Optimizing Diverse Information Exposure in Social Graphs

Jonathan Colin
Université Paris-Saclay, CNRS, LISN
Gif-sur-Yvette, France
jonathan.colin@universite-paris-saclay.fr

Silviu Maniu
Univ. Grenoble Alpes, CNRS, Grenoble INP, LIG
Grenoble, France
silviu.maniu@univ-grenoble-alpes.fr

***Abstract*—The popularity of online social networks and the social interactions they allow has brought great benefits in terms of ease of communication, allowing them to hold a major role in the dissemination and consumption of information. Users of media can be exposed to a wide range of opinions, either actively or passively. Recommendation systems have been developed to steer users towards like-minded content, to the detriment of new, niche, or diverse content. This can lead to fake news, filter bubbles, and opinion polarization.**

In this paper, we introduce a framework to achieve better diversity in social networks, by formulating *information exposure diversity* as an optimization problem in which local modifications on the graph, via edge additions, have the objective of maximizing a target diversity metric from the point of view of an user in the network. We formalize the notion of information exposure linking it to well-studied models in the literature, and provide several algorithms for solving this problem, by leveraging gradient descent-based approaches and greedy algorithms. We show experimentally that our algorithms achieve better diversity measures than state-of-the-art algorithms, on a varied set of real-world graphs.

***Index Terms*—graph optimization, diversity and polarization, social networks, graph data mining**

I. Introduction

Social media is now ubiquitous in our society and has become the main source of information for an increasing number of people; in this sense, it holds a major role in the broadcast of information [1]. At the same time, exposure to social media effectively influences our opinions either *actively* (when we consume media) or even *passively* (when we browse content). This means that any content that a user merely *sees* (or is *exposed* to) also influences their beliefs, giving undue power to the algorithms that recommend content on online platforms, whose aim is to *maximize user engagement* – by only providing content that is similar to the users' current interests, while actively *hiding* content that is too dissimilar. Due to time constraints or even simplicity of use, user often rely on recommendations only to discover new interests, opinions or contents; in turn, these recommendation algorithm do not change the *exposure to information* of the user.

All these dynamics, contribute to an environment that naturally facilitates the spread of *fake news*, as no counter points are readily available and the shock value maintains user attention; as well as the formation of *echo chambers* when users tend to aggregate around common interests, or *filter bubbles* when recommender systems are over tuned towards similarity-based recommendations. This is problematic as this can lead towards segregation of opinions and mass polarization.

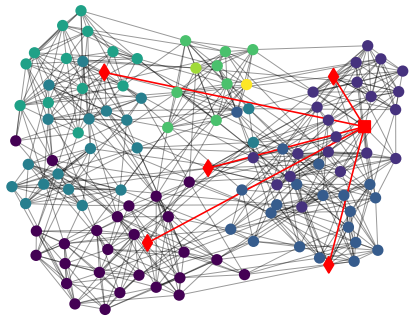

Fig. 1: Example of diverse recommendations (Football dataset). Node colors represent communities as computed by a greedy modularity algorithm. The red square is the target node and red diamonds are diverse recommendations, i.e., from other communities.

Research to measure the full extent of the possible harmful consequences is ongoing [2]–[5]; such research is important, especially since society is slowly letting go of traditional news outlets in favour of news propagated through social media [4], [6]. Algorithmically, ways to optimize diversity and reduce polarization and disagreement in social graphs have been proposed: either based on a global measure of polarization [7]–[9], optimizing exposure to a variety of opinions [10], or by assuming that reducing overall distances in the graphs will increase diversity [11].

Most of these approaches rely on reconfiguring the graph at a *global* level, thus assuming that one can make changes affecting any user in the graph. We believe that *local* optimization, i.e., from the point of view of a single user, may allow to provide more personalized recommendations, minimize the global effect in the graph (indeed, other users might not desire diversity); in a word, increase effectiveness and chance of adoption. To achieve this, we need to provide complementary semantics to how to find new content in social networks. Such semantics can be based on maximizing a measure of *diversity* of the content that a given user is exposed to; indeed,

if we consider that opinions are clustered into *communities* one way to see diversity is to promote the exposure of users to content from communities different from their own. An example is given in Fig. 1: when the node identified by the red square gets linked to the red diamond nodes — part of other communities —, it is directly exposed to more, diverse, sources of information. Intuitively, it also is a way to increase the *entropy of the information exposure* from the node's point of view.

Contributions: In this paper, we present an approach to optimizing the *diversity* of *information exposure* in the context of social networks based on a general formulation of exposure as the distribution of opinions (represented as non necessarily distinct communities) a given node u is exposed to. This allows us to: (i) be flexible in the definition of the exposure function, and we take inspiration from well-studied approaches such as the *polarization* of opinions in the graph [7], but also from classic random-walk and distance-based metrics; (ii) going beyond polarization, to model diversity of exposure over more than 2 polarized opinions; and (iii) define diversity not only as targeting uniform distributions of opinions but *any* distribution of opinions – crucially allowing us to incorporate user interests in the optimization.

We formulate the setting of opinions as representing partitions (or communities) in the graph, define the diversity of exposure as a function of the distribution of opinions, and link it with entropy-based measures in Section III. We present the matrix formulation of the problem in Section IV and provide a gradient descent algorithm, along with two reasonable greedy-based heuristics. We compare our algorithms to algorithms in the state-of-the-art on a varied set of real-world graphs, covering ground-truth and inferred communities along with non-discrete ones, in Section V.

II. Related Work

Our formulation of exposure of information is similar to [10] where maximizing the exposure to information is also the objective. The model that is studied is the independent-cascade model inspired by influence maximization strategies [12]. Originally aimed at political or marketing campaigns, these approaches can also be extended to the goal of maximizing information spread [13]–[15]. Our approach is based on a variety of models, more notably the Friedkin-Johnsen opinion dynamics model [16], probably the most widely used opinion dynamics model in the social media research sphere. It uses an aggregation similar to that of the (personalized) PageRank algorithm [17] to model the opinion of a node as a function of its neighbors' opinions.

Based on the Friedkin-Johnsen model, research has focused on the effects and mechanisms of *filter bubbles*, *echo chambers* or misinformation spread [3] in social media [1], [5], [18]–[21]. Subsequently, when these processes are in play, a lot of research has focused on the subject of polarization [22], [23], its measurement [24]–[26], and potential algorithms that can alleviate its effects [27]–[29]. Others have also focused on minimizing disagreement [7]–[9]. While effective, most of these approaches only capture online interaction through the lens of a single dimension of opinions, either ranging from -1 to 1, or 0 to 1. While effective for polarization minimization, it does not allow for multiple poles of opinions, arguably a more realistic scenario.

In our model, we consider a multitude of opinions as represented by communities of a graph. As such, our problem resembles that of fairness and diversity in recommender systems: optimizing a range of opinions is akin to presenting fair or diverse recommendations. Indeed, no viewpoint should be privileged or censored (with the exception of harmful discourse). However, there is a difference in how the results are achieved. *Fairness* aims at making all candidates equally viable by removing biases [30]–[32] that may appear in the recommendation production process [33], [34]. *Diversity* aims at producing a wide range of dissimilar recommendations [35]–[37]. These approaches tend to disregard social dynamics and can be conflicting with regular user behavior as they are neither based on any opinion dynamics or cascade model. Another approach is to maximize *serendipity*. Even though definitions of serendipity vary [38], [39], they seem consistent on two main properties: interest and novelty. Content that is unknown (so novel if recommended) yet interesting to a user is considered serendipitous. However, concrete results in this area are quite sparse. Indeed, there is a difficulty of measuring serendipity. Current approaches rely on user inputs; these are quite rare commodities on online platforms as users usually dislike having to share personal thought on a given post – consider the number of likes/dislikes and comments compared to number of views on videos on e.g., YouTube. Moreover, datasets that provide such information are far and few between making benchmarks and comparisons difficult.

