

Exposure Inequality in People Recommender Systems: The Long-Term Effects

Francesco Fabbri,^{1,2} Maria Luisa Croci,³ Francesco Bonchi,^{4,2} Carlos Castillo^{1,5}

¹Pompeu Fabra University, Barcelona, Spain

²Eurecat, Barcelona, Spain

³Sapienza University of Rome, Italy

⁴Centai, Turin, Italy

⁵ICREA, Barcelona, Spain.

francesco.fabbri@eurecat.org, marialuisa.croci@gmail.com, francesco.bonchi@eurecat.org, carlos.castillo@upf.edu

Abstract

People recommender systems may affect the exposure that users receive in social networking platforms, influencing attention dynamics and potentially strengthening pre-existing inequalities that disproportionately affect certain groups.

In this paper we introduce a model to simulate the feedback loop created by multiple rounds of interactions between users and a link recommender in a social network. This allows us to study the long-term consequences of those particular recommendation algorithms. Our model is equipped with several parameters to control: (i) the level of homophily in the network, (ii) the relative size of the groups, (iii) the choice among several state-of-the-art link recommenders, and (iv) the choice among three different stochastic user behavior models, that decide which recommendations are accepted or rejected.

Our extensive experimentation with the proposed model shows that a minority group, if homophilic enough, can get a disproportionate advantage in exposure from all link recommenders. Instead, when it is heterophilic, it gets underexposed. Moreover, while the homophily level of the minority affects the speed of the growth of the disparate exposure, the relative size of the minority affects the magnitude of the effect. Finally, link recommenders strengthen exposure inequalities at the individual level, exacerbating the “*rich-get-richer*” effect: this happens for both the minority and the majority class and independently of their level of homophily.

Introduction

Contact recommender algorithms (e.g., “*People You May Know*” in Facebook or “*Who to Follow*” in Twitter) are recognized as key components in any on-line social networking platform: they help the users extending their network faster, thus driving engagement and loyalty (Ricci, Rokach, and Shapira 2011; Barbieri, Bonchi, and Manco 2014; Xie et al. 2016; Aiello and Barbieri 2017). As they affect the “social capital” –the number of followers one user has on these platforms– these algorithms determine to a great extent the exposure and the amount of attention a user receives, potentially strengthening pre-existing inequalities and societal biases. Effects such as the “*algorithmic glass ceiling*” (Stoica, Riederer, and Chaintreau 2018) and more generally,

inequalities in node rankings within and between groups of users (Karimi et al. 2018, Fabbri et al. 2020), are phenomena that may be exacerbated by the recommenders, especially in presence of an homophilic behaviour. Given that people recommenders are adopted not only in recreational social networks, but also in spheres such as employment (Heap et al. 2014; Liu et al. 2016b,a; Ha-Thuc et al. 2016; Domeniconi et al. 2016; Guy, Ronen, and Wilcox 2009) and education (Vassileva, McCalla, and Greer 2016; Zhang et al. 2016), studying their potential biases is of great significance.

In our previous work we highlighted the harmful consequences of link recommenders after one round of recommendations (Fabbri et al. 2020). Such a static picture can be limited as it does not study the consequential effects of the user behaviour which, by accepting or rejecting the recommendations, can determine the future structure of the social network and thus the exposure distribution. Specifically, multiple interactions between users and recommendation algorithms tend to nourish a *feedback loop*: i.e., the output generated by the recommendation algorithm is then fed as future input for the next training of the recommender. In our setting, the recommended new links which are accepted, modify the structure of the network, thus constituting the input for the next cycle of link suggestions. In the context of items recommendation, recent simulation-based studies interested in the side-effects of collaborative filtering algorithms, show how a similar feedback loop (Mansoury et al. 2020) impacts over user preferences, stimulating the *popularity bias* (Yao et al. 2021). Those works underline the importance of providing models able to simulate the potential scenarios which may be otherwise difficult to investigate.

Following a similar approach, but focusing on people recommendations, in this paper we tackle the following research question: “*which impact can link recommendation algorithms have over the network structure and user exposure along multiple rounds of recommendations?*”

Our contribution towards answering this question is a model able to simulate the long-term consequences of the injection of new recommended links into the network, reproducing the feedback loop triggered by the multiple interactions between users and the link recommender. Figure 1 provides a bird’s eye view of our proposed simulation model.

Starting with a social network where the nodes are partitioned in subgroups, e.g., by means of protected attributes

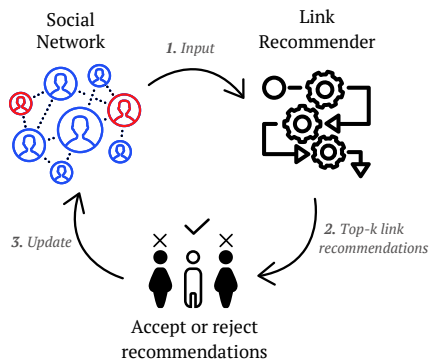


Figure 1: Bird’s eye view of the simulation framework.

such as gender or race (§), a set of different link recommenders are applied to the network to provide, at each iteration, k link recommendations to each user (§). The user at this stage may then decide to accept or reject the recommendation. This decision is governed by three different stochastic user behavior models (§). The rejected links are then discarded, while the new ones are included in the social graph. The new augmented graph will then be the input to the next round of training of the recommender. Despite our models does not consider the organic growth of the social graph, the simulations show that the injection of new links proposed by the recommendation algorithms can move the social graph far from the initial configuration.

Contributions and findings. In this paper we propose a simulation model able to utilize several network configurations, user behaviors, and recommendation models in order to study the long-term effects of people-recommender systems in social networks. We quantify the long-term disparate exposure generated by different initial network topologies, minority size and homophily level, and using different state-of-the-art link recommenders, and different stochastic user behavior models. Our work confirms and extends the preliminary theoretical insights provided by Stoica, Riederer, and Chaintreau (2018) and the empirical results of our previous work (Fabbri et al. 2020), which was limited to one single round of recommendations.

Our findings are summarised as follows:

- Confirming the theoretical findings of Stoica, Riederer, and Chaintreau (2018), our experiments show that, if the minority class is homophilic enough, it can get an advantage in exposure from all link recommenders. If the minority is heterophilic instead, it gets underexposed.
- While the previous observation is robust to all the recommenders, the speed and magnitude of the disparate exposure along time differ across recommenders.
- While the homophily of the minority affects the speed of the growth of disparate exposure, the size of the minority affects its magnitude.
- The user behavior model (how recommendations are accepted or declined) does not impact significantly the evolution of exposure as much as the initial network configuration and the algorithm do.

- Some recommenders can strengthen exposure inequalities at the individual level: after a few iterations, most of the links are recommended towards a small subset of “super-star” nodes. This happens for both the minority and the majority class and independently of their level of homophily. Hence, in the long-term, the “rich-get-richer” effect is exacerbated .

Related Work

In this section we discuss the literature most related to our work. We divide the presentation into two topics: work dealing with inequalities in social networks, and simulation-based studies in recommender systems.

