Maximizing the Probability of Fixation in the Positional Voter Model

Petros Petsinis, Andreas Pavlogiannis, Panagiotis Karras

Aarhus University, Denmark
{petsinis,pavlogiannis,panos}@cs.au.dk

Abstract

The Voter model is a well-studied stochastic process that models the invasion of a *novel trait* A (e.g., a new opinion, social meme, genetic mutation, magnetic spin) in a network of individuals (agents, people, genes, particles) carrying an existing *resident trait* B. Individuals change traits by occasionally sampling the trait of a neighbor, while an *invasion bias* $\delta \geq 0$ expresses the stochastic preference to adopt the novel trait A over the resident trait B. The strength of an invasion is measured by the probability that eventually the whole population adopts trait A, i.e., the *fixation probability*. In more realistic settings, however, the invasion bias is not ubiquitous, but rather manifested only in parts of the network. For instance, when modeling the spread of a social trait, the invasion bias represents *localized* incentives. In this paper, we generalize the standard biased Voter model to the *positional* Voter model, in which the invasion bias is effectuated only on an arbitrary subset of the network nodes, called *biased nodes*. We study the ensuing optimization problem, which is, given a budget k, to choose k biased nodes so as to maximize the fixation probability of a randomly occurring invasion. We show that the problem is $\mathbf{NP}$-hard both for finite δ and when $\delta \to \infty$ (strong bias), while the objective function is not submodular in either setting, indicating strong computational hardness. On the other hand, we show that, when $\delta \to 0$ (weak bias), we can obtain a tight approximation in $\mathcal{O}(n^{2\omega})$ time, where ω is the matrix-multiplication exponent. We complement our theoretical results with an experimental evaluation of some proposed heuristics.

Introduction

Several real-world phenomena involve the emergence and spread of novel traits in populations of interacting individuals of various kinds. For example, such phenomena may concern the propagation of new information in a social network, the sweep of a novel mutation in a genetically homogeneous population, the rise and resolution of spatial conflict, and the diffusion of atomic properties in interacting particle systems. Network science collectively studies such phenomena as *diffusion processes*, whereby the system of interacting individuals is represented as a network and the corresponding process defines the (stochastic, in general) dynamics of local trait spread, from an individual to its neighbors. These dynamics can vary drastically from one application domain to another, and have been studied extensively, e.g., in the cases of influence and information cascades in social networks (Domingos and Richardson 2001; Kempe, Kleinberg, and Tardos 2003; Mossel and Roch 2007; Zhang et al. 2020), epidemic spread (Kermack, McKendrick, and Walker 1927; Newman 2002), genetic variation in structured populations (Moran 1958; Lieberman, Hauert, and Nowak 2005; Pavlogiannis et al. 2018; Tkadlec et al. 2021), and game-theoretic models of antagonistic interaction (Ohtsuki et al. 2006; Ibsen-Jensen, Chatterjee, and Nowak 2015).

One of the most fundamental diffusion processes is the *Voter model*. While introduced to study particle interactions (Liggett and Liggett 1985) and territorial conflict (Clifford and Sudbury 1973), its elegance and simplicity have rendered it applicable to a wide variety of other domains (often under other names, such as *imitation updating* or *death-birth dynamics*), including the spread of social traits (Even-Dar and Shapira 2011; Castellano, Fortunato, and Loreto 2009; Bhat and Redner 2019; Durocher et al. 2022), robot coordination and swarm intelligence (Talamali et al. 2021), and evolutionary dynamics (Antal, Redner, and Sood 2006; Ohtsuki et al. 2006; Hindersin and Traulsen 2015; Allen et al. 2017; Tkadlec et al. 2020; Allen et al. 2020).

The voter process starts with a homogeneous population of agents (aka *voters*) scattered over an interaction network and carrying a *resident trait* B. A *novel trait* A invades the population by initially appearing on some random agent(s), and gets diffused in the network by local stochastic updates: each agent occasionally wakes up and updates its own trait by randomly sampling the trait of a neighbor. In general, the trait A is associated with an *invasion bias* $\delta \geq 0$, which quantifies the stochastic preference of an agent to choose A over B while sampling the traits of its neighbors. The process eventually reaches an *absorption state* (aka *consensus state* (Even-Dar and Shapira 2011)), in which A either *fixates* in the population or *goes extinct*. The key quantity of interest is the probability of *fixation*, which depends on the network structure and the bias δ. A key difference with standard cascade models (Kempe, Kleinberg, and Tardos 2005) is that the Voter process is *non-progressive*, meaning that the spread of A can both increase and shrink over time. It can thus express settings such as switching of opinions.

In realistic situations, the invasion bias is not ubiquitous throughout the network, but rather present only in parts. For instance, in the spread of social traits, the invasion bias represents incentives that are naturally local to certain members of the population. Similarly, in the diffusion of magnetic spins, the invasion bias typically stands for the presence of a magnetic field that is felt locally in certain areas. In this paper, we generalize the standard Voter model to faithfully express the locality of such effects, leading to the *positional Voter model*, in which the invasion bias is only present in a (arbitrary) subset of network nodes, called *biased nodes*, and analyze the computational properties thereof.

The positional Voter model gives rise to the following optimization problem: given a budget k, which k nodes should we bias to maximize the fixation probability of the novel trait? This problem has a natural modeling motivation. For instance, ad placement can tune consumers more receptive to viral (word-of-mouth) marketing (Barbieri, Bonchi, and Manco 2013; Barbieri and Bonchi 2014; Zhang et al. 2020). Feature selection of a product can increase its appeal to strategic agents and hence maximize its dissemination (Ivanov et al. 2017). Expert placement can strengthen the robustness of a social network under adversarial attacks (Alon et al. 2015). Nutrient placement can be utilized to increase the territorial spread of a certain organism in ecological networks (Brendborg et al. 2022). In all these settings, optimization can be modeled as selecting a biased set of network nodes/agents, on which one trait has a propagation advantage over the other.

Our contributions. In this paper we introduce the positional Voter model, and studied the associated optimization problem. Our main results are as follows.

1. We show that computing the fixation probability on undirected networks admits a fully-polynomial-time approximation scheme (FPRAS, Theorem 1).
2. On the negative side, we show that the optimization problem is $\mathbf{NP}$-hard both for finite δ (general setting) and as $\delta \to \infty$ (strong bias, Theorem 2), while the objective function is not submodular in either setting (Theorem 3).
3. We show that, when the network has self-loops (capturing the effect that an individual might choose to remain to their trait), the objective function becomes submodular as $\delta \to \infty$, hence the optimization problem can be efficiently approximated within a factor $1 - 1/e$ (Theorem 4).
4. We show that, as $\delta \to 0$ (weak bias), we can obtain a tight approximation in $\mathcal{O}(n^{2\omega})$ time, where ω is the matrix-multiplication exponent (Theorem 5).
5. Lastly, we propose and experimentally evaluate a number of heuristics for maximizing the fixation probability.

