_{i \in \mathcal{S}'} \pi_i$. Second, by limiting the impact on the spectrum of $\mathbf{A}$.

We now demonstrate that minimizing $I(G_{S'})$ is approximately equivalent to minimizing the eigenvector centrality of $\mathcal{S}'$. Let $\mathbf{P}^t$ be the global configuration of the graph at time step t , where P_i^t is the probability that node i is infected (e.g., compromised with malware). Following Miegheem et al. [22], ignoring higher-order terms and taking the time step to be infinitesimally small, the dynamics of P_i^t is modeled as the following:

$$(3.2) \quad \frac{dP_i^t}{dt} = \sum_{j \in \mathcal{V}} \beta \tilde{\mathbf{A}}_{ij} P_j^t - \delta P_i^t.$$

Here, we can think of the two terms on the right side as two competing forces. The first term is the force

contributed by the infected neighbors of node i (which increases P_i^t), while the second term is the force due to i 's self recovery (which decreases P_i^t). Rewriting in matrix notation yields:

$$(3.3) \quad \frac{d\mathbf{P}^t}{dt} = [\beta \tilde{\mathbf{A}} - \delta \mathbf{I}] \mathbf{P}^t,$$

which gives a linear approximation to the non-linear dynamical system proposed in [5]. The steady state $\boldsymbol{\pi}$ must satisfy $[\beta \tilde{\mathbf{A}} - \delta \mathbf{I}] \boldsymbol{\pi} = \mathbf{0}$, which is equivalent to $\tilde{\mathbf{A}} \boldsymbol{\pi} = (\delta/\beta) \boldsymbol{\pi}$. Suppose $\lambda_1(\tilde{\mathbf{A}}) = \delta/\beta$, and $\boldsymbol{\pi}$ is the corresponding eigenvector. Let $\tilde{\mathbf{v}}_1$ be the unit eigenvector associated with $\lambda_1(\tilde{\mathbf{A}})$. Let $\sigma(\mathcal{S}) = \sum_{j \in \mathcal{S}} \tilde{v}_1[j]$ be the eigenvector centrality of $\mathcal{S}$. Noting that $\boldsymbol{\pi}$ may differ from $\tilde{\mathbf{v}}_1$ by up to a multiplicative constant c , the impact on $G_{S'}$ can be approximated as:

$$(3.4) \quad I(G_{S'}) = \sum_{j \in \mathcal{S}'} \pi_j \approx c \sum_{j \in \mathcal{S}'} \tilde{v}_1[j] = c(1 - \sigma(\mathcal{S})),$$

where the last equality is because $\mathcal{S}$ and $\mathcal{S}'$ are disjoint and $\tilde{\mathbf{v}}_1$ is a unit vector. Thus, minimizing the impact on $G_{S'}$ is approximately equivalent to maximizing the eigenvector centrality of $\mathcal{S}$.

Recall that to have an epidemic spread, one needs $\lambda_1(\tilde{\mathbf{A}}) \geq \delta/\beta$. Here, we assumed $\lambda_1(\tilde{\mathbf{A}}) = \delta/\beta$. In Section 6, we demonstrate that our analysis yields an approach that is effective even when this assumption fails to hold (i.e., when $\lambda_1(\tilde{\mathbf{A}}) > \delta/\beta$).

Now we focus on limiting the impact on the spectrum of $\mathbf{A}$.⁵ Let $\epsilon > 0$ be the attacker's budget. Formally, this notion is captured through the following constraints:

$$(3.5) \quad |\lambda_i(\tilde{\mathbf{A}}) - \lambda_i(\mathbf{A})| \leq \epsilon, \quad i = 1, \dots, n.$$

In summary, the principal aims to (i) maximize the impact on G_S through maximizing $\lambda_1(\tilde{\mathbf{A}}_S)$ while (ii) limiting the impact on $G_{S'}$ by maximizing the eigenvector centrality $\sigma(\mathcal{S})$, and satisfying Eq. (3.5). Formally, the principal aims to solve the following optimization problem:

$$(3.6) \quad \begin{aligned} \max_{\tilde{\mathbf{A}}} \quad & \alpha_1 \lambda_1(\tilde{\mathbf{A}}_S) + \alpha_2 \sigma(\mathcal{S}) + \alpha_3 \phi(\mathcal{S}) \\ \text{s.t.} \quad & \tilde{\mathbf{A}} \in \mathcal{P} = \left\{ \tilde{\mathbf{A}} \mid |\lambda_i(\tilde{\mathbf{A}}) - \lambda_i(\mathbf{A})| \leq \epsilon, i = 1, \dots, n, \right. \\ & \left. \tilde{\mathbf{A}} = \tilde{\mathbf{A}}^\top, \tilde{\mathbf{A}}_{ii} = 0, \forall i = 1, \dots, n \right\}, \end{aligned}$$

where the relative importance of the terms is balanced by the nonnegative constants $\alpha_1, \alpha_2, \alpha_3$, and the restrictions $\tilde{\mathbf{A}} = \tilde{\mathbf{A}}^\top$ and $\tilde{\mathbf{A}}_{ii} = 0, \forall i = 1, \dots, n$ ensure that $\tilde{\mathbf{A}}$ is a valid adjacency matrix.

⁴If G_S is not connected, we may replace $\lambda_1(\tilde{\mathbf{A}}_S)$ by the largest eigenvalue of the largest connected component of G_S .

⁵In cybersecurity, there are natural interpretations of an attack's stealth. For further details, see the extended version at <https://arxiv.org/abs/2008.05589>.

4 POTION-ALG : Proposed Algorithm

To solve the optimization problem in Eq. (3.6), a natural approach would be to use a form of projected gradient ascent. There are, however, two major hurdles to this basic approach: 1) the objective function involves terms that do not have an explicit functional representation in the decision variables, and 2) the projection step is quite expensive, as it involves projecting into a spectral norm ball, which entails an expensive SVD operation [17]. We address these challenges in Algorithm 1, which is our gradient-based solution to the attacker's optimization problem as described in Eq. (3.6).

Algorithm 1 POTION-ALG

```

1: Input:  $\mathbf{A}, \epsilon, \{\eta_i\}_{i=1} \triangleright \{\eta_i\}_{i=1}$  is a schedule of step sizes
2: Initialize:  $i = 1, \tilde{\mathbf{A}}_1 = \mathbf{A}, B_1 = 0 \triangleright B_i$ : the amount of
   budget used just before step  $i$ 
3: while True do
4:   Set  $\Delta_i$  to the gradient of  $\alpha_1 \lambda_1(\tilde{\mathbf{A}}_S) + \alpha_2 \sigma(\mathcal{S}) + \alpha_3 \phi(\mathcal{S})$ 
   w.r.t. to  $\tilde{\mathbf{A}}_i$ 
5:   Set the diagonal entries of  $\Delta_i$  to zeros
6:   if  $\|\Delta_i\| = 0$  then  $\triangleright$  a local optimum is found
7:     return  $\tilde{\mathbf{A}}_i$ 
8:   end if
9:   if  $B_i + \|\eta_i \Delta_i\|_2 \leq \epsilon$  then  $\triangleright$  one-step look ahead
10:     $\tilde{\mathbf{A}}_{i+1} = \tilde{\mathbf{A}}_i + \eta_i \Delta_i, B_{i+1} = B_i + \|\eta_i \Delta_i\|_2, i = i+1$ 
11:   else
12:     return  $\tilde{\mathbf{A}}_i$ 
13:   end if
14: end while

```

A key step of Algorithm 1 is line 4, where we compute the gradient of the attacker's utility function with respect to $\tilde{\mathbf{A}}$. This gradient involves terms that do not have an explicit functional form in terms of the decision variable, and we deal with each of these in turn.

