

REPUBLISH: Reducing the Polarized Bubble Radius with Link Insertions

Extended version

Shahrzad Haddadan
Dept. of Computer Science
& Data Science Initiative
Brown University
Providence, RI, USA
shahrzad_haddadan@brown.edu

Cristina Menghini
DIAG
Sapienza University
Rome, Italy
menghini@diag.uniroma1.it

Matteo Riondato
Dept. of Computer Science
Amherst College
Amherst, MA, USA
mriondato@amherst.edu

Eli Upfal
Dept. of Computer Science
Brown University
Providence, RI, USA
eli@cs.brown.edu

Democracy begins in conversation — John Dewey (attr.)

ABSTRACT

The topology of the hyperlink graph among pages expressing different opinions may influence the exposure of readers to diverse content. Structural bias may trap a reader in a “polarized” bubble with no access to other opinions. We model readers’ behavior as random walks. A node is in a “polarized” bubble if the expected length of a random walk from it to a page of different opinion is large. The structural bias of a graph is the sum of the radii of highly-polarized bubbles. We study the problem of decreasing the structural bias through edge insertions. “Healing” all nodes with high polarized bubble radius is hard to approximate within a logarithmic factor, so we focus on finding the best k edges to insert to maximally reduce the structural bias. We present REPUBLISH, an algorithm that leverages a variant of the random walk closeness centrality to select the edges to insert. REPUBLISH obtains, under mild conditions, a constant-factor approximation. It reduces the structural bias faster than existing edge-recommendation methods, including some designed to reduce the polarization of a graph.

CCS CONCEPTS

• **Theory of computation** → *Random walks and Markov chains*;
Graph algorithms analysis; • **Information systems** → *Social networks*; **Social recommendation**.

KEYWORDS

Bias, Fairness, Polarization

ACM Reference Format:

Shahrzad Haddadan, Cristina Menghini, Matteo Riondato, and Eli Upfal. 2021. REPUBLISH: Reducing the Polarized Bubble Radius with Link Insertions: Extended version. In *Proceedings of the Fourteenth ACM International Conference on Web Search and Data Mining (WSDM '21)*, March 8–12, 2021, Virtual Event, Israel. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3437963.3441825>

WSDM '21, March 8–12, 2021, Virtual Event, Israel

© 2021 Association for Computing Machinery.

This is the author’s version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *Proceedings of the Fourteenth ACM International Conference on Web Search and Data Mining (WSDM '21)*, March 8–12, 2021, Virtual Event, Israel, <https://doi.org/10.1145/3437963.3441825>.

1 INTRODUCTION

The World Wide Web often contains thousands or even millions of pages on every topic, covering the whole spectrum of opinions. Exposure to *diverse content* is necessary to obtain a complete picture about a topic. This exposure depends on the hyperlinks connecting the pages to each other. It can be argued that enabling easier access to diverse content improves society as it creates a more informed and less polarized general public [11]. Indeed politicians have strongly promoted and even requested that audiences are exposed to varied content [35].

The fact that diverse information is easily *available* does not imply that *exploring* such diverse information is easy. Rather, echo chambers and polarization on (social) media and blogs [1, 17, 25] keep the user in a *homogeneous bubble*, exposing them only to agreeable information [8], and leading to conflicts between users in different bubbles [18, 34].

A web user can freely click on any hyperlink on the page they are currently visiting, but the choice of which hyperlinks to include in the page is with the website owner or editor, who, if not careful, may stop the user from being exposed to diverse opinions. In other words, the hyperlink topology of a website may suffer from *structural bias* that traps the user in a bubble of one-sided content without them knowing [44, 55]. For example, structural bias on topic-induced networks, such as Wikipedia topic-induced subgraphs, prevents users from building a well-rounded knowledge about the topic. On query-/user-induced recommendation networks such as those on Amazon and YouTube, structural bias hinders the discovery of diversified content, reducing serendipity [3, 30, 33]. Structural bias thus limits the user’s freedom while navigating the Web.

Socially-minded website editors would try to *minimize such structural bias* by adding appropriate links to pages in the website. As an editor can only modify a few pages and add only a few links to each edited page, they need to carefully choose what page p to edit, and what pages to link to from p . The goal of our work is to develop an algorithm that can give editors recommendations for what links to add in order to reduce the structural bias. There are key technical challenges that must be solved in order to give effective link recommendations: 1. quantify the structural bias of a page p , i.e., how hard it is to reach pages of a different opinion from p ; and 2. decide which of these pages should be linked from p . Existing approaches to link recommendation, such as those based on vertex

similarity, fall short at this task because they are oblivious to the network structural bias. Their recommendations often *increase* the bias, rather than decreasing it [48], especially for highly controversial topics such as political blogs (see also our experimental results in Sect. 6). Thus, there is a need for a radically different approach to suggest links that decrease the structural bias.

Contributions. We study the problem of reducing the structural bias of a graph by adding edges, and propose an algorithm, REP-BUBBLIK, to suggest such edges. Our contributions are the following.

- We consider directed graphs with vertices of two colors, representing a network of webpages on the same topic, with the two colors identifying the two opposite opinions on the topic, and edges representing links between pages. We define the (*Polarized*) *Bubble Radius* (BR) of a vertex p as a novel measure to quantify the structural bias of p (see Def. 4.1), based on a task-specific variant of the hitting time for random walks, which models the navigation of a user on the web [23, 24]. The BR is the expected number of steps to go from p to a page of different opinion, and can be easily estimated with a sampling-based approach with probabilistic guarantees (Lemma 4.3), which enables us to tackle the first of the key challenges.
- We define the *structural bias* of a graph G as the sum of the BRs of vertices with high BR (Eq. 1). Completely removing the bias is APX-hard by reduction from set cover (see Lemma 4.5). We therefore state the *k-edge structural bias decrease* problem as the task of finding the set of k pairs of vertices of different color such that adding the edge between the vertices in each pair would *maximally* decrease the structural bias, over all possible sets of k pairs (see Prob. 2 and Thm. 5.1). This problem connects two areas: link recommendation and polarization reduction.
- We present REP-BUBBLIK, an efficient approximation algorithm for the *k-edge structural bias decrease* problem, that recommends the addition of k edges between vertices of different color. Under mild conditions, the resulting decrease of the structural bias is within a constant factor of the optimal. Website editors have limited control on the probability that a newly added edge will be traversed by the users, so our algorithm makes no assumption or impose any restriction on it, as this probability is essentially external. At the core of REP-BUBBLIK is an analysis of the submodularity of the objective function (see Lemma 5.6), combined with the use of a task-specific variant of random-walk closeness [62], a well-established centrality measure. REP-BUBBLIK requires good estimations of the random walk closeness, so we also give an approximation algorithm for this quantity (see Lemma 3.1).
- We evaluate REP-BUBBLIK on eight real datasets. We compare it to baselines and existing methods for edge recommendation either designed with the goal of reducing the controversy of a graphs [26] or with the more general purpose of completing the network’s link structure [31]. Our algorithm leads to a faster reduction of the average BR (i.e., requiring fewer edge insertions) than existing contributions.

The proofs of our theoretical results are in App. A.

2 RELATED WORK

Polarization has long been studied in political science [32, 59], and the recent diffusion of (micro-) blog and social media platforms brought the issue to the attention of the broad computer science community. Many works focused on showing the existence of polarization on these platforms [1, 17, 18, 25, 45], and on modeling, quantifying, and reducing polarization [2, 7, 9, 16, 26–29, 37, 38, 40, 40, 41, 47–49], or the glass ceiling effect [56–58]. The literature is rich, to the point that times seems ripe for an in-depth survey on the topic. Due to space limitations, we discuss here only the relationship between our work and the most relevant algorithmic contributions to polarization reduction [7, 9, 16, 26, 28, 41, 43, 44, 48, 58].

A first important difference of our work with respect to most previous contributions is that they consider a network of *users*, with edges representing notions such as friendship or endorsement (e.g., retweets) [7, 9, 16, 26, 28, 41, 48, 58]. We focus instead on networks of *content*, such as web pages linked to each other, or products that are connected when similar. This deep difference makes our contribution quite orthogonal to the ones in these previous works: we focus on the polarization that is introduced by the topology of the network, rather than on the polarizing effect of content on users or on the effect of users on each other. We believe both aspects are important, but the structural bias we focus on has only been subject to few studies [43, 44]. These works, relying on the notion of weighted reciprocity, propose a static and dynamic analysis of structural bias on Wikipedia. The measure of structural bias we use is not tailored to a specific website.

A second relevant difference from many previous works is that we consider the “opinion” of a page (i.e., a vertex) to be fixed, as it depends on its content, while many previous contributions consider different models of user opinion dynamics [20, 46] to study the evolution of such opinions as the users are exposed to different content or recommended different friendships. The problem of recommending changes to the content of a page to modify the opinion expressed in it is interesting but outside the scope of our work. Instead, we focus on recommending the addition of links between pages, to reduce the structural bias.

An interesting line of work studies how to reduce polarization in the content seen by the users, by adapting information diffusion approaches through better selection of the seed set for cascades [7, 9, 28, 41, 57], or by directly acting on recommendation systems [54]. These methods can not be adapted to the problem we study, as they do not act on the graph of content, but on that of users.

The most similar methods to ours are those that act on the structure of the graph [16, 26, 48, 58], although as we mentioned, they consider a network of users, not of content. Musco et al. [48] propose a network-design approach: they aim to find the best set of edges between vertices such that the resulting graph would minimize both disagreement and polarization. Rather than a “design-from-scratch” approach, which seems mostly of theoretical relevance, we consider instead a practical incremental approach that suggests modifications to an existing network. Like us, Garimella et al. [26] consider a graph polarization measure based on random walks [29]. This measure essentially quantifies the probability that a user of one opinion is exposed to content from a user of a different opinion, thanks to a chain of retweets (represented by the

random walks). The measure is based on a variant of personalized PageRank for sets of users with different opinions. The task requires to recommend new edges, i.e., retweets, to increase this probability. Our measure of structural bias is instead defined on the basis of the (Polarized) Bubble Radius (BR) (Def. 4.1), which is a vertex-dependent measure that represents the expected number of steps, for a user starting at the page represented by vertex v , to reach, with a random walk, a vertex with color different from v , representing a page expressing a different opinion. Our measure is appropriate for our task of suggesting new edges to make it easier for user to reach pages of different opinions. In Sect. 6 we compare our approach to that of Garimella et al. [26].

An important line of work in graph analysis and mining looked at manipulating the topology to modify different interesting characteristic quantities of the graph, such as shortest paths and related measures [22, 50, 51, 53], various forms of centrality [4, 13, 19, 39, 42, 52, 61], and more [5, 6, 14, 60, 63]. Despite the fact that we consider a specific centrality to choose the source of the added edges, these methods cannot be used to solve our task of interest.

Another body of work related to ours are those which estimate graph properties using random walks [10, 12, 15, 21]. The studied properties are not defined based on random walks, rather random walks are used as a tool to estimate them. Here based on random walks, we define a new property for networks: *the structural bias*, and we use random walks to estimate it.

3 PRELIMINARIES

Let $G = (V, E)$ be a directed weighted graph with $|V| = n$ vertices, such that no vertex $v \in V$ has only incoming edges and no outgoing edges. V is partitioned in two disjoint sets R and B (i.e., $R \cap B = \emptyset$ and $R \cup B = V$), called “red” or “blue” vertices, respectively. We denote the color of a vertex v by $c(v)$ and its opposite color by $\bar{c}(v)$. The sets of all *other* vertices of the same color as v is denoted as C_v and the sets of all vertices of color different than v is denoted as \bar{C}_v .

The edge weights are *transition probabilities*, as follows. Let M be a $n \times n$ right-stochastic *transition matrix* associated to G , i.e., a matrix such that each entry $m_{i,j}$ is a probability, with $m_{i,j} = 0$ if $(i, j) \notin E$, and such that $\sum_{j=1}^n m_{i,j} = 1$.