Due to limited space, some proofs appear in the full version of this paper (Petsinis, Pavlogiannis, and Karras 2022).

Preliminaries

Network structures. We consider a population of n agents spread over the nodes of a (generally, directed and weighted) graph $G = (V, E, w)$, where V is a set of n nodes, $E \subseteq V \times V$ is a set of edges capturing interactions between the agents, and $w\colon E \to \mathbb{R}_{>0}$ is a weight function mapping each edge (u, v) to a real number $w(u, v)$ denoting the strength of interaction of u with v. We denote by $\mathrm{in}(u) = \{v \in V : (v, u) \in E\}$ the set of incoming neighbors of node u. The *(in)-degree* of node u is $\mathrm{d}(u) = 1/|\mathrm{in}(u)|$. For the Voter process to be well-defined on G, G should be *strongly connected*. In some cases, we consider *undirected* and *unweighted* graphs, meaning that (i) E is symmetric; and (ii) for all $(u, v) \in E$, $w(u, v) = 1$; this setting captures networks in which interactions are bidirectional and each node v is equally influenced by each neighbor.

The standard Voter model. A *configuration* $X \subseteq V$ represents the set of agents that carry the trait A. The trait A is associated with an *invasion bias* $\delta \geq 0$ (with $\delta = 0$ denoting no bias, or the *neutral setting*). Given a configuration X, the *influence strength* of node v is defined as

$$f_X(v) = \begin{cases} 1 + \delta, & \text{if } v \in X \\ 1, & \text{otherwise.} \end{cases} \tag{1}$$

Let $\mathcal{X}_t \subseteq V$ be a random variable representing the configuration at time t. The Voter (or *death-birth*) process is a discrete-time stochastic process $\{\mathcal{X}_t\}$, $t \geq 0$ that models the invasion of a *novel trait* A on a homogeneous population of agents carrying a *resident trait* B, scattered over the nodes of a graph $G = (V, E, w)$. Initially, the trait A appears uniformly at random on one agent, i.e., $\mathbb{P}[\mathcal{X}_0 = \{u\}] = 1/n$ $\forall u \in V$. Given the configuration $\mathcal{X}_t = X$ at time t, the next configuration $\mathcal{X}_{t+1}$ at time $t+1$ is determined by a sequence of two stochastic events:

1. *Death:* an agent u dies with probability $1/n$.
2. *Birth:* a neighbor $v \in \mathrm{in}(u)$ sets u's trait with probability

$$\frac{f_X(v) \cdot w(v, u)}{\sum_{x \in \mathrm{in}(u)} f_X(x) \cdot w(x, u)} \tag{2}$$

In effect, the set of agents carrying A may grow or shrink at any given step. In general, we may have $(u, u) \in E$, expressing the event that agent u stays at its current trait.

The positional Voter model. To capture cases where the bias is manifested *in parts* of the population, we extend the standard Voter model to its *positional* variant as follows:

1. The network structure comprises a component $S \subseteq V$, the subset of agents that are biased to the invasion of A.
2. The influence strength of agent v is *conditioned* on whether the neighbor u that v is attempting to influence is biased. Formally, we replace $f_X(v)$ with:

$$f_X^S(v|u) = \begin{cases} 1 + \delta, & \text{if } v \in X \text{ and } u \in S \\ 1, & \text{otherwise.} \end{cases} \tag{3}$$

We retrieve the standard Voter model by setting $S = V$, i.e., uniformly present invasion bias. Fig. 1 shows the process.

Fixation. In the long run, the Voter process reaches a homogeneous state almost surely; that is, with probability 1, there exists a time t such that $\mathcal{X}_t \in \{\emptyset, V\}$. If $\mathcal{X}_t = V$, the novel trait has *fixated* in G, otherwise it has *gone extinct*. Given a configuration X, a biased set S and a bias δ, we denote the probability that A fixates when starting from nodes X as:

$$\mathrm{fp}(G^S, \delta, X) = \mathbb{P}[\exists t \geq 0 \colon \mathcal{X}_t = V \mid \mathcal{X}_0 = X]$$

As the invasion starts on a random node, the *fixation probability* of A in G is the average fixation probability over V:

$$\mathrm{fp}(G^S, \delta) = \frac{1}{n} \sum_{u \in V} \mathrm{fp}(G^S, \delta, \{u\}) \,.$$

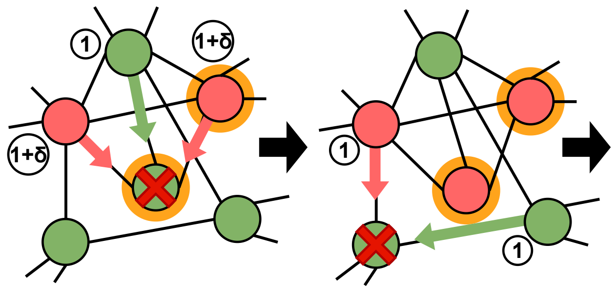

Figure 1: Two steps of the positional Voter process. Red (resp., green) nodes carry the mutant trait A (resp., resident trait B); orange circles mark biased nodes.

When $\delta = 0$, the set S is inconsequential, hence the positional Voter model reduces to the standard model, for which $\mathrm{fp}(G^S, \delta) = {}^1\!/_n$ (McAvoy and Allen 2021). When G is undirected and $S = V$ (i.e., in the standard Voter model), $\mathrm{fp}(G^S, \delta)$ admits a fully polynomial randomized approximation scheme (FPRAS) via Monte-Carlo simulations (Durocher et al. 2022). For arbitrary graphs, however (that may have directed edges or non-uniform edge weights), the complexity of computing $\mathrm{fp}(G^S, \delta)$ is open. We later show that $\mathrm{fp}(G^S, \delta)$ admits an FPRAS for undirected graphs even under the positional Voter model (i.e., for *arbitrary* S).

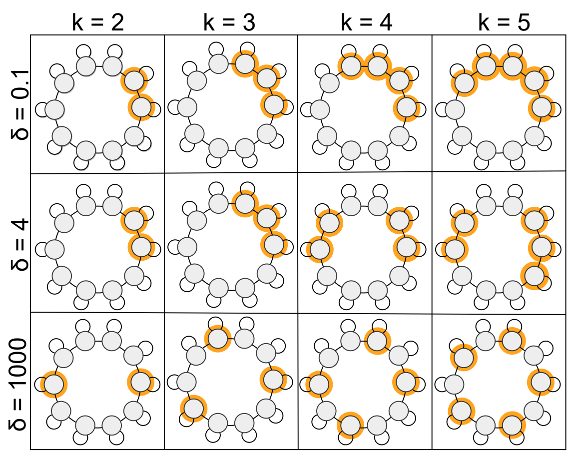

Figure 2: Optimal biased node sets S (in orange) on a cycle graph for different values of bias δ and budget k; for small δ (top), the optimal biased nodes are consecutive, yet for large δ (bottom), they are spaced apart. For intermediate values (middle), the optimal strategy varies depending on k.