Inequalities in social networks. In our previous work (Fabbri et al. 2020) we observed, in a “static” single round of recommendations, that homophily is a driving force in shifting visibility distribution. In particular, we introduced the concept of disparate visibility in a bi-populated network, showing how effects such bias in rankings and rich-get-richer can get amplified by homophilic networks. The main limitation of our previous study is that it looks at one single round of recommendations, missing the long-term effects.

Stoica, Riederer, and Chaintreau (2018) study the existence of a *glass-ceiling* (an invisible barrier that prevents women from rising in higher rankings) effect amplified by the combination of organic growth and a random walk recommenders. Their preliminary analytical findings inspired us to study, through simulations, different recommenders and distinct network topologies. The main limitation of (Stoica, Riederer, and Chaintreau 2018) is that its theoretical findings require many strong assumptions to hold, such as, the existence of a power inequality between groups in-degree distributions and the use of a specific, random-walk based recommender.

Lee et al. (2019) show that the perceptions about the size of minority groups in social networks can be biased, often exhibiting systematic over- or underestimation. Moreover, these biases can be explained by the level of homophily and by the size of the minority class. Our work, is inspired by their insights, extending the analysis of the inequalities while the network is injected with new links driven by the recommendation output. We confirm their observations, showing how the recommender algorithms can introduce even more inequality along the time.

Tsioutsoulouklis et al. (2020) propose methods for fairness-aware link analysis, introducing techniques able to mitigate unfairness generated by Pagerank. Later, in §, we will show that another popular random-walk based recommender (i.e., SALSA) can increase the unfairness in visibility in the long run, thus confirming the need to devise methods able to mitigate these effects.

Simulation-based studies in recommender systems. Cinus et al. (2022) combine link recommendation and opinion-dynamics models in a simulation-based framework, to assess the effect of people recommenders on the evolution of opinions in a social network. They show that, if the initial network exhibits high level of homophily, people recommenders can help creating echo chambers and polarization.

In the context of collaborative-filtering-based methods, Mansoury et al. (2020) show that popularity bias can be stimulated by feedback loop, where popular items tend to obtain more and more interactions if generated through recommendations. Jiang et al. (2019) propose a theoretical framework to model the effects of “*filter bubble*”, i.e., the tendency of the recommendation algorithm to drive the preferences of the user towards a limited amount of items. Recently, Yao et al. (2021) propose a simulation model for measuring the impact of recommender systems over time, analyzing the changes in the user experience with an application designed for food recommender system.

Our work is motivated by the importance of studying algorithmic bias in recommendations and rankings in the long term, i.e., beyond the single round of algorithmic intervention. In this regard, Ge et al. (2021) have recently introduced the problem of *long-term fairness*, designing also solutions able to account for algorithmic unfairness in the long-term in movies recommendations. Shi et al. (2017) propose a simulation model able to include multiple recommender systems combined with different users choice models, proving that the rich-get-richer effect tends to increase over time, stimulated by the algorithm. In our study we analyze the evolution of rich-get-richer effect in social networks, fueled by the edges created thanks to the recommendation algorithms.

Nguyen et al. (2014) show how in the case of MovieLens data, recommendations generated through a collaborative filtering approach have not strengthened the filter bubble effect. In our paper we go in the opposite direction, showing how homophily may generate biased recommendations, towards a smaller set of recommendations, reducing the diversity of those. Sun et al. (2019) propose a mitigation strategy to reduce popularity bias in recommendations through different methods based on active-learning. The methods proposed are aimed at reducing popularity bias, which in our setting can be related to rich-get-richer effect. Although this paper may be used to propose mitigation strategies to reduce inequality in exposure, the main weakness of their work regards the lack of analysis of different input data distribution. This kind of analysis may help to understand which kind of distribution of user-item interactions may benefit more from their method.

Model

We consider a social graph whose nodes are partitioned by demographics (e.g. gender, age or other characteristics). More formally, let $G = (V, E, \ell)$ be the social graph, where V is the set of nodes, $E \subseteq V \times V$ is the set of directed edges, such that an edge $(u, v) \in E$ indicates the fact that u follows v . Finally $\ell : V \rightarrow \{V_m, V_M\}$ is a labeling function assigning each node to either the minority (V_m) or the majority (V_M) class (with $|V_m| < |V_M|$). We denote by $s_m = |V_m|/|V|$ the fraction of nodes belonging to the class less represented in the network, i.e., the *minority*, and by s_M the fraction of nodes belonging to the majority

Homophily. To capture the bias in the distribution of the edges towards each group, we introduce a measure of homophily, expressed as *the tendency of people in a group to*

connect to individuals in the same group. We model the homophily as the portion of edges distributed within the same group discounted by the fraction observed in a random configuration (Fabbri et al. 2020). More formally:

$$h_i = \frac{|E_{ii}|}{|E_i|} - s_i \quad (1)$$

where $E_{ii} = \{(u, v) \in E | u \in V_i \wedge v \in V_i\}$ and $E_i = \{(u, v) \in E | u \in V_i\}$. This measure ranges in the interval $(-s_i, 1 - s_i]$. A group is called *homophilic* if the tendency to connect to nodes of the same group is stronger than expected ($h_i > 0$), *heterophilic* when this tendency is weaker than expected ($h_i < 0$), and *neutral* if the number of edges within the group is comparable to the relative size of the group ($h_i = 0$).

Step-by-step. In our simulations we reproduce the multiple interactions between the users and the recommendation algorithm, where at each round, a set of new links is recommended to a portion of users randomly sampled from the graph. This sampling represents the fact that only a set of users are online at a certain time, helping reproducing a more realistic scenario. Then the users accepts or rejects the recommendations according to a given stochastic user behavior model. This process is iterated for a given number of iterations. The graph grows accordingly to the new accepted recommendations: neither organic growth, nor edge removal are considered. Table 1 summarizes the simulation process step-by step.

1. **Input.** We start with an initial network configuration with specified levels of homophily and size of the minority class (how this initial configuration is generated by modifying a real-world social graph is presented in §). We also set parameters such as the number of recommendations k that a user receives in a round, the number of iterations T , the fraction α of users to sample, the link-recommendation algorithm A (presented in §), and the stochastic user behavior model B (discussed in §).
2. **Recommendation round.** A link recommender model is trained over the current social graph by the algorithm A . A portion α of users is sampled from the network. Those sampled users receive their top- k recommendations each. The recommendations are links never recommended before and are generated from the set of missing edges at distance two (e.g. “friends of friends”).
3. **Graph update.** Each user decide to accept or reject each of the k recommended links, according to the model B . The social graph is thus updated by adding the newly accepted links. Each link rejected at this stage is discarded and never recommended again.
4. **Repeat.** Steps 2 and 3 are repeated T times.

Table 1: Simulation steps.

For all the results that we report in § we use $T = 20$, $\alpha = 20\%$ and $k = 3$.

