Do Recommender Systems Make Social Media More Susceptible to Misinformation Spreaders?

Antonela Tommasel
antonela.tommasel@isistan.unicen.edu.ar
ISISTAN (CONICET/UNCPBA)
Tandil, Argentina

Filippo Menczer
fil@iu.edu
Observatory on Social Media, Indiana University
Bloomington, USA

ABSTRACT

Recommender systems are central to online information consumption and user-decision processes, as they help users find relevant information and establish new social relationships. However, recommenders could also (unintendedly) help propagate misinformation and increase the social influence of the spreading it. In this context, we study the impact of friend recommender systems on the social influence of misinformation spreaders on Twitter. To this end, we applied several user recommenders to a COVID-19 misinformation data collection. Then, we explore what-if scenarios to simulate changes in user misinformation spreading behaviour as an effect of the interactions in the recommended network. Our study shows that recommenders can indeed affect how misinformation spreaders interact with other users and influence them.

CCS CONCEPTS

• **Information systems → Social networking sites**; **Recommender systems**; • **Computing methodologies → Simulation evaluation**.

KEYWORDS

link prediction, social media, misinformation, diffusion models

ACM Reference Format:
Antonela Tommasel and Filippo Menczer. 2022. Do Recommender Systems Make Social Media More Susceptible to Misinformation Spreaders?. In *Sixteenth ACM Conference on Recommender Systems (RecSys '22), September 18–23, 2022, Seattle, WA, USA*. ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3523227.3551473

1 INTRODUCTION

In the last few years, social media (as well as other online platforms) has seen an increment in the spreading of fake news, rumours, hoaxes, and other forms of misinformation, as well as the proliferation of hate speech, incitement to violence, harassment and other forms of abusive behaviours [18]. This situation has become a critical problem with negative real-world consequences, ranging from weakening democratic systems to public health issues [25]. For example, the COVID-19 pandemic caused an increased need for trustworthy information in response to the highly emotional and uncertain situation. This led to misinformation linked to health recommendations [27], which could reduce the credibility of scientific evidence and affect compliance to evidence-based policy interventions, with long-term and potentially deadly consequences [12].

As mediators of online information consumption, recommenders are affected by the proliferation of low-quality and undesired content, serving as an unintended means for their spread and massive amplification while reducing the quality of predictions. Recently, recommenders have undergone criticism for inducing the creation of filter bubbles, echo chambers, and facilitating opinion manipulation [7]. Similarly, users' vulnerability to dis/misinformation can be fostered by data, algorithms, and interaction biases, which limit users' openness to contrasting points of view. While studies regarding how recommenders select the information users are exposed to are common, the studies of how recommenders influence misinformation spreading are limited.

In this work, we contribute to the study of *whether recommender systems make social media more susceptible to misinformation spreaders by amplifying their influence*. To this end, we simulate changes in user misinformation spreading behaviour as an effect of the interactions with their network. First, we apply different *recommenders* to define what-if scenarios of how the structure of social interactions would evolve based on the obtained recommendations. Then, we *simulate a discrete opinion dynamic model* over the derived networks to assess how recommendations would affect misinformation spreaders' influence. To support our proposal, we conducted a simulation over a COVID-19 misinformation Twitter data collection. Our study shows that recommenders had a differentiated impact on misinformation spreader influence, leading to networks of different propagation characteristics. Particularly, the analyses showed that it is not enough that recommendations include many spreaders, but they also need to be well connected to affect spreading.

2 RELATED WORKS

The initial step for mitigating misinformation is understanding how it propagates [1]. However, despite the increasing attention to misinformation detection, the study of how social media contributes to opinion formation has received comparatively less attention [4]. Closely related to this study are works analyzing the effect of recommenders on several phenomena [4, 6, 8, 22]. Fernández et al. [8] studied the impact of news recommenders on misinformation spread. Spread was statically measured based on the recommended items. Results showed that recommenders suffering from popularity bias were more likely to recommend misinformation.