Optimization. In the positional Voter model, $\mathrm{fp}(G^S, \delta)$ depends not only on the network structure G and invasion bias δ, but also on the set of biased nodes S. As Fig. 2 illustrates, this dependency can be quite intricate, as the optimal strategy for choosing S may vary depending on its size and the value of δ. In effect, the positional Voter model naturally gives rise to the following *optimization* problem: Given a budget $k \in \mathbb{N}$, which subset $S \subseteq V$ of at most k nodes should we bias, to maximize the fixation probability? Formally, we seek a set $S^* \subseteq V$ such that

$$S^* = \underset{S \subseteq V, |S| \leq k}{\arg\max}\ \mathrm{fp}(G^S, \delta)\ . \tag{4}$$

As we will show, $\mathrm{fp}(G^S, \delta)$ is monotonic on S for all δ, thus the condition $|S| \leq k$ reduces to $|S| = k$. We also consider the two extreme cases of $\delta \to \infty$ (*strong bias*) and $\delta \to 0$ (*weak bias*). In the case of strong bias, we define

$$\mathrm{fp}^\infty(G^S) = \lim_{\delta \to \infty} \mathrm{fp}(G^S, \delta),$$

thus the optimization objective of Eq. (4) becomes

$$S^* = \underset{S \subseteq V, |S| = k}{\arg\max}\ \mathrm{fp}^\infty(G^S)\ . \tag{5}$$

In the case of weak bias, if $\delta = 0$, we have $\mathrm{fp}(G^S, 0) = {}^1\!/_n$, hence the fixation probability is independent of S. However, as $\delta \to 0$ but remains positive, different biased sets will yield different fixation probability. In this case we work with the Taylor expansion of $\mathrm{fp}(G^S, \delta)$ around 0, and write

$$\mathrm{fp}(G^S, \delta) = \frac{1}{n} + \delta \cdot \mathrm{fp}'(G^S, 0) + \mathcal{O}(\delta^2), \tag{6}$$

where $\mathrm{fp}'(G^S, 0) = \frac{\mathrm{d}}{\mathrm{d}\,\delta}\Big|_{\delta=0} \mathrm{fp}(G^S, \delta)$. As $\delta \to 0$, the lower-order terms $\mathcal{O}(\delta^2)$ approach 0 faster than the second term. Hence, for sufficiently small positive bias δ, the optimal placement of bias in the network is the one maximizing the derivative $\mathrm{fp}'(G^S, 0)$. In effect, the optimization objective for weak invasion bias becomes

$$S^* = \underset{S \subseteq V, |S| = k}{\arg\max}\ \mathrm{fp}'(G^S, 0)\ . \tag{7}$$

Fig. 3 illustrates the behavior of $\mathrm{fp}(G^S, \delta)$ for $k = 2$ and various values of δ on a small graph.

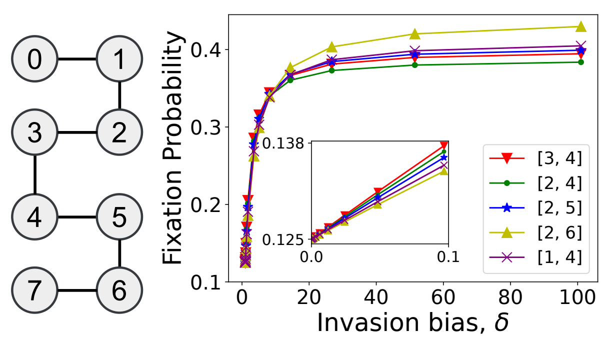

Figure 3: Fixation probabilities for different biased sets of size $k = 2$; when $\delta \in [0, 0.1]$ the fixation probability is roughly linearly dependent on δ; the set $\{2, 6\}$ is worst for small δ but best for large δ.

Computing the Fixation Probability

In this section we compute the fixation probability $\mathrm{fp}(G^S, \delta)$ for undirected G. In a key step, we show that the expected time $\mathrm{T}(G^S, \delta)$ until the positional Voter process reaches a homogeneous state is polynomial in n. In particular, we obtain the following lemma that generalizes a similar result for the standard Voter model (Durocher et al. 2022).

Lemma 1. *Given an undirected graph G with n nodes, a set $S \subseteq V$ and some $\delta \geq 1$, we have* $\mathrm{T}(G^S, \delta) \leq n^5$.

In effect, we obtain an FPRAS by simulating the process sufficiently many times, and reporting the empirical average number of fixation occurrences as the fixation probability. We thus arrive at the following theorem.

Theorem 1. *Given a connected undirected graph $G = (V, E)$, a set $S \subseteq V$ and a real number $\delta \geq 0$, the function* $\mathrm{fp}(G^S, \delta)$ *admits a FPRAS.*

Hardness of Optimization

We now study the optimization problem for the positional Voter model, and show that it is $\mathbf{NP}$-hard. We first examine the process with strong bias, $\delta \to \infty$, running on undirected, regular graphs where nodes have self-loops. Our first observation is that, due to self-loops, if the process reaches a configuration X with $X \cap S \neq \emptyset$, then fixation is guaranteed.

Lemma 2. *Consider an undirected graph G with self-loops and biased set S, and let X be an arbitrary configuration. If $X \cap S \neq \emptyset$, then* $\mathrm{fp}^{\infty}(G^S, X) = 1$.

Proof. Consider any node $u = X \cap S$. For any two configurations X_1 and X_2 with $u \in X_1$ and $u \notin X_2$, we have

$$\mathbb{P}[\mathcal{X}_{t+1} = X_2 | \mathcal{X}_t = X_1] \leq \frac{1}{n}\frac{1}{\delta} \to 0 \quad \text{as} \quad \delta \to \infty$$

On the other hand, with probability at least $(1/n)^n > 0$, the process reaches fixation within n steps when starting from any non-empty configuration. Indeed, while $\emptyset \subset X \subset V$, with probability at least $1/n$, a B-node u with an A-neighbor is chosen for replacement, and such A-neighbor propagates its trait to u with probability 1 as $\delta \to \infty$. Thus, the probability that we reach fixation before we reach a configuration X_2 with $X_2 \cap S = \emptyset$ approaches 1 as $\delta \to \infty$. ☐

Our second observation relies on Lemma 2 to argue that, for an undirected, regular graph G with self-loops, when the budget k is sufficiently large, the optimal choice for the biased set S forms a vertex cover of G.