First, consider the gradient of the normalized cut $\phi(\mathcal{S})$ w.r.t. $\tilde{\mathbf{A}}$. Let $\mathbf{x}_S$ be the characteristic vector of $\mathcal{S}$, that is $x_S[i] = 1$ iff $i \in \mathcal{S}$. Let $\tilde{\mathbf{D}}$ be the diagonal degree matrix $\tilde{D}_{ii} = \sum_j \tilde{A}_{ij}$, and let $\tilde{\mathbf{L}} = \tilde{\mathbf{A}} - \tilde{\mathbf{D}}$ be the Laplacian matrix. Using $\tilde{\mathbf{D}}$ and $\tilde{\mathbf{L}}$ to express $\text{vol}(\mathcal{S})$ and $\text{cut}(\mathcal{S}, \mathcal{S}')$, respectively, we have:

$$(4.7) \quad \phi(\mathcal{S}) = \mathbf{x}_S^\top \tilde{\mathbf{L}} \mathbf{x}_S \left(\frac{1}{\mathbf{x}_S^\top \tilde{\mathbf{D}} \mathbf{x}_S} + \frac{1}{\mathbf{x}_{S'}^\top \tilde{\mathbf{D}} \mathbf{x}_{S'}} \right).$$

Clearly, Eq. (4.7) is a differentiable function of $\tilde{\mathbf{A}}$. Computing its gradient $\nabla_{\tilde{\mathbf{A}}} \phi(\mathcal{S})$ can then be handled by automatic differentiation tools such as PyTorch [26].

Next, we compute the gradient of $\lambda_1(\tilde{\mathbf{A}}_S)$ w.r.t. $\tilde{\mathbf{A}}$. A standard way to compute $\lambda_1(\tilde{\mathbf{A}}_S)$ is by using SVD. However, this is both prohibitively expensive ($O(n^3)$),

and does not provide us with the necessary gradient information. Instead, we use the power method [11] to compute $\lambda_1(\tilde{\mathbf{A}}_S)$. Let $\mathbf{v}_S$ be the eigenvector associated with the largest eigenvalue $\lambda_1(\tilde{\mathbf{A}}_S)$. Using Rayleigh quotients [35], we can compute $\lambda_1(\tilde{\mathbf{A}}_S)$ as follows:

$$(4.8a) \quad \mathbf{v}_S = \arg \max_{\|\mathbf{x}\|_2=1} \mathbf{x}^\top \tilde{\mathbf{A}}_S \mathbf{x}$$

$$(4.8b) \quad \lambda_1(\tilde{\mathbf{A}}_S) = \mathbf{v}_S^\top \tilde{\mathbf{A}}_S \mathbf{v}_S.$$

Thus, when $\mathbf{v}_S$ is known, the computation of $\lambda_1(\tilde{\mathbf{A}}_S)$ reduces to matrix multiplications. In addition, $\tilde{\mathbf{A}}_S$ is usually sparse, so we can leverage sparse matrix multiplication to speed up the computation.

The remaining challenge is that $\mathbf{v}_S$ is an optimal solution of an optimization problem, and we need an explicit derivative of it. Fortunately, our problem has a special structure that we exploit to obtain an approximation of the derivative of $\mathbf{v}_S$. From our experiments we find that $\tilde{G}_S$ is nearly always connected. This means that the largest eigenvalue of $\tilde{\mathbf{A}}_S$ is simple. In addition, due to the Perron–Frobenius theorem, the absolute value of the largest eigenvalue is strictly greater than the absolute values of others, i.e., $|\lambda_1(\tilde{\mathbf{A}}_S)| > |\lambda_k(\tilde{\mathbf{A}}_S)|$ for all $k \neq 1$. Under these conditions, we can use the power method to estimate $\mathbf{v}_S$ by repeating the formula: $\tilde{\mathbf{v}}_S^{(t+1)} = \tilde{\mathbf{A}}_S \tilde{\mathbf{v}}_S^{(t)} / \|\tilde{\mathbf{A}}_S \tilde{\mathbf{v}}_S^{(t)}\|_2$. The ℓ_2 -norm distance between $\tilde{\mathbf{v}}_S^k$ and $\mathbf{v}_S$ decreases in a rate $O(\rho^k)$ [11], where $\rho < 1$. In our experiments we found $k = 50$ is enough to give a high-quality estimation for a graph with 986 nodes. Intuitively, we are using a sequence of differentiable operations to approximate the argmax operation. Therefore the computation of $\nabla_{\tilde{\mathbf{A}}} \lambda_1(\tilde{\mathbf{A}}_S)$ can be handled by PyTorch.

We use the same machinery to compute $\nabla_{\tilde{\mathbf{A}}} \sigma(\mathcal{S})$. First, we write $\sigma(\mathcal{S})$ in matrix notation:

$$(4.9) \quad \sigma(\mathcal{S}) = \mathbf{v}^\top \mathbf{x}_S,$$

where $\mathbf{v}$ is the unit eigenvector associated with $\lambda_1(\tilde{\mathbf{A}})$. Then we apply the power method to compute $\mathbf{v}$. Finally, $\sigma(\mathcal{S})$ is just a linear function of $\mathbf{v}$. All of these operations are differentiable, and the computation of $\nabla_{\tilde{\mathbf{A}}} \sigma(\mathcal{S})$ is handled by PyTorch.

We next address the challenge imposed by the constraints (3.5), which can result in a computationally challenging projection step which can also significantly harm solution quality. We address this challenge as follows. Given a real symmetric matrix $\mathbf{X}$, let $\|\mathbf{X}\|_2$ denote its spectral norm. To satisfy Eq. (3.5), we use the following result

from pseudospectrum theory (see [36], Theorem 2.2):

$$(4.10) \quad |\lambda_i(\tilde{\mathbf{A}}) - \lambda_i(\mathbf{A})| \leq \epsilon, i = 1, \dots, n \iff \|\tilde{\mathbf{A}} - \mathbf{A}\|_2 \leq \epsilon$$

Since $\Delta = \tilde{\mathbf{A}} - \mathbf{A}$ is real and symmetric, we have $\|\Delta\|_2 = \max\{|\lambda_1(\Delta)|, |\lambda_n(\Delta)|\}$ and $-\lambda_n(\Delta) = \lambda_1(-\Delta)$, which leads to:

$$(4.11) \quad \tilde{\mathbf{A}} \text{ satisfies Eq. (3.5)} \iff \max\{|\lambda_1(\Delta)|, |\lambda_1(-\Delta)|\} \leq \epsilon.$$

This equivalence allows the attacker to check whether she is within budget simply by evaluating $\max\{|\lambda_1(\Delta)|, |\lambda_1(-\Delta)|\}$, i.e., computing the largest eigenvalue of a real symmetric matrix, which can be computed efficiently using, e.g., the power method [11].

Our algorithm leverages this connection as follows. Line 9 in Algorithm 1 is a one step look-ahead, which ensures that the perturbation Δ_i is only added to $\tilde{\mathbf{A}}_i$ when there is enough budget. Recall from Section 3.2 that $\|\Delta_i\|_2 = \max\{|\lambda_1(\Delta_i)|, |\lambda_1(-\Delta_i)|\}$. Thus this step requires us to compute $\lambda_1(\Delta_i)$ and $\lambda_1(-\Delta_i)$, using again the power method. Line 10 tracks the amount of budget used so far. We now show that the output of Algorithm 1 always returns a feasible solution. Suppose Algorithm 1 terminates after $k > 1$ iterations. This means $B_k + \|\eta_k \Delta_k\|_2 > \epsilon$ and $B_k \leq \epsilon$. In other words $B_k = \sum_{i=1}^{k-1} \|\eta_i \Delta_i\|_2 \leq \epsilon$. Note that the total amount of perturbation added to $\mathbf{A}$ is $\Delta = \sum_{i=1}^{k-1} \eta_i \Delta_i$. The triangle inequality implies $\|\Delta\|_2 \leq \epsilon$.

For each iteration of Algorithm 1, the most computationally expensive components are the power method and matrix multiplication. Let m be the number of nonzeros in $\tilde{\mathbf{A}}_i$; if the graph is unweighted then m is the number of edges at this iteration. By leveraging the sparseness exhibited in $\tilde{\mathbf{A}}_i$, the power method runs in $O(m)$ and the matrix multiplications cost $O(mn)$. Thus, the time complexity of each iteration is $O(mn)$, which significantly improves the $O(n^3)$ time complexity of SVD that would otherwise be needed.

Recall that our model for targeted diffusion is applicable to both weighted and unweighted graphs. For weighted graphs, the attacker modifies the weights on existing edges. For unweighted graphs, the attacker adds new edges or deletes existing edges from the graph. The main difference between the two settings is that the latter needs a rounding heuristic to convert a matrix with fractional entries to a binary adjacency matrix. We discuss this heuristic below.

