

# Joint Non-negative Matrix Factorization for Learning Ideological Leaning on Twitter

Preethi Lahoti  
Max Planck institute for Informatics  
Saarbruecken, Germany  
plahoti@mpi-inf.mpg.de

Kiran Garimella  
Aalto University  
Helsinki, Finland  
kiran.garimella@aalto.fi

Aristides Gionis  
Aalto University  
Helsinki, Finland  
aristides.gionis@aalto.fi

## ABSTRACT

People are shifting from traditional news sources to online news at an incredibly fast rate. However, the technology behind online news consumption promotes content that confirms the users' existing point of view. This phenomenon has led to polarization of opinions and intolerance towards opposing views. Thus, a key problem is to model information filter bubbles on social media and design methods to eliminate them. In this paper, we use a machine-learning approach to learn a *liberal-conservative* ideology space on Twitter, and show how we can use the learned latent space to tackle the filter bubble problem.

We model the problem of learning the *liberal-conservative* ideology space of social media users and media sources as a constrained non-negative matrix-factorization problem. Our model incorporates the social-network structure and content-consumption information in a joint factorization problem with shared latent factors. We validate our model and solution on a real-world Twitter dataset consisting of controversial topics, and show that we are able to separate users by ideology with over 90% *purity*. When applied to media sources, our approach estimates ideology scores that are highly correlated (Pearson correlation 0.9) with ground-truth ideology scores. Finally, we demonstrate the utility of our model in real-world scenarios, by illustrating how the learned ideology latent space can be used to develop exploratory and interactive interfaces that can help users in diffusing their information filter bubble.

## 1 INTRODUCTION

Social media and the web have provided a foundation where users can easily access diverse information from around the world. However, over the years, various factors, such as user homophily (social network structure), and algorithmic filtering (e.g., news feeds and recommendations) have narrowed the content that a user consumes. As an example, imagine two users of opposite ideological stances (liberal and conservative). Though the two users may be looking at the same topic (e.g., a presidential debate), they might be seeing completely different viewpoints due to the diverse network surrounding and different content sources they get their information from. Consequently, users on different ends of the ideological spectrum live in their own information bubbles [32], oblivious to the views on the other side and creating their own world-view of truth. This phenomenon has led to the polarization of viewpoints, intolerance towards opposing views, and ideological segregation [37].

Many studies suggest that increasingly users live in their echo chambers [36] and polarization of the public has intensified [11].

In this paper, we propose a principled approach to infer the *ideological stances* (also known as *ideology* or *polarity* or *leaning*) of both the *users* in a social network, and the *media sources* that provide news content in the network. Our approach is based on a non-negative matrix-factorization model, which jointly decomposes the social network of users and the content they consume in a *shared latent space*. Learning the ideological stances of social media users and content sources is an important step in building useful tools that can make users aware of their informational bias and consequently, help in reducing those biases, thus mitigating the increasing polarization in the society.

Existing approaches to identifying the ideological leaning of users either (i) require large amounts of manually annotated data [10, 33]; or (ii) consider only the structure of social ties [4] or user interactions [16, 38]; or (iii) analyze only the content shared by users [34]. Each of these families of approaches has its own limitations. First, obtaining manual annotations is both expensive and time consuming. Second, when using only the structure of social ties or interactions, it is assumed that users with the same ideology are more likely to interact with or follow each other. However, social networks are extremely sparse and noisy, and users with the same ideological leaning may never interact with or follow each other, while users with different ideological leaning may still interact due to different reasons. Last, inferring ideological leaning by using only content is a challenging task and prone to errors due to the inherent complexity of natural-language understanding.

To overcome these problems, we propose an unsupervised approach that uses simultaneously the network structure and the information about content shared by users. We are motivated by the observation that a user's ideological stance on a topic depends on both the surrounding network structure as well as the content sources that users consume their information from. Thus, we exploit the inherent connection between the two data types: Users are not only more likely to interact with or follow like-minded users, but also to share content of aligned ideological leaning. Taking into account the full available information results in more accurate estimation of ideological leaning both for users and content sources.

In addition, we formulate the task of learning ideological leaning as a joint matrix factorization problem in which users and sources are represented in a shared latent space. This formulation allows us to identify the relationship between data points of the two types, which is particularly useful for visualizing social-media users and sources in a common space, and building applications for exploring the ideological landscape in one's media neighborhood, or making

recommendations to escape the filter bubble. In fact, we demonstrate concretely how to use the learned ideology latent space to develop exploration and recommendation tools (Section 6).

Experiments comparing our approach to the state of the art show the benefit of such a principled joint approach. Our method is able to separate users into ideological clusters with over 90% purity. When applied to media sources, our approach estimates ideology scores that are highly correlated with ground-truth ideology scores.

