What Happened on Election Night?!
Assessing the Effect of Recommender Systems on the Evolution of User-Content Network Polarization

Final Report

Akhil Prakash*
akhilp@stanford.edu

Rohisha Adke*
radke@stanford.edu

Stephen Koo*
sckoo@cs.stanford.edu

December 12, 2016

1 Introduction

The last few years have seen an increasing number of researchers and commentators concerned over political polarization in the United States. The extremity of the spectrum of beliefs has increased, and individuals are more likely to not know people who do not share their beliefs than ever before.

At the same time, social media (Facebook, Twitter) has been increasing in popularity and usage. The amount of content available to each individual has exploded, and to sift through it all, users need the help of automated recommendation engines. Increasingly, in addition to the browsing behavior of users on the sites being influenced by recommendation engines, the content recommended to users affects other aspects of the lives of individuals, such as the news articles people read and products people purchase. Gone are the days when all individuals in a local, spatially-based community would receive one physical newspaper, and all have common ground upon which to base their understanding of the world. Instead, these recommendation engines shape the content each person sees, and the way we subsequently view our world. Individuals on both sides of the political spectrum can be getting their information from completely different sources [9] [8] [5] [13].

Researchers and social media companies have put forth both the theory that polarization has been increasing as a result of social media and their recommendation engines, as well as the theory that social media and recommendation engines are not responsible for the increasing polarization of the American public (and could even help expose individuals to a broader range of perspectives than they would see otherwise). We have assessed these theories using a graph theoretic framework to lend analytical rigor to the question of whether recommender systems are changing the structure of networks, in terms of group polarization, over time. To do this, we built a framework that simulates the evolution of a bipartite graph of users and the articles they read. Using this experimental framework, we have evaluated the effect that different recommender systems have on the polarization of users, in terms of the set of articles that users have read by the end of the simulation.

2 Related Work

2.1 Political Polarization, Social Media, Information Sources

In 2014, a survey of 10,000 Americans conducted by the Pew Research Center found that the ideological polarization of individuals has increased over time [5]. The study tracked responses to 10 questions about political values over a period of 20 years. Respondents were categorized along a spectrum of political values ranging from consistently conservative to consistently liberal. The proportion of Democrats that gave uniformly liberal responses and the proportion of Republicans that gave uniformly conservative responses reached an all-time high by the end of the study in 2014. The study also found that individuals are likely to discuss politics only with people who already share their viewpoints. These results suggest that recommendation engines which only show users results posted by their connections will skew towards showing users information their friends, and therefore, they, already believe.

The Pew Research Center additionally provides us with the percentage of each political group who says they trust each of a set of common news sources (Fox News, NYT, etc.) [5], which we use as the basis for a couple components of our simulated model, as described in the

* All authors contributed equally to this project.

following sections.

Researchers have proposed differing theories regarding the effect of social media on political polarization. Barberá et. al. postulated that social media reduces political polarization because the weak ties in networks tend to be politically heterogeneous, so people are exposed to a range of political opinions, beyond those of their close friends alone [2]. On the other hand, Fernbach et. al. believe that political polarization comes from a lack of understanding; people suffer from "the illusion of explanatory depths," which means that people think they understand things when they actually do not [7]. Thus, the way to combat polarization is not to question views directly, but to ask them how they believe the policy would work. When people realize that they do not understand as much as they previously thought, they become more moderate in their views, and open to new ideas.

2.2 Effects of Recommender Systems

In its most common formulation, a recommender system tries to predict the rating that a user would give to an item. Social media, e-commerce, and other platforms use these predictions to select items to display to their users, in the hope that showing a user items that they are most likely to rate highly would encourage continued consumption of content, merchandise, etc. The simplest recommender system would recommend the most popular items to every user. However, most platforms employ personalized recommendation algorithms, which take into account the user's past ratings, past activity, personal preferences, and even context-sensitive details such as the user's current geographic location.

2.2.1 Anchoring effect

The most salient effect of recommendations is the so-called "anchoring effect," a human cognitive bias toward making decisions based only on the first piece of information offered. Adomavicius et. al. performed a laboratory experiment to measure the influence of the anchoring effect on users' ratings of jokes [1]. The experiment demonstrated that positive perturbations to the recommended ratings led to statistically significant positive rating drifts and the negative perturbations led to commensurate negative rating drifts in the other direction. Adomavicius et. al. also evaluated a couple methods for removing this bias from the ratings, finding that designing a bias-aware user interface is more effective than post-hoc ratings adjustment.

2.2.2 Recommender systems and network evolution

Zhao et. al. use simulations to investigate the effect of recommendation-induced bias over time on network structure [4]. In particular, they model user ratings as a bipartite graph, with undirected edges representing users' ratings connecting user nodes to movie nodes. They simulate the effect of a recommender systems on the evolution of the network, adding edges to the graph deterministically based on the recommendations outputted by the recommender systems. In their analysis, Zhao et. al. measure item degree, item degree heterogeneity, number of square motifs, and clustering coefficient over the iterations in their simulation, and find that all the recommendation algorithms lengthen the tails of the item degree distributions, increase item degree heterogeneity (which measures global diversity), increase the number of squares, and increase the average clustering coefficient.

