The Role of Polarization and Diversity in Social Networks

A Thesis

submitted to the designated
by the General Assembly of Special Composition
of the Department of Computer Science and Engineering
Examination Committee

by

Ioannis Kouvatis

in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE IN COMPUTER SCIENCE

WITH SPECIALIZATION
IN SOFTWARE

University of Ioannina
March 2018

Examining Committee:

- **Evaggelia Pitoura**, Professor, Department of Computer Science and Engineering, University of Ioannina (Supervisor)
- **Panayiotis Tsaparas**, Associate Professor, Department of Computer Science and Engineering, University of Ioannina
- **Nikos Mamoulis**, Associate Professor, Department of Computer Science and Engineering, University of Ioannina

Dedication

Dedicated to my family and friends.

Acknowledgements

Foremost, I would like to express my sincere gratitude to my advisor Professor Evaggelia Pitoura and my co-advisor Associate Professor Panayiotis Tsaparas for the continuous support of my MSc study. Their guidance helped me in all the time of research and writing of this thesis.

I am also grateful to the last member of my thesis committee, Associate Professor Nikos Mamoulis, for his encouragement and insightful comments.

Last but not least, I would like to thank my family and friends. Their great help and support kept me going through all these years. Without them this thesis could not have been possible.

Table of Contents

List of Figures

List of Tables

List of Algorithms

Abstract

Ioannis Kouvatis, M.Sc. in Computer Science, Department of Computer Science and Engineering, University of Ioannina, Greece, March 2018.
The Role of Polarization and Diversity in Social Networks.
Advisor: Evaggelia Pitoura, Professor.

In this thesis we study the role of polarization and diversity in social networks. More specifically, our work consists of three parts. The first part is about identifying whether a topic is controversial and causes people to split into opposing sides. To achieve this, we use an existing model from the literature. The second part examines the occurence of homophily in these topics. Finally the third part deals with the diversity of users and networks. We propose methods to quantify diversity, by taking into account not only the structure of a network, but also the content that users have shared. Based on information about diversity, we focus on the existence of 'echo chambers' in these networks. Moreover, we show that there exists a negative correlation between polarity and diversity in both user and network levels. Finally, we propose algorithmic techniques to recommend content to a user, so that the user is exposed to views of the opposing sides and increase diversity of a topic. We experiment on real data collected from Twitter, and we demonstrate that our methods perform well in practice.

Εκτεταμενη Περιληψη

Ιωάννης Κουβάτης, Μ.Δ.Ε. στην Πληροφορική, Τμήμα Μηχανικών Η/Υ και Πληροφορικής, Πανεπιστήμιο Ιωαννίνων, Μάρτιος 2018.
Ο Ρόλος της Πόλωσης και της Ποικιλομορφίας στα Κοινωνικά Δίκτυα.
Επιβλέπων: Ευαγγελία Πιτουρά, Καθηγήτρια.

Τα κοινωνικά δίκτυα έχουν γνωρίσει τεράστια ανάπτυξη τα τελευταία χρόνια. Παρά το γεγονός πως δημιουργήθηκαν έχοντας ως στοχο να φέρουν ανθρώπους από όλον τον κόσμο πιο κοντά, μέσω της εύκολης ανταλλαγής απόψεων και της πρόσβασής τους σε άφθονη πληροφορία, φαίνεται πως είναι υπεύθυνα για την εμφάνιση ανεπιθύμητων φαινομένων, όπως το φαινόμενο της πόλωσης και της έλλειψης ποικιλομορφίας. Οι χρήστες των κοινωνικών δικτύων προτιμάνε να συναναστρέφονται περισσότερο με άτομα που έχουν τις ίδιες ιδέες, αντιλήψεις και νοοτροπίες, παρά με άτομα τα οποία θα μοιραστούν μαζί τους διαφορετικές απόψεις σχετικές με ένα θέμα συζήτησης. Με αυτόν τον τρόπο το χάσμα μεταξύ των ατόμων αντί να γεφυρωθεί γίνεται ακόμα πιο έντονο. Αυτό μπορεί να επιφέρει επιζήμιες συνέπειες στις κοινωνίες, κάνοντας τες λιγότερο δημοκρατικές και ανοιχτές σε αλλαγές και διαφορετικές γνώμες.

Σε αυτή την εργασία μελετάμε το ρόλο της πόλωσης και της ποικιλομορφίας στα κοινωνικά δίκτυα. Πιο συγκεκριμένα, η δουλειά μας χωρίζεται σε τρία μέρη. Το πρώτο μέρος αφορά την αναγνώριση των θεμάτων που είναι αμφιλεγόμενα και χωρίζουν τον κόσμο σε αντίπαλα στρατόπεδα, ανάλογα με την άποψη τους. Μετράμε το μέγεθος της πόλωσης που δημιουργούν, χρησιμοποιώντας μία μέθοδο από τη βιβλιογραφία που λαμβάνει υπόψιν της τη δομή του δικτύου. Το δεύτερο μέρος έχει να κάνει με την ύπαρξη ομοφυλίας (homophily) στα παραπάνω θέματα. Το τρίτο και τελευταίο μέρος πραγματεύεται την ποικιλομορφία των χρηστών και των δικτύων. Προτείνουμε μεθόδους για να μετρήσουμε την ποικιλομορφία, λογαριάζοντας τη δομή του δικτύου, καθώς και το περιεχόμενο που οι χρήστες έχουν μοιραστεί. Έχοντας λοιπόν ορίσει την ποικιλομορφία, σαν πρώτη εφαρμογή μπορούμε να διαπιστώσουμε την ύπαρξη θαλάμων αντήχησης (echo chambers) στα δίκτυα μας. Επιπλέον, μελετάμε τη σχέση που έχουν η πόλωση και η ποικιλομορφία. Δείχνουμε ότι υπάρχει μια αρνητική συσχέτιση μεταξύ αυτών των μεγεθών, τόσο σε επίπεδο χρήστη, όσο και σε επίπεδο δικτύου. Τέλος, παρουσιάζουμε αλγοριθμικές τεχνικές, ώστε να προτείνουμε περιεχόμενο άλλων χρηστών σε ένα χρήστη, με σκοπό την έκθεση του σε απόψεις διαφορετικές από τις δικές του και την αύξηση της ποικιλομορφίας του σχετικά με ένα θέμα. Πειραματιζόμαστε με πραγματικά δεδομένα από την πλατφόρμα του Twitter, και αποδεικνύουμε οτι οι μέθοδοι μας δουλεύουν καλά στην πράξη.

Chapter 1

Introduction

1.1 Scope

Polarization is a phenomenon where people are divided into opposing subgroups because of differences between them. It occurs quite commonly in society, due to the fact that people tend to interact more often with like-minded individuals, rather than with individuals that share different points of view about a theme. This situation is widely known as homophily [1] [2].

Given the widespread diffusion of online social networks and social media (such as Facebook, Twitter etc.), polarization over a variety of themes has been largely increased. Although social networks have the potential to make people more diverse by exposing them to different ideas and viewpoints, the easy access in plenty of information causes the opposite effect [3] [4].

Social media cause the creation of echo chambers, a situation in which information, ideas and beliefs of a user are reinforced and endorsed by the people that have the same point of view [5]. This leads people to encounter less diverse opinions about a theme, as they become unexposed to arguments from different directions. The existence of these echo chambers blocks the democratic process and the functioning of society at large, as they cultivate isolation and misunderstanding across people of the society. This problem is augmented when social media recommend content based on what users already know, on what are their current connections, and what people that are similar to them have done and liked before. This selective exposure, known as 'filter bubble', instead of making people having more diverse views, leads to increased polarization [6] [7].

Polarization mostly appears in themes from the political domain [8] [9], but also occurs in themes from other domains (music, sports, movies, religion etc.) [10]. This separates individuals into two or more opposing partitions which have less communication. This phenomenon may have a detrimental impact on societies, by making them less democratic and diverse. Thus, it is very important to burst this filter bubble effect and find ways to expose users to content that opposes their beliefs, in order to escape their echo chamber, increase their diversity for a theme and reduce polarization [11].

In this thesis, we study the co-occurence of polarization and diversity in social networks on themes from different domains. We consider not only network (induced by a theme) features, but also content features, exploiting what users have shared in Twitter platform. Our study consists of three parts.

(i) In the first part we deal with identifying which themes are the most controversial and we measure the amount of controversy for each one of them using an existing method from the literature [12]. This method takes into account only the (clustered) structure of the network.

(ii) In the second part, we examine the homophily that exists in different themes, between intra-cluster and inter-cluster users, by considering the content that the users have produced. We also define user level homophily and address the correlation between polarity and homophily of users. We show that there exists a low positive correlation between these two.

(iii) Finally, in the third part we deal with the diversity of network and users. We propose ways to quantify diversity, by taking into consideration the content of users, as well as the structure of the network. Based on information about diversity of users we focus on the occurence of echo chambers in these networks. Furthermore, we study the correlation between diversity and polarity both at the user and at the network level. Finally, we propose algorithmic techniques to recommend content -in our case tweets related to the specific hashtag- to a user, so that if the user endorses them, his diversity will increase.

We cast the last problem as an edge recommendation problem to the retweet graph. A new edge is created, when a new retweet appears between two users. So, we want to recommend to a user, tweets of a set of users that are more likely to by approved and retweeted by him, containing contrarian content and make him open to different viewpoints and have a more diverse opinion about a theme.

We mainly use the retweet graph as opposed to the followers graph, since endorsement features, like retweets, are more useful than a simple connection between two users that follow each other. An edge that exists between two users in the retweet graph is a strong indication that these two users have a common point of view for a specific topic t. On the other hand, a connection between two users in a follow network does not mean that these users agree on this topic.

For example, a user u may follow user v because they agree on another topic t', but they do not agree for topic t. So the edge that may be created in the graph will not be helpful. Due to the fact that we are interested in studying specific themes separately, we find that recommending edges to the retweet graph is the most effective approach for our problem. Besides that, recommending and connecting two users u, v in the follow graph does not guarantee us that user u will see and endorse the tweets of user v for this specific theme, because u may not be interested in them.

To study the aforementioned problems, we perform a case study that uses real data retrieved from Twitter. We choose ten different themes (hashtags), some from the political domain and some from other domains. We also expect some of them to be more controversial and some of them less controversial.

In summary, in this thesis we make the following contributions:

- We measure the amount of controversy for each network created by a theme using an existing model [12] and we evaluate the results by exploiting the labels we have for each user.

- By taking into account the content (tweets) of each user, we consider the homophily between users in the same side of controversy and between users in different sides. We also define user level homophily.

- We define and quantify diversity of users and network, considering the content of users and the structure of the networks. Based on this definition we study the existence of echo chambers.

- We focus on the correlation between polarity and diversity on user and network level. Results show that there exists a negative correlation between these two.
- We propose two algorithms that recommend content of users to a user in order to improve his diversity.
- We experiment on real data collected from Twitter and we demonstrate that our algorithms perform well in practice not only in terms of diversity, but also in terms of polarity.

1.2 Thesis Outline

The remaining of this thesis is organized as follows: In Chapter 2 we provide related work that has to do with measuring and reducing polarization, as well as studies that deal with diversity and recommendation algorithms. In Chapter 3 we provide information about the data we collected for our experiments and we describe the first part of our study, that is measuring controversy of themes by implementing the three stages pipeline [12]. We also calculate user level polarity and we provide statistics about each one of these themes. In Chapter 4 we provide details about the text mining techniques we used to label users' tweets. We measure the homophily between intra cluster users and inter cluster users, we define user level homophily and we depict some results that show the Polarity-Homophily correlation. Finally, in Chapter 5 we define and quantify diversity, we focus on the existence of echo chambers, we consider the correlation between diversity and polarity and we propose two algorithms that recommend users' content to a user, to improve his diversity. We conclude the thesis and we provide some future work in Chapter 6

Chapter 2

Related Work

In this Chapter, we present the state of the art in the related literature on the topic of this thesis. First, we present case studies about polarization and diversity in social networks and then we compare our work with these studies.

2.1 Quantifying and reducing polarization

Garimella et al. in [12] performed a study of controversy detection taking into advantage the structure of social media networks. They focused on comparing themes in any domain. To quantify controversy, they introduced a graph based pipeline that consists of three stages:
(i) Build a conversation graph about a theme
(ii) Partition the conversation graph to find potential sides of controversy
(iii) Measure controversy of characteristics of the graph

They performed a comparison of controversy measures and they showed that their random walk based measure outperforms the existing ones. They also explored alternative approaches to measure controversy that use only the content of the discussion, rather than the structure of user interactions. They tested two types of features extracted from the content: bag of words representation and sentiment analysis.

Except of this study, that measures the global controversy of a theme, the same authors introduced polarity [11] in a user level. They proposed a metric that indicates how polarized a user is, that is how involved he is in his partition. We adopt this approach to measure the polarity of users in our themes.

In the same paper, the authors described some algorithmic techniques in order to reduce controversy. They represented the discussion of a controversial topic using an endorsement graph and cast their problem as an edge recommendation problem on this graph. They want to find edges that bridge the opposing sides and produce the biggest reduction in the controversy score. They also considered the acceptance probability, that is how likely the edge is to appear in the endorsement graph.

The algorithm that they propose works as follows: For each edge that connects two high degree vertices of the two sides, it computes how much RWC metric is reduced when that edge is added to the original graph. Finally it selects the top k edges that lead to the lowest score when added to the graph individually. Finding only the top k edges is not enough, because the probability of an edge to appear in the graph differs. So they introduced a probabilistic model that takes into account the acceptance probability in order to consider edges that minimize the RWC score in expectation. They further improved their algorithm to consider this probability.

A similar approach introduced by Parotsidis et al. in [13]. Instead of suggesting links that are most likely to be adopted, they aimed at solving the problem of finding a list of recommended links for a specific user that if adopted, they would lead to the best possible improvement of the user centrality. This problem is NP-hard, so they introduced a greedy approach with a constant approximation ratio. Except from greedy, they described three more heuristic approaches for this problem, to examine the trade-off between efficiency and effectiveness of their methods.

Matakos et al. [14] in 2017 studied the problem of measuring polarization of opinions in social networks. They introduced the polarization index which quantifies polarization in networks, taking into consideration the opinions of individuals. Then, they focused on reducing opinion polarization by convincing people to adopt a more neutral stand on controversial themes.

In [15] and [16] Garimella et. al proposed an algorithmic solution to the problem of reducing polarization. Their goal was to reduce user level polarity by exposing users to opposing points of view. In order to create the recommendation lists for users they took into account several factors such as the reduction of user polarization score (they consider items shared by high degree nodes), the exclusivity of item on either side, the acceptance probability, the topic diversity, and the popularity of item on either side.

Garimella et. al [6] studied the degree to which echo chambers exist in political discussions on Twitter and how they are structured. They considered the interplay of two components: The opinion of a user (echo) and the place that allows its exposure (chamber). They introduced a production and consumption measure in order to distinguish the different roles (partisans, bipartisans, gatekeepers) of users in the network.

They showed that partisan users (that produce one sided content) enjoy a higher appreciation than bipartisan users (that produce diverse content) as measured by network and content features. In addition, they explored features of gatekeepers users (that consume diverse content, but produce one sided content) and examined how they differ from the other users. Finally they built a classification model that that predicts if a user is a partisan or a gatekeeper.

Garrido et al. in [17] focused on connecting people of opposing views. They took into advantage the partial homophily of users to suggest agreeable content from individuals that share opposite views on sensitive issues. They used an organic design to create a data portrait for each user which displays their preferences on topics and tweets. Based on these portraits they recommended tweets between users that are dissimilar. They found that these recommendations result in negative emotional effects. They evaluated their methodology by performing a case study.

2.2 Diversity

Cai-Nicolas Ziegler et al. in 2005 [18] designed topic diversification, an algorithmic framework to balance and diversify user's recommendation list. Their goal was to improve real user satisfaction by creating recommendation lists that are not only accurate, but cover all user's extent of interest in specific topics.

They introduced a new similarity metric between recommendation lists, the intra-list similarity, to estimate the topical diversity of recommendation lists and prove the efficiency of topic diversification method, which is designed in order to reduce intra-list similarity. This means, that if intra-list similarity is low, then the recommendation lists are more diverse.

They showed that precision and recall metrics used for both item-based and user-based collaborative filtering techniques may be accurate, but they do not involve other factors such as the users' perceived list diversity.

Drosou et al. [19] in 2010 classified the different definitions of the diversification problem and comparately studied algorithms and metrics for result diversification. The same authors [20] proposed a new definition of diversity, based on similarity and coverage, called DisC diversity.

In [21] Van Dang and W. Bruce Croft presented a technique which results to a list with more diverse search results with respect to the topics that a query covers. They claimed that the number of documents that are retrieved on each topic should be proportional to the topic popularity.

They proposed an effectiveness measure for proportionality (Cumulative Proportionality (CPR)), which is based on metrics that are used for evaluating outcomes of elections. They also proposed a framework to optimize the proportionality in search results. This technique determines in which position a topic should be placed to maintain the best overall proportionality. Then, on this topic, the best document is placed in the corresponding position.

Garimella et al. [22] addressed the problem of finding a set of users in social networks that belong to opposing viewpoints, so that the overall information exposure for the opposing sides is balaced. They characterized the hardness of the problem and they proposed approximation algorithms.