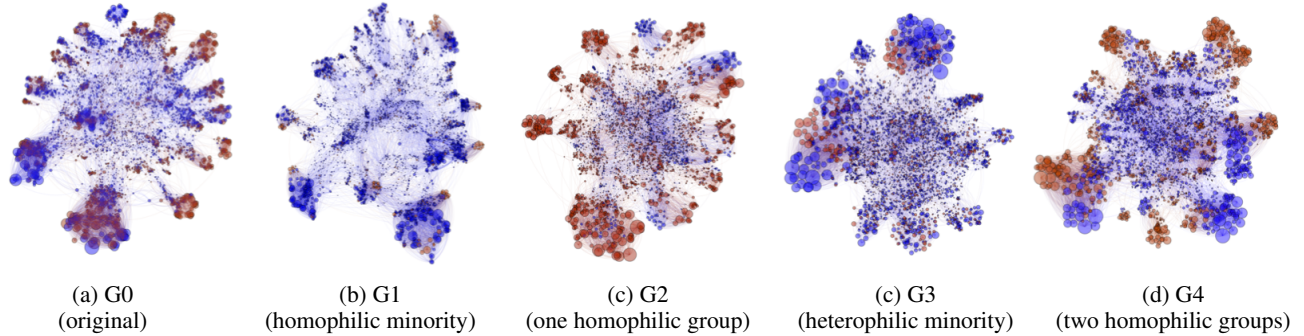


Figure 2: Representation of a sample of each generated network, where the minority is indicated in red, while the majority in blue. Each sample considers 5,000 nodes and those are the ones with highest degree in each group. Specifically for each group, a total of $s_i \times 5,000$ ($i \in \{m, M\}$) nodes are sampled.

We next present in more details the various key components of our model.

Initial Network Configuration

In order to control the level of homophily and the size of the minority class, while keeping a realistic network structure, we propose a novel data generation process which, starting with a real-world bi-populated network, performs just the minimum amount of node class-swappings and link rewirings to match the requested levels of homophily and size of the minority class. In this regard, our networks are *semi-synthetic*. The process is explained in details next.

Our starting real-world network comes from Tuenti, a social network popular few years ago in Spain, which was known as “the Spanish Facebook”. The dataset includes demographic information about users as gender and age (Laniano et al. 2016). The network has 8,983,560 nodes (users) and 17,830,103 edges, where a generic edge (u, v) indicates a user posting on another users’ *wall*. Along the paper, we use a sample of the original network used in our previous work (Fabbri et al. 2020). This sample, which is also the one from which we derive the other configurations, is given by partitioning the users by age (16 as cut-off). We call it **G0**: it contains 500,000 nodes and 2,813,744 edges, with a relevant minority ($s_m = 0.30$) and both groups (majority and minority) appearing to have some level of homophily. Starting from this network, we generate 4 different semi-synthetic networks, ranging different values of h and s . More in details, let V_m and V_M be respectively the set of minority’s and majority’s node in the input network and $N_m = |V_m|$, $N_M = |V_M|$. Let also h_m, h_M indicate the actual homophily of each class, and finally s_m and s_M are the relative size of both classes. The generation process takes as input the desired level of homophily for both classes, denoted h_m^*, h_M^* and the desired proportion for the minority class s_m^* , and works as follows:

1. **Change minority-majority size.** Let $N_m^* = (N_m + N_M)s_m^*$. If $N_m^* < N_m$, then $N_m - N_m^*$ nodes are sampled at random from the minority class V_m and their label is flipped to the majority. Otherwise, we sample of

Graph	Scenario	s_m	h_m
G0	<i>original</i>	0.3	0.42
G1	<i>different sizes + homophilic minority</i>	0.1	0.4
G2	<i>same sizes + homophilic minority</i>	0.45	0.5
G3	<i>different sizes + heterophilic minority</i>	0.3	-0.25
G4	<i>different sizes + homophilic groups</i>	0.3	0.6

Table 2: Table summarizing information about the generated graphs. For each one we have: i) name, ii) scenario characterizing the network

$N_m^* - N_m$ nodes are extracted from the majority V_M and their label flipped to the minority.

2. **Change homophily.** For each group $i \in \{m, M\}$ we first compute the difference between the initial and the final homophily $|h_i^* - h_i| = B_i$. Then depending on the sign of the difference $h_i^* - h_i$ we define which edges need to be rewired. Rewiring an edge (u, v) means substituting (u, v) with an edge (u, w) such that $\ell(v) \neq \ell(w)$. If $h_i^* - h_i > 0$ a sample of edges belonging to $E_{ij} = \{(u, v) \in E | u \in V_i \wedge v \in V_j\}$ is selected and rewired towards nodes in V_i . In this way we increase the set of nodes in E_{ii} , reaching the requested level of homophily. Viceversa, if $h_i^* - h_i < 0$ the operation is the opposite: old edges belonging to the subset E_{ii} are rewired towards nodes in V_j . In both cases, the final amount of edges rewired is $E_i \times B_i$.

Since some recent literature has shown that small subpopulations within a social network can impact the whole graph (Stoica, Riederer, and Chaintreau 2018; Fabbri et al. 2020; Karimi et al. 2018), we generate networks with biased distributions for the minorities. Only for one case, to have a comprehensive analysis, we modify the homophily level in both minority and majority groups.

More in details, these are the configurations we focus on:

- **G1.** To analyze the effect of a **small homophilic minority** we generate a graph with $s_m = 0.1$ and $h_m = 0.4$, with a neutral majority.

- **G2.** To emphasize the role of homophily we also generate a graph with **comparable sizes** between the two groups ($s_m = 0.45$) with the minority strongly homophilic ($h_m = 0.5$) and a majority still neutral.
- **G3.** This configuration is the unique with a **small heterophilic minority** ($h_m = -0.25$) (and neutral majority).
- **G4.** The the final configuration has **both subpopulations homophilic**. In particular, we keep a small minority ($s_m = 0.3$) and both groups with high level of homophily ($h_m = 0.6$ and $h_M = 0.2$).

$G1$ and $G2$ are a useful comparison against $G0$, since they present comparable level of homophily but different sizes of the minority, while $G3$ is useful to explore the heterophilic case and $G4$ resembles a scenario quite common in contexts where phenomena such as polarization and filter bubbles drive the network formation (Garimella et al. 2018). Table 2 summarizes the five networks used in our analysis, while Figure 2 depicts a sample of each network.

Link Recommenders

Link recommendation algorithms are selected accordingly to state-of-the-art performance and popularity in the literature (Li, Fang, and Sheng 2017, Fabbri et al. 2020, Stoica, Riederer, and Chaintreau 2018). In accordance with our data model, all the methods recommend directed links (i.e., who to follow).

ADA: Network Topology Based. Among the different heuristics which aim to define similarity between nodes looking at the graph topology, we select the directed version of the *Adamic-Adar coefficient* (for short “ADA” in the rest of the paper), method that penalizes connections with high degree nodes (Zhao et al. 2020).

SLS: Random Walks Based. As representative of random-walks based approaches, we use SALSA (Stochastic Approach for Link-Structure Analysis; “SLS” in the rest of the paper), which is at the basis of the *who-to-follow* recommender at Twitter (Su, Sharma, and Goel 2016). Recommendation of a generic link is defined as the probability of the source node to jump to the target one, rather than to any other node in the graph.

