

Effective submodularity of influence maximization on temporal networksŞirag Erkol ¹, Dario Mazzilli ², and Filippo Radicchi ¹¹*Center for Complex Networks and Systems Research, Luddy School of Informatics, Computing, and Engineering, Indiana University, Bloomington, Indiana 47408, USA*²*Enrico Fermi Research Center, Via Panisperna 89 A, Rome, Italy*

(Received 11 May 2022; accepted 16 August 2022; published 2 September 2022)

We study influence maximization on temporal networks. This is a special setting where the influence function is not submodular, and there is no optimality guarantee for solutions achieved via greedy optimization. We perform an exhaustive analysis on both real and synthetic networks. We show that the influence function of randomly sampled sets of seeds often violates the necessary conditions for submodularity. However, when sets of seeds are selected according to the greedy optimization strategy, the influence function behaves effectively as a submodular function. Specifically, violations of the necessary conditions for submodularity are never observed in real networks, and only rarely in synthetic ones. The direct comparison with exact solutions obtained via brute-force search indicates that the greedy strategy provides approximate solutions that are well within the optimality gap guaranteed for strictly submodular functions. Greedy optimization appears, therefore, to be an effective strategy for the maximization of influence on temporal networks.

DOI: [10.1103/PhysRevE.106.034301](https://doi.org/10.1103/PhysRevE.106.034301)**I. INTRODUCTION**

Influence maximization is the optimization problem aiming to find the most influential nodes in a network [1,2]. The influence of a set of nodes is measured in terms of the size of the outbreak that the nodes generate when used as initial seeds for a spreading process occurring on the network. Optimization is generally constrained by the number of initial spreaders.

Influence maximization is a NP-hard problem [2]. The number of candidates to the solution increases exponentially with the number of seeds, so applying brute-force search to find the exact solution is possible for very small networks only. Many attempts to find approximate solutions to the problem in a computationally feasible way are present in the literature. Some approaches consist of using greedy algorithms [2–8]. Greedy optimization has provable performance bounds thanks to the fact that the outbreak size of a spreading process occurring on a static network is a submodular function [9]. Roughly speaking, a set function is submodular if the addition one element to the input set generates an increment of the value returned by the function that decreases as the size of the input set increases. Unfortunately, greedy optimization algorithms are computationally expensive, and thus other more efficient approaches use centrality metrics to approximate the influence of nodes [10–15]. The drawback of centrality-based approaches is the lack of provable bounds of performance as they do not directly optimize the outbreak size of the spreading process.

The influence maximization problem is traditionally studied on static networks, i.e., the topology is given and does not change as spreading unfolds. However, there are many real-life networks where interactions between nodes might

happen only for a given period of time [16]. If the timescales of the structural evolution and of spreading dynamics are similar, then changes in network topology cannot be ignored as they can dramatically affect the outcome of the spreading process [17–20]. So far, just a few papers have considered the influence maximization problem on temporal networks [21–24]. Osawa *et al.* studied the influence maximization problem on temporal networks for the susceptible-infected (SI) model [25]. They proposed an alternative to the greedy algorithm, and they showed that their method is very effective in correctly identifying top spreaders in networks with community structure. Michalski *et al.* analyzed the effect of the size of time steps in temporal networks, including the aggregate version of the network [26]. They observed that using the aggregate version of the temporal network gives suboptimal results. Gayraud *et al.* studied the independent cascade and the linear threshold models, showing that the influence function is not submodular [27]. They further showed that the activation time of nodes greatly affects the size of the final outbreak, and that delaying the activation of a node can increase its effective influence.

In a recent paper, we studied the influence maximization problem under the susceptible-infected-recovered (SIR) model on temporal networks [28]. We performed a systematic analysis on 12 real-world temporal networks and analyzed the performances of different approximation methods that have different levels of knowledge on the network topology and dynamics. We found that complete knowledge of a network, but in an aggregate way, is not helpful in solving the problem effectively. On the other hand, knowledge of the initial stages of the network helps to find good solutions to the influence maximization problem. The influence function of the SIR model on temporal networks is not submodular, except

for the trivial case when the model is equivalent to the SI model. However, the solutions provided by the greedy algorithm proved to be good upper bounds for the performances of other methods to identify influential spreaders. This fact suggests that even though there is no theoretical proof of the performance of the greedy algorithm, in practice the optimization strategy is effective in approximating solutions to the influence maximization problem. This brings the questions of how often the condition for submodularity is violated and how far the solutions of the greedy algorithm are from the ground-truth optimum.

We answer the above questions in this paper. In particular, we provide evidence for an effective submodular behavior of the influence function under greedy optimization. Results of our analysis show that the condition for submodularity is violated frequently for randomly selected seeds; however, when seeds are selected with the greedy algorithm, the frequency drops to almost zero, especially for real-world networks. Also, we show that the solutions of the greedy algorithm have a performance very close to the optimal solution found with brute-force search.

II. METHODS

A. Temporal networks

A temporal network is a collection of T ordered network layers, namely, $A(0), A(1), \dots, A(t), \dots, A(T-1)$, each representing the topology of the system at a specific time. All layers of a temporal network are composed of the same N nodes, with labels uniquely identifying the nodes across the layers. The one above is just one possible definition of a temporal network; other definitions consider the evolution to happen in continuous rather than discrete time [16].

There are $R(N, T) = 2^T \binom{N}{2}$ total possible temporal networks with N nodes and T layers. There are, in fact, $2^{\binom{N}{2}}$ different labeled networks with N nodes; those networks can be permuted in $(2^{\binom{N}{2}})^T$ ways, where permutations also include repetitions.

We consider exhaustive enumerations of all possible temporal networks in only one experiment. Clearly, we choose very small values of the parameters N and T to make the enumeration computationally feasible. In all other experiments, the space of potential temporal networks is sampled by either constructing synthetic models or leveraging real data.

1. Synthetic temporal networks

We generate random temporal networks with correlated layers. Specifically, the layer $A(t+1)$ of the temporal network is obtained by copying all the edges in the layer $A(t)$, and then shuffling with probability r the end points of each individual edge with those of another random edge. For instance, if the edge (i, j) undergoes shuffling, we first select at random another edge (v, w) . We then verify that the edges (i, w) and (v, j) are not yet present in the network. If they are present, we select another edge (v, w) and repeat the operation. Otherwise, we shuffle their end points in the sense that the edges (i, j) and (v, w) are removed from the network, and they are replaced by the edges (i, w) and (v, j) . The above procedure of shuffling the end points of edges keeps the degree

sequence of the network layers unchanged. For $r = 0$, we have $A(t+1) = A(t)$ for all t , i.e., layers are perfectly correlated and the temporal network is essentially a static network. For $r = 1$, all edges are surely shuffled, so no correlation exists between $A(t+1)$ and $A(t)$ except for the fact that they have the same degree sequence. In our experiments, we start by generating the first layer $A(0)$ according to the Erdős-Rényi model with average degree k .

Also, we consider a simple uncorrelated model where all layers are generated according to the Erdős-Rényi model with average degree k . Networks of this type are statistically equivalent to those created according to the model above with reshuffling probability $r = 1$, with the only difference that whereas the average degree of the network is invariant across network layers, the degree sequence is not preserved across network layers. In all cases where the value of the parameter r is not specified, we will take advantage of this simple model to generate uncorrelated random temporal networks.

We note that the hypothesis of having an average degree that is invariant across layers is not necessarily a characteristic of real-world temporal networks [28]. To study the problem of influence maximization in realistic settings, we generate temporal networks directly from real data without the need to make any assumption, as described below.