Lemma 3. *Let $G = (V, E)$ be an undirected, d-regular graph with self-loops, and $S \subseteq V$ a biased set. Then* $\mathrm{fp}^{\infty}(G^S) \geq \frac{|S|/n+d}{1+d}$ *iff S is a vertex cover of G.*

Proof. Since G is connected and has self-loops, we have $d \geq 2$. Due to the uniform initial placement of trait A, the probability that it lands on a biased node is $|S|/n$. Let Y be the set of nodes in $V \setminus S$ that have at least one neighbor *not* in S. By Lemma 2 we have

$$\mathrm{fp}^{\infty}(G^S) = \frac{|S|}{n} + \frac{n - |S| - |Y|}{n} p + \frac{1}{n}\sum_{u \in Y} \mathrm{fp}^{\infty}(G^S, \{u\}),$$

where p is the probability that a node $u \in V \setminus S$ with initial trait A and whose neighbors are all in S propagates A to any of those neighbors before replacing its own trait. Let p_1 and p_2 be the probabilities that u propagates to, and gets replaced by, any of its $d-1$ neighbors, respectively, and we have $p = \frac{p_1}{p_1+p_2}$. Moreover, $p_1 = \frac{d-1}{n}1$ (once a neighbor v dies, u propagates its trait to v with probability 1), and $p_2 = \frac{1}{n}\frac{d-1}{d}$, leading to $p = \frac{d}{d+1}$. If S is a vertex cover of G, then $Y = \emptyset$, hence $\mathrm{fp}^{\infty}(G^S) = \frac{|S|/n+d}{1+d}$. On the other hand, if S is not a vertex cover of G, then $|Y| \geq 1$. Consider a node $u \in Y$ and let $v \in \mathrm{in}(u) \setminus S$. Observe that $\mathrm{fp}^{\infty}(G^S, \{u, v\}) < 1$, since with probability at least $(\frac{1}{n}\frac{1}{d})^2$ the traits on u and v get successively replaced by trait B. It follows that $\mathrm{fp}^{\infty}(G^S) < \frac{|S|/n+d}{1+d}$, as desired. ☐

Since vertex cover is $\mathbf{NP}$-hard on regular graphs (Feige 2003), Lemma 3 implies $\mathbf{NP}$-hardness for maximizing $\mathrm{fp}^{\infty}(G^S)$: given a budget k, G has a vertex cover of size k iff $\max_{S \subseteq V, |S|=k} \mathrm{fp}^{\infty}(G^S) = \frac{\frac{|S|}{n}+d}{1+d}$. Moreover, the continuity of $\mathrm{fp}(G^S, \delta)$ as a function of δ implies hardness for finite δ too. We thus arrive at the following theorem.

Theorem 2. *The problem of maximizing* $\mathrm{fp}(G^S, \delta)$ *and* $\mathrm{fp}^{\infty}(G^S)$ *in the positional Voter model is* $\mathbf{NP}$*-hard.*

Monotonicity and Submodularity

We now turn our attention to the monotonicity and (conditional) submodularity properties of the fixation probability. Our first lemma formally establishes the intuition that, as we increase the set of biased nodes S or the invasion bias δ, the chances that the novel trait A fixates do not worsen.

Lemma 4. *For a graph G, biased sets $S_1, S_2 \subseteq V$ with $S_1 \subseteq S_2$ and biases $\delta_1, \delta_2 \geq 0$ with $\delta_1 \leq \delta_2$, we have* $\mathrm{fp}(G^{S_1}(\delta_1)) \leq \mathrm{fp}(G^{S_2}(\delta_2))$.

Proof. Consider the two respective Voter processes $\mathcal{M}_1 = \mathcal{X}_0^1, \mathcal{X}_1^1, \dots$ and $\mathcal{M}_2 = \mathcal{X}_0^2, \mathcal{X}_1^2, \dots$. We establish a coupling between $\mathcal{M}_1$ and $\mathcal{M}_2$ that guarantees $\mathcal{X}_t^1 \subseteq \mathcal{X}_t^2$, for all t, which proves the lemma.

Indeed, assume that the two processes are in configurations $\mathcal{X}_t^1 = X^1$ and $\mathcal{X}_t^2 = X_2$, with $X_1 \subseteq X_2$. We choose the same node u to be replaced in the two processes, uniformly at random. Observe that, since $S_1 \subseteq S_2$ and $\delta_1 \leq \delta_2$, the probability p_2 that v is replaced by an A-neighbor in $\mathcal{M}_2$ is at least as large as the corresponding probability p_1 in $\mathcal{M}_1$. Thus, with probability $p_2 - p_1 \geq 0$, we replace u in $\mathcal{M}_2$ with one of its A-neighbors, leading to a configuration $X_2' \supseteq X_2 \supseteq X_1$. With probability p_1, we replace u in both $\mathcal{M}_1$ and $\mathcal{M}_2$ with one of its A-neighbors, leading to configurations X_1' and X_2' with $X_2' \supseteq X_1'$. Lastly, with probability $1 - p_1$, we replace u in both $\mathcal{M}_1$ and $\mathcal{M}_2$ with one of its B-neighbors, again leading to configurations X_1' and X_2' with $X_2' \supseteq X_1'$. The desired result follows. ☐

Submodularity. A real-valued set function f is called submodular if for any two sets S_1, S_2, we have

$$f(S_1) + f(S_2) \geq f(S_1 \cup S_2) + f(S_1 \cap S_2)\,. \qquad (8)$$

Submodularity captures a property of diminishing returns, whereby the contribution of a node $u \in S_1$ to the value of $f(S_1)$ decreases as S_1 increases. The maximization of monotone submodular functions is efficiently approximable, even though it might be intractable to achieve the maximum value. In light of Theorem 2 Lemma 4, it is natural to hope for the submodularity of the fixation probability as a means to at least approximate our maximization problem efficiently. Unfortunately, as the next theorem shows, neither $\mathrm{fp}(G^S, \delta)$ nor $\mathrm{fp}^\infty(G^S)$ exhibit submodularity.