All our analysis is done on Twitter, though the methods generalize to any other social network. In the rest of the paper, for the sake of clarity we use Twitter-specific nomenclature (e.g., retweets, follow, etc.).

In summary, our contributions are as follows:

- We present a principled approach to jointly compute the ideology scores of both users and content sources by formulating the problem as a joint constrained matrix-factorization task. This formulation allows us to cluster the two different data types (“user” and “source”) at a time and obtain the relationship between them. We apply our learned ideological latent space to the problems of ‘who’ and ‘what’ to recommend to users so as to reduce polarization and diffuse their information filter bubble.
- To the best of our knowledge, this is the first work to compute ideological stances for *both* users and content sources, simultaneously, based on both the network structure and the users’ interaction with sources. Note that other approaches have used network and content information together to learn ideology-leaning scores of users [29]. However, in this work we use both data types to learn ideology-leaning scores about both of them in a common latent space.
- We provide an extensive experimental evaluation, presenting both qualitative and quantitative results on real-world Twitter data. Our method shows promise when compared to existing state of the art approaches.

## 2 RELATED WORK

In this paper, we propose an approach that can identify the ideological leaning of users on Twitter. We define ideology based on the policy dimension that articulates a user’s political preference. Our definition is inspired by work in political-science literature, such as work by Bafumi et al. [2], where ideology is defined as “a line whose left end is understood to reflect an extremely liberal position and whose right end corresponds to extreme conservatism.”

**Estimation of user ideology.** Traditionally, the most common sources for estimating ideology comprised of behavioral data generated from roll call votes [35], co-sponsorship records [1], or political contributions [7]. These datasets were often only available for the political elite, like members of congress, and hence getting such estimates for a large population of ordinary citizens was difficult, if not impossible.

With the proliferation of social media platforms, behavioral data started being available at an individual level and researchers have tried to use such data for identifying political ideology for social media users at scale. Initial work started with supervised methods [10, 33] for predicting a (binary) political alignment of

users on Twitter. Though these works report accuracies over 90%, Cohen and Ruths [9] warns about the limitations of such approaches and their dependence on politically active users.

Unsupervised approaches have also been proposed, mainly based on the structure of user interests [24], social connections [4], and interactions [6, 16, 38]. The main idea behind these methods is that users typically either surround themselves (follow/friend) with other users who are similar in their ideology (homophily), or interact with others (retweet/like) similar to them.

Perhaps the closest approach to this paper is the work by Lu et al. [29] (BIASWATCH), who seek to identify the *bias* of a user on a topic by combining their retweet and content networks, where a content network is obtained based on the similarity of users tweets. BIASWATCH, however, assumes the presence of a set of labeled *bias anchors* (seed hashtags), making it not completely unsupervised. Second, fusing the content and retweet networks is somewhat arbitrary, since there is no common underlying principle that holds the two networks together and hence a graph that results from such a merger contains different types of edges (multigraph) simply merged together. We compared our approach with BIASWATCH in Section 5 and show that our method outperforms their bias scores.

**Estimation of source ideology.** Polarization and bias of media outlets has existed long before the time of internet, however, data availability makes it easier to study and quantify nowadays. Mitchell et al. [30] study the media habits of American public using a large scale survey and show how fractured the media production and consumption have become in the recent years. Groseclose and Milyo [22] propose a method to estimate the ideology of various media sources by comparing the number of citations to think tanks and policy groups to those of Congress members. For a complete survey of methods on measuring ideology of media and media bias, please refer to Groeling [21].

**Reducing polarization.** The problem of reducing polarization and diffusing information bubbles on social media has been tackled before. Most papers approach the problem by recommending content outside a user’s filter bubble. Studies have looked at what to recommend [20], how to recommend [26, 27] and to whom to recommend such content [17, 19].

Based on the papers discussed above, we make the following observations:

- a. Most existing approaches only consider one dimension (content or network) for computing user ideology. We propose a principled approach that can compute a user’s ideology by taking into account *both* content and network.
- b. We compute the ideology of users and sources simultaneously. No existing approach, to the best of our knowledge, proposes such a method to do it simultaneously.
- c. Most approaches for reducing polarization do not take a user’s *choice* into account. Our methods (in Section 6) can help users diffuse their information bubbles based on their own choices.

## 3 DATA MODEL AND PROBLEM FORMULATION

In this section we discuss the problem setting, introduce the notation, and formally define the problem we consider.

### 3.1 Data model

**The social graph.** We consider a social graph of Twitter users  $G = (U, E, w)$ , where  $U$  represents the set of users,  $E$  is the set of edges representing social interactions between the users, and  $w : E \rightarrow \mathbb{R}$  is a weighting function that assigns real-valued weights to each edge in  $G$ .