ALS: Collaborative Filtering Based. Connections among nodes can be considered as implicit feedback in a collaborative filtering approach. An Alternating Least Squares algorithm (“ALS” in the rest of the paper) is selected to perform recommendations (Hu, Koren, and Volinsky 2008). New links are suggested based on latent features extracted from the adjacency matrix.

RND: Random baseline. As baseline, we consider a random recommender (“RND” in the rest of the paper), which picks recommendations uniformly at random from the candidate nodes at distance 2.

User Behavior Models

In order to simulate the user feedback on the received recommendation, we consider three stochastic user behavior models. The first two are adapted from a recent work simulating user-item interactions (Yao et al. 2021), while the third

one defines acceptance probability biased by the position in the ranked list of recommendations. Through these stochastic choice models the users add in expectation one edge per recommended list.

- **B-LZY - Lazy.** The user accepts directly the first recommendation:

$$P(u \text{ selects } v \text{ at position } i) = \begin{cases} 1 & \text{if } i = 1 \\ 0 & \text{otherwise} \end{cases}$$

- **B-RND - Random.** The user picks in the top-k, one recommendation uniformly at random:

$$P(u \text{ selects } v \text{ at position } i) = \frac{1}{k}$$

- **B-PSB - Position Biased.** This policy refers to the idea of having user choices biased by the position bias of rankings, where the user may accept or reject the recommendation with probability based on its position (Craswell et al. 2008). Hence, higher ranked suggestions are more likely to be chosen:

$$P(u \text{ selects } v \text{ at position } i) = \frac{1/\log(i+1)}{\sum_{j=1}^k 1/\log(j+1)}.$$

- **B-MIX - Mixed.** In order to evaluate how heterogeneous user behaviors may affect the exposure distribution, we also include B-MIX, which is a combination of the previous policies. Specifically, at each iteration, each user first picks, uniformly at random, one of the three strategies above, then follows it.

B-PSB model resembles the classical position bias, observed as key factor for predicting clickthrough rates in search engines (Craswell et al. 2008; Joachims 2002). The other two user behavior resembles two extreme situations: i) B-RND is the one less dependent by the order of the recommendation list; while B-LZY represents the case in which the user relies completely on the order imposed by the recommendation algorithm.

Results

In this section we present the results of our experiments, focusing on the key measure that we call *exposure* of the minority, which is simply defined as *the portion of total number of recommendations which suggest a node of the minority*, and denoted \mathcal{E}_m . Note that the total number of recommendations is constant and corresponds to $k|V^\alpha|$. Our code is publicly available at <https://github.com/FraFabbri/long-term-inequality-recsys>.

Exposure in The Long-Run

Figure 3 shows the trend of the exposure of the minority for each of the four networks ($G0$ is omitted for space limitation as it presents results almost indistinguishable from $G4$). For each experiment, we track the exposure and the percentage of new edges added at each iteration with respect the original network. The dashed line represents in each plot the relative

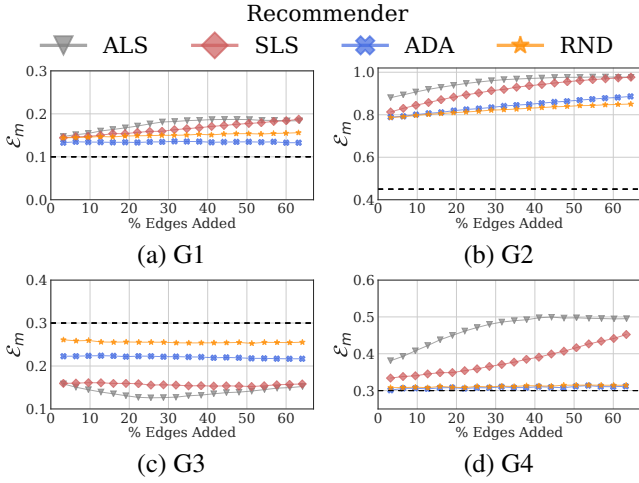


Figure 3: Exposure of the minority (\mathcal{E}_m) along time, for different recommenders and one fixed user behavior (B-PSB).

size of the minority and the user choice model is fixed to B-PSB. In cases when the minority is homophilic (G1, G2 and G4), generating recommendations through ALS and SALS leads to a positive trend of growth for the exposure of the minority. For the other two recommenders (ADA and RND) the effects described above are still present but less visible.

For the case in which the minority is heterophilic (G3) the exposure decreases weakly, slightly benefitting the majority. This is the only case when the exposure distributed to the minority is less than its relative size. It is also evident how the collaborative filtering approach (ALS) and the random walk based model (SLS) contribute more to reduce the exposure allocated to the minority with respect the other two models (ADA and RND). For all the networks characterized by an homophilic minority (G1, G2 and G4), the growth in the case of the collaborative filtering approach (ALS) is faster in the first steps and then stabilizes to a steady-state to the rest of iterations. While, for the random walk solution (SLS), the trend starts at similar values of exposure, but then grows constantly.

Observation 1 *The disparate exposure grows after each iteration in favour of the minority, when it is homophilic. On the other hand, an heterophilic behavior of the minority does not impact abruptly its exposure. When both groups are homophilic, the recommender still increases the exposure of the minority. The severity of all those effects is stronger when using ALS and SLS and weaker for ADA and RND.*

To further investigate the differences in growth of exposure, we track $\mathcal{E}_t/\mathcal{E}_1$ for $t \in \{2, \dots, T\}$, which is relative quantity of exposure measured with the respect to the first iteration. In order to analyze the transition phases previously mentioned, we focus respectively on the iterations $t \in \{2, 10, 20\}$. In Figure 5 we plot those values on the y -axis and the iterations on the x -axis. As suggested by the previous plots, with the ALS recommender, when the mi-

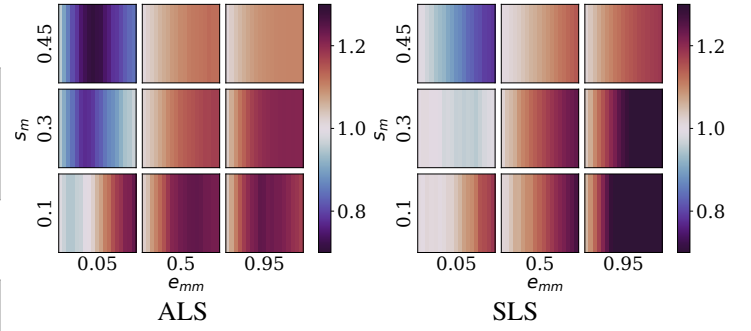


Figure 4: Heatmaps describing the evolution of $\mathcal{E}_t/\mathcal{E}_1$ in $T = 20$ iterations, computed over 9 configurations which are small variants of G1, G2 and G3: $e_{mm} \in \{0.05, 0.5, 0.95\}$ (x -axis) and $s_m \in \{0.1, 0.3, 0.45\}$ (y -axis), all having neutral majority $h_M = 0$. ALS recommender (left-hand side), SLS recommender (right-hand side).

nority is homophilic its exposure tends to increase faster in the first iterations to then stabilize. On the other hand, SLS presents a continuous increase, without slowing down the process after 20 iterations. The stronger growth comes from cases where the differences in sizes between minority and majority is relevant ($s_m = 0.3$) and the minorities are homophilic (G0, G1 and G4). This means that, even when also the majority is homophilic, the effect is still present, showing again that having both groups homophilic does not imply a benefit for the majority. As already seen in Figure 3 and as expected, ADA and RND do not produce much exposure disparity, even a slight advantage for the minority class can be observed for the cases in which the minority is homophilic.