Several studies focused on the impact of friend recommenders [4, 6, 22]. Fabbri et al. [6] studied the impact on user exposure, which might influence attention dynamics and reinforce the inequalities affecting certain groups (e.g., gender). Results showed that minority

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
RecSys '22, September 18–23, 2022, Seattle, WA, USA
© 2022 Association for Computing Machinery.
ACM ISBN 978-1-4503-9278-5/22/09…$15.00
https://doi.org/10.1145/3523227.3551473

and homophilic groups could get a disproportionate advantage in exposure, while non-homophilic groups get underexposed. Adapted to our study, this could imply that if misinformation spreaders originally belong to a tight group, they might get over-recommended. Finally, exposure did not seem to be significantly affected by how users decided which recommendations to accept. Instead, it was affected by the initial network configuration and the recommender. This might have negative consequences as user actions are not considered relevant for network evolution.

Santos et al. [22] and Cinus et al. [4] assessed the effect of friend recommenders in echo chamber evolution. Cinus et al. [4] combined traditional topological and opinion-guided recommenders with two opinion models. Results showed that recommenders could contribute to echo chambers in the presence of more homophilic than non-homophilic links. In addition, the impact of recommenders was negligible if the initial network was already segregated in polarized communities. Finally, Santos et al. [22] evaluated the effect of topological recommenders over synthetic networks. During simulation, opinions were updated based on the average influence of neighbours. Results showed an increment of opinion polarization, measured as the opinions' standard deviation at the end of the simulation. Unlike Fabbri et al. [6], network density was controlled by removing one edge for each newly added one.

3 TASK AND METHODS

To assess the effect of recommenders over misinformation spreaders' influence, we defined a 4-step pipeline (Figure 1).

3.1 Data collection

Evaluation was based on `FibVid` [17][1], a COVID-related misinformation dataset comprising news claims appearing in Politifact and Snopes. From each claim, the authors extracted keywords that were searched on Twitter to retrieve the associated content. Tweets were labelled according to the matching Politifact or Snopes label of the associated news claim. The collection comprised 772 COVID-related claims associated with 112, 433 tweets shared during 2020, and 24, 340 users. We retrieved all tweets using the *Faking it!* tool[2].

Based on the Politifact and Snopes labels of the associated news claims, 26% and 74% of tweets were labeled as authentic and fake content, respectively. Then, we used these labels to determine whether users were misinformation spreaders. For each user, we computed the proportion of shared tweets deemed fake, and if it was higher than a certain threshold, we considered the user a misinformation spreader. We adopted a conservative definition of spreaders, setting the threshold to 0.5.

3.2 Network definition

Given the dynamic content-based nature of social interactions, instead of focusing on the topological follower/followee graph, we focused on the conversational interaction network. We built a user graph in which nodes represent users, and edges represent reply, mention, quote, or retweet actions. Edges are directed to represent the flow of conversation. In the case of replies, there is an edge from user A to user B if A replied to B. For mentions, there is an edge from A to B if A mentioned B. For quotes, there is an edge from A to B if A quoted B. For retweets, there is an edge from A to B if B retweeted A. To ensure that the user graph is not disconnected, we kept users (and their tweets) in the largest connected component that had shared more than one tweet. Then, from the original user set, we kept 20, 154 users. On average, each user shared 4.9 (±17.25) tweets and interacted 4.4 (±17.4) user relations[3].

The network was temporally split into training and test. User interactions were sorted according to their date; the first 70% were used for the training set, while the remaining 30% were used as test set. Based on the defined threshold, in the training network, 41% of users were deemed as spreaders, 49% non-spreaders, and the remaining 10% were neutral (i.e., their score was equal to the threshold). When considering the entire network, the spreader/non-spreader distribution changed. As users shared more tweets, the proportion of users that could be considered spreaders increased (i.e., the tendency of users to share misinformation increased), reaching 80% of spreaders, 10% of non-spreaders, and 10% of neutral users. Only 2% reverted their misinformation spreading behaviour.