We are interested in random walks on the graph G using the transition matrix M . Intuitively, a random walk starting at a vertex v explores the graph by choosing at each step an outgoing edge from the current vertex, with probability equal to the weight of such edge, independently from previous choices. Let $S \subseteq V$ and $v \in V$. Let $T_v(S)$ be the random variable indicating the first instant when a random walk from v hits (i.e., reaches) any vertex in S . The quantity $\mathbb{E}_G [T_v(S)]$ is known as the *hitting time of S from v* , where the expectation is over the space of all random walks on G starting from v , with transition probabilities given by M .

Variants of random walks, such as random walks with restarts or with back button, are widespread models for network exploration [23, 24]. It is realistic to assume that there is an upper bound t , which we call the *exploration factor*, on the length of a walk performed by the users. For example, we can assume that there is an upper limit on the number of pages that a user will visit one after the other in a browsing session. The value of the parameter t can be derived, for example, from traces of visits. In most practical cases, t

is likely to be bounded by a polylogarithmic quantity in the number of nodes, if not a constant.

For a random walk starting from $v \in V$, given a set $S \subseteq V$, we define the random variable $T_v^t(S)$ as $\min\{t, T_v(S)\}$. This variable is more appropriate for measuring the length of browsing sessions, which have bounded length, than the unbounded length classically used when discussing random walks.

For a graph Z , any vertex u , and any set S of vertices, let $u \xrightarrow[Z]{\text{cond}} S$, denote the event that a random walk in Z from u hits a vertex in S without first visiting any vertex in \bar{C}_u and while satisfying the condition *cond* on the number of steps needed to hit S . For example, $u \xrightarrow[Z]{<t} S$ is the event that a random walk in Z from u hits a vertex in S in *less than t steps*, without first visiting any vertex in \bar{C}_u . We denote the complementary event as $u \not\xrightarrow[Z]{\text{cond}} S$.

3.1 Random-Walk Closeness Centrality

We adapt the definition of the standard random-walk closeness centrality [62] to bounded random walks so that the contribution to the centrality of v by vertices that do not reach v in less than t' steps (in expectation) is zero, for any t' .

Random-walk closeness centrality (bounded form). For a vertex $v \in V$, and any t' , the t' -bounded Random Walk Closeness Centrality (RWCC) measure with respect to subset $S \subseteq V$ is

$$\begin{aligned} r^{t'}(v; S) &\doteq \frac{1}{|S|} \sum_{w \in S} \left(t' - \mathbb{E}_G \left[T_w^{t'}(v) \right] \right) \\ &= \frac{1}{|S|} \sum_{w \in S} \sum_{i=1}^{t'} (t' - i) \mathbb{P} \left(w \xrightarrow[G]{\bar{c}(v)} v \right). \end{aligned}$$

Computing the exact RWCC is expensive. To estimate $r^{t'}(v; S)$, we pick z vertices $\{w_i\}_{i=1}^z$ u.a.r. from S , and run some κ random walks to obtain an estimate \bar{h}_{w_i} of $\mathbb{E}_G [T_{w_i}^{t'}(v)]$ for each w_i . The quantity $\bar{r}(v) \doteq t' - 1/z \sum_{i=1}^z \bar{h}_{w_i}$ is a good approximation of $r^{t'}(v; S)$.

LEMMA 3.1. *Let $z \geq (t'/2\epsilon)^2 \delta^{-1}$. Then*

$$\mathbb{P} \left(\left| \bar{r}(v) - r^{t'}(v; S) \right| \geq \epsilon \right) \leq \delta.$$

4 BUBBLE RADIUS AND STRUCTURAL BIAS

We introduce the (Polarized) Bubble Radius to quantify how *likely* users starting their random walk on a vertex $v \in V$ of one color, are to hit a vertex of the other color in at most t steps.

Definition 4.1. The (Polarized) Bubble Radius (BR) $B_G^t(v)$ of v with exploration parameter t is

$$B_G^t(v) \doteq \mathbb{E}_G \left[T_v^t(\bar{C}_v) \right].$$

A random walk starting at a vertex v with high BR is unlikely to hit a vertex in \bar{C}_v in fewer-than-or-exactly t steps. The following lemma formalizes this idea on common models for web browsing (random walks with restarts or with back button [23, 24]).

LEMMA 4.2. *Let $r \in \mathbb{N}$, and consider a user who starts their random walk at $v \in V$ and may either restart their walk from v or hit the back button up to r times. Let \mathcal{T}_v be the random variable denoting*

the number of steps such user takes to hit a vertex in \bar{C}_v . If $B_G^t(v) \geq t(1 - 1/8r)$, then $\mathbb{P}(\mathcal{T}_v \leq t/2) \leq 1/4$. If instead $B_G^t(v) \leq b$ for some $b > 0$, then $\mathbb{P}(\mathcal{T}_v > 4br) \leq 1/4$.

Given t , it is easy to estimate $B_G^t(v)$ for each vertex $v \in V$ by sampling random walks from v . The following result, whose proof uses the Hoeffding's bound and the union bound, shows the trade-off between the number of sampled random walks and the accuracy in estimating the BR of v .

LEMMA 4.3. For each $v \in V$, let $w_1^{(v)}, w_2^{(v)}, \dots, w_r^{(v)}$ be r random walks from v and stopped either when they hit a vertex of color $\bar{c}(v)$ or when they run for t steps, whichever happens first. For $i = 1, \dots, r$, let $b_i^{(v)}$ be the length of random walk $w_i^{(v)}$. Let

$$\bar{B}(v) \doteq \frac{1}{r} \sum_{i=1}^r b_i^{(v)}.$$

Let $\varepsilon, \delta \in (0, 1)$. If $r \geq \frac{t^2}{\varepsilon^2} \ln \frac{2n}{\delta}$, then

$$\mathbb{P}(\exists v \in V \text{ s.t. } |B_G^t(v) - \bar{B}(v)| > \varepsilon) < \delta,$$

where the probability is over the choice of the random walks.

In the rest of the work, we assume for simplicity to have access to the exact BR of every vertex. The above result makes this assumption reasonable because computing approximations of extremely high quality is relatively inexpensive.

The structural bias. On the basis of the BR, we define two sets of vertices: *cosmopolitan* and *parochial*. Given two reals b and r with $1 \leq b < r \leq t$, the set $\mathcal{Z}(G)$ of *cosmopolitan* vertices contains all and only the vertices in G with BR at most b , and the set $\mathcal{P}(G)$ of *parochial* vertices contains all and only the vertices in G with BR at least r . For ease of notation, we do not include b and r in the notation for $\mathcal{Z}(G)$ and $\mathcal{P}(G)$. In the rest of this work, we assume for simplicity $b = 2$ and $r = t/2$, but this assumption can be easily removed. $\mathcal{Z}(G)$ and $\mathcal{P}(G)$ are *disjoint*, but they do not necessarily form a partitioning of V . We will often consider the partitioning of $\mathcal{P}(G)$ by color, i.e., the two sets $\mathcal{P}_R(G)$ and $\mathcal{P}_B(G)$, containing the parochial vertices of color R or B respectively.

Definition 4.4. The *structural bias* $\rho(G)$ of G is the sum of the BRs of the parochial nodes of G , i.e.,

$$\rho(G) \doteq \sum_{v \in \mathcal{P}(G)} B_G^t(v). \quad (1)$$

It is reasonable to consider only the parochial nodes in the definition of structural bias because they are the ones such that a random walk from them is very unlikely to hit any vertex of color different than the starting vertex (see also Lemma 4.2).

Our goal in this work is to find a set of edges with extrema of different color whose addition to G would decrease the structural bias of the network. It is reasonable to only consider edge with extrema of different color, as they are always preferable (i.e., will result in a higher decrease of the structural bias) than edges with monochromatic extrema: the addition of the new edge can only have positive impact on the parochial vertices of the same color as the source, and has no impact on the parochial vertices of the other color. If we could add *any number* of such edges to G , it would be easy to bring the structural bias of G to zero, as there would be no

parochial nodes left. This assumption is not realistic: the number of links that a website editor can add to a single page and to the whole graph is limited by many factors, such as the fact that a human-readable page cannot have too many links, and the fact that the editor can only spend a limited time on this activity. Nevertheless, ideally one would want to solve the following problem.

Problem 1. Given a color $C \in \{R, B\}$, find the smallest set A of pairs of distinct edges $(v, w) \notin E$ with $c(v) = C$ and $c(w) \neq C$ such that, for the graph $G_{\text{new}} = (V, E \cup A)$ it holds $\mathcal{P}_C(G_{\text{new}}) = \emptyset$.

LEMMA 4.5. *Problem 1 is NP-hard and APX-hard.*

5 REDUCING THE BR WITH INSERTIONS

Since Prob. 1 is hard to even approximate (Lemma 4.5), we seek to answer a close relative (Prob. 2). We first introduce a set of measures to capture the change in the BRs of the (original) parochial nodes of G after edge insertions. Let G_{new} be obtained from G by inserting a set Σ of directed edges between nodes of different colors, with each inserted edge $e = (v, w)$ having weight m_e (also denoted as m_{vw}). For a set U of vertices, we define the *gain* of U due to Σ as

$$\Delta(G, U, \Sigma, \{m_e\}_{e \in \Sigma}, t') \doteq \frac{1}{|U|} \sum_{u \in U} \left(B_G^{t'}(u) - B_{G_{\text{new}}}^{t'}(u) \right).$$

When adding an edge to the graph, we also have to decide its weight. It seems excessive to assume complete freedom in choosing the weight. We make the assumption that the weight m_{vw} of an edge (v, w) that we would like to add is given to us by an oracle which computes m_{vw} only as a function of v and of information *local* to v (e.g., its out-degree) obtained from G and potentially a set of other edges (and their weights) that we want to add from v . The weight m_{vw} is the probability that a random walk arriving at v will move to w in the next step. When adding (v, w) with weight m_{vw} , the other edges outgoing from v have their weights multiplied by $1 - m_{vw}$ to ensure that the sum of the weights of the edges leaving v is 1.

The problem we want to solve then is the following.

Problem 2. In graph G , let C be either blue or red. Find a set $\Sigma = \{(v_i, w_i)\}_{i=1}^k$ of k edges whose source vertices all have color C and all destination vertices have the other color, that maximizes $\Delta(G, \mathcal{P}_C(G), \Sigma, \{m_e\}_{e \in \Sigma}, t)$.

REPUBBLIK (Alg. 1) is our algorithm to approximate Prob. 2. Before describing it in detail, we give an intuition of its workings, and present the theoretical results that guided its design. Specifically, since our objective function is *monotonic and submodular* (Lemma 5.6), we can greedily choose the edges to be added one by one. Due to our oracle assumption on the weights, any vertex of color different than the source can be picked as the target of the added edge, so the problem essentially *reduces to finding the sources for the edges to be added*. Lemma 5.4 quantifies the gain when picking each source according to a specific measure depending on the bounded RWCC and on the oracle-given weight that only depends on the source. In Lemma 5.5 we show that under mild conditions this choice is constantly close to an optimal choice. The following theorem states the approximation qualities of REPUBBLIK.

THEOREM 5.1. *Let Σ be the output of REPUBBLIK and OPT be the optimal solution to Prob. 2. Let $\Delta_\Sigma = \Delta(G, \mathcal{P}_C(G), \Sigma, \{m_e\}_{e \in \Sigma}, t)$*

Then

$$\Delta(G, \mathcal{P}_C(G), \text{OPT}, \{m_e\}_{e \in \text{OPT}}, t) \leq (4\gamma(G) + 1) \left(1 + \frac{1}{e}\right) \Delta_\Sigma,$$

where $\gamma(G)$ is the maximum over all $u \in V$ of sum of the probabilities, for $i = 0, \dots, t-1$, that a random walk starting at u visits u at step i without first visiting a vertex in \bar{C}_u (see also (2)), which is a constant for many graphs.

We now proceed towards presenting lemmas which together provide a proof for Thm. 5.1. Lemmas 3.1 and 4.3 provide bounds of order $\Theta(nt^2)$ on the runtime of the pre-processing phases of REPBUBLIK. Therefore for small values of t , REPBUBLIK is more efficient than algorithms that compute hitting times using the Laplacian, which need $\Omega(n^3)$ steps.