We also represent the graph  $G$  by its adjacency matrix  $A \in \mathbb{R}^{n \times n}$ . In particular, it is  $A(u, v) = w_i$  for each edge  $e_i = (u, v, w_i) \in E$ , and  $A(u, v) = 0$  if there is no edge between  $u$  and  $v$ . Since the social graph  $G$  may not represent a symmetric relation, the matrix  $A$  is not necessarily symmetric.

To define the edges of the graph  $G$  one can consider social links between users based on their retweet and follow networks. Two commonly-used options are to consider an edge  $e_i = (u, v, w_i)$  if (i) users  $u$  and  $v$  follow a common set of users, and  $w_i$  is the *number of common users*; or if (ii) user  $u$  retweets user  $v$ , and  $w_i$  is the *number of retweets*. In the first case the social graph is symmetric, while in the second case it is not.

**Content sources.** Users join a social network, such as Twitter, to discover interesting content and promote content they endorse. We model the presence of content in the network by considering a set of content sources  $S$ . The set  $S$  represents a set of items that are shared in the network. The number of sources is denoted by  $|S| = m$ . We associate with each user  $u$  the subset of sources in  $S$  that the user shares (i.e., posts) in the network. Thus the overall user activity in the network is denoted by a matrix  $C \in \mathbb{R}^{n \times m}$ , where  $C(u, s)$  denotes the number of times that user  $u \in U$  shares source  $s \in S$ .

We derive the content features (i.e., items of the set  $S$ ) based on users' tweets. Specifically, we experiment with two variants of feature sets: (i) extracted urls from the tweets; and (ii) hostnames of urls extracted from the tweets. In other words, in the latter case, the content is aggregated by their source of authorship (i.e., various news media channels).

For the presentation of our method, we also need the following definitions.

**Affinity matrix.** Given a matrix  $X$ , we define the *affinity matrix of the rows of  $X$*  to be  $X_{rn}X_{rn}^T$ , where  $X_{rn}$  is row normalized  $X$ . That is, the affinity matrix of rows of  $X$  is formed by the cosine similarities of all pairs of rows of  $X$ . Similarly, we define the *affinity matrix of the columns of  $X$*  to be  $X_{cn}^T X_{cn}$ , where  $X_{cn}$  is column normalized  $X$ .

**Laplacian matrix.** Given a square and symmetric matrix  $X$ , we consider its Laplacian to be the matrix  $L = D - X$ , where  $D$  is a diagonal matrix whose  $(i, i)$  entry is the  $i$ -th row-sum of  $X$ .

### 3.2 Motivation of the approach

Our goal in this paper is to learn the ideological leaning of users and content sources on Twitter. The underlying motivation is that learning the ideological leaning of Twitter users and sources is an important step in building useful tools that can help users perceive their informational bias and consequently improve their news diet.

A simple approach to identify ideological leaning of users is to consider only the social graph and apply one of the many community-detection algorithms. The intuition here is that users with the same ideological leaning are more likely to interact with each other, and

thus, to form graph communities. The drawback of this simple approach is that community detection on real-world social graphs is an extremely challenging task, due to sparsity and overlapping communities. For instance, two users  $u$  and  $v$  can have the same ideological leaning even though they do not have any social interaction, or they have different topical interests.

Similarly, one simple approach to attempt identifying the ideological leaning of sources is to cluster them using deep-NLP or semantic-analysis techniques. However, this approach is again prone to errors due to the inherent complexity of the text-analysis task. Furthermore, it ignores the rich user information about how content is shared in the social network.

In contrast to these simple techniques, which rely on one-sided clustering of either users or content, with no association between them, the proposed approach seeks to *jointly* learn the ideological leaning of users and content sources.

Combining different data types in a unified learning framework has several advantages. First, we exploit the inherent connection between the two data types: users are not only more likely to interact with, or follow, like-minded users, but also to share content of aligned ideological leaning. Consequently, taking into account the full information on social structure and content will result in better clustering performance. Furthermore, considering both data types simultaneously allows us to learn the ideological leanings of users and content sources in a *shared latent space*. This means that not only do we separate users and sources into ideological clusters, but we also identify the *relationship* between the clusters of the two data types. The applications we present in the Section 6 rely heavily on the ability to represent users and content sources in a shared latent space.

### 3.3 Problem formulation

To learn the latent space of the input data (users and sources) we use non-negative matrix factorization (NMF) techniques. In particular, we propose a joint matrix factorization formulation, which exploits the duality between user and source clustering.

First, assume that the ideological leaning is represented by  $k$  factors (dimensions). Our model is described by two components. The first component, represented by a  $n \times k$  matrix  $U$ , captures user information. In particular, the entry  $(u, i)$  of the matrix  $U$  represents the degree to which user  $u$  aligns with ideology factor  $i$ .

