Estimating the Political Orientation of Twitter Users Using Network Embedding Algorithms

Morteza Shahrezaye
University of St.Gallen

Miriam Meckel
University of St.Gallen

Simon Hegelich
Technical University of Munich

Estimating the political orientation of citizens has always been a crucial task in communication as well as political science studies. In this study, advanced network analysis tools are developed to tackle this task. Specifically, using network embedding algorithms, friendship networks are embedded into lower-dimensional Euclidean space while preserving specific topological features of social networks. The resulting embedded vectors are then used to estimate political orientation of Twitter users. It is also shown that these numerical representations can be used to estimate other user traits. The developed tools are applied to a benchmark dataset as well as a dataset developed by the authors. Our model decreased the mean absolute error of the state-of-the-art predictions on the benchmark income dataset by 15%. The developed tools have multiple use cases, for example, studying echo chambers and political communication on OSNs and in marketing campaigns to estimate user's preferences.

Keywords: network embedding, social networks, polarization, microtargeting

INTRODUCTION

The large-scale data produced on Online Social Networks (OSNs) combined with proper computational tools gives computational social scientists the opportunity to study different social phenomena on both individual and group levels (Lazer, Brewer, Christakis, Fowler, & King, 2009; King, 2014). Many studies have tried to estimate users' traits including political orientation (Shahrezaye, Papakyriakopoulos, Serrano, & Hegelich, 2019a; Barberá, 2015), geolocation, age (Schwartz et al., 2013) and gender (Burger, Henderson, Kim, Zarrella, 2011), personal characteristics (Kosinski, Stillwell, & Graepel, 2013), income (Preoţiuc-Pietro, Volkova, Lampos, Bachrach, & Aletras, 2015; Preoţiuc-Pietro, Lampos, & Aletras, 2015; Hasanuzzaman, Kamila, Kaur, Saha, & Ekbal, 2017; Aletras & Chamberlain, 2018), and socioeconomic class by using the data produced on OSNs. This data might include a range of inputs from meta and textual data, friendship networks, and uploaded pictures and videos.

On the other side, studies have shown that the algorithms implemented by OSNs to facilitate communication among online users lead to formation of echo chambers and exacerbation of political polarization (Bakshy, Messing, & Adamic, 2015; Rau & Stier, 2019; Pariser, 2011; Nguyen, Hui, Harper, Terveen, & Konstan, 2014; Bozdag & van den Hoven, 2015). Echo chambers occur when users on OSNs are mostly exposed to opinions that conform to their own worldview (Warner, 2010). Political polarization increases when political discussions on social media platforms increase the distance between the users' political orientations (Sunstein, 2007). There are contradictory efforts to quantitatively prove or disprove the existence of echo chambers and political polarization on OSNs like Twitter and Facebook (Bakshy et al., 2015; Rau & Stier, 2019; Bright, 2017; Barberá, Jost, Nagler, Tucker, & Bonneau, 2015). To study these two effects, one must be able to accurately estimate political orientation of online users.

Friendship networks combined with meta and textual data can help in estimating users' features and traits. Friendship networks contain latent users' features that are formed due to the homophily phenomenon. Homophily refers to the fact that individuals are, on average, socially closer to people with similar opinions than to those with opposing opinions (Lazarsfeld, Merton, et al., 1954). Echo chambers and political polarization unfold on OSNs because of the homophily effect (Sunstein, 2007; Bright, 2018).

In this paper, we focus on the friendship networks in order to estimate political orientation of online users. There are three reasons why we chose friendship networks as the only input to our algorithm. (1) About one third of all OSN users are mostly inactive and they do not have any textual posts on these platforms (Liu, Kliman-Silver, & Mislove, 2014). Therefore, if we used textual data as an input, these users would have to be excluded from the analysis. (2) Because of privacy concerns, OSNs offer limited amount of data to researchers, thus it is reasonable to develop computational tools that do not require huge amounts of data (Perrin, 2018; Isaak & Hanna, 2018). (3) The newly developed network embedding algorithms empower the researchers to efficiently analyze large scale networks.