Theorem 3. *The following assertions hold:*
1. $\mathrm{fp}(G^S, \delta)$ *is not submodular, and this holds also on graphs with self loops.*
2. $\mathrm{fp}^\infty(G^S)$ *is not submodular in general.*

Proof. Our proof is via counterexamples, shown in Fig. 4.
1. Consider the cycle graph of four nodes (Fig. 4, left), with $S_1 = \{0, 2\}$ and $S_2 = \{1, 3\}$. We calculate
$$\mathrm{fp}(G^{S_1}, 0.1) = \mathrm{fp}(G^{S_2}, 0.1) \leq 0.26194, \text{ while}$$
$$\mathrm{fp}(G^{S_1 \cup S_2}, 0.1) \geq 0.274 \text{ and } \mathrm{fp}(G^{S_1 \cap S_2}, 0.1) \geq 0.25$$
Now consider the same graph with additional self-loops (Fig. 4, middle). We calculate
$$\mathrm{fp}(G^{S_1}, 0.1) = \mathrm{fp}(G^{S_2}, 0.1) \leq 0.2702, \text{ while}$$
$$\mathrm{fp}(G^{S_1 \cup S_2}, 0.1) \geq 0.2909 \text{ and } \mathrm{fp}(G^{S_1 \cap S_2}, 0.1) \geq 0.25$$
2. Consider the wheel graph of 9 nodes (Fig. 4, right), with $S_1 = \{1\}$ and $S_2 = \{5\}$. We calculate
$$\mathrm{fp}^\infty(G^{S_1}) = \mathrm{fp}^\infty(G^{S_2}) \leq 0.19, \text{ while}$$
$$\mathrm{fp}^\infty(G^{S_1 \cup S_2}) \geq 0.27 \text{ and } \mathrm{fp}^\infty(G^{S_1 \cap S_2}) \geq 0.111$$

All three aforementioned examples violate Eq. (8). □

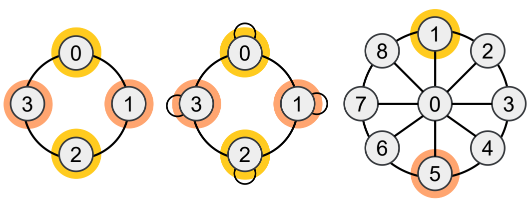

Figure 4: Counterexamples to submodularity.

Theorem 3 distinguishes the positional Voter and Moran models, as for the latter, $\mathrm{fp}^\infty(G^S)$ is submodular on arbitrary graphs (Brendborg et al. 2022). Also, note the asymmetry between $\mathrm{fp}(G^S, \delta)$ and $\mathrm{fp}^\infty(G^S)$ in Theorem 3 with regards to graphs having self loops. As we show next, $\mathrm{fp}^\infty(G^S)$ becomes submodular on graphs with self loops.

Lemma 5. *For any undirected graph G with self-loops, $\mathrm{fp}^\infty(G^S)$ is submodular.*

Proof. A finite trajectory $\mathcal{T}$ is either a node $\mathcal{T} = u_0$ or a sequence $\mathcal{T} = u_0, (u_1, v_1), (u_2, v_2), \dots, (u_m, v_m), u_{m+1}$, representing that
1. trait A starts from u_0,
2. for each time $t \in \{1, \dots, m\}$, node u_t adopts the trait of its neighbor v_t, and
3. u_{m+1} is chosen for death at time $m + 1$.

A prefix $\mathcal{T}'$ of $\mathcal{T}$ is $\mathcal{T}' = u_0$, if $\mathcal{T} = u_0$, or $\mathcal{T}' = u_0, (u_1, v_1), \dots, (u_{t-1}, v_{t-1}), u_t$, for some $t \in [m + 1]$. Thus $\mathcal{T}'$ follows $\mathcal{T}$ until the death-step of the $(t+1)$-th event. We say that $\mathcal{T}$ is *minimal and fixating* for S if either $\mathcal{T} = u_0$ with $u_0 \in S$, or
1. $u_0 \notin S$;
2. for each $t \in \{1, \dots, m\}$, we have that either $u_t \notin S$, or all neighbors of u_t are B-nodes after executing $u_0, (u_1, v_1), \dots (u_t, v_t)$; and
3. $u_{m+1} \in S$ and u_{m+1} has at least one A-neighbor after executing $u_0, (u_1, v_1), (u_2, v_2), \dots, (u_m, v_m)$.

In other words, the last step of $\mathcal{T}$ is the first that makes a biased node adopt the novel trait A. Every trajectory that leads to the fixation of A in G^S has a *minimal and fixating* prefix. By Lemma 2, the opposite is also true: every minimal and fixating trajectory eventually leads to fixation. Thus, we can compute $\mathrm{fp}^\infty(G^S)$ by summing the probabilities of occurrence of each minimal and fixating trajectory $\mathcal{T}$.

Moreover, the probability of a minimal and fixating $\mathcal{T}$ occurring is independent of S: the steps u_0 and u_{m+1} have probability of $1/n$ (since the initial placement of A is uniform, and each node is chosen for death also uniformly), while each step (u_t, v_t), for $t \in \{1, \dots, m\}$ has probability $\frac{1}{n \cdot \mathrm{d}(u_t)}$. Thus, to arrive at the submodularity of $\mathrm{fp}^\infty(G^S)$, it suffices to argue that, for any two biased sets S_1, S_2,
1. if $\mathcal{T}$ is minimal and fixating for $S_1 \cup S_2$, then it is minimal and fixating for at least one of S_1, S_2, and
2. if $\mathcal{T}$ is minimal and fixating for $S_1 \cap S_2$, then it has prefixes that are minimal and fixating for both S_1 and S_2.

Indeed, for Item 1, if $\mathcal{T} = u_0$, then $u_0 \in S_1 \cup S_2$ thus clearly $\mathcal{T}$ is minimal and fixating for at least one of S_1 and S_2. Similarly, if $\mathcal{T} = u_0, (u_1, v_1), (u_2, v_2), \dots, (u_m, v_m), u_{m+1}$, then u_{m+1} has an A-neighbor in $S_1 \cup S_2$. Thus $\mathcal{T}$ is minimal and fixating for at least one of S_1 and S_2, as further, no earlier node u_t, for $t \in [m]$ could have an A-neighbor in $S_1 \cup S_2$.

Similarly for Item 2, if $\mathcal{T} = u_0$, then $u_0 \in S_1 \cap S_2$, and thus $\mathcal{T}$ is also minimal and fixating for both S_1 and S_2. On the other hand, if $\mathcal{T} = u_0, (u_1, v_1), (u_2, v_2), \dots, (u_m, v_m), u_{m+1}$, then u_{m+1} has an A-neighbor in $S_1 \cap S_2$. Hence, for each set $Y \in \{S_1, S_2\}$, either some earlier node u_t has an A-neighbor in Y (and thus the prefix $\mathcal{T}' = u_0, (u_1, v_1), \dots, (u_{t-1}, v_{t-1}), u_t$ is minimal and fixating for Y), or $\mathcal{T}$ is fixating for Y at u_{m+1}. The desired result follows. □

The monotonicity and submodularity properties of Lemma 4 and Lemma 5 lead to the following theorem (Nemhauser, Wolsey, and Fisher 1978).