2. Real-world temporal networks

We use 12 empirical datasets containing time-stamped interactions between pairs of nodes. Some datasets contain only bidirectional interactions, while others consist of only unidirectional interactions. In both cases, we treat the interactions as undirected. To construct the temporal networks, we follow the exact same procedure as followed in [28]. Given the dataset, we divide the interactions into temporal windows of equal length. All interactions in a single slice are aggregated to create a temporal network layer that is undirected and unweighted. After the procedure is applied to each slice, we end up with T layers of networks that form the temporal network as a whole. The networks used are listed in Table I. For more details of the procedure of temporal network construction, see Ref. [28].

B. Spreading dynamics

We consider the discrete version of the susceptible-infected-recovered (SIR) model for spreading dynamics. We use the model in the same way as already done in Ref. [28]. However, to properly compare the importance of the initial conditions on the long-term behavior of the model, we generate individual instances of the SIR model in a slightly different manner. Under this procedure, the difference in impact of different initial conditions is measured on identical, deterministic dynamical systems. We then average the difference over multiple individual instances of the SIR model.

Before starting any dynamics on the network, we generate the q th instance of the SIR model with parameters λ and μ by determining the propensity of individual edges to spread the infection and the propensity of individual nodes to recover at particular instants of time. Specifically, for each edge (i, j) appearing at time t , we set the spreading propensity of the edge $\rho_{(i,j)}^{(q)}(t) = 1$ with probability λ ; otherwise, we set

TABLE I. Real-world temporal networks. List of the empirical datasets used to construct temporal networks. From left to right, we report the name of the dataset, the number of network layers T , the number of nodes N in the network, the reference to the paper(s) where the data were first considered, and the label of the dataset as used in the analysis of Fig. 7.

Dataset	T	N	Ref.	Label
Email, dept. 1	18	309	[29]	a
Email, dept. 2	18	162	[29]	b
Email, dept. 3	18	89	[29]	c
Email, dept. 4	18	142	[29]	d
High school, 2011	11	126	[30]	e
High school, 2012	21	180	[30]	f
High school, 2013	14	327	[31]	g
Hospital ward	20	75	[32]	h
Hypertext, 2009	11	113	[33]	i
Primary school	11	242	[34,35]	j
Workplace	20	92	[36]	k
Workplace-2	20	217	[37]	l

$\rho_{(i,j)}^{(q)}(t) = 0$. Also, we set the recovery propensity of node i at time t as $\rho_i^{(q)}(t) = 1$ with probability μ , and $\rho_i^{(q)}(t) = 0$ otherwise.

Once propensities of edges and nodes are set for all temporal layers of the network, SIR dynamics can be run in a deterministic fashion, starting from the initial condition $\vec{\sigma}(t=0) = [\sigma_1(t=0), \dots, \sigma_N(t=0)]$, where $\sigma_i(0) = S, I$, or R . Please note that the initial condition does not depend on the actual realization of the SIR model. The rules that determine the dynamics for $t > 0$ are as follows. Indicate with $\sigma_i^{(q)}(t)$ the dynamical state of node i at time t in the q th realization of the SIR model. We have that

$$\sigma_i^{(q)}(t+1) = S \quad \text{if} \quad \begin{cases} \sigma_i^{(q)}(t) = S \\ \bigwedge_{\#j} | \sigma_j^{(q)}(t) = I \wedge \rho_{(i,j)}^{(q)}(t) = 1 \end{cases}, \quad (1)$$

meaning that node i remains in the state S if it does not get in active contact with any infected neighbor. Also, we have that

$$\sigma_i^{(q)}(t+1) = I \quad \text{if} \quad \sigma_i^{(q)}(t) = I \wedge \rho_i^{(q)}(t) = 0 \quad (2)$$

and

$$\sigma_i^{(q)}(t+1) = I \quad \text{if} \quad \begin{cases} \sigma_i^{(q)}(t) = S \\ \bigwedge_{\exists j} | \sigma_j^{(q)}(t) = I \wedge \rho_{(i,j)}^{(q)}(t) = 1 \end{cases}. \quad (3)$$

Equation (2) describes the case of a node already infected that does not recover. Equation (3) accounts instead for the change of the dynamical state of the node i getting infected because of an active contact with at least one infected neighbor. Finally, we have that

$$\sigma_i^{(q)}(t+1) = R \quad \text{if} \quad \begin{cases} \sigma_i^{(q)}(t) = I \wedge \rho_i^{(q)}(t) = 1 \\ \bigvee \\ \sigma_i^{(q)}(t) = R \end{cases}, \quad (4)$$

i.e., node i recovers if infected and prone to recovery at time t , or it does not change its state if it already recovered. After

all of the above operations are executed for all nodes i , time increases as $t \rightarrow t+1$.

Assuming that network evolution and spreading dynamics happen in discrete time and precisely at the same timescale simplifies the numerical and analytical analysis of the dynamical system. We expect results obtained under these simplifications to also be valid for temporal networks evolving in continuous time as long as the duration of individual edges is sufficiently homogeneous and provided that the SIR model is reformulated in continuous time [38]. We recognize, however, that our framework would likely fail to properly describe SIR dynamics happening on temporal networks characterized by heterogeneous activity of the edges.

In our experiments, we restrict our attention to initial conditions where all nodes are in the susceptible state, except the nodes in the seed set \mathcal{X} which are in the infected state, i.e., $\sigma_i(t=0) = I$ if $i \in \mathcal{X}$ and $\sigma_i(t=0) = S$ if $i \notin \mathcal{X}$. Starting from such initial configurations, the dynamics in a network with T layers is simulated until the stage T . The outbreak size $f^{(q)}(\mathcal{X})$ at the end of the q th realization of the process is calculated as the total number of infected and recovered nodes at T , i.e.,

$$f^{(q)}(\mathcal{X}) = \frac{1}{N} \sum_{i=1}^N [\mathbb{1}_{\sigma_i^{(q)}(T),I} + \mathbb{1}_{\sigma_i^{(q)}(T),R}],$$

where $\mathbb{1}_{x,y}$ is the identity operator, i.e., $\mathbb{1}_{x,y} = 1$ if $x = y$ and $\mathbb{1}_{x,y} = 0$ otherwise.

We indicate the marginal gain in influence of adding node v to the seed set \mathcal{X} as

$$f_{\mathcal{X}}^{(q)}(v) = f^{(q)}(\mathcal{X} \cup \{v\}) - f^{(q)}(\mathcal{X}). \quad (5)$$

We estimate the influence of the set \mathcal{X} by taking the average value of the outbreak size over Q independent realizations of the SIR model, i.e.,

$$\langle f(\mathcal{X}) \rangle = \frac{1}{Q} \sum_{q=1}^Q f^{(q)}(\mathcal{X}). \quad (6)$$

C. Influence maximization

Influence maximization is defined as the constrained optimization problem

$$\mathcal{X}^*(M) = \arg \max_{|\mathcal{X}|=M} \langle f(\mathcal{X}) \rangle. \quad (7)$$

Essentially, the goal of the problem is finding the set of seeds $\mathcal{X}^*(M)$ of size M corresponding to the maximum value of the influence function.

Exact solutions of the problem are obtainable via brute-force search over all possible $\binom{N}{M}$ ways of choosing M seed nodes out of N total nodes in the network. For each of these candidate sets, the influence function of Eq. (6) should be evaluated. Clearly, the brute-force search can only be applied on relatively small networks and small seed set sizes. In most of the practical settings, the solution of the problem of Eq. (7) can only be approximated.