The original network topology and homophily (i.e., the tendency of users to interact with other similar users) can affect the final characteristics of the network and condition the effect of recommenders [6]. Then, following Garimella et al. [9] and Cota et al. [5], we quantified the polarization of the selected users relying on the relation between users' score and the score of the users with whom they interacted to determine the homophily levels of each group. Positive Pearson correlations were found for non-spreaders. Spreaders show a negligible negative correlation. Correlations for non-spreaders were higher with replied than mentioned users. On average, 27% of spreaders' interactions were with other spreaders, with 17% of interactions with users with higher or equal scores. Conversely, non-spreaders interacted with users on a broader range of scores, accounting for only 16% of interactions with spreaders. This shows that both spreaders and non-spreaders tended to interact more with non-spreaders, with non-spreaders showing a higher homophilic behaviour than spreaders.

3.3 Making recommendations

Several recommenders were considered for recommending links in the defined conversational interaction network. First, two trivial and non-personalized baselines: i) **popularity** (i.e., users were recommended based on their degree), and ii) **random**, as a lower bound reference. Second, three traditional user recommenders: i) **Topological** is based on Resource Allocation [20]. ii) **Content** is based on the cosine similarity of the centroid of the Word2Vec vectors of users' shared tweets. iii) **Friend-of-Friends (FoF)** is based on the triadic closure principle by which friends of friends also tend to become friends [13]. Then, users at a 2-hop distance were recommended according to their popularity. Finally, **Implicit MF** [15], a top-performing matrix factorization technique tailored for implicit feedback settings. All recommendations were performed over the same temporal data partitions. A cut-off threshold was defined to recommend the top-k users, with $k = 10$. Recommendations were deemed correct if they appeared in the test set.

Evaluation metrics. While the performance of recommenders is not the main focus of this study, unlike other works in the literature [4, 6, 8, 22], we consider it an important aspect to analyze.

[1]The collection is originally available at https://doi.org/10.5281/zenodo.4441377
[2]The tool is publicly available at: https://github.com/knife982000/FakingIt

[3]The final retrieved set of tweet IDs and their metadata are available at: https://github.com/tommantonela/recsys2022-spreader-recommendation

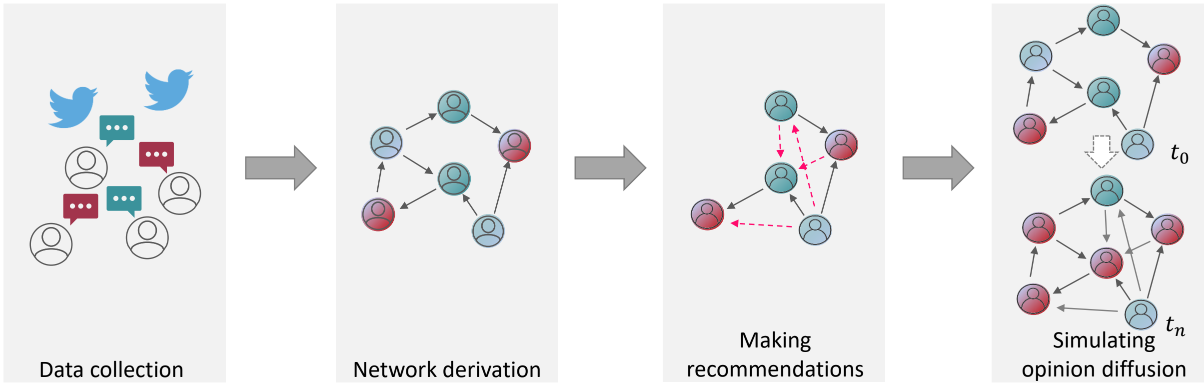

Figure 1: Phases of the proposed pipeline

This does not aim to endorse recommendations but check whether they would be close to the actual user interactions. If recommendations do not follow (even loosely) them, then the resulting network structure might not be representative of user behaviour.