Many of the previously mentioned algorithms are also based on a greedy structure which has guarantees when coupled with sub-modular objectives. In this work, we take inspiration from [7], [9] which use a matrix representation of the social network in order to use gradient descent algorithms enabled by the convex nature of the formulation of their problem or of the feasible set. Abebe et. al [40] show that interesting results can also be obtained in a non-convex environment. One strong advantage of such approach is the absence of needing to iterate through a candidate set of recommendations, which saves us the issue of how to conceive it and the overall computational cost.

III. Preliminaries

We represent the social network as a connected directed graph $G = (V, E)$, where V is the set of nodes and $E \subseteq V \times V$ the set of edges, $n = |V|$, $m = |E|$. We use the standard definitions of the adjacency matrix $A(G)$ (with $A_{ij} > 0$ if $(i, j) \in E$ and 0 otherwise) and the degree matrix $D(G)$ (with $D_{ii} = \sum_j A_{ij}$). We denote by $\mathcal{N}_u$ the set of neighbor indices of u in G. When unambiguous, the subscript G is omitted.

To represent the fact that the social network is organized in communities (possibly having different semantics), we define the notion of a *graph partition* $\mathcal{P}$:

Definition 1 (Graph Partition)**.** *A graph partition of a graph G, $\mathcal{P}_G$, is a set of subsets of nodes $\mathcal{X}_i \subset V$:*

$$\mathcal{P}_G = \left\{ \mathcal{X}_i \mid \mathcal{X}_i \subset V, \mathcal{X}_i \neq \emptyset, \bigcup^{i} \mathcal{X}_i = V \right\}. \quad (1)$$

The set of partitions to which a node u belongs is denoted as:

$$\mathcal{P}_G(u) = \{\mathcal{X}_i \mid u \in \mathcal{X}_i\}. \quad (2)$$

Note that this definition allows for overlapping partitions. In matrix form, we represent the partitions as a column-stochastic matrix P of dimension $p \times n$, with the property that $\sum_i P_{ij} = 1,\ \forall j$. The column encode the distribution of communities of opinions an individual node has.

A. Node Exposure

We define next the concept of *exposure to information* and give some examples of how it can be measured in a graph.

Definition 2 (Node Exposure)**.** *The* node exposure *function is a function $\mathcal{E} : V \to \mathbb{R}^n$, representing the distribution of exposure of information from all nodes $v \in V$; moreover, we require that this is a distribution, i.e., $\mathbf{1}^\top \cdot \mathcal{E}(u) = 1,\ \forall u \in V$.*

From node exposure, we can naturally extend it to a *partition exposure* function, $\mathcal{E}_P : V \to \mathbb{R}^p$, such that

$$\mathcal{E}(u \mid \mathcal{P}_G) = \mathcal{E}(u) \cdot P.$$

It is easy to check by the properties of $\mathcal{E}$ and $\mathcal{P}_G$ that $\mathcal{E}(u \mid \mathcal{P}_G)$ is also a distribution.

Exposure Functions: There exist various ways of defining the exposure function $\mathcal{E}$, and we present here some options.

Personalized Random Walk (PPR) This is the classic PageRank random walk process [17], rooted at a node u using transition weight matrix $W_u = \alpha D^{-1}A + (1-\alpha)(\mathbf{1}_n \otimes \mathbf{u})$, where $\mathbf{u}$ is the one-hot vector of u, and $\otimes$ is the Kronecker product. The exposure distribution is then the stationary distribution of this process:

$$\mathcal{E}(u) = W_u^\infty = (1-\alpha)\left(I - \alpha D^{-1}A\right)^{-1}\mathbf{u},$$

where $\mathbf{u}$ is the one-hot encoding of u.

Friedkin-Johnsen Model (F-J) This is a model of opinion dynamics in social networks [16], where the opinion of a node u is an aggregation of opinions of neighbors. It is equivalent to a PPR on the transpose graph, if we consider that the initial vector of opinions is $\mathbf{u}$:

$$\mathcal{E}(u) = W_u^\infty = (1-\alpha)\left(I - \alpha {D^\top}^{-1} A^\top\right)^{-1}\mathbf{u}.$$

Note that when $\alpha = 0.5$ and $D = I$, this is equivalent to the stationary distribution found in [7], [9].

Breadth-First (BFS) This is a simple exposure model, where the exposure of a node u is inversely proportional to the depth at which each node v is found in the process of a BFS search.

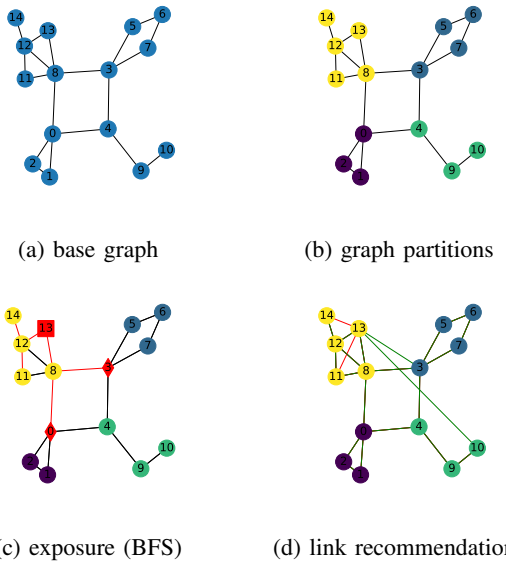

(a) base graph (b) graph partitions

(c) exposure (BFS) (d) link recommendations

Fig. 2: Example social network: partitions, vision, and recommendations.

Example 1. *Fig. 2a exemplifies a small social network of 15 nodes. In this example we consider the point of view of node $u = 13$. The network is divided into 4 partitions as shown in Fig. 2b – in this case, generated using the Louvain modularity algorithm. We exemplify node 13 and the BFS exposure (cut-off at depth 2 for legibility reasons), shown in Fig. 2c. The resulting distribution over the 4 partitions is: $\mathcal{E}(13 \mid \mathcal{P}_G) = \{0.125, 0.125, 0, 0.750\}$. This is because in the search, the blue partition is reached once (node 3 at depth 2), the purple partition once (node 0 at depth 2), the green partition never, and the "home" yellow partition 4 times (nodes 8, and 12 at depth 1, and 12 and 14 at depth 2).*

B. Diversity of Exposure

Given the probability distribution of the exposure function $\mathcal{E}$ we can now define a diversity measure function $\sigma\left(\mathcal{E}\left(u \mid \mathcal{P}_G\right)\right) \in \mathbb{R}$ which quantifies — given a node u in a graph G partitioned according to $\mathcal{P}_G$ — how diverse u's exposure $\mathcal{E}(u \mid \mathcal{P}_G)$ is.

Intuitively, if we view the partitions in $\mathcal{P}_G$ as representing communities of opinion, the aim of diversity of exposure measure σ is to be as close as possible to a distribution in which all opinions / partitions are equally represented.