Su et. al. also investigated the effect of recommendations on network structure and found that Twitter's "Who to Follow" feature benefited more popular users in terms of new followers substantially more than less popular users [16]. The authors of this paper had the unique advantage of access to Twitter subscription data both prior to and following the introduction of the "Who to Follow" feature — which is a link-prediction system that recommends new subscriptions to users — creating the perfect "natural" experiment. The network they analyzed is a directed graph of nodes representing users and directed edges representing subscriptions, or "follows." This paper rigorously demonstrates, both theoretically and empirically, that the standard friend-of-friend algorithm on which "Who to Follow" is based results in a notable "rich get richer" effect: users with more followers are more likely to get a greater number of additional followers due to the recommendations than are users with fewer followers.

3 Methods

To perform our experiment, we built a simulation framework to initialize user nodes and a "friend graph" from a dataset, provide recommendations to users, simulate subsequent user behavior, evaluate the polarization of the graph at each iteration, and evaluate the final structure of the bipartite graph, as well as the collapsed user graph, at the end of the simulation.[1]

The steps of a simulation "run" are:

I) Network initialization (from data)

II) For each iteration:

  i) Create new articles

[1] The code and data for our system can be found at https://github.com/akhilprakash/CS224_Project

ii) Have some users go online

iii) Recommend articles to each online user

iv) Simulate users liking or not liking each article

v) Calculate metrics that are evaluated at each iteration

III) At the end of the simulation, calculate the metrics that are evaluated on end graphs

Each simulation run is parameterized by many options an experimenter chooses between for each of the steps above. The parameters that can be set by the experimenter include: the number of iterations to run the simulation for, dataset to initialize the friend graph from, method for initializing the political orientations of users given a friend graph, number of new articles that are "created" and introduced into the system at each iteration, number of users that are "online" at each iteration, primary recommender system being tested, baseline recommender system to run in the background alongside the primary system being tested, function for calculating the probability that a given user will like an article from a specific source, number of recommendations provided to each online user at each iteration, and metrics to use to evaluate the graph at each iteration, and/or the end of the simulation.

After our initial exploration, we set some parameters as constant throughout all experiments we ran, in order to be able to perform stable and realistic experiments in terms of the parameters we were interested in varying. The parameters we set as constant were:

- Number of iterations to run the simulation for: 100
- Dataset to initialize friend graph: one of either the collaboration network of arXiv general relativity [11], or Zachary's Karate Club friendship network [14]
- Number of new articles that are "created" and introduced into the system at each iteration: 5
- Number of users that are "online" at each iteration: 500
- Baseline recommender system to run in the background alongside the primary system being tested: a random recommender system
- Number of recommendations provided to each online user at each iteration: 30 (10 from baseline recommender, and 20 from primary system being tested)

Below, we elaborate on the details of each option for each parameter that we chose to vary, as well as the metrics we used to evaluate our results. To run our experiments, the parameters we varied were:

- Method for initializing the political orientations of users given a friend graph
- Primary recommender system being tested
- Function for calculating the probability that a given user will like an article from a specific source

Varying these parameters allowed us to test the effect of each recommender system on our polarization metrics under different assumptions one could make about the political orientations of users in a friend graph (whether people tend to be friends with users of similar political orientation, or not), as well as how the political orientation of a user influences the likelihood that the user will like an article from a specific source (whether the political orientation does, or does not, influence the probability that a user will like an article from a specific source).

I) Network Initialization

We validated our results by initializing our network of users from two existing datasets: Zachary's Karate Club friend graph[14], and the network of collaboration of researchers in general relativity and quantum cosmology [11] (network has an edge between users if they have co-authored at least one paper). Each of these graphs have different qualities; most notably, the former has 34 nodes, while the latter has 5,424. Each can be used as a sample "friend" graph in the context of our simulation; the former is a friend graph, while the latter provides a network of colleagues.

For each experiment "run", we select one of these two seed networks of nodes and edges, then perform one of the two procedures below to initialize the political orientation of each user node. The first initialization strategy (random initialization) allows us to run our experiment under the assumption that individuals are equally likely to be friends with individuals of any political orientation, while the second allows us to run our experiment under the assumption that individuals are more likely to be friends with individuals that have similar political orientations to themselves.

Random Initialization

With the random initialization option, each user is assigned a political orientation by sampling from an orientation distribution (collected in a Gallup poll of the American public)[10]:

Political Orientation	Probability of user having this political orientation
-2	0.1
-1	0.2
0	0.4
1	0.2
2	0.1

Each assignment is therefore made independently of the underlying friend network.