In addition to relevance-based metrics (*precision*, *nDCG*), recommendations were assessed based on intra-list dissimilarities [24] in terms of *diversity* (i.e., differences between recommended users) and *novelty* (i.e., differences between the recommended and the already known users). Dissimilarities were measured by the euclidean distance of structural and content-based vectors. Content-based distance was computed based on the Word2Vec representation of users' tweets. For computing structural distances, users were represented by a vector defining their interaction rate with users from the different communities [11] discovered using the Louvain algorithm [2]. As relevance and diversity metrics, which target user satisfaction, are independent of whether the recommendations include spreaders, we assessed *user exposure to misinformation*, measuring the ratio of recommended users who were misinformation spreaders.

3.4 Simulating opinion diffusion

Information or opinion propagation models aim to simulate (and estimate) how information/opinions are propagated over networks through user interactions. The process takes as input the training user graph (i.e., the one used as the base for recommendations) and the recommendations. In each iteration, a random user is selected, and the next recommendation in their ranking is added to the graph. In accordance to the attention budget phenomenon [16], for each newly added edge, we remove users' eldest neighbour. This allows discarding any effect associated with graph densification [4, 22]. Then, the update rule of the selected model is applied, and opinions are modified accordingly. We assume that users accept all recommendations. Although this assumption might be unrealistic [23], it allows observing how recommendations could shape the network. Once all recommendations for every user are considered, the simulation continues to let any long-term effect of the final graph emerge. Iterations are stopped once a stable state is reached.

Simulations were based on the **Q-voter** model [3][4], which represents a generalization of the original voter model [14], in which the opinions of any given user change at random times under the influence of the opinion of one neighbour. Similarly, in the Q-voter model, in each iteration, a random user interacts with q of their neighbours. Then, if all q neighbours share the same opinion, the user changes their opinion accordingly. Otherwise, the user keeps their opinion. Opinions are defined as whether users are misinformation spreaders or not. The number of neighbours q was selected according to a preliminary simulation over the training user graph.

Evaluation metrics. To quantify the effect of recommenders on spreader propagation, we compared the characteristics of the induced network obtained at the end of the simulation with the real edge and spreader distribution network based on: i) *Percentage of misinformation spreaders*, i.e., the percentage of users at the end of the simulation that are spreaders. ii) *Clustering coefficient* measures how connected user's neighbours are to each other. Centrality metrics do not only define how nodes are connected but also how actively they communicate and how they could influence others [26]. iii) *User interaction polarization*, we quantify the polarization of users in terms of the Pearson correlation between users' spreadness state and their neighbours state. iv) *Random Walk Controversy Score (RWC)* [10] measures polarization by computing the probability of a random user being exposed to users from the other group. High scores indicate a low probability of groups interacting, implying that the two groups are well separated and thus polarized. Scores close to 0 indicate a similar probability of users interacting with others from either group. As the selected model has random components (the selection of the user to update and the selected Q neighbours), we run each model 50 times.

4 ANALYSIS

This section presents the results of the conducted analyses aiming at answering: **RQ1:** How do recommenders contribute to misinformation spreaders recommendations? **RQ2:** How do recommenders contribute to amplifying the influence of misinformation spreaders?

Recommendation evaluation

Table 1 presents the recommendation results. For each metric, we report the average across all users and the standard deviation. Diversity and novelty for the base graph were computed considering the test set as the recommended users. In general, recommenders achieved low precision and moderate nDCG. The highest relevance results were obtained with the topology-based recommenders. Popularity achieved the highest nDCG. Except for the random and content-based recommenders, statistically significant differences were observed (with an alpha of 0.01) between the results obtained for the spreader and non-spreader groups.