The second component, represented by a  $m \times k$  matrix  $V$ , captures source information: the entry  $(s, i)$  of the matrix  $V$  represents the degree to which source item  $s$  aligns with ideology factor  $i$ .

To determine the user and source ideology clusters, we decompose the two input matrices  $A$  and  $C$ , using the components  $U$  and  $V$  as latent factors. For decomposing  $A$  and  $C$  so as to capture ideology clusters we require the following constraints.

#### Partitioning constraints:

1. Users in a user-cluster interact with each other more often than with users outside the cluster.
2. Users in the same user-cluster post content from the same content-cluster.
3. Content in the same content-cluster is posted by users in the same user-cluster.

#### Co-partitioning constraints:

4. Users in a user-cluster share more articles from their corresponding content-cluster than from other content-clusters.
5. Content in a content-cluster is shared by more users from their corresponding user-cluster than from other user-clusters.

**Non-negative matrix factorization for co-clustering:** For a given input data matrix  $\mathbf{X}$ , the bi-orthogonal non-negative 3-factor decomposition (ONMTF) [13], formulated as  $\mathbf{X} \approx \mathbf{W}\mathbf{H}\mathbf{Z}^T$ , provides a good framework for simultaneously clustering the rows and the columns of  $\mathbf{X}$ . The left factor  $\mathbf{W}$  provides a clustering of the rows of  $\mathbf{X}$ , while the right matrix  $\mathbf{Z}$  provides a clustering of the columns of  $\mathbf{X}$ . The middle factor  $\mathbf{H}$  provides association between the clusters and additional degrees of freedom.

**PROBLEM 1.** (ONMTF [13]) *Given an  $n \times m$  matrix  $\mathbf{X}$ , and integer  $k$ , with  $k \ll n, m$ , find non-negative matrices  $\mathbf{W}$ ,  $\mathbf{H}$ , and  $\mathbf{Z}$ , of respective dimensions  $n \times k$ ,  $k \times k$ , and  $k \times m$ , so as to*

$$\begin{aligned} & \text{minimize } \|\mathbf{X} - \mathbf{W}\mathbf{H}\mathbf{Z}^T\|_F^2 \\ & \text{subject to } \mathbf{W} \geq 0, \mathbf{H} \geq 0, \mathbf{Z} \geq 0, \\ & \text{and } \mathbf{W}\mathbf{W}^T = \mathbf{Z}\mathbf{Z}^T = \mathbf{I}. \end{aligned}$$

As shown by Ding et al. [13], the bi-orthogonality constraints provide an interpretation of the ONMTF problem as a simultaneous clustering of the rows and columns of the input data matrix.

In our approach we use ONMTF to decompose the input matrices (social-graph matrix  $\mathbf{A}$  and user-source matrix  $\mathbf{C}$ ) using the latent factors  $\mathbf{U}$  and  $\mathbf{V}$ .

First, the user-source matrix is decomposed by  $\mathbf{C} \approx \mathbf{U}\mathbf{H}_s\mathbf{V}^T$ , subject to orthogonality constraints for  $\mathbf{U}$  and  $\mathbf{V}$ .

For the social-graph matrix  $\mathbf{A}$ , as both the rows and columns represent users, we equate the left and right factors and require  $\mathbf{A} \approx \mathbf{U}\mathbf{H}_u\mathbf{U}^T$ , subject to orthogonality constraints for  $\mathbf{U}$ .

Note that as  $\mathbf{H}_u$  is not necessarily symmetric, the decomposition  $\mathbf{U}\mathbf{H}_u\mathbf{U}^T$  can produce a non-symmetric matrix. Furthermore, the formulation can capture link transitivity [40]: consider a path  $u_i \rightarrow u_k \rightarrow u_j$ . A non-symmetric factorization  $\mathbf{A} \approx \mathbf{X}\mathbf{Z}^T$  represents the values of  $\mathbf{A}$  as links from set of users to a set of objects, say,  $O = \{o_i\}$ . Hence, it would split the path ( $u_i \rightarrow u_k \rightarrow u_j$ ) into two parts  $u_i \rightarrow o_k$  and  $u_k \rightarrow o_j$ , which is a misinterpretation of the original path. Whereas,  $\mathbf{A} \approx \mathbf{U}\mathbf{H}_u\mathbf{U}^T$  considers the links to be amongst the same set of objects. Hence, the transitive link  $u_i \rightarrow u_k \rightarrow u_j$  is correctly captured by the latent factors in  $\mathbf{U}$ .