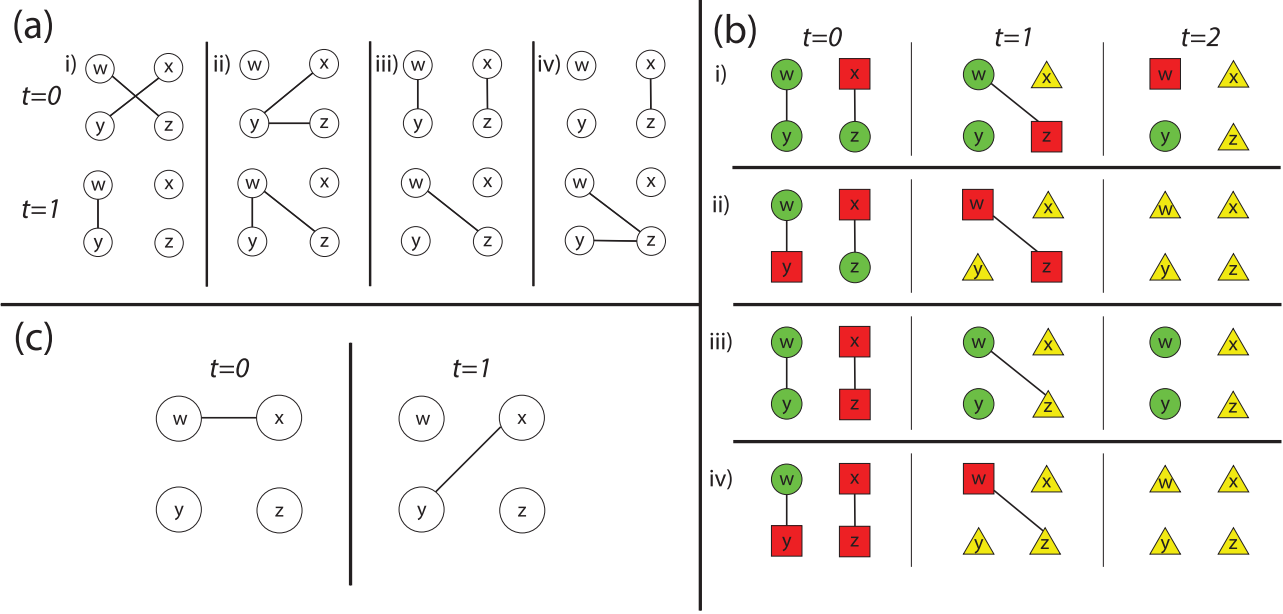


FIG. 1. Violation of the necessary conditions for the submodularity of the influence function on temporal networks. For simplicity, we consider the deterministic case of the SIR model where the probabilities of infection and recovery are $\lambda = \mu = 1$. (a) We display four possible scenarios for the marginal gain of adding node v to sets \mathcal{A} and \mathcal{B} , where $\mathcal{A} = \{x\}$, $\mathcal{B} = \{x, y\}$, $v = z$. The cases are separated on the basis of the marginal gains being negative or not. The marginal gains in the four scenarios are (i) $f_{\mathcal{A}}(v) = 1$, $f_{\mathcal{B}}(v) = 2$, (ii) $f_{\mathcal{A}}(v) = 1$, $f_{\mathcal{B}}(v) = -1$, (iii) $f_{\mathcal{A}}(v) = -1$, $f_{\mathcal{B}}(v) = 0$, and (iv) $f_{\mathcal{A}}(v) = -2$, $f_{\mathcal{B}}(v) = -1$. The inequality (10) is violated in (i), (iii), and (iv). (b) A counterexample for the submodularity of the influence function. Let $\mathcal{A} = \{x\}$, $\mathcal{B} = \{x, y\}$, $v = z$. Green circles denote nodes in the susceptible state, red squares denote nodes in the infected state, and yellow triangles denote nodes in the recovered state. (c) A counterexample for the γ -weakly submodularity of the influence function. Let $\mathcal{A} = \{w\}$, $\mathcal{B} = \{x, y, z\}$.

D. Greedy optimization

Approximate solutions to the problem can be obtained using a greedy optimization algorithm, defined as follows. Representing the seed set at stage m of the algorithm as $\mathcal{G}_m = \{g_1, g_2, \dots, g_m\}$, we initialize $\mathcal{G}_0 = \emptyset$. The seed at stage $m > 0$ is chosen as

$$g_m = \arg \max_{x \notin \mathcal{G}_{m-1}} \langle f_{\mathcal{G}_{m-1}}(x) \rangle. \quad (8)$$

The algorithm of Eq. (8) provides us with an approximation of the solution of Eq. (7) after $m = M$ iterations. As Eq. (8) tells us, at each stage of the greedy selection algorithm, the node giving rise, on average, to the largest marginal gain in the influence function is selected. The average value of the marginal gain of each node $x \notin \mathcal{G}_{m-1}$ is numerically estimated from Eq. (6) by running independent simulations.

On static networks, the influence function of Eq. (6) is a submodular function with non-negative marginal gains. These two properties guarantee that the solution provided by the greedy algorithm is at max $1 - 1/e$ times away from the ground-truth optimal solution [9]. On temporal networks, such an optimality bound is guaranteed only for $\mu = 0$. For $\mu > 0$, the influence function is not necessarily a submodular function with non-negative marginal gains, and hence there is no guarantee on the optimality gap for the solutions obtained via the greedy algorithm [27].

III. SUBMODULARITY OF THE INFLUENCE FUNCTION

In the following sections, we analyze the details for the violation of the condition for submodularity in temporal networks. In particular, we show why the influence function is neither submodular nor γ -weakly submodular.

A. Submodularity

In their work, Kempe *et al.* showed that the influence function of Eq. (6) on static networks has non-negative marginal gains and is a submodular function [2].

The influence function has non-negative marginal gains if

$$f_{\mathcal{A}}(v) \geq 0, \quad (9)$$

for any set \mathcal{A} and for any node v . Influence is a submodular function if it satisfies the submodularity or “diminishing returns” condition, i.e.,

$$f_{\mathcal{A}}(v) \geq f_{\mathcal{B}}(v), \quad (10)$$

for all nodes $v \notin \mathcal{B}$ and for all sets of nodes $\mathcal{A} \subseteq \mathcal{B}$. This means that the marginal gain obtained by adding node v to the set \mathcal{A} , a subset of \mathcal{B} , must be greater than or equal to the marginal gain obtained by adding node v to the set \mathcal{B} .

As we already mentioned, properties (9) and (10) hold for the SIR model on static networks. They hold for temporal networks too as long as the recovery probability is $\mu = 0$. However, they are generally not valid on temporal networks for $\mu > 0$. Violations of the condition of Eq. (10) may happen in three main ways, as illustrated in Fig. 1(a). Four scenarios are realized depending on whether the marginal gains $f_{\mathcal{A}}(v)$

and $f_{\mathcal{B}}(v)$ are non-negative or negative. For $f_{\mathcal{A}}(v) \geq 0$ and $f_{\mathcal{B}}(v) < 0$ in (ii), the diminishing returns condition holds; for $f_{\mathcal{A}}(v) < 0$ and $f_{\mathcal{B}}(v) \geq 0$ in (iii), the diminishing returns condition is violated. In the other two scenarios illustrated in (i) and (iv), it depends on the actual values of the marginal gain.