For any vertex v , and $0 \leq i \leq t$, let $\Psi_v(i)$ be the probability that a random walk (in G) from v visits v at step i before reaching a vertex in \bar{C}_v (it holds $\Psi_v(0) = 1$ and $\Psi_v(1) = 0$ for every v). For any $t' \leq t$, let

$$\mathcal{F}_{t'}(v) = \sum_{i=0}^{t'-1} \Psi_v(i). \quad (2)$$

The following lemma shows upper and lower bounds to the change in the bubble radius of a vertex when an new edge from it is added to the graph.

LEMMA 5.2. *Let $v \in \mathcal{P}(G)$, $w \in \bar{C}_v$ and $t' \leq t$. Let G_{new} be the graph obtained after adding $e = (v, w)$ to G , with weight m_e . The gain $\Delta(G, v, e, m_e, t')$ is such that*

$$\left(B_G^{t'}(v) - 1\right) m_e \leq \Delta(G, v, e, m_e, t') \leq \mathcal{F}_{t'}(v) \left(B_G^{t'}(v) - 1\right) m_e.$$

Decreasing the BR of v decreases the BRs of vertices in C_v close to v , and thus the whole network. Lemma 5.3 quantifies this change.

LEMMA 5.3. *Let $e = (v, w)$ be the edge with weight m_e added to G to obtain G_{new} . For any other vertex $u \in \mathcal{P}_{C_v}(G)$, it holds*

$$\Delta(G, u, e, m_e, t) = \sum_{i=1}^{t-2} \left(\Delta(G, v, e, m_e, t-i) \mathbb{P} \left(u \overset{\substack{=i \\ \rightsquigarrow \\ G}}{\rightsquigarrow} v \right) \right).$$

Recall that our greedy choice is to identify a node v that maximizes the gain $\Delta(G, \mathcal{P}(G), (v, w), m_v, t)$ where w is any vertex in \bar{C}_v . Lemma 5.3 suggests that a good candidate v is a vertex that is likely to be reached by short random walks from many other vertices in $\mathcal{P}_{C_v}(G)$, a property that is captured by the bounded RWCC $r^{t-2}(v; \mathcal{P}_{C_v}(G))$ (Sect. 3.1).

Now, we first quantify the gain for adding an edge from any vertex with RWCC c (Lemma 5.4). Then we show that under mild conditions on the return time of vertices we get a constant approximation by greedily choosing a vertex with maximum RWCC $\times m_v$ (Lemma 5.5).

LEMMA 5.4. *Let $v \in \mathcal{P}(G)$. Let $w \in \bar{C}_v$, and assume to add the edge $e = (v, w)$ with weight m_e . It holds*

$$\Delta(G, \mathcal{P}_{C_v}(G), e, m_e, t) \geq \frac{m_e}{2} r^{t-2}(v; \mathcal{P}_{C_v}(G)).$$

This lemma suggests that inserting edges from a vertex v with the highest value of $m_v r^{t-2}(v; \mathcal{P}_{C_v}(G))$ may result in a larger improvement in the objective function than if we chose a different source. In the next lemma we compare the effect of choosing such sources to the effect of an optimal choice.

Algorithm 1 REPBUBLIK

```

1: Input: Graph  $G = (V, E)$ , desired insertions  $k_C$ , oracle  $\mathcal{W}_G : V \times 2^{V \times V} \rightarrow [0, 1]$ ,  $C \in \{R, B\}$ .
2: Output: Set  $\Sigma_C$  of  $k_C$  edges to be inserted, with their weights.
3:  $\Sigma_C \leftarrow \emptyset$ 
4: for  $i = 1$  to  $k_C$  do
5:    $P \leftarrow \text{computeParochials}(G \cup \Sigma_C, C)$ 
6:    $\mathcal{R} \leftarrow \text{computeRWCentrality}(P, G \cup \Sigma_C)$ 
7:    $v_i \leftarrow \text{argmax}_{v \in P} \mathcal{R}(v) \times \mathcal{W}_G(v, \Sigma_C)$ 
8:    $u_i \leftarrow \text{arbitrary in } \bar{C}_{v_i}$ 
9:    $\Sigma_C \leftarrow \Sigma_C \cup \{(v_i, u_i)\}$ 
10: end for
11: return  $\Sigma_C$ 

```

LEMMA 5.5. *Consider the set $\mathcal{P}_C(G)$ where C is either color. Among all vertices in $\mathcal{P}_C(G)$ let v and opt be*

$$\begin{aligned} \text{opt} &= \text{arg max}_{u \in \mathcal{P}_C(G)} \Delta(G, \mathcal{P}_C(G), e_u, m_u, t), \\ v &= \text{arg max}_{u \in \mathcal{P}_C(G)} m_u r^{t-2}(u; \mathcal{P}_C(G)), \end{aligned}$$

where e_u is any potentially inserted edge connecting u to \bar{C}_u , and m_u is its weight.¹

It holds

$$\Delta(G, \mathcal{P}_C(G), e_{\text{opt}}, m_{\text{opt}}, t) \leq (4\gamma(G) + 1) \Delta(G, \mathcal{P}(G), e_v, m_v, t),$$

where $\gamma(G) = \max_{u \in G} \mathcal{F}_t(u)$.

If the probability of getting back to u in less than t steps is less than α for some constant α then $\gamma(G) \leq \alpha$. This assumption is realistic since t is usually small and the return time to u is often much larger than t .

Finally, we show that the gain function is monotonic and sub-modular.

LEMMA 5.6. *Let C be either blue or red and $v, u \in \mathcal{P}_C(G)$, and $w_v, w_u \in \bar{C}_v$, such that $e_v = (v, w_v)$ and $e_u = (u, w_u)$ are not existing edges. Let $\Sigma = \{e_v, e_u\}$. It holds*

$$\Delta(G, \mathcal{P}_C(G), e_v, m_{e_v}, t) \leq \Delta(G, \mathcal{P}_C(G), \Sigma, \{m_e\}_{e \in \Sigma}, t), \quad (3)$$

and

$$\begin{aligned} \Delta(G, \mathcal{P}_C(G), \Sigma, \{m_e\}_{e \in \Sigma}, t) &\leq \Delta(G, \mathcal{P}_C(G), e_v, m_{e_v}, t) \\ &\quad + \Delta(G, \mathcal{P}_C(G), e_u, m_{e_u}, t). \end{aligned} \quad (4)$$

We are now ready to prove Thm. 5.1.

PROOF OF THM. 5.1. Lemma 5.6 shows the monotonicity and sub-modularity of the objective function. Thus, a greedy algorithm that picks, iteratively, the k best choices over all parochial vertices of color C as the sources of the added edges, will result in a $(1 + 1/e)$ -approximation. Lemmas 5.4 and 5.5 show that by choosing a vertex v maximizing $m_v r^{t-2}(v; \mathcal{P}_C(G))$ among all parochial vertices of color C , we obtain a vertex such that the gain when adding an edge from this source is a $4\gamma(G) + 1$ -approximation to the greedy choice. Thus, the correctness of our algorithm is concluded by putting these lemmas together. \square

¹Our assumption on the oracle giving the weight ensures that m_u only depends on u , not on the target of e_u .

We can now give the details to REPUBLIC. The algorithm takes as input the graph G , the number k_C of desired edge insertions, the oracle \mathcal{W} that determines the weights of the new edges, and the set of nodes C . It first creates the empty set Σ_C that will store the edges to be added and then enters a for loop to be repeated for k_C times. At every iteration of the loop, it first computes the BR of every node in C in the graph (denoted in the pseudocode as $G \cup \Sigma_C$) obtained by adding to G the edges currently in Σ_C (with their weights obtained from the oracle \mathcal{W}_G) (in practice, the BR is computed using the approximation algorithm outlined in Lemma 4.3). Thanks to this computation, the algorithm obtains (line 5) the set P of parochial nodes in this graph (at the first iteration of the loop $P = \mathcal{P}_C(G)$). It then obtains the centralities values $r^{t-2}(v; P)$ of every node $v \in P$ (in practice, using the approximation algorithm outlined in Lemma 3.1), storing them in a dictionary \mathcal{R} (line 6). The algorithm then selects the node $v_i \in P$ associated to the maximum quantity $\mathcal{R}(v_i) \times \mathcal{W}_G(v_i, \Sigma_C)$, and arbitrarily picks a node u_i of the opposite color of v_i (i.e., of the color other than C). The directed edge (v_i, u_i) is added to the set Σ_C (lines 7–9). After k_C iterations of the loop, the algorithm returns Σ_C , together with the weights obtained from the oracle.

REPUBLIC would require a re-computation of the BRs and of the centralities of all vertices, at every iteration of the loop, which would require to run a very large number of random walks, making it computationally very expensive. We now propose a more practical alternative REPUBLIC+, at the price of losing the approximation guarantees. REPUBLIC+ only computes $\mathcal{P}_C(G)$ and \mathcal{R} before entering the for loop, and uses the same values throughout its execution, but trades off the consequences of this choice by adding a penalty factor to the objective function involved in the selection of the source vertices for the edges to be added. Specifically, REPUBLIC+ chooses v_i (line 7) by maximizing the quantity $\mathcal{R}(v) \times \mathcal{W}_G(v, \Sigma_C) / \eta_v$, where η_v is a penalty factor equals to one plus the number of edges with source v in Σ_C (thus at iteration 1, $\eta_v = 1$ for every node). This penalty factor favours the insertion of edges from nodes that have not yet been altered. Consequently, it indirectly (1) handles the possibility that nodes with new edges are no longer parochial, thus we want to avoid to keep adding edges to them; and (2) avoids that the new edges are added from a restricted set of nodes, limiting the positive effect of the insertions on $\Delta(G, \mathcal{P}_C(G), \Sigma_C, \{m_e\}_{e \in \Sigma}, t')$.

6 EXPERIMENTAL EVALUATION

The goal of our experimental evaluation is to understand how the addition of the set $\Sigma = \Sigma_R \cup \Sigma_B$ of $K = k_R + k_B$ edges output by REPUBLIC+, run separately with $C = R$ and B , affects the structural bias of the network, by computing the gain in the structural bias reduction. In particular, we measure the gain with $\Delta(G, \Sigma)$, introduced in Sect. 5, used here with a simpler notation. We also measure the change $|\mathcal{P}(G)| - |\mathcal{P}(G_{\text{new}})|$ after adding Σ .

Baselines. We compare REPUBLIC+ to three different baselines (i.e., simplified variants of REPUBLIC+) and to two existing algorithms, described in the following. The first baseline, *PureRandom* (PR) selects the source, and the target, nodes of the new edges uniformly at random from the set $\mathcal{P}_C(G)$ and \bar{C} , respectively. The second baseline *Random Top-N Central Nodes (N-RCN)*, given a parameter $N \in (0, 100)$, sorts the nodes in $\mathcal{P}_C(G)$ by descending

centrality, and picks, uniformly at random, k_C edges with source in the top- N percent of nodes in $\mathcal{P}_C(G)$. The last baseline, *Random Top-N Weighted Central Nodes (N-RWCN)*, differs from N-RCN as the nodes in $\mathcal{P}_C(G)$ are sorted in descending order by $\mathcal{R}(v) \times m_{v,u}$.

We compare REPUBLIC+ also to two existing methods, ROV [26], and node2vec [31]. The ROV algorithm outputs a set of k edges to be added to G to minimize the controversy score (RWC) [29]. The RWC is a metric that characterizes how controversial a topic is by capturing how well separated the two colors are. ROV considers as candidates the edges between the high-degree vertices of each color [26, Algorithm 1]. These edges are sorted by descending impact on the graph controversy score, and the top- k edges are added to the graph. The objective of the comparison between ROV and REPUBLIC+ is to verify whether an algorithm developed to minimize the RWC can be used to minimize the structural bias. *node2vec* is a graph embedding technique that encodes a network in a low-dimensional space retaining characteristics like the nodes' similarity [31]. The generation of the embedding is based on random walks. One of the main applications of node2vec is to employ the embedding as the feature space to train link recommendation algorithms. The goal of comparing node2vec to REPUBLIC+ is to understand how the predictions of widely-used link recommendation algorithms affect the network's structural bias. In the experiments, we create for each network a 128-dimensional space, then we train a logistic regression (avg. AUC 85%) over these features, and we predict the existence probabilities of edges from $\mathcal{P}(G)$. We add to the graph the top k edges according to these probabilities.