After running Algorithm 1, we obtain a perturbed matrix $\tilde{\mathbf{A}}$ with fractional entries. For unweighted graphs, a rounding heuristic is needed to convert $\tilde{\mathbf{A}}$ to a valid adjacency matrix. Let $\mathcal{D} = \{(i, j) | \tilde{A}_{i,j} \neq A_{i,j}\}$ be the set of candidate edges that will be added or deleted from G . For each edge $(i, j) \in \mathcal{D}$ define the score $s_{(i,j)} = |\tilde{A}_{i,j} - A_{i,j}|$. Intuitively, $s_{(i,j)}$ indicates the impact that adding or deleting the edge has on the principal's utility. Next, we iteratively modify G , by adding or deleting edges in $\mathcal{D}$, starting with the one with the largest $s_{(i,j)}$. The modification process stops when the budget is exhausted, which results in the desired binary adjacency matrix. For weighted graphs, let $C = \max_{i,j} A_{i,j}$ and normalize each entry by C , that is $A_{i,j}/C$. We run Algorithm 1 on the normalized adjacency matrix, which results in $\tilde{\mathbf{A}}$. The desired adjacency matrix is obtained by multiplying each $\tilde{A}_{i,j}$ by C , $C\tilde{A}_{i,j}$. If integer weights are desired (e.g., the number of packages transmitted between two computers), a final rounding step is applied. Our experimental results show that the rounding heuristic is effective in practice.

5 Certified Robustness

This section addresses the following question: what are the limits on the attacker's ability to successfully accomplish her attack? More precisely, we now seek to identify necessary conditions on the attack budget ϵ so the attack succeeds; conversely, we can view a given graph to be *certified* to be robust to attacks that use a smaller budget than the one required.

Let $\text{TargetDiff}(\mathcal{S}, G, \epsilon)$ be an instance of the targeted diffusion problem with target subset $\mathcal{S}$, underlying graph G and budget ϵ . The attacker is successful on an instance $\text{TargetDiff}(\mathcal{S}, G, \epsilon)$ if she is able to modify G into $\tilde{G}$ within budget ϵ such that $I(\tilde{G}_{\mathcal{S}}) > I(G_{\mathcal{S}})$. We now derive a necessary condition for successful attacks, in the form of a lower bound on ϵ .

To derive the necessary condition on ϵ , we use our experimental observation that in successful attacks the degrees of nodes in the targeted subgraph $G_{\mathcal{S}}$ always increase. This is intuitive: a denser subgraph $G_{\mathcal{S}}$ will tend to increase the propensity of the diffusion (e.g., of malware) to spread within it, which is one of our explicit objectives. Let d_i (resp. $\tilde{d}_i$) be the degree of node i before (resp. after) graph modification. We assume if an attack is successful, the degrees of nodes in $G_{\mathcal{S}}$ are increased, i.e., $\tilde{d}_i \geq d_i$ for $i \in \mathcal{S}$.

Now, observe that computing the exact value of $I(G_{\mathcal{S}})$ is intractable, since the exact computation of π_i is prohibitive (see, e.g., [22], Section IV.B). Miegheem et al. [22] proposed a simple yet effective estimator

for π_i to be $1 - \delta/(\beta d_i)$. The estimator works in the regime $\delta/\beta \leq d_{\min}$, where $d_{\min}$ is the minimum degree of G . Consequently, an estimator for $I(G_S)$ is $\hat{I}(G_S) = \sum_{i \in S} 1 - \delta/(\beta d_i)$. We focus on the setting where the estimation error is bounded by a small number, i.e., $|\hat{I}(G_S) - I(G_S)| \leq \tau$. Note that τ can be estimated from historical diffusion data. The formal statement of the necessary condition is in Theorem 5.1.⁶

THEOREM 5.1. *Given an instance $\text{TargetDiff}(\mathcal{S}, G, \epsilon)$, $I(G_S)$ is estimated by $\hat{I}(G_S) = \sum_{i \in S} 1 - \delta/(\beta d_i)$. Suppose we have an upper bound $|\hat{I}(G_S) - I(G_S)| \leq \tau$, the degrees of nodes in $\mathcal{S}$ are increased, i.e., $\tilde{d}_i \geq d_i$ for $i \in \mathcal{S}$, and $\delta/\beta \leq d_{\min}$. In order to have $I(\tilde{G}_S) - I(G_S) > 2\tau$, the budget ϵ must satisfy:*

$$(5.12) \quad \epsilon \geq \sqrt{\frac{|\mathcal{S}|}{n}} \left(\frac{\sum_{i \in \mathcal{S}} d_i^2}{|\mathcal{S}|} - \frac{(\sum_{i \in \mathcal{S}} d_i)^2}{|S|^2} \right)^{1/2}.$$

The quantity inside the square root is always nonnegative due to Jensen’s inequality. The lower bound involves only structural properties of the graph (node degrees and the size of $\mathcal{S}$) and thus can be easily computed given an arbitrary graph. As mentioned above, we can view this lower bound as a robustness certificate, or guarantee for the given graph. It guarantees, in particular, that when the budget is below the lower bound, the total probability of “infection” (e.g., malware infection) in G_S cannot be increased by more than 2τ . In the special case of perfect estimation ($\tau = 0$), it implies impossibility of increasing the susceptibility of G_S to targeted diffusion.

The proof of Theorem 5.1 does not depend on the specific objective function proposed in this paper. Consequently, the certificate is not specific to our particular objective function. Further, the lower bound is independent of the values of δ and β , as long as $\delta/\beta \leq d_{\min}$.

We briefly discuss the settings where the robustness guarantee is most applicable. First, the estimation for the infected ratio on G_S is accurate, i.e., $|\hat{I}(G_S) - I(G_S)| \leq \tau$ and τ is small. According to Miegheem et al. [22], this usually happens on graphs with small degree variation. Another setting is where the degrees of nodes in G_S increase as a result of the attack, which is both natural and empirically founded, as we mentioned earlier. We provide experimental results on synthetic networks to verify the robustness guarantee in Section 6.

6 Experiments and Discussion

This section presents experimental results on three real-world datasets: an email network, an airport network, and a brain network.⁷

For each network we run POTION with hyperparameters $\alpha_1 = \alpha_2 = \alpha_3 = 1/3$, which encodes that the attacker’s objectives are equally important. To study how the attacker’s effectiveness changes with respect to her budget, we set $\epsilon = \gamma \lambda_1(\mathbf{A})$ and vary γ from 10% to 50%. A single initially infected node is selected uniformly at random.

Recall we use G and $\tilde{G}$ to denote the original and the modified graphs, respectively. We simulate the spreading dynamics 2000 times on both G and $\tilde{G}$. For unweighted graphs the recovery rate δ and transmission rate β are set to 0.24, 0.06, resp.; for weighted graphs we set $\delta = 0.24$ and $\beta = 0.2$. The spreading dynamics converges exponentially fast to the steady state: empirically, we found 30 time steps to be enough to reach the steady state in most cases. When the simulation finishes, we extract the number of nodes that are “infected”. We use I_{original} and I_{modified} to represent the fractions of infected nodes on G and $\tilde{G}$, resp.

We use two other algorithms as baselines for comparison, which we call **deg** and **gel**. The two algorithms work by alternating between modifying G_S and modifying $G_{S'}$ until the budget is spent. When modifying G_S , **deg** chooses the edge (i, j) with the maximum value of $d_i + d_j$. Algorithm **gel** is based on [32], and chooses the edge (i, j) with maximum *eigenscore*, defined as $u[i]v[j]$, where u, v are the left and right principal eigenvectors of $\mathbf{A}$, respectively. This edge is chosen from among those edges that are absent (if the graph is unweighted) or present (if the graph is weighted). When modifying $G_{S'}$, these baselines choose an existing edge of to remove (if unweighted) or decrease its weight (if weighted) with the highest value of $d_i + d_j$ or $u[i]v[j]$, respectively.

Unweighted Graphs: We consider (the largest connected component of) an email network [18] that has 986 nodes. An edge (i, j) indicates that there were email exchanges between nodes i and j . This data set contains ground-truth labels to indicate which community a node belongs to. We pick a community with 15 nodes as $\mathcal{S}$; the results for communities with other sizes are similar.