**Combining link and content information** is achieved by using a common latent factor  $\mathbf{U}$  and formulating a joint factorization problem asking to minimize

$$\|\mathbf{A} - \mathbf{U}\mathbf{H}_u\mathbf{U}^T\|_F^2 + \|\mathbf{C} - \mathbf{U}\mathbf{H}_s\mathbf{V}^T\|_F^2, \quad (1)$$

subject to non-negativity of  $\mathbf{U}$ ,  $\mathbf{H}_u$ ,  $\mathbf{V}$ ,  $\mathbf{H}_s$ , and orthogonality of  $\mathbf{U}$  and  $\mathbf{V}$ . This approach is inspired by the formulation of Zhu et al. [40] for classifying web-pages by exploiting both content and link information.

**Graph regularization.** From a geometric perspective, a dataset can be viewed as a set of data points on a continuous manifold. The task of clustering is to find these intrinsic manifolds in the data. However, the clustering formulation of ONMTF fails to consider this geometric structure in the data. To address this problem, Cai et al. [8] introduced a graph-regularized NMF based on the *manifold*

*assumption* that if two data points  $\mathbf{x}_i, \mathbf{x}_j$  are close in the input data matrix  $\mathbf{X}$ , their projections  $\mathbf{u}_i$  and  $\mathbf{u}_j$  in the new basis  $\mathbf{U}$  are also close. This is formulated by seeking to minimize

$$\frac{1}{2} \sum_{i,j} \|\mathbf{u}_i - \mathbf{u}_j\|^2 \mathbf{W}_{ij} = \text{tr}(\mathbf{U}^T \mathbf{L} \mathbf{U}),$$

where  $\mathbf{W}$  is the affinity matrix of rows of  $\mathbf{X}$  (see Section 3.1), and  $\mathbf{L}$  is the laplacian of  $\mathbf{W}$ .

Motivated by the duality between row and column manifolds, Gu and Zhou [23] proposed a dual-manifold regularized co-clustering method as a decomposition of an input matrix  $\mathbf{X}$ , by asking to minimize

$$\|\mathbf{X} - \mathbf{W}\mathbf{H}\mathbf{Z}^T\|_F^2 + \alpha \cdot \text{tr}(\mathbf{W}^T \mathbf{L}_w \mathbf{W}) + \beta \cdot \text{tr}(\mathbf{Z}^T \mathbf{L}_z \mathbf{Z}),$$

subject to non-negativity constraints of the factor matrices. Here  $\mathbf{L}_w$  and  $\mathbf{L}_z$  are the Laplacians on the affinity matrices of rows and columns of  $\mathbf{X}$ , respectively. As before, the matrix  $\mathbf{H}$  captures the association between row and column clusters.

We extend our joint factorization model (Equation 1) by including dual-graph regularization constraints on users and sources. We refer to this problem by IFD, for *ideology factor decomposition*. The formal problem definition is the following.

**PROBLEM 2 (IFD).** *Given a user-user social matrix  $\mathbf{A}$  of dimension  $n \times n$ , a user-source matrix  $\mathbf{C}$  of dimension  $n \times m$ , an integer  $k$ , with  $k \ll n, m$ , and regularization parameters  $\alpha$  and  $\beta$ , find factors  $\mathbf{U}$ ,  $\mathbf{H}_u$ ,  $\mathbf{V}$ ,  $\mathbf{H}_s$ , of dimensions  $n \times k$ ,  $k \times k$ ,  $m \times k$ ,  $n \times k$ , respectively, so as to*

$$\begin{aligned} & \text{minimize } \|\mathbf{A} - \mathbf{U}\mathbf{H}_u\mathbf{U}^T\|_F^2 + \|\mathbf{C} - \mathbf{U}\mathbf{H}_s\mathbf{V}^T\|_F^2 \\ & \quad + \alpha \cdot \text{tr}(\mathbf{U}^T \mathbf{L}_u \mathbf{U}) + \beta \cdot \text{tr}(\mathbf{V}^T \mathbf{L}_s \mathbf{V}) \quad (2) \\ & \text{subject to } \mathbf{U} \geq 0, \mathbf{H}_u \geq 0, \mathbf{V} \geq 0, \mathbf{H}_s \geq 0, \\ & \text{and } \mathbf{U}\mathbf{U}^T = \mathbf{V}\mathbf{V}^T = \mathbf{I}, \end{aligned}$$

where  $\mathbf{L}_u$  and  $\mathbf{L}_s$  are the Laplacians of the affinity matrices of rows (users) and columns (sources) of  $\mathbf{C}$ , respectively.

In the IFD problem we factorize  $\mathbf{A}$  and  $\mathbf{C}$  jointly based on the dual manifold assumption i.e., both users and content share the same latent space and the cluster labels of users are smooth with respect to the content manifold, while the cluster labels of content are smooth with respect to the user manifold. To apply these manifold constraints, we consider affinity matrices for users and content sources. While there are many ways to construct such an affinity matrices, in our experiments we consider the affinity matrices of the rows (for users) and columns (for sources) of the  $\mathbf{C}$  matrix.