Datasets. We create graphs obtained from *Wikipedia*, *Amazon*² and *PolBlogs*³. Table 1 shows the relevant statistics.

From *Wikipedia* we consider four bi-partitioned subgraphs related to controversial topics: *politics*, *abortion*, *guns* and *sociology* [44]. Each node in the graph is a page, and is assigned to one color according to Wikipedia's categorization. Directed edges denote links, and are weighted using Wikipedia's clickstream data.⁴

The *Amazon* dataset contains metadata about *books* [36]. Given two book categories, the vertices are all the items in those categories, colored accordingly. There is a directed edge (u, v) if v appears in the list of items similar to u . The edge is weighted by v 's sales rank.⁵ We built three graphs by considering pairs of the following categories: *Mathematics & Technology (MaTe)*, *History of Technology & Military Science (MiHi)*, and *Mathematics & Astronomy (MaAs)*.

The *Political Blogs* dataset is a directed network of hyperlinks between weblogs on US politics [1]. Each node represents a blog and is colored according to its political leaning. Links between blogs were automatically extracted from a crawl of the front page of the blog and represent the edges of the graph. Each edge (v, u) has weight proportional to the out-degree of v .

Setup. Given a network, we run REPUBLIC and the other algorithms on that network for increasing values of K , with $K = 1, 2, 4, 6, \dots, 400$ or 2000 for larger graphs (*Sociology* and *Politics*). These values of K represent only a small percentage of the set of

²<https://snap.stanford.edu/data/amazon-meta.html>

³<http://www-personal.umich.edu/~mejn/netdata/>

⁴<https://dumps.wikimedia.org/other/clickstream/>

⁵ Amazon sales rank is a metric of the relationship among products within one category based on their sales performance. It expresses how well a product is selling relative to other products in the same category.

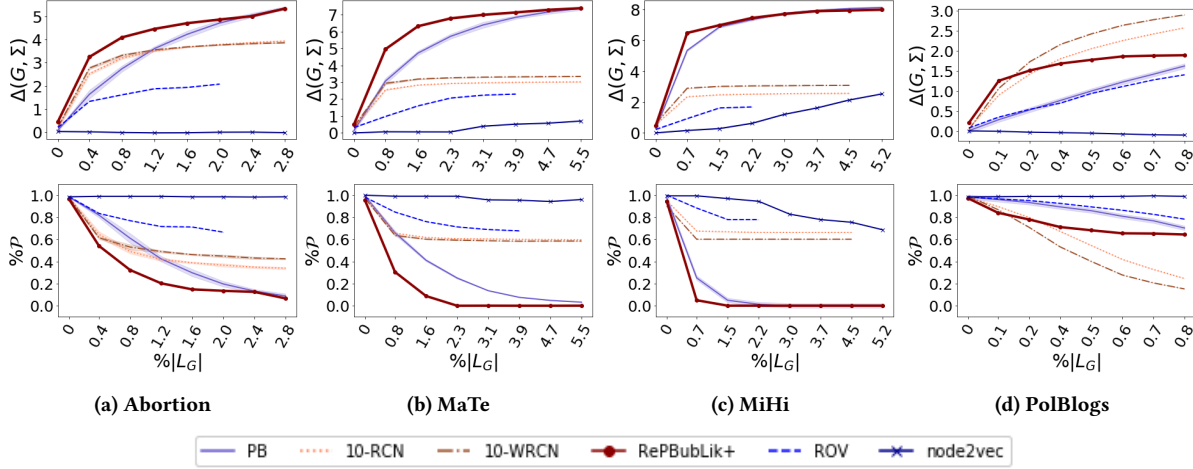


Figure 1: The first row shows the $\Delta(G, \Sigma)$ (y-axis) for increasing value of k , reported in terms of $\%L_G$, the union of possible edges across $\mathcal{P}_C(G)$ and \bar{C} for $C \in R, B$, (x-axis) for each algorithm. Higher values of Δ show more significant reduction of the structural bias. In the second row, we show the percentage of nodes that are still parochial, $\%P = \frac{|\mathcal{P}(G)| - |\mathcal{P}(G_{\text{new}})|}{|\mathcal{P}(G)|}$ after k additions.

Wikipedia							
Topic	R	B	E _{R→B}	E _{B→R}	E	$\%P_R(G)$	$\%P_B(G)$
Abort.	208	413	80	170	1911	85.56	89.20
Guns	142	118	72	79	723	82.95	71.69
Pol.	10347	10129	17452	16484	141486	25.97	42.36
Sociol.	602	2283	284	192	10514	91.32	96.36
Amazon							
Topic	R	B	E _{R→B}	E _{B→R}	E	$\%P_R(G)$	$\%P_B(G)$
MaTe	827	566	25	42	675	90.91	79.63
MiHi	446	405	66	63	482	58.33	63.46
MaAs	827	294	11	6	680	97.31	95.15
PolBlogs							
Topic	R	B	E _{R→B}	E _{B→R}	E	$\%P_R(G)$	$\%P_B(G)$
Politics	545	488	902	781	17348	87.71	90.37

Table 1: Networks’ statistics. The notation is consistent with the rest of the paper.

possible edges to insert and correspond to the total number of edges to add to the graph. Once we set the value of K , accordingly, we allocate k_B and k_R of the K edge insertions to each color proportionally to the sum of the BRs of the parochial vertices in each color. In particular, we define $Y_C = \sum_{v \in \mathcal{P}_C(G)} B_G^t(v)$, for $C \in R, B$, then $k_B = \left\lceil k \frac{Y_B}{Y_B + Y_R} \right\rceil$ and $k_R = K - k_B$. This allocation strategy is a simple but reasonable heuristic that ensures that more edges are added from nodes whose color is more parochial.

We assign the weight $m_{v,u} = 1/(d(v)+1)$ to the added edge (v, u) , where $d(v)$ is the out-degree of v before the insertion, and then we re-normalize the weights of the other edges by multiplying each of them by $1 - m_{v,u}$. Furthermore, we set $r = 5$ and $b = 2$. Moreover, for the algorithms picking the top- N central nodes $N = 10$ To

account for variability of the algorithm, we run them 10 times. The variance of the results is low, overall.

Reproducibility. The code for our experiments is available from <https://github.com/CriMenghini/RePBubLik>.

Experiment results. In Fig. 1, the plots in the first row show how the structural bias is affected by the insertion of an incrementally larger set of edges, while the ones on the second row show the reduction in the number of parochial nodes. Each curve in the plot illustrates the gain by a different algorithm. We can draw the following observations. (1) REPBUBLIK+ performs better than the baselines and the competitors, especially after the insertion of a few edges, as they obtain much larger gain with fewer insertions, i.e., the average BR of parochial nodes decreases faster requiring less modifications. (2) N-RCN, N-WRC, and ROV after a certain point become flat. (3) Overall, REPBUBLIK+ is the best algorithm. (4) The values of REPBUBLIK+ and PR converge, at different speed, to the same value when we add more edges. (5) node2vec, in the best cases, shows little improvement of the structural bias that, in the remaining cases, stays flat or even increases. We now explain these behaviours using the plots on the second row of Fig. 1.

(1) REPBUBLIK+ chooses edges that directly affect the BR of central nodes and, with a chain effect, the BR of nodes connected to them. More central are the nodes we attach the edges to, higher the structural bias drop is. In fact, it follows, as shown for all the networks, that the addition of even small set of edges is very effective. Additionally, we observe that the structural bias reduction corresponds to a significant drop of the number of parochial nodes.

(2) N-RCN, N-WRC, and ROV attach edges only to a subset of $\mathcal{P}(G)$ and as k increases, so does the probability of adding multiple edges to the same nodes. These facts imply respectively that, especially on disconnected graphs (see MiHi in Fig. 1c), the addition of edges may affect few nodes, and that even the insertion of more edges does not modify the set of nodes on which the new edges have effect. Thus, the curves of N-RCN, N-WRCN and ROV reach

an early saturation that expresses the scarce impact of subsequent edge additions. This explanation is confirmed by the percentage of parochial nodes, which does not decrease after the saturation point. Furthermore, the ROV shows a stepping behaviour due to it selecting edges between high-degree central nodes that minimize the RWC without imposing diversity constraints on nodes. And resulting in many selected edges being attached to the same node. Last, we see that on *Polblogs* the best algorithms are N-RCN, N-WRCN. This surprising superiority of the random approaches can be explained by the fact that *Polblogs* is a connected graph, thus edges added to the top-central nodes potentially affect all the nodes in $\mathcal{P}(G)$. Thus, even when N-RCN and N-WRCN add multiple new edges to the same set of nodes, Δ continues to increase.

(3) REPUBLIC+ shows a consistent behaviour, indeed it increases the gain faster than other methods, requiring fewer insertions. The penalty factor η allows the algorithm to diversify the set of nodes to which the new edges attach, raising the chances of lowering the BR of a larger number of parochial nodes, thus increasing the gain. This feature is important especially on disconnected graphs, where the vertices in tiny connected components always have lower centrality compared to those in huge ones. More importantly, we observe that the size of $\mathcal{P}(G)$ is often reduced to 0: REPUBLIC+ is able to “heal” all the bad vertices, and if we measured the structural bias on the obtained graph it would be zero.

(4) The variants of REPUBLIC: REPUBLIC+ and PR, pick edges from the same candidate set, thus the more edges they can pick, the more likely they choose edges with similar effect, thus the average parochial nodes’ BR converges. This is the main explanation why the random algorithm performs so well.

(5) Generally, link recommendation algorithms tend to suggest edges between similar nodes. Node2vec captures this similarity through the nodes’ neighborhood. In this context, graphs partitions have high within- and low between-density. Nodes in the same partition then lie close in the embedding space. Edges suggested by node2vec with high probability connect nodes close to each other in the embedding, which often are in the same partition. Thus, node2vec has a hard time reducing the structural bias, and in some cases increases it.

Plots for *guns*, *sociology*, *politics* and *MaAs* show similar behaviour and can be found in App. A.

7 CONCLUSION

We presented REPUBLIC, an algorithm that reduces the structural bias of a graph by adding k edges. Thanks to the monotonicity and submodularity of the objective function, REPUBLIC is able to return a constant-factor approximation using a greedy approach based on a task-specific variant of the random walk closeness centrality. The results of our experimental evaluation show that the edge insertions suggested by REPUBLIC result in a much quicker decrease of the structural bias than existing methods and reasonable baselines.

The functionality of REPUBLIC relies on the existence of an oracle receiving the network and a page in it as input and outputting the transition probabilities of potentially added links to the input page. We leave the question of designing an algorithm which learns such probabilities from data as future direction of this work.

ACKNOWLEDGMENTS

Shahzad Haddadan was supported by NSF Award CCF-1740741. Part of Cristina Menghini’s work was done while visiting Brown University and is supported by the ERC Advanced Grant 788893 AMDROMA@. Matteo Riondato is supported in part by National Science Foundation award IIS-2006765. Eli Upfal was supported in part by NSF awards RI-1813444, and CCF-1740741. We thank an anonymous reviewer for correcting one of our lemmas.