The overall effectiveness of our approach is shown in Figure 2, top. The difference $I_{\text{modified}} - I_{\text{original}}$ of the

⁶Due to brevity the proof is in Appendix C at <https://arxiv.org/abs/2008.05589>.

⁷Due to brevity, additional results on real and synthetic networks are at <https://arxiv.org/abs/2008.05589>.

impact on the modified and original graphs is shown for G_S (red line) and $G_{S'}$ (purple line), respectively. As γ gets larger, the impact on G_S increases, while the impact on $G_{S'}$ is under control, which demonstrates that the proposed approach is highly effective at both increasing the impact of diffusion on the targeted subgraph, and at the same time preventing the impact on the remaining graph.

Weighted Graphs: We consider an airport network and a brain network. The airport network [24] was collected from the website of Bureau of Transportation Statistics of the U.S., where the nodes represent all of the 1572 airports in the U.S. and the weights on edges encode the number of passengers traveled between two airports in 2010. We scaled the weights on the airport network to $[0, 1]$. The targeted set S was chosen by first sampling a node i uniformly at random, and then setting S to be i and all its neighbors. We report experimental results for an S with 60 nodes. The brain network [8] consists of 638 nodes where each node corresponds to a region in human brain. An edge between nodes i and j indicates that the two regions have co-activated on some tasks. The weight on the edge quantifies the strength of the co-activation estimated by the Jaccard index. The weights on edges lie in $[0, 1]$. The 638 regions are categorized into four areas: default mode, visual, fronto-parietal, and central. Each area is responsible for some functionality of human. We select 100 nodes from the central area as the targeted set S . The results for the airport (resp. brain) network are at the center (resp. right) column of Figure 2. The overall trend is similar to that of the email network.

Comparison against Baselines: The comparisons against the baselines are shown in Figure 2, middle and bottom rows. The middle row shows the infectious ratios within the targeted subgraphs. It is clear that our algorithm is more effective at increasing the infectious ratios than the baselines. The bottom row shows the infectious ratios within the non-targeted subgraphs. The magnitudes of the differences are negligible, although in some cases our algorithm is significantly better than the baselines (e.g., on airport network when $\gamma = 0.4$).

Verify the Certified Robustness: We run experiments on synthetic networks to verify the certified robustness; the synthetic networks include Barabási-Albert (BA) [4], Watts-Strogatz [40], and Block Two-level Erdős-Rényi (BTER) networks [30]. We use the same experimental setup as described above. Fig. 1 shows the difference of infectious ratios on the modified and original graphs (within targeted subgraphs), as a function of the attacker's budget ϵ . The vertical dashed lines are the




Figure 1: Certified robustness results. Dashed lines mark the lower bounds from Eq. (5.12). Solid lines represent infectious ratios within targeted subgraphs.




Figure 2: POTION effectively achieves targeted diffusion (**top**) in G_S (red line) without affecting $G_{S'}$ (purple); higher is better. Comparison against deg and gel baselines in G_S (**middle**; higher is better) and $G_{S'}$ (**bottom**; lower is better).

lower bounds on the budget computed using Eq. (5.12). Note that when the budget is less than the lower bound, the differences are close to zero, which means that the network is robust against targeted diffusion.

Running Time: The running time of Algorithm 1 on the three real-world networks is showed in Figure 3. Each point in the figure is the average running time over 10 trials. Intuitively, as the budget γ increases the attacker needs to search a larger space, therefore the running time increases. The numbers of nodes and edges of the three networks are in Table 1.




Figure 3: Running time on the real-world networks.

	Email	Airport	Brain
#nodes	986	1572	638
#edges	16064	17214	18625

Table 1: Statistics of the real-world networks.

7 Conclusion

Diffusion control on network has attracted much attention, however, most studies focus on diffusion over the entire network. We address the problem of *targeted* diffusion attack on networks. We present a combination of modeling and algorithmic advances to systematically address this problem. On the modeling side, we present a novel model called POTION that optimizes graph structure to affect such targeted diffusion attacks, which preserves structural properties of the graph. On the algorithmic side, we design an efficient algorithm named POTION-ALG by leveraging Rayleigh quotients and pseudospectrum theory, which is scalable to real-world graphs. We also derive a condition to certify whether a network is robust against a broad class of targeted diffusion. Our experiments on both synthetic and real-world networks show that the model is highly effective in implementing the targeted diffusion attack.

Acknowledgement

SY and YV were partially supported by the National Science Foundation (grants IIS-1903207 and IIS-1910392) and Army Research Office (grants W911NF1810208 and W911NF1910241). LT and TER were supported in part by the National Science Foundation (IIS-1741197) and by the Combat Capabilities Development Command Army Research Laboratory (under Cooperative Agreement Number W911NF-13-2-0045). The authors would like to thank the anonymous reviewers and Chloe Wohlgemuth for their helpful comments.

References

- [1] Victor Amelkin and Ambuj K. Singh. Fighting opinion control in social networks via link recommendation. In *KDD*, pages 677–685. ACM, 2019.
- [2] Konstantin Avrachenkov and Nelly Litvak. The effect of new links on google pagerank. *Stochastic Models*, 22(2):319–331, 2006.
- [3] Norman TJ Bailey et al. *The mathematical theory of infectious diseases and its applications*. Charles Griffin & Company Ltd, 1975.
- [4] Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.
- [5] Deepayan Chakrabarti, Yang Wang, Chenxi Wang, Jure Leskovec, and Christos Faloutsos. Epidemic thresholds in real networks. *ACM Trans. Inf. Syst. Secur.*, 10(4): 1:1–1:26, 2008.
- [6] Chen Chen, Hanghang Tong, B. Aditya Prakash, Charalampos E. Tsourakakis, Tina Eliassi-Rad, Christos Faloutsos, and Duen Horng Chau. Node immunization on large graphs: Theory and algorithms. *TKDE*, 28(1):113–126, 2016.
- [7] Wei Chen, Yajun Wang, and Siyu Yang. Efficient influence maximization in social networks. In *KDD*, pages 199–208. ACM, 2009.
- [8] Nicolas A Crossley, Andrea Mechelli, Petra E Vértés, Toby T Winton-Brown, Ameera X Patel, Cedric E Ginestet, Philip McGuire, and Edward T Bullmore. Cognitive relevance of the community structure of the human brain functional coactivation network. *PNAS*, 110(28):11583–11588, 2013.
- [9] Ernesto Estrada. ‘Hubs-repelling’ Laplacian and related diffusion on graphs/networks. *Linear Algebra Appl.*, 2020.
- [10] Pablo Fleurquin, José J Ramasco, and Victor M Eguiluz. Systemic delay propagation in the US airport network. *Scientific Reports*, 3:1159, 2013.
- [11] Gene Golub and Charles Van Loan. *Matrix computations*. Johns Hopkins Studies in Mathematical Sciences, 1996.
- [12] Nika Haghtalab, Aron Laszka, Ariel D. Procaccia, Yevgeniy Vorobeychik, and Xenofon Koutsoukos. Monitoring stealthy diffusions. *Knowledge and Information Systems*, 2017.
- [13] Christopher Ho, Mykel J. Kochenderfer, Vineet Mehta, and Rajmonda S. Caceres. Control of epidemics on graphs. In *CDC*, pages 4202–4207. IEEE, 2015.
- [14] David Kempe, Jon M. Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *KDD*, pages 137–146. ACM, 2003.
- [15] David Kempe, Sixie Yu, and Yevgeniy Vorobeychik. Inducing equilibria in networked public goods games through network structure modification. In *AAMAS*, pages 611–619, 2020.
- [16] Long T. Le, Tina Eliassi-Rad, and Hanghang Tong. MET: A fast algorithm for minimizing propagation in large graphs with small eigen-gaps. In *SDM*, pages 694–702. SIAM, 2015.
- [17] Stamatios Lefkimmiatis, John Paul Ward, and Michael Unser. Hessian Schatten-norm regularization for linear inverse problems. *IEEE Trans. Image Process.*, 22(5): 1873–1888, 2013.