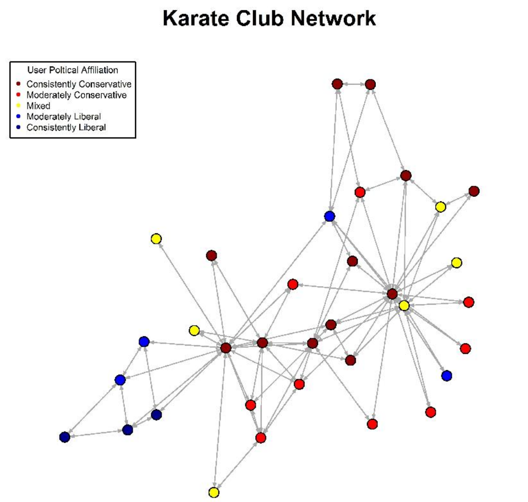

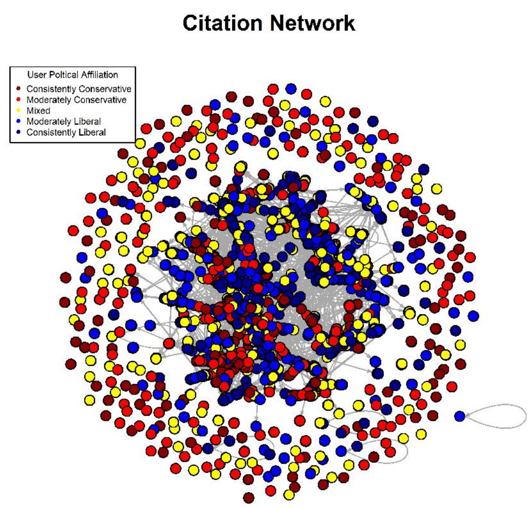

Figure 1: Karate Club and Citation Network graphs after propagation initialization. Red nodes are conservative. Blue nodes are liberal. Yellow nodes are mixed.

Propagation Initialization

The propagation initialization option propagates political orientation from each node to its neighbors, similarly to how a contagion model works. For each connected component, we identify the longest shortest path in the friend graph. Without loss of generality, we assign the source node to 2 and the destination node to -2. These nodes will be on opposite ends of the graph. From there, each node has a certain probability of propagating its political orientation (o), $o-1$, or $o+1$ to its neighbors, as long as each would be within $[-2,2]$. Thus, friends only have a political difference of 1. There is a decay factor as o propagates, so that we see the full spectrum of political orientations across the graph.

i) New articles created

At each iteration we introduce five new articles into the simulation. Each new article is assigned a source (e.g. CNN, Fox News, NYT), selected uniformly at random from the set of 36 sources considered in the Pew study on the degree to which each political group trusts common news sources [5]. Each article is also assigned a "lifetime" in number of iterations, sampled from an exponential distribution, such that the article can no longer be recommended to users after its lifetime, or period of relevance, has passed. In our experiments, this exponential distribution is parameterized such that articles "live" on average for one tenth of the entire simulation, i.e. 10 iterations.

ii) Certain users go online

In most social networks, nodes "wake up" to create edges in time intervals following a power law distribution with an exponential cutoff, parameterized by the degree of the node [12]:

$$\Pr(\delta_d) \propto \delta_d^{-\alpha} e^{-\beta_d} \tag{1}$$

where δ_d is the time interval between now and the next time a node with degree d will wake up, and α and β_d are parameters of the distribution. Thus, nodes with higher degree are more likely to create edges more often than nodes with lower degree. In our simulation, we set $\alpha = 0.8$ and $\beta_d = 0.5d$, then sample δ_d from each user's distribution. To ensure that enough edges are created in our graph during a simulation, we simply select the 500 users with largest δ_d to go "online." For the collaboration network, 500 amounts to about 10% of the total set of users.

iii) Recommend articles to online users

The experimenter can choose one of six implemented recommender systems to use for a specific simulation run. For each user that "goes online," the chosen recommender system selects a set of $N = 20$ articles that the user has not liked yet to "display" for the user.

We implemented three simple recommender systems: the "Random" recommender recommends articles uniformly randomly, and the "Popular" recommender selects the articles with the most likes to display for the user. The "Content-Based" recommender uses empirical trust statistics collected in the 2014 Pew study [5]: for each major news source and political orientation, the study measured the percentage of people with the given political orientation that indicated they trust the given news source. The "Content-Based" recommender simply selects the articles that come from sources with the greatest trust percentages, conditioned on the user's political orientation. On some iterations, it is possible that an online user has liked too many articles, such that there are not enough articles in the entire system to meet the quota of N recommendations; in this case, these recommender systems simply select as many articles as possible.

We also implemented two graph-based recommender systems: the "Friend-Based" recommender randomly selects articles uniformly from the articles that the user's immediate friends have liked in previous iterations. If the user's friends have not liked enough articles to meet the quota of N articles, the remaining recommendations will be filled out by either the "Random" recommender or the "Content-Based" recommender; we label these composite recommender system as "Friend-Based with Random Default" and "Friend-Based with Content-Based Default," respectively. The "Collaborative Filtering" (CF) recommender implements an item-based collaborative filtering method [6]. In our implementation, the similarity between each pair of articles (x, y) is computed as:

$$s_{xy} = \frac{|L_x \cap L_y|}{|L_x \cup L_y|} \tag{2}$$

Where L_x and L_y are the set of users that liked the articles x and y, respectively. Then, a score for each article x is computed as the sum of the similarities between the candidate article and the articles that the user u has liked:

$$\mathrm{score}_u(x) = \sum_{y \in \tilde{L}_u} s_{xy} \tag{3}$$

where $\tilde{L}_u$ is the set of articles that user u has liked. The CF recommender then selects the articles with the highest scores.

At each iteration, we also recommend $N_0 = 10$ articles using the "Random" recommender to give new articles a chance to receive likes, to handle the fact that most of the other recommender systems will never recommend an article that does not have any likes yet.

iv) Simulate users liking recommended articles

Each simulation is parameterized by one of three "pLike" models, which determines the probability that a given user will "like" a recommended article.