Eytan Bakshy et al. in [3] performed a study on Facebook. They measured the ideological homophily in friend networks and studied how much heterogeneous friends could potentially expose individuals to cross cutting content. They also measured the extent to which individuals encounter more or less diverse content in Facebook's algorithmically ranked news feed. Finally they studied user's choices, that is the click through to ideologically discordant content. For this study they constructed a deidentified dataset that contains users who report their ideological affiliation and stories that are classified as hard and soft, depending on their content.

In [23], Munson et al. studied the cost of presenting (with sorting and highlighting techniques) content to users that is opposite to their political beliefs. They focused on the satisfaction of individuals when they are exposed to content agreeable or not. They distinguished users into two categories: the diversity seeking users, that they are open to consume contrarian opinions, and the challenge-averse users whose satisfaction is decreased if they are exposed to diverse opinions.

Munson et al. [24] created a browser extension that informs users about the political lean of their weekly and all time reading behaviors. This feedback leads to a move towards balanced exposure. Users visit more often ideologically opposing or centrist sites.

2.3 Comparison to the state of the art

Although our work relates to a point to the studies mentioned above, there exist some key differences. First of all, for the problems of polarization, homophily and diversity we take advantage both the structure of the network and the content of the users. While these works focus on minimizing other measures of polarization, our goal is to maximize the diversity of users and networks. So, we introduce a novel way to define and quantify diversity and we propose two algorithms to maximize it. Furthermore, to the best of our knowledge, we are the first to study the occurence of echo chambers considering diversity. Finally, we focus on the correlation between diversity and polarity of both user and network level.

Chapter 3

Quantifying Controversy

In the first part of this thesis, we deal with identifying which themes are the most controversial and we measure the amount of controversy for each one of them. So, we adopt an existing model that was introduced by [12] that takes into account the structure of the networks and produces a score of how controversial a theme is. In the following subsections we provide information about the data we collected for our study, as well as we present each one of the three stages of the pipeline. We also compute user level polarity using the definition introduced in [11].

3.1 Data

The data that we use are collected from Twitter using the Search API. We obtained all Clinton (almost 8 million) and Trump (almost 12 million) followers. Among them, we chose a subset (1.6 million users for Clinton and 2.4 million users for Trump) maintaining their ratio the same, and retrieved all of their tweets that were published between October 1st and November 8th (USA presidential elections). Over that period, we collected about 300.000.000 tweets. We must mention that there is no overlap between Clinton and Trump users, that is a user follows either Clinton or Trump.

From all the tweets collected, we obtained the most frequent hashtags (themes). Among them, we chose themes that we thought it would be interesting to explore, such as trending themes from the political domain dealing with the presidential elections in USA(#PresidentialDebate, #Elections2016, #USElection2016, #Politics) and themes from other domains, such as music, tvseries etc (#ClimateChange, #Supernatural, #NowPlaying, #WorldSeries). Additionally, we expected some of them to be more controversial, such as #PresidentialDebate and some other less, like #NowPlaying. Finally we combined two different datasets, by concatenating their corresponding networks. We combined two political controversial themes (PresidentialDebate, Elections2016) and two themes from different domains (PresidentialDebate, Supernatural).

3.2 First Stage: Build Endorsement Graphs

Having collected all the tweets that are associated with the themes we choose to study and their additional information (Name, ScreenName, UserID, FollowersCount, FriendsCount, Location, Description, CreatedAt, StatusID, Language, Place, RetweetCount, FavoriteCount, Text), we create the endorsement graphs between users, which is the first stage of the pipeline. Although we study controversy only in the retweet networks, we provide the definitions for both retweet and follow graphs.

Retweet Graph $\mathrm{G}_{RT} = (V, E)$: V is the set of users who generated content relevant to a theme (hashtag). E is the set of *undirected* edges, that show endorsement between users. Typically, retweets are used as endorsements, so an edge between two users, means that there is at least one retweet between them.

Follow Graph $\mathrm{G}_{F} = (V, E)$: V is the set of users who generated content relevant to a theme (hashtag). E is the set of *directed* edges that show connection between users. There is an edge between users u and v if u follows v or vice-versa.

Tables 3.1 and 3.2 provide statistics about the retweet and follow networks we constructed respectively, such as the number of the nodes (users) and the number of edges (retweet or follow interactions). As we expect, retweet networks are smaller in size and less sparse than follow networks. We work on the biggest connected components (BCC) of retweet networks.

Hashtag	Nodes	Edges	BCC Nodes	BCC Edges
ClimateChange	20763	30629	16072	22237
Elections2016	28328	39536	20648	24384
NowPlaying	13729	21438	5919	8183
Politics	45459	148031	41341	64993
PresidentialDebate	26438	30069	19495	25489
PresDeb-Elections2016	50539	59641	37984	51223
PresDeb-Supernatural	31206	43825	26989	38452
Supernatural	8056	29643	7607	12776
USElection2016	18477	20334	13435	15650
WorldSeries	89281	198902	85387	154588

Table 3.1: Statistics of the 10 retweet networks.

Hashtag	Nodes	Edges	BCC Nodes	BCC Edges
ClimateChange	20763	305223	20702	305198
Elections2016	28328	239125	27817	238853
NowPlaying	13729	26099	13431	25955
Politics	45459	1180931	45418	1180908
PresidentialDebate	26438	353506	25894	353118
PresDeb-Elections2016	50539	555603	50149	555185
PresDeb-Supernatural	31206	347859	30974	347684
Supernatural	8056	50097	7992	50065
USElection2016	18477	45563	18103	45355
WorldSeries	89281	683748	88893	683513

Table 3.2: Statistics of the 10 follow networks.

3.3 Second Stage: Graph Partitioning

We deal with the second stage of the pipeline that has to do with the partitioning of the graph. We use two different approaches: (i) An approach that partitions the graph based on structural properties (Metis Partitioning) [25] and (ii) An approach that partitions the graph based on the labels that we have for users (Label Partitioning).

3.3.1 Metis Partitioning

For each network we create two communities using the Metis algorithm (Metis Partitioning). Since we have the ground truth about which person (Clinton or Trump) each user supports, for each of the two communities we count the number of Clinton and Trump users that belong to each community. There also exist some additional users that we do not have information about their support. In Tables 3.3-3.12 we provide information about the labels of users in the partitions created by Metis.

We observe label homophily in the datasets that we expect to be controversial (Tables 3.7, 3.4, 3.11, 3.8). Users with the same label tend to interact more with each other. The one community has a majority of Trump's followers and the other community has a majority of Clinton's supporters. However, there also exists a small percentage of users that support the opposing side, or support neither of the two sides.

We can not draw any conclusion in the datasets that we expect to be less controversial (Tables 3.3, 3.10, 3.6, 3.5, 3.12). This is expected, because datasets like Supernatural, ClimateChange, NowPlaying and WorldSeries are completely independent of the political label a user has.

ClimateChange	Trump	Clinton	No Label	Total
Community1	870 (11%)	4736 (59%)	2429 (30%)	8035
Community2	2352 (29%)	3963 (49%)	1722 (22%)	8037

Table 3.3: Statistics of the 2 communities by METIS for ClimateChange dataset

Elections2016	Trump	Clinton	No Label	Total
Community1	1715 (16%)	7168 (68%)	1647 (16%)	10530
Community2	7863 (78%)	1158 (11%)	1097 (11%)	10118

Table 3.4: Statistics of the 2 communities by METIS for Elections2016 dataset

NowPlaying	Trump	Clinton	No Label	Total
Community1	841 (28%)	966 (33%)	1162 (39%)	2969
Community2	375 (13%)	1723 (58%)	852 (29%)	2950

Table 3.5: Statistics of the 2 communities by METIS for NowPlaying dataset

Politics	Trump	Clinton	No Label	Total
Community1	227 (1%)	1025 (5%)	19267 (94%)	20519
Community2	5176 (25%)	1393 (7%)	14253 (68%)	20822

Table 3.6: Statistics of the 2 communities made by METIS for Politics dataset

PresidentialDebate	Trump	Clinton	No Label	Total
Community1	7242 (76%)	937 (10%)	1310 (14%)	9489
Community2	1000 (10%)	7547 (75%)	1459 (15%)	10006

Table 3.7: Statistics of the 2 communities by METIS for PresidentialDebate dataset

PresDeb-Elections2016	Trump	Clinton	No Label	Total
Community1	13378 (72%)	2306 (13%)	2793 (15%)	18477
Community2	2355 (12%)	13719 (70%)	3433 (18%)	19507

Table 3.8: Statistics of the 2 communities by METIS for PresidentialDebate-Elections2016 dataset

PresDeb-Supernatural	Trump	Clinton	No Label	Total
Community1	7492 (56%)	4084 (31%)	1718 (13%)	13294
Community2	1597 (12%)	5824 (43%)	6274 (45%)	13695

Table 3.9: Statistics of the 2 communities made by METIS for PresidentialDebate-Supernatural dataset

Supernatural	Trump	Clinton	No Label	Total
Community1	128 (3%)	103 (3%)	3569 (94%)	3800
Community2	754 (20%)	1339 (35%)	1714 (45%)	3807

Table 3.10: Statistics of the 2 communities made by METIS for Supernatural dataset

USElection2016	Trump	Clinton	No Label	Total
Community1	4256 (63%)	1689 (25%)	803 (12%)	6748
Community2	1421 (21%)	4392 (66%)	874 (13%)	6687

Table 3.11: Statistics of the 2 communities made by METIS for USElection2016 dataset

WorldSeries	Trump	Clinton	No Label	Total
Community1	17914 (43%)	15735 (34%)	7924 (23%)	41573
Community2	22984 (50%)	12548 (27%)	8282 (23%)	43814

Table 3.12: Statistics of the 2 communities by METIS for WorldSeries dataset

3.3.2 Label Partitioning

Subsequently, based on the information about the support of the majority of the users, we create communities based on the labels of each user (Label Partitioning). The one community contains only Clinton's followers and the other contains only Trump's followers. We do this, because we believe that for the political datasets, users are divided based on their political affiliation. The users that we have no information about their support are omitted. More details about this partitioning method are depicted in Tables A.1 - A.10 in Appendix A.1.

3.4 Third Stage: Random Walk Results

Finally, we proceed to the third and final stage of the pipeline. We measure the controversy using the random walk metric that captures the intuition of how likely a random user is to be exposed to content of the opposing side. Controversy score indicates the difference of the probability of two random walks to begin and end in the same partition and the probability to begin and end in different partitions of a graph (we should denote that a random walk ends when it reaches a high degree node). The closer to 1 the score is, the higher the difference of these two probabilities is (probability to end in the same partition is much bigger than to end in different partitions). Similarly, the closer to 0 the score is, the lower the difference of these two probabilities is (probability to end in the same partition is almost equal to the probability to end in crossing sides).

In Table 3.13 we provide the random walk scores in descending order for each of the two partitioning methods. We observe that the partitioning made by Metis is more effective than labels partitioning, as it produces a higher score. This is due to the fact that, although Metis applies a better clustering to the graph, the two of the clusters consist of some users that should be in the other cluster concerning their labels. If we wanted a perfect clustering, the Metis partitioning technique should result to users with only one label per partition. For example, in Tables 3.7, 3.4 and 3.11 we can see that there is a minority of 'crossing' users that belong to wrong cluster concerning their label.

Using Metis, five of the datasets (PresidentialDebate, Elections2016, Supernatural, USElection2016 and NowPlaying) have a higher random walk score, which means that these themes are considered as more controversial than the others. Although in most cases the two partitioning methods result in similar scores (both methods produce high scores for the political datasets), there are some themes like Supernatural, NowPlaying and WorldSeries which produce fairly low controversy scores with Label partitioning method.

This is due to the fact that in some cases the labels (C or T) of the users may be independent of the theme that is discussed. For example, Supernatural may be a controversial theme, but is completely independent of which political label a user has. This means that if we partition the graph using these labels, the probability to end in different partitions is almost equal as this to end in the same partition. In addition, we observe that PresidentialDebate-Elections2016 dataset that created by the combination of the two controversy themes, leads to a controversial network.

Finally, two very interesting themes is USElection2016 and NowPlaying. While USElection using Metis partitioning, has a fairly big controversy score, we observe that using Label partitioning the controversy score is low. This opposes to the statistics in Tables 3.11 and A.9, as we can see that the number of the supporters of Clinton and Trump in every community is almost equal. Additionaly, we observe that although we expected NowPlaying to be a non controversial theme, it has a very high controversy score. This is due to the fact that users tend to retweet songs that they like and ignore songs from other genres of music.

Dataset	Metis RW Score	Labels RW Score
NowPlaying	0.89	0.12
PresidentialDebate	0.65	0.43
Elections2016	0.64	0.44
PresDeb-Supernatural	0.63	0.42
PresDeb-Elections2016	0.60	0.40
Supernatural	0.59	0.04
USElection2016	0.56	0.27
ClimateChange	0.44	0.30
Politics	0.43	0.39
WorldSeries	0.38	0.08

Table 3.13: Polarization score using Random Walk metric for communities made by Metis and Labels

3.5 User Level Polarity

The above three-stage pipeline leads to a score that indicates how controversial a theme is. Here, we proceed to a user level polarity computation based again on random walks.

Let X^+ and Y^+ be the two sets of k highest degree vertices on each side of the controversy. Intuitively, a vertex is assigned a score of higher absolute value (closer to -1 and 1), if, compared to other vertices in the graph, it takes a very different time to reach a high degree vertex on either side.

Specifically, for each vertex $u \in V$, we consider a random walk that starts at u and estimate the expected number of steps, l_u^X before the random walk reaches any high degree vertex in X^+. Considering the distribution of values of l_u^X across all vertices $u \in V$, $p^X(u)$ is defined as the fraction of vertices $v \in V$ with $l_v^X < l_u^X$. $p^Y(u)$ is defined similarly. The polarity score of each user is then defined as:

$$R_u = p^X(u) - p^Y(u) \in [-1, 1]$$

High absolute values (close to -1 or 1) indicate that the user clearly belongs to one side of the controversy, while central values (close to 0) indicate that the user is almost in the middle of the two sides.

In Appendix A.2 we provide the distribution of the polarity values (Figure A.1). We plot polarity (X axis, 50 equal bins) in terms of their frequency (Y axis) on each dataset. We observe that in all datasets we have many polarity values that are close to -1 or 1. This is mostly observed in the most controversial datasets. However, in datasets like ClimateChange, Politics and WorldSeries, which are less controversial, we notice that there also exist many users that have polarity values close to 0.

3.6 Labeling Communities

The abovementioned methods take into consideration only the structure of the networks to produce network and users level polarity scores. We want to exploit the tweets of the users to see the content that they deal with. We perform a text analysis in the datasets for the tweets of every partition, for the two partitioning methods.

We labeled the partitions with the most frequent hashtags and its most important words using the tf-idf measure. However, the language that is used in tweets makes text mining not an easy task. As we can see in Tables 3.14 and 3.15, where we depict the top 10 keywords and hashtags for the dataset USElection2016, there are many common words (e.g. breathe, asleep, etc.) and hashtags (e.g. electionnight, electionday, etc.) between the communities made by Metis and Label partitioning. As we cannot draw any conclusions about the content of each community, we omit the results for the rest of the datasets.

Met1 #	Met2 #	Met1 words	Met2 words
electionnight	trump	breathe	successful
trump	electionnight	asleep	hidden
trumpwins	putin	monkeys	active
donaldtrump	breaking	ears	loose
ripamerica	electionday	americaa	overwhelming
presidenttrump	donaldtrump	deciding	code
electionday	trumppresident	negative	headache
electionsnight	presidenttrump	sandmonkey	label
breaking	brexit	trolling	neck
usadecides	hillary	bothered	uh

Table 3.14: Most frequent hashtags and words in every community made by Metis for USElection2016

Lab1 #	Lab2 #	Lab1 words	Lab2 words
trump	electionight	asleep	breathe
electionight	ripamerica	greetings	bleach
putin	electionday	loose	hidden
breaking	election2016	america	active
electionday	usadecides	monkeys	spare
donaldtrump	merkel	christmas	ears
trumpwins	breaking	headache	deciding
trumppresident	hillaryclinton	successful	trolling
brexit	myvote2016	fuckedup	disturbed

Table 3.15: Most frequent hashtags and words in every community made by Labels for USElection2016

Chapter 4

Existence of Homophily

In this Chapter we describe the text mining technique we used for dealing with the content of the users. Subsequently, we use some network and content properties to study the existence of homophily in the two sides of controversy. We expect that the content of the users that belong to the same communities is more similar than the content of the users that belong to different communities . We define homophily in a user level and we focus on the correlation between homophily and polarity values.

4.1 Topic Models

We use topic modeling techniques to find out the different topics (aspects) that a user deals with in his tweets, which are related to a specific hashtag. We use this approach because we believe that will help us understand and organize our large collections of tweets better than tfidf or most frequent hashtags methods. In general, topic models help us find a group of words from a collection of documents that best represents the information of the collection.

We use the LDA topic modeling technique. In LDA, documents are represented as a mixture of topics that spit out words with certain probabilities. The model proposes that each word in the document is attributable to one of the topics of the document. LDA discovers the different topics that the documents represent and how much of each topic is present in a document. Each document is characterized by a multinomial distribution over topics and each topic is in turn characterized by a multinomial distribution over words.