The embedding algorithm we developed exhausts the topological information of friendship networks and homophily effect while remaining computationally efficient. The following are the contributions of this article:

- To the best of our knowledge, we developed the first network embedding algorithm applicable on friendship networks with the aim of predicting political orientation of online users.
- The suggested algorithm leverages the specific features of social networks to identify political orientation of users even if they do not show any political activity on the platform.
- We applied the method on a benchmark income dataset from Twitter as well as on a large pool of Twitter users active within German politics and improved the state-of-the-art prediction error by 15%.
- The tools developed can be applied to an extensive range of OSN platforms like Twitter, Facebook, and Instagram. In addition, these tools can be used to estimate other users' traits for marketing campaigns.

RELATED WORK

Estimating Political Orientation of Users, Echo Chambers, and Political Polarization

With the introduction of blogs and OSNs, researchers extended the concept political polarization to the digital world (Hill & Hughes, 1998; Adamic & Glance, 2005; Dahlberg, 2007; Papacharissi, 2002). Barberá et al. (2015) analyzed 3.8 million Twitter users and 150 million tweets and found evidence of echo chambers on Twitter. Bright (2017) in a large study of 90 political parties in 23 different countries showed the existence of echo chambers on Twitter. He further discovered that more extremist parties are more prone to form echo chambers.

While some studies blame the algorithms implemented by OSNs for exacerbating the echo chambers and political polarization among online users (Pariser, 2011; Nguyen et al., 2014; Bozdag & van den Hoven, 2015), there are contradicting studies showing that these two effects do not exist on OSNs or if they exist then they do not get stronger due to the algorithms (Bakshy et al., 2015; Rau & Stier, 2019; Dubois &

Blank, 2018). DiMaggio et al. (1996) is among the first to introduce quantitative tools to measure political orientation among groups of people. They defined political polarization as the distance between the political orientations of different people. Therefore, to better study and monitor political polarization on OSNs, it is necessary to efficiently estimate the political orientation of online users. Barberá (2015) employed Bayesian ideal point estimation. Under the assumption that social networks are homophilic, he estimated users' political orientation based on which political actors they followed. The underlying method requires the users to follow a list of politicians, otherwise, the method fails to estimate the political orientation. On the other, Shahrezaye et al. (2019a) applied metric learning algorithms combined with label propagation methods on friendship networks to estimate the political orientation of online users. The advantage of their method is that it can be applied on the majority of online users even if they do not follow any political actors or have not posted any textual data. Shahrezaye et al. (2019b) introduced a projection method for bipartite endorsement networks combined with information theory tools in order to capture the dynamics of political polarization between different political parties on Facebook.

Network Embedding

OSNs offer multiple communication channels to their users. Communication on these platforms leads to creation of social networks in the form of nodes and links between users. Due to the homophily phenomenon and cognitive biases of human behavior, social networks exhibit two distinctive features when compared to other complex biological and technical networks: (1) assortative mixing and (2) local clustering (Newman, 2002). While the former property refers to the positive degree correlation between neighboring nodes, the latter property indicates that social networks consist of many smaller subclusters. These two features make the study of social networks more complex than that of other networks. Any effort to analyze these networks must take these features into consideration.

A phenomenal recent development in the study of networks is the advances in network embedding algorithms. The basic idea is that understanding and analyzing networks in their initial form is a complex task. Therefore, it is useful to transform the network being studied into a low-dimensional Euclidean representation in such a way that preserves initial network features. In other words, the main goal of network embedding algorithms is to find a vector representation of nodes in continuous Euclidean space such that the topology of the network can be recovered from it. Resulting vector representations can then be employed for many high-level tasks including: (1) link prediction (Al Hasan & Zaki, 2011), (2) clustering, and (3) node classification (Neville & Jensen, 2000).