[4] The model was implemented as in https://github.com/GiulioRossetti/ndlib

	Precision	nDCG	Content-based dissimilarities Diversity	Content-based dissimilarities Novelty	Structural dissimilarities Diversity	Structural dissimilarities Novelty	Misinformation Exposure
base graph	-	-	0.077 ± 0.141	0.294 ± 0.085	0.118 ± 0.217	0.461 ± 0.147	0.18 ± 0.353
Random	0.1 ± 0.01	0.401 ± 0.141	0.373 ± 0.033	0.343 ± 0.046	0.599 ± 0.046	0.547 ± 0.078	0.382 ± 0.154
Popularity	0.111 ± 0.038	0.626 ± 0.254	0.248 ± 0.114	0.301 ± 0.06	0.373 ± 0.18	0.431 ± 0.116	0.164 ± 0.22
Friend-of-Friends	0.203 ± 0.204	0.581 ± 0.186	0.218 ± 0.005	0.258 ± 0.057	0.366 ± 0.011	0.427 ± 0.093	0.196 ± 0.022
Topology - Resource Allocation	0.231 ± 0.193	0.553 ± 0.235	0.248 ± 0.114	0.301 ± 0.06	0.373 ± 0.18	0.431 ± 0.116	0.164 ± 0.22
Content-based	0.1 ± 0.021	0.5 ± 0.13	0.381 ± 0.035	0.346 ± 0.045	0.601 ± 0.045	0.548 ± 0.078	0.371 ± 0.155
Implicit MF	0.109 ± 0.029	0.558 ± 0.266	0.324 ± 0.078	0.347 ± 0.085	0.483 ± 0.073	0.467 ± 0.097	0.161 ± 0.123

Table 1: Relevance, diversity and novelty recommendation results comparison for $k = 10$.

Generally, recommendations' structural diversity/novelty were higher than content-based diversity/novelty. Structural diversity/novelty indicate that recommended users belonged to communities that users had not yet explored. These recommendations could affect network rewiring during simulations as new edges might connect far away users. Techniques achieving high relevance also achieved low diversity/novelty scores, as the popularity and topological recommenders. Conversely, recommenders achieving low relevance also achieved the highest diversity/novelty as the random and content-based recommenders. The fact that recommendations based on the shared content increased content-based diversity/novelty could imply that users interact with others sharing semantically different content, which might be related to the observed cross-group interactions.

As expected, the random recommender achieved the highest coverage (i.e., the proportion of users that were recommended from the whole pool of users). In contrast, the lowest coverage was observed for recommenders with low diversity/novelty (e.g., popularity). Content-based recommendations achieved a coverage of 33%, closely followed by Implicit-MF and FoF. Thus, recommenders with low coverage foster the densification of certain network regions, while the others foster a more balanced distribution of edges.

Regarding user exposure to misinformation, considering similar proportions of spreaders and non-spreaders, recommendations included less than 40% of misinformation spreaders. The last interactions of the base graph included, on average, only 17% of spreaders. Scores showed that the most popular recommended users were mostly non-spreaders. In general, the ratio of recommended spreaders was higher for spreaders than for non-spreaders, which implies that spreaders might be inserted in echo chambers that are strengthened as an effect of recommendations.

In summary, the best relevance performing recommenders were the ones recommending the fewest spreaders, while recommenders with increasing diversity/novelty tended to recommend the highest ratio of spreaders.

Opinion propagation evaluation

Table 2 presents the simulation results. For comparison purposes, we include the scores obtained when simulating the propagation over the base graph plus the edges in the test partition. None of the recommenders greatly increased the proportion of spreaders in the network Instead, content-based and random decreased it. In the case of random, spreaders almost disappear from the network. These recommenders achieved both the largest number of recommended spreaders and the largest number of users becoming spreaders for at least one iteration, converting 52% of users. However, as they also favoured the interactions with novel and far away users, interactions with spreaders might not have been concentrated in the same neighbourhood, reducing their joint influence.

Most recommenders contributed to increase the proportion of users interacting with spreaders, with popularity accounting for the 95% of users with a spreader in their neigbourhood. On the other hand, random greatly decreased the proportion for non-spreaders, and increased it for spreaders, inducing segregation in the network. This is confirmed by the interaction correlation [5, 9], which showed that popularity and topology increased the correlation for spreaders (i.e, spreaders are surrounded by other spreaders) and decreased it for non-spreaders (i.e., there were more spreaders in their neighbourhood). For non-spreaders, a similar effect is observed for content and random recommenders.