Theorem 4. *Given an undirected graph G with self loops and integer k, let S^* be the biased set that maximizes $\mathrm{fp}^\infty(G^S)$, and S' the biased set constructed by a greedy algorithm opting for maximal gains in each step. Then* $\mathrm{fp}^\infty(G^{S'}) \geq (1 - \frac{1}{e})\, \mathrm{fp}^\infty(G^{S^*})$.

Optimization for Weak Bias

In this section we turn our attention to the case of weak bias. Recall that our goal in this setting is to maximize $\mathrm{fp}'(G^S, 0)$, i.e., the derivative of the fixation probability evaluated at $\delta = 0$. We show that the problem can be solved efficiently on the class of graphs that have symmetric edge weights (i.e., $w(u, v) = w(v, u)$ for all $u, v \in V$). We arrive at our result by extending the weak-selection method that was developed recently for the basic Voter model in the context of evolutionary dynamics (Allen et al. 2020, 2021).

Consider a symmetric graph $G = (V, E)$ and a biased set S. Given a node $i \in V$, we write $\lambda_i = 1$ to denote that $i \in S$, and $\lambda_i = 0$ otherwise. Given two nodes $i, j \in V$, we let $p_{ij} = \frac{w(i,j)}{\sum_{l \in V} w(i,l)}$ be the probability that a 1-step random walk starting in i ends in j. We also let $b_{ij}^{(2)} = \sum_{l \in V} \lambda_l p_{il} \sum_{j \in V} p_{lj}$ be the probability that a 2-step random walk that starts in i, goes through a biased node l and ends in j. The following is the key result in this section.

Lemma 6. *Consider a symmetric graph $G = (V, E)$, and arbitrary biased set S. We have $\mathrm{fp}'(G^S, 0) = \frac{1}{n}\sum_{i,j \in V} \pi_i b_{ij}^{(2)} \psi_{ij}$, where $\{\pi_i | i \in V\}$ is the solution to the linear system*

$$\pi_i = \left(1 - \sum_{j \in V} p_{ji}\right) \pi_i + \sum_{j \in V} p_{ij} \pi_j, \quad \forall i \in V$$

$$\sum_{i \in V} \pi_i = 1 \tag{9}$$

and $\{\psi_{ij} | (i, j) \in E\}$ is the solution to the linear system

$$\psi_{ij} = \begin{cases} \frac{1 + \sum_{l \in V}(p_{il}\psi_{lj} + p_{jl}\psi_{il})}{2} & j \neq i \\ 0 & \textit{otherwise} \end{cases} \tag{10}$$

Although Lemma 6 might look somewhat arbitrary at first, the expressions in it have a natural interpretation. We provide this interpretation here, while we refer to the full version of the paper (Petsinis, Pavlogiannis, and Karras 2022) for the detailed proof.

The quantities π_i express the probability that the novel trait A fixates when $\delta = 0$ and the invasion starts from node i, and are known to follow Eq. (9) (Allen et al. 2021). Since $\delta = 0$, the two traits A and B are indistinguishable by the Voter process. Note that, eventually, the whole population will adopt the initial trait of some node i, which leads to $\sum_i \pi_i = 1$. The first part of Eq. (9) expresses the fact that the current trait of node i can fixate by either (i) node i not adopting a new trait in the current round, and having its trait fixating from the next round on (first term), or (ii) node i propagating its trait to some neighbor j in the current round, and having j's trait fixate from that round on.

The quantities ψ_{ij} express the average time throughout the Voter process that nodes i and j spend carrying traits A and B, respectively. First, note that if $i = j$ (second case of Eq. (10)), then clearly $\psi_{ij} = 0$, as i and j always have the same trait. Now, focusing on the first case of Eq. (10), the term 1 in the numerator captures the case that the invasion starts at node i, in which case indeed i and j carry traits A and B respectively. The second term in the numerator captures the evolution of ψ_{ij} throughout the process, and is obtained similarly to the π_i's above. Indeed, given a current configuration X, in the next round i and j carry traits A and B, respectively, if either (i) node i adopts the trait of some node l, while l and j have traits A and B, respectively (term $p_{il}\psi_{lj}$), or (ii) node j adopts the trait of some node l, while i and l have traits A and B, respectively (term $p_{lj}\psi_{il}$). The denominator 2 in Eq. (10) is a normalizing term.

Now we turn our attention to the expression

$$\mathrm{fp}'(G^S, 0) = \frac{1}{n}\sum_{i,j \in V} \pi_i b_{ij}^{(2)} \psi_{ij} = \frac{1}{n}\sum_{i \in V} \pi_i \sum_{j \in V} b_{ij}^{(2)} \psi_{ij}.$$

Operationally, this expression can be interpreted as follows: with probability $1/n$, the invasion of A starts at node i. Then the contribution of the bias δ to the fixation of i's trait A is multiplicative to the baseline fixation probability π_i of i under $\delta = 0$. The multiplicative factor $\sum_{j \in V} b_{ij}^{(2)} \psi_{ij}$ can be understood by expanding $b_{ij}^{(2)}$, thereby rewriting as

$$\sum_{j \in V} b_{ij}^{(2)} \psi_{ij} = \sum_{l \in V} \lambda_l p_{il} \sum_{j \in V} p_{lj} \psi_{il}$$

Node i benefits by carrying the resident trait A whenever a neighbor thereof, l, is chosen for death. Still, the bias δ has an effect if and only if $l \in S$ (hence the factor λ_l). Moreover, even when $l \in S$, the benefit of i is further proportional to the chance ψ_{ij} that i carries trait A while j carries trait B (summed over all neighbors j of l), since, when j also carries trait A, it cancels the advantage that i has due to the bias δ.

We can now obtain our main result under the weak bias.

Theorem 5. *Given a symmetric graph G of n nodes, the maximization of $\mathrm{fp}'(G^S, 0)$ can be done in $\mathcal{O}(n^{2\omega})$ time, where ω is the matrix-multiplication exponent.*

Proof. We use standard algorithms to solve the linear systems of Lemma 6 and compute the quantities π_i and ψ_{ij}. In particular, Eq. (9) contains n unknowns and can be solved in $\mathcal{O}(n^\omega)$ time, while Eq. (10) contains n^2 unknowns and can be solved in $\mathcal{O}(n^{2\omega})$. Then, by Lemma 6, we have

$$\mathrm{fp}'(G^S, 0) = \frac{1}{n}\sum_{i,j \in V} \pi_i \cdot b_{ij}^{(2)} \cdot \psi_{ij} = \frac{1}{n}\sum_{i \in V} \pi_i \sum_{j \in V} b_{ij}^{(2)} \psi_{il} =$$

$$= \frac{1}{n}\sum_{i \in V} \pi_i \sum_{j \in V}\sum_{l \in V} \lambda_j p_{ij} p_{jl} \psi_{il} = \frac{1}{n}\sum_{j \in V} \lambda_j h(j), \tag{11}$$

where $h(j) = \sum_{i,l \in V} \pi_i p_{ij} p_{jl} \psi_{il}$. Hence, the optimal biased set S is consisted of the top-k nodes by $h(\cdot)$ value. □

Experiments

We present experimental results for the proposed algorithms and additional heuristics, on 100 randomly-selected strongly connected components of real-life social and community networks (Peixoto 2020), with varying the number of nodes in the range $[20, 130]$. For each graph we use a budget k corresponding to the 10%, 30% and 50% of the graph's size. We evaluate 7 algorithms that are often used in network analysis.