Uniform pLike

The uniform pLike model assumes that a user will like a recommended article with a fixed probability of 20%, irrespective of the article's source, or the user's political orientation.

Empirical pLike

The empirical pLike model assumes that a user with political orientation π will like a recommended article from source σ with a probability equal to the empirical percentage of people with political orientation π that trust σ, as reported by the Pew study [5].

Individual pLike

The individual pLike model reintroduces variance to the mean "likelihood of trusting source given political orientation" reported in the Pew study by, for each user and article, sampling the probability that the user will like the article from a normal distribution centered around the empirical model's probability of liking, with a variance that is parameterized by the user's political orientation $(1 - \frac{|\text{user political orientation}|}{2} + 0.01)$. This gives us a more realistic pLike, in which each individual user has a unique probability of liking each article, which is influenced but not determined by the user's political orientation.

v) Run iteration-specific metrics

At each iteration, we compute and plot metrics which allow us to assess the evolution of the graph over time.

Item-degree heterogeneity

The item-degree heterogeneity (IDH) measures the unevenness of the degree distribution of the articles in the user-article graph [18], and is computed as the ratio of the mean square degree and the square of the mean degree:

$$\mathrm{IDH}(G) = \frac{\sum_i k_i^2}{\left(\sum_i k_i\right)^2} \tag{4}$$

where k_i is the degree of the ith article. Note that a larger IDH indicates a more uneven distribution of degree among the nodes in the graph.

Number of squares

The local clustering coefficient as defined by Watts and Strogatz [17] is zero for every node in a bipartite user-article graph. Opsahl et. al. generalize the clustering coefficient for bipartite graphs by considering square motifs rather than triangle motifs [15]. Conceptually, square motif represents a "reinforcement" pattern, where two users both like two of the same articles. The full definition of this global clustering coefficient is defined as the ratio of the number of squares to the number of paths of length 4; however, computing the latter is costly and inefficient, thus we just use the number of squares (NoS) as a rough proxy for the clustering coefficient. Although the NoS cannot serve as an absolute measure of clustering, we use it for relative comparisons of our simulations, since the node and edge sets grow roughly at the same rate across all of our experiments. We estimate the NoS of the user-article graph using the first 100 eigenvalues of its adjacency matrix [3]:

$$\mathrm{NoS}(G) = \frac{1}{12}\sum_{i=1}^{100}\lambda_i^4 \tag{5}$$

where λ_i is the ith largest eigenvalue of the adjacency matrix.

Average standard deviation of readership

We also compute the standard deviation of the political orientations of the users who liked each article, and compute the mean of this value across all articles. This average standard deviation of readership (ASR) serves as a measure of the "echo chamber effect": a lower ASR indicates that the readership of each article has a more uneven distribution of political orientations, i.e. some articles are mostly liked by conservative users, and others are mostly liked by liberal users, leading to different silos of exposure to opinion. To ensure a reliable ASR, we compute the average over only articles that were liked by at least five users.

For reference, some values the standard deviation of readership of an article could have include i) 2.83, which occurs when two users, one of political orientation 2, and one of -2, like an article, ii) 1.41, if an article has an infinite number of likes from all types of users, iii) 2, if an article has an infinite number of likes from users of political learning 2 and an infinite number of likes from users of political leaning -2.

III) Calculate metrics on network's final realization

Hierarchical Clustering (run on user-user graph)

An option for post-simulation analysis of network polarization we explored was clustering on the nodes of the network, after collapsing the bipartite graph of users and articles read into a user-user graph, where the edge weight between two users indicated the number of articles that both users read. This allowed us to assess the polarization in the base of articles and knowledge shared between each pair of users. One clustering method we experimented with was hierarchical clustering to cluster the user nodes into groups. From the relative ability, or inability, of the clustering algorithm to identify meaningful clusters, we can infer the possible existence of polarization in the graph; if there are identifiable clusters of users that read similar sets of articles, it is more likely that readers are polarized in terms of the articles upon which they are basing their knowledge of the world.