Entropy: A reasonable candidate measure is then **Shannon's entropy** — which is maximized when the distribution is equi-probable:

$$\sigma\left(\mathcal{E}\left(u \mid \mathcal{P}_G\right)\right) := \mathcal{H}\left(\mathcal{E}\left(u \mid \mathcal{P}_G\right)\right) = -\sum_{\mathcal{X}_i \in \mathcal{P}_G} p(\mathcal{X}_i)\log(p(\mathcal{X}_i)). \quad (3)$$

Entropy makes intuitive sense because of its aim to measure the amount of information in a message. Entropy-based measure have been used before to measure diversity [35] and

fairness [32], [41]. Moreover, [20] showed experimentally that the entropy can also quantify polarization in a graph.

Cross-entropy: Another way to look at the problem is that we are aiming to minimize the distribution distance to a reference desired distribution Q. This can occur when we would like to minimize the exposure to e.g., toxic partitions while keeping the other partitions diverse. This can be measured using the **KL divergence** between $\mathcal{E}$ and Q:

$$\sigma_Q\left(\mathcal{E}\left(u \mid \mathcal{P}_G\right)\right) := \mathcal{D}_{\mathrm{KL}}\left(Q||\mathcal{E}\left(u \mid \mathcal{P}_G\right)\right) = \\ = -\sum_{\mathcal{X}_i \in \mathcal{P}_G} q(\mathcal{X}_i) \log\left(\frac{q(\mathcal{X}_i)}{p(\mathcal{X}_i)}\right). \tag{4}$$

When Q is the equi-probable distribution, Eq. (4) is equivalent to Eq. (3).

Example 2. *We return to the example in Fig. 2. The BFS vision distribution is associated to an entropy value of* 0.736*. When one recommends using classic, triadic closure based, recommendation algorithms, node* 13*'s community becomes denser (edges* $13-14$ *and* $13-11$*, in red in Figure 2d), resulting in a more "biased" exposure:* $\mathcal{E}(13 \mid \mathcal{P}_G) = \{0.1, 0.1, 0, 0.8\}$ *and a smaller entropy value of* 0.639*. On the other hand, when recommending by optimizing entropy-based measures (edges* $13-10$ *and* $13-3$*, in green in Figure 2d) we get a more balanced vision:* $\mathcal{E}(13 \mid \mathcal{P}_G) = \{0.145, 0.285, 0.285, 0.285\}$ *and a higher entropy of* 1.351*. Moreover, we can see that this is achieved by linking* 13*'s original community, yellow, with the green and blue communities.*

IV. Problem and Algorithms

The studied problem can then be stated as follows:

Problem 1 (Single-Node Exposure Diversity Maximization)**.** *Given a graph* $G = (V, E)$*, a node* $u \in V$*, a partition* $\mathcal{P}_G$*, an exposure function* $\mathcal{E}(u)$ maximize *the diversity measure* $\sigma\left(\mathcal{E}\left(u \mid \mathcal{P}_G\right)\right)$ *for a fixed number* k *of* edge additions *to the graph starting at* u*.*

We know that both entropy and KL divergence are maximized at the equi-probable or the Q distributions, respectively. Assuming that we have a set of *feasible graphs* $\mathcal{G}$ (defined later), the optimization problem can be stated as the minimization of the distance between the exposure distribution and these distributions:

$$\arg\min_{G \in \mathcal{G}} ||\mathbf{d}||^2 = ||\mathcal{E}_G\left(u \mid \mathcal{P}\right) - \mathbf{q}||^2 = ||\mathcal{E}_G\left(u\right) \cdot P \cdot \mathbf{u} - \mathbf{q}||^2, \tag{5}$$

where $\mathcal{E}_G$ denotes the exposure function computed on the graph G, $\mathbf{u}$ the one-hot encoding of u, and $\mathbf{q}$ is the desired target distribution.

Note that the formulation in Eq. 5 is similar to the formulation for minimizing polarization in a Friedkin-Johnsen model as the distance between the vector of opinions and the average opinion in the graph, as in [8], [9], [26]. We claim that our formulation is more general, as minimizing the distance to the average opinion risks decreasing polarization but increasing extreme opinions, e.g.,when the average opinion is extreme (note that this can also occur when the average opinion is 0-centered *a posteriori*, as claimed in [8]).

Properties of the objective function: The objective is *non-monotone*. This means that the greedy algorithm is not guaranteed to find the optimal solution; indeed, as we will show in the experimental results, there comes a point of diminishing returns when adding edges to the graph actually *decreases* the diversity of exposure.

Moreover, the objective is not always convex. On the one hand, for the PPR and F-J functions, the objective is *convex*, as it has the same form as the polarization minimization in [7]. On the other hand, for the BFS function – and when expressing the exposure function as a sum of powers of the adjacency matrix – one can see that, from depth $d = 3$ on, the objective is not convex due to at least the term A^3.

A. Convex Set of Feasible Graphs

For all exposure functions, we can apply gradient descent methods. Consider the following set of graphs:

$$\mathcal{G}_u(A) = \{G' \mid A'_{ij} > 0, \forall (i,j) \in E(G), \\ A'_{ij} = 0, \forall (i,j) \notin E(G), i \neq u, \\ \sum_j A'_{ij} = 1, \forall i \in V(G)\}, \tag{6}$$

i.e., all graphs that have their edge weights normalized, and for which $D = I$. Moreover, we allow only graphs that change only in row u (adding new links from u to other nodes). It is straightforward to show that this set is convex: a linear combination of row-stochastic matrices is also row-stochastic; it is also the set defined in [7], [9], restricted to u:

Proposition 1. *The set of feasible graphs* $\mathcal{G}_u(A)$ *is convex.*

By Eq. 5 and 6 we can adapt gradient descent algorithms for Problem 1. For this, we have to define the gradient of the objective function, and the projection step.

Gradient Step: The derivative of the objective depends on $\frac{\partial \mathcal{E}}{\partial A}$ in the following way:

Proposition 2. *Assuming that* $\mathcal{E}$ *is a derivable function of the adjacency matrix* A*, the gradient of the objective function is given by:*

$$\frac{\partial \mathcal{O}}{\partial A} = 2\mathbf{d}\left(\mathbf{u}^\top \otimes P\right)\frac{\partial \mathcal{E}}{\partial A}. \tag{7}$$

Proof Sketch. Follows from writing $\mathcal{O}$ in algebraic form, $\mathcal{O} = \mathbf{d}^\top\mathbf{d}$ and then using the product rule of matrix calculus. □

This general form allows us to "plug-in" any exposure function and its derivative (or simply provide a numerical approximation thereof). For PPR and F-J exposure:

Proposition 3. *The gradient of the objective function for the* ***PPR*** *and* ***F-J*** *exposure functions is given by:*

$$\frac{\partial \mathcal{E}}{\partial A} = -\alpha(1-\alpha)(W_\infty^{-\top} \otimes W_\infty^{-1}),$$

where $W_\infty = I - \alpha A$ *for PPR, and* $W_\infty = I - \alpha A^\top$ *for F-J.*

Proof Sketch. We use the facts that:

$$\begin{aligned} \mathrm{d}W_\infty &= -\alpha \mathrm{d}A, \\ \mathrm{d}\mathcal{E} &= -(1-\alpha)(W_\infty^{-1}\mathrm{d}W_\infty W_\infty^{-1}) \\ &= (1-\alpha)(W_\infty^{-\top} \otimes W_\infty^{-1})\mathrm{d}W_\infty. \end{aligned}$$

The final formula follows. □

We can express the BFS exposure function as a sum of powers of A, limited to a depth of 3 for practical reasons:

$$\mathcal{E}^{(\mathrm{BFS})} = \sum_{i=1}^{3} A^i = A(I - A^3)(I - A)^{-1}.$$

Proposition 4. *The gradient of the objective function for the **BFS** exposure function at depth 3 is given by:*

$$\begin{aligned} \frac{\partial \mathcal{E}}{\partial A} &= I + (I \otimes A + A^T \otimes I) + \\ &\quad + ({A^T}^2 \otimes I + A^T \otimes A + I \otimes A^2). \end{aligned}$$

Proof Sketch. This follows from developing the $\mathcal{E} = \sum_{i=1}^{3} A^i$ as $\mathrm{d}\mathcal{E} = \mathrm{d}A + \mathrm{d}A^2 + \mathrm{d}A^3$ and applying the rules of derivatives of powers of matrices. □

Note that we could straightforwardly expand this to higher powers, but it would require both more involved formulas and would cost more computation time.