REFERENCES

- [1] Lada A. Adamic and Natalie Glance. 2005. The Political Blogosphere and the 2004 U.S. Election: Divided They Blog. In *Proceedings of the 3rd International Workshop on Link Discovery* (Chicago, Illinois) (*LinkKDD '05*). Association for Computing Machinery, New York, NY, USA, 36–43. <https://doi.org/10.1145/1134271.1134277>
- [2] Leman Akoglu. 2014. Quantifying political polarity based on bipartite opinion networks. In *Eighth International AAAI Conference on Weblogs and Social Media*.
- [3] Aris Anagnostopoulos, Luca Becchetti, Adriano Fazzone, Cristina Menghini, and Chris Schwegelshohn. 2020. Principal Fairness: Removing Bias via Projections. arXiv:1905.13651 [cs.DS]
- [4] Eugenio Angriman, Alexander van der Grinten, Aleksandar Bojchevski, Daniel Zügner, Stephan Günnemann, and Henning Meyerhenke. 2020. Group Centrality Maximization for Large-scale Graphs. In *2020 Proceedings of the Twenty-Second Workshop on Algorithm Engineering and Experiments (ALENEX)*.
- [5] Francesca Arrigo and Michele Benzi. 2016. Edge modification criteria for enhancing the communicability of digraphs. *SIAM J. Matrix Anal. Appl.* 37, 1 (2016), 443–468.
- [6] Francesca Arrigo and Michele Benzi. 2016. Updating and downdating techniques for optimizing network communicability. *SIAM Journal on Scientific Computing* 38, 1 (2016), B25–B49.
- [7] Cigdem Aslay, Antonis Matakos, Esther Galbrun, and Aristides Gionis. 2018. Maximizing the Diversity of Exposure in a Social Network. In *2018 IEEE International Conference on Data Mining (ICDM)*, 863–868.
- [8] Eytan Bakshy, Solomon Messing, and Lada A Adamic. 2015. Exposure to ideologically diverse news and opinion on Facebook. *Science* 348, 6239 (2015), 1130–1132.
- [9] Ruben Becker, Federico Corò, Gianlorenzo D’Angelo, and Hugo Gilbert. 2020. Balancing Spreads of Influence in a Social Network. *Proceedings of the AAAI Conference on Artificial Intelligence* 34, 01 (2020), 3–10.
- [10] Anna Ben-Hamou, Roberto I. Oliveira, and Yuval Peres. 2018. Estimating Graph Parameters via Random Walks with Restarts. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms* (New Orleans, Louisiana) (*SODA '18*). Society for Industrial and Applied Mathematics, USA, 1702–1714.
- [11] Seyla Benhabib. 1996. Toward a deliberative model of democratic legitimacy. In *Democracy and difference: Contesting the boundaries of the political*. Princeton University Press, Princeton, NJ, 67–94.
- [12] Suman K. Bera and C. Seshadhri. 2020. How to Count Triangles, without Seeing the Whole Graph. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (Virtual Event, CA, USA) (*KDD '20*). Association for Computing Machinery, New York, NY, USA, 306–316. <https://doi.org/10.1145/3394486.3403073>
- [13] Elisabetta Bergamini, Pierluigi Crescenzi, Gianlorenzo D’Angelo, Henning Meyerhenke, Lorenzo Severini, and Yllka Velaj. 2018. Improving the betweenness centrality of a node by adding links. *Journal of Experimental Algorithmics (JEA)* 23 (2018), 1–32.
- [14] Hau Chan, Leman Akoglu, and Hanghang Tong. 2014. Make it or break it: Manipulating robustness in large networks. In *Proceedings of the 2014 SIAM International Conference on Data Mining*, SIAM, 325–333.
- [15] Flavio Chierichetti and Shahzad Haddadan. 2018. On the Complexity of Sampling Vertices Uniformly from a Graph. In *45th International Colloquium on Automata, Languages, and Programming* (Prague, Czech Republic) (*ICALP 2018*).
- [16] Uthsav Chitra and Christopher Musco. 2020. Analyzing the Impact of Filter Bubbles on Social Network Polarization. In *Proceedings of the 13th International Conference on Web Search and Data Mining*. ACM.
- [17] Michael D Conover, Jacob Ratkiewicz, Matthew Francisco, Bruno Gonçalves, Filippo Menczer, and Alessandro Flammini. 2011. Political polarization on Twitter. In *Fifth international AAAI conference on weblogs and social media*.
- [18] Alessandro Cossard, Gianmarco De Francisci Morales, Kyriaki Kalimeri, Yelena Mejova, Daniela Paolotti, and Michele Starnini. 2020. Falling into the Echo Chamber: The Italian Vaccination Debate on Twitter. In *Proceedings of the International AAAI Conference on Web and Social Media*.
- [19] Gianlorenzo D’Angelo, Martin Olsen, and Lorenzo Severini. 2019. Coverage centrality maximization in undirected networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 501–508.
- [20] Abhimanyu Das, Sreenivas Gollapudi, and Kamesh Munagala. 2014. Modeling opinion dynamics in social networks. In *Proceedings of the 7th ACM international*

- conference on Web search and data mining. 403–412.
- [21] Anirban Dasgupta, Ravi Kumar, and Tamas Sarlos. 2014. On Estimating the Average Degree. In *Proceedings of the 23rd International Conference on World Wide Web (Seoul, Korea) (WWW '14)*. Association for Computing Machinery, New York, NY, USA, 795–806. <https://doi.org/10.1145/2566486.2568019>
- [22] Erik D Demaine and Morteza Zadimoghaddam. 2010. Minimizing the diameter of a network using shortcut edges. In *Scandinavian Workshop on Algorithm Theory*. Springer, 420–431.
- [23] Ioana Dumitriu, Prasad Tetali, and Peter Winkler. 2003. On Playing Golf with Two Balls. *SIAM J. Discrete Math.* 16 (2003), 604–615.
- [24] Ronald Fagin, Anna Karlin, Jon Kleinberg, Prabhakar Raghavan, Sridhar Rajagopalan, Ronitt Rubinfeld, and Andrew Tomkins. 2001. Random Walks with "Back Buttons". *The Annals of Applied Probability* 11 (06 2001).
- [25] Seth Flaxman, Sharad Goel, and Justin M Rao. 2016. Filter bubbles, echo chambers, and online news consumption. *Public opinion quarterly* 80, S1 (2016), 298–320.
- [26] Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. 2017. Reducing Controversy by Connecting Opposing Views. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining (WSDM '17)*.
- [27] Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. 2018. Political discourse on social media: Echo chambers, gatekeepers, and the price of bipartisanship. In *Proceedings of the 2018 World Wide Web Conference*. 913–922.
- [28] Kiran Garimella, Aristides Gionis, Nikos Parotsidis, and Nikolaj Tatti. 2017. Balancing information exposure in social networks. In *Advances in Neural Information Processing Systems*. 4663–4671.
- [29] Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. 2018. Quantifying controversy on social media. *ACM Transactions on Social Computing* (2018).
- [30] Mouzhi Ge, Carla Delgado-Battenfeld, and Dietmar Jannach. 2010. Beyond Accuracy: Evaluating Recommender Systems by Coverage and Serendipity. In *Proceedings of the Fourth ACM Conference on Recommender Systems (RecSys '10)*. 257–260.
- [31] Aditya Grover and Jure Leskovec. 2016. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*. 855–864.
- [32] Daniel J. Isenberg. 1986. Group polarization: A critical review and meta-analysis. *Journal of personality and social psychology* 50, 6 (1986), 1141.
- [33] Denis Kotkov, Jari Veijalainen, and Shuaiqiang Wang. 2016. Challenges of serendipity in recommender systems. In *WEBIST 2016: Proceedings of the 12th International conference on web information systems and technologies*.
- [34] Srijan Kumar, William L. Hamilton, Jure Leskovec, and Dan Jurafsky. 2018. Community interaction and conflict on the web. In *Proceedings of the 2018 World Wide Web Conference*. 933–943.
- [35] Rob LeFebvre. 2017. Obama Foundation taps social media to fight online echo chambers. (2017).
- [36] Jure Leskovec, Lada A Adamic, and Bernardo A Huberman. 2007. The dynamics of viral marketing. *ACM Transactions on the Web (TWEB)* 1, 1 (2007), 5–es.
- [37] Q Vera Liao and Wai-Tat Fu. 2014. Can you hear me now? Mitigating the echo chamber effect by source position indicators. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*. 184–196.
- [38] Q. Vera Liao and Wai-Tat Fu. 2014. Expert voices in echo chambers: effects of source expertise indicators on exposure to diverse opinions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2745–2754.
- [39] Ahmad Mahmood, Charalampos E Tsourakakis, and Eli Upfal. 2016. Scalable betweenness centrality maximization via sampling. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- [40] Antonis Matakos, Evimaria Terzi, and Panayiotis Tsaparas. 2017. Measuring and moderating opinion polarization in social networks. *Data Mining and Knowledge Discovery* 31 (2017), 1480–1505.
- [41] Antonis Matakos, Sijing Tu, and Aristides Gionis. 2020. Tell me something my friends do not know: diversity maximization in social networks. *Knowledge and Information Systems* 9 (2020), 3697–3726.
- [42] Sourav Medya, Arlei Silva, Ambuj Singh, Prithwish Basu, and Ananthram Swami. 2018. Group centrality maximization via network design. In *Proceedings of the 2018 SIAM International Conference on Data Mining*. SIAM, 126–134.
- [43] C. Menghini, A. Anagnostopoulos, and E. Upfal. 2019. Wikipedia Polarization and Its Effects on Navigation Paths. In *2019 IEEE International Conference on Big Data (Big Data)*. 6154–6156.
- [44] Cristina Menghini, Aris Anagnostopoulos, and Eli Upfal. 2020. Wikipedia's Network Bias on Controversial Topics. <https://arxiv.org/abs/2007.08197>
- [45] Alfredo Jose Morales, Javier Borondo, Juan Carlos Losada, and Rosa M. Benito. 2015. Measuring political polarization: Twitter shows the two sides of Venezuela. *Chaos: An Interdisciplinary Journal of Nonlinear Science* 25, 3 (2015), 033114.
- [46] Elchanaan Mossel and Omer Tamuz. 2017. Opinion exchange dynamics. *Probability Surveys* 14 (2017), 155–204.
- [47] Sean A Munson, Stephanie Y. Lee, and Paul Resnick. 2013. Encouraging reading of diverse political viewpoints with a browser widget. In *Seventh International AAAI Conference on Weblogs and Social Media*.
- [48] Cameron Musco, Christopher Musco, and Charalampos E. Tsourakakis. 2018. Minimizing Polarization and Disagreement in Social Networks. In *Proceedings of the 2018 World Wide Web Conference on World Wide Web - WWW '18*.
- [49] Matti Nelimarkka, Salla-Maaria Laaksonen, and Bryan Semaan. 2018. Social media is polarized, social media is polarized: towards a new design agenda for mitigating polarization. In *Proceedings of the 2018 Designing Interactive Systems Conference*. 957–970.
- [50] Manos Papagelis, Francesco Bonchi, and Aristides Gionis. 2011. Suggesting ghost edges for a smaller world. In *Proceedings of the 20th ACM international conference on Information and knowledge management*. 2305–2308.
- [51] Nikos Parotsidis, Evaggelia Pitoura, and Panayiotis Tsaparas. 2015. Selecting shortcuts for a smaller world. In *Proceedings of the 2015 SIAM International Conference on Data Mining*. SIAM, 28–36.
- [52] Nikos Parotsidis, Evaggelia Pitoura, and Panayiotis Tsaparas. 2016. Centrality-aware link recommendations. In *Proceedings of the Ninth ACM International Conference on Web Search and Data Mining*. 503–512.
- [53] Senni Perumal, Prithwish Basu, and Ziyu Guan. 2013. Minimizing eccentricity in composite networks via constrained edge additions. In *MILCOM 2013-2013 IEEE Military Communications Conference*. 1894–1899.
- [54] Bashir Rastegarpanah, Krishna P. Gummadi, and Mark Crovella. 2019. Fighting Fire with Fire: Using Antidote Data to Improve Polarization and Fairness of Recommender Systems. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining (WSDM '19)*.
- [55] Manoel Horta Ribeiro, Raphael Ottoni, Robert West, Virgílio A. F. Almeida, and Wagner Meira. 2020. Auditing Radicalization Pathways on YouTube. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAT* '20)*. 131–141.
- [56] Ana-Andreea Stoica and Augustin Chaintreau. 2019. Hegemony in Social Media and the effect of recommendations. In *Companion Proceedings of The 2019 World Wide Web Conference*.
- [57] Ana-Andreea Stoica, Jessy Xinyi Han, and Augustin Chaintreau. 2020. Seeding Network Influence in Biased Networks and the Benefits of Diversity. In *Proceedings of The Web Conference 2020*. ACM.
- [58] Ana-Andreea Stoica, Christopher Riederer, and Augustin Chaintreau. 2018. Algorithmic Glass Ceiling in Social Networks. In *Proceedings of the 2018 World Wide Web Conference*. ACM Press.
- [59] Cass R. Sunstein. 2002. The Law of Group Polarization. *Journal of Political Philosophy* 10, 2 (2002), 175–195.
- [60] Hanghang Tong, B Aditya Prakash, Tina Eliassi-Rad, Michalis Faloutsos, and Christos Faloutsos. 2012. Gelling, and melting, large graphs by edge manipulation. In *Proceedings of the 21st ACM international conference on Information and knowledge management*. 245–254.
- [61] Tomasz Waś, Marcin Waniek, Talal Rahwan, and Tomasz Michalak. 2020. The Manipulability of Centrality Measures—An Axiomatic Approach. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*. 1467–1475.
- [62] Scott White and Padhraic Smyth. 2003. Algorithms for Estimating Relative Importance in Networks. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '03)*. 266–275.
- [63] An Zeng, Linyuan Lü, and Tao Zhou. 2012. Manipulating directed networks for better synchronization. *New Journal of Physics* 14, 8 (2012), 083006.