Finally, it is instructive to review how IFD addresses the *partitioning constraints* 1–3 and *co-partitioning constraints* 4–5: the bi-orthogonality constraints on the tri-factorization problem provide a clustering interpretation for the rows and columns of the input matrices, while the manifold constraints provide a geometric interpretation to the discovered latent space. The correspondance between the user clusters and source clusters is achieved by using a shared latent factor.

### 3.4 Solving the optimization problem

Following the standard theory of constrained optimization, we solve Problem 2 by introducing Lagrangian multipliers  $\Lambda$  (a symmetric

matrix of size  $k \times k$ ) and minimizing the Lagrangian function

$$\begin{aligned} L = & \|A - U\mathbf{H}_u\mathbf{U}^T\|_F^2 + \|C - U\mathbf{H}_s\mathbf{V}^T\|_F^2 \\ & + \alpha \cdot \text{tr}(\mathbf{U}^T\mathbf{L}_u\mathbf{U}) + \beta \cdot \text{tr}(\mathbf{V}^T\mathbf{L}_s\mathbf{V}) \\ & \text{tr}(\Lambda(\mathbf{U}\mathbf{U}^T - \mathbf{I})) + \text{tr}(\Lambda(\mathbf{V}\mathbf{V}^T - \mathbf{I})). \end{aligned} \quad (3)$$

We can compute the gradient of  $L$  with respect to  $\mathbf{U}$ ,  $\mathbf{V}$ ,  $\mathbf{H}_u$ , and  $\mathbf{H}_s$ . A locally-optimal solution for Problem 2 can be found using an iterative-update algorithm, similar to the one proposed by Ding et al. [13].

The update rules are as follows:

$$\mathbf{U} \leftarrow \mathbf{U} \cdot \sqrt{\frac{A\mathbf{U}\mathbf{H}_1^T + C\mathbf{V}\mathbf{H}_3^T + \alpha \cdot S_u\mathbf{U}}{U\mathbf{H}_1\mathbf{U}^T U\mathbf{H}_1^T + U\mathbf{H}_3\mathbf{V}^T V\mathbf{H}_3^T + \alpha D_u\mathbf{U} + U\lambda_u}} \quad (4)$$

$$\mathbf{V} \leftarrow \mathbf{V} \cdot \sqrt{\frac{C^T U\mathbf{H}_3 + \beta \cdot S_d\mathbf{V}}{\beta \cdot D_d\mathbf{V} + V\mathbf{H}_3\mathbf{U}^T U\mathbf{H}_3 + V\lambda_v}} \quad (5)$$

$$\mathbf{H}_1 \leftarrow \mathbf{H}_1 \cdot \sqrt{\frac{U^T A\mathbf{U}}{U^T U\mathbf{H}_1\mathbf{U}^T U}} \quad (6)$$

$$\mathbf{H}_3 \leftarrow \mathbf{H}_3 \cdot \sqrt{\frac{U^T C\mathbf{V}}{U^T U\mathbf{H}_3\mathbf{V}^T V}} \quad (7)$$

where

$$\begin{aligned} \lambda_u &= U^T A\mathbf{U}\mathbf{H}_1^T + U^T C\mathbf{V}\mathbf{H}_3^T - \alpha U^T \cdot L_u\mathbf{U} - \mathbf{H}_1\mathbf{U}^T U\mathbf{H}_1^T - \mathbf{H}_3\mathbf{V}^T V\mathbf{H}_3^T \\ \lambda_v &= V^T C^T U\mathbf{H}_3 - \beta V^T \cdot L_d\mathbf{V} - \mathbf{H}_3\mathbf{U}^T U\mathbf{H}_3 \end{aligned}$$

## 4 ESTIMATING IDEOLOGICAL LEANING

In this section we discuss how to use the latent factors  $\mathbf{U}$  and  $\mathbf{V}$  computed using IFD in order to estimate ideological leaning scores for Twitter users and media channels (sources). Our approach utilizes the probabilistic model of NMF factorizations, discussed in Yoo and Choi [39], to derive a probabilistic interpretation of latent factors. We also discuss how to derive a hard cross-ideological separation (hard clustering) from the latent factors.


As proposed in the literature [25, 39], the latent factors  $\mathbf{U}$  and  $\mathbf{V}$  have the following probabilistic interpretation:

- the entry  $(i, \ell)$  of matrix  $\mathbf{U}$  indicates the degree to which user  $i$  belongs to the user-cluster  $\ell$ ; and
- the entry  $(j, \ell)$  of matrix  $\mathbf{V}$  indicates the degree to which media source  $j$  belongs to the content-cluster  $\ell$ .