Once we have the matrix-function gradients, we can directly apply projected gradient descent algorithms, alternating gradient descent and projection in the feasible set (i.e., a modified graph). Algorithm 1 details the steps.

Algorithm 1: DESCENTDIVERSE

Data: graph G, target node u, exposure function $\mathcal{E}$, graph partition $\mathcal{P}_G$, target distribution $\mathbf{q}$, number of edges k, rate η

$G^{(1)} \leftarrow G$, $A^{(1)} \leftarrow A$;
for $i \in \{1, \dots, k\}$ **do**
  compute $\mathcal{O}_i = ||\mathcal{E}_{G^{(i)}}(u \mid \mathcal{P}) - \mathbf{q}||^2$;
  gradient: $\Delta A^{(i)} = \frac{\partial \mathcal{O}_i}{\partial A}$ according to Eq. 7;
  descent: $A' \leftarrow A^{(i)} - \eta \Delta A^{(i)}$;
  projection: $\{v_i, A'_u\} \leftarrow$ PROJECT(u, A'_u);
  add to graph: $G^{(i+1)} \leftarrow G^{(i)} + (u, v_i)$, $A^{(i+1)} = A$, $A^{(i+1)}_u = A'_u$;
end
return $\{v_1, \dots, v_k\}$

Algorithm 2: PROJECT

Data: target node u, set of neighbor indices $\mathcal{N}_u$, u's row in the modified adjacency matrix A'_u

get the candidate set of nodes which have positive values: $\mathcal{C}(u) = \{v \mid v \notin \mathcal{N}_u, A'_{uv} > 0\}$;
get $v = \arg\max_{v \in \mathcal{C}(u)} A'_{uv}$;
set all other edges to 0: $A'_{uj} = \mathbf{0}$, $\forall j \notin \mathcal{N}_u \cup \{v\}$;
re-normalize the edge weights: $A'_{ui} = 1/|\mathcal{N}_u \cup \{v\}|$;
return $\{v, A'_u\}$;

Projected gradient descent algorithms are used in online convex optimization, and are know to have a regret of $O(\sqrt{k})$ for $\eta = O(1/\sqrt{k})$ [42], *if* the function $\mathcal{O}$ is convex. For us, this is not generally the case, but we show that DESCENTDIVERSE is effective in practice.

Projection Step: The projection step ensures that we only add edges from the candidate node u, and that the edge weights are normalized, thus keeping the modified graph in the feasible set $\mathcal{G}_u$. The steps are detailed in Algorithm 2.

Optimizing the Gradient Computation: The gradient computation is the main bottleneck of Algorithm 1. However, we can note that the derivative of the objective function depends on the whole adjacency matrix A but we only care about the final distribution for node u. The Kronecker product $(\mathbf{u}^\top \otimes P)$ is in fact a (very) sparse matrix, only having P for the coordinates corresponding to u. Moreover, we can simplify each Kronecker product by identifying all indices corresponding to the target node u:

- $(\mathbf{u}^\top \otimes P)$ simply becomes P at indices $[u, u]$
- for **PPR** and **F-J**, we only need the square sub-matrix of size n at $[un, un]$ in the Kronecker product $Q^{-\top} \otimes Q^{-1}$ which becomes $Q^{-\top}_{u,u} Q^{-1}$, hence the local gradient at u becomes (abusing notation):
$$\frac{\partial \mathcal{E}}{\partial A}[u] = -\alpha(1-\alpha)(I - \alpha A)^{-T}_{u,u}(I - \alpha A)^{-1};$$
- for **BFS**, the Kronecker products simplify naturally if we only consider the local gradient at u: (i) $I \otimes A$ becomes A, (ii) $I \otimes A^2$ becomes A^2, (iii) $A^T \otimes I$ and $A^T \otimes A$ disappear in the sum because A's diagonal values are zeros, and (iv) ${A^T}^2 \otimes I$ becomes ${A^T}^2_{u,u}$, leading to:
$$\frac{\partial \mathcal{E}}{\partial A}[u] = I + A + A^2 + {A^T}^2_{u,u}.$$

In order to avoid computing the inverse, we use the Biconjugate gradient stabilized method (BiCGSTAB) to approximate it. This approximation is given by solving for $\mathbf{x}$:

$$W\mathbf{x} = P^+\mathbf{q} + \mathbf{u}, \tag{8}$$

where W is either $I - \alpha A$ or $I - \alpha A^T$, P^+ is the Moore-Penrose inverse of P, and $\mathbf{q}$ the target distribution. The resulting solution corresponds directly to the exposure of node u. This change of shape, from matrix $n \times n$ to a vector of length n, entails minimal changes to Algorithm 1, where the projected gradient descent is performed directly on the node's exposure. No changes are need in the the projection step, Algorithm 2.

Computational complexity: For the DESCENTDIVERSE algorithm, the first computation of the inverse matrix for the PPR and F-J exposure is unavoidable and has a complexity of $O(n^3)$, but the subsequent updates have a complexity of $O(n^2)$, as we only need to update the row and column of the target node u. For the BFS exposure function at depth 3, the initial complexity is also $O(n^3)$, but subsequent computations are $\mathcal{O}(n^2)$, as we only need to update the row and column of the target node u and the candidate node c. The overall complexity is thus $O(n^3 + kn^2)$.

The overall complexity of solving the linear equation system in Eq. 8 using the BiCGSTAB iterative method depends on the size of the system, the density of the graph and the convergence rate which depends on the condition number of $I - \alpha A$. We used the default and standard tolerance of 10^{-6} meaning that the residual error is lower than 10^{-6}. Hence, when using the BiCGSTAB method, the complexity of the gradient computation depends on the sparsity of the networks is $O(kRm)$, where R is the number of rounds to convergence.

B. Candidate Set of Edges

Another option, but which has a high computational cost, is to evaluate the objective over a set of *candidate edges* $\mathcal{C}$, which is potentially as large as all nodes $v \notin \mathcal{N}_u$, where $\mathcal{N}_u$ is the set of neighbors of u.

Under this formulation, Problem 1 can be stated as maximizing over all possible configurations of k edges in $\mathcal{C}$. Our objective function is not general monotone or sub-modular, so we rely on heuristics which add one recommendation at a time. We present here two alternatives to solve this problem: a "partition boosting" algorithm and a greedy algorithm.

Partition Boosting: This algorithm uses the intuition that, at any step, the best edge to add is one to the partition that is "farthest" from the objective $\mathbf{q}$, but only if it is under-represented[1]. Stated otherwise, we choose some partition from the set $\{\mathcal{X} \mid d_i < 0, \mathcal{X} \in \mathcal{P}_G\}$. We also assume that the set of candidates is the union of candidates per partition, i.e., $\mathcal{C} = \cup_i \mathcal{C}(\mathcal{X}_i)$.

Once the partition is chosen, we add the edge in $\mathcal{C}(\mathcal{X}_i)$ that maximizes the objective, and repeat the process until k edges are added, as detailed in Algorithm 3.