Observation 2 *Different recommenders exhibit different influence on exposure along time. ALS increases exposure inequality in the first iterations, then stabilizing in a steady state. SLS instead keeps increasing disparate exposure constantly.*

We further extend this analysis, exploring a wider range of initial configurations, with the aim of disentangling the effects of size and homophily along time. For this purpose, as the size of the minority s_m is part of the definition of the homophily h_m (Eq. 1), we use directly the fraction of edges which, starting from the minority, remain in the minority: i.e.,

$$e_{mm} = \frac{|E_{mm}|}{|E_m|}$$

In particular, we produce 9 configurations which are small variants of G1, G2 and G3: $e_{mm} \in \{0.05, 0.5, 0.95\}$ and $s_m \in \{0.1, 0.3, 0.45\}$ (all having neutral majority $h_M = 0$).

Each box of the heatmaps in Figure 4 represents, for one configuration, the evolution of $\mathcal{E}_t/\mathcal{E}_1$ along $T = 20$ iterations.

Analyzing the two heatmaps, comparing the boxes by columns and posing the attention on a single row, we observe that both effects already observed in the previous ex-

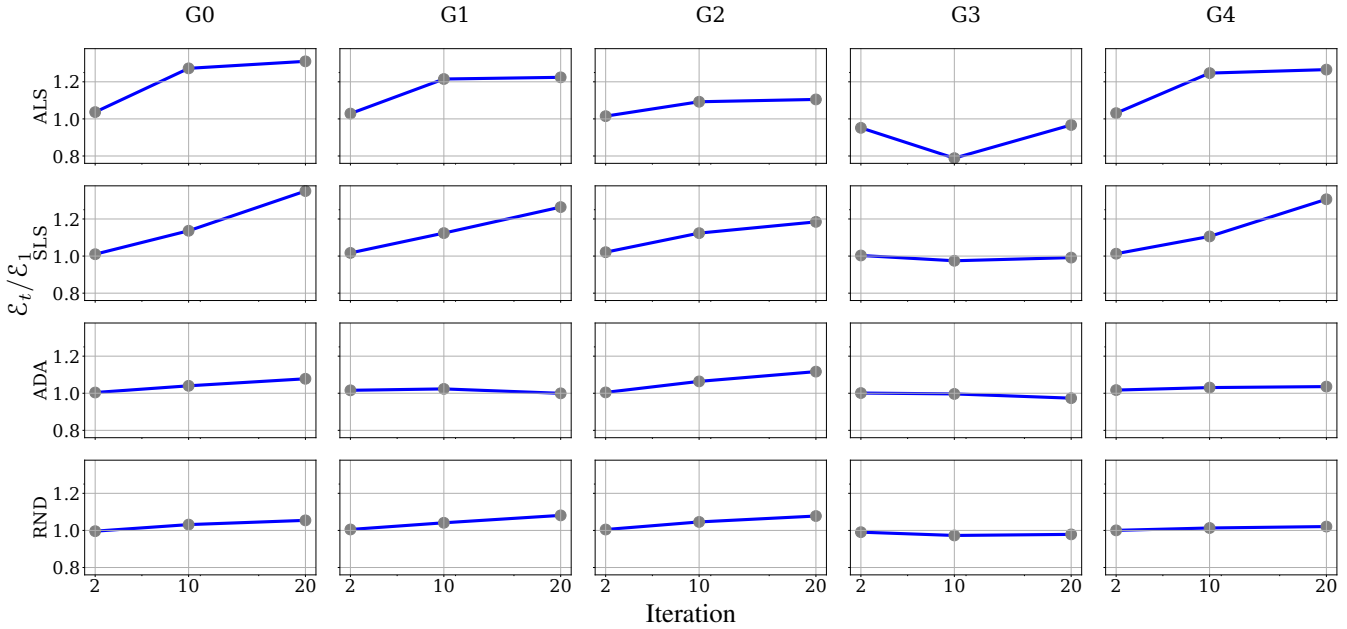


Figure 5: Evolution of exposure relative to the one observed at first iteration $\mathcal{E}_t/\mathcal{E}_1$, after 1, 10 and 20 iterations.

periments, i.e. the steady-state generated by ALS and the constant growth caused by SLS, change in terms of severity but not in timing. This means that the variation of $\mathcal{E}_t/\mathcal{E}_1$ can be less or more severe, depending on the distribution of e_{mm} , but the pace to which process evolves is the same. Analyzing the heatmaps by rows, and posing the attention on a single column, is evident how the size of the minority (represented here by s_{mm}) can influence the pace of the effects but not the range of values (color intensity) of $\mathcal{E}_t/\mathcal{E}_1$.

Observation 3 *The homophily of minority can impact the speed at which the growth of exposure disparity occurs. On the other hand, the severity of this effect is mostly determined by the size of the minority.*

Effect of User Behavior Models

In all the experiments presented so far we were adopting the B-PSB (position bias) user behaviour model. We next analyze the effect of different user behaviour models. Figure 6 reports the exposure of the minority tracked under three different policies on G0. Each plot represents a recommender and each line in the plots represents the trend of \mathcal{E}_m for one user behavior. For all the plots there is not such a significant difference in trends between models. This means that in circumstances where user behavior is either homogenous (B-LZY, B-PSB and B-RND) or heterogeneous (B-MIX), and the organic growth of the network is not considered, the effect of the recommenders are consistent.

Observation 4 *The different user behaviour models do not impact the exposure in our simulations as much as the type of recommender system and the initial configuration of the network do.*

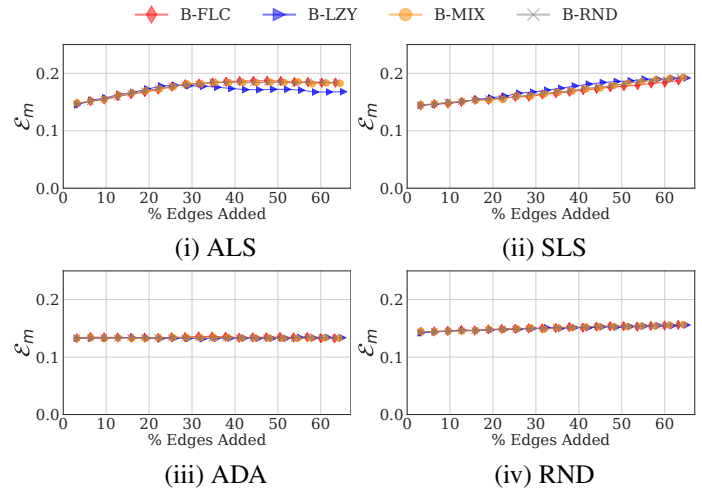


Figure 6: Exposure of minority when using different acceptance policies, running on G0.