In the context of our problem setting, we are interested in identifying two main ideologies, *liberal* and *conservative*, and thus, we set the number of latent dimensions equal to 2 ( $k = 2$ ). It follows that each user and each media source are represented by a 2-dimensional vector  $(x, y)$  in the latent space — a row of  $\mathbf{U}$  for users, and a row of  $\mathbf{V}$  for media sources. Thus, user and source ideology hard clusters are derived as  $\arg \max U_{ij}$  and  $\arg \max V_{ij}$ , respectively. To estimate a single score for users and media sources, we compute the angle of line defined by the center of origin and the latent vector  $(x, y)$ , normalized to be between 0 and 1, i.e.,<sup>1</sup>

$$i(x, y) = \frac{\theta}{\pi/2} = \frac{\arctan(y/x)}{\pi/2}. \quad (8)$$

<sup>1</sup>Recall that  $\arctan(0) = 0$  and  $\lim_{z \rightarrow \infty} \arctan(z) = \pi/2$ .



(a) original latent space (b) transformed coordinate space

**Figure 1: Projection of a subset data points in the learned ideology latent space and the transformed ideology-popularity coordinate space.**

We also compute the magnitude of the latent vector  $(x, y)$

$$\rho(x, y) = \sqrt{x^2 + y^2}, \quad (9)$$

which, as can be shown easily, represents the popularity of the corresponding user or source — in particular, it is correlated with the number of re-tweets and follows of a given user, and the number of tweets containing a given source. Figure 1 visualizes a subset of real data points projected in the original latent space and their transformation to the corresponding ideology/popularity co-ordinate space according to the aforementioned computations.

In summary, given a user or media source represented by a latent vector  $(x, y)$ , we can estimate a single ideological leaning score by Equation (8), as well as its popularity by Equation (9). When using more than 2 dimensions for the latent ideology space ( $k > 2$ ) it is not possible to estimate ideology scores with a single number, but we can still handle the user and source representations by standard vector operations.

## 5 EXPERIMENTS

### 5.1 Dataset

Our dataset is collected using Twitter’s streaming API from 2011 to 2016, by filtering for keywords related to three popular controversial topics: *gun control*, *abortion* and *obamacare*. We use the list of keywords proposed by Lu et al. [29] to filter the tweets related to these topics. We only consider users who tweeted about all three topics at least once, obtaining a set of  $n = 6391$  users, and collect all their tweets, which gave us 19 million tweets.

As discussed in Section 3, we consider two variants of matrix  $\mathbf{A}$  (*retweet* and *follow*) and two variants of matrix  $\mathbf{C}$  (*urls* and *hostnames*). We observe that the *follow* user-user matrix  $\mathbf{A}$  along with *hostname* user-content matrix  $\mathbf{C}$  gives significantly superior results compared to all other variants, thus, all subsequent results use that variant. We omit results with the other variants due to lack of space.

### 5.2 Ground truth

**Ideology scores for sources:** We collect ground truth for news-media channels from multiple studies in the literature: (i) 500 most shared news domains on Facebook [3] (ii) 100 most visited domains in Bing toolbar [15] and (iii) 27 domains from an offline survey and

webpage visit data [18]. Each of these scores roughly measures the fraction of views/shares/clicks by a conservative user. We map all scores in the  $[0, 1]$  range, 1 being conservative. For the domains listed in multiple lists we compute the ideology by averaging the scores. We also remove domains that are not necessarily news sources (e.g., *wikipedia.org*, *reddit.com*, etc.). In total, we collect 559 news domains with ground-truth ideology scores. We refer to this dataset as CONTENT ground truth.

**Ideology scores for Twitter users:** We use two different ground-truth scores for users: (i) BARBERA: ideology score estimated by Barberá et al. [5], which applies Bayesian ideal point estimate on nearly 12 million Twitter users, and (ii) AVG\_CONTENT: average ground-truth ideology scores of the sources tweeted by the user.

**Popularity scores for sources:** We use the aggregated number of tweets about each news media channel in the collected data set as a proxy for the popularity of the source.

**Popularity scores for users:** Since the collection of users is a random set of people on user, we do not have any ground truth for popularity of Twitter users.

### 5.3 Baseline algorithms

We compare our method with three types of methods for ideology detection: network-only, content-only, and a combination of network and content.