The inspection of Fig. 1(a) reveals that violations of the conditions necessary for the submodularity property are caused by recovered nodes blocking the paths of future infections. In static networks, a recovered node would already have exploited any possible paths to infect its neighbors. The difference in temporal networks is that the neighbors of a node change in time. This means that the infection and recovery time of a node are fundamental factors that determine which paths are effectively used by the infection to propagate. In particular, an early infection of a node may be detrimental for the long-term fate of the spreading process just because, once recovered, the node may block paths that would have been otherwise available if the node was still in the susceptible state.

For example, in Fig. 1(b), we set $\mathcal{A} = \{x\}$, $\mathcal{B} = \{x, y\}$, and $v = z$. We also set $\lambda = \mu = 1$ for simplicity. Each row shows the initial conditions at $t = 0$ for the seed sets \mathcal{A} , $\mathcal{A} \cup v$, \mathcal{B} , and $\mathcal{B} \cup v$, respectively. Infections and recoveries occur at $t = 0$ and $t = 1$, and the final configuration is seen at $t = 2$, where no more infections or recoveries happen. From Fig. 1(b), we can see that $f(\mathcal{A}) = 3$ in (i), $f(\mathcal{B}) = 4$ in (ii), $f(\mathcal{A} \cup v) = 2$ in (iii), and $f(\mathcal{B} \cup v) = 4$ in (iv). Here, the condition of Eq. (10) is violated. The main reason for the violation is the premature infection of node z . The marginal gain of adding node z to set \mathcal{A} is $f_{\mathcal{A}}(z) = -1$. z recovered at $t = 1$, thus is not able to infect w . The infection could have occurred if the node was susceptible at the beginning of dynamics. So, by adding node z to the set \mathcal{A} , the path to infecting node w is blocked by the premature infection and recovery of node z , thus decreasing the total outbreak size at the end of the dynamics.

A toy example of how multiple, misplaced seeds may have dramatic consequences on the size of the outbreak is provided in Fig. 2. This is a rather specific and unrealistic example, where a single edge is present at each layer. We expect, however, that the very same issue, although in more complicated forms, is at the basis of violations of the necessary conditions for the submodularity of the influence function in real temporal networks.

B. γ -weakly submodularity

After showing that the influence function is not submodular, we can test weaker definitions of submodularity that would give us looser optimality gaps. In their work, Santiago and Yoshida defined γ -weakly submodularity for nonsubmodular functions [39]. In their definition, a function is γ -weakly submodular when

$$\sum_{v \in \mathcal{B}} f_{\mathcal{A}}(v) \geq \min \left\{ \gamma f_{\mathcal{A}}(\mathcal{B}), \frac{1}{\gamma} f_{\mathcal{A}}(\mathcal{B}) \right\} \quad (11)$$

for any disjoint sets of nodes \mathcal{A} and \mathcal{B} . In the above inequality, $0 < \gamma \leq 1$. When the inequality holds, it is possible to find a solution with a so-called randomized greedy algorithm resulting in an optimality gap equal to $1 - \gamma e^{-1/\gamma}$.

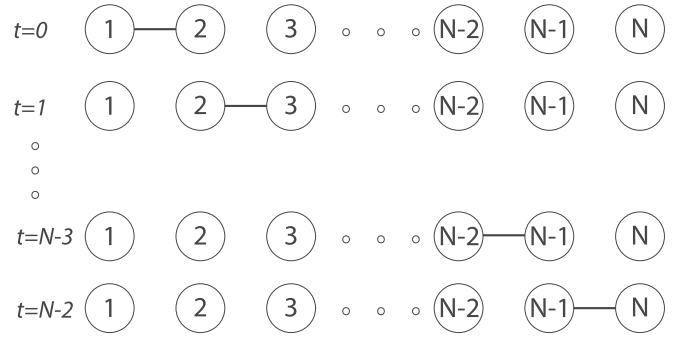


FIG. 2. Blocking paths of future infections with multiple seeds. We display a toy network where increasing the number of seeds have catastrophic effects on the outbreak size of the spreading process. For simplicity, we consider the deterministic case of the SIR model where the probabilities of infection and recovery are $\lambda = \mu = 1$. Setting node 1 as the only seed of the process leads to maximum spread in the network, i.e., $f(\{1\}) = 1$. However, adding another node $i > 1$ to the seed set generates a reduction in the influence function, i.e., $f(\{1, i\}) = i/N$.

Unfortunately, the influence function of the SIR model on temporal networks is not γ -weakly submodular. An example of the violation of the condition of Eq. (11) is shown in Fig. 1(c). We select $\mathcal{A} = \{w\}$, $\mathcal{B} = \{x, y, z\}$ and $\lambda = \mu = 1$. The marginal gains are $f_{\mathcal{A}}(\mathcal{B}) = 1$, $f_{\mathcal{A}}(x) = -1$, $f_{\mathcal{A}}(y) = 0$, and $f_{\mathcal{A}}(z) = 1$. The left-hand side of the inequality (11) reads $\sum_{v \in \mathcal{B}} f_{\mathcal{A}}(v) = 0$. This means that for the inequality to hold, we need $\gamma = 0$ or $\gamma = \infty$ because $f_{\mathcal{A}}(\mathcal{B}) = 1$. However, the definition of γ -weakly submodularity requires $0 < \gamma \leq 1$, meaning that the inequality does not hold.

IV. RESULTS

There are temporal networks where the inequality (10) is violated and the influence function is not submodular. In these situations, the greedy algorithm does not provide any guarantee on the optimality of its solutions. However, it is still possible that the algorithm provides solutions close enough to the ground-truth optimum. To investigate this property, we perform a systematic analysis on synthetic and real-world networks.

A. Synthetic temporal networks

1. Random selection of seeds

We calculate the frequency of violations of the inequality (10) as

$$g(\mathcal{A}, \mathcal{B}, v) = \frac{1}{Q} \sum_{q=1}^Q H[f_{\mathcal{B}}^{(q)}(v) - f_{\mathcal{A}}^{(q)}(v)], \quad (12)$$

where $\mathcal{A} \subset \mathcal{B}$ and $v \notin \mathcal{B}$, and $H(x)$ is the Heaviside step function, i.e., $H(x) = 1$ if $x > 0$ and $H(x) = 0$ otherwise. Please note that $f_{\mathcal{A}}^{(q)}(v)$ and $f_{\mathcal{B}}^{(q)}(v)$ are both computed on the same q th instance of the SIR model.