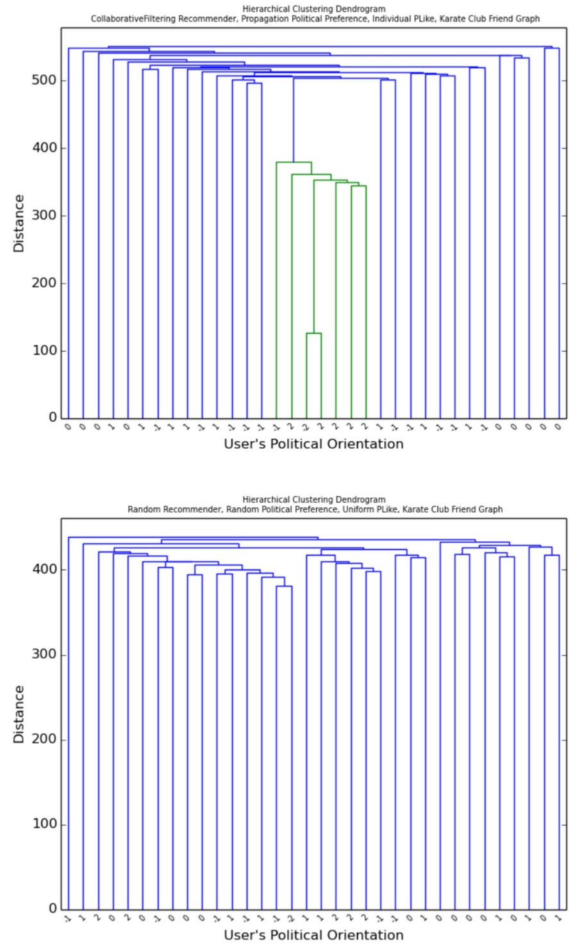

Figure 2: Dendrograms resulting from hierarchical clustering of the users by the Jaccard similarities between their like sets.

In Figure 2, we see two different dendrograms produced by different experiments on the Karate Club network. In one, the clusters identified by hierarchical clustering reflect the political orientation of users in the cluster, while in the other, cluster membership does not tell us much about the likely political orientation of a user. This means that polarization in terms of articles read by users of certain political orientations is more likely in the first simulation than the second.

Article readership distribution

For each simulation run, in addition to the collapsed standard deviation of the political orientation of the users that like each article, we assess the full distribution of user political orientations (Figure 3).

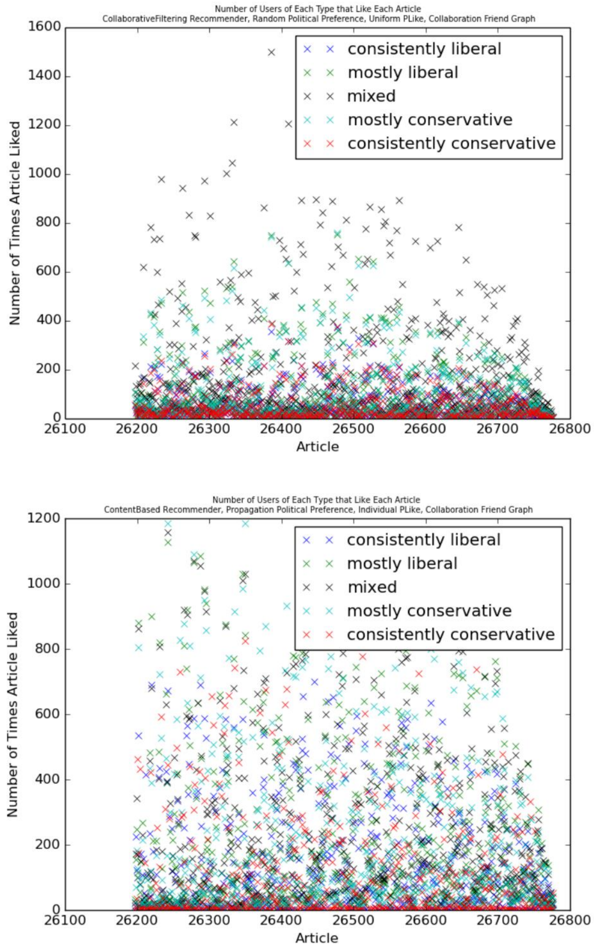

Figure 3

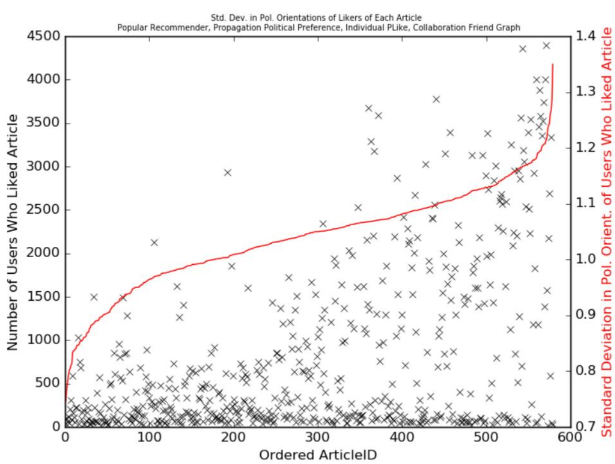

Figure 4

Miscellaneous graph statistics

To assess the simulation and effect of the parameters chosen for the experiment, we also compute various statistics such as the number of articles liked by each user, number of times each article was liked, and number of users of each political orientation in the initialized graph (Figures 4 and 5).