- [18] Jure Leskovec, Jon M. Kleinberg, and Christos Faloutsos. Graph evolution: Densification and shrinking diameters. *ACM Trans. Knowl. Discov. Data*, 1(1):2, 2007.
- [19] Jure Leskovec, Mary McGlohon, Christos Faloutsos, Natalie S. Glance, and Matthew Hurst. Patterns of cascading behavior in large blog graphs. In *SDM*, pages 551–556. SIAM, 2007.
- [20] Jure Leskovec, Lars Backstrom, and Jon M. Kleinberg. Meme-tracking and the dynamics of the news cycle. In *KDD*, pages 497–506. ACM, 2009.
- [21] Marina Meila and Jianbo Shi. Learning segmentation by random walks. In *NIPS*, pages 873–879. MIT Press, 2000.
- [22] Piet Van Mieghem, Jasmina Omic, and Robert E. Kooij. Virus spread in networks. *IEEE/ACM Trans. Netw.*, 17(1):1–14, 2009.
- [23] Adilson E Motter and Ying-Cheng Lai. Cascade-based attacks on complex networks. *Phys. Rev. E*, 66(6):065102, 2002.
- [24] Tore Opsahl. US airport network traffic data in 2010. <https://toreopsahl.com/2011/08/12/why-anchorage-is-not-that-important-binary-ties-and-sample-size-estimation/>, 2010.
- [25] Romualdo Pastor-Satorras, Claudio Castellano, Piet Van Mieghem, and Alessandro Vespignani. Epidemic processes in complex networks. *Reviews of modern physics*, 87(3):925, 2015.
- [26] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch, 2017.
- [27] B. Aditya Prakash, Deepayan Chakrabarti, Nicholas Valler, Michalis Faloutsos, and Christos Faloutsos. Threshold conditions for arbitrary cascade models on arbitrary networks. *Knowl. Inf. Syst.*, 33(3):549–575, 2012.
- [28] Victor M. Preciado, Michael Zargham, Chinwendu Enyioha, Ali Jadbabaie, and George J. Pappas. Optimal vaccine allocation to control epidemic outbreaks in arbitrary networks. In *CDC*, pages 7486–7491. IEEE, 2013.
- [29] Sudip Saha, Abhijin Adiga, B. Aditya Prakash, and Anil Kumar S. Vullikanti. Approximation algorithms for reducing the spectral radius to control epidemic spread. In *SDM*, pages 568–576. SIAM, 2015.
- [30] Comandur Seshadhri, Tamara G Kolda, and Ali Pinar. Community structure and scale-free collections of erdős-rényi graphs. *Phys. Rev. E*, 85(5):056109, 2012.
- [31] Hanghang Tong, B. Aditya Prakash, Charalampos E. Tsourakakis, Tina Eliassi-Rad, Christos Faloutsos, and Duen Horng Chau. On the vulnerability of large graphs. In *ICDM*, pages 1091–1096. IEEE Computer Society, 2010.
- [32] Hanghang Tong, B. Aditya Prakash, Tina Eliassi-Rad, Michalis Faloutsos, and Christos Faloutsos. Gelling, and melting, large graphs by edge manipulation. In *CIKM*, pages 245–254. ACM, 2012.
- [33] Hanghang Tong, B. Aditya Prakash, Tina Eliassi-Rad, Michalis Faloutsos, and Christos Faloutsos. Gelling, and melting, large graphs by edge manipulation. In *CIKM*, pages 245–254. ACM, 2012.
- [34] Leo Torres, Kevin S Chan, Hanghang Tong, and Tina Eliassi-Rad. Node immunization with non-backtracking eigenvalues. *arXiv preprint arXiv:2002.12309*, 2020.
- [35] Lloyd N Trefethen and David Bau III. *Numerical linear algebra*, volume 50. SIAM, 1997.
- [36] Lloyd N Trefethen and Mark Embree. *Spectra and pseudospectra: the behavior of nonnormal matrices and operators*. Princeton University Press, 2005.
- [37] Piet Van Mieghem, Dragan Stevanović, Fernando Kuipers, Cong Li, Ruud Van De Bovenkamp, Daijie Tan, and Shihong Wang. Decreasing the spectral radius of a graph by link removals. *Phys. Rev. E*, 84(1):016101, 2011.
- [38] Tan Van Vu and Yoshihiko Hasegawa. Diffusion-dynamics laws in stochastic reaction networks. *Phys. Rev. E*, 99:012416, 2019.
- [39] Zhihui Wang and Thomas S. Deisboeck. Dynamic targeting in cancer treatment. *Frontiers in Physiology*, 10:1–9, 2019.
- [40] Duncan J Watts and Steven H Strogatz. Collective dynamics of small-world networks. *Nature*, 393(6684):440, 1998.
- [41] Yang Yang, Takashi Nishikawa, and Adilson E Motter. Small vulnerable sets determine large network cascades in power grids. *Science*, 358(6365):eaan3184, 2017.
- [42] Sixie Yu and Yevgeniy Vorobeychik. Removing malicious nodes from networks. In *AAMAS*, pages 314–322, 2019.
- [43] Haifeng Zhang, Yevgeniy Vorobeychik, Joshua Letchford, and Kiran Lakkaraju. Data-driven agent-based modeling, with application to rooftop solar adoption. *JAAMAS*, 30(6):1023–1049, 2016.
- [44] Kai Zhou, Tomasz P. Michalak, Marcin Waniek, Talal Rahwan, and Yevgeniy Vorobeychik. Attacking similarity-based link prediction in social networks. In *AAMAS*, pages 305–313, 2019.

Appendix

A Generalization to Other Diffusion Dynamics

In this section we discuss generalization of the targeted diffusion model, i.e., Eq. (3.6), to other common diffusion dynamics. The fundamental question is: does the heuristic encoded by the model apply to other scenarios with different diffusion dynamics (e.g., SIR or SEIR)?

First, the feasible region of the model is independent of the diffusion dynamics, as it is only related to the spectral properties of the underlying graph. Thus, the structural (i.e., spectra, degree sequence, and triangle distribution) preserving properties of the diffusion model generalize to other diffusion dynamics. Next, recall that the objective function of the model is the following

$$\alpha_1 \lambda_1(\tilde{\mathbf{A}}_S) + \alpha_2 \sigma(\mathcal{S}) + \alpha_3 \phi(\mathcal{S}).$$

The third term is the normalized cut, which only depends on structural properties of the underlying graph, so it generalizes to any other diffusion dynamics. The first term also generalizes to many common diffusion dynamics, including SIR and SEIR, as their epidemic thresholds are known to also be determined by the largest eigenvalue of the underlying adjacency matrix [27]. The only exception is the second term $\sigma(\mathcal{S})$, that is, limiting the impact on non-targeted subset through maximizing the eigencentality of the targeted subset. This is because the rationale of maximizing $\sigma(\mathcal{S})$ depends on the steady state of the diffusion dynamics. Here, the steady state is where in the long run a constant (in average) fraction of infected nodes exist. However, both SIR and SEIR have been shown without a steady state, as in the long run all nodes will be in the recovered state (i.e., immune to the diffusion) [25].

Finally, the certified robustness in Section 5 generalizes to other diffusion dynamics, as its proof only depends on the spectral properties of the underlying graph.

B Degree Sequence and Triangles

We now show that satisfying the restrictions Eq (3.5) implies that certain structural properties of the graph will be perturbed by only a small amount.

Indeed, the principal's action has mild impact on the degree sequence of G . Let $\mathbf{d} = \mathbf{A}\mathbf{1}$ be the vector whose i -th entry is the degree of the i -th node in the original graph, and similarly, let $\tilde{\mathbf{d}} = \tilde{\mathbf{A}}\mathbf{1}$ be the degree sequence after the perturbation.