Rich-get-richer Effect

After having analyzed the inequality in exposure at the group level, we now focus on the in-degree distribution at the individual level, focusing on the relationship with the popularity of the nodes (number of followers or in-degree).

As already observed in the literature, new links injected in the network can alter the inequality in the distribution of in-degree (Su, Sharma, and Goel 2016). For this reason, we study here the evolution of the rich-get-richer effect within the two groups, focusing the attention on how the in-degree of nodes is altered by the recommended output in the long-

term. We compute the Gini coefficient to analyze the level of concentration of in-degree within the two group of nodes, after each iteration. Although it was introduced in economics to measure the income or wealth inequality, Gini coefficient is widely used to measure inequalities in general (Halfman and Leydesdorff 2010). It is defined as follows:

$$G = \frac{1}{N} \left(N + 1 - 2 \frac{\sum_{i=1}^N (N + 1 - i) y_i}{\sum_{i=1}^N y_i} \right).$$

In our context N is the number of nodes and y_i is the in-degree of the i -th node, which has been indexed in ascending order by $y_i \leq y_{i+1}$. The index ranges from zero to one, where if all nodes receive the same amount of quantity (in-degree) then it is 0, and 1 if only one node receive the total amount. Thus, the higher the coefficient is, the higher the inequality distribution is as well. Figure 7 reports the Gini index in the long-run. Each row indicate the network, each column of plots refers to a group (minority or majority) and each line shows the Gini index after each iteration, when using all the four recommenders and one user behavior ($B - PSB$).

For all the networks we can observe a rich-get-richer effect in both the minority and the majority class: inequality of in-degree, as expressed by the Gini index, keeps growing, meaning that the high-degree nodes keep receiving more and more recommendations. Among the recommenders, ALS and SLS present a faster growth of the in-degree inequality, while ADA and RND are by far slower. When the two groups have comparable size and only one of them is homophilic (G2), the non-homophilic group gets a less severe effect. It is also evident that when the minority is heterophilic, the more impacted group is the majority, which even if not presenting biased preferences (either homophilic or heterophilic), experiences a stronger positive trend for the growth of Gini index.

The homophily level of one group impacts, not only the inequality in in-degree distribution inside the group, but indirectly it affects also the inequality in the rest of the graph.

After having observed an exacerbation of the rich-get-richer effect in the long-run, we next monitor the distribution of exposure between nodes in the two groups. We study how subset of nodes, grouped by different in-degree, can be exposed differently in the long-run. In Figure 8 we show the cumulative distribution of exposure accumulated by nodes ordered by their in-degree. In particular, each bar is divided by colors, where starting from the bottom, it represents the subset of nodes having at most the correspondent in-degree. This means that, for example, the first three colors (from the bottom) represents the first 5% of the nodes having the highest in-degree. On the y -axis we track the fraction of visibility accumulated by the nodes.

For sake of space, in Figure 8 we focus only on two recommenders, but results are consistent also with the other two models. In particular, we have similar findings for ALS and SLS, while ADA results really close to RND. Figure 8 shows that, with graphs presenting either a homophilic or a heterophilic minority (G1, G3), only a subset of nodes receives most of the exposure produced. What is also evident

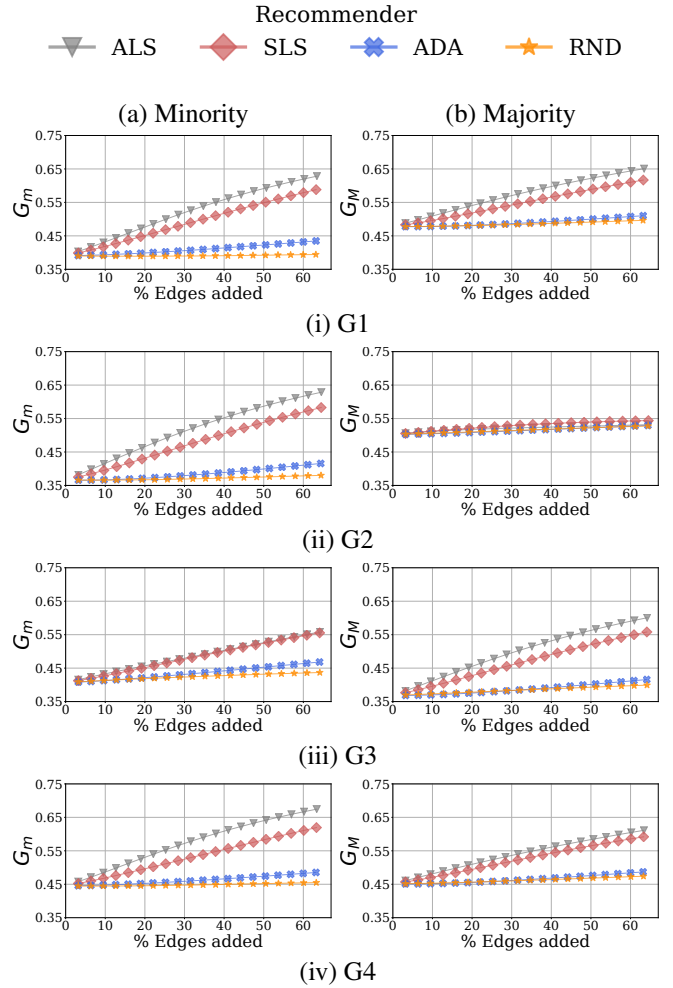


Figure 7: Gini coefficient computed on the in-degree of both minority and majority, for all the recommenders and networks, with B-PSB.

is that after each iteration, the number of nodes getting most of the exposure become smaller, confirming some recent analytical results that point out how rankings can be biased towards few individuals getting most of the exposure (Germano, Gómez, and Le Mens 2019). Specifically, this effect, in the long-term, results faster for two specific groups: the homophilic minority in G1 and the non-homophilic majority in G3. In the first case, after only 5 iterations the nodes belonging to the top-1% acquires more than 75% of exposure. While in G3, despite the majority group not being biased, after 15 iterations, only the top-3% of nodes is recommended. Moreover, in both graphs, ADA does not present the same increase in disparity in the long-term, but still, only a small fraction of users (20%) receives consistently the 75% of the exposure in both groups.