More specifically, we use the MALLET topic model package [26]. It takes as input the number of topics that we want it to produce and the set of tweets of each user (each user's tweets are considered as one document). It performs a preprocessing of the tweets and trains the topics to find the top most important in the whole corpus of documents. As an output, it produces a vector for each user, that provides information about the proportion of each topic that he covers in his tweets.

More formally: Each user u in the graph (G_{RT} or G_F) is represented as a vector $T(u) = \{f_1, f_2, ..., f_k\}$, where k is the different number of topics that exist in the tweets related to a hashtag, and f_i for $i \in [1, k]$, is the proportion of each of these topics that the user covers in his tweets.

4.2 Homophily in the Communities

In Section 3.3.1 we had a first contact with homophily that exists in the partitions made by Metis, by considering the labels of users. We observed that most users with the same label belong to the same partition. This only applies to political datasets, because we only have information about the political labels. We want to study homophily in themes of any domain based on users' vectors $T(u)$ and not on their labels.

So, in this section we measure homophily between users that belong to the same or different communities (produced by Metis partitioning method), by focusing on the content that users have shared in their tweets. We expect that in more controversial themes, there exists homophily in the communities, that is the vectors of the users that belong to the same communities are more similar than the vectors of the users that belong to different communities. Furthermore, we do not expect this to be valid for less controversial topics.

Based on the information about the vectors $T(u)$ for each user u, we calculate the similarity between them, using the cosine similarity metric between their vectors. We define this as

$$SIM(u, v) = cos(T(u), T(v))$$

We must note that, to discover the topics, we take into consideration only the users' tweets that contain the specific hashtag, not all their tweets. Due to the fact that the number of topics must be given as input to the model, we experiment with different numbers. When we give as input 20 topics, we observe that some of the topics are repeated and when we give as input 5, the topics are not so distinguishable. So we ended up giving as input 10 topics.

To measure homophily we choose two approaches: **(i)** We compute the average similarity between every pair of users that belong to the same community and the average similarity between all pairs of users that belong to different communities. **(ii)** We compute the average similarity between pairs of directly connected users (neighborhood of users) that belong to the same community and between users that belong to different communities. The results of these two aproaches are depicted in figures 4.1 and 4.2 respectively.

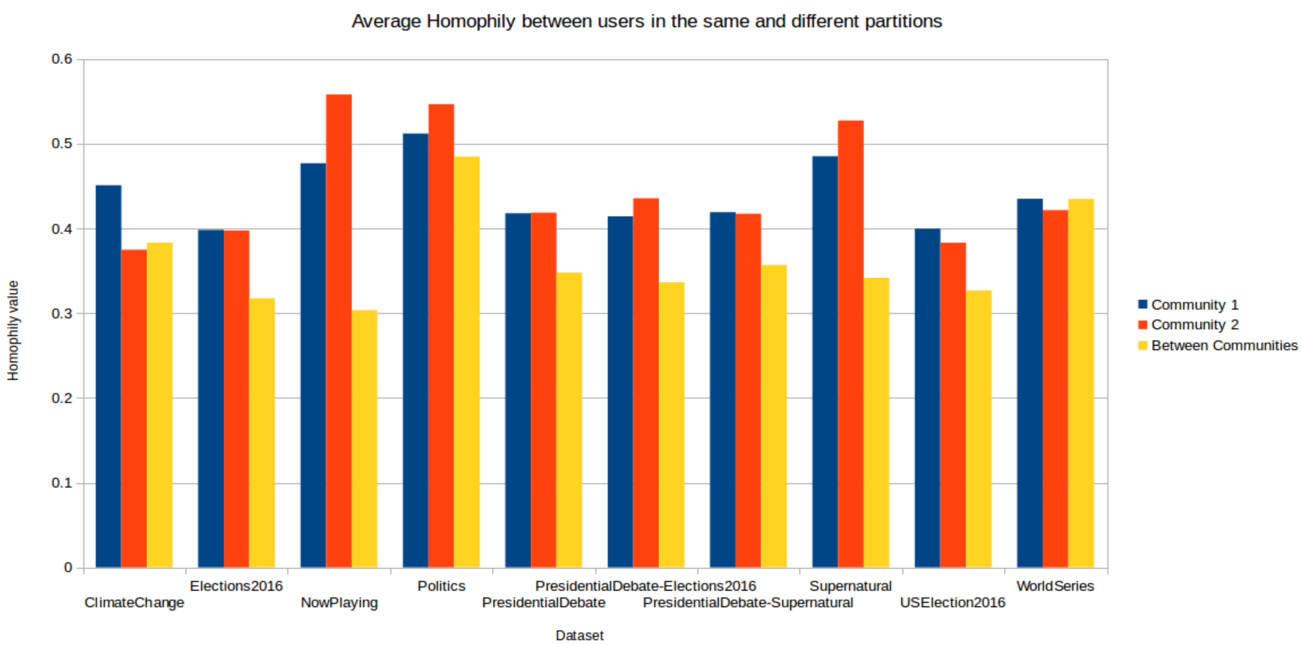

Figure 4.1: Average homophily between users in the same and different partitions

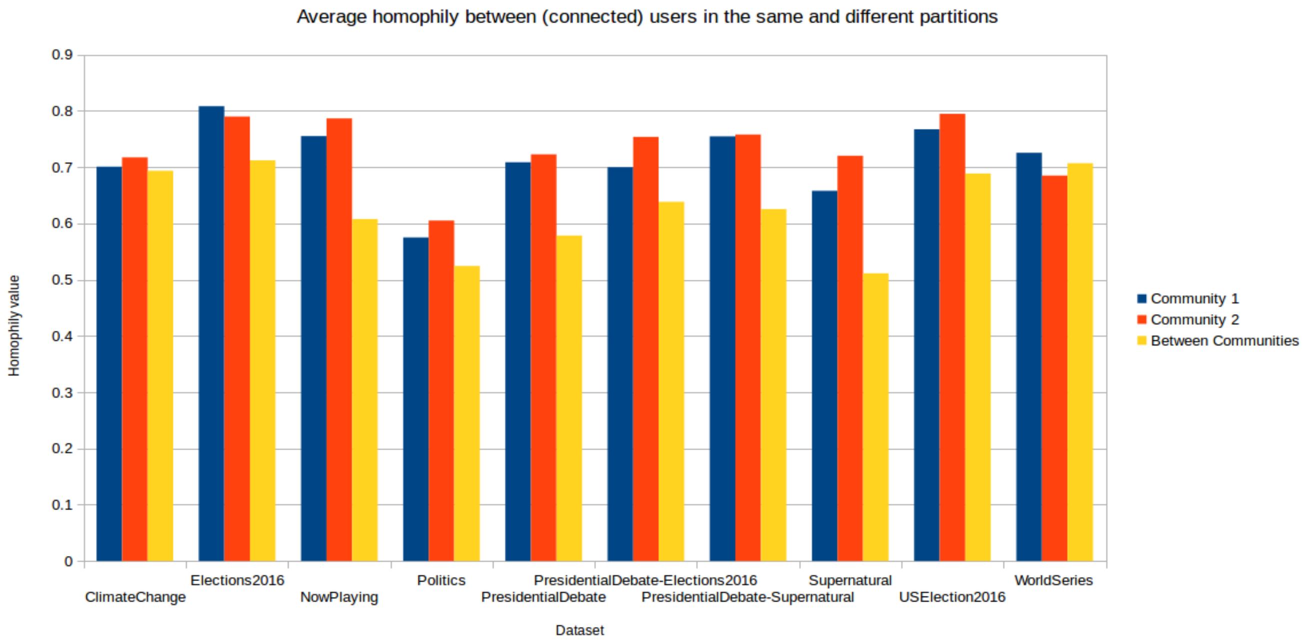

Figure 4.2: Average homophily between connected users in the same and different partitions

We notice that, in both of the approaches, for the datasets that have a bigger controversial score (e.g NowPlaying, Elections2016, Supernatural, etc.), the users that belong to the same community are more similar than the users that belong to different communities. We observe that there is a difference between the homophily values of community1 and community2 and the homophily values between the two communities.

This difference does not apply in the case of less controversial topics, such as ClimateChange and WorldSeries. We can see that the homophily value between users that belong to different communities is very close to this of the homophily between users in the same communities.

Another observation is that when we use the second approach, the homophily values are bigger than the corresponding values of the first approach. This is expected, because we deal with a retweet graph and among directly connected users there exist many common tweets (retweets).

In Appendix B.1 we provide the exact homophily values for all datasets for both of the approaches.

So we can say that the assumption we made, about the similarity of users in controversial and non controversial themes, applies here. The results show that in controversial themes, intra-cluster users are more similar than inter-cluster users. So, the similarity between users is correlated to a point with how controversial a theme is.

4.3 Homophily-Polarity Correlation

In this section we perform some experiments to check if there exists a correlation between user level polarity and homophily values. We define a user level homophily metric defined as the difference of the similarity of a user u with users from the same partition and his similarity with users from the opposing side. Formally:

$$h(u) = avg(Sim(u_X)) - avg(Sim(u_Y))$$

X is the partition that user u belongs, Y is the opposing partition and u_X, u_Y are the users from the corresponding partitions. This definition may be also applied for the second homophily approach, that considers only the directly connected users of user u. In this case user level homophily is denoted as $h_N(u)$ and u_X and u_Y are the neighbors of u that belong to partition X and Y respectively. It is clear that $h(u)$ and $h_N(u)$ take values in the interval $[-1, 1]$.

More specifically, we examine the polarity and homophily correlation. We expect that users that have a bigger homophily value will be more polarized and users that have a smaller homophily value will belong almost to the middle of the two sides and will have low polarity values.

By plotting homophily and polarity values, we can not draw any conclusion about their correlation, as the plots that are produced are quite noisy. So, we use Pearson correlation between homophily and absolute polarity values. The scores produced by this technique are depicted in Table 4.1.

As we can see in Table 4.1, in all of the datasets, except WorldSeries, there exists a positive correlation between homophily and polarity values (especially in datasets like NowPlaying, Politics and PresidentialDebate-Elections2016, we observe a fairly high correlation score). This is something we expected. Users that belong to extreme ends of the partitions have a higher homophily value, that is they are more similar

Dataset	Pearson ($h(u)$, $p(u)$)	Pearson ($h_N(u)$, $p(u)$)
ClimateChange	0.03569848	0.15191591
Elections2016	0.22453888	0.2184552
NowPlaying	0.42520734	0.32789913
Politics	0.40169776	0.12532427
PresidentialDebate	0.09495357	0.09632558
PresDeb-Elections2016	0.33530687	0.19794588
PresDeb-Supernatural	0.28182292	0.17580384
Supernatural	0.17235736	0.07022321
USElection2016	0.25204647	0.07112477
WorldSeries	-0.06392817	0.05662398

Table 4.1: Pearson scores for homophily-polarity correlation

with users that belong to the same partition and are less similar with users from the other side. Subsequently, users that belong close to the middle of the two sides have homophily score close to 0 (or slightly negative).

Furthermore, we observe that using the second approach (third column), we get lower scores for the most of the datasets. This is expected, because users that are directly connected have more common tweets. We also notice that all the datasets have a positive score.

Chapter 5

Diversity

In this chapter we define user and network level diversity. Based on these definitions, we examine the occurence of echo chambers in all of the networks. Moreover we study the correlation between diversity and polarity values (in both user and network levels). Furthermore, we propose two algorithms that recommend content to a user, to improve his diversity. We demonstrate that these two methods perform well in practice, not only in terms of diversity of users, but also in terms of their polarity.

5.1 Diversity Definitions

Before we proceed to the definitions of both user and network level diversity, we need to consider the probability of a user u_i to be exposed to content of a user u_j. In this way, we focus on the expected content that a user receives from his wider neighborhood. So, we denote as P_{ij}, the probability of u_i to reach content of u_j. To estimate these probabilities we perform random walks with restarts.

In order to define diversity of a user u_i we take into account both the users' vectors $T(u)$ and the aforementioned probabilities. We want the distribution of the expected content that u_i receives from his wide neighborhood to be close to uniform. This means that the frequency of every topic from the content he receives is close to $1/k$, where k is the number of topics. So, diversity $D(u_i)$ of user u_i is defined as follows:

$$D(u_i) = Entropy(\sum_{j=1, j \neq i}^{|V|} P_{ij} T(u_j))$$

To obtain values between 0 and 1, we normalize the score by dividing with maximum diversity ($log_2 k$). Values close to 1 indicate that a user is very diverse (deals with all the topics with almost the same frequency) and values close to 0 indicate that a user is not diverse at all. This definition can be applied either on retweet networks or follow networks. We also provide a definition for a network level diversity: We say that diversity of a network is defined as the average diversity of users:

$$D_N = Avg_{u \in V} D(u)$$

We can limit this definition to measure diversity considering only the (directly connected) neighbors of a user. Then we will have the following perspectives:

(i) We assume that a user u is diverse, if the content that he receives from his direct (retweet) neighborhood is diverse. This means that he gets access to all points of view about a theme, from the people that has *endorsed*. We define $T_R(u) = Avg_{w \in N(u)}(T(w))$, that is the average (aggregation) of the vectors of neighbors of u in G_{RT}. So we define diversity $D_R(u)$ of a user u as the entropy of $T_R(u)$ vector: $D_R(u) = Entropy(T_R(u))$. We normalize diversity by dividing with the maximum diversity score ($log_2 k$) so as to obtain values between 0 and 1. In this case, network diversity is denoted as D_{NR}.

(ii) We assume that a user u is diverse, if the content that he recieves from his direct (follow) neighborhood is diverse. This means that he gets access to all points of view about a theme, from the people that he *follows*. Similarly, we define $T_F(u) = Avg_{w \in N(u)}(T(w))$, that is the average (aggregation) of the vectors of neighbors of u in G_F. In this case, diversity $D_F(u)$ of a user u is defined as the entropy of $T_F(u)$ vector: $D_F(u) = Entropy(T_F(u))$. Moreover, network diversity is expressed as D_{NF}.

(iii) There is another point of view, where we consider only the content that user u has in his retweets and not the content that he recieves from his (wider or direct) neighborhood. Similarly, we define diversity $D_P(u)$ as the entropy of $T(u)$ vector. So $D_P(u) = Entropy(T(u))$. Again, we normalize diversity to $[0, 1]$. We express network diversity as D_{NP}.

We generally prefer the first two approaches, because they take into account not only the content that users produce, but also the structure of the network (content of the neighborhood of a user). In addition, we want our users to have a diverse neighborhood that they will have access to all points of view about a theme and then decide which tweets fit to his beliefs and retweet them. So actually we are more interested in user's vector $T_R(u)$ and $T_F(u)$ than in user's vector $T(u)$.

5.2 Echo Chambers

Based on the aforementioned diversity perspectives, we focus on the existence of echo chambers in the networks by considering diversity of users. We expect that echo chambers will occur in all of the themes we study, especially in the controversial ones. There is no formal definition for echo chambers in the literature so we propose a definition that exploits diversity of users.

We say that echo chambers exist in a network if: (i) the average diversity of users is low and (ii) the content of users is similar to the content of their neighborhood. The first condition means that a user's neighborhood (chamber) does not provide him all the points of view about a theme and he is not exposed to different ideas and viewpoints. The second condition is for ensuring that the opinions of a user (echo) is in line with the opinions of his neighborhood.

To ensure that diversity of users is low, we compare the average diversity score (in both follow and retweet networks) of users, with their average diversity score if they had random neighbors. For each one of the themes we plot the actual average diversity of users and the random average diversity of users, in terms of user's degree. In all themes we observe that the individuals that users choose to connect, provide them less information, than the information that random users would have provided them. So we can say that the first condition fo the existence of echo chambers is true. For every dataset, we depict their plots in Appendix C.1.1 (follow) and Appendix C.1.2 (Retweet). Furthermore we observe from these plots, that the actual diversity that users have, tend to increase, as the number of neighbors increase. This means that having many neighbors, leads to exposure to more points of view.

Similarly, in order to see if the second condition is true, we compare the average similarity between the content that users produce and the content that they receive from the neighbors, with the average similarity between the content that users produce and the content that they would have received from random nodes. In Appendix C.1.1 and Appendix C.1.2 we depict these similarities for all datasets. We observe that users tend to interact more with users that are more similar to them. So, the second condition is also true.

We also measure the extent that echo chambers exist in the networks. We produce a score of how echo chambered a theme is, that is the proportion of users from the network that both the conditions are true. These scores are depicted in Table 5.1.

Dataset	Echo Chambers Score (G_F)	Echo Chambers Score (G_{RT})
ClimateChange	76.5%	82.0%
Elections2016	90.7%	83.4%
NowPlaying	85.4%	74.2%
Politics	83.5%	79.8%
PresidentialDebate	89.6%	84.7%
PresDeb-Elections2016	88.6%	88.3%
PresDeb-Supernatural	88.3%	75.7%
Supernatural	75.9%	84.2%
USElection2016	87.4%	79.3%
WorldSeries	80.3%	82.2%

Table 5.1: Echo Chamber scores for Retweet and Follow networks

5.3 Diversity-Polarity Correlation

In this section we focus on the correlation between the polarity and diversity values of users for the ten different datasets. In 5.3.2 we also measure the correlation of polarity and diversity in a global level.