In Fig. 3(a), we consider all the $R(N = 4, T = 3) = 262, 144$ possible temporal networks that can be formed with $N = 4$ nodes and $T = 3$ layers. Indicating the four nodes of




FIG. 3. Violations of the submodularity condition on temporal networks. (a) We display the frequency of violations of the diminishing returns inequality, i.e., g as defined in Eq. (12), on random synthetic networks of size $N = 4$ as a function of the SIR parameters λ and μ . In the computation of Eq. (12), the sets \mathcal{A} and \mathcal{B} , and node v are selected randomly with $|\mathcal{A}| = 1$, $|\mathcal{B}| = 2$, and $\mathcal{A} \subset \mathcal{B}$. The simulations are run on a network with $N = 4$ and $T = 3$, and all possible configurations with this specific parameters have been used for the experiments. (b) Same as in (a), but only on a random temporal network with $N = 100$, $T = 10$, $k = 5$, and $r = 1$. Results are averaged over 50 networks. The dashed black line shows the critical threshold values $\lambda_c(\mu)$ averaged over 10 networks. (c) Same as in (a), but for the real-world temporal network “High school, 2012.” The dashed black line shows the critical threshold values $\lambda_c(\mu)$. (d) Same as in (a), but for the real-world temporal network “Hypertext, 2009.”

the network as v , w , x , and y , we set $\mathcal{A} = \{y\}$ and $\mathcal{B} = \{x, y\}$. Please note that since we are considering all possible networks, the choice of the sets is irrelevant for our purposes. For a given network, we compute Eq. (12) over $Q = 100$ SIR model instances for each combination of μ and λ values. We then take the average value of $g(\mathcal{A}, \mathcal{B}, v)$ over all possible networks to generate the heat map of Fig. 3(a). For $\lambda = 0.00$, the influence function is submodular, there is no spreading, and $f_{\mathcal{A}}(v) = f_{\mathcal{B}}(v) = 1$. For $\mu = 0.00$, the spreading model becomes equivalent to the SI model, which displays a submodular influence function in temporal networks. For other λ and μ values, violations of the inequality (10) are observed; the maximum frequency of violations is registered for $\mu = 1.00$ and $\lambda = 0.90$.

We repeat a similar analysis on random temporal networks with $N = 100$, $T = 10$, $k = 5$, and $r = 1$. Results are reported in Fig. 3(b). For each run of the SIR model, we select three nodes at random, namely x , y , and v . We compose the sets $\mathcal{A} = \{y\}$ and $\mathcal{B} = \{x, y\}$, and compute Eq. (12). Results in the figure are obtained by averaging $g(\mathcal{A}, \mathcal{B}, v)$ over 40 runs of the SIR model, and over 50 realizations of the random temporal network. The pattern revealed from the figure is similar to one obtained from the exhaustive analysis of Fig. 3(a).




FIG. 4. Violations of the submodularity condition in synthetic temporal networks. In all panels, unless stated otherwise, we consider 10 000 networks composed of $N = 200$ nodes, average degree $k = 5$, total number of temporal layers $T = 10$, and probability of edge shuffle between consecutive layers $r = 0.2$. SIR parameters are $\lambda = \mu = 1.00$. (a) We display g in Eq. (12) as a function of N . (b) We display Eq. (12) as a function of T for different values of N . (c) We display Eq. (12) as a function of k . (d) We display Eq. (12) as a function of r .

Finally, in Figs. 3(c) and 3(d), we report results for two real-world temporal networks. In this specific case, we still select nodes at random, namely, x , y , and v , to compose the sets $\mathcal{A} = \{y\}$ and $\mathcal{B} = \{x, y\}$. We compute Eq. (12) using $Q = 2000$ SIR simulations. For each realization, we randomly select the nodes x , y , and v . The pattern observed is similar to those of the previous two cases, although we observe less violations than in the case of random temporal networks. Also, we observe that the probability to observe a violation of the submodularity inequality is much higher in the supercritical regime of the SIR process on the network is supercritical if $\lambda > \lambda_c(\mu)$, and subcritical for $\lambda < \lambda_c(\mu)$. Here, $\lambda_c(\mu)$ is the critical value of the spreading probability for a given value of the recovery probability μ ; see the Appendix for details.

We systematically study violations of the submodularity condition in random temporal networks. In Fig. 4, we display the frequency of violations for different N , T , k , and r values. In each of our experiments, we first generate a network with a given set of parameters. Then, we select three nodes, namely, x , y , and v , at random. We form the sets $\mathcal{A} = \{y\}$ and $\mathcal{B} = \{x, y\}$, and use these sets together with node v in Eq. (12) to tell whether or not the diminishing returns inequality is violated. Please note that we consider $\mu = \lambda = 1.00$ in this set of experiments, so only one realization of the SIR model is possible. We consider 10 000 networks and record the average number of Eq. (12) over such an ensemble. From Fig. 4, we see that as N increases, the frequency of violation increases. For increasing T , we observe a phase transition from almost no violations to a non-null plateau value. For increasing k , there is a value where the frequency of violations of the diminishing returns property reaches a maximum; instead, when the

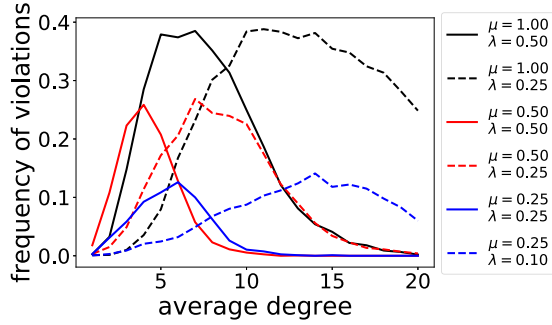


FIG. 5. Violations of the condition for marginal gain in synthetic temporal networks. We measure the frequency of violations of the condition of marginal gain for the influence function using \tilde{g} as defined in Eq. (13). We analyze random temporal networks composed of $N = 100$ nodes and $T = 10$ layers. We consider different values of the average degree k , and of the SIR parameters λ and μ . In the estimation of Eq. (13), we select \mathcal{A} and v randomly, and we average over $Q = 40$ instances of the SIR model. We further take the average over 50 temporal networks.

network is either too sparse or too dense, the frequency drops to zero. As r increases, we see an increase in the violation frequency. Note that for $r = 0$, all the layers of the network are the same, meaning that the network is static and the influence function is submodular.

In Fig. 4, we reshuffle edges with probability $r = 0.2$. However, a similar phenomenology can be obtained by creating layers independently (see Fig. 9). Given the similarity of the results, from now on, we focus our attention on the simpler model where random temporal networks are composed of layers generated independently.

Also, we measure the frequency of violations of the inequality (9) as

$$\tilde{g}(\mathcal{A}, v) = 1 - \frac{1}{Q} \sum_{q=1}^Q H[f_{\mathcal{A}}^{(q)}(v)]. \quad (13)$$

The above equation quantifies how often the addition of the node v to the set \mathcal{A} generates a marginal loss in the influence function. Results of our analysis are reported in Fig. 10. We set $\lambda = \mu = 1.00$, and consider $Q = 1$ SIR simulations; in each simulation, we select at random two nodes, namely, x and v . We set $\mathcal{A} = \{x\}$. We estimate \tilde{g} by taking the average over 10 000 different networks; for each network, we sample \mathcal{A} and v once. We observe that the frequency for observing a marginal loss by adding a node follows a very similar behavior as the frequency of violations of the submodularity inequality; see Fig. 4. We take advantage of this similarity and focus our attention on marginal loss cases from now on, rather than measuring violations of the submodularity inequality. This allows us to alleviate some computational burden without altering the conclusions of the numerical analysis. We can, in fact, safely assume that the pattern of violations of the inequality (9) is similar to the pattern of violations of the condition of Eq. (10).

All the results presented until now used $\lambda = \mu = 1.00$. In Fig. 5, we show results valid for $\lambda < 1.00$ and $\mu \leq 1.00$. We see that the frequencies of inequality violations decrease as the recovery probability μ decreases, indicating that with a




FIG. 6. Violations of the condition for marginal gain in synthetic temporal networks under greedy selection. We measure the frequency of marginal losses using \tilde{g} as defined in Eq. (13) on random temporal networks for different sizes of \mathcal{A} . In all panels, unless stated otherwise, we consider networks composed of $N = 100$ nodes, $T = 10$ layers, and average degree $k = 2.5$. (a) We display Eq. (13) as a function of N . (b) We display Eq. (13) as a function of T . (c) We display Eq. (13) as a function of k .

low recovery probability, it is less likely to observe violations of the inequality (9).