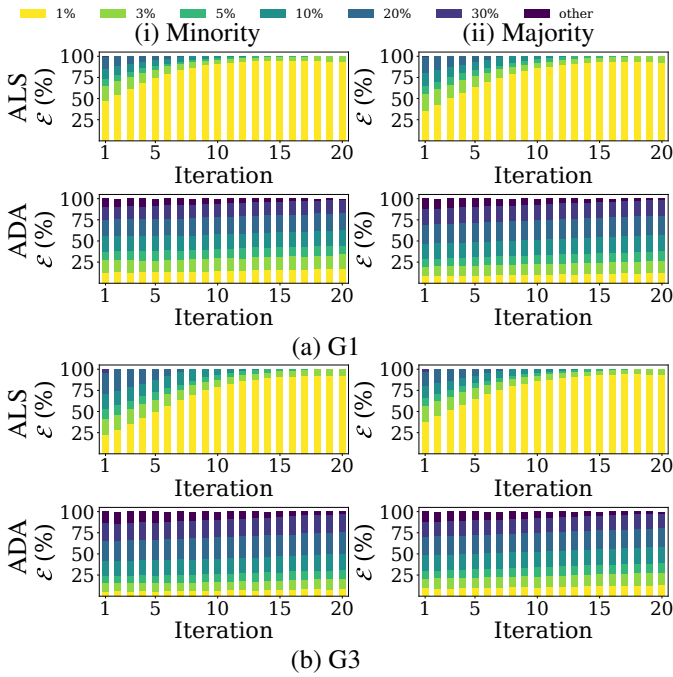


Figure 8: Distribution of Exposure among nodes, where each color delimitates the % of nodes with highest degree (best seen in color).

Observation 5 *When the minority presents non-neutral preferences (either homophilic or heterophilic), ALS and SLS can increase disparity in both exposure and in-degree: a small subset nodes benefits in terms of exposure by the injection of new links, and those are also the ones with highest degree. The cardinality of this subset of nodes becomes smaller after each iteration.*

Model Evaluation

In the experiments seen so far, we used the same configuration of k and α . To analyze how those input parameters may affect the simulation outcome, we next produce configurations presenting more sparse interactions (smaller α) and longer lists of recommendations (larger k). In Fig. 9 we present the exposure of the minority for G1, simulating the user behavior using B-PSB. In the first figure (Fig. 9a) we tune different values of α , which the smaller, the less the number of users sampled to submit new recommendations. The plot shows how the exposure tends to growth, as expected, but with a slower pace. This parameter can be tuned looking at interactions generated by the social media platform over time. Here the length of recommendations is fixed to $k = 3$. In the case of longer lists of recommendations, the length of recommendation output impacts even less on the final output (Fig. 9b). The effect observed with the smaller recommendation list ($k = 3$) presents a trend close to all the other configurations. In this case, α is fixed to the usual value of 0.2.

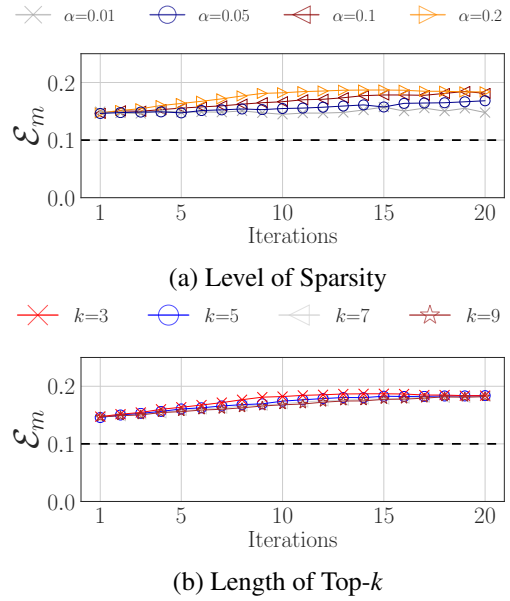


Figure 9: Testing different values of α and k on G1, with B-PSB.

Limitations and Implications

The main goal of this work is to improve our comprehension of the long-term consequences on the disparate exposure of a minority in the recommendations provided by people recommender systems in a social network. In this endeavor we need to take care of the interplay between the recommender algorithm, user behavior in accepting the recommendations, and pre-existing conditions in the network (e.g., the existence of an homophilic minority). However, disentangling the consequences of these single components remains challenging. Hence, we next discuss some limitations of our study, proposing potential extensions for future investigation.

Our analysis shows how the initial level of homophily within a subpopulation in the graph can drive exposure inequalities that grow over time: this is obtained without considering organic growth of the network (i.e., new links are created only if recommended by the algorithm) and assuming a homogenous user behavior for accepting or rejecting the link recommendations. In the future, we plan to extend our framework, including organic network growth as part of the simulation and allowing the users to have heterogeneous behaviors. In this way, we may study how homophily evolving over time may impact the recommendation output. Moreover, the user behavior models analyzed here do not consider homophily as a potential factor. In our future effort we plan to investigate how homophily can impact user choices when accepting or rejecting algorithmic recommendations.

Our work analyzes the impact of human biases, such as homophilic behavior, and link recommender algorithms on the disparate exposure of a minority at the level of the whole network. A more fine-grain analysis at the mesoscale level

of communities or subgraphs might be useful to better understand the phenomena at play.

Our work raises critical observations about the long-term consequences of algorithms in online social networking platforms and hints the need to design algorithms which keeps in consideration existing biases and aim at mitigating the long-term consequences, instead of exacerbating them. The challenge is the usual trade-off between generating recommendations which are still relevant for the users, while being able to mitigate the disparate exposure between groups even after repeated round of recommendations.

Moreover, designers of social network platforms, may take benefit of using simulation framework to infer potential harmful scenarios: the framework we present in this work can be easily adapted to test new recommendation algorithms, e.g., to shed the light on the consequences of introducing new features, before the deployment.

Finally, the results of the present study highlight the urgency to include link recommendation algorithms among the key elements when modeling network dynamics.

Acknowledgments

This work has been partially supported by: the HUMANIT programme (Human Behaviour and Machine Intelligence), Joint Research Centre, European Commission; "la Caixa" Foundation (ID 100010434), under the agreement LCF/PR/PR16/51110009; and the EU-funded "So-BigData++" project, under Grant Agreement 871042. This work has been partially supported by the Catalan Government through the funding grant ACCIÓ-Eurecat (Project PRIVany-nom).

