When Do Ranking Algorithms Reinforce Inequalities in Directed Social Networks?

Lisette Espín-Noboa [1 2] **Fariba Karimi** [1 2] **Markus Strohmaier** [3] **Claudia Wagner** [1 2]

Abstract

Despite treating each individual equally, ranking algorithms do not guarantee a fair distribution of opportunities among minorities in a social network. To that end, we study two well-known ranking algorithms, namely PageRank and Who-to-follow (WTF), and show the inequalities on their *rank distribution*, and *the visibility of minorities among their top-k ranks* on six real-world networks. Moreover, we propose a generative network model to demonstrate that *homophily* and *preferential attachment* are the two main mechanisms that can explain these inequalities. Our simulations show that regardless of the level of homophily among minorities, they are always over-represented when majorities are heterophilic. Our findings shed light on the role played by majorities and minorities in a social network to reduce inequalities without algorithmic intervention.

1. Introduction

One of the most challenging problems facing our society today is the growing prevalence of inequality. In social groups, ranging from small online and offline communities to entire nations, many undesirable social outcomes can result from lopsided distributions of opportunities, income, and wealth. Nowadays search engines and recommender systems are increasingly used for various applications such as whom to follow in a social network. Typically, these applications use ranking algorithms to order items (e.g., people you may know) based on "importance" or "relevance", and may therefore impact inequality by discriminating certain items or groups in the top of the rank.

By way of example, consider the network shown in Figure 1(a), where orange nodes represent minorities, and blue nodes majorities. In that linkage state (i.e., only majority nodes link to minority nodes), we see that minorities are over-represented (y-axis) in the top rank of PageRank (x-axis). However, this ranking drastically varies if the linkage is reversed or modified, Figure 1(b, c, d).

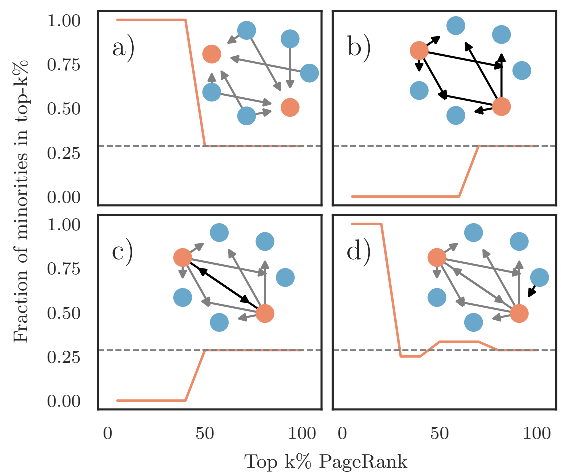

Figure 1. **Example of ranking bias:** Using PageRank we compute the fraction of minorities (y-axis) in the top-k% rank (x-axis). Networks contain 7 nodes; 2 orange (minority) and 5 blue (majority). Horizontal dotted lines represent the actual fraction of minorities. Network linkage differ as follows. a) Only majority nodes point to minority nodes; e.g., *fans following celebrities*. b) Only minorities link to majorities; e.g., *women citing men in scientific articles*. c) Similar to -b-, including links within minorities. d) Similar to -c-, including one link from majority to minority. We see how the ranking of minorities changes depending on how minorities and majorities are connected. The extreme case in -b- can be improved or reversed by connecting nodes to minorities.

Inequality, however, can manifest on a vertical and on a horizontal dimension. *Vertical inequality* refers to the skewness of a distribution and is often measured with the Gini coefficient or similar measures of variation, while *horizontal inequality* is concerned with differences between groups of people that share ascriptive characteristics such as gender or ethnicity. In this work we study inequality on both dimensions in parallel and aim to unveil the mechanisms that explain vertical and horizontal inequality with respect to PageRank (Page et al., 1999) and Who-To-Follow (WTF) (Gupta et al., 2013), two very popular ranking algorithms.

[1]GESIS [2]University of Koblenz-Landau [3]RWTH Aachen University. Correspondence to: <Lisette.Espin@gesis.org>.

Proceedings of the 37^{th} International Conference on Machine Learning, Vienna, Austria, PMLR 108, 2020.