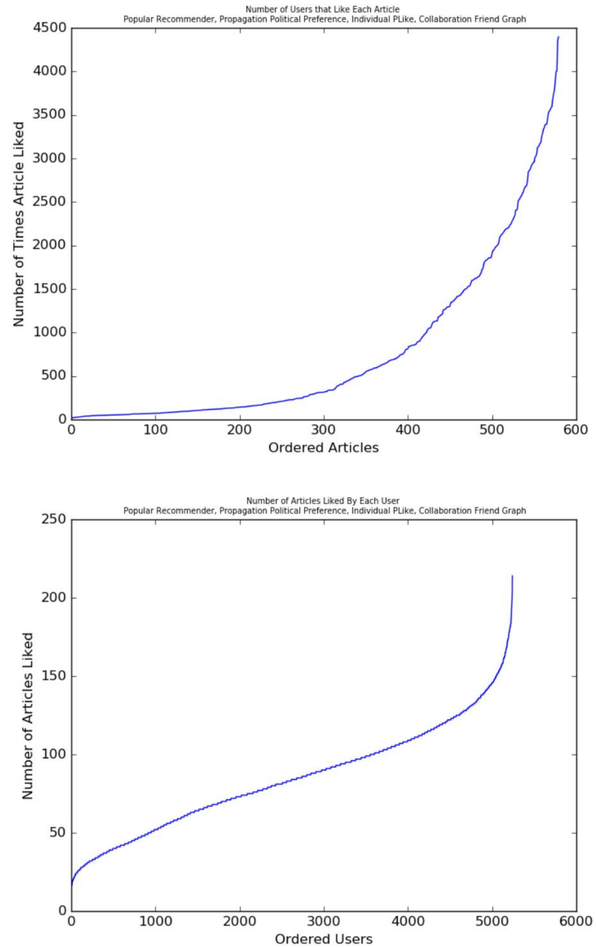

Figure 5

Proportion of Likes for each Article by User Political Leaning

In order to visualize the distribution of user likes on the articles of high variance, we create stacked bar charts where each bar represents a single article, and we can see what proportion of the likes from each article came from a specific user political leaning (Figure 6). We can also do the same thing for sources instead of articles and figure out which sources are creating the most polarizing content and which are creating the least polarizing content. We subset and look at only the 30 articles with the highest variance.

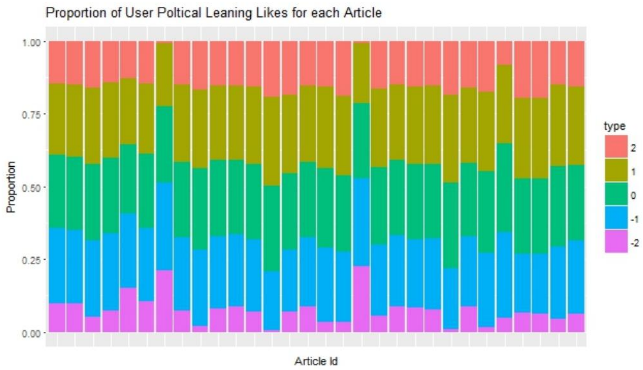

Figure 6: This is for the content based recommender system, with propagation initialization, and individual pLike. We see that most of the distributions are roughly the same except there are a few conservative articles.

4 Results

As described above, to do our analysis, we ran a simulation for each of the 36 possible combinations of the network initialization methods, recommender systems, and pLike models. Below, we detail the results of comparing our metrics across these 36 combinations.

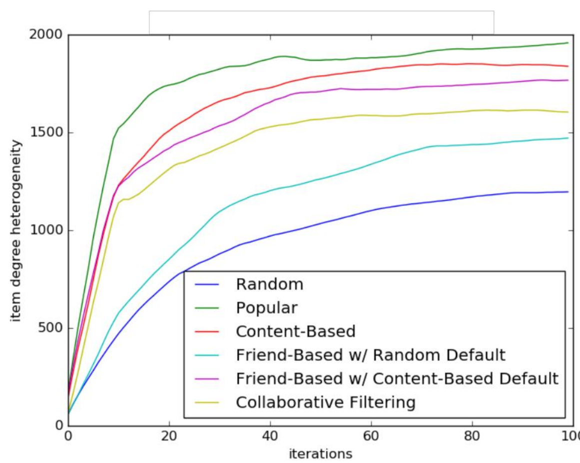

Figure 7: A comparison of the evolution of item-degree heterogeneity over the course of simulations using propagated political orientation initialization and the Individual pLike model.

We observed that while the choice of pLike model shifts the scale and magnitude of the item-degree heterogeneity, the ordering of the effects of the recommender systems is generally consistent regardless of the other parameters. We found that simulations that use the popularity-based recommender system consistently lead to some of the highest values of IDH (Figure 7). This is probably because showing only the most popular articles to every user will lead to a "rich get richer" effect where articles with the most likes will continue accumulating more likes, leading to a highly skewed article degree distribution. This is similar to the findings of Zhao et. al. [16]. As expected, the Random recommender system also consistently leads to the lowest IDH values. While IDH tells us about the skewness in the distribution of number of likes each article gets, the metric alone does not tell us about article readership polarization, since number of likes alone does not tell us which users were liking an article.