Regarding clustering coefficient, significant increments were observed for topology and popularity recommenders, while significant decrements were observed for content and random. This confirms that the interactions added by content and random do not contribute to increasing user centrality. Also, for those recommenders, the clustering coefficient for non-spreaders was higher than for spreaders, indicating that non-spreaders had a more central role in the network, while spreaders might not be well connected, although scattered across the whole network.

Finally, according to RWC, FoF increased group segregation, while the other recommenders caused users to mix. While reducing polarity (as defined by the metric) might seem good, it also implies a higher mixture of user interactions, which might increase spreaders reachability. When analyzing the individual groups, interactions between non-spreaders seemed to consolidate. All recommenders increased the likelihood of spreaders interacting, and perhaps influencing, non-spreaders. Conversely, popularity increased the likelihood of non-spreaders interacting with spreaders.

Observations indicate that recommending a large number of spreaders does not directly lead to a high conversion rate. Instead, recommended spreaders also need to be well connected with their neighbours to affect spreading.

The number of converted users seemed insufficient to characterize spreading dynamics. Generally, recommenders with low coverage and low diversity/novelty strengthened the influence and centrality of the small set of users involved in more paths around the network. However, at the same time, their influence is mainly circumscribed to a “small" network region and can hardly spread.

Recommenders diversifying interactions and fostering connections with users in other network regions seemed to have a stronger effect on spreaders presence and interaction dynamics.

	% spreaders	Clustering Coefficient	User interaction polarization	Random Walk Controversy Score
base graph	0.345 ± 0.002	0.033 ± 0.119	0.26 ± 0.001	0.193 ± 0.002
Random	0.042 ± 0	0.003 ± 0.028	0.19 ± 0.001	0.147 ± 0.041
Popularity	0.366 ± 0	0.71 ± 0.31	0.16 ± 0	0.34 ± 0.023
Friend-of-Friends	0.367 ± 0	0.137 ± 0.233	0.245 ± 0	0.363 ± 0.002
Topology - Resource Allocation	0.368 ± 0	0.069 ± 0.167	0.203 ± 0	0.445 ± 0.032
Content-based	0.2 ± 0.001	0.01 ± 0.06	0.173 ± 0	0.143 ± 0.001
Implicit MF	0.369 ± 0	0.033 ± 0.12	0.214 ± 0	0.226 ± 0.021

Table 2: Results for the simulation of opinion propagation

In this scenario, network topology and rewiring seem to be the greatest drivers for opinion spreading. Then, for those cases that foster interactions with far away users, it would be interesting to see how spreading evolves under different rewiring conditions and levels of network densification.

5 CONCLUSIONS

We presented a preliminary exploration to better understand how user recommenders affect network dynamics in terms of misinformation spreader distribution and influence. Our study brings to attention the potential implications of recommenders in network evolution and dynamics, serving as a basis to study other related polarization phenomena (e.g., echo chambers and filter bubbles). Also, simulations could help evaluate potential scenarios to test new or modified recommenders and assess their effects before deployment.

The study presents some limitations. First, recommenders are agnostic of the particularities of recommended users. For example, if groups were inverted, recommenders would have fostered the propagation of non-spreaders instead of spreaders. Then, recommenders might need to be enriched with information regarding user trustworthiness. Second, the used data collection was topically focused and sparse. Fabbri et al. [6] highlighted that recommenders' effects varied based on the initial network characteristics. Then, the study should be replicated considering collections of varying number, density and misinformation spreader distribution. Third, this preliminary evaluation included a reduced set of recommenders and a unique opinion propagation model. More recommenders and propagation models should be included in the analysis to prove observations generalizability.

Several aspects could be tackled in future works. First, the evaluation should include additional recommenders, data collections, and opinion propagation models. Particularly, the definition of spreader could be relaxed to evaluate continuous opinion models. Second, in a diversity-enabling scenario, as recommendations can be perceived as interventions to the organic network dynamics, recommenders should limit the number of recommendations needed to increase diversity [21]. Analogously, recommenders could explore which recommendations cause the largest effect on misinformation spreading to avoid them or reduce their relevance, among other possibilities. Third, we adopted a simple rewiring strategy in which their eldest edge was removed for each added edge to a user. Additional strategies could be considered, and even a probability of edge removal could be introduced to simulate densification scenarios [19].