2. Greedy selection of seeds

All results we obtained so far indicate that violations of the inequalities (9) and (10) may occur frequently if tested for sets of randomly chosen nodes. It is, however, natural to ask whether such an observation is also valid when nodes are selected according to the greedy optimization protocol of Eq. (8). In our tests, we measure the marginal gain obtained by adding the m th node at the m th stage of the greedy algorithm. We perform the tests on 2000 random temporal networks for various values of the parameters N , T , and k . SIR simulations are performed for $\lambda = \mu = 1.00$. Results are obtained by averaging Eq. (13) over the various network realizations and are reported in Fig 6. The frequency of violations of the inequality (9) under greedy selection decreases significantly compared to the case where nodes are randomly selected. For the initial stages of the greedy algorithm, the frequency is almost zero; we only start seeing non-null cases of marginal loss in late stages of the algorithm. We do not show the results, but we observe that for $\lambda < 1.00$ and $\mu < 1.00$, the frequency of marginal loss becomes almost zero when selecting up to 10% of all nodes as initial spreaders.

B. Real-world temporal networks

We analyze the frequency of violations of condition of Eq. (9) on real temporal networks. We conduct our analysis on the networks shown in Table I. Results of our analysis are reported in Fig. 7. We observe the same phenomenology as for the case of random temporal networks. If initial spreaders are selected randomly, then the number of configurations for which inequality (9) does not hold is not negligible. On the other hand, if spreaders are chosen according to the greedy strategy, then cases where the inequality (9) is violated are almost nonexistent. This finding suggests that even though the

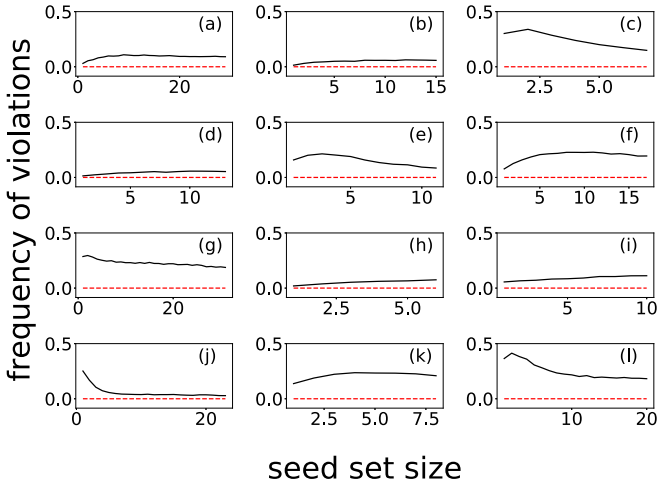


FIG. 7. Violations of the condition for marginal gain in real temporal networks. We measure the frequency of marginal losses using \bar{g} as defined in Eq. (13) on the real-world temporal networks listed in Table I. Results are displayed as full black curves. Labels of the various panels reflect those appearing in the table. The SIR parameters are $\lambda = \mu = 1.00$, so that $Q = 1$ in Eq. (13). We select \mathcal{A} and v randomly 10 000 times and display the average value of the violation of marginal gains. The red dashed curves represent the frequency values when seeds are selected according to greedy optimization.

right conditions to apply greedy optimization are not satisfied, in practice the greedy algorithm might still work as intended, finding solutions close to the ground-truth optimum.

C. Greedy maximization against brute-force optimization

After showing that the greedy algorithm finds solutions characterized by negligible violations of the condition (9), we also want to see how close greedy solutions are to the ground-truth optimal solution. To this end, we apply brute-force optimization to find seed sets of small size in both random and real temporal networks. We then compare these solutions to those obtained using greedy optimization and random selection. Results from all our tests are summarized in Tables II and III; results for a few sample cases are reported in Fig. 8. In general, we observe that the greedy algorithm is almost optimal. In particular, the performance of the greedy algorithm is far above the $1 - 1/e$ bound that would be valid if the influence was indeed a submodular function with non-negative marginal gains. At the same time, we see that random selection performs poorly, generating outbreak sizes well below the optimality bound.

V. CONCLUSION

In spreading processes occurring on temporal networks, the influence function is not a submodular function, and yet greedy optimization provides performances far better than those of other heuristic methods [28]. Here, we investigated the reasons behind such effectiveness of the greedy strategy. We measured violations of the necessary conditions for the submodularity property. We observed that the premature in-

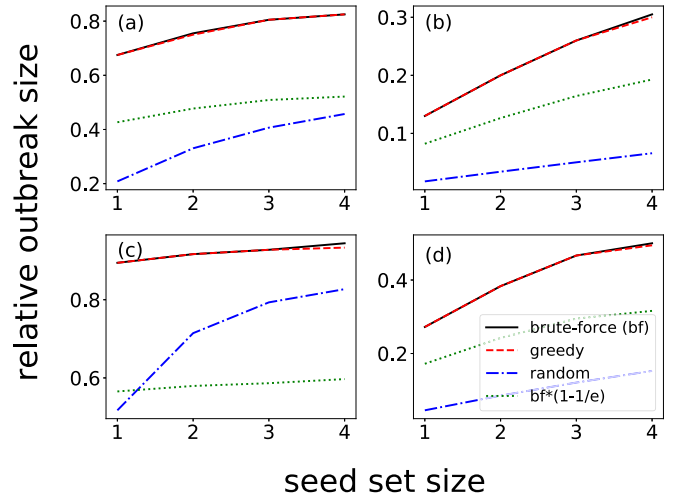


FIG. 8. Optimal selection of influential spreaders in temporal networks. (a) We display the influence function of Eq. (6) as a function of the size of the set of influential spreaders \mathcal{X} . Different methods for the identification of the influential spreaders are used, either brute-force search (black), greedy optimization (red), or random selection (blue). We also display the $1 - 1/e$ bound from the brute-force solution (green). Results are valid for a random temporal network composed of $N = 200$ nodes, $T = 10$ layers, and average degree $k = 1.5$. SIR parameters are $\lambda = \mu = 1.00$. (b) Same as in (a), but for $\lambda = 0.25$ and $\mu = 0.50$. (c) Same as in (a), but for the real-world temporal network “High school, 2012.” (d) Same as in (c), but for $\lambda = 0.10$ and $\mu = 0.25$.

fection and recovery of some nodes may have detrimental effects on the outbreak size, which in turn bring violations of the necessary conditions for the submodularity property. In our systematic analysis, we showed that violations occur frequently if seeds are selected at random. When seeds are selected according to the greedy optimization protocol, we observe that the frequency of violations is much lower in random temporal networks and nonexistent in real-world temporal networks. This finding suggests that even though the function to be optimized does not satisfy the strict definition of a submodular function, in practice it behaves as an effectively submodular function under greedy optimization. As a matter of fact, solutions found by the greedy algorithm are as good as expected for optimization problems involving truly, mathematically speaking, submodular functions. To actually test this hypothesis, we compared greedy solutions to ground-truth solutions in small networks and for small sizes of the seed sets. We observed that in all considered cases, the performance of the greedy algorithm is very close to the optimal solution. On average, the greedy algorithm has a performance of, at worst, 97% in random temporal networks and, at worst, 98% in real-world temporal networks. This is much higher than the optimality guarantee of 63% for the greedy algorithm in static networks and also much better than the performance that can be achieved with randomly selected nodes. The fact that greedy optimization generates quasioptimal solutions to the influence maximization problem on temporal networks is associated with its ability of avoiding the selection of seeds whose premature recovery would block future infection paths. We believe that this ability is not related to the specific pro-

tol of optimization, rather to the fact that the optimization procedure relies on direct measurements of the influence function. We expect that any other optimization algorithm that uses dynamical information should be able to achieve similar performance by learning about the blocking effect of some nodes via measurements of the influence function. Also, the great effectiveness displayed by the greedy optimization strategy might indicate the effective existence of a tight lower bound on the performance of the greedy algorithm in solving the influence maximization problem on temporal networks. We leave to future research the challenging task to formulate a theory for such a performance guarantee.