Network embedding algorithms can be categorized into three main classes: (a) factorization-based, (b) random walk based, and (c) deep learning-based (Goyal & Ferrara, 2018). The tool we developed in this study is an instance of random walk-based algorithms. This class of network embedding algorithms is based on DeepWalk algorithm introduced by Perozzi et al. (2014), in which the authors applied multiple random walks of fixed length on each node of the network and applied the famous word2vec model (Mikolov, Sutskever, Chen, Corrado, & Dean, 2013) on them. Grover and Leskovec (2016) introduced the node2vec algorithm as their modification to DeepWalk. This algorithm replaced normal random walks with biased random walks in order to produce a richer sequence. Many other modifications to DeepWalk are implemented with the aim of gaining higher performances based on the network structure and features (Perozzi, Kulkarni, & Skiena, 2016; Yang, Tang, & Cohen, 2015; Li, Zhu, & Zhang, 2016; Yang, Cohen, & Salakhutdinov, 2016; Ahmed et al., 2018).

Aletras and Chamberlain (2018) applied the DeepWalk model on Twitter friendship networks to estimate the socioeconomic attributes of Twitter users. Their embedded representation of users outperformed conventional predictive models. However, that model has multiple shortcomings: (1) It has applied random walks that are tied to the users' IDs. This makes the model not generalizable to new users unless the new users are connected to the same users that have been used in the training process. (2) Users on OSNs might follow certain politicians only to have access to news and not because they really support them. Therefore, "there is no guarantee that similar nodes are surrounded by similar context obtained using random walks on graphs [...]. Proximity does not guarantee similarity." (Ahmed et al., 2018) (3) When applied to social networks, DeepWalk leads to rigid embeddings that are not transferable across networks.

This is because "unlike words in languages that are universal with semantics and meaning independent of the corpus of documents, vertex ids obtained by random walks on graphs are not universal and are only meaningful within a particular graph". (Ahmed et al., 2018) (4) DeepWalk is not computationally scalable to real-life online social networks with millions of nodes.

METHODOLOGY

The main idea behind our algorithm is applying the random walks on node types but not on the node IDs. One can define the type function in multiple ways. To minimize the use of meta data due to privacy concerns and to emphasize the importance of the friendship network's topology, we define node degrees as the type function. Two main benefits of this choice are: (1) While using random walks to generate sequence of neighboring nodes captures the subclusters of social networks, (2) defining node degrees as the type function, captures the assortative mixing feature of social networks.

To define our embedding, we employ Role2Vec model introduced in Nesreen et al. (2018). The friendship network $G = (V, E)$, where V is the set of nodes (users) and $E \subseteq V \times V$ is the set of edges (friendship statuses) is given. If two users $(v_i, v_j) \in E$ are friends, then $(v_i, v_j) \in E$. For each node $\mathrm{v_i} \in \mathrm{V}$ let $\Gamma(\mathrm{v_i})$ be the set of direct neighbors of $\mathrm{v_i}$, and thus $|\Gamma(\mathrm{v_i})|$ is the degree of node $\mathrm{v_i}$. We further define the type of each node as its degree, $\Phi(\mathrm{v_i}) = |\Gamma(\mathrm{v_i})|$. The set of possible node degrees is given as

$$\{\Gamma_1, \ldots, \Gamma_\mathrm{M}\}$$

with $\mathrm{M} \ll |\mathrm{V}|$.

The standard DeepWalk framework has four main elements (Perozzi et al., 2014): (1) γ random walks of length t on each node of the graph G. A sample random walk would be

$$\mathrm{v_0}, \mathrm{v_1}, \ldots, \mathrm{v_t}$$

where $\mathrm{v_i}$ represents the ID of a node. (2) A context function based on the corpus of the random walks defines a set of context nodes, c_i, for each node $\mathrm{v_i} \in \mathrm{V}$. (3) The conditional probability distribution of observing the context nodes in a window of length w given each node in the random walks based on the Skipgram model of word2vec (Mikolov et al., 2013):

$$\forall\, \mathrm{v_i} \in \mathrm{V}: \mathrm{x_{c_i}} | \mathrm{x_i} \sim \mathbb{P}$$

(4) The embedding $\alpha_\mathrm{i} \in \mathbb{R}^\mathrm{d}$ and context vectors β_i as the parameters that are defined for each $\mathrm{v_i} \in \mathrm{V}$.