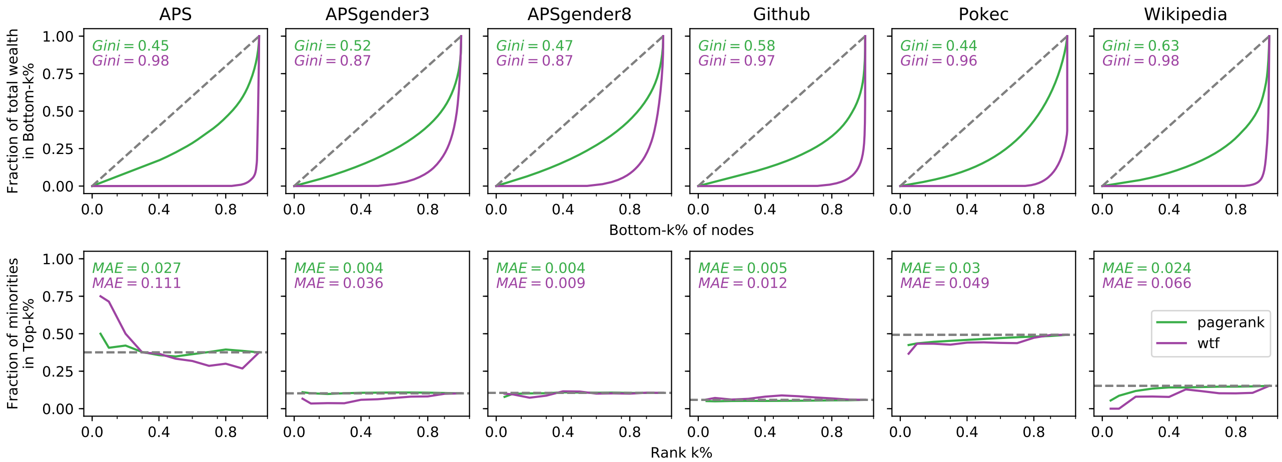

Figure 2. **Ranking inequalities:** Given six different networks, we rank all nodes using two well-known ranking algorithms: PageRank and Who-to-Follow (WTF). Next, we compute the inequalities in the top of their ranks as follows: Vertical inequality (top row): Inequality on the rank distribution using Lorenz curve and Gini coefficient. Horizontal inequality (bottom row): Visibility of minorities among the top-k rank. We see that WTF is more prone to reinforce inequalities in both dimensions than PageRank.

Since these algorithms harness structural information, they can alter the visibility of protected groups (e.g., minorities based on gender), and consequently reinforce societal issues such as the invisibility syndrome (Franklin & Boyd-Franklin, 2000) and the glass ceiling effect (Cotter et al., 2001). In fact, previous studies have shown that certain social mechanisms such as homophily can alter the structure of social *undirected* networks and consequently affect their degree rank (Karimi et al., 2018).

2. Experiments

Empirical approach: We introduce six real-world directed networks that contain nodes with a binary attribute (e.g., gender). We denote nodes, from the attribute that is less representative, as a minority group. We rank nodes using PageRank and WTF, and compute the inequalities on their ranks. As shown in Figure 2, we see that: (i) Both ranking algorithms reinforce vertical inequality. However, WTF is prone to higher levels. (ii) Horizontal inequality is mostly replicated by PageRank and reinforced by WTF.

Simulation: We have seen that PageRank and WTF ranking results might benefit some nodes over others in the top ranks. This advantage is not a random event and can be explained by who connects to whom and in which direction; recall that both techniques rely purely on the structure of the network. Moreover, social networks are complex systems where different groups of nodes co-exist with different preferences and different mechanisms of tie formation. Thus, we propose a network model that replicates the main mechanisms of social networks, i.e., directed links, homophily, preferential attachment, and group size, and show how these mechanisms can explain ranking inequalities.

Network model: Based on the preferential attachment model with homophily (Karimi et al., 2018), we propose a growth model for *directed* networks with tunable homophily and group size. We assign a homophily value to each dyad based on a pre-defined homophily parameter h that ranges from 0.0 to 1.0. If the homophily value is high, that means that two nodes of the same class are attracted to each other more often than two nodes of different class. In order to generate directed links, we assign an activity score (Perra et al., 2012) to each node. This score determines the probability that an existing node becomes active to create additional links to other nodes. This model ensures that the in-degree and out-degree distribution of nodes follow a seemingly power-law distribution that have been observed in many large social networks (Voitalov et al., 2018).

3. Conclusions

In this paper, we propose a directed network model that simulates preferential attachment, homophily, and group size in the tie formation process. Furthermore, we systematically study how relative size differences between groups of nodes, with different levels of homophily, impact the ranking of nodes in PageRank and Who-To-Follow (WTF) using synthetic and real-world social networks. Our results suggest that regardless of the level of homophily among minorities, they are always over-represented when majorities are heterophilic. This effect intensifies when minorities are either heterophilic and very small, or homophilic. Although this holds for both ranking algorithms, WTF is more sensitive to these inequalities. This means that not only people's own behavior can lead to inequalities, but also recommender systems, such as WTF, can intensify and reinforce such inequalities by never recommending certain groups of nodes.

References

Cotter, D. A., Hermsen, J. M., Ovadia, S., and Vanneman, R. The Glass Ceiling Effect*. *Social Forces*, 80(2):655–681, 12 2001. ISSN 0037-7732. doi: 10.1353/sof.2001.0091. URL https://doi.org/10.1353/sof.2001.0091.

Franklin, A. J. and Boyd-Franklin, N. Invisibility syndrome: A clinical model of the effects of racism on african-american males. *American Journal of Orthopsychiatry*, 70(1):33–41, 2000.

Gupta, P., Goel, A., Lin, J., Sharma, A., Wang, D., and Zadeh, R. Wtf: The who to follow service at twitter. In *Proceedings of the 22nd international conference on World Wide Web*, pp. 505–514, 2013.

Karimi, F., Génois, M., Wagner, C., Singer, P., and Strohmaier, M. Homophily influences ranking of minorities in social networks. *Scientific reports*, 8, 2018.

Page, L., Brin, S., Motwani, R., and Winograd, T. The pagerank citation ranking: Bringing order to the web. Technical report, Stanford InfoLab, 1999.

Perra, N., Gonçalves, B., Pastor-Satorras, R., and Vespignani, A. Activity driven modeling of time varying networks. *Scientific reports*, 2:469, 2012.

Voitalov, I., van der Hoorn, P., van der Hofstad, R., and Krioukov, D. Scale-free networks well done. *arXiv preprint arXiv:1811.02071*, 2018.