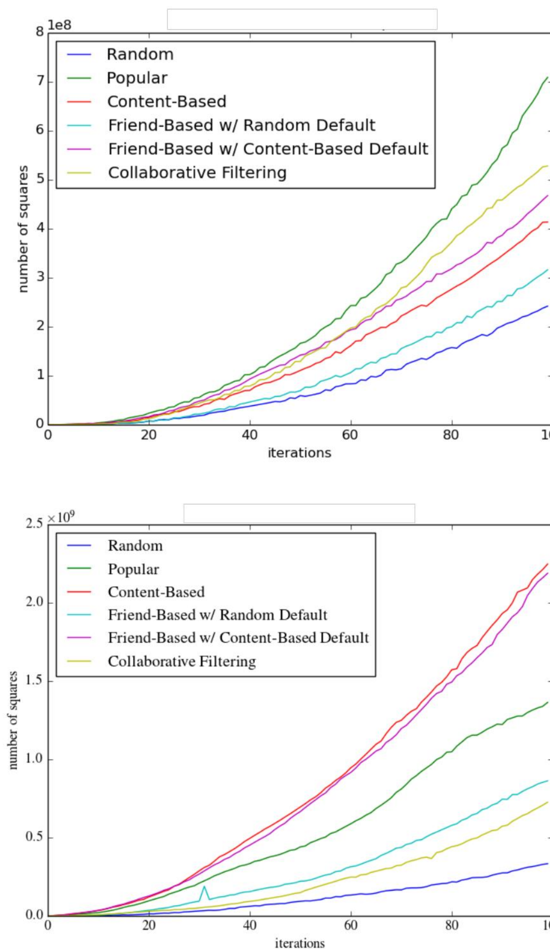

Figure 8: A comparison of the evolution of number of squares over the course of simulations using random political orientation initialization with the Uniform or the Empirical pLike model.

The number of squares illustrates interesting differences in the relative effect of the recommender systems under different conditions. For example, we see that the "Popular" recommender leads to the greatest growth in NoS under the Uniform pLike model, but the "Content-Based" recommender leads to the greatest growth in NoS under the Empirical pLike model (Figure 8). Since the "Popular" recommender frequently and repeatedly selects the same set of recommendations, we end up with a situation where a large population of the users all like the same set of the most popular articles, leading to a high NoS. However, when the users have more biased preferences (as under the Empirical and Individual pLike models), a higher NoS can be achieved by properly tailoring recommendations to the users' preferences (as does the "Content-Based" recommender), encouraging people of similar political orientations to like the same common articles. However, similarly to IDH, a high NoS also does not serve as an accurate indication of polarization, since a local cluster could be composed of users with a wider or narrower range of political orientations while still contributing to the overall clustering coefficient.

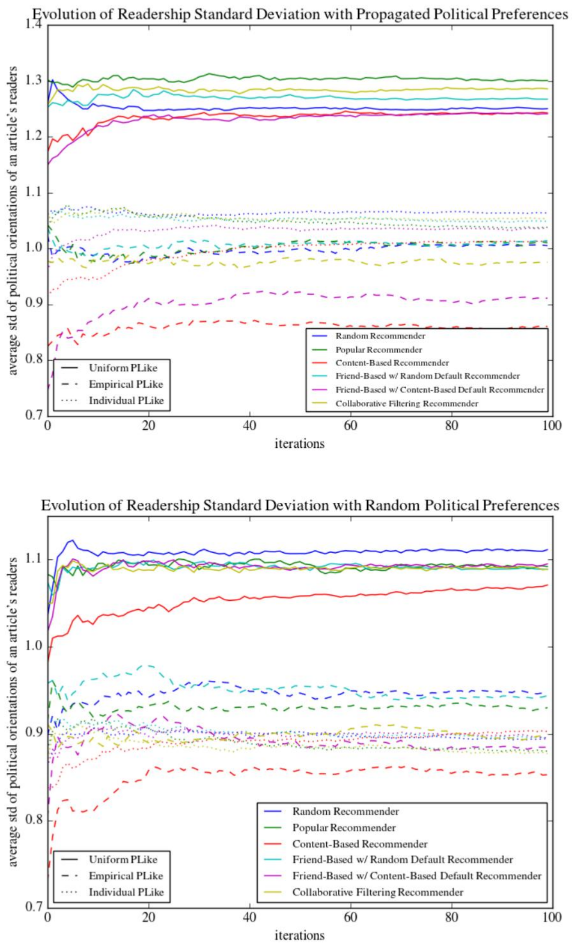

Figure 9: A comparison of the evolution of the ASR over the course of simulations using either (a) propagated or (b) random political orientation initialization.

The standard deviation in political orientations of users that liked each article is most telling regarding the effect of each of our experiment's parameters on final polarization in the articles users of each political orientation read in the course of the simulation. In the figures above, we compare the average standard deviation across all articles at each iteration of each simulation. As one would expect, we note that when users have a Uniform pLike function, for a specific network initialization strategy, we see the greatest variance among the orientations of users that like each article — this is because the probability that a user will like an article does not depend upon that user's political orientation.

Striking in Figure 9 is the effect that pLike has, no matter the recommender system, on article readership polarization. Varying recommender systems does make a difference in how polarization evolves, but resulting effect size is on the same order of the magnitude as (if not smaller than) that of the choice of pLike model.

Tempering this observation is the curve that results from the simulation run with the "Content-Based" recommender with a network initialized randomly, and uniform pLike across users (red line in Figure 9b). We see that the ASR during this experiment dips nearly as low as the standard deviations we see using a variety of other recommender systems, the empirical/individual pLikes, and propagated network initialization. This is likely because in this experiment, the "Content-Based" recommender is optimizing for an empirical pLike even though users actually have a uniform pLike function, and are not influenced by their political orientations when deciding whether to like an article. This tells us that even if users have a uniform and unbiased probability of liking any article, when they are only shown articles that the recommender system thinks they will like because of their presumed "type" (i.e. political orientation), the users' liking behavior will end up appearing to be biased by their political orientation. This serves as a cautionary tale for those who wish to empirically determine whether a user's orientation influences liking behavior using systems in which users were only shown a biased set of articles as options to begin with, as is true for each system involving a non-random recommender system.