1. Random: k nodes uniformly at random.
2. Closeness: Top-k nodes by closeness centrality.
3. Betweenness: Top-k nodes by betweenness centrality.
4. Harmonic: Top-k nodes by harmonic centrality.
5. Degree: Top-k nodes by degree.

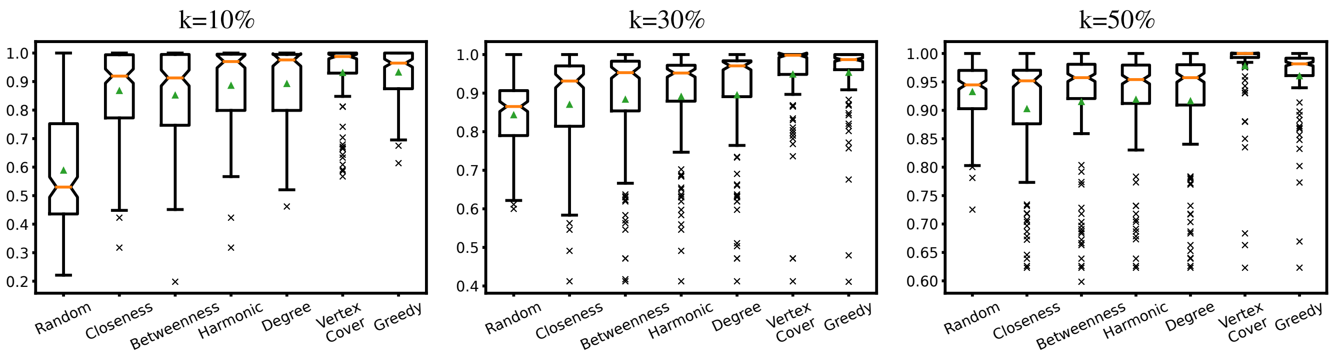

Figure 5: Performance under strong bias ($\delta \to \infty$)

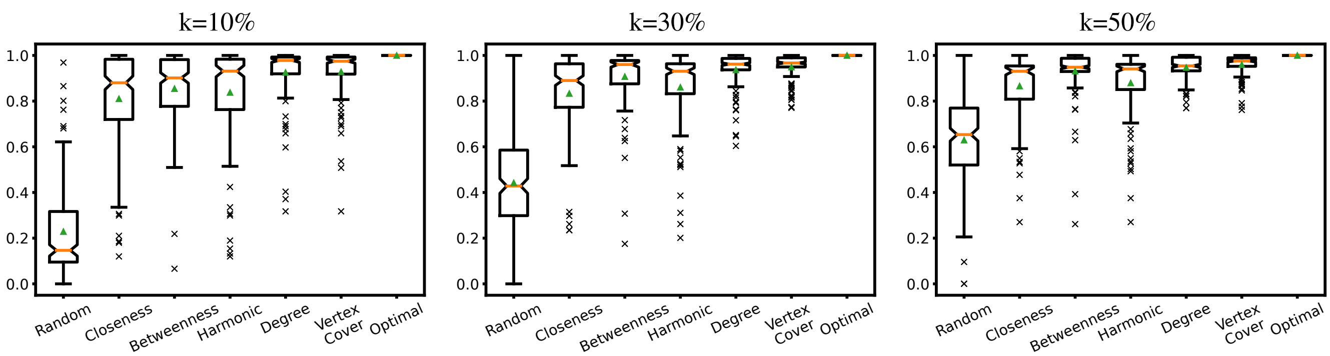

Figure 6: Performance under weak bias ($\delta \to 0$)

6. Vertex Cover: A greedy algorithm that seeks the set S maximizing covered edges, motivated by Lemma 3, by which vertex cover yields an optimal set under strong bias on regular graphs with self-loops; it sequentially selects the node with the maximum marginal gain wrt the number of edges having at least one endpoint in S.
7. Greedy: A greedy algorithm that seeks the set S that maximizes the fixation probability. The algorithm sequentially selects the node with the maximal marginal gain wrt fixation probability. To evaluate the gain in each iteration we simulate the process many times to achieve high precision. This algorithm is motivated by Theorem 1, which indicates that the greedy algorithm provides an $(1 - 1/e)$-approximation guarantee on graphs with self-loops.

While Random places the nodes uniformly all over the network, Closeness and Harmonic centralities consider the distance of a node to all other nodes. Betweenness Centrality indicates the importance of a node as a bridge, while Degree only considers the 1-hop neighbors. Vertex Cover considers edge coverage rather than an individual score. Lastly, the Greedy approach is informed of the positional Voter process. We report relative performance, dividing each fixation probability by the maximum achieved in a given network.

Strong bias ($\delta \to \infty$). Fig. 5 illustrates the distribution of the normalized results over all graphs. We first observe that Greedy and Vertex Cover perform best for all size constraints and have similar behavior. This high performance of Greedy and Vertex Cover is expected given the theoretical guarantees from Theorem 4 and Lemma 3, respectively. For small k, the problem is more challenging, as Random performs poorly. As k increases, Random is likely to cover all the graph, and the precise selection is less important.

Weak bias ($\delta \to 0$). By Theorem 5, to maximize the fixation probability, we need to maximize its derivative $\mathrm{fp}'(G^S, 0)$. We find the optimal value by solving the linear equation system described in Lemma 6. Fig. 6 presents the normalized results over all graphs. We see that Degree and Vertex Cover perform better than centrality algorithms. Our intuition is that nodes with many neighbors provide good invasion hubs, as they propagate their trait more frequently. Lastly, the Random algorithm under weak selection is ineffective.

Conclusion

We introduced the *positional Voter model*, which generalizes the standard Voter model to express *localized* effects that bias the invasion of novel traits. The new model raises the optimization problem of maximizing the fixation probability by distributing such effects in the network. A number of theoretical questions remain open with respect to this problem. Can we achieve efficient approximations, despite the lack of submodularity with finite bias δ, or for the case of strong bias, under which submodularity holds for graphs with self-loops? Does the tractability under weak bias extend to nonsymmetric graphs (i.e., when generally $w(u, v) \neq w(v, u)$?

Acknowledgments

This work was supported by DFF (Project 9041-00382B) and Villum Fonden (Project VIL42117).