5.3.1 User Level Correlation

We expect that users who are extremely polarized (values close to 1 and -1) will be less diverse than users who are not polarized (values close to 0).

By plotting user polarity in terms of user diversity ($D_F(u)$ or $D_R(u)$) we saw that it is not easy to conclude if there exists a correlation between these two, because the resulted plots where quite noisy. On the other hand in some datasets like Climate-Change, Elections2016 and PresidentialDebate-Elections2016 we observed that there exists an amount of very polarized users that have lower diversity than the other users.

To have a clear indication on how correlated polarity and diversity are, we use pearson correlation on diversity values and absolute polarity values. The results are shown in Table 5.2. We can see that in all datasets (especially in PresidentialDebate-Elections2016 and Elections2016 where we observe fairly high values) there exists a negative correlation between user polarity and user diversity. This means that users that have higher polarity values are less diverse than users that have lower polarity values.

In Table 5.2 we also observe the pearson correlation values if we consider user's vector $T_P(u)$ (not neighbors' vectors) for the definition of diversity. In this case we see that the correlation is not so obvious. We have smaller correlation values and in datasets like NowPlaying and WorldSeries we observe positive values. The same is also observed if we consider diversity values $D_F(u)$ in the follow networks.

The number of stars next to the numbers in Table 5.2 shows the importance of each correlation (p-value = 0.05). One star (*) indicates a non significant correlation and two stars (**) indicate a significant correlation.

Dataset	Pearson ($D_R(u)$)	Pearson ($D_P(u)$)	Pearson ($D_F(u)$)
ClimateChange	-0.28288167 (**)	-0.18307296 (**)	-0.08197743 (**)
Elections2016	-0.45254164 (**)	-0.25277629 (**)	-0.16158856 (**)
NowPlaying	-0.05830189 (**)	0.01560889 (*)	-0.12782654 (**)
Politics	-0.20885943 (*)	-0.10502326 (**)	0.03864223 (**)
PresidentialDebate	-0.29607188 (**)	-0.22095027 (**)	-0.099514605 (**)
PresDeb-Elections2016	-0.37005492 (*)	-0.19613803 (**)	-0.14915447 (**)
PresDeb-Supernatural	-0.15735304 (**)	-0.16097161 (**)	-0.14417582 (**)
Supernatural	-0.17763765 (**)	-0.12623966 (**)	-0.04725573 (**)
USElection2016	-0.1127577 (**)	-0.18065104 (*)	-0.0905457 (*)
WorldSeries	-0.02123373 (**)	0.05507899 (*)	-0.09873747 (**)

Table 5.2: Pearson Correlation Between Polarity and Diversity for each of the datasets (user level)

5.3.2 Global Level Correlation

In Table 5.3 we depict the global diversity of network (average diversity of users) in terms of network polarity score for each one of the datasets. We observe that there exists a fairly strong correlation between these two. The higher the polarity score of a theme is, the lower the average diversity of the users of the network is. So we can say that low diversity values is an indication that a theme is controversial. We measure this correlation by applying Pearson correlation: $R(D_{NR}, Controversy\ score) = -0.8316$.

By using the other diversity definition (taking into account $T_P(u)$) we have a Pearson score: $R(D_{NP}, Controversy\ score) = -0.523$. This, also is a high score, but is lower that the first. Finally, using $D_F(u)$ (average diversity of the followers) we obtain score: $R(D_{NF}, Controversy\ score) = -0.6627$

Another important observation resulting from Table 5.3 is that users are on average more diverse in Follow networks than in Retweet networks. This is due to the fact that they have more neighbors and they are exposed to more points of view. Also, we expected that the diversity of content that users produce, D_{NP}, would be lower than D_{NR} and D_{NF}. This is not true, as we can see that although D_{NP} is lower than D_{NF}, it is higher than D_{NR}, meaning that users's content is more diverse than the content they recieve from the users they endorse.

Dataset	D_{NR}	D_{NP}	D_{NF}	Polarity
ClimateChange	0.6817	0.7327	0.8476	0.44
Elections2016	0.5902	0.7191	0.8519	0.64
NowPlaying	0.6068	0.716	0.7361	0.89
Politics	0.6954	0.7374	0.8338	0.43
PresidentialDebate	0.6098	0.7196	0.8128	0.65
PresDeb-Elections2016	0.6224	0.7178	0.8101	0.60
PresDeb-Supernatural	0.6112	0.7108	0.8046	0.63
Supernatural	0.6578	0.6821	0.7468	0.59
USElection2016	0.6494	0.7077	0.8269	0.56
WorldSeries	0.7237	0.7537	0.8288	0.38

Table 5.3: Polarity and Diversity values for each of the datasets (network level)

5.4 Improving Diversity

The problem we study in this section is to increase diversity of users, burst the filter bubble and eliminate echo chamber phenomenon. By doing this, we ensure that they will have access to all points of view and will be informed about opinions that oppose their beliefs. So our problem is to recommend to a user u, k edges so that his diversity value is maximized. In section 5.4.2 we propose two recommendation algorithms that aim to increase diversity of users and expected diversity of users by recommending them content of other users.

Problem 1. *Given a Graph $G = (V, E)$, a user u, a positive integer k and a candidate set $R = \{(e, p_e)\colon\ e \notin E,\ p_e \in [0, 1]\}$ of probabilistic edges incident to u recommended by a link prediction algorithm, select a set $S \subseteq R$ of size k to recommend to u, such that $D(u)$ is maximized.*

5.4.1 Link Prediction

The two recommendation algorithms that follow in Section 5.4.2 take into consideration users that are most likely to be endorsed by u. We create the set R that contains n of these users. We use several link prediction methods from the literature to find this set. We evaluate these link prediction methods to see which of them fits better in our endorsement graph. We create the ROC plot, that plots the true positive rate against the false positive rate for the different link prediction methods. The results of these link prediction methods are obtained using linkpred library [1]. As we can see by the ROC plot in Figure 5.1, the most effective method for our case is the 'CommonNeighbor' method, as it has the biggest area under its curve. So, we choose this method to find set R.

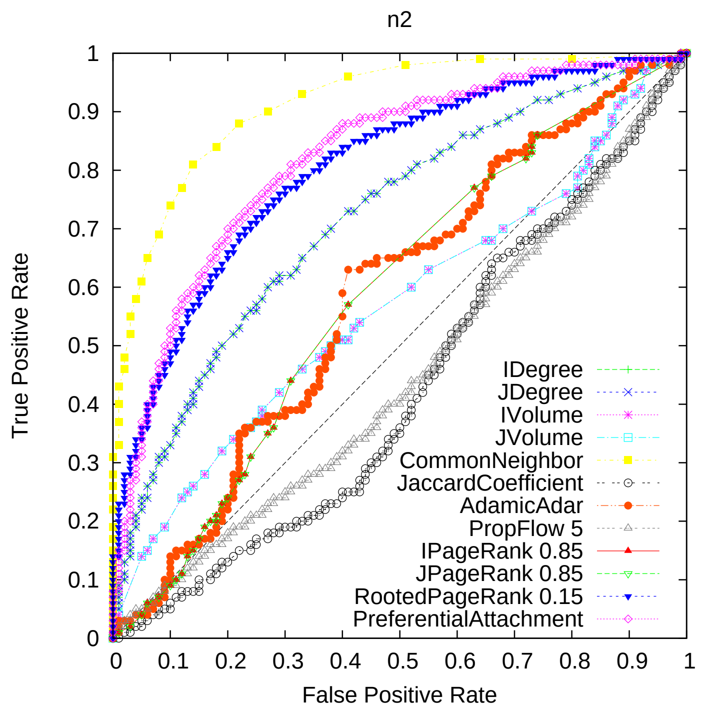

Figure 5.1: Roc Plot for 12 different link prediction methods

5.4.2 Algorithms

In this Section we provide two methods that aim to increase diversity $D_N(u)$ (or $D_F(u)$) of a user u. We propose two greedy algorithms that recommend k edges to a user u that increase the objective function $D_N(u)$ (or $D_F(u)$). The outline of the two algorithms is the following:

[1]https://github.com/rafguns/linkpred/

Algorithm 5.1 Greedy Diverse algorithm.

1: **Input**: User u, Set R of n users predicted for u.
2: **Output**: Recommendation list.
3: Initialize $out \leftarrow \emptyset$.
4: **for** $1 \leq i \leq k$ **do**
5:     **for** all users $v \in R$ **do**
6:         Compute $\delta D(u)_{u \to v}$, the increase of diversity of u if $u \to v$ is added
7:     **end for**
8:     $out = out \cup argmax_v(\delta D(u)_{u \to v})$
9:     $R = R / argmax_v(\delta D(u)_{u \to v})$
10: **end for**
11: **return** out

Algorithm 5.2 Greedy Diverse (expected increase) algorithm.

1: **Input**: User u, Set R of n users predicted for u.
2: **Output**: Recommendation list.
3: Initialize $out \leftarrow \emptyset$.
4: **for** $1 \leq i \leq k$ **do**
5:     **for** all users $v \in R$ **do**
6:         Compute $\delta D(u)_{u \to v}$, the **expected** increase of diversity of u if $u \to v$ is added
7:     **end for**
8:     $out = out \cup argmax_v(\delta D(u)_{u \to v})$
9:     $R = R / argmax_v(\delta D(u)_{u \to v})$
10: **end for**
11: **return** out

Both algorithms build the solution incrementally. They take as input the user u that we target, along with the list R that contains the predicted users for u.

Algorithm 5.1 does not consider the probability of an edge to be materialized. In each iteration it selects the node that if added, it will lead to maximum increase of user's diversity without taking into consideration how possible is the recommendation to be accepted by u. This is the critical difference between Algorithm 5.1 and Algorithm 5.2 and is shown in line 6.

Algorithm 5.2 takes into account the probability of an edge to be accepted by the user. It selects the user that leads to the best expected decrease of diversity. We should mention that there is a natural tradeoff between accuracy (how likely an edge is to appear) and the utility of recommendations (improvement of diversity). Links that connect users that are close to each other and have many common neighbors are most possible to materialize than links that connect far-away nodes but increase diversity. Having this in mind, we expect Algorithm 5.1 to produce better results in terms of actual diversity increase than Algortihm 5.2.

5.4.3 Performance of Algorithms

We firstly apply our algorithms to a set of 500 random users. We compare the output of the two algorithms with the output of two alternative methods: (i) top k recommended users from set R and (ii) k random recommended users from R. For these experiments, the set R is constructed using the Common Neighbors link prediction method. We should also note that we recommend 5 users for each user.

The score that is returned from Common Neighbors method does not give the probability for an edge to be accepted that we need for Algortihm 2. We need a method to transform these scores into probabilistic scores that yield from 0 to 1. So we do the following procedure: We normalize the scores by dividing with the maximum score over all predictions. We then apply a function $f(x)$ to an edge with normalized score x, in order to shape this probability. The function that we use for this experiment is $f(x) = x$.

We also performed the same experiment using other parameters. For example we used Adamic Adar link prediction for constructing R. We tried different number of users to recommend (different values of k). We also tried $f(x) = 0.4log_{10}x + 1$ as function for estimating the probability. Due to the fact that the results were almost identical, we do not provide them.

The Entropy increase for these methods is depicted in Table 5.4 and 5.5. We see that our methods performs better than the two other methods. We also observe that the two other methods perform almost the same. This means that top k from recommended, recommends random users in terms of diversity. Finally, as we mentioned above, it is observed that there exists a tradeoff between accuracy and the utility of recommendations, as we see that Algorithm 5.2 performs slightly worse than Algorithm 5.1. We must clarify that the third column of Table 5.4 represents the actual diversity after applying Algorithm 5.2, even though it is based on a different criterion (the expected increase).

Dataset	Greedy	Greedy(Expected)	Top k from R	Random from R
ClimateChange	26.10%	24.02%	12.21%	12.79%
Elections2016	30.88%	27.67%	16.11%	16.84%
NowPlaying	21.57%	19.28%	7.58%	7.93%
Politics	25.36%	22.95%	15.72%	16.64%
PresidentialDebate	30.16%	28.83%	14.86%	15.04%
PresDeb-Elections2016	27.93%	25.57%	13.58%	14.35%
PresDeb-Supernatural	25.13%	21.52%	12.39%	11.94%
Supernatural	18.38%	14.89%	9.41%	8.81%
USElection2016	22.15%	20.01%	13.62%	14.03%
WorldSeries	21.32%	18.26%	15.47%	15.81%

Table 5.4: Comparison of Greedy algorithm with two alternative methods in terms of diversity increase

Dataset	Greedy(Expected)	Top k from R	Random from R
ClimateChange	8.72%	5.92%	3.16%
Elections2016	13.23%	9.48%	4.39%
NowPlaying	6.19%	4.03%	2.42%
Politics	4.34%	2.69%	1.86%
PresidentialDebate	7.57%	5.01%	2.07%
PresDeb-Elections2016	6.74%	4.28%	2.94%
PresDeb-Supernatural	8.72%	6.12%	3.14%
Supernatural	5.81%	3.89%	2.57%
USElection2016	9.11%	6.74%	3.96%
WorldSeries	9.51%	6.91%	4.18%

Table 5.5: Comparison of Greedy (expected) algorithm with two alternative methods in terms of expected diversity increase

Subsequently, we focus on which users improve more their diversity. So, we study average diversity increase in terms of degree of users. We expect that users that have lower degree will be easier to increase their diversity (we can easily balance the content he recieves from his neighbors), than users who have many neighbors. Moreover, we observed in the previous sections, that low degree users have lower diversity than high degree users. We plot the correlation between average increase of diversity and degree of users. The plots are shown in Figure C.3 in Appendix C.2 and prove that the assumption we made is true. We also focus on the diversity increase in terms of polarity values. Unfortunately no correlation between these two exists, so we omit the results.

5.4.4 Performance in terms of Polarity

In previous sections we showed the effectiveness of our algorithms in terms of diversity increase. We want to test our algorithms in terms of user polarity, in order to see if the polarity values of the users we target are decreased. Results indicate that the polarity of the users we targeted (the same random users we tagreted above) is reduced on average for the both algorithms. We depict the results in Table 5.6.

Dataset	Avg Polarity Decrease (Greedy)	Avg Polarity Decrease (Expected)
ClimateChange	14.9%	15.2%
Elections2016	9%	11.5%
NowPlaying	11.6%	9.2%
Politics	34.4%	29.6%
PresidentialDebate	8.2%	12.9%
PresDeb-Elections2016	15.7%	14.7%
PresDeb-Supernatural	22.5%	21.6%
Supernatural	42.6%	44.6%
USElection2016	7%	9.2%
WorldSeries	1%	2.3%

Table 5.6: Average polarity decrease after the application of the two algorithms.

We also focus on the controversy score of the whole network after our recommendations. After applying the two methods (in total 2.500 new edges) we observe that we have a slight reduction in this score. We see the results in Table 5.7.

Dataset	Contr. score after Greedy	Contr. score after Expected	Contr. Before
ClimateChange	38%	39%	44%
Elections2016	57%	59%	64%
NowPlaying	81%	80%	89%
Politics	41%	39%	43%
PresidentialDebate	60%	61%	65%
PresDeb-Elections2016	56%	53%	60%
PresDeb-Supernatural	58%	60%	63%
Supernatural	53%	54%	59%
USElection2016	49%	51%	56%
WorldSeries	34%	33%	38%

Table 5.7: Controversy score after the appliance of the two algorithms.

5.4.5 Comparison to the State of the art

We compare our algorithms (in terms of controversy score) with the two algorithms proposed in [11], named ROV and ROV-AP. The difference between these two algorithms with ours, is that their objective function is the reduction of the controversy score (and the expected controversy score respectively) of the whole network (by recommending edges), while ours, make user level recommendations.

In order to make the comparison possible and fair, we find the users that the edges they recommend connect and then we make one recommendation for each one of them. In Tables 5.8 and 5.9 we depict the results after making 500 recommendations. As we expect ROV and ROV-AP perform very well in terms of controversy score, because this is the objective function that they aim to minimize. We observe that, although our algorithms are not designed to optimize this objective function, they perform quite well in terms of controversy score, as we see that it is reduced in all datasets.

We also compare our algorithms with ROV and ROV-AP in terms of diversity increase (Tables 5.10 and 5.11) . The scores depict the average diversity increase for the 500 users that we made recommendations. We see that in all datasets, Greedy and Greedy Expected outperform ROV and ROV-AP. We also observe that in some cases the average diversity increase is negative after the recommendation made by

ROV and ROV-AP.