ACKNOWLEDGMENTS

Ş.E. and F.R. acknowledge support by the Air Force Office of Scientific Research (Grant No. FA9550-21-1-0446); F.R. acknowledges support by the Army Research Office (Grant No. W911NF-21-1-0194). The funders had no role in study design, data collection, and analysis, decision to publish, or any opinions, findings, and conclusions or recommendations expressed in the manuscript.

APPENDIX A: FINDING THE CRITICAL THRESHOLD

The critical threshold value λ_c of a temporal network is a function of the recovery probability μ , denoted as $\lambda_c(\mu)$. In order to estimate $\lambda_c(\mu)$, we start the spreading process from each single node in the first layer of the temporal network and calculate the outbreak size. We repeat this process 500 times for each node for a range of values for the infection probability λ . The λ value that gives the maximum for the ratio between the standard deviation and mean of the outbreak sizes is selected as the critical threshold value λ_c for a given recovery probability μ .




FIG. 9. Violations of the submodularity condition on synthetic temporal networks. Frequency of violations of the inequality (10) on synthetic network models with $N = 100$, $T = 10$, $k = 5$ unless stated otherwise, and SIR parameters $\lambda = \mu = 1.00$. (a) Equation (12) as a function of N . Results obtained for $r = 1.0$ (black curve) are compared to the results obtained when layers in the temporal network are created independently (red curve). (b) Same as in (a), but we display Eq. (12) as a function of T . (c) Same as in (a), but for Eq. (12) as a function of k .

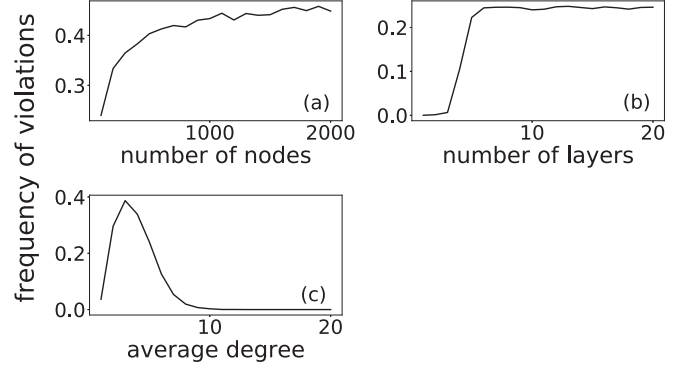


FIG. 10. Violations of the condition for marginal gain in synthetic temporal networks. Frequency of marginal loss cases with random seeds on random temporal networks for $N = 100$, $T = 10$, $k = 5$ unless stated otherwise, and SIR parameters $\lambda = \mu = 1.00$. (a) Equation (13) as a function of N . (b) Equation (13) as a function of T . (c) Equation (13) as a function of k .

TABLE II. Greedy selection of optimal spreaders in real temporal networks. We report the average outbreak size of the seeds found by the greedy algorithm on various networks relative to the optimal solution found with brute-force optimization. The results are averaged over all real-world temporal networks listed in Table I. We excluded “Email, dept. 1” and “High school, 2013” due to their size.

λ	μ	$ \mathcal{A} = 2$	$ \mathcal{A} = 3$	$ \mathcal{A} = 4$	$ \mathcal{A} = 5$
1.00	1.00	99.0%	98.8%	98.4%	98.2%
0.50	1.00	99.6%	99.3%	99.1%	98.7%
0.25	1.00	100.0%	99.3%	98.9%	98.5%
0.50	0.50	99.7%	99.6%	99.5%	99.4%
0.25	0.50	99.6%	99.6%	98.9%	98.8%
0.10	0.50	98.8%	98.3%	98.6%	98.5%
0.25	0.25	98.9%	99.3%	99.1%	99.1%
0.10	0.25	99.6%	99.1%	99.3%	99.3%
0.05	0.25	99.2%	99.0%	98.6%	98.7%

TABLE III. Greedy selection of optimal spreaders in synthetic temporal networks. We report the average outbreak size of the seeds found by the greedy algorithm on random temporal networks relative to the optimal solution found with brute-force optimization. The results are averaged over random temporal networks created with all combinations of the parameters $N = 200$, $T \in \{5, 10\}$, and $k \in \{1.3, 1.4, 1.5, 1.6, 1.7, 2.0, 2.5\}$.

λ	μ	$ \mathcal{A} = 2$	$ \mathcal{A} = 3$	$ \mathcal{A} = 4$
1.00	1.00	99.2%	98.4%	98.1%
0.50	1.00	97.9%	97.6%	97.1%
0.25	1.00	100.0%	100.0%	100.0%
0.50	0.50	99.6%	99.4%	98.7%
0.25	0.50	99.7%	99.3%	98.9%
0.10	0.50	100.0%	100.0%	100.0%
0.25	0.25	99.4%	99.4%	99.4%
0.10	0.25	99.7%	99.8%	99.8%
0.050	0.25	100.0%	100.0%	100.0%

APPENDIX B: ADDITIONAL RESULTS

In Fig. 9, we report the frequency of violations of the inequality (10) on synthetic network models for different values of the parameters N , T , and k . The spreading process is initiated by randomly selected seeds. Specifically, we compare the results for a temporal network created with $r = 0.2$ against a network created with independent layers. In Fig. 10, we show

the results for the frequency of violations of the inequality (9) on synthetic network models for different values of the parameters N , T , and k . Also here, the spreading process is initiated by randomly selected seeds. In Table II, we report the average performance of the greedy algorithm relative to brute-force optimization in real-world temporal networks. In Table III, we report the same results for random temporal networks.