Assuming the context nodes are conditionally independent, the goal of the DeepWalk is to model $\mathbb{P}[\mathrm{x_{c_i}} | \mathrm{x_i}] = \prod_{\mathrm{j} \in \mathrm{c_i}} \mathbb{P}[\mathrm{x_j} | \mathrm{x_i}]$ with

$$\mathbb{P}\left[\mathrm{x_j} \middle| \mathrm{x_i}\right] = \frac{\exp\{\alpha_\mathrm{i} \cdot \beta_\mathrm{j}\}}{\sum_{v_k \in V} \exp\{\alpha_i \cdot \beta_k\}} \tag{1}$$

where $\cdot$ represents the inner product operator. The objective function of the DeepWalk is then defined as

$$\mathcal{L}(\alpha, \beta) = \sum_{\mathrm{i}=1}^{|\mathrm{V}|} \log \mathbb{P}\left[\mathrm{x_{c_i}} \middle| \mathrm{x_i}\right] \tag{2}$$

Finally, this optimization problem is efficiently solved using the Noise Contrastive Estimation (NCE) algorithm (Mikolov et al., 2013).

The resulting embedding vectors mostly capture the proximity between neighboring nodes. In other words, the nodes that are close to each other in the initial network would possess close vectors in the Euclidean space (Goyal & Ferrara, 2018). To further expand the DeepWalk model, more precisely, to exhaust the topological features of the network and to take the magnitude of node activity into consideration, we use the framework of Nesreen et al. (2018) and suggest replacing the node IDs in the random walks from

$$v_0, v_1, \dots, v_t$$

to the random walks of the degrees as type function or

$$\Gamma\,(v_0), \Gamma(v_1), \dots, \Gamma(v_t)$$

This leads to an embedding model that maps nodes to Euclidean vectors based on the degree distribution of their neighbors (capturing assortative mixing) and their position in the network (capturing subclusters). This model, relates each node degree to the degrees of the corresponding context nodes:

$$\mathbb{P}\left[\Gamma(x_{c_i})\,\middle|\Gamma(x_i)\right] = \prod_{j \in c_i} \mathbb{P}\left[\Gamma(x_j)|\Gamma(x_i)\right] \tag{3}$$

This modification leads to embedding vectors that preserve more information about the network structure. This embedding method is less complex than DeepWalk but still follows Markov chain random walks (Nesreen et al., 2018). One can also introduce multiple attributes alongside the node degrees in order to increase the quality of the embedding vectors.

DATA & RESULTS

Income Dataset

We applied the proposed model on two benchmark datasets. The first dataset maps 5,191 Twitter users to their income level (Preoţiuc-Pietro, Volkova, et al., 2015). Multiple methods have been applied to train models to estimate the income variable (Preoţiuc-Pietro, Volkova, et al., 2015; Preoţiuc-Pietro, Lampos, & Aletras, 2015; Aletras & Chamberlain, 2018; Hasanuzzaman et al., 2017). Preoţiuc-Pietro et al. (2015) used comprehensive feature engineering methods to create different features using the available meta and textual data. Hasanuzzaman et al. (2017) used temporal analysis of tweeting behavior to estimate the income variable. Aletras and Chamberlain (2018) developed two different models. The first model relies only on the friendship network and the second model on the friendship network combined with the textual data.

At the time of our analysis only 5,186 of the Twitter users were still available. We downloaded the complete list of their friends (the users they follow) and formed the friendship network. The complete network had 4,754,888 nodes out of which 5,186 were labeled. To estimate the effect of network size and also to improve computational efficiency, we pruned the network in three steps removing the nodes of degree less than 100, 30 and 10 (Table 1). We set the embedding dimension to 128 and ran 50 walks of length 80 on each node. After generating the embeddings using the networks, we applied a gradient boosting regression with 5-fold cross-validation on the labeled data in order to measure the performance (Table 2). One can observe that our model decreases the MAE about 15% compared to the state-of-the-art model.