5 Conclusion and Next Steps

We began this paper hoping to evaluate how recommender systems influence polarization of a user network over time. Our analysis has led us to conclude that, while the recommender system is important in determining polarization, equally important are the assumptions an experimenter makes about how political orientations are distributed across a network to begin with, and particularly, how the political orientation of a user influences, or does not influence, the probability that he/she will like an article.

As discussed in the analysis section, we were also interested to find that, though no recommender system can induce a readership diversity as high as that resulting from a uniform pLike, forcing the "wrong" recommender system (system that is optimizing for a non-random pLike) on a population with a uniform pLike can still induce a readership diversity as low as what we see with a non-uniform pLike.

In the future, the simulation framework we built could be made available to researchers and members of the public to test out their own theories on the effect of online systems on user polarization and network properties. Our simulation framework supports a wide variety of parameters, including some not mentioned in this report, leaving a significant amount of unexplored territory in the space of possibilities. The framework could be made further nuanced by removing some of the assumptions we made to make the project tractable (i.e. we conflate reading, liking, and sharing). Real life is not made up of normally distributed randomness, and further incorporating empirical data into various components of our framework would allow us to better model this reality.

An important takeaway from this project, and intriguing candidate for future research, is the importance of the assumption about each user's probability of liking an article — and how this changes over time, with each article a user is exposed to. Users are not static objects that like or do not like articles, but rather, human beings influenced by content, and evolving along with the ideas to which they are exposed. However, as long as we treat users as static and unchanging, we will continue recommending articles based only on their historical behavior, and our simulation frameworks will lack the ability to model critical nuances of human behavior. Perhaps the arguments made by those creating online systems that users will "like what they like" are most revealing of their perception of users as fixed actors instead of as influenceable people reading and changing their beliefs as a result of what they read over time. Future work on this topic should strive to measure and model this effect of reading articles on users, and determine ways for recommender systems to incorporate a user growth and learning model into their recommendations.

References

[1] Gediminas Adomavicius, Jesse Bockstedt, Shawn Curley, and Jingjing Zhang. De-biasing user preference ratings in recommender systems. In *RecSys 2014 Workshop on Interfaces and Human Decision Making for Recommender Systems (IntRS 2014), Foster City, CA, USA*, pages 2–9, 2014.

[2] Pablo Barberá, John T Jost, Jonathan Nagler, Joshua A Tucker, and Richard Bonneau. Tweeting from left to right is online political communication more than an echo chamber? *Psychological science*, page 0956797615594620, 2015.

[3] Steven Kay Butler. *Eigenvalues and structures of graphs*. ProQuest, 2008.

[4] Zhao Dan-Dan, Zeng An, Shang Ming-Sheng, and Gao Jian. Long-term effects of recommendation on the evolution of online systems. *Chinese Physics Letters*, 30(11):118901, 2013.

[5] Michael Dimock, Jocelyn Kiley, Scott Keeter, and Carroll Doherty. Political polarization in the american public. *Pew Research Center*, 2014.

[6] Michael D Ekstrand, John T Riedl, and Joseph A Konstan. Collaborative filtering recommender systems. *Foundations and Trends in Human-Computer Interaction*, 4(2):81–173, 2011.

[7] Philip M Fernbach, Todd Rogers, Craig R Fox, and Steven A Sloman. Political extremism is supported by an illusion of understanding. *Psychological science*, page 0956797612464058, 2013.

[8] Sounman Hong and Sun Hyoung Kim. Political polarization on twitter: Implications for the use of social media in digital governments. *Government Information Quarterly*, 2016.

[9] Shanto Iyengar and Kyu S Hahn. Red media, blue media: Evidence of ideological selectivity in media use. *Journal of Communication*, 59(1):19–39, 2009.

[10] Jeffrey M. Jones. Democratic, republican identification near historical lows. 2016.

[11] Jure Leskovec. *https://snap.stanford.edu/data/*.

[12] Jure Leskovec, Lars Backstrom, Ravi Kumar, and Andrew Tomkins. Microscopic evolution of social networks. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 462–470. ACM, 2008.

[13] Amy Mitchell, Jeffrey Gottfried, Jocelyn Kiley, and Katerina Eva Matsa. Political polarization & media habits. *Online abrufbar unter: http://www. journalism. org/2014/10/21/political-polarization-media-habits*, 2014.

[14] University of California at Irvine. *https://networkdata.ics.uci.edu/data.php?id=105*.

[15] Tore Opsahl. tnet» two-mode networks» clustering.

[16] Jessica Su, Aneesh Sharma, and Sharad Goel. The effect of recommendations on network structure. In *Proceedings of the 25th International Conference on World Wide Web*, pages 1157–1167. International World Wide Web Conferences Steering Committee, 2016.

[17] Duncan J Watts and Steven H Strogatz. Collective dynamics of 'small-world'networks. *nature*, 393(6684):440–442, 1998.

[18] Qian-Ming Zhang, An Zeng, and Ming-Sheng Shang. Extracting the information backbone in online system. *PloS one*, 8(5):e62624, 2013.