References

- Aiello, L. M.; and Barbieri, N. 2017. Evolution of Ego-networks in Social Media with Link Recommendations. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining (WSDM)*, 111–120.
- Barbieri, N.; Bonchi, F.; and Manco, G. 2014. Who to follow and why: link prediction with explanations. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 1266–1275.
- Cinus, F.; Minici, M.; Monti, C.; and Bonchi, F. 2022. The Effect of People Recommenders on Echo Chambers and Polarization. In *Proceedings of the Sixteenth International AAAI Conference on Web and Social Media, (ICWSM)*.
- Craswell, N.; Zoeter, O.; Taylor, M.; and Ramsey, B. 2008. An Experimental Comparison of Click Position-Bias Models. In *Proceedings of the 2008 International Conference on Web Search and Data Mining (WSDM)*, 87–94.
- Domeniconi, G.; Moro, G.; Pagliarani, A.; Pasini, K.; and Pasolini, R. 2016. Job Recommendation from Semantic Similarity of LinkedIn Users' Skills. In *Proceedings of the 5th International Conference on Pattern Recognition Applications and Methods*, 270–277.
- Fabbi, F.; Bonchi, F.; Boratto, L.; and Castillo, C. 2020. The Effect of Homophily on Disparate Visibility of Minorities in People Recommender Systems. In *Proceedings of the Fourteenth International AAAI Conference on Web and Social Media, (ICWSM)*, 165–175.
- Garimella, K.; Morales, G. D. F.; Gionis, A.; and Mathioudakis, M. 2018. Quantifying controversy on social media. *ACM Transactions on Social Computing* 1(1): 1–27.
- Ge, Y.; Liu, S.; Gao, R.; Xian, Y.; Li, Y.; Zhao, X.; Pei, C.; Sun, F.; Ge, J.; Ou, W.; and et al. 2021. Towards Long-term Fairness in Recommendation. *Proceedings of the 14th ACM International Conference on Web Search and Data Mining (WSDM)* 445–453.
- Germano, F.; Gómez, V.; and Le Mens, G. 2019. The few-get-richer: a surprising consequence of popularity-based rankings? In *Proceedings of the Web Conference 2019 (WWW)*, 2764–2770.
- Guy, I.; Ronen, I.; and Wilcox, E. 2009. Do you know?: recommending people to invite into your social network. In *Proceedings of the 14th International Conference on Intelligent User Interfaces*, 77–86.
- Ha-Thuc, V.; Venkataraman, G.; Rodriguez, M.; Sinha, S.; Sundaram, S.; and Guo, L. 2016. Personalized Expertise Search at LinkedIn. In *2015 IEEE International Conference on Big Data*, 1238–1247.
- Halffman, W.; and Leydesdorff, L. 2010. Is Inequality Among Universities Increasing? Gini Coefficients and the Elusive Rise of Elite Universities. *Minerva* 48(1): 55–72.
- Heap, B.; Krzywicki, A.; Wobcke, W.; Bain, M.; and Compton, P. 2014. Combining Career Progression and Profile Matching in a Job Recommender System. In *Trends in Artificial Intelligence - 13th Pacific Rim International Conference on Artificial Intelligence*, 396–408.
- Hu, Y.; Koren, Y.; and Volinsky, C. 2008. Collaborative filtering for implicit feedback datasets. In *2008 IEEE International Conference on Data Mining (ICDM)*, 263–272.
- Jiang, R.; Chiappa, S.; Lattimore, T.; György, A.; and Kohli, P. 2019. Degenerate feedback loops in recommender systems. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, 383–390.
- Joachims, T. 2002. Optimizing search engines using click-through data. In *Proceedings of the eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 133–142.
- Karimi, F.; Génois, M.; Wagner, C.; Singer, P.; and Strohmaier, M. 2018. Homophily influences ranking of minorities in social networks. *Scientific reports* 8(1): 1–12.
- Laniado, D.; Volkovich, Y.; Kappler, K.; and Kaltenbrunner, A. 2016. Gender homophily in online dyadic and triadic relationships. *EPJ Data Sci.* 5(1): 19.
- Lee, E.; Karimi, F.; Wagner, C.; Jo, H.-H.; Strohmaier, M.; and Galesic, M. 2019. Homophily and minority-group size explain perception biases in social networks. *Nature human behaviour* 3(10): 1078–1087.
- Li, Z.; Fang, X.; and Sheng, O. R. L. 2017. A survey of link recommendation for social networks: methods, theoretical foundations, and future research directions. *ACM Transactions on Management Information Systems* 9(1): 1–26.

- Liu, R.; Ouyang, Y.; Rong, W.; Song, X.; Tang, C.; and Xiong, Z. 2016a. Rating Prediction Based Job Recommendation Service for College Students. In *Proceedings of the 16th International Conference on Computational Science and Its Applications*, 453–467.
- Liu, R.; Ouyang, Y.; Rong, W.; Song, X.; Xie, W.; and Xiong, Z. 2016b. Employer Oriented Recruitment Recommender Service for University Students. In *Proceedings of the 12th International Conference on Intelligent Computing Methodologies*, 811–823.
- Mansoury, M.; Abdollahpouri, H.; Pechenizkiy, M.; Mobasher, B.; and Burke, R. 2020. Feedback Loop and Bias Amplification in Recommender Systems. In *Proceedings of the 29th ACM International Conference on Information and Knowledge Management (CIKM)*, 2145–2148.
- Nguyen, T. T.; Hui, P.-M.; Harper, F. M.; Terveen, L.; and Konstan, J. A. 2014. Exploring the filter bubble: the effect of using recommender systems on content diversity. In *Proceedings of the Web Conference 2014 (WWW)*, 677–686.
- Ricci, F.; Rokach, L.; and Shapira, B. 2011. Introduction to recommender systems handbook. In *Recommender systems handbook*, 1–35. Springer.
- Shi, X.; Shang, M.-S.; Luo, X.; Khushnood, A.; and Li, J. 2017. Long-term effects of user preference-oriented recommendation method on the evolution of online system. *Physica A: Statistical Mechanics and its Applications* 467: 490–498.
- Stoica, A.-A.; Riederer, C.; and Chaintreau, A. 2018. Algorithmic Glass Ceiling in Social Networks: The Effects of Social Recommendations on Network Diversity. In *Proceedings of the Web Conference 2018 (WWW)*, 923–932.
- Su, J.; Sharma, A.; and Goel, S. 2016. The Effect of Recommendations on Network Structure. In *Proceedings of the Web Conference 2016 (WWW)*, 1157–1167.
- Sun, W.; Khenissi, S.; Nasraoui, O.; and Shafto, P. 2019. Debiasing the human-recommender system feedback loop in collaborative filtering. In *Companion Proceedings of the Web Conference 2019 (WWW)*, 645–651.
- Tsioutsoulis, S.; Pitoura, E.; Tsaparas, P.; Kleftakis, I.; and Mamoulis, N. 2020. Fairness-Aware PageRank. In *Proceedings of the Web Conference 2021 (WWW)*, 3815–3826.
- Vassileva, J.; McCalla, G. I.; and Greer, J. E. 2016. From Small Seeds Grow Fruitful Trees: How the PHelpS Peer Help System Stimulated a Diverse and Innovative Research Agenda over 15 Years. *International Journal of Artificial Intelligence in Education* 26(1): 431–447.
- Xie, M.; Yin, H.; Wang, H.; Xu, F.; Chen, W.; and Wang, S. 2016. Learning graph-based poi embedding for location-based recommendation. In *Proceedings of the 25th ACM International Conference on Information and Knowledge Management (CIKM)*, 15–24.
- Yao, S.; Halpern, Y.; Thain, N.; Wang, X.; Lee, K.; Prost, F.; Chi, E. H.; Chen, J.; and Beutel, A. 2021. Measuring Recommender System Effects with Simulated Users. *arXiv e-prints* arXiv:2101.04526.
- Zhang, M.; Ma, J.; Liu, Z.; Sun, J.; and Silva, T. 2016. A research analytics framework-supported recommendation approach for supervisor selection. *British Journal of Educational Technology*, 47(2): 403–420.
- Zhao, X.; Ji, X.; Liu, S.; and He, Z. 2020. Link Prediction Based on Information Preference Connection for Directed Network. In *2020 IEEE International Conference on Smart Internet of Things (SmartIoT)*, 183–189.