TABLE 1
STATISTICS OF THE INCOME NETWORK

	number of nodes	number of edges	remaining labeled nodes
complete network	4,754,888	7,865,568	5,186
threshold at 10 (net10)	62,069	1,273,435	4,442
threshold at 30 (net30)	13,102	542,308	4,407
threshold at 100 (net100)	5,221	177,859	4,216

TABLE 2
MODEL PERFORMANCE ON THE INCOME NETWORK

Preotiuc-Pietro et al. (Preoţiuc-Pietro, Volkova, et al., 2015)		
MAE		**ρ**
9528		0.59
Hasanuzzaman et al. (Hasanuzzaman et al., 2017)		
MAE		**ρ**
10235		0.51
Aletras, Nikolaos and Chamberlain (Aletras & Chamberlain, 2018)		
network	**MAE**	**ρ**
Friendship network	9048	0.62
Friendship network + text	9072	0.64
ours		
network	**MAE**	**ρ**
net10	7764	0.66
net30	8146	0.64
net100	9518	0.52

Political Parties' Dataset

The second dataset we used is the one introduced in Shahrezaye et al. (2019a). This dataset contains a list of 11,293 Twitter users who have retweeted at least three tweets from members of one of the German political parties and no retweets from the members of other parties. We downloaded complete friendship network of these users using Twitter API. Then we followed the same procedure as we did for income data and pruned the networks. We split the data to test and train data and reported out-of-the-bag prediction accuracies (Table 3). Specifically, we trained a 5-fold Radial-Basis SVM model on the training data and measured the performances on the test split. The performance of the test data prediction using the three pruned networks are presented in Table 4.

TABLE 3
STATISTICS OF THE POLITICAL PARTIES' NETWORK

	Number of nodes	**Number of edges**	**Number of training data**	**Number of test data**
threshold at 10 (net10)	103,097	4,374,278	7,638	849
threshold at 30 (net30)	37,604	3,331,524	7,396	822
threshold at 100 (net100)	14,577	2,184,323	6,421	714

TABLE 4
MODEL'S PERFORMANCE ON THE POLITICAL PARTIES' NETWORK

	f1-score (macro average)	**accuracy**
threshold at 10 (net10)	0.64	0.69
threshold at 30 (net30)	0.68	0.74
threshold at 100 (net100)	0.66	0.73

DISCUSSION

In this paper we employed a network embedding algorithm and introduced a new type function to preserve most of the topological information in social networks. The input to the developed model is only the friendship network. Model's performance on two different Twitter datasets is reported. For the income dataset our model decreased the mean absolute error of the state-of-the-art predictions by 15%. For the political parties' dataset, which is a 7-class classification problem, our embeddings gained the accuracy of 74% on the test data.

We argue that the reason for superior performance of the developed model is exhausting the fact that the online social networks contain high levels of homophily: similar users tendto be connected with higher probabilities. Binding the embedding model to node degrees but not node IDs allows the learned embeddings to average the information over the whole network while taking the special features of social networks into consideration. The two direct benefits of this approach are: (1) The resulting embeddings captures the subclusters of social networks, and (2) defining node degrees as the type function, captures the assortative mixing feature of social networks.

This study provides valuable insights for possible future research studies. Firstly, a recent big concern is the effect of fake news, trolls, bots, and automated accounts on political discussion on OSNs (Shahrezaye, Meckel, Steinacker, & Suter, 2020; Papakyriakopoulos, Shahrezaye, Serrano, & Hegelich, 2019). One of the relevant challenges is to classify users to humans/bots or credible/not credible. These questions can be accurately answered using the newly developed method. Secondly, the performance of the model is accomplished using minimal input data and without using any meta or textual data. This raises the question how to address the data-privacy concerns on OSNs. In other words, are the users safe from microtargeting efforts given the current policies or maybe we need new regulations to protect the online users (Papakyriakopoulos, Hegelich, Shahrezaye, & Serrano, 2018).

REFERENCES

Adamic, L.A., & Glance, N. (2005). The political blogosphere and the 2004 US election: Divided they blog. *Proceedings of the 3rd International Workshop on Link Discovery*, (S. 36–43).