A MISSING PROOFS

We present here the proofs missing from the main body. For convenience, we repeat the statements of the lemmas.

LEMMA 3.1. *Let $z \geq (t'/2\epsilon)^2 \delta^{-1}$. Then*

$$\mathbb{P} \left(\left| \bar{F}(v) - r'(v; S) \right| \geq \epsilon \right) \leq \delta .$$

PROOF OF LEMMA 3.1. We can write

$$r'(v; S) = t' - \frac{1}{|S|} \sum_{w \in S} \mathbb{E}_G \left[T'_w(v) \right] .$$

We apply Chebyshev's inequality to the r.v. $1/z \sum_{i=1}^z \bar{h}_{w_i}$, to bound the deviation from its expectation

$$\frac{1}{|S|} \sum_{w \in S} \mathbb{E}_G \left[T'_w(v) \right] .$$

To get an upper bound to the variance of this r.v., we use the fact that the r.v.'s \bar{h}_{w_i} , $i = 1, \dots, z$, are independent, and, from Popoviciu's

inequality, the fact that each has a variance at most $t^2/4$, as $\bar{h}_{w_i} \in [0, t']$. \square

The following result is used in the proof of Lemma 4.2.

LEMMA A.1 (MARKOV INEQUALITY FOR BOUNDED RANDOM VARIABLES). *Let X be a random variable satisfying $0 \leq X \leq t$. We have:*

$$\mathbb{P}(X \leq k) \leq \frac{t - \mathbb{E}[X]}{t - k} .$$

PROOF. It holds

$$\begin{aligned} \mathbb{E}[X] &= \int_0^k xp(x)dx + \int_k^t xp(x)dx \\ &\leq k(1 - \mathbb{P}(X \geq k)) + t\mathbb{P}(X \geq k) . \end{aligned}$$

Thus,

$$\mathbb{P}(X \geq k) \geq \frac{\mathbb{E}[X] - k}{t - k} ,$$

and

$$\mathbb{P}(X \leq k) = 1 - \mathbb{P}(X \geq k) \leq 1 - \frac{\mathbb{E}[X] - k}{t - k} = \frac{t - \mathbb{E}[X]}{t - k} . \quad \square$$

LEMMA 4.2. *Let $r \in \mathbb{N}$, and consider a user who starts their random walk at $v \in V$ and may either restart their walk from v or hit the back button up to r times. Let \mathcal{T}_v be the random variable denoting the number of steps such user takes to hit a vertex in \bar{C}_v . If $B_G^t(v) \geq t(1 - 1/8r)$, then $\mathbb{P}(\mathcal{T}_v \leq t/2) \leq 1/4$. If instead $B_G^t(v) \leq b$ for some $b > 0$, then $\mathbb{P}(\mathcal{T}_v > 4br) \leq 1/4$.*

PROOF OF LEMMA 4.2. Assume first that $B_G^t(v) \geq t(1 - 1/8r)$. Consider a set of r independent random walkers, w_1, \dots, w_r , each starting from v . We can see the trace of the partial walks taken by our random walker with restarts as the union of the traces of these walkers. The event $\mathcal{E}' \doteq \mathcal{T}_v \leq t/2$ is a strict subset of the event $\mathcal{E}'' \doteq$ "there is (at least) a walker w_i for which $T_v^t \leq t/2$ ", as the condition in \mathcal{E}' implies the condition in \mathcal{E}'' , but not vice versa. Thus, $\mathbb{P}(\mathcal{E}') < \mathbb{P}(\mathcal{E}'')$. By Lemma A.1 we have, for each walker, that

$$\mathbb{P}\left(T_v^t \leq \frac{t}{2}\right) \leq \frac{t - \mathbb{E}[T_v^t(S)]}{t - \frac{t}{2}} \leq \frac{\frac{t}{8r}}{\frac{t}{2}} \leq \frac{1}{4r} .$$

Thus, using the union bound over the r walkers, we get $\Pr(\mathcal{E}'') \leq 1/4$. Equivalently $\mathbb{P}(\mathcal{T}_v \leq t/2) \leq 1/4$.

For the case when $B_G^t(v) \leq b$, using Markov inequality we get $\mathbb{P}(\mathcal{T}_v > 4br) \leq 1/4$. \square

LEMMA 4.5. *Problem 1 is NP-hard and APX-hard.*

PROOF OF LEMMA 4.5. We show an approximation-preserving polynomial time reduction from the minimum set cover problem to Prob. 1. Our reduction does *not* change the cost of the optimal solution, thus maintaining, in addition to NP-hardness, the APX-hardness.

Let $U = \{u_1, u_2, \dots, u_n\}$ be a domain and let $S_1, S_2, \dots, S_m \subseteq U$ be an instance of the set cover problem. We construct an instance of Prob. 1 as follows. Fix $t \geq 3$. Let V be union of the following sets: U , $S = \{s_i\}_{i=1}^m$ representing the sets, $T = \bigcup_{j=1}^m T_j$ where each T_j is a set of $\lceil t/2 \rceil - 1$ distinct vertices, and $\{g\}$. Assume all vertices except g have color red and g is blue. For each $i \in [n]$ and $j \in [m]$, place an edge from u_i to s_j if and only if $u_i \in S_j$. For each $j \in [m]$, using the vertices in T_j , place a path of length $\lceil t/2 \rceil - 1$

going from s_j to g . For each $1 \leq j \leq m$, it holds $B_G^t(s_j) = \lceil t/2 \rceil - 1$, and for each $1 \leq i \leq n$,

$$B_G^t(u_i) = \frac{1}{|\{j : u_i \in S_j\}|} \sum_{j \text{ s.t. } u_i \in S_j} B_G^t(s_j) + 1 = \lceil t/2 \rceil .$$

Clearly the bubble radius of vertices in T is strictly less than $t/2$. Thus the parochial vertices are all and only those in U . Assume there is a polynomial-time algorithm for Prob. 1. For any (optimal) solution $\Sigma \subseteq V \times V$, it holds $B_{G_{\text{new}}}^t(u_i) < t/2$ if and only if Σ contains an edge whose source is in $\{u_i\} \cup \bigcup_{j \text{ s.t. } u_i \in S_j} (\{s_j\} \cup T_j)$, for each $i \in [n]$. The source vertices of the edges in Σ must be distinct, as any solution containing two edges originating from the same vertex cannot be optimal. Denote with Z the set of the source vertices of the edges in Σ . Consider now the solution Σ' obtained by changing (in polynomial time) Σ as follows: 1. each edge in Σ whose source is in T_i is modified to have source s_i , for each $i \in [n]$; and 2. each edge in Σ whose source is $u \in U$ is changed to have source s_j where j is such that $u \in S_j$. Clearly Σ' is still an (optimal) solution to Prob. 1. Let OPT be the set of source vertices of the edges in Σ' . Clearly it must be $\text{OPT} \subseteq S$. We now show that Σ' is an (optimal) solution to Prob. 1 if and only if OPT is such that $\{S_j : s_j \in \text{OPT}\}$ is a minimum set cover for the considered instance. It is evident that $\{S_j : s_j \in \text{OPT}\}$ is a set cover, which can be obtained in polynomial time from Σ' . We now show that this set cover is minimal. Consider now any set cover $Y \subseteq \{S_1, \dots, S_m\}$, and consider the set of edges $\{(s_i, g) : s_i \in Y\}$. Adding these edges to G would result in all the vertices in U to no longer be parochial. This holds in particular for any *minimal* set cover Y , from which we can create an (optimal) solution Σ_Y to Prob. 1. Thus we found a bijection between (optimal) solutions to Prob. 1 and minimal set covers for the considered instance, and computing one from the other can be done in polynomial time, showing the NP-hardness of Prob. 1. The APX-hardness follows because, for any minimum set cover Y , the corresponding optimal solution Σ_Y to Prob. 1, built as above, is such that $|\Sigma_Y| = |Y|$, thus if we had a constant-factor polynomial-time approximation algorithm for Prob. 1 we would have an algorithm with the same properties for the minimum set cover problem. \square

LEMMA 5.2. *Let $v \in \mathcal{P}(G)$, $w \in \bar{C}_v$ and $t' \leq t$. Let G_{new} be the graph obtained after adding $e = (v, w)$ to G , with weight m_e . The gain $\Delta(G, v, e, m_e, t')$ is such that*

$$\left(B_G^t(v) - 1\right) m_e \leq \Delta(G, v, e, m_e, t') \leq \mathcal{F}_{t'}(v) \left(B_G^t(v) - 1\right) m_e .$$

PROOF OF LEMMA 5.2. Consider the probability space of all random walks starting from v in G_{new} and G . We introduce a coupling between these two probability spaces as follows: consider a walk in G_{new} and couple every step of it to an identical step in G . If a walk in G_{new} never traverses (v, w) then the gain function is zero as it gets coupled to the identical walk in G . Assume that the walk in G_{new} traverses (v, w) at the i th step without first visiting a vertex in \bar{C}_v . Before traversing (v, w) , the two identical walks in G_{new} and G have the same probabilities and the above coupling works. We partition the state space by conditioning on the step i as follows:

Let \mathcal{E}_i , $1 \leq i \leq t'$, be the event that the walk in G_{new} traverses (v, w) at step i . Consider all such walks, at step $i-1$ these walks need one more steps to reach the other color, and they are coupled to

walks in G which in expectation need $B_G^{t'-i+1}(v)$ steps to reach \bar{C}_v . Thus, assuming \mathcal{E}_i , the gain in bubble radius is equal to $B_G^{t'-i+1}(v) - 1$.