Algorithm 3: PARTITIONBOOSTINGDIVERSE

Data: graph G, target node u, exposure function $\mathcal{E}$, graph partition $\mathcal{P}_G$, target distribution $\mathbf{q}$, number of edges k, candidates $\mathcal{C} \subseteq V$

for $i \in \{1, \dots, k\}$ **do**
  $\mathcal{X}_i = \text{CHOOSEPARTITION}(\mathbf{d})$;
  $v_i = \arg\min_{v \in \mathcal{C}(\mathcal{X}_i)} ||\mathcal{E}_{G+(u,v)}(u \mid \mathcal{P}_G) - \mathbf{q}||^2$;
  add edge (u, v_i) to G: $G \leftarrow G + (u, v)$;
end
return $\{v_1, \dots, v_k\}$

There are several ways to choose the partition in CHOOSEPARTITION, and we present three alternatives that we have evaluated experimentally:

1) *the most under-represented partition*: $\mathcal{X}_i = \arg\min_i d_i$;
2) *draw randomly from the under-represented partitions*: $\mathcal{X}_i \sim \mathcal{U}(\{\mathcal{X} \mid d_i < 0\})$;
3) *draw from a soft-max distribution*: $\mathcal{X}_i \sim \text{SoftMax}(\mathbf{d})$.

Greedy Algorithm: The objective function that we have defined, including the entropy and the KL-divergence, is not monotone. There is hence no hope of having an approximation algorithm using the greedy algorithm *in general*. However, it remains an interesting heuristic: at each step, we add the edge that maximizes the marginal gain in the objective function, as detailed in Algorithm 4.

Algorithm 4: GREEDYDIVERSE

Data: graph G, target node u, exposure function $\mathcal{E}$, graph partition $\mathcal{P}_G$, target distribution $\mathbf{q}$, number of edges k, candidates $\mathcal{C} \subseteq V$

for $i \in \{1, \dots, k\}$ **do**
  $v_i = \arg\min_{v \in \mathcal{C}} ||\mathcal{E}_{G+(u,v)}(u \mid \mathcal{P}_G) - \mathbf{q}||^2$;
  add edge (u, v_i) to G: $G \leftarrow G + (u, v)$;
end
return $\{v_1, \dots, v_k\}$

Optimizing the computation: Instead of computing the full objective function at each step, we can compute incrementally the changes in the exposure function.

For **BFS** exposure, instead of computing powers of the matrix A, we can simply update the matrix A considering only the node u and the index of the candidate that has been added in the previous step, c. Once the initial $\mathcal{E}$ has been computed, we can accurately estimate the changes in the exposure function by taking advantage of the following intuitions: (i) at hop 2 we have to add the neighbors of u to c and vice versa, by transitivity, i.e., add the vector A_u to the column A_c and the column $A_c^\top$ to the row A_u; and (ii) at hop 3 we have to add the neighbors of the neighbors of u to the neighbors of c and vice versa.

For **PPR** and **F-J** exposure, we can incrementally update the inverses using the Sherman-Morrison formula.

Computational complexity: For the PARTITIONBOOSTINGDIVERSE and GREEDYDIVERSE algorithms, the complexity is $\mathcal{O}(kn^2)$, as we need to compute the objective function for each candidate node v, by using the optimizations that have been described above. Moreover, we have to do it for each of the candidates at each step of the algorithm. However, the PARTITIONBOOSTINGDIVERSE algorithm only needs to compute for at most $\mathcal{C}_{\max} := \max_i |\mathcal{C}(\mathcal{X}_i)|$ candidates. Hence the complexity is $O(n^3 + k\mathcal{C}_{\max} n^2)$ for PARTITIONBOOSTINGDIVERSE and $O(n^3 + k\mathcal{C}n^2)$ for GREEDYDIVERSE.

V. EXPERIMENTS

A. Experimental Setup

Implementation: We implemented our experiments in Python using the `NetworkX` package for graph manipulation and the computational packages `numpy` and `scipy.sparse`. The GREEDYDIVERSE and PARTITIONBOOSTINGDIVERSE algorithms support multiprocessing for faster running time.[2].

Datasets: We evaluate our approach on several real graph datasets, described in Table I.

The *MovieLens* dataset is built from the Movielens100k dataset, in which we consider only movies. Each movie has a

[1] An interesting option would be to also allow removing edges from over-represented partitions, but this has obvious drawbacks in practice.

[2] https://github.com/Jonathan-COLIN/Optimizing-Diverse-Information-Exposure-in-Social-Graphs

TABLE I: Dataset details. For the dataset marked with †, the communities represent the number of node features, and for datasets marked ‡ the number of communities found by the Louvain algorithm.

Dataset	Nodes	Edges	Existing	#Comm.
Polbooks	105	441	✓	3
Football [43]	115	613	✓	12
MovieLens [44]	1,680	264,718	✗	19†
Reddit [45]	34,671	123,570	✗	54‡
Facebook [46]	63,392	816,886	✗	74‡
Amazon [47]	334,863	925,872	✗	230‡

distribution over 19 movie genres. We assign an edge between two movies if they share more than two thirds of their movie genres. After this linking process, each node's genre vector is updated as the average of all genres of its neighbors. Using these vectors we can then build the partition matrix P, where each row corresponds to a genre and each column to the genre vector of a node. This allows us to exemplify the functioning of our algorithms on *non-discrete partition*.

The *Football* and *Polbooks* datasets have ground-truth communities, respectively the football conferences and the political alignment of the books. Additionally, for all datasets without ground truth we can infer a partition using a community detection algorithm, we used Louvain's algorithm for greedy modularity [48]. For all other datasets, we use the Louvain community detection algorithm to infer the partitions.

a) Target node selection: The optimization process depends greatly on the initial situation of the target node. Indeed, high degree nodes have very different exposures compared to very low degree nodes. As such, in the following, target nodes are selected according to their degree. We batch nodes into three groups of high, medium and low degree and select an equal amount of nodes among these three groups.

b) Generating Non-Uniform Target Distributions: The target distribution $\mathbf{q}$ is the main factor in how the exposure is optimized. Indeed, a standard uniform distribution would lead to an equal representation of each partition and a unit vector used as a distribution would lead to only promoting a single partition. To show the power of *non-uniform target distributions*, we can adapt $\mathbf{q}$ to take user interests into account. We assume that for all users, their initial exposure before the optimization process represents their own interest. From the initial exposure distribution, we scale this distribution by adding the uniform distribution and a normal distribution centered at u's partition. This allows us to promote the partitions that are closer to the user's initial interest and reduce the probability of partitions that are further away. We do this for all datasets.

Targeting non-uniform partitions rather than uniform partitions allows us to present recommendation that are less "extreme" in their diversity, and which would allow the user to be gently nudged towards a more diverse exposure. Furthermore, if a given partition is recognized as harmful, we can improve the target distribution by setting the corresponding value to a very small value.

c) Baselines: We compare our approach to the following baseline algorithms: random assignment of edges (RANDOM), triadic closure (recommending based on common neighbors, TRIADICCLOSURE), polarization and disagreement minimization algorithms in the state-of-the-art, and a global and local diameter minimization algorithm. We compare the baselines with the GREEDY, PARTITIONBOOSTING and DESCENT algorithms.