Dataset	Contr. score after Greedy	Contr. score after ROV	Contr. Before
ClimateChange	42%	38%	44%
Elections2016	61%	57%	64%
NowPlaying	82%	76%	89%
Politics	42%	39%	43%
PresidentialDebate	61%	58%	65%
PresDeb-Elections2016	57%	53%	68%
PresDeb-Supernatural	59%	54%	63%
Supernatural	54%	51%	59%
USElection2016	52%	50%	56%
WorldSeries	36%	29%	38%

Table 5.8: Comparison of Greedy with ROV in terms of controversy score

Dataset	Contr. score after Expected	Contr. score after ROV-AP	Contr. Before
ClimateChange	42%	39%	44%
Elections2016	62%	58%	64%
NowPlaying	84%	78%	89%
Politics	41%	38%	43%
PresidentialDebate	61%	58%	65%
PresDeb-Elections2016	56%	54%	68%
PresDeb-Supernatural	60%	55%	63%
Supernatural	55%	50%	59%
USElection2016	52%	49%	56%
WorldSeries	36%	31%	38%

Table 5.9: Comparison of Greedy (Expected) with ROV-AP in terms of controversy score

Dataset	Diversity Increase after Greedy	Diversity Increase after ROV
ClimateChange	2.93%	-0.91%
Elections2016	2.74%	0.46%
NowPlaying	9.89%	1.23%
Politics	6.51%	0.98%
PresidentialDebate	2.76%	0.42%
PresDeb-Elections2016	4.92%	-0.22%
PresDeb-Supernatural	3.98%	0.26%
Supernatural	6.21%	0.35%
USElection2016	3.77%	0.46%
WorldSeries	7.24%	0.18%

Table 5.10: Comparison of Greedy with ROV in terms of diversity increase

Dataset	Diversity Increase after Expected	Diversity Increase after ROV-AP
ClimateChange	2.48%	0.38%
Elections2016	1.68%	0.17%
NowPlaying	8.71%	0.24%
Politics	6.22%	0.18%
PresidentialDebate	1.92%	1.17%
PresDeb-Elections2016	3.74%	0.35%
PresDeb-Supernatural	3.18%	0.48%
Supernatural	4.82%	-0.18%
USElection2016	2.15%	0.68%
WorldSeries	5.59%	-0.24%

Table 5.11: Comparison of Greedy (Expected) with ROV-AP in terms of diversity increase

Chapter 6

Conclusions and Future Work

6.1 Conclusions

In this thesis, we studied the co-occurence of polarization and diversity in social networks on themes from different domains, by exploiting both network and content features. We firslty applied an existing model, to identify which of the themes we study are more controversial and cause people to split into opposing sides. This method considers only the structure of the network.

We focused on the content of the users to study homophily between the opposing partitions. We found that in more controversial themes, the similarity of users that belong to the same community is bigger than this of users that belong to different communities. We also defined a user level homophily metric and we examined the correlation between this and polarity of users.

Subsequently, we defined and measured diversity in user and network level. Based on these definitions we investigated the existence of echo chambers in all of the networks. We also addressed the correlation between polarity and diversity. Results showed that there exists a negative correlation between these two. Finally we proposed two algorithms that recommend content to users, so to improve their diversity. We experimented on real data collected from Twitter and demonstrated that our algorithms perform well not only in terms of diversity, but also in terms of polarity in user and network level.

6.2 Future Work

In our work, we mainly focused on the retweet datasets. An alternative approach would be to use the follow networks to find the opposing partitions, provide statistics about them and calculate their controversy scores. We could also apply the two recommendation algorithms to these networks, with the purpose to connect users with follow relationships.

As far as the content of users is concerned, in this study we rely on topic models as the main text mining technique. We could use alternative techniques for this task, as characterizing users with their most frequent hashtags or with their most important words using (weighted) tf-idf.

Additionally, the two algorithms described in Section 5.4.2 aim to increase diversity of individual users. We could propose an algorithm that its objective function is the increase of network diversity. This, is not a simple task, as we would have to focus on the users that we should target (increase their diversity) in order to maximize network diversity.

Finally, another approach in the recommendations would be to recommend specific tweets (not all tweets) to users that are more possible to be accepted by them. In this way, we increase the probability of links (retweets) to be materialized.

Bibliography

[1] A. Bessi, F. Petroni, M. Del Vicario, F. Zollo, A. Anagnostopoulos, A. Scala, G. Caldarelli, and W. Quattrociocchi, "Viral misinformation: The role of homophily and polarization," in *Proceedings of the 24th International Conference on World Wide Web*, WWW '15 Companion, (New York, NY, USA), pp. 355–356, ACM, 2015.

[2] P. Dandekar, A. Goel, and D. Lee, "Biased assimilation, homophily and the dynamics of polarization," *CoRR*, vol. abs/1209.5998, 2012.

[3] E. Bakshy, S. Messing, and L. A. Adamic, "Exposure to ideologically diverse news and opinion on facebook," *Science*, vol. 348, no. 6239, pp. 1130–1132, 2015.

[4] A. Bessi, F. Zollo, M. D. Vicario, M. Puliga, A. Scala, G. Caldarelli, B. Uzzi, and W. Quattrociocchi, "Users polarization on facebook and youtube," *CoRR*, vol. abs/1604.02705, 2016.

[5] R. K. Garrett, "Echo chambers online?: Politically motivated selective exposure among internet news users1," *Journal of Computer-Mediated Communication*, vol. 14, no. 2, pp. 265–285, 2009.

[6] K. Garimella, G. D. F. Morales, A. Gionis, and M. Mathioudakis, "Political discourse on social media: Echo chambers, gatekeepers, and the price of bipartisanship," *CoRR*, vol. abs/1801.01665, 2018.

[7] E. Pariser, *The Filter Bubble: What the Internet Is Hiding from You*. Penguin Group , The, 2011.

[8] M. Conover, J. Ratkiewicz, M. Francisco, B. Gonçalves, A. Flammini, and F. Menczer, "Political polarization on twitter," in *Proc. 5th International AAAI Conference on Weblogs and Social Media (ICWSM)*, 2011.

[9] P. Barberá, J. T. Jost, J. Nagler, J. A. Tucker, and R. Bonneau, "Tweeting from left to right: Is online political communication more than an echo chamber?," *Psychological science*, vol. 26 10, pp. 1531–42, 2015.

[10] P. Barberá, J. T. Jost, J. Nagler, J. A. Tucker, and R. Bonneau, "Tweeting from left to right: Is online political communication more than an echo chamber?," *Psychological Science*, vol. 26, no. 10, pp. 1531–1542, 2015. PMID: 26297377.

[11] K. Garimella, G. D. F. Morales, A. Gionis, and M. Mathioudakis, "Reducing controversy by connecting opposing views," in *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining, WSDM 2017, Cambridge, United Kingdom, February 6-10, 2017*, pp. 81–90, 2017.

[12] K. Garimella, G. D. F. Morales, A. Gionis, and M. Mathioudakis, "Quantifying controversy in social media," in *Proceedings of the Ninth ACM International Conference on Web Search and Data Mining, San Francisco, CA, USA, February 22-25, 2016*, pp. 33–42, 2016.

[13] N. Parotsidis, E. Pitoura, and P. Tsaparas, "Centrality-aware link recommendations," in *Proceedings of the Ninth ACM International Conference on Web Search and Data Mining, San Francisco, CA, USA, February 22-25, 2016*, pp. 503–512, 2016.

[14] A. Matakos, E. Terzi, and P. Tsaparas, "Measuring and moderating opinion polarization in social networks," *Data Mining and Knowledge Discovery*, vol. 31, pp. 1480–1505, Sep 2017.

[15] K. Garimella, G. D. F. Morales, A. Gionis, and M. Mathioudakis, "Factors in recommending contrarian content on social media," in *Proceedings of the 2017 ACM on Web Science Conference, WebSci 2017, Troy, NY, USA, June 25 - 28, 2017*, pp. 263–266, 2017.

[16] K. Garimella, G. De Francisc iMorales, A. Gionis, and M. Mathioudakis, "Mary, mary, quite contrary: Exposing twitter users to contrarian news," in *Proceedings of the 26th International Conference on World Wide Web Companion*, WWW '17 Companion, (Republic and Canton of Geneva, Switzerland), pp. 201–205, International World Wide Web Conferences Steering Committee, 2017.

[17] E. Graells-Garrido, M. Lalmas, and D. Quercia, "Data portraits: Connecting people of opposing views," *CoRR*, vol. abs/1311.4658, 2013.

[18] C. Ziegler, S. M. McNee, J. A. Konstan, and G. Lausen, "Improving recommendation lists through topic diversification," in *Proceedings of the 14th international conference on World Wide Web, WWW 2005, Chiba, Japan, May 10-14, 2005*, pp. 22–32, 2005.

[19] M. Drosou and E. Pitoura, "Search result diversification," *SIGMOD Rec.*, vol. 39, pp. 41–47, Sept. 2010.

[20] M. Drosou and E. Pitoura, "Disc diversity: Result diversification based on dissimilarity and coverage," *CoRR*, vol. abs/1208.3533, 2012.

[21] V. Dang and W. B. Croft, "Diversity by proportionality: an election-based approach to search result diversification," in *The 35th International ACM SIGIR conference on research and development in Information Retrieval, SIGIR '12, Portland, OR, USA, August 12-16, 2012*, pp. 65–74, 2012.

[22] K. Garimella, A. Gionis, N. Parotsidis, and N. Tatti, "Balancing information exposure in social networks," *CoRR*, vol. abs/1709.01491, 2017.

[23] S. A. Munson and P. Resnick, "Presenting diverse political opinions: How and how much," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '10, (New York, NY, USA), pp. 1457–1466, ACM, 2010.

[24] S. Munson, S. Lee, and P. Resnick, *Encouraging reading of diverse political viewpoints with a browser widget*, pp. 419–428. AAAI press, 2013.

[25] G. Karypis and V. Kumar, "Metis – unstructured graph partitioning and sparse matrix ordering system, version 2.0," 01 1995.

[26] A. K. McCallum, "Mallet: A machine learning for language toolkit." http://mallet.cs.umass.edu, 2002.

Appendix A

Quantifying Controversy (Appendix)

A.1 Label Partitioning Statistics

ClimateChange	Trump	Clinton	No Label	Total
Community1	3222	0	0	3222
Community2	0	8699	0	8699

Table A.1: Statistics of the 2 communities made by Labels for ClimateChange

Elections2016	Trump	Clinton	No Label	Total
Community1	9578	0	0	9578
Community2	0	8326	0	8326

Table A.2: Statistics of the 2 communities made by Labels for Elections2016

NowPlaying	Trump	Clinton	No Label	Total
Community1	1216	0	0	1216
Community2	0	2689	0	2689

Table A.3: Statistics of the 2 communities made by Labels for NowPlaying

Politics	Trump	Clinton	No Label	Total
Community1	5403	0	0	5403
Community2	0	2418	0	2418

Table A.4: Statistics of the 2 communities made by Labels for Politics

PresidentialDebate	Trump	Clinton	No Label	Total
Community1	8242	0	0	8242
Community2	0	8484	0	8484

Table A.5: Statistics of the 2 communities made by Labels for PresidentialDebate

PresDeb-Elections2016	Trump	Clinton	No Label	Total
Community1	15733	0	0	15733
Community2	0	16025	0	14574

Table A.6: Statistics of the 2 communities made by Labels for PresidentialDebate-Elections2016

PresDeb-Supernatural	Trump	Clinton	No Label	Total
Community1	9089	0	0	9089
Community2	0	9908	0	9073

Table A.7: Statistics of the 2 communities made by Labels for PresidentialDebate-Supernatural

Supernatural	Trump	Clinton	No Label	Total
Community1	882	0	0	882
Community2	0	1442	0	1442

Table A.8: Statistics of the 2 communities made by Labels for Supernatural

USElection2016	Trump	Clinton	No Label	Total
Community1	5677	0	0	5677
Community2	0	6081	0	6081

Table A.9: Statistics of the 2 communities made by Labels for USElection2016

WorldSeries	Trump	Clinton	No Label	Total
Community1	40898	0	0	40898
Community2	0	28283	0	28283