PROPOSITION B.1. *The degree sequence of G before and after the perturbation satisfies:*

$$(2.13) \quad \|\tilde{\mathbf{d}} - \mathbf{d}\|_2 \leq \sqrt{n}\epsilon.$$

Proof.

$$(2.14) \quad \begin{aligned} \|\tilde{\mathbf{d}} - \mathbf{d}\|_2 &= \|\tilde{\mathbf{A}}\mathbf{1} - \mathbf{A}\mathbf{1}\|_2 \stackrel{(a)}{=} \|\Delta\mathbf{1}\|_2 = \|\Delta(1/\sqrt{n})\mathbf{1}\sqrt{n}\|_2 \\ &\leq \sqrt{n} \max_{\|x\|=1} \|\Delta x\|_2 \\ &\stackrel{(b)}{=} \sqrt{n}\|\Delta\|_2 \\ &\leq \sqrt{n}\epsilon, \end{aligned}$$

where (a) is due to the fact that $\tilde{\mathbf{A}} - \mathbf{A} = \Delta$, and (b) comes from the definition of spectral norm. $\square$

A direct corollary of Proposition B.1 concerns the average degree of G .

COROLLARY B.1. *The average degree of G after the perturbation is within ϵ of the average degree before the perturbation:*

$$(2.15) \quad \left| d_{avg}(G; \tilde{\mathbf{A}}) - d_{avg}(G; \mathbf{A}) \right| \leq \epsilon.$$

Proof. Note that $d_{avg}(\tilde{G}) = \frac{1^\top \tilde{\mathbf{d}}}{n}$ and $d_{avg}(G) = \frac{1^\top \mathbf{d}}{n}$. Thus we have:

$$(2.16) \quad \begin{aligned} \left| 1^\top \tilde{\mathbf{d}}/n - 1^\top \mathbf{d}/n \right| &= (1/n) \left| 1^\top (\tilde{\mathbf{d}} - \mathbf{d}) \right| \\ &\leq (1/n) \|1\|_2 \cdot \|\tilde{\mathbf{d}} - \mathbf{d}\|_2 \\ &= \epsilon. \end{aligned}$$

$\square$

Next, we perform a similar analysis for the number of triangles before and after the perturbation.

PROPOSITION B.2. *Assume G is unweighted with m edges and T triangles. Suppose the number of triangles after the perturbation is $\tilde{T}$. Then we have*

$$(2.17) \quad |T - \tilde{T}| \leq \epsilon m,$$

where the estimate is correct up to a first order approximation.

Proof. Since G is unweighted, we have $T = \text{Tr}(\mathbf{A}^3)/6$, where Tr is the trace operator. The restrictions Eq. (3.5) guarantee that we can write $\lambda_i(\tilde{\mathbf{A}}) = \lambda_i(\mathbf{A}) + \eta_i\epsilon$, where $\eta_i \in [-1, 1]$. Thus,

$$(2.18) \quad 6\tilde{T} = \sum_i \lambda_i(\tilde{\mathbf{A}}^3) = \sum_i \lambda_i(\tilde{\mathbf{A}})^3 = \sum_i (\lambda_i(\mathbf{A}) + \eta_i\epsilon)^3.$$

Expanding the cube and neglecting the terms of higher order in η_i , we have

$$(2.19) \quad 6\tilde{T} \approx \sum_i \lambda_i(\mathbf{A})^3 + 3\epsilon \sum_i \lambda_i(\mathbf{A})^2 \eta_i = 6T + 3\epsilon \sum_i \lambda_i(\mathbf{A})^2 \eta_i.$$

And thus

$$(2.20) \quad \begin{aligned} 2|\tilde{T} - T| &\approx \epsilon \left| \sum_i \lambda_i(\mathbf{A})^2 \eta_i \right| \leq \epsilon \sum_i |\lambda_i(\mathbf{A})^2 \eta_i| \\ &\leq \epsilon \sum_i \lambda_i(\mathbf{A})^2 \\ &= \epsilon \text{Tr}(\mathbf{A}^2). \end{aligned}$$

Since $\mathbf{A}$ is symmetric and binary, we have $\text{Tr}(\mathbf{A}^2) = 2m$, where m is the number of edges in G . $\square$

In what follows we present experimental results to show that the spectra and degree sequences do not change a lot due to the targeted diffusion. The spectra and degree sequences of G and $\tilde{G}$ are showed in Figure 4, in which the top row presents the spectra with the eigenvalues as ranked in descending order, and the bottom row presents the degree sequences. The three columns (from left to right) correspond to the email network, the airport network, and the brain network, respectively. The parameter γ is set to 0.5, the most powerful principal.

The Email Network: From Figure 4 (top row), the eigenvalues with large value admit the largest deviation, while the bottom row of that figure shows that the degree sequence is not significantly affected by the targeted diffusion. In fact, the change to the original degree sequence is mild, and a student's t-test cannot differentiate the modified degree sequence from the original one (p-value=0.081).

The Airport and The Brain Networks: The airport network is directed and each node pair is associated with two edges in opposite directions. We convert the network to an undirected one by substituting an undirected edge for the two edges. The weight of the undirected edge is the sum of the weights on the two edges. The spectra and degree sequences of G and $\tilde{G}$ are showed in the last two columns of Figure 4. As we can see from Figure 4 (top row), the graph spectrum is again nearly preserved, except the eigenvalues with small values admit some deviation. Similarly, the modified degree sequences cannot be differentiated from the original one by student's t-tests (airport: p-value=0.4969, brain: p-value=0.9919).




Figure 4: Spectra (top) and degree sequences (bottom) for the **left**: the email network; **middle**: the airport network; and **right**: the brain network. The hyper-parameters are set to $\alpha_1 = \alpha_2 = \alpha_3 = 1/3$.

C Proof of Theorem 5.1

Proof. From the discussion in the main paper, an instance $\text{TargetDiff}(\mathcal{S}, G, \epsilon)$ can be encoded by the following meta

model:

$$(3.21) \quad \begin{aligned} \max_{\tilde{\mathbf{A}}} \quad & I(\tilde{G}_{\mathcal{S}}) - I(G_{\mathcal{S}}) \\ \text{s.t.} \quad & \tilde{\mathbf{A}} \in \mathcal{P}. \end{aligned}$$

As discussed in [22] (see Section IV.B), computing the exact value of $I(G_{\mathcal{S}})$ is intractable, since the exact computation of π_i is challenging. Our model Eq. (3.6) can be thought of as a tractable proxy to the meta model. An estimation to $I(G_{\mathcal{S}})$ is given in [22], i.e., $\hat{I}(G_{\mathcal{S}}) = \sum_{i \in \mathcal{S}} 1 - \delta/(\beta d_i)$, where d_i is the degree of node i in G . The estimator works in the region $\delta/\beta \leq d_{\min}$, where $d_{\min}$ is the minimum degree of G . When the estimation is reasonably good, that is $|\hat{I}(G_{\mathcal{S}}) - I(G_{\mathcal{S}})| \leq \tau$, we have the following relation:

$$(3.22) \quad \begin{aligned} I(\tilde{G}_{\mathcal{S}}) - I(G_{\mathcal{S}}) > 2\tau &\implies (\hat{I}(\tilde{G}_{\mathcal{S}}) + \tau) - (\hat{I}(G_{\mathcal{S}}) - \tau) > 2\tau \\ &\implies \hat{I}(\tilde{G}_{\mathcal{S}}) - \hat{I}(G_{\mathcal{S}}) > 0. \end{aligned}$$

Thus, in what follows we focus on deriving the necessary condition for $\hat{I}(\tilde{G}_{\mathcal{S}}) - \hat{I}(G_{\mathcal{S}}) > 0$, which directly translates to the necessary condition for $I(\tilde{G}_{\mathcal{S}}) - I(G_{\mathcal{S}}) > 2\tau$.