Ahmed, N.K., Rossi, R., Lee, J.B., Willke, T.L., Zhou, R., Kong, X., & Eldardiry, H. (2018). *Learning Role-Based Graph Embeddings*.

Al Hasan, M., & Zaki, M.J. (2011). A survey of link prediction in social networks. In *Social Network Data Analytics* (S. 243–275). Springer.

Aletras, N., & Chamberlain, B.P. (2018). Predicting Twitter User Socioeconomic Attributes with Network and Language Information. *Proceedings of the 29th on Hypertext and Social Media* (S. 20–24). New York, NY, USA: ACM. doi:10.1145/3209542.3209577

Bakshy, E., Messing, S., & Adamic, L.A. (2015). Exposure to ideologically diverse news and opinion on Facebook. *Science*, *348*, 1130–1132. doi:10.1126/science.aaa1160

Barberá, P. (2015). Birds of the Same Feather Tweet Together: Bayesian Ideal Point Estimation Using Twitter Data. *Political Analysis*, *23*, 76–91. doi:10.1093/pan/mpu011

Barberá, P., Jost, J.T., Nagler, J., Tucker, J.A., & Bonneau, R. (2015). Tweeting From Left to Right: Is Online Political Communication More Than an Echo Chamber? *Psychological Science*, *26*, 1531-1542. doi:10.1177/0956797615594620

Barnett, G.A. (1974). *Social System Homophily as a Function of Communication*.

Bernstein, B.B. (2003). *Class, codes and control: Applied studies towards a sociology of language* (Bd. 2). Psychology Press.

Bozdag, E., & van den Hoven, J. (2015). Breaking the filter bubble: Democracy and design. *Ethics and Information Technology*, *17*, 249–265.

Bright, J. (1 2018). Explaining the Emergence of Political Fragmentation on Social Media: The Role of Ideology and Extremism. *Journal of Computer-Mediated Communication*, *23*, 17-33. doi:10.1093/jcmc/zmx002

Bright, J. (2017). *Explaining the emergence of echo chambers on social media: The role of ideology and extremism*. SSRN 2839728.

Burger, J.D., Henderson, J., Kim, G., & Zarrella, G. (2011). Discriminating gender on Twitter. *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, (S.1301–1309).

Cappella, J.N. (10 2017). Vectors into the Future of Mass and Interpersonal Communication Research: Big Data, Social Media, and Computational Social Science. *Human Communication Research*, *43*, 545-558. doi:10.1111/hcre.12114

Dahlberg, L. (2007). Rethinking the fragmentation of the cyberpublic: From consensus to contestation. *New Media & Society*, *9*, 827-847. doi:10.1177/1461444807081228

DiMaggio, P., Evans, J., & Bryson, B. (1996). Have American's social attitudes become more polarized? *American Journal of Sociology*, *102*, 690–755.

Dubois, E., & Blank, G. (2018). The echo chamber is overstated: The moderating effect of political interest and diverse media. *Information, Communication & Society*, *21*, 729–745.

Garcia, D., Abisheva, A., Schweighofer, S., Serdült, U., & Schweitzer, F. (2015). Ideological and temporal components of network polarization in online political participatory media. *Policy & Internet*, *7*, 46–79.

Gayo-Avello, D., Metaxas, P.T., & Mustafaraj, E. (2011). Limits of electoral predictions using twitter. *Fifth International AAAI Conference on Weblogs and Social Media*.

Goyal, P., & Ferrara, E. (2018, July). Graph embedding techniques, applications, and performance: A survey. *Knowledge-Based Systems*, *151*, 78–94. doi:10.1016/j.knosys.2018.03.022

Grover, A., & Leskovec, J. (2016). node2vec: Scalable feature learning for networks. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, (S. 855–864).

Hanusch, F., & Nölleke, D. (2019). Journalistic Homophily on Social Media. *Digital Journalism*, *7*, 22-44. doi:10.1080/21670811.2018.1436977

Hasanuzzaman, M., Kamila, S., Kaur, M., Saha, S., & Ekbal, A. (2017). *Temporal orientation of tweets for predicting income of users*.