Table A.10: Statistics of the 2 communities made by Labels for WorldSeries

A.2 User Polarity Distribution

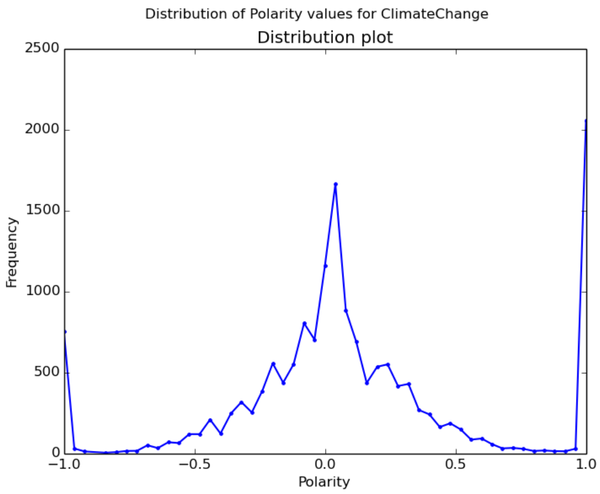

(a) ClimateChange

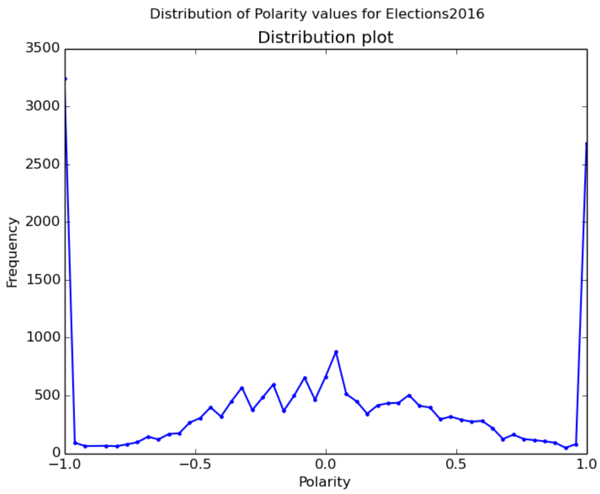

(b) Elections2016

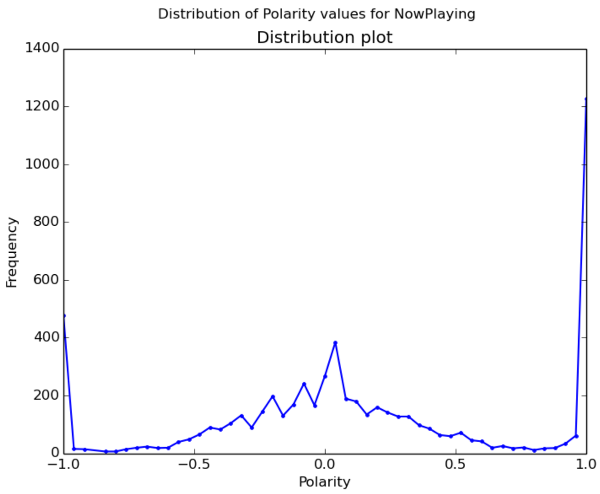

(c) NowPlaying

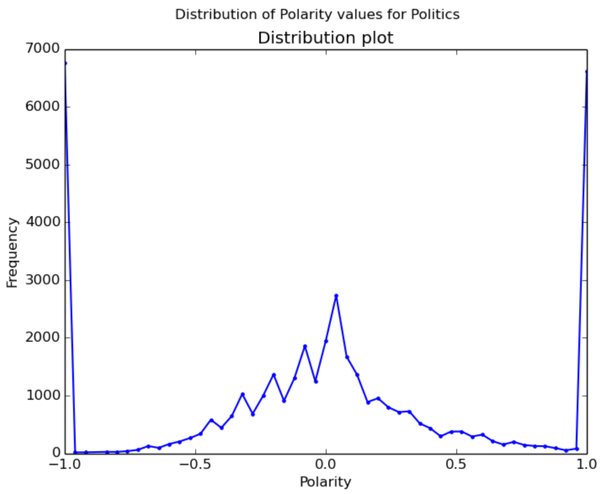

(d) Politics

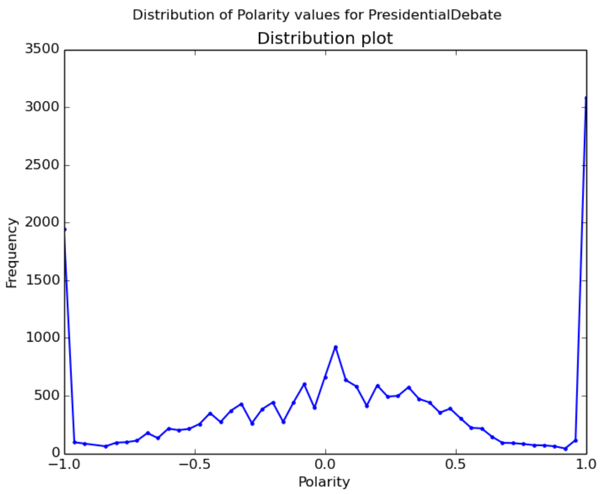

(e) PresidentialDebate

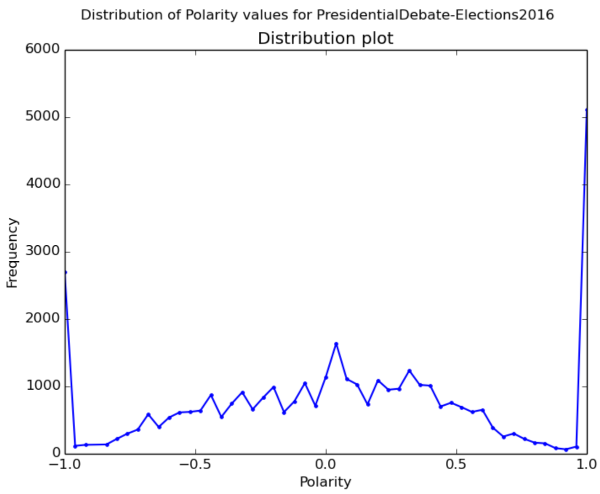

(f) PresidentialDebate-Elections2016

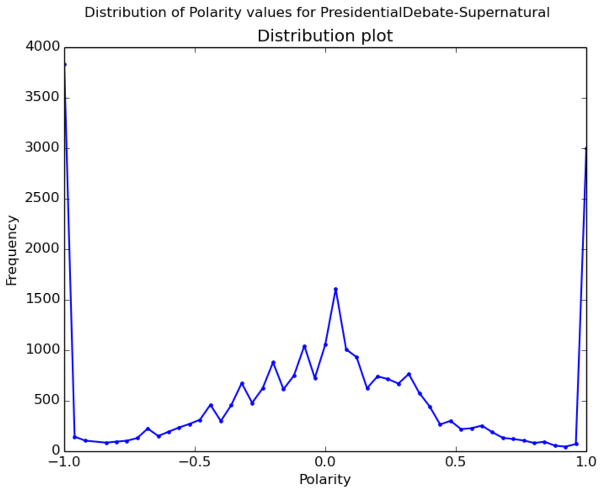

(g) PresidentialDebate-Supernatural

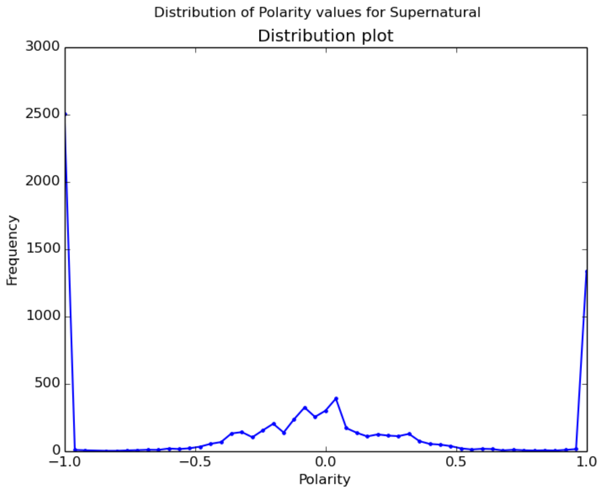

(h) Supernatural

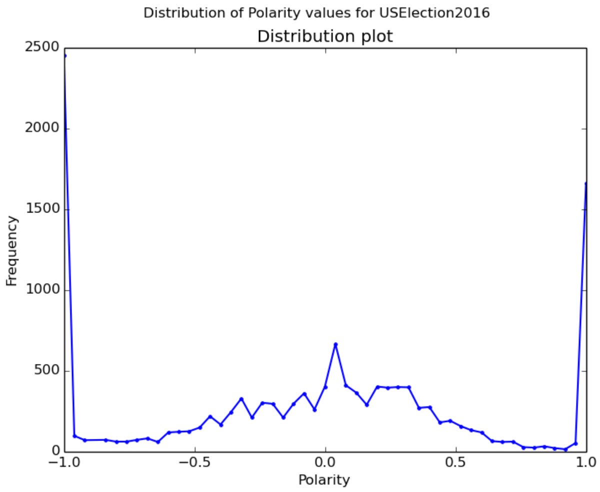

(i) USElection2016

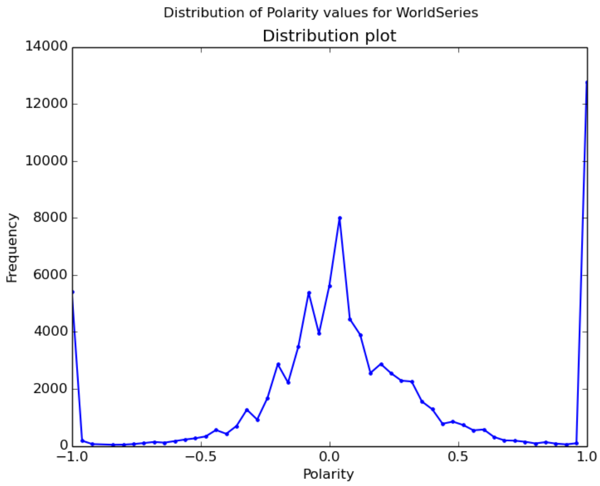

(j) WorldSeries

Figure A.1: User Polarity Distribution for all datasets

Appendix B

Existence of Homophily (Appendix)

B.1 Exact Homophily Values

Dataset	Homophily Community1	Homophily Community2	Homophily Between Communities
ClimateChange	0.4510421898	0.3749230906	0.383262213
Elections2016	0.3983516549	0.3976081204	0.3175865515
NowPlaying	0.4769082542	0.5582042715	0.3035974189
Politics	0.5119812405	0.5466053906	0.4848204283
PresidentialDebate	0.4179055335	0.4187091046	0.3478762187
PresDeb-Elections2016	0.4142522234	0.4356810529	0.3365259613
PresDeb-Supernatural	0.4191931204	0.4173896609	0.356933239
Supernatural	0.4851803951	0.5274347166	0.3417553803
USElection2016	0.3997878685	0.3831825595	0.3268656435
WorldSeries	0.435008893	0.4216456607	0.4348975728

Table B.1: Average homophily between users for all datasets

Dataset	Homophily Community1	Homophily Community2	Homophily Between Communities
ClimateChange	0.700495809	0.717032269	0.693340082
Elections2016	0.8080472113	0.7895968303	0.7119085195
NowPlaying	0.7548664851	0.7863247537	0.6074519313
Politics	0.5748736027	0.6049755362	0.5241896523
PresidentialDebate	0.7084070501	0.7224552397	0.5781789603
PresDeb-Elections2016	0.6996445027	0.7534317942	0.638121211
PresDeb-Supernatural	0.7543444609	0.7577103915	0.625310137
Supernatural	0.6576469896	0.7199796597	0.5111153005
USElection2016	0.767028945	0.7943995776	0.6884301264
WorldSeries	0.7252123061	0.6847923139	0.7067481008

Table B.2: Average homophily between directly connected users for all datasets

Appendix C

Diversity (Appendix)