Using the law of total expectation and summing over all $1 \leq i \leq t'$, we can write

$$\Delta(G, v, (v, w), m_{vw}, t') = \sum_{i=1}^{t'} \left(B_G^{t'-i+1}(v) - 1 \right) \mathbb{P}(\mathcal{E}_i) . \quad (5)$$

The left hand side follows from the fact that $\mathbb{P}(\mathcal{E}_1) = m_{vw}$ and that $B_G^j(v) \geq 1$ for any $1 \leq j \leq t'$. The right-hand side is concluded from the fact that $B_G^{t'-i+1}(v) \leq B_G^t(v)$ and that

$$\sum_{i=1}^{t'} \mathbb{P}(\mathcal{E}_i) = \sum_{i=0}^{t'-1} \Psi_v(i) m_{vw} = \mathcal{F}_{t'}(v) m_{vw} .$$

□

LEMMA 5.3. *Let $e = (v, w)$ be the edge with weight m_e added to G to obtain G_{new} . For any other vertex $u \in \mathcal{P}_{C_v}(G)$, it holds*

$$\Delta(G, u, e, m_e, t) = \sum_{i=1}^{t-2} \left(\Delta(G, v, e, m_e, t-i) \mathbb{P}\left(u \overset{=i}{\rightsquigarrow} v\right) \right) .$$

PROOF OF LEMMA 5.3. Using the law of total expectation, for any graph Z , it holds

$$\begin{aligned} B_Z^t(u) &= \left(\sum_{i=1}^{t-1} \left(i + B_Z^{t-i}(v) \right) \mathbb{P}\left(u \overset{=i}{\rightsquigarrow} v\right) \right) \\ &\quad + \mathbb{E}_Z \left[T_u^t(\bar{C}_v) \mid u \overset{<t}{\rightsquigarrow} v \right] \mathbb{P}\left(u \overset{<t}{\rightsquigarrow} v\right) . \end{aligned}$$

Between G and G_{new} , we are only adding an outgoing edge from v and modifying the weights of the edges outgoing from v , so

$$\begin{aligned} \mathbb{E}_G \left[T_u^t(\bar{C}_v) \mid u \overset{<t}{\rightsquigarrow} v \right] &= \mathbb{E}_{G_{\text{new}}} \left[T_u^t(\bar{C}_v) \mid u \overset{<t}{\rightsquigarrow} v \right] , \\ \mathbb{P}\left(u \overset{<t}{\rightsquigarrow} v\right) &= \mathbb{P}\left(u \overset{<t}{\rightsquigarrow} v\right)_{G_{\text{new}}} , \text{ and } \mathbb{P}\left(u \overset{=i}{\rightsquigarrow} v\right) = \mathbb{P}\left(u \overset{=i}{\rightsquigarrow} v\right)_{G_{\text{new}}} . \end{aligned}$$

Therefore,

$$\begin{aligned} \Delta(G, u, (v, w), m_v, t) &\doteq B_G^t(u) - B_{G_{\text{new}}}^t(u) \\ &= \sum_{i=1}^{t-1} (\Delta(G, v, (v, w), m_v, t-i)) \mathbb{P}\left(u \overset{=i}{\rightsquigarrow} v\right) \\ &= \sum_{i=1}^{t-2} (\Delta(G, v, (v, w), m_v, t-i)) \mathbb{P}\left(u \overset{=i}{\rightsquigarrow} v\right) . \end{aligned}$$

The last step follows from the fact that $\Delta(G, v, (v, w), m_v, 1) = 0$ because $B_Z^1(u) = 1$ for every vertex u of any graph Z . □

We need the following technical result in successive proofs.

LEMMA A.2. *If $B_G^t(v) \geq t/2$ then, for any $t' \leq t$, it holds*

$$\frac{t'}{2} \leq B_G^{t'}(v) \leq t' .$$

PROOF. The rightmost inequality is straightforward from the definition of $B_G^{t'}(v)$. Expanding the definition of $B_G^t(v)$, it holds, for any $t' < t$,

$$\begin{aligned} B_G^t(v) &= t \mathbb{P}\left(v \overset{\geq t}{\rightsquigarrow} \bar{C}_v\right) + \sum_{i=1}^{t-1} i \mathbb{P}\left(v \overset{=i}{\rightsquigarrow} \bar{C}_v\right) \\ &\leq t \mathbb{P}\left(v \overset{\geq t}{\rightsquigarrow} \bar{C}_v\right) + t \mathbb{P}\left(v \overset{t' \leq \cdot \leq t}{\rightsquigarrow} \bar{C}_v\right) + \sum_{i=1}^{t'-1} i \mathbb{P}\left(v \overset{=i}{\rightsquigarrow} \bar{C}_v\right) \\ &= t \mathbb{P}\left(v \overset{\geq t'}{\rightsquigarrow} \bar{C}_v\right) + \sum_{i=1}^{t'-1} i \mathbb{P}\left(v \overset{=i}{\rightsquigarrow} \bar{C}_v\right) . \end{aligned}$$

Thus, since the l.h.s. is at least $t/2$,

$$\frac{t}{2} \leq t \mathbb{P}\left(v \overset{\geq t'}{\rightsquigarrow} \bar{C}_v\right) + \sum_{i=1}^{t'-1} i \mathbb{P}\left(v \overset{=i}{\rightsquigarrow} \bar{C}_v\right) ,$$

i.e.,

$$\frac{1}{2} - \frac{1}{t} \sum_{i=1}^{t'-1} i \mathbb{P}\left(v \overset{=i}{\rightsquigarrow} \bar{C}_v\right) \leq \mathbb{P}\left(v \overset{\geq t'}{\rightsquigarrow} \bar{C}_v\right) .$$

By expanding $B_G^{t'}(v)$ in a similar way, and plugging in the last inequality above, we get

$$\begin{aligned} B_G^{t'}(v) &= t' \mathbb{P}\left(v \overset{\geq t'}{\rightsquigarrow} \bar{C}_v\right) + \sum_{i=1}^{t'-1} i \mathbb{P}\left(v \overset{=i}{\rightsquigarrow} \bar{C}_v\right) \\ &\geq t' \left(\frac{1}{2} - \frac{1}{t} \sum_{i=1}^{t'-1} i \mathbb{P}\left(v \overset{=i}{\rightsquigarrow} \bar{C}_v\right) \right) + \sum_{i=1}^{t'-1} i \mathbb{P}\left(v \overset{=i}{\rightsquigarrow} \bar{C}_v\right) \\ &= \frac{t'}{2} + \underbrace{\left(1 - \frac{t'}{t} \right) \sum_{i=1}^{t'-1} i \mathbb{P}\left(v \overset{=i}{\rightsquigarrow} \bar{C}_v\right)}_{\geq 0} \\ &\geq \frac{t'}{2} . \end{aligned} \quad \square$$

LEMMA 5.4. *Let $v \in \mathcal{P}(G)$. Let $w \in \bar{C}_v$, and assume to add the edge $e = (v, w)$ with weight m_e . It holds*

$$\Delta(G, \mathcal{P}_{C_v}(G), e, m_e, t) \geq \frac{m_e}{2} t^{t-2} (v; \mathcal{P}_{C_v}(G)) .$$

PROOF OF LEMMA 5.4. Using Lemma 5.3, we get

$$\begin{aligned} \Delta(G, \mathcal{P}_{C_v}(G), e, m_e, t) &= \\ \frac{1}{|\mathcal{P}_{C_v}(G)|} \sum_{u \in \mathcal{P}_{C_v}(G)} \sum_{i=1}^{t-2} \Delta(G, v, e, m_e, t-i) \mathbb{P}\left(u \overset{=i}{\rightsquigarrow} v\right) . \end{aligned} \quad (6)$$

It holds from Lemmas 5.2 and A.2 that

$$\Delta(G, v, e, m_e, t') \geq \left(\frac{t'}{2} - 1 \right) m_e \text{ for every } 1 \leq t' \leq t .$$

Using this fact, we can continue from (6) as follows

$$\begin{aligned} & \Delta(G, \mathcal{P}_{C_v}(G), e, m_e, t) \\ & \geq \frac{1}{|\mathcal{P}_{C_v}(G)|} \sum_{u \in \mathcal{P}_{C_v}(G)} \sum_{i=1}^{t-2} \binom{t-i}{2} m_e \mathbb{P} \left(u \overset{=i}{\rightsquigarrow}_G v \right) \\ & = \frac{m_e}{2} \frac{1}{|\mathcal{P}_{C_v}(G)|} \sum_{u \in \mathcal{P}_{C_v}(G)} \sum_{i=1}^{t-2} (t-i-2) \mathbb{P} \left(u \overset{=i}{\rightsquigarrow}_G v \right), \\ & \quad \underbrace{\hspace{10em}}_{r^{t-2}(v; \mathcal{P}_{C_v}(G))} \end{aligned}$$

which concludes the proof. \square

LEMMA 5.5. Consider the set $\mathcal{P}_C(G)$ where C is either color. Among all vertices in $\mathcal{P}_C(G)$ let v and opt be

$$\begin{aligned} \text{opt} &= \arg \max_{u \in \mathcal{P}_C(G)} \Delta(G, \mathcal{P}_C(G), e_u, m_u, t), \\ v &= \arg \max_{u \in \mathcal{P}_C(G)} m_u r^{t-2}(u; \mathcal{P}_C(G)), \end{aligned}$$

where e_u is any potentially inserted edge connecting u to \bar{C}_u , and m_u is its weight.⁶

It holds

$$\Delta(G, \mathcal{P}_C(G), e_{\text{opt}}, m_{\text{opt}}, t) \leq (4\gamma(G) + 1) \Delta(G, \mathcal{P}(G), e_v, m_v, t),$$

where $\gamma(G) = \max_{u \in G} \mathcal{F}_t(u)$.

PROOF OF LEMMA 5.5. It follows from Lemma 5.2 that, for any t' ,

$$\begin{aligned} \Delta(G, \text{opt}, e_{\text{opt}}, m_{\text{opt}}, t') &\leq \mathcal{F}_{t'}(u) \left(B_G^t(\text{opt}) - 1 \right) m_{\text{opt}} \\ &\leq (t' - 1) m_{\text{opt}} \mathcal{F}_{t'}(\text{opt}). \end{aligned}$$

By applying Lemma 5.3 first, and then the above inequality, we get

$$\begin{aligned} & \Delta(G, \mathcal{P}_C(G), e_{\text{opt}}, m_{\text{opt}}, t) \\ & \leq \frac{1}{|\mathcal{P}_C(G)|} \sum_{u \in \mathcal{P}_C(G)} \sum_{i=1}^{t-2} (\Delta(G, \text{opt}, e_{\text{opt}}, m_{\text{opt}}, t-i)) \mathbb{P} \left(u \overset{=i}{\rightsquigarrow}_G v \right) \\ & \leq \frac{1}{|\mathcal{P}_C(G)|} \sum_{u \in \mathcal{P}_C(G)} \sum_{i=1}^{t-3} (\Delta(G, \text{opt}, e_{\text{opt}}, m_{\text{opt}}, t-i)) \mathbb{P} \left(u \overset{=i}{\rightsquigarrow}_G v \right) \\ & \quad + \Delta(G, \text{opt}, e_{\text{opt}}, m_{\text{opt}}, 2) \mathbb{P} \left(u \overset{=t-2}{\rightsquigarrow}_G v \right) \\ & \leq \frac{1}{|\mathcal{P}_C(G)|} \sum_{u \in \mathcal{P}_C(G)} \sum_{i=1}^{t-3} (t-1-i) m_{\text{opt}} \mathcal{F}_t(\text{opt}) \mathbb{P} \left(u \overset{=i}{\rightsquigarrow}_G v \right) + 1 \\ & \leq \frac{1}{|\mathcal{P}_C(G)|} \sum_{u \in \mathcal{P}_C(G)} \sum_{i=1}^{t-2} 2(t-2-i) m_{\text{opt}} \mathcal{F}_t(\text{opt}) \mathbb{P} \left(u \overset{=i}{\rightsquigarrow}_G v \right) + 1 \\ & \leq 2m_{\text{opt}} r^{t-2}(\text{opt}; \bar{C}_{\text{opt}}) \mathcal{F}_t(\text{opt}) + 1 \leq 2m_v r^{t-2}(v; \bar{C}_{\text{opt}}) \gamma(G) + 1 \\ & \leq (4\gamma(G) + 1) \Delta(G, \mathcal{P}_C(G), e_v, m_v, t), \end{aligned}$$

where the last step follows from Lemma 5.4. \square

LEMMA 5.6. Let C be either blue or red and $v, u \in \mathcal{P}_C(G)$, and $w_v, w_u \in \bar{C}_v$, such that $e_v = (v, w_v)$ and $e_u = (u, w_u)$ are not existing edges. Let $\Sigma = \{e_v, e_u\}$. It holds

$$\Delta(G, \mathcal{P}_C(G), e_v, m_{e_v}, t) \leq \Delta(G, \mathcal{P}_C(G), \Sigma, \{m_e\}_{e \in \Sigma}, t), \quad (3)$$

⁶Our assumption on the oracle giving the weight ensures that m_u only depends on u , not on the target of e_u .

and

$$\begin{aligned} \Delta(G, \mathcal{P}_C(G), \Sigma, \{m_e\}_{e \in \Sigma}, t) &\leq \Delta(G, \mathcal{P}_C(G), e_v, m_{e_v}, t) \\ &\quad + \Delta(G, \mathcal{P}_C(G), e_u, m_{e_u}, t). \quad (4) \end{aligned}$$

PROOF OF LEMMA 5.6. Let G_v be the graph after adding only the edge e_v , G_u be the graph after only adding the edge e_u , and G_{vu} be the graph after adding both edges.