We used the Polarization and Disagreement minimization algorithm presented in [8], in particular their simpleGreedy method. Given a candidate set of edges, it iteratively select the most promising one with regards to a standard measure of Polarization and Disagreement. Following their implementation we initialize our opinions uniformly between -1 and 1. Since this is a global algorithm, operating on the entire graph, we constrain the candidate set to only include edges starting from the parameter node u. Additionally, we also compare our approach to a graph diameter minimization algorithm as presented in [11]. Their global method identifies center nodes of the graph and connects them together, we adapted it to always include the target node in the set of center nodes. Their local method solves a decision problem using linear programming in order to identify nodes to add. We identify these algorithms by GLOBALDIAMETERREDUCTION and SINGLESOURCEDIAMETERREDUCTION respectively.

All experiments choose 30 target nodes, equally distributed among high-degree, medium-degree and low-degree nodes in their respective graph. We set $\alpha = 0.05$ and $\eta = 0.1$.

B. Results

Ground-truth communities: We start by presenting the results of the graphs for which we have ground-truth communities, *Polbooks* and *Football*, for both uniform and preferential target distributions, seen in Figure 3. Generally, we notice that DESCENTDIVERSE manages to achieve its objective faster than other approaches. Moreover, as expected, TRIADICCLOSURE is among the least effective approaches. Interestingly, RANDOM is quite competitive (and always better than TRADICCLOSURE), which can be attributed to the relatively small sizes of the graphs. In a number of cases, the PARTITIONBOOSTINGDIVERSE does not manage to optimize past a certain point; indeed, once it can no longer find a partition to boost (it looks only for under-represented ones), it is liable to get "stuck" for the rest of the run. On the other hand, the diameter reduction approaches of [11], especially SINGLEDIAMETERREDUCTION, are very competitive, even for preferential distributions.

One interesting, and potentially important, finding can be seen in the results on the BFS exposure function. One can notice that there exists a point in the process where adding new edges becomes counter-productive. This makes intuitive sense, seeing that the objectives are non-monotonous. Moreover, it means that one cannot indiscriminately add edges to maximize diversity, but that there is an optimal number of edges, which will depend on the graph and the node itself.

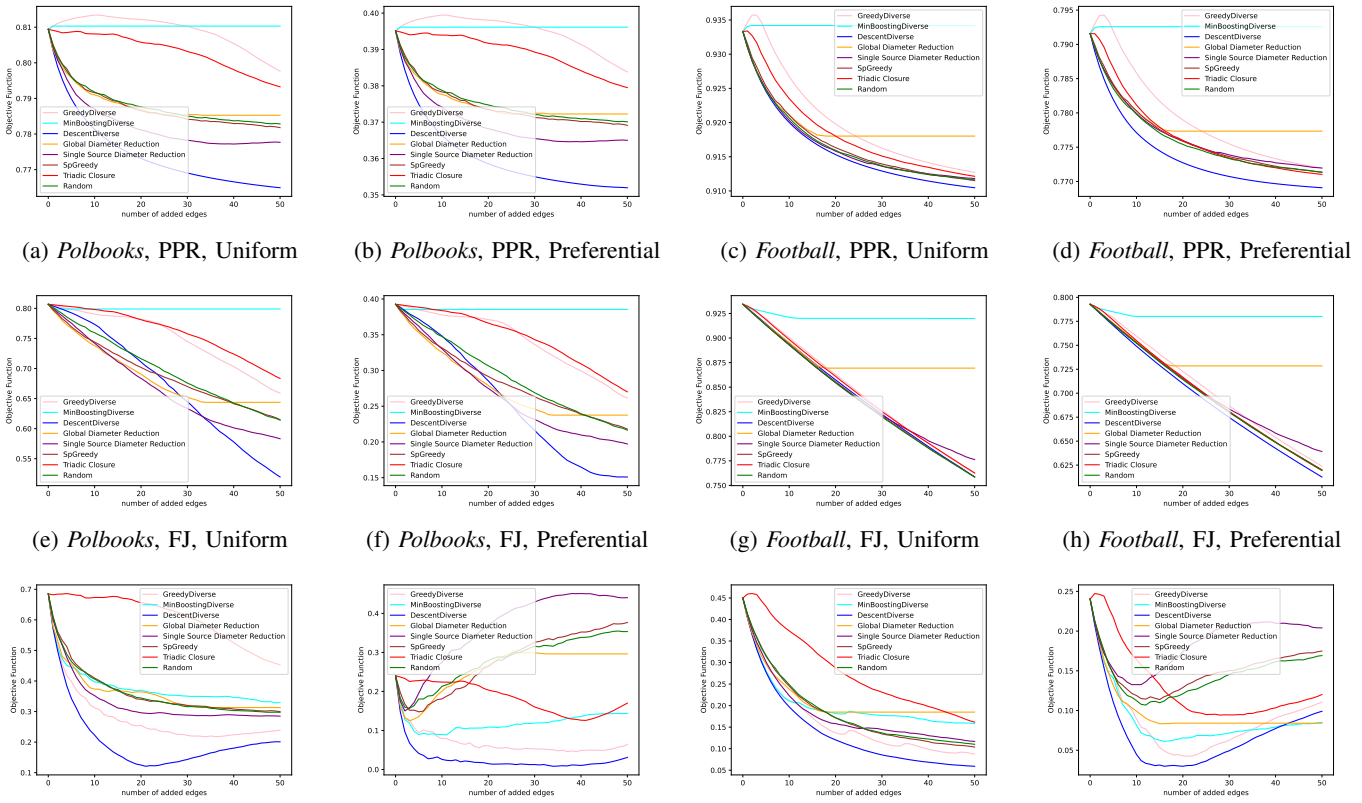

(a) *Polbooks*, PPR, Uniform (b) *Polbooks*, PPR, Preferential (c) *Football*, PPR, Uniform (d) *Football*, PPR, Preferential

(e) *Polbooks*, FJ, Uniform (f) *Polbooks*, FJ, Preferential (g) *Football*, FJ, Uniform (h) *Football*, FJ, Preferential

(i) *Polbooks*, BFS, Uniform (j) *Polbooks*, BFS, Preferential (k) *Football*, BFS, Uniform (l) *Football*, BFS, Preferential

Fig. 3: Objective vs. number of edge additions, ground truth communities.

Inferred communities: Where no ground truth is available, we generated partitions using a community detection algorithm. The results comparing RANDOM and TRIADICCLOSURE to DESCENTDIVERSE are shown in Fig. 4. The results are similar to those on graph having ground truth communities, with DESCENTDIVERSE being the most effective at the end of the recommendation process (even if, in some cases, it is not the best at the start). Interestingly, recommending in the "classic" approach, TRIADICCLOSURE, is almost always worse than RANDOM. This makes some intuitive sense as recommending based on common neighbors would get the user "stuck" in the same community, whereas recommending random links at least gives a chance of escaping their current community and thus increasing the diversity of exposure.

Non-discrete communities: Our formulation also allows for optimizing in cases where partition are non-discrete or overlapping. As detailed above, we adapted the *Movielens* dataset so that each movie is defined by a distribution over its movie genres, and optimize for uniform or preferential exposure. The results are shown in Figure 5. Again, the DESCENTDIVERSE algorithm is the most effective, and TRIADICCLOSE the least effective, in line with the other results.

Running time: The running time of the algorithms, in seconds per iteration, is shown in Fig. 6. Our algorithms (DESCENTDIVERSE, PARTITIONBOOSTINGDIVERSE, GREEDYDIVERSE) have similar running times. The SPGREEDY algorithm is the slowest, but note that we have implemented the exact algorithm in [8], SimpleGreedy, not the fastest, approximate, FastGreedy. The diameter reduction algorithms are fast, and provide a reasonable performance/efficiency trade-off, but only for small graphs. For graphs of a few thousand nodes on, the baselines timed-out but our optimized algorithm is still capable of running on million-node graphs. One interesting results is the peak observed around 10^5 nodes: this is because the dataset used there (not presented here but available in the supplementary material online) is a denser graph, so BiCGSTAB requires more iterations to converge.