C.1 Echo Chambers

C.1.1 Follow Networks

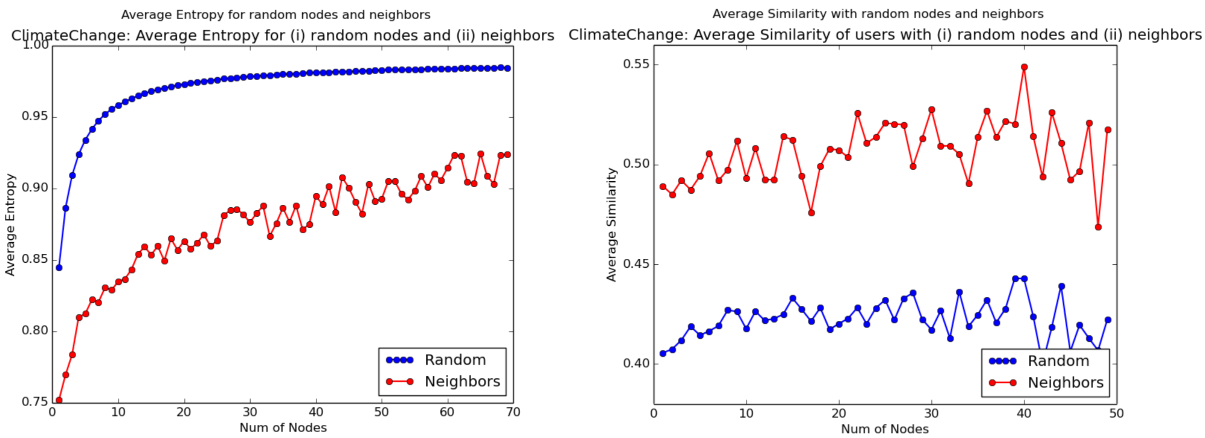

(a) ClimateChange Diversity

(b) ClimateChange Similarity

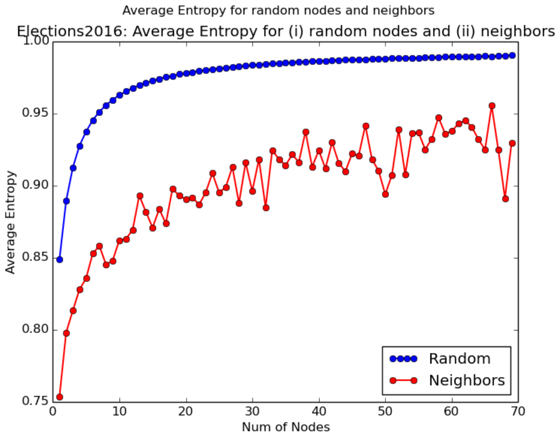

(c) Elections2016 Diversity

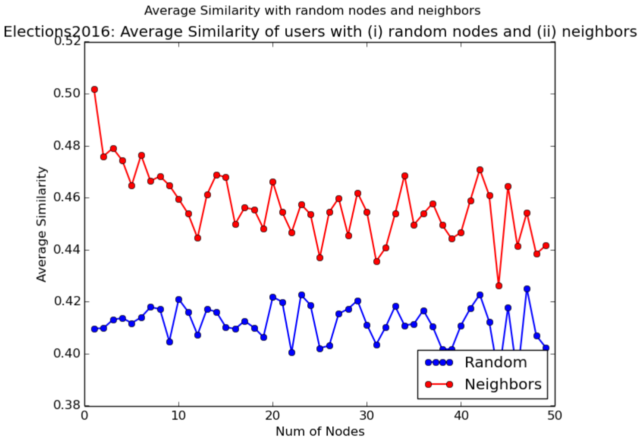

(d) Elections2016 Similarity

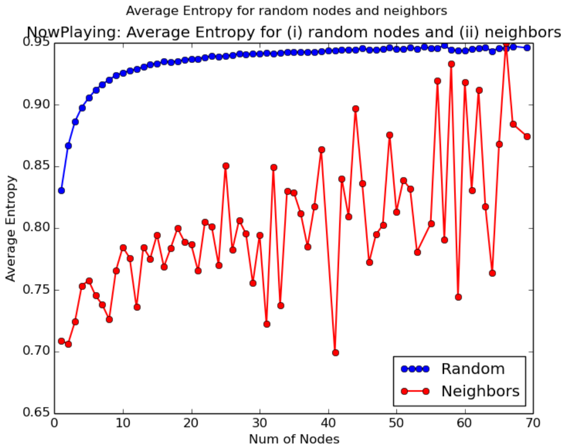

(e) NowPlaying Diversity

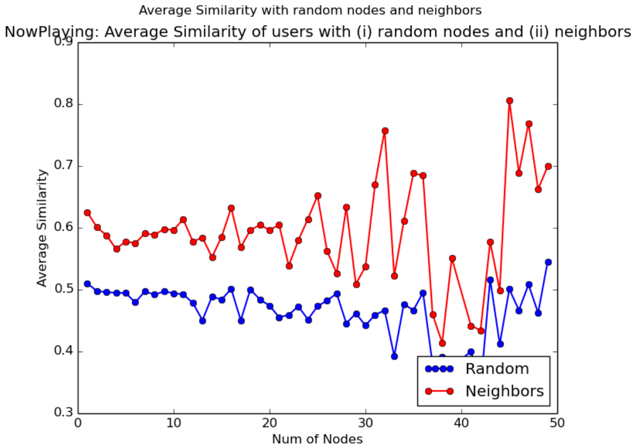

(f) NowPlaying Similarity

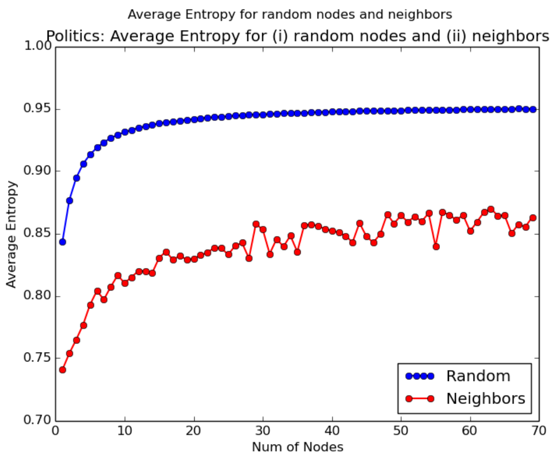

(g) Politics Diversity

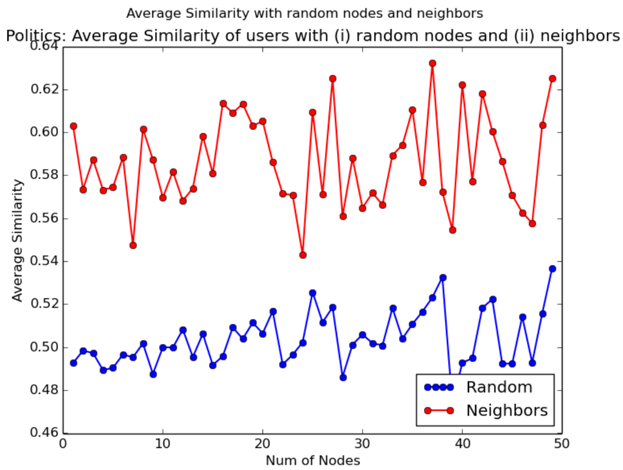

(h) Politics Similarity

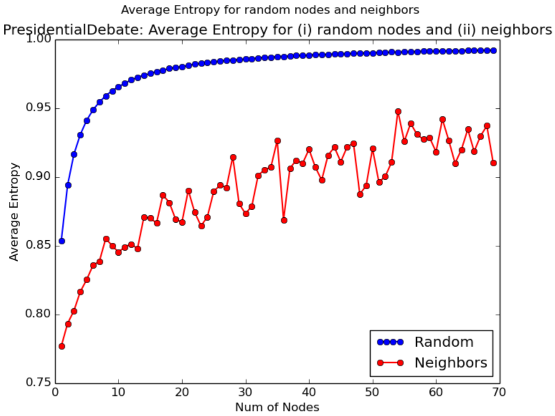

(i) PresidentialDebate Diversity

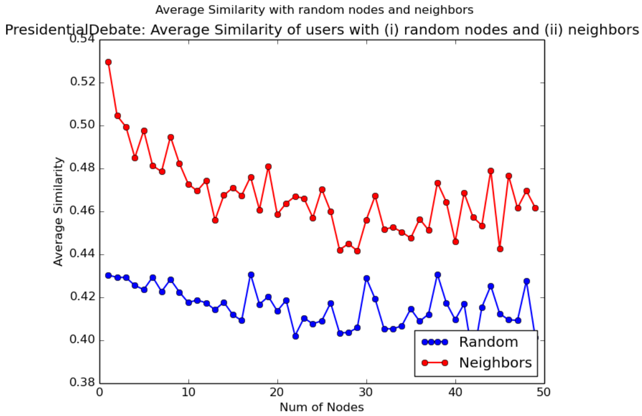

(j) PresidentialDebate Similarity

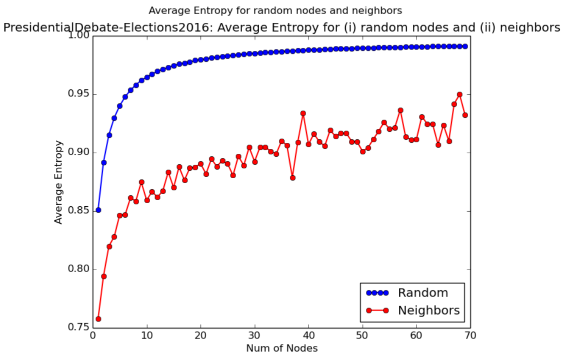

(k) PresidentialDebate-Elections2016 Diversity

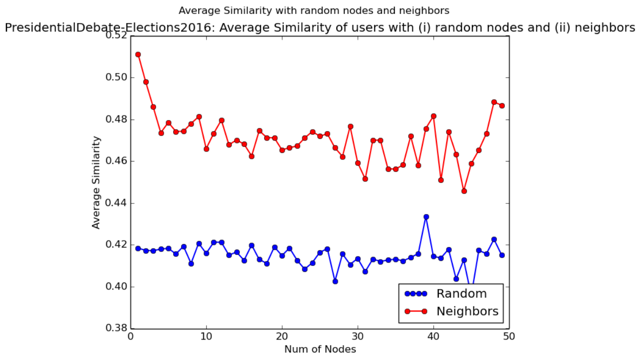

(l) PresidentialDebate-Elections2016 Similarity

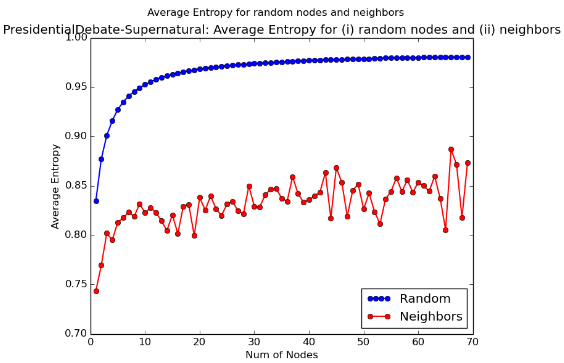

(m) PresidentialDebate-Supernatural Diversity

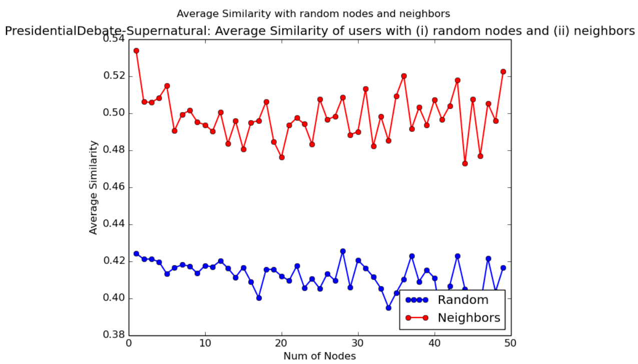

(n) PresidentialDebate-Supernatural Similarity

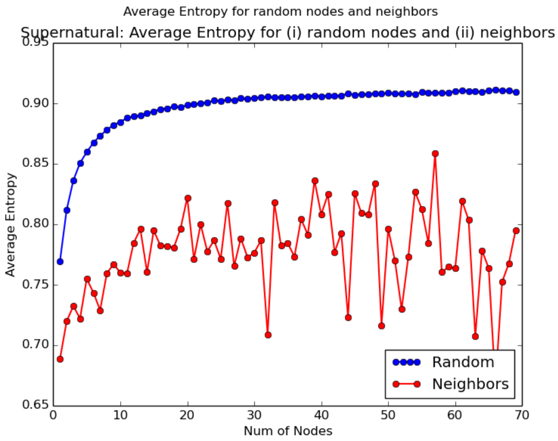

(o) Supernatural Diversity

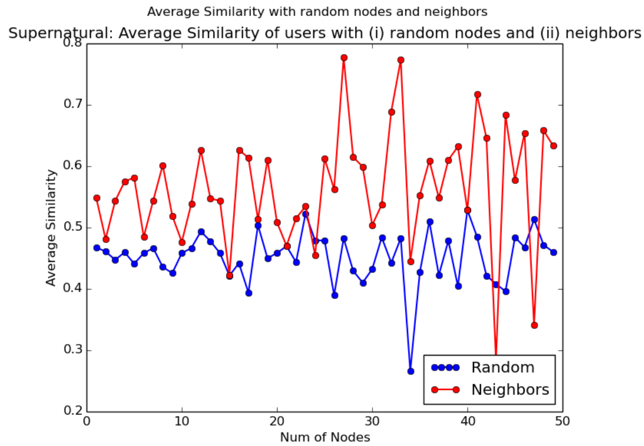

(p) Supernatural Similarity

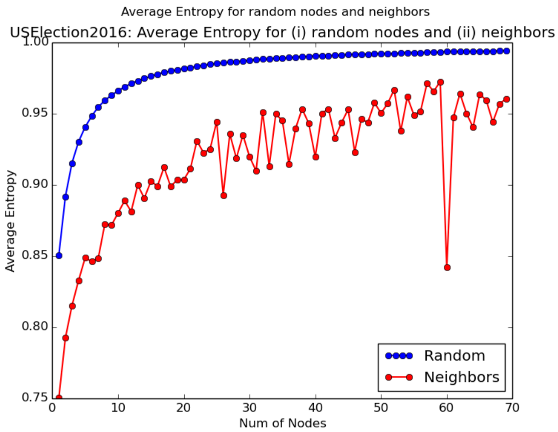

(q) USElection2016 Diversity

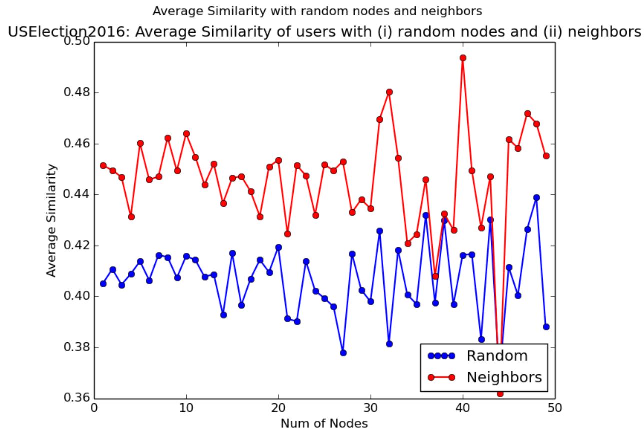

(r) USElection2016 Similarity

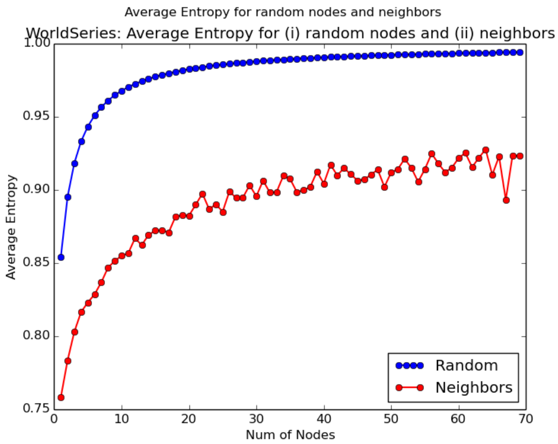

(s) WorldSeries Diversity

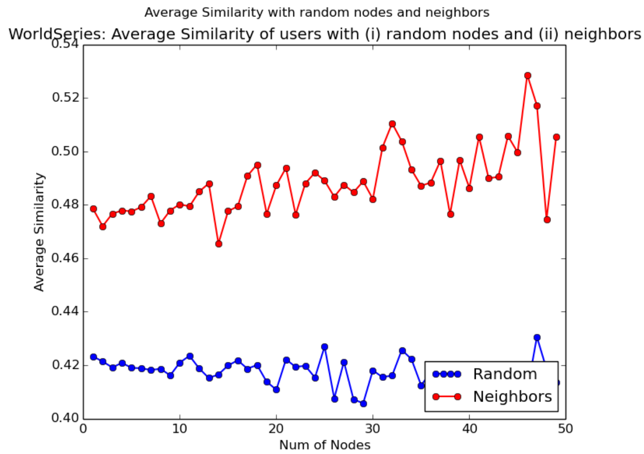

(t) WorldSeries Similarity

Figure C.1: Existence of echo chambers in follow networks (Left plots: Average Diversity, Right plots: Average Similarity)