Suppose there exists an adjacency matrix $\tilde{\mathbf{A}}^* \in \mathcal{P}$ such that $I(\tilde{G}_{\mathcal{S}}) - I(G_{\mathcal{S}}) > 2\tau$. This indicates that the corresponding instance $\text{TargetDiff}(\mathcal{S}, G, \epsilon)$ is successful. Consequently, it follows that $\hat{I}(\tilde{G}_{\mathcal{S}}) - \hat{I}(G_{\mathcal{S}}) > 0$. Recall that $\hat{I}(\tilde{G}_{\mathcal{S}}) - \hat{I}(G_{\mathcal{S}}) = \frac{\delta}{\beta} \sum_{i \in \mathcal{S}} (\frac{1}{\tilde{d}_i} - \frac{1}{d_i})$. Let $\tilde{\mathbf{d}}_{\mathcal{S}} \in \mathbb{R}^{|\mathcal{S}|}$ represent the degree sequence of nodes in $\mathcal{S}$. Due to Proposition B.1 we have $\|\tilde{\mathbf{d}}_{\mathcal{S}} - \mathbf{d}_{\mathcal{S}}\|_2^2 \leq \|\tilde{\mathbf{d}} - \mathbf{d}\|_2^2 \leq n\epsilon^2$. Consider the optimization problem in Eq. (3.23), where the objective function is $\hat{I}(\tilde{G}_{\mathcal{S}}) - \hat{I}(G_{\mathcal{S}})$ (up to a multiplicative factor). The last constraint follows from the assumption that for a successful instance the degrees of nodes in $\mathcal{S}$ increase. The fact that $\hat{I}(\tilde{G}_{\mathcal{S}}) - \hat{I}(G_{\mathcal{S}}) > 0$ implies that the optimal solution of Eq. (3.23) exists and the associated objective value is greater than zero.

$$(3.23) \quad \begin{aligned} \max_{\tilde{\mathbf{d}}_{\mathcal{S}}} \quad & \sum_{i \in \mathcal{S}} \left(\frac{1}{\tilde{d}_i} - \frac{1}{d_i} \right) \\ \text{s.t.} \quad & \|\tilde{\mathbf{d}}_{\mathcal{S}} - \mathbf{d}_{\mathcal{S}}\|_2^2 \leq n\epsilon^2 \\ & \tilde{\mathbf{d}}_{\mathcal{S}} \succeq \mathbf{d}_{\mathcal{S}}. \end{aligned}$$

Denote the feasible region of the above optimization problem as $\mathcal{M}$. Note that $\mathcal{M}$ is a convex set since it is the intersection of two convex sets.

The objective function is concave in $\tilde{\mathbf{d}}_{\mathcal{S}}$, since it is twice differentiable on the feasible region $\mathcal{M}$ and the Hessian matrix is negative definite; the Hessian matrix is a diagonal matrix with the i -th diagonal element being $-2/\tilde{d}_i^3$. Thus, Eq. (3.23) is a convex optimization problem. Note that the Slater's condition is satisfied (e.g., with $\tilde{\mathbf{d}}_{\mathcal{S}} = \mathbf{d}_{\mathcal{S}}$), which indicates that strong duality holds. Thus, the KKT conditions are satisfied at *any* primal and dual optimal solutions.

For convenience, in what follows we use d_i (resp. $\tilde{d}_i$) to represent the degree of a node $i \in \mathcal{S}$ before (resp. after) graph modification. The Lagrange function of Eq. (3.23) is: (3.24)

$$\mathcal{L}(\tilde{\mathbf{d}}_{\mathcal{S}}, \lambda, \beta) = \sum_{i \in \mathcal{S}} \left(\frac{1}{d_i} - \frac{1}{\tilde{d}_i} \right) + \lambda \left(n\epsilon^2 - \|\tilde{\mathbf{d}}_{\mathcal{S}} - \mathbf{d}_{\mathcal{S}}\|_2^2 \right) + \beta^\top (\tilde{\mathbf{d}}_{\mathcal{S}} - \mathbf{d}_{\mathcal{S}}),$$

where $\lambda \geq 0$ and $\beta \succeq \mathbf{0}$ are Lagrangian multipliers. Recall that the degrees of nodes in $\mathcal{S}$ are increased, i.e., $\tilde{d}_i \geq d_i$ for all $i \in \mathcal{S}$. Note that for a node $i \in \mathcal{S}$ such that $\tilde{d}_i = d_i$, we let the corresponding $\beta_i = 0$. Thus, by complementary slackness, we have $\beta_i = 0$ for all $i \in \mathcal{S}$. The gradient of $\mathcal{L}(\tilde{\mathbf{d}}_{\mathcal{S}}, \lambda, \beta)$ w.r.t. $\tilde{d}_i$ becomes:

$$\frac{\partial \mathcal{L}}{\partial \tilde{d}_i} = \frac{1}{\tilde{d}_i^2} - 2\lambda \tilde{d}_i + \beta_i = \frac{1}{\tilde{d}_i^2} - 2\lambda \tilde{d}_i.$$

Setting the gradient to zero leads to:

$$\tilde{d}_i = \left(\frac{1}{2\lambda} \right)^{1/3}, \forall i \in \mathcal{S}.$$

Since the optimal solution exists, we have $\lambda \neq 0$. By complementary slackness we have $\lambda (n\epsilon^2 - \|\tilde{\mathbf{d}}_{\mathcal{S}} - \mathbf{d}_{\mathcal{S}}\|_2^2) = 0$, which indicates:

$$n\epsilon^2 = \|\tilde{\mathbf{d}}_{\mathcal{S}} - \mathbf{d}_{\mathcal{S}}\|_2^2.$$

Expand the above equation:

$$\begin{aligned} n\epsilon^2 &= \|\tilde{\mathbf{d}}_{\mathcal{S}} - \mathbf{d}_{\mathcal{S}}\|_2^2 \\ &= \sum_{i \in \mathcal{S}} (\tilde{d}_i - d_i)^2 \\ (3.25) \quad &= \sum_{i \in \mathcal{S}} \left(\left(\frac{1}{2\lambda} \right)^{1/3} - d_i \right)^2. \end{aligned}$$

Substitute $\left(\frac{1}{2\lambda} \right)^{1/3}$ with a variable x and re-arrange the above equation:

$$x^2 - \frac{2 \sum_{i \in \mathcal{S}} d_i}{|\mathcal{S}|} x + \frac{\sum_{i \in \mathcal{S}} d_i^2 - n\epsilon^2}{|\mathcal{S}|} = 0.$$

According to vieta theorem, a necessary condition that we can solve for $x \in \mathbb{R}$ from the above equation is:

$$\left(\frac{2 \sum_{i \in \mathcal{S}} d_i}{|\mathcal{S}|} \right)^2 - 4 \left(\frac{\sum_{i \in \mathcal{S}} d_i^2 - n\epsilon^2}{|\mathcal{S}|} \right) \geq 0,$$

which leads to:

$$\epsilon \geq \sqrt{\frac{|\mathcal{S}|}{n}} \left(\frac{\sum_{i \in \mathcal{S}} d_i^2}{|\mathcal{S}|} - \frac{(\sum_{i \in \mathcal{S}} d_i)^2}{|\mathcal{S}|^2} \right)^{1/2}.$$

□

D Additional Results on Real Networks

We run three experiments to show the effectiveness of each term in the objective function of POTION. The results are showed in Figure 5. The three rows (from top to bottom) correspond to experimental results on the email, airport and brain networks.

The first column corresponds to the first experiment, which is to show that maximizing $\lambda_1(\tilde{\mathbf{A}}_{\mathcal{S}})$ leads to higher infected ratios. The hyper-parameters are set to $\alpha_1 = 1/3, \alpha_2 = 0$, and $\alpha_3 = 1/3$ (the hyper-parameters do not need to sum to one). The labels of the y -axis become $I_{\text{modified}}^{\mathcal{S}} - I_{\text{original}}^{\mathcal{S}}$, which highlights that the infected ratios are for the targeted subgraph $G_{\mathcal{S}}$ (the higher the better). Note that α_2 is set to zero in order to avoid the coupling between the eigenvector centrality of $\mathcal{S}$ and $\lambda_1(\tilde{\mathbf{A}}_{\mathcal{S}})$. From the plot it is clear that maximizing $\lambda_1(\tilde{\mathbf{A}}_{\mathcal{S}})$ is important to increase the infected ratios within $G_{\mathcal{S}}$ (the blue line). Note that solely maximizing the normalized cut of $\mathcal{S}$ may backfire (the red line), as a large portion of edges are deleted from $G_{\mathcal{S}}$ when γ increases.

The second column is to show the effectiveness of limiting the impact on $G_{\mathcal{S}'}$ by maximizing the eigenvector centrality of $\mathcal{S}$. The y -axis represents $I_{\text{modified}}^{\mathcal{S}'} - I_{\text{original}}^{\mathcal{S}'}$, which highlights that the infected ratios are for the non-targeted subgraph $G_{\mathcal{S}'}$ (the lower the better). The plot shows that the impact on $G_{\mathcal{S}'}$ is well limited; the effectiveness is most significant when $\gamma > 40\%$.