-
- [1] P. Domingos and M. Richardson, Mining the network value of customers, in *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, California, USA* (ACM, New York, 2001), pp. 57–66.
- [2] D. Kempe, J. Kleinberg, and É. Tardos, Maximizing the spread of influence through a social network, in *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Washington, D.C.* (ACM, New York, 2003), pp. 137–146.
- [3] J. Leskovec, A. Krause, C. Guestrin, C. Faloutsos, J. VanBriesen, and N. Glance, Cost-effective outbreak detection in networks, in *Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Jose, California, USA* (ACM, New York, 2007), pp. 420–429.
- [4] W. Chen, Y. Wang, and S. Yang, Efficient influence maximization in social networks, in *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Paris, France* (ACM, New York, 2009), pp. 199–208.
- [5] H. T. Nguyen, M. T. Thai, and T. N. Dinh, Stop-and-stare: Optimal sampling algorithms for viral marketing in billion-scale networks, in *Proceedings of the 2016 International Conference on Management of Data, San Francisco, California, USA* (ACM, New York, 2016), pp. 695–710.
- [6] A. Goyal, W. Lu, and L. V. Lakshmanan, Celf++ optimizing the greedy algorithm for influence maximization in social networks, in *Proceedings of the 20th International Conference Companion on World Wide Web, Hyderabad, India* (ACM, New York, 2011), pp. 47–48.
- [7] Y. Tang, Y. Shi, and X. Xiao, Influence maximization in near-linear time: A martingale approach, in *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data, Melbourne, Victoria, Australia* (ACM, New York, 2015), pp. 1539–1554.
- [8] M. Minutoli, M. Halappanavar, A. Kalyanaraman, A. Sathanur, R. McClure, and J. McDermott, Fast and scalable implementations of influence maximization algorithms, in *2019 IEEE International Conference on Cluster Computing (CLUSTER), Albuquerque, NM, USA* (IEEE, Piscataway, NJ, 2019), pp. 1–12.
- [9] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher, An analysis of approximations for maximizing submodular set functions-i, *Math. Program.* **14**, 265 (1978).
- [10] Ş. Erkol, C. Castellano, and F. Radicchi, Systematic comparison between methods for the detection of influential spreaders in complex networks, *Sci. Rep.* **9**, 15095 (2019).
- [11] M. Kitsak, L. K. Gallos, S. Havlin, F. Liljeros, L. Muchnik, H. E. Stanley, and H. A. Makse, Identification of influential spreaders in complex networks, *Nat. Phys.* **6**, 888 (2010).
- [12] G. F. de Arruda, A. L. Barbieri, P. M. Rodríguez, F. A. Rodrigues, Y. Moreno, and L. F. Costa, Role of centrality for the identification of influential spreaders in complex networks, *Phys. Rev. E* **90**, 032812 (2014).
- [13] F. Morone and H. A. Makse, Influence maximization in complex networks through optimal percolation, *Nature (London)* **524**, 65 (2015).
- [14] L. Lü, D. Chen, X.-L. Ren, Q.-M. Zhang, Y.-C. Zhang, and T. Zhou, Vital nodes identification in complex networks, *Phys. Rep.* **650**, 1 (2016).
- [15] P.-L. Yang, G.-Q. Xu, Q. Yu, and J.-W. Guo, An adaptive heuristic clustering algorithm for influence maximization in complex networks, *Chaos: Interdis. J. Nonlin. Sci.* **30**, 093106 (2020).
- [16] P. Holme and J. Saramäki, Temporal networks, *Phys. Rep.* **519**, 97 (2012).
- [17] B. A. Prakash, H. Tong, N. Valler, M. Faloutsos, and C. Faloutsos, Virus propagation on time-varying networks: Theory and immunization algorithms, *Machine Learning and Knowledge Discovery in Databases*, edited by J. L. Balczar, F. Bonchi, A. Gionis, and M. Sebag, Vol. 6323 (Springer, New York, 2010), pp. 99–114.
- [18] M. Karsai, M. Kivela, R. K. Pan, K. Kaski, J. Kertész, A.-L. Barabási, and J. Saramäki, Small but slow world: How network topology and burstiness slow down spreading, *Phys. Rev. E* **83**, 025102(R) (2011).
- [19] N. Perra, B. Gonçalves, R. Pastor-Satorras, and A. Vespignani, Activity driven modeling of time varying networks, *Sci. Rep.* **2**, 469 (2012).
- [20] E. Valdano, L. Ferreri, C. Poletto, and V. Colizza, Analytical Computation of the Epidemic Threshold on Temporal Networks, *Phys. Rev. X* **5**, 021005 (2015).
- [21] T. Murata and H. Koga, Extended methods for influence maximization in dynamic networks, *Comput. Social Networks* **5**, 8 (2018).
- [22] M. Han, M. Yan, Z. Cai, Y. Li, X. Cai, and J. Yu, Influence maximization by probing partial communities in dynamic online social networks, *Transact. Emerg. Telecommun. Technol.* **28**, e3054 (2017).
- [23] H. Zhuang, Y. Sun, J. Tang, J. Zhang, and X. Sun, Influence maximization in dynamic social networks, in *2013 IEEE 13th International Conference on Data Mining* (IEEE, New York, 2013), pp. 1313–1318.
- [24] R. Michalski, J. Jankowski, and P. Bródka, Effective influence spreading in temporal networks with sequential seeding, *IEEE Access* **8**, 151208 (2020).

- [25] S. Osawa and T. Murata, Selecting seed nodes for influence maximization in dynamic networks, *Complex Networks VI*, edited by G. Mangioni, F. Simini, S. Uzzo, and D. Wang (Springer, New York, 2015), pp. 91–98.
- [26] R. Michalski, T. Kajdanowicz, P. Bródka, and P. Kazienko, Seed selection for spread of influence in social networks: Temporal vs. static approach, *New Gener. Comput.* **32**, 213 (2014).
- [27] N. T. Gayraud, E. Pitoura, and P. Tsaparas, Diffusion maximization in evolving social networks, in *Proceedings of the 2015 ACM Conference on Online Social Networks* (2015), pp. 125–135.
- [28] Ş. Erkol, D. Mazzilli, and F. Radicchi, Influence maximization on temporal networks, *Phys. Rev. E* **102**, 042307 (2020).
- [29] A. Paranjape, A. R. Benson, and J. Leskovec, Motifs in temporal networks, in *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining, Cambridge, United Kingdom* (ACM, New York, 2017), pp. 601–610.
- [30] J. Fournet and A. Barrat, Contact patterns among high school students, *PLoS One* **9**, e107878 (2014).
- [31] R. Mastrandrea, J. Fournet, and A. Barrat, Contact patterns in a high school: A comparison between data collected using wearable sensors, contact diaries and friendship surveys, *PLoS One* **10**, e0136497 (2015).
- [32] P. Vanhems, A. Barrat, C. Cattuto, J.-F. Pinton, N. Khanafer, C. Régis, B.-A. Kim, B. Comte, and N. Voirin, Estimating potential infection transmission routes in hospital wards using wearable proximity sensors, *PLoS One* **8**, e73970 (2013).
- [33] L. Isella, J. Stehlé, A. Barrat, C. Cattuto, J.-F. Pinton, and W. Van den Broeck, What’s in a crowd? Analysis of face-to-face behavioral networks, *J. Theor. Biol.* **271**, 166 (2011).
- [34] V. Gemmetto, A. Barrat, and C. Cattuto, Mitigation of infectious disease at school: Targeted class closure vs school closure, *BMC Infect. Dis.* **14**, 695 (2014).
- [35] J. Stehlé, N. Voirin, A. Barrat, C. Cattuto, L. Isella, J.-F. Pinton, M. Quaggiotto, W. Van den Broeck, C. Régis, B. Lina *et al.*, High-resolution measurements of face-to-face contact patterns in a primary school, *PLoS One* **6**, e23176 (2011).
- [36] M. Génois, C. L. Vestergaard, J. Fournet, A. Panisson, I. Bonmarin, and A. Barrat, Data on face-to-face contacts in an office building suggest a low-cost vaccination strategy based on community linkers, *Network Sci.* **3**, 326 (2015).
- [37] M. Génois and A. Barrat, Can co-location be used as a proxy for face-to-face contacts? *EPJ Data Sci.* **7**, 11 (2018).
- [38] R. Pastor-Satorras, C. Castellano, P. Van Mieghem, and A. Vespignani, Epidemic processes in complex networks, *Rev. Mod. Phys.* **87**, 925 (2015).
- [39] R. Santiago and Y. Yoshida, Weakly submodular function maximization using local submodularity ratio, [arXiv:2004.14650](https://arxiv.org/abs/2004.14650).