ETHICS STATEMENT

Research is based on publicly available Twitter data initially collected and tagged by third parties. No user identity was used or disclosed in the analysis. As per Twitter TOS, the shared graphs only include the IDs involved in the interactions. The analysis performed aims at showing the inadvertent effects of recommenders on network shaping, and the amplification of misinformation spreaders' influence. Nonetheless, it can suffer from bias stemming, for example, from the data collection and tagging process. In this sense, biases should be considered before applying any derived result from this study in real-world settings.

REFERENCES

[1] Banafsheh Behzad, Bhavana Bheem, Daniela Elizondo, Deyana Marsh, and Susan Martonosi. 2021. Prevalence and Propagation of Fake News. *arXiv preprint arXiv:2106.09586* (2021).

[2] Vincent D Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. 2008. Fast unfolding of communities in large networks. *Journal of statistical mechanics: theory and experiment* 2008, 10 (2008), P10008.

[3] Claudio Castellano, Miguel A Muñoz, and Romualdo Pastor-Satorras. 2009. Nonlinear q-voter model. *Physical Review E* 80, 4 (2009), 041129.

[4] Federico Cinus, Marco Minici, Corrado Monti, and Francesco Bonchi. 2022. The Effect of People Recommenders on Echo Chambers and Polarization. *Proceedings of the International AAAI Conference on Web and Social Media* 16, 1 (May 2022), 90–101. https://ojs.aaai.org/index.php/ICWSM/article/view/19275

[5] Wesley Cota, Silvio C Ferreira, Romualdo Pastor-Satorras, and Michele Starnini. 2019. Quantifying echo chamber effects in information spreading over political communication networks. *EPJ Data Science* 8, 1 (2019), 1–13.

[6] Francesco Fabbri, Maria Luisa Croci, Francesco Bonchi, and Carlos Castillo. 2022. Exposure Inequality in People Recommender Systems: The Long-Term Effects. *Proceedings of the International AAAI Conference on Web and Social Media* 16, 1 (May 2022), 194–204. https://ojs.aaai.org/index.php/ICWSM/article/view/19284

[7] Miriam Fernandez and Alejandro Bellogín. 2020. Recommender Systems and Misinformation: The Problem or the Solution?. In *Proceedings of the Workshop on Online Misinformation- and Harm-Aware Recommender Systems (OHARS 2020)*. CEUR, Virtual Event, Brazil, 40–50.

[8] Miriam Fernández, Alejandro Bellogín, and Iván Cantador. 2021. Analysing the effect of recommendation algorithms on the amplification of misinformation. *arXiv preprint arXiv:2103.14748* (2021).

[9] Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. 2018. Political Discourse on Social Media: Echo Chambers, Gatekeepers, and the Price of Bipartisanship. In *Proceedings of the 2018 World Wide Web Conference* (Lyon, France) *(WWW '18)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 913–922. https://doi.org/10.1145/3178876.3186139

[10] Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. 2018. Quantifying controversy on social media. *ACM Transactions on Social Computing* 1, 1 (2018), 1–27.

[11] Quentin Grossetti, Cédric du Mouza, and Nicolas Travers. 2020. Community-based recommendations on Twitter: avoiding the filter bubble. In *International Conference on Web Information Systems Engineering*. Springer, Springer International Publishing, Cham, 212–227.

[12] Kris Hartley and Minh Khuong Vu. 2020. Fighting fake news in the COVID-19 era: policy insights from an equilibrium model. *Policy Sciences* 53, 4 (2020), 735–758.

[13] Paul W Holland and Samuel Leinhardt. 1971. Transitivity in structural models of small groups. *Comparative group studies* 2, 2 (1971), 107–124.

[14] Richard A. Holley and Thomas M. Liggett. 1975. Ergodic Theorems for Weakly Interacting Infinite Systems and the Voter Model. *The Annals of Probability* 3, 4 (1975), 643–663. http://www.jstor.org/stable/2959329

[15] Yifan Hu, Yehuda Koren, and Chris Volinsky. 2008. Collaborative filtering for implicit feedback datasets. In *2008 Eighth IEEE international conference on data mining*. IEEE, 263–272.