VI. CONCLUSIONS AND FUTURE WORK

We presented in this paper a formulation of diversity optimization when keeping into account a node's exposure to information, expressed as a distribution over partitions of opinions in the graph. Our formulation allows to express the problem of local edge additions as an optimization problem, over objectives derived from entropy-based measures. We presented several algorithms to solve the problem, and showed that they outperform state-of-the-art algorithms in terms of diversity measures, and that they are scalable to large graphs.

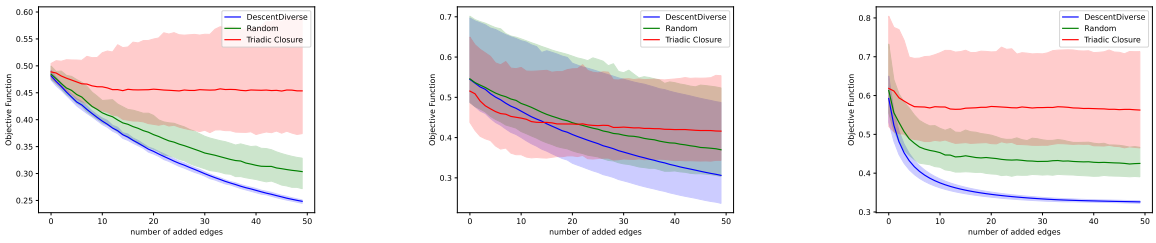

(a) *Reddit*, PPR (b) *Reddit*, F-J (c) *Reddit*, BFS

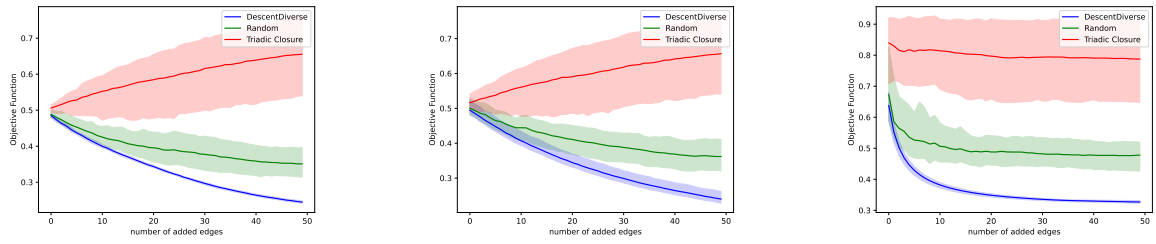

(d) *Facebook*, PPR (e) *Facebook*, F-J (f) *Facebook*, BFS

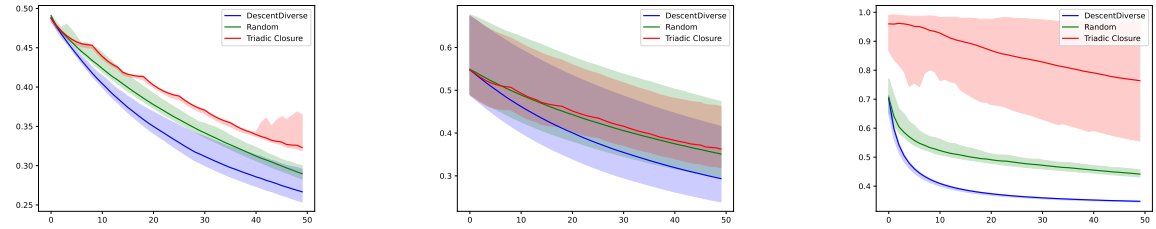

(g) *Amazon*, PPR (h) *Amazon*, F-J (i) *Amazon*, BFS

Fig. 4: Objective vs. number of edge additions, inferred communities, uniform target partition.

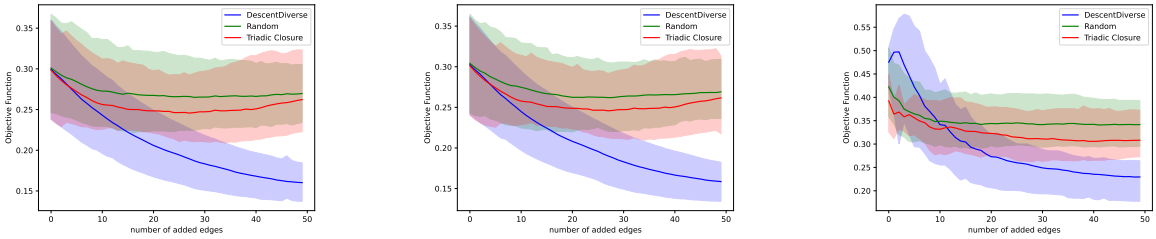

(a) *MovieLens*, PPR (b) *MovieLens*, F-J (c) *MovieLens*, BFS

Fig. 5: Objective vs. number of edge additions, *MovieLens*, non-discrete communities.

The fact that our formulation of the objective function is non-monotone opens up the interesting problem of finding the optimal number of edges to add to maximize diversity. Moreover, one should account for the fact that some nodes might not belong to any opinion partition; it is not immediately apparent what is the best semantics to take into account such nodes.

References

[1] R. Levy, "Social Media, News Consumption, and Polarization: Evidence from a Field Experiment," *American Economic Review*, 2021.

[2] T. T. Hills, "The Dark Side of Information Proliferation," *Perspectives on Psychological Science*, 2019.

[3] A. K. Ghoshal, N. Das, and S. Das, "Influence of community structure on misinformation containment in online social networks," *Knowledge-Based Systems*, 2021.

[4] M. Cinelli, W. Quattrociocchi, A. Galeazzi, C. M. Valensise, E. Brugnoli, A. L. Schmidt, P. Zola, F. Zollo, and A. Scala, "The COVID-19 social media infodemic," *Scientific Reports*, 2020.

[5] A. Ross Arguedas, C. Robertson, R. Fletcher, and R. Nielsen, *Echo chambers, filter bubbles, and polarisation: A literature review*. Reuters Institute for the Study of Journalism, 2022.

[6] A. J. Stewart, M. Mosleh, M. Diakonova, A. A. Arechar, D. G. Rand, and J. B. Plotkin, "Information gerrymandering and undemocratic decisions," *Nature*, 2019.

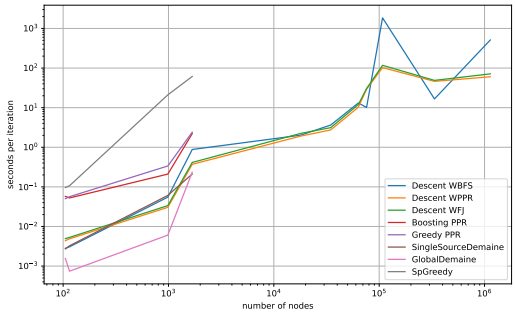

Fig. 6: Running time, in seconds per iteration (node additions), in function of graph size.

[7] C. Musco, C. Musco, and C. E. Tsourakakis, "Minimizing Polarization and Disagreement in Social Networks," in *Proceedings of the 2018 World Wide Web Conference on World Wide Web - WWW '18*, 2018.

[8] L. Zhu, Q. Bao, and Z. Zhang, "Minimizing polarization and disagreement in social networks via link recommendation," *Advances in Neural Information Processing Systems*, vol. 34, 2021.