The last column is to show the effectiveness of maximizing the normalized cut of $\mathcal{S}$. We set $\alpha_2 = 0$ to avoid the effect of maximizing the eigenvector centrality of $\mathcal{S}$. Observe that maximizing the normalized cut of $\mathcal{S}$ is effective in increasing the infected ratio within $G_{\mathcal{S}}$ only for the email network. This suggests that on weighted graphs normalized cut might not be a good heuristic to increase the centrality of $\mathcal{S}$.

E Additional Results on Synthetic Networks

In this section we show experimental results on synthetic unweighted graphs with 375 nodes. We focus on three classes of networks: Barabasi-Albert (BA), Watts-Strogatz, and BTER [30]. BA is characterized by its power-law degree distribution [4]. Watts-Strogatz is well-known for its local clustering in a way as to qualitatively resemble real networks [40]. BTER are generative network models that can be calibrated to match real-world networks, in particular, to reproduce the community structures [30].

The experimental setup is similar to the setup for the email network, except for a few changes. First, the experimental results for each class of the synthetic networks are averaged over 30 randomly generated network topologies. Another difference lies in how the targeted set $\mathcal{S}$ is selected. For each randomly generated network, the targeted set $\mathcal{S}$ is selected as the node whose degree is the 90 percentile of the degree sequence, and its neighbors. Some statistics of the synthetic networks are summarized in Table 2. Recall that $\delta = 0.24$ and




Figure 5: Experiments showing the model’s effectiveness. **Top:** the email network; **Middle:** the airport network; **Bottom:** the brain network. The three columns (from left to right) show the effectiveness of: 1) maximizing $\lambda_1(\mathbf{A}_S)$; 2) Maximizing the eigenvector centrality of $\mathcal{S}$; 3) maximizing the normalized cut of $\mathcal{S}$.

$\beta = 0.06$. The experimental results are showed in Figure 6. The conclusion derived from Figure 6 is similar to that of the email network. It is worth pointing out that maximizing the normalized cut of $\mathcal{S}$ is effective on BA networks, while for other network it may backfire.

	BA	Watts-Strogatz	BTER
$ \mathcal{S} $	17.5	12	20.03
d_{min}	9.86	10	11.69
density	0.02	0.03	0.03
average degree	9.87	10	11.5
average clustering coeff.	0.08	0.35	0.05

Table 2: Statistics of synthetic networks.

F Additional Results for Different Values of δ and β

In the main paper, $\delta = 0.24$ and $\beta = 0.2$ for the airport and brain networks, while $\delta = 0.24$ and $\beta = 0.06$ for the email network. The ratio δ/β is 1.2 for the former two networks, while 4 for the latter. In what follows we explore the effectiveness of our model in different regimes of δ/β . For the airport and brain networks, we present results for $(\delta = 0.5, \beta = 0.1)$ and $(\delta = 0.3, \beta = 0.5)$. The former (resp. latter) corresponds to the regime above (resp. below) 1.2. The results for the airport network are showed in Figure 9, and the results for the brain network are in Figure 10. For the email network we present results for $(\delta = 0.5, \beta = 0.1)$ and $(\delta = 0.3, \beta = 0.5)$, also corresponds to the regime above and below the original ratio respectively. The results are showed in Figure 11. The conclusions are consistent with that presented in the main paper.

G Results for Random Walk Based Spreading Dynamics

We simulate random walk based spreading dynamics on the original and the modified networks. Although POTION is motivated from the analysis of SIS spreading dynamics, the simulation results show that it is capable of achieving targeted diffusion when the underlying spreading dynamics is based on random walk. Random walk has extensive use in machine learning, data mining, security, ranking, etc. [? ? ? ? ?]. We focus on two variants of random walks: random walk with restart (RWR, a.k.a. personalized PageRank) and PageRank. The former has been widely used in data mining and security applications [? ?]. The latter is a powerful tool to measure the “importance” of nodes in a network [?]. We run POTION on the Airport, Brain, and Email networks, with the same targeted subgraphs as in previous experiments. The trade-off parameters are set to $(\alpha_1 = 1/3, \alpha_2 = 1/3, \alpha_3 = 1/3)$.

For the RWR dynamics, the starting node of a random walk is picked uniformly at random from the non-targeted subgraph $G_{S'}$. The restart probability c is set to 0.05 – i.e., at each time step the RWR restarts from the starting node with probability 0.05. The RWR dynamics is simulated until convergence,⁸ which gives us a rank vector $\mathbf{r} \in \mathbb{R}_+^n$ over the nodes for the given starting node. The sum of the sub-vector $\mathbf{r}[\mathcal{S}]$ (resp. $\mathbf{r}[\mathcal{S}']$) is the probability that a random walk lands in the targeted subgraph G_S (resp. non-targeted subgraph $G_{S'}$), which quantifies the impact on G_S (resp. $G_{S'}$). Figure 7 shows the experimental results here. The left column represents the landing probability on the targeted subgraph G_S . It is clear that the probability is higher when the underlying graph is modified by POTION (although the difference is only statistically significant on the Email network). The right column is the landing probability on the non-targeted subgraph $G_{S'}$. The probability does not increase, which is desired as we would like to limit the impact on $G_{S'}$.

For the PageRank dynamics, the starting node is picked from the node set $\mathcal{V}$ uniformly at random. The restart probability c is set to 0.1 – i.e., at each time step the PageRank dynamics restarts with probability 0.1 from a node (not necessarily the starting node) picked from $\mathcal{V}$ uniformly at random. When the simulation is finished the PageRank gives a vector $\mathbf{r} \in \mathbb{R}_+^n$ indicating how “important” each node is. Intuitively, $\mathbf{r}$ specifies a ranking of the nodes in $\mathcal{V}$ – i.e., a node $i \in \mathcal{V}$ is ranked higher when $r[i]$ is larger. We use the sum of the sub-vector $\mathbf{r}[\mathcal{S}]$ (resp. $\mathbf{r}[\mathcal{S}']$) to quantify the impact on G_S (resp. $G_{S'}$). Other experimental setup is the same as the setup for the RWR dynamics. Figure 8 shows these results. The left column indicates the ranking of the nodes in G_S . It is clear that the ranking is boosted and the increase is statistically significant. The right column shows that the ranking of the nodes in $G_{S'}$ is not increased, as desired.

⁸We are guaranteed convergence since the Markov transition matrix of the network is stochastic, irreducible, and aperiodic.




Figure 6: Experimental results on synthetic networks. The first column shows the overall effectiveness of the threat model. Remaining columns show the effectiveness of: 1) maximizing $\lambda_1(\hat{A}_S)$; 2) Maximizing the eigenvector centrality of S ; 3) maximizing the normalized cut of S . **Top:** BA; **Middle:** Watts-Strogatz; **Bottom:** BTER




Figure 7: Landing probability on (left) the targeted subgraph G_S and (right) the non-targeted subgraph $G_{S'}$. Modifications by POTION increase the probabilities of landing in the targeted subgraph G_S (left column) and do not increase the probabilities of landing in the non-targeted subgraph $G_{S'}$ (right column), as desired. **Top:** Email; **Middle:** Brain; **Bottom:** Airport.




Figure 8: Ranking of the nodes in (left) the targeted subgraph G_S and (right) the non-targeted subgraph $G_{S'}$. Modifications by POTION increase the probabilities of landing in the targeted subgraph G_S (left column) and do not increase the probabilities of landing in the non-targeted subgraph $G_{S'}$ (right column), as desired. **Top:** Email; **Middle:** Brain; **Bottom:** Airport.




Figure 9: Experimental results for different δ and β values on the airport network. **Top:** $\delta = 0.5, \beta = 0.1$; **Bottom:** $\delta = 0.3, \beta = 0.5$.




Figure 10: Experimental results for different δ and β values on the brain network. **Top:** $\delta = 0.5, \beta = 0.1$; **Bottom:** $\delta = 0.3, \beta = 0.5$.




Figure 11: Experimental results for different δ and β values on the email network. **Top:** $\delta = 0.5, \beta = 0.1$; **Bottom:** $\delta = 0.3, \beta = 0.5$.