We first show the monotonicity of the objective function, i.e., that (3) holds. For any $w \in \mathcal{P}_C(G)$, it holds

$$\begin{aligned} \Delta(G, w, e_v, m_v, t) &\doteq B_G^t(w) - B_{G_v}^t(w) \\ &\leq B_G^t(w) - B_{G_{vu}}^t(w) \\ &\doteq \Delta(G, w, \{e_v, e_u\}, \{m_v, m_u\}, t) \end{aligned}$$

because $B_{G_v}^t(w) \geq B_{G_{vu}}^t(w)$, as adding an edge from u , which is in C_w , to a vertex in \bar{C}_w cannot increase the bubble radius of w . The result generalizes to (3) in a straightforward way.

We now show the sub-modularity of the objective function, i.e., that (4) holds. We start by showing that, for $w \in \mathcal{P}_C(G)$, it holds

$$\begin{aligned} \Delta(G, w, \{e_v, e_u\}, \{m_v, m_u\}, t) &\leq \Delta(G, w, e_v, m_v, t) \\ &\quad + \Delta(G, w, e_u, m_u, t). \end{aligned}$$

With an expansion of the definition and a slight rearrangement of the terms, the above inequality is equivalent to

$$\underbrace{B_{G_v}^t(w) - B_{G_{vu}}^t(w)}_{\Delta(G_v, w, e_u, m_u, t)} \leq \underbrace{B_G^t(w) - B_{G_u}^t(w)}_{\Delta(G, w, e_u, m_u, t)},$$

i.e., the gain of adding the same edge (in this case e_u) is smaller when the edge is added to a graph (in this case G_v) that has a superset of the edges (compared to G).

Consider all the walks from w that either pass through v or u . Among such walks, let \mathcal{E}_v be the event of seeing v first and \mathcal{E}_u be the event of seeing u first. If a walk does not pass through either v or u , its probability of hitting the other color is the same in all graphs, as the graphs differ only in the outgoing edges from these two nodes and their weights. For the same reason, $\mathbb{P}(\mathcal{E}_v)$ and $\mathbb{P}(\mathcal{E}_u)$ do not change across the graphs. Thus,

$$\begin{aligned} & B_{G_v}^t(w) - B_{G_{vu}}^t(w) = \\ & (\mathbb{E}_{G_v} [T_w^t(\bar{C}_w) | \mathcal{E}_v] - \mathbb{E}_{G_{vu}} [T_w^t(\bar{C}_w) | \mathcal{E}_v]) \mathbb{P}(\mathcal{E}_v) \\ & + (\mathbb{E}_{G_v} [T_w^t(\bar{C}_w) | \mathcal{E}_u] - \mathbb{E}_{G_{vu}} [T_w^t(\bar{C}_w) | \mathcal{E}_u]) \mathbb{P}(\mathcal{E}_u). \end{aligned}$$

Similarly,

$$\begin{aligned} & B_G^t(w) - B_{G_u}^t(w) = \\ & (\mathbb{E}_G [T_w^t(\bar{C}_w) | \mathcal{E}_v] - \mathbb{E}_{G_u} [T_w^t(\bar{C}_w) | \mathcal{E}_v]) \mathbb{P}(\mathcal{E}_v) \\ & + (\mathbb{E}_G [T_w^t(\bar{C}_w) | \mathcal{E}_u] - \mathbb{E}_{G_u} [T_w^t(\bar{C}_w) | \mathcal{E}_u]) \mathbb{P}(\mathcal{E}_u). \end{aligned}$$

We want to show that it holds

$$\begin{aligned} & \mathbb{E}_{G_v} [T_w^t(\bar{C}_w) | \mathcal{E}_v] - \mathbb{E}_{G_{vu}} [T_w^t(\bar{C}_w) | \mathcal{E}_v] \\ & \leq \mathbb{E}_G [T_w^t(\bar{C}_w) | \mathcal{E}_v] - \mathbb{E}_{G_u} [T_w^t(\bar{C}_w) | \mathcal{E}_v]. \quad (7) \end{aligned}$$

and

$$\begin{aligned} & \mathbb{E}_{G_v} [T_w^t(\bar{C}_w) | \mathcal{E}_u] - \mathbb{E}_{G_{vu}} [T_w^t(\bar{C}_w) | \mathcal{E}_u] \\ & \leq \mathbb{E}_G [T_w^t(\bar{C}_w) | \mathcal{E}_u] - \mathbb{E}_{G_u} [T_w^t(\bar{C}_w) | \mathcal{E}_u]. \quad (8) \end{aligned}$$

We can write

$$\begin{aligned} & \mathbb{E}_G \left[T_w^t (\bar{C}_w) \mid \mathcal{E}_v \right] \\ &= \sum_{i=1}^t \left(i + \mathbb{E}_G \left[T_v^{t-i} (\bar{C}_v) \right] \right) \mathbb{P} \left(w \overset{\bar{i}}{\rightsquigarrow} v \mid \mathcal{E}_v \right) . \end{aligned}$$

The probability on the right is the same on all graphs. Similar expressions hold for

$$\mathbb{E}_{G_v} \left[T_w^t (\bar{C}_w) \mid \mathcal{E}_v \right], \mathbb{E}_{G_u} \left[T_w^t (\bar{C}_w) \mid \mathcal{E}_v \right], \mathbb{E}_{G_{vu}} \left[T_w^t (\bar{C}_w) \mid \mathcal{E}_v \right],$$

and when conditioning on \mathcal{E}_u . To prove (7) and (8), we now show that, for every $t' \leq t$, it holds

$$\mathbb{E}_{G_v} \left[T_v^{t'} (\bar{C}_v) \right] - \mathbb{E}_{G_{vu}} \left[T_v^{t'} (\bar{C}_v) \right] \leq \mathbb{E}_G \left[T_v^{t'} (\bar{C}_v) \right] - \mathbb{E}_{G_u} \left[T_v^{t'} (\bar{C}_v) \right], \quad (9)$$

and

$$\mathbb{E}_{G_v} \left[T_u^{t'} (\bar{C}_u) \right] - \mathbb{E}_{G_{vu}} \left[T_u^{t'} (\bar{C}_u) \right] \leq \mathbb{E}_G \left[T_u^{t'} (\bar{C}_u) \right] - \mathbb{E}_{G_u} \left[T_u^{t'} (\bar{C}_u) \right].$$

We focus on showing (9), as the same steps, with simple modifications, can be followed to show the other inequality. For $Z \in \{G, G_u, G_v, G_{vu}\}$, let \mathcal{A}_Z be the event that a random walk starting

at v reaches u in at most t steps before visiting any vertex in \bar{C}_v , and let $\bar{\mathcal{A}}_Z$ be the complementary event. It holds

$$\mathbb{P}(\mathcal{A}_{G_v}) = \mathbb{P}(\mathcal{A}_{G_{vu}}) \leq \mathbb{P}(\mathcal{A}_G) = \mathbb{P}(\mathcal{A}_{G_u}),$$

due to the insertion of e_v . It also holds

$$\mathbb{E}_{G_v} \left[T_v^{t'} (\bar{C}_v) \mid \bar{\mathcal{A}}_{G_v} \right] = \mathbb{E}_{G_{vu}} \left[T_v^{t'} (\bar{C}_v) \mid \bar{\mathcal{A}}_{G_{vu}} \right],$$

and

$$\mathbb{E}_G \left[T_v^{t'} (\bar{C}_v) \mid \bar{\mathcal{A}}_G \right] = \mathbb{E}_{G_u} \left[T_v^{t'} (\bar{C}_v) \mid \bar{\mathcal{A}}_{G_u} \right],$$

Using the law of total expectation (across \mathcal{A}_Z and $\bar{\mathcal{A}}_Z$) and applying these facts, we can rewrite (9) as

$$\begin{aligned} & \left(\mathbb{E}_{G_v} \left[T_v^{t'} (\bar{C}_v) \mid \mathcal{A}_{G_v} \right] - \mathbb{E}_{G_{vu}} \left[T_v^{t'} (\bar{C}_v) \mid \mathcal{A}_{G_{vu}} \right] \right) \mathbb{P}(\mathcal{A}_{G_v}) \\ & \leq \left(\mathbb{E}_G \left[T_v^{t'} (\bar{C}_v) \mid \mathcal{A}_G \right] - \mathbb{E}_{G_u} \left[T_v^{t'} (\bar{C}_v) \mid \mathcal{A}_{G_u} \right] \right) \mathbb{P}(\mathcal{A}_G) . \end{aligned}$$

The differences between parentheses have the same value, as their corresponding terms have the same values. The inequality holds because $\mathbb{P}(\mathcal{A}_{G_v}) \leq \mathbb{P}(\mathcal{A}_G)$ due to the insertion of e_v in G to obtain G_v . □

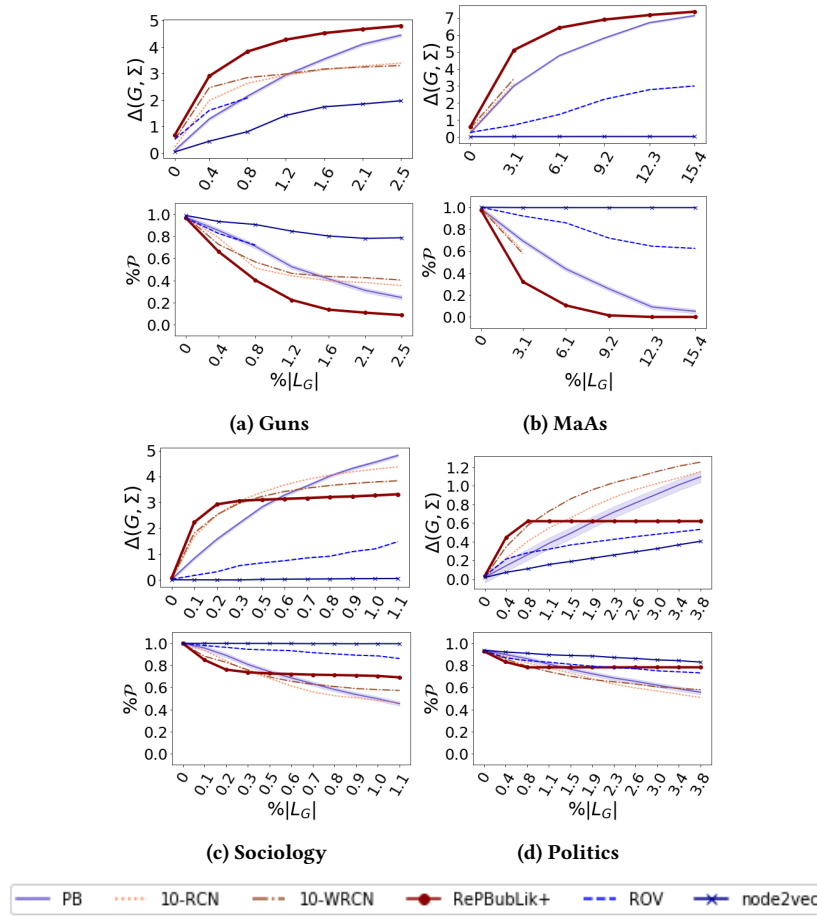


Figure 2: The first row shows the $\Delta(G, \Sigma)$ (y-axis) for increasing value of k , reported in terms of $\%L_G$, the union of possible edges across $\mathcal{P}_C(G)$ and \hat{C} for $C \in R, B$, (x-axis) for each algorithm. Higher values of Δ show more significant reduction of the structural bias. In the second row, we show the percentage of nodes that are still parochial, $\%P = \frac{|\mathcal{P}(G)| - |\mathcal{P}(G_{\text{new}})|}{|\mathcal{P}(G)|}$ after k additions.