[9] F. Cinus, A. Gionis, and F. Bonchi, "Rebalancing Social Feed to Minimize Polarization and Disagreement," in *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, 2023.

[10] A. Matakos, C. Aslay, E. Galbrun, and A. Gionis, "Maximizing the Diversity of Exposure in a Social Network," *IEEE Transactions on Knowledge and Data Engineering*, vol. 34, 2022.

[11] E. D. Demaine and M. Zadimoghaddam, "Minimizing the Diameter of a Network Using Shortcut Edges," in *Algorithm Theory - SWAT 2010*, Berlin, Heidelberg, 2010.

[12] M. Azaouzi, W. Mnasri, and L. Ben Romdhane, "New trends in influence maximization models," *Computer Science Review*, 2021.

[13] S. Tu, C. Aslay, and A. Gionis, "Co-exposure maximization in online social networks," *Advances in Neural Information Processing Systems*, 2020.

[14] K. Garimella, A. Gionis, N. Parotsidis, and N. Tatti, "Balancing information exposure in social networks," *Advances in neural information processing systems*, 2017.

[15] J. Yang, Z. Wang, X. Rui, Y. Chai, P. S. Yu, and L. Sun, "Triadic Closure Sensitive Influence Maximization," *ACM Transactions on Knowledge Discovery from Data*, vol. 17, 2023.

[16] N. E. Friedkin and E. Johnsen, "Social Influence and Opinions," 1990.

[17] L. Page, "The pagerank citation ranking: Bringing order to the web. technical report," *Stanford Digital Library Technologies Project, 1998*, 1998.

[18] A. Efstratiou, J. Blackburn, T. Caulfield, G. Stringhini, S. Zannettou, and E. De Cristofaro, "Non-polar Opposites: Analyzing the Relationship between Echo Chambers and Hostile Intergroup Interactions on Reddit," *Proceedings of the International AAAI Conference on Web and Social Media*, 2023.

[19] M. Hohmann, K. Devriendt, and M. Coscia, "Quantifying ideological polarization on a network using generalized Euclidean distance," *Science Advances*, 2023.

[20] K. Sasahara, W. Chen, H. Peng, G. L. Ciampaglia, A. Flammini, and F. Menczer, "Social influence and unfollowing accelerate the emergence of echo chambers," *Journal of Computational Social Science*, vol. 4, 2021.

[21] T. Evans and F. Fu, "Opinion formation on dynamic networks: identifying conditions for the emergence of partisan echo chambers," *Royal Society Open Science*, vol. 5, 2018.

[22] H. A. Prasetya and T. Murata, "A model of opinion and propagation structure polarization in social media," *Computational Social Networks*, 2020.

[23] H. Z. Brooks and M. A. Porter, "A model for the influence of media on the ideology of content in online social networks," *Physical Review Research*, vol. 2, 2020.

[24] F. P. Santos, Y. Lelkes, and S. A. Levin, "Link recommendation algorithms and dynamics of polarization in online social networks," *Proceedings of the National Academy of Sciences*, vol. 118, 2021.

[25] Y. Lelkes, "Mass Polarization: Manifestations and Measurements," *Public Opinion Quarterly*, 2016.

[26] N. Bhalla, A. Lechowicz, and C. Musco, "Local edge dynamics and opinion polarization," in *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*, 2023, pp. 6–14.

[27] K. Garimella, G. De Francisci Morales, A. Gionis, and M. Mathioudakis, "Reducing Controversy by Connecting Opposing Views," in *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, 2017.

[28] F. Fabbri, Y. Wang, F. Bonchi, C. Castillo, and M. Mathioudakis, "Rewiring What-to-Watch-Next Recommendations to Reduce Radicalization Pathways," in *Proceedings of the ACM Web Conference 2022*, Virtual Event, Lyon France, 2022.

[29] R. Interian, J. R. Moreno, and C. C. Ribeiro, "Polarization reduction by minimum-cardinality edge additions: Complexity and integer programming approaches," *International Transactions in Operational Research*, vol. 28, 2021.

[30] J. Chen, H. Dong, X. Wang, F. Feng, M. Wang, and X. He, "Bias and Debias in Recommender System: A Survey and Future Directions," *ACM Transactions on Information Systems*, 2023.

[31] B. Friedman and C. College, "Bias in computer systems," *ACM Transactions on Information Systems*, 1996.

[32] Y. Wang, W. Ma, M. Zhang, Y. Liu, and S. Ma, "A Survey on the Fairness of Recommender Systems," *ACM Transactions on Information Systems*, 2023.

[33] G. K. Patro, A. Biswas, N. Ganguly, K. P. Gummadi, and A. Chakraborty, "FairRec: Two-Sided Fairness for Personalized Recommendations in Two-Sided Platforms," in *Proceedings of The Web Conference 2020*, 2020.

[34] Y. Li, H. Chen, Z. Fu, Y. Ge, and Y. Zhang, "User-oriented Fairness in Recommendation," in *Proceedings of the Web Conference 2021*, 2021.

[35] M. Kunaver and T. Požrl, "Diversity in recommender systems – A survey," *Knowledge-Based Systems*, 2017.

[36] J. Li, T. Cai, K. Deng, X. Wang, T. Sellis, and F. Xia, "Community-diversified influence maximization in social networks," *Information Systems*, 2020.

[37] E. Isufi, M. Pocchiari, and A. Hanjalic, "Accuracy-diversity trade-off in recommender systems via graph convolutions," *Information Processing & Management*, 2021.

[38] D. Kotkov, J. Veijalainen, and S. Wang, "How does serendipity affect diversity in recommender systems? A serendipity-oriented greedy algorithm," *Computing*, 2020.

[39] R. J. Ziarani and R. Ravanmehr, "Serendipity in Recommender Systems: A Systematic Literature Review," *Journal of Computer Science and Technology*, vol. 36, 2021.

[40] R. Abebe, T.-H. H. Chan, J. Kleinberg, Z. Liang, D. Parkes, M. Sozio, and C. E. Tsourakakis, "Opinion Dynamics Optimization by Varying Susceptibility to Persuasion via Non-Convex Local Search," *ACM Transactions on Knowledge Discovery from Data*, 2022.

[41] E. Pitoura, K. Stefanidis, and G. Koutrika, "Fairness in rankings and recommendations: an overview," *The VLDB Journal*, 2022.

[42] M. Zinkevich, "Online convex programming and generalized infinitesimal gradient ascent," in *International Conference on Machine Learning*, 2003, p. 928–935.

[43] M. Girvan and M. E. J. Newman, "Community structure in social and biological networks," *Proceedings of the National Academy of Sciences*, Jun. 2002.

[44] F. M. Harper and J. A. Konstan, "The MovieLens Datasets: History and Context," *ACM Transactions on Interactive Intelligent Systems*, 2015.

[45] S. Kumar, W. L. Hamilton, J. Leskovec, and D. Jurafsky, "Community Interaction and Conflict on the Web," in *Proceedings of the 2018 World Wide Web Conference on World Wide Web - WWW '18*, 2018.

[46] B. Viswanath, A. Mislove, M. Cha, and K. P. Gummadi, "On the evolution of user interaction in Facebook," in *Proceedings of the 2nd ACM workshop on Online social networks*, 2009.

[47] J. Yang and J. Leskovec, "Defining and evaluating network communities based on ground-truth," in *ACM SIGKDD Workshop on Mining Data Semantics*, ser. MDS, 2012.

[48] V. D. Blondel, J.-L. Guillaume, R. Lambiotte, and E. Lefebvre, "Fast unfolding of communities in large networks," *Journal of Statistical Mechanics: Theory and Experiment*, 2008.