[16] Bernardo A Huberman, Daniel M Romero, and Fang Wu. 2008. Social networks that matter: Twitter under the microscope. *arXiv preprint arXiv:0812.1045* (2008).

[17] Jisu Kim, Jihwan Aum, SangEun Lee, Yeonju Jang, Eunil Park, and Daejin Choi. 2021. FibVID: Comprehensive fake news diffusion dataset during the COVID-19 period. *Telematics and Informatics* 64 (2021), 101688. https://doi.org/10.1016/j.tele.2021.101688

[18] David M. J. Lazer, Matthew A. Baum, Yochai Benkler, Adam J. Berinsky, Kelly M. Greenhill, Filippo Menczer, Miriam J. Metzger, Brendan Nyhan, Gordon Pennycook, David Rothschild, Michael Schudson, Steven A. Sloman, Cass R. Sunstein, Emily A. Thorson, Duncan J. Watts, and Jonathan L. Zittrain. 2018. The science of fake news. *Science* 359, 6380 (2018), 1094–1096. https://doi.org/10.1126/science.aao2998

[19] Jure Leskovec, Jon Kleinberg, and Christos Faloutsos. 2005. Graphs over Time: Densification Laws, Shrinking Diameters and Possible Explanations. In *Proceedings of the Eleventh ACM SIGKDD International Conference on Knowledge Discovery in Data Mining* (Chicago, Illinois, USA) *(KDD '05)*. Association for Computing Machinery, New York, NY, USA, 177–187. https://doi.org/10.1145/1081870.1081893

[20] David Liben-Nowell and Jon Kleinberg. 2007. The link-prediction problem for social networks. *Journal of the American society for information science and technology* 58, 7 (2007), 1019–1031.

[21] Antonis Matakos, Sijing Tu, and Aristides Gionis. 2020. Tell me something my friends do not know: Diversity maximization in social networks. *Knowledge and Information Systems* 62, 9 (2020), 3697–3726.

[22] Fernando P Santos, Yphtach Lelkes, and Simon A Levin. 2021. Link recommendation algorithms and dynamics of polarization in online social networks. *Proceedings of the National Academy of Sciences* 118, 50 (2021), e2102141118.

[23] Javier Sanz-Cruzado and Pablo Castells. 2018. Enhancing Structural Diversity in Social Networks by Recommending Weak Ties. In *Proceedings of the 12th ACM Conference on Recommender Systems* (Vancouver, British Columbia, Canada) *(RecSys '18)*. Association for Computing Machinery, New York, NY, USA, 233–241. https://doi.org/10.1145/3240323.3240371

[24] Javier Sanz-Cruzado, Sofía M. Pepa, and Pablo Castells. 2018. Structural Novelty and Diversity in Link Prediction. In *Companion Proceedings of the The Web Conference 2018* (Lyon, France) *(WWW '18)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 1347–1351. https://doi.org/10.1145/3184558.3191576

[25] C. Shao, G.L. Ciampaglia, O. Varol, A. Flammini, and F. Menczer. 2017. The spread of fake news by social bots. *CoRR* abs/1707.07592 (2017). arXiv:1707.07592

[26] Zelin Wan, Yash Mahajan, Beom Woo Kang, Terrence J. Moore, and Jin-Hee Cho. 2021. A Survey on Centrality Metrics and Their Network Resilience Analysis. *IEEE Access* 9 (2021), 104773–104819. https://doi.org/10.1109/access.2021.3094196

[27] Kai-Cheng Yang, Francesco Pierri, Pik-Mai Hui, David Axelrod, Christopher Torres-Lugo, John Bryden, and Filippo Menczer. 2021. The COVID-19 Infodemic: Twitter versus Facebook. *Big Data & Society* 8, 1 (2021), 20539517211013861. https://doi.org/10.1177/20539517211013861