C.1.2 Retweet Networks

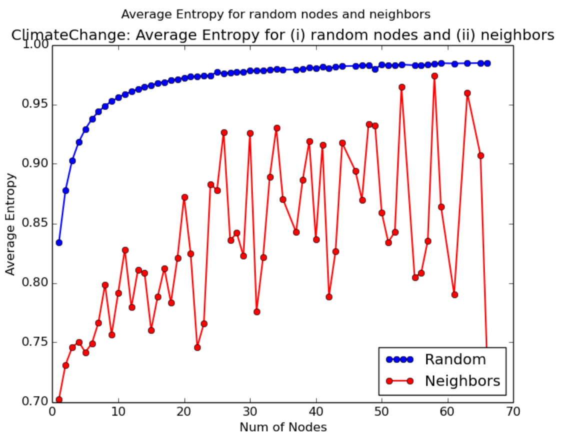

(a) ClimateChange Diversity

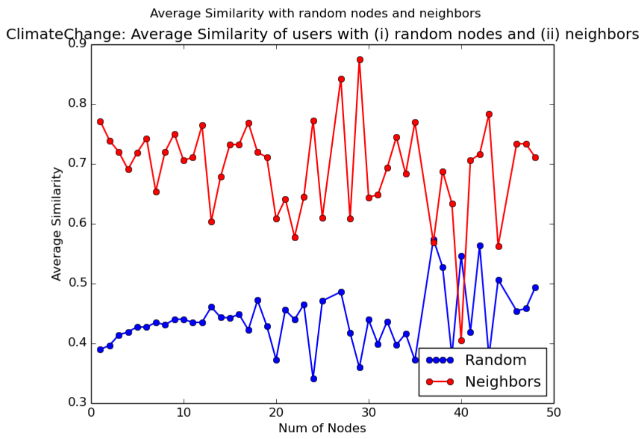

(b) ClimateChange Similarity

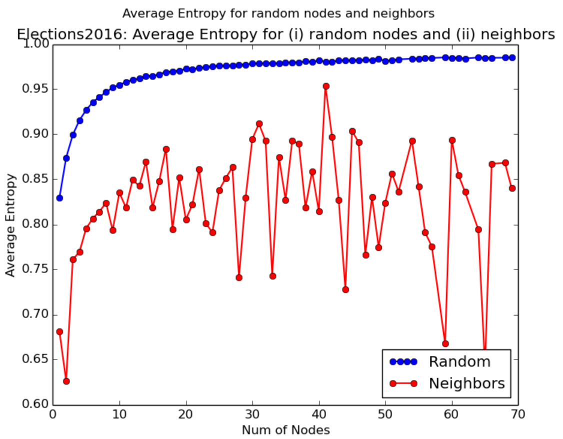

(c) Elections2016 Diversity

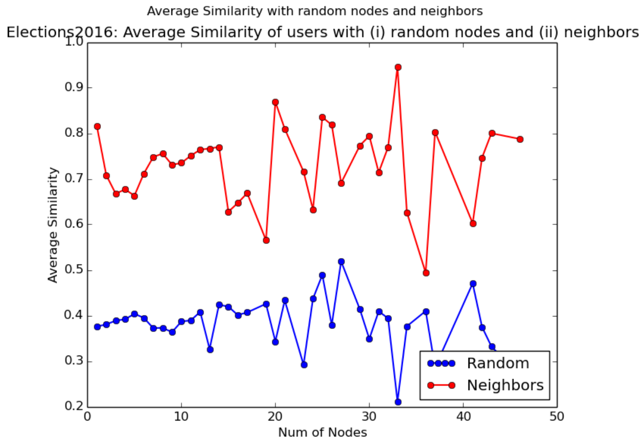

(d) Elections2016 Similarity

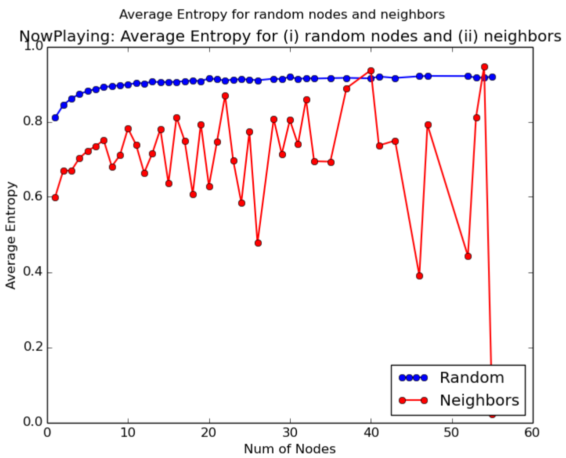

(e) NowPlaying Diversity

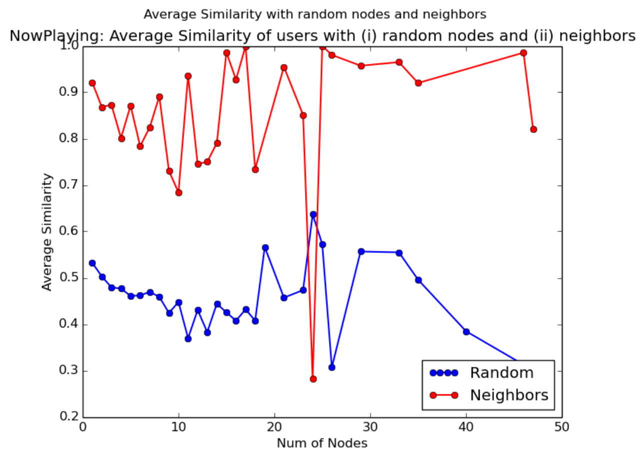

(f) NowPlaying Similarity

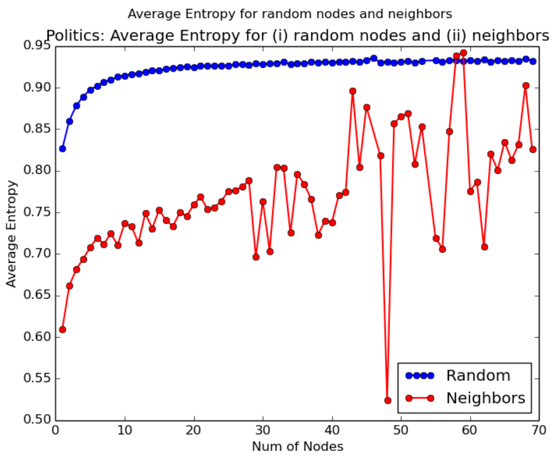

(g) Politics Diversity

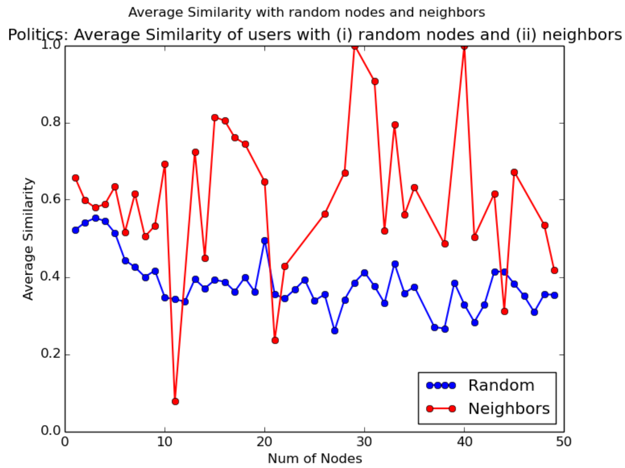

(h) Politics Similarity

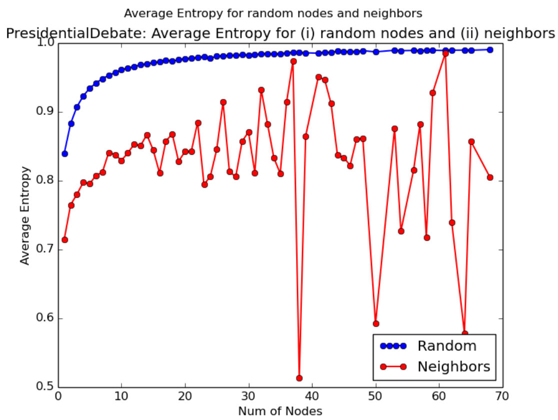

(i) PresidentialDebate Diversity

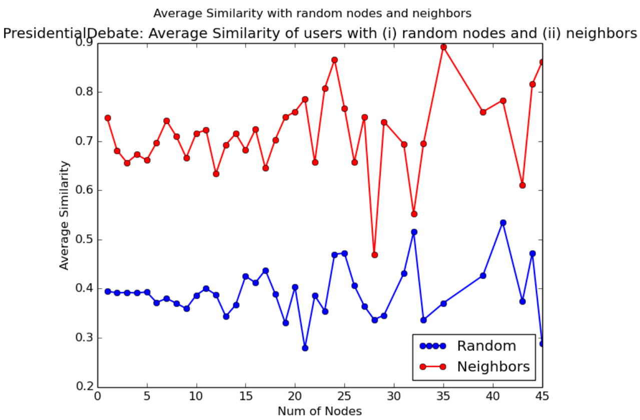

(j) PresidentialDebate Similarity

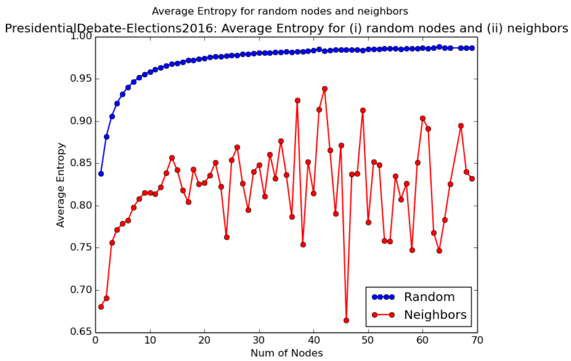

(k) PresidentialDebate-Elections2016 Diversity

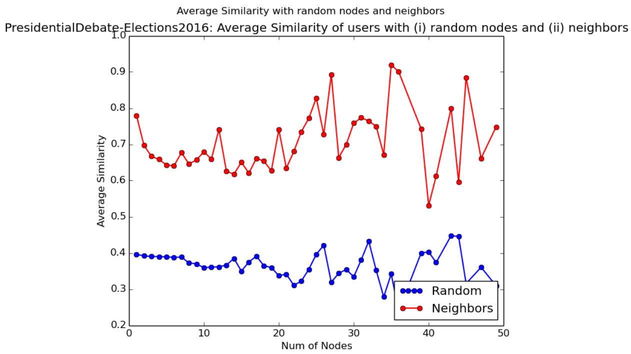

(l) PresidentialDebate-Elections2016 Similarity

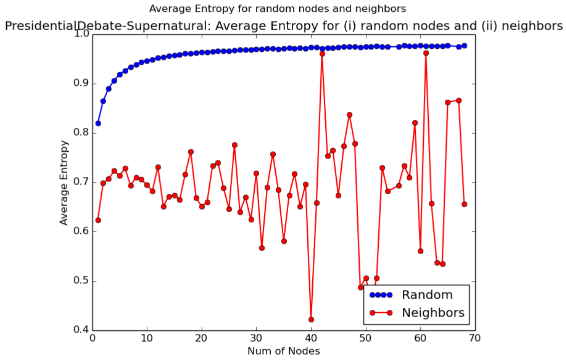

(m) PresidentialDebate-Supernatural Diversity

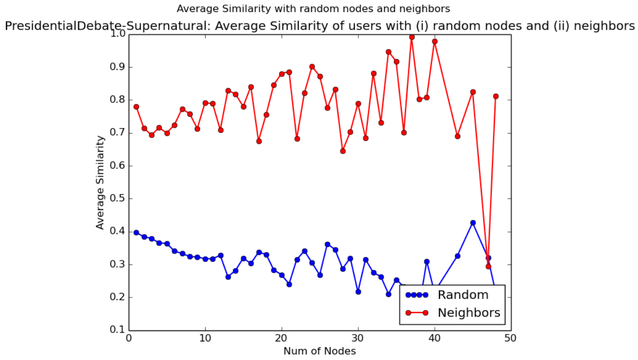

(n) PresidentialDebate-Supernatural Similarity

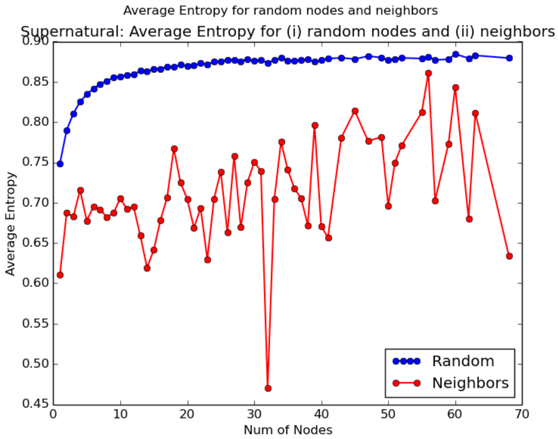

(o) Supernatural Diversity

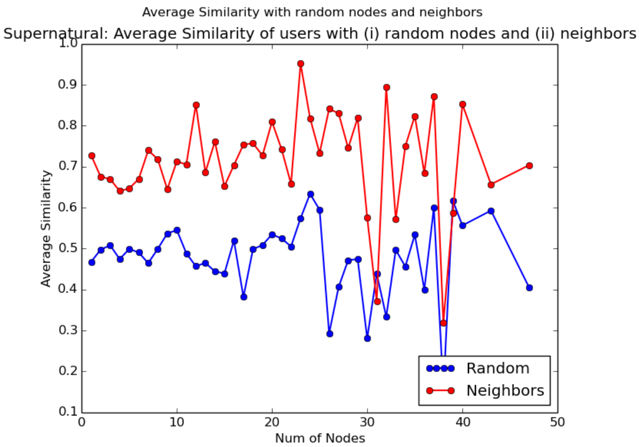

(p) Supernatural Similarity

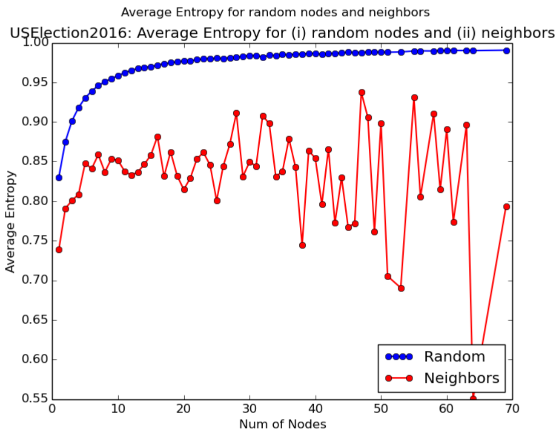

(q) USElection2016 Diversity

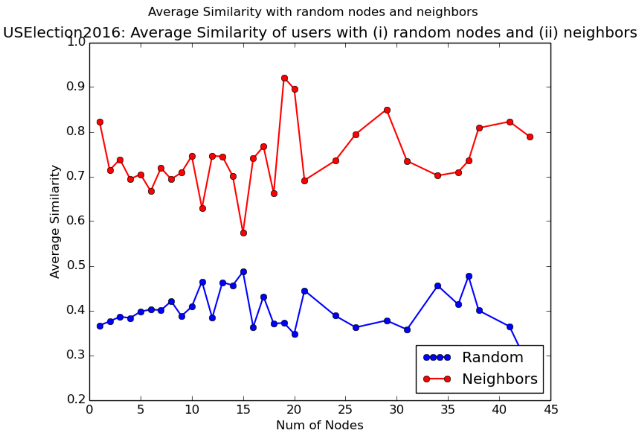

(r) USElection2016 Similarity

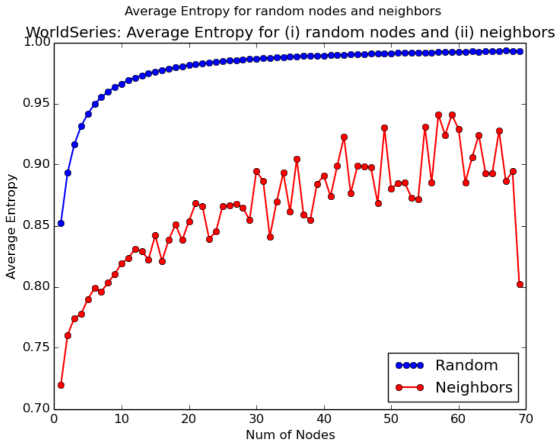

(s) WorldSeries Diversity

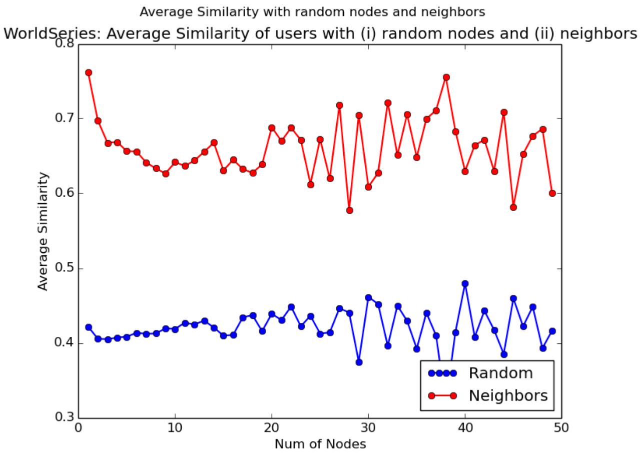

(t) WorldSeries Similarity

Figure C.2: Existence of echo chambers in retweet networks (Left plots: Average Diversity, Right plots: Average Similarity)

C.2 Diversity Increase

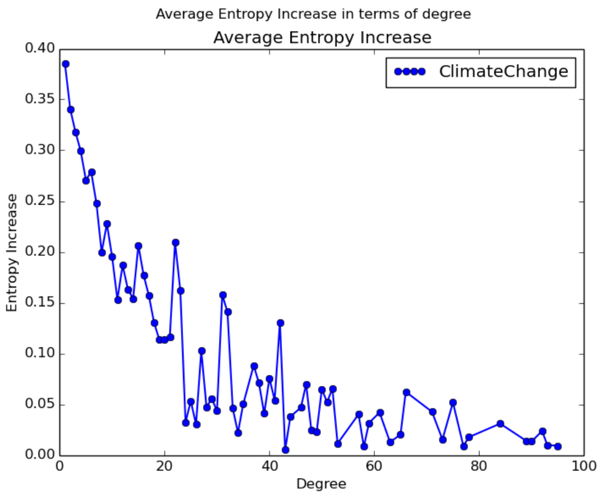

(a) ClimateChange

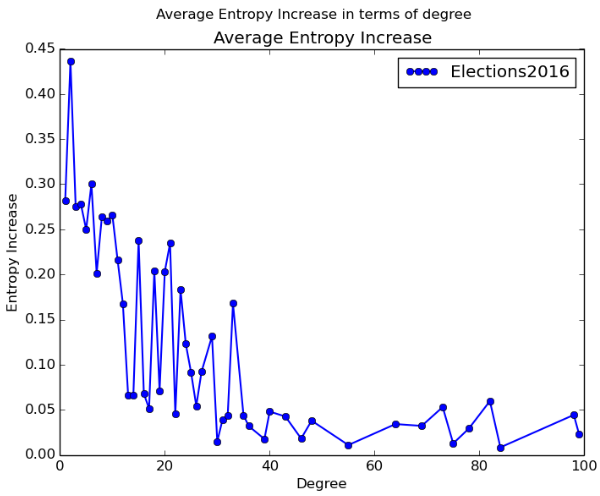

(b) Elections2016

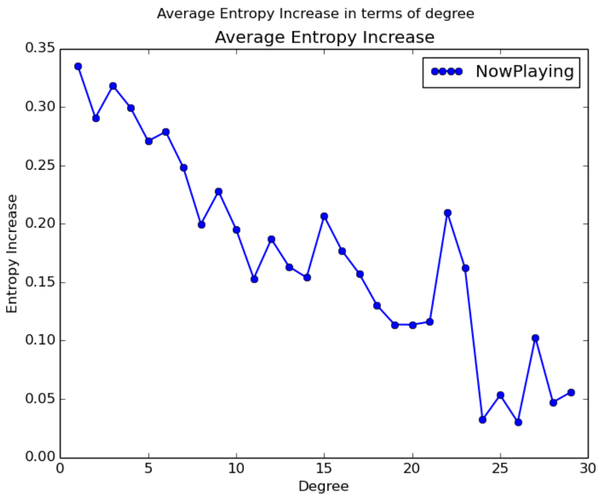

(c) NowPlaying

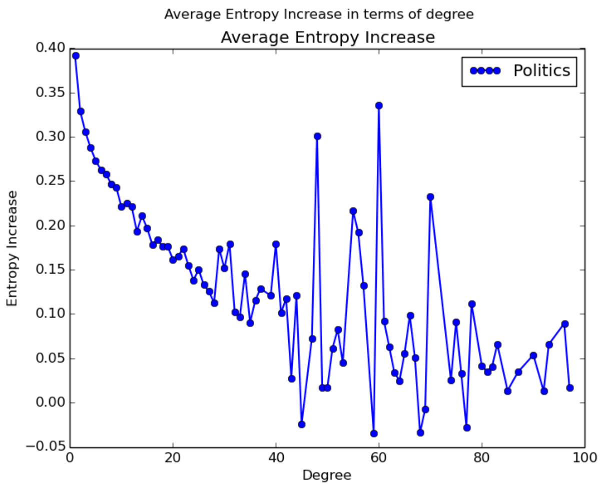

(d) Politics

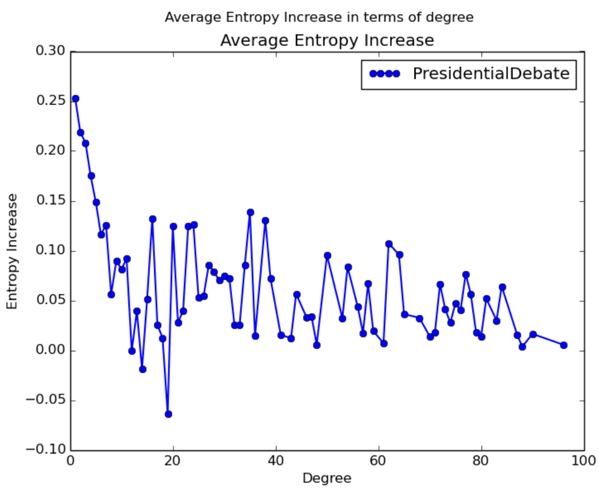

(e) PresidentialDebate

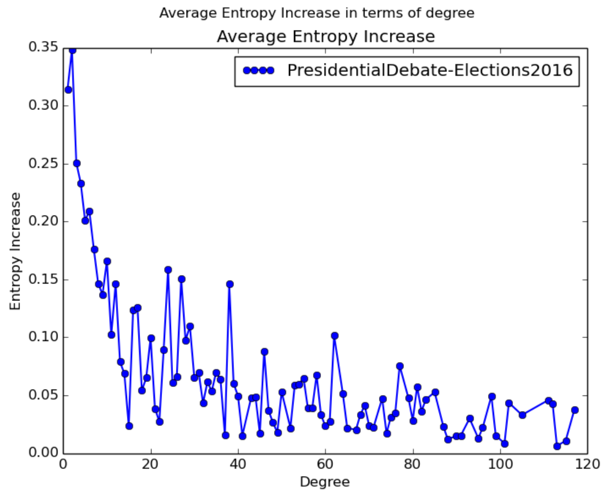

(f) PresidentialDebate-Elections2016

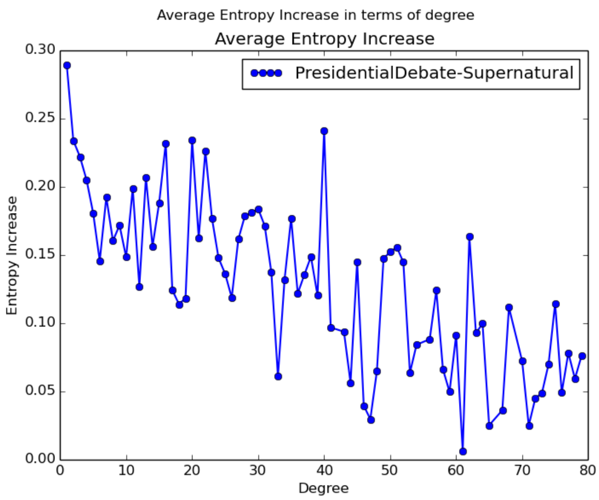

(g) PresidentialDebate-Supernatural

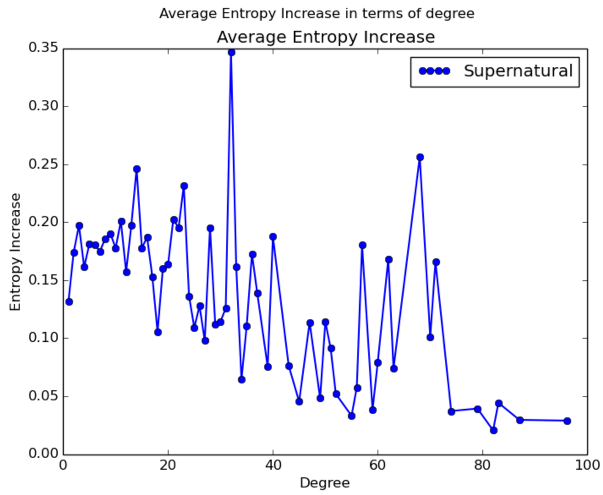

(h) Supernatural

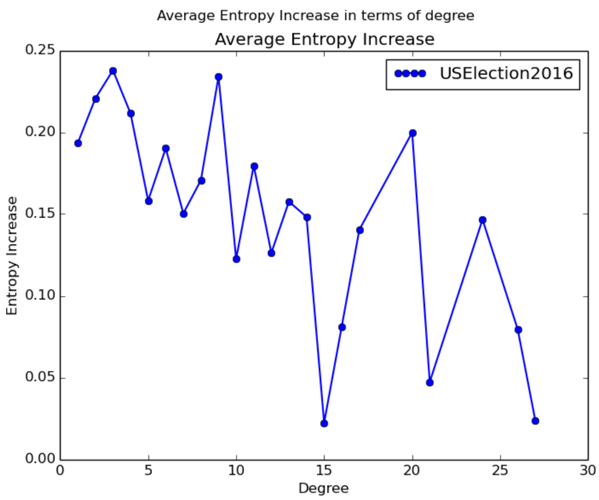

(i) USElection2016

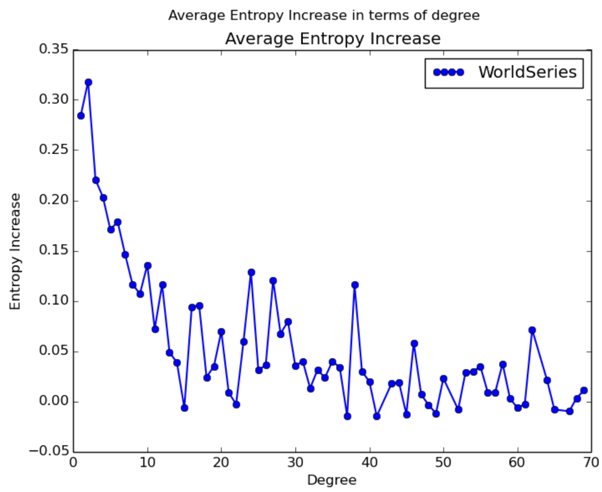

(j) WorldSeries

Figure C.3: Average diversity increase in terms of Degree

Short Biography

Ioannis Kouvatis was born in Ioannina, Greece in 1992. In 2010 he enrolled in the Computer Science and Engineering department in University of Ioannina, where he received his BSc degree in 2016. The same year, in continuation of his studies, he became an MSc student at the same institution, under the supervision of Professor Evaggelia Pitoura. His main research interests are in the area of Data Mining and its applications on Social Networks.