Hegelich, S., & Janetzko, D. (2016). Are social bots on Twitter political actors? Empirical evidence from a Ukrainian social botnet. *Tenth International AAAI Conference on Web and Social Media*.

Hill, K.A., & Hughes, J.E. (1998). *Cyberpolitics: Citizen activism in the age of the Internet*. Rowman & Littlefield Publishers, Inc.

Isaak, J., & Hanna, M.J. (2018). User Data Privacy: Facebook, Cambridge Analytica, and Privacy Protection. *Computer*, *51*, 56–59.

Kandel, D.B. (1978). Homophily, selection, and socialization in adolescent friendships. *American Journal of Sociology*, *84*, 427–436.

Kendall, K.E., & Paine, S.C. (1995). Political images and voting decisions. *Candidate Images in Presidential Elections*, pp. 19–35.

King, G. (2014). Restructuring the social sciences: reflections from Harvard's Institute for Quantitative Social Science. *PS: Political Science & Politics*, *47*, 165–172.

Kosinski, M., Stillwell, D., & Graepel, T. (2013). Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences*, *110*, 5802–5805.

Lazarsfeld, P.F., & Merton, R.K. (1954). Friendship as a social process: A substantive and methodological analysis. *Freedom and Control in Modern Society*, *18*, 18–66.

Lazer, D., Brewer, D., Christakis, N., Fowler, J., & King, G. (2009). Life in the network: The coming age of computational social. *Science*, *323*, 721–723.

Li, J., Zhu, J., & Zhang, B. (2016). Discriminative deep random walk for network classification. *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics* (Volume 1: Long Papers), (S. 1004–1013).

Liu, Y., Kliman-Silver, C., & Mislove, A. (2014). The tweets they are a-changin': Evolution of Twitter users and behavior. *Eighth International AAAI Conference on Weblogs and Social Media*.

McCombs, M. (2018). *Setting the agenda: Mass media and public opinion*. John Wiley & Sons.

Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. *Advances in Neural Information Processing Systems*, (S. 3111–3119).

Morstatter, F., Pfeffer, J., Liu, H., & Carley, K.M. (2013). Is the sample good enough? Comparing data from twitter's streaming api with twitter's firehose. *Seventh International AAAI Conference on Weblogs and Social Media.*

Neville, J., & Jensen, D. (2000). Iterative classification in relational data. Proc. *AAAI-2000 Workshop on Learning Statistical Models from Relational Data*, (S. 13–20).

Nguyen, T.T., Hui, P-M., Harper, F.M., Terveen, L., & Konstan, J.A. (2014). Exploring the filter bubble: The effect of using recommender systems on content diversity. *Proceedings of the 23rd International Conference on World Wide Web*, (S. 677–686).

Papacharissi, Z. (2002). The virtual sphere: The internet as a public sphere. *New Media & Society*, *4*, 9-27. doi:10.1177/14614440222226244

Papakyriakopoulos, O., Hegelich, S., Shahrezaye, M., & Serrano, J.C. (2018). Social media and microtargeting: Political data processing and the consequences for Germany. *Big Data & Society*, *5*, 2053951718811844.

Papakyriakopoulos, O., Shahrezaye, M., Serrano, J.C., & Hegelich, S. (4 2019). Distorting Political Communication: The Effect of Hyperactive Users in Online Social Networks. *IEEE INFOCOM 2019 - IEEE Conference on Computer Communications Workshops* (INFOCOM WKSHPS), (S. 157-164). doi:10.1109/INFCOMW.2019.8845094

Pariser, E. (2011). *The filter bubble: What the Internet is hiding from you*. Penguin UK.

Perozzi, B., Al-Rfou, R., & Skiena, S. (2014). DeepWalk: Online Learning of Social Representations. *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (S. 701–710). New York, NY, USA: ACM. doi:10.1145/2623330.2623732

Perozzi, B., Kulkarni, V., & Skiena, S. (2016). *Walklets: Multiscale graph embeddings for interpretable network classification*. arXiv preprint arXiv:1605.02115.