References

Allen, B.; Lippner, G.; Chen, Y.-T.; Fotouhi, B.; Momeni, N.; Yau, S.-T.; and Nowak, M. A. 2017. Evolutionary dynamics on any population structure. *Nature*, 544(7649): 227–230.

Allen, B.; Sample, C.; Jencks, R.; Withers, J.; Steinhagen, P.; Brizuela, L.; Kolodny, J.; Parke, D.; Lippner, G.; and Dementieva, Y. A. 2020. Transient amplifiers of selection and reducers of fixation for death-Birth updating on graphs. *PLOS Computational Biology*, 16(1): 1–20.

Allen, B.; Sample, C.; Steinhagen, P.; Shapiro, J.; King, M.; Hedspeth, T.; and Goncalves, M. 2021. Fixation probabilities in graph-structured populations under weak selection. *PLoS computational biology*, 17(2): e1008695.

Alon, N.; Feldman, M.; Lev, O.; and Tennenholtz, M. 2015. How Robust is the Wisdom of the Crowds? In *Twenty-Fourth International Joint Conference on Artificial Intelligence*.

Antal, T.; Redner, S.; and Sood, V. 2006. Evolutionary Dynamics on Degree-Heterogeneous Graphs. *Phys. Rev. Lett.*, 96: 188104.

Barbieri, N.; and Bonchi, F. 2014. Influence maximization with viral product design. In *Proceedings of the 2014 SIAM International Conference on Data Mining*, 55–63. SIAM.

Barbieri, N.; Bonchi, F.; and Manco, G. 2013. Topic-aware social influence propagation models. *Knowledge and information systems*, 37(3): 555–584.

Bhat, D.; and Redner, S. 2019. Nonuniversal opinion dynamics driven by opposing external influences. *Phys. Rev. E*, 100: 050301.

Brendborg, J.; Karras, P.; Pavlogiannis, A.; Rasmussen, A. U.; and Tkadlec, J. 2022. Fixation Maximization in the Positional Moran Process. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(9): 9304–9312.

Castellano, C.; Fortunato, S.; and Loreto, V. 2009. Statistical physics of social dynamics. *Reviews of modern physics*, 81(2): 591.

Clifford, P.; and Sudbury, A. 1973. A model for spatial conflict. *Biometrika*, 60(3): 581–588.

Domingos, P.; and Richardson, M. 2001. Mining the Network Value of Customers. In *ACM SIGKDD*, KDD '01, 57–66. ISBN 158113391X.

Durocher, L.; Karras, P.; Pavlogiannis, A.; and Tkadlec, J. 2022. Invasion Dynamics in the Biased Voter Process. In Raedt, L. D., ed., *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, 265–271. International Joint Conferences on Artificial Intelligence Organization. Main Track.

Even-Dar, E.; and Shapira, A. 2011. A note on maximizing the spread of influence in social networks. *Information Processing Letters*, 111(4): 184–187.

Feige, U. 2003. Vertex cover is hardest to approximate on regular graphs. *Technical Report MCS03–15*.

Hindersin, L.; and Traulsen, A. 2015. Most undirected random graphs are amplifiers of selection for birth-death dynamics, but suppressors of selection for death-birth dynamics. *PLoS Comput Biol*, 11(11): e1004437.

Ibsen-Jensen, R.; Chatterjee, K.; and Nowak, M. A. 2015. Computational complexity of ecological and evolutionary spatial dynamics. *Proceedings of the National Academy of Sciences*, 112(51): 15636–15641.

Ivanov, S.; Theocharidis, K.; Terrovitis, M.; and Karras, P. 2017. Content recommendation for viral social influence. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 565–574.

Kempe, D.; Kleinberg, J.; and Tardos, É. 2003. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, 137–146.

Kempe, D.; Kleinberg, J.; and Tardos, É. 2005. Influential nodes in a diffusion model for social networks. In *International Colloquium on Automata, Languages, and Programming*, 1127–1138. Springer.

Kermack, W. O.; McKendrick, A. G.; and Walker, G. T. 1927. A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, 115(772): 700–721.

Lieberman, E.; Hauert, C.; and Nowak, M. A. 2005. Evolutionary dynamics on graphs. *Nature*, 433(7023): 312–316.

Liggett, T. M.; and Liggett, T. M. 1985. *Interacting particle systems*, volume 2. Springer.

McAvoy, A.; and Allen, B. 2021. Fixation probabilities in evolutionary dynamics under weak selection. *Journal of Mathematical Biology*, 82(3): 1–41.

Moran, P. A. P. 1958. Random processes in genetics. *Mathematical Proceedings of the Cambridge Philosophical Society*, 54(1): 60–71.

Mossel, E.; and Roch, S. 2007. On the Submodularity of Influence in Social Networks. In *STOC*, 128–134. ISBN 9781595936318.

Nemhauser, G. L.; Wolsey, L. A.; and Fisher, M. L. 1978. An analysis of approximations for maximizing submodular set functions—I. *Mathematical programming*, 14(1): 265–294.

Newman, M. E. J. 2002. Spread of epidemic disease on networks. *Phys. Rev. E*, 66: 016128.

Ohtsuki, H.; Hauert, C.; Lieberman, E.; and Nowak, M. A. 2006. A simple rule for the evolution of cooperation on graphs and social networks. *Nature*, 441(7092): 502–505.

Pavlogiannis, A.; Tkadlec, J.; Chatterjee, K.; and Nowak, M. A. 2018. Construction of arbitrarily strong amplifiers of natural selection using evolutionary graph theory. *Communications Biology*, 1(1): 71.

Peixoto, T. P. 2020. Netzschleuder: the network catalogue, repository and centrifuge. https://networks.skewed.de. Accessed: 2022-07-30.

Petsinis, P.; Pavlogiannis, A.; and Karras, P. 2022. Maximizing the Probability of Fixation in the Positional Voter Model. *arXiv preprint arXiv:2211.14676*.

Talamali, M. S.; Saha, A.; Marshall, J. A.; and Reina, A. 2021. When less is more: Robot swarms adapt better to changes with constrained communication. *Science Robotics*, 6(56): eabf1416.

Tkadlec, J.; Pavlogiannis, A.; Chatterjee, K.; and Nowak, M. A. 2020. Limits on amplifiers of natural selection under death-Birth updating. *PLoS computational biology*, 16(1): e1007494.

Tkadlec, J.; Pavlogiannis, A.; Chatterjee, K.; and Nowak, M. A. 2021. Fast and strong amplifiers of natural selection. *Nature Communications*, 12(1): 4009.

Zhang, K.; Zhou, J.; Tao, D.; Karras, P.; Li, Q.; and Xiong, H. 2020. Geodemographic influence maximization. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2764–2774.