Perrin, A. (2018). *Americans are changing their relationship with Facebook*. Pew Research Center.

Preoţiuc-Pietro, D., Lampos, V., & Aletras, N. (2015). An analysis of the user occupational class through Twitter content. *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing* (Volume 1: Long Papers), (S. 1754–1764).

Preoţiuc-Pietro, D., Volkova, S., Lampos, V., Bachrach, Y., & Aletras, N. (2015). Studying user income through language, behaviour and affect in social media. *PloS One*, *10*, e0138717.

Quattrociocchi, W., Scala, A., & Sunstein, C.R. (2016). *Echo chambers on Facebook*. SSRN 2795110.

Rau, J.P., & Stier, S. (2019, September 1). Die Echokammer-Hypothese: Fragmentierung der Öffentlichkeit und politische Polarisierung durch digitale Medien? *Zeitschrift für Vergleichende Politikwissenschaft*, *13*, 399–417. doi:10.1007/s12286-019-00429-1

Rogers, E.M., & Bhowmik, D.K. (1970). Homophily-heterophily: Relational concepts for communication research. *Public Opinion Quarterly*, *34*, 523–538.

Schwartz, H.A., Eichstaedt, J.C., Kern, M.L., Dziurzynski, L., Ramones, S.M., Agrawal, M., . . . Ungar, L.H. (2013). Personality, gender, and age in the language of social media: The open-vocabulary approach. *PloS One*, *8*, e73791.

Shahrezaye, M., Papakyriakopoulos, O., Serrano, J. C., & Hegelich, S. (2019). Estimating the Political Orientation of Twitter Users in Homophilic Networks. *Proceedings of the Symposium Interpretable AI for Well-being: Understanding Cognitive Bias and Social Embeddedness*.

Shahrezaye, M., Papakyriakopoulos, O., Serrano, J.C., & Hegelich, S. (2019). Measuring the Ease of Communication in Bipartite Social Endorsement Networks: A Proxy to Study the Dynamics of Political Polarization. *Proceedings of the 10th International Conference on Social Media and Society* (S. 158–165). New York, NY, USA: ACM. doi:10.1145/3328529.3328556

Shrum, W., Cheek, N.H., Jr., & MacD, S. (1988). Friendship in school: Gender and racial homophily. *Sociology of Education*, pp. 227–239.

Shu, K., Sliva, A., Wang, S., Tang, J., & Liu, H. (2017). Fake news detection on social media: A data mining perspective. *ACM SIGKDD Explorations Newsletter*, *19*, 22–36.
Simpson, I.H., & England, P. (1981). Conjugal work roles and marital solidarity. *Journal of Family Issues*, *2*, 180–204.
Sunstein, C.R. (2007). *Republic.com 2.0*. Princeton University Press. Von. Retrieved from http://www.jstor.org/stable/j.ctt7tbsw abgerufen
Tan, Q., Liu, N., & Hu, X. (4 2019). Deep Representation Learning for Social Network Analysis. *Frontiers in Big Data*, *2*. doi:10.3389/fdata.2019.00002
Varol, O., Ferrara, E., Davis, C.A., Menczer, F., & Flammini, A. (2017). Online human-bot interactions: Detection, estimation, and characterization. *Eleventh international AAAI Conference on Web and Social Media*.
Warner, B.R. (2010). Segmenting the electorate: The effects of exposure to political extremism online. *Communication Studies*, *61*, 430–444.
Yamaguchi, K. (1990). Homophily and social distance in the choice of multiple friends an analysis based on conditionally symmetric log-bilinear association models. *Journal of the American Statistical Association*, *85*, 356–366.
Yang, Z., Cohen, W.W., & Salakhutdinov, R. (2016). *Revisiting semi-supervised learning with graph embeddings*. arXiv preprint arXiv:1603.08861.
Yang, Z., Tang, J., & Cohen, W. (2015). *Multi-Modal Bayesian Embeddings for Learning Social Knowledge Graphs*.