**Network-only:** We consider two types of network-only methods: (i) NMF-based methods that can provide a continuous ideology score for a user between 0 and 1; and (ii) other methods that only produce binary labels for ideology (a user is either liberal or conservative). We use symmetric NMF (NMF-SYMM) [13], a 3-factor NMF shown to be equivalent to normalized-cut spectral clustering [12], RETWEET a method based on partitioning the retweet graph [16] and FOLLOW a graph partitioning approach on the follow network. In order to construct a source-source relationship matrix, we use  $C^T C$ . It is noteworthy that network-only methods perform only one-side clustering — one data type at a time. Hence, we need to apply the methods separately for users and content sources. As such, network-only methods do not provide any information about the correspondence between the two clusterings. Further, RETWEET and FOLLOW return only binary labels, hence we do not use this baseline for comparing ideology scores.

**Content-only:** We use orthogonal NMF tri-factorization (ON-MTF), a co-clustering approach [13], and dual manifold co-clustering (DMCC) [23]. In these methods the bipartite content matrix  $C$  is used to co-cluster the rows (users) and columns (sources) of the matrix simultaneously using bi-orthogonality and graph-regularization constraints.

**Network and content:** We compute ideology scores of Twitter users estimated by Kulshrestha et al. [24] (KULSHRESTHA) and Lu et al. [29] (BIASWATCH).

**Proposed methods:** We use the proposed method IFD, and a variant of IFD without graph-regularization constraints (IFD-NGR). We initialized  $U$  and  $V$  randomly from a uniform distribution in  $[0, 1]$  and  $H_u$  and  $H_s$  as identity matrices of size  $k$ . Parameters  $\alpha$  and  $\beta$  are chosen using grid search. Additional details on various approaches tried for parameter initialization, parameter tuning,

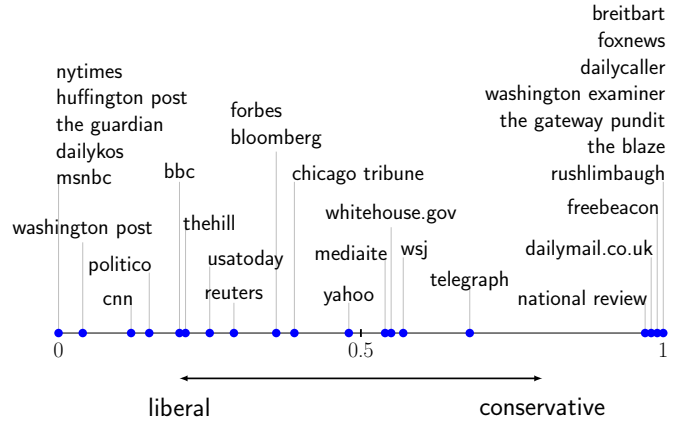


Figure 2: Popular media outlets and their ideology leaning scores computed by our method.

and stability of the algorithms with respect to the parameters are omitted due to lack of space, and will be provided in the full version of the paper.

### 5.4 Experimental setup

**Evaluation measures.** We perform two types of qualitative evaluation tasks: (i) quality of ideological cluster separation (into *liberal* or *conservative* clusters) and (ii) correlation between the computed ideology scores and ground-truth scores. In order to evaluate cluster separation, we measure PURITY, adjusted Rand index (ARI), adjusted mutual information (AMI), and normalized mutual information (NMI) between the clusters detected by the algorithm and the set of ground-truth communities derived by separating users at ideology score threshold at 0.5. In order to measure correlation between the computed ideology scores and the ground-truth scores, we use Pearson mutual correlation coefficient (CORR).

### 5.5 Results

**Ideology estimates for users and sources.** At a first look, the user ideology scores seem intuitive with the top liberal users being @barackobama (score: 0.0), @berniesanders (0.0), @thedemocrats (0.0) and top conservative users @tedcruz (score: 0.99), @seanhannity (0.99), and @davidlimbaugh (0.9). Figure 2 shows popular news-media outlets and their ideology leaning scores computed by our method. We observe that the position of the news sources is as expected: Liberal-leaning news outlets (e.g., *nytimes*, *washington post*, *the guardian*) are on the left, and conservative news outlets (e.g., *fox news*, *breitbart*, *rushlimbaugh*) on the right. This is also consistent with the survey-based results found by [30]. While it is easy to identify the extreme left and right, it is more difficult to identify the neutral users and sources (like *yahoo*, *mediaite*, *whitehouse.gov*, etc), which, in fact, is the most important subset of users and sources to tackle the information filter-bubble issue.

**Evaluation of clustering and ideology scores.** We compare the proposed methods with the baselines on (i) quality of ideological cluster separation (PURITY, ARI, AMI and NMI) and (ii) correlation (CORR) between computed ideology score and ground-truth

scores. The results of both evaluation tasks are listed in Table 1. Following are some noteworthy observations:

(i) IFD has the best performance among all the methods for both user as well as source clustering. We observe that combining network and content information gives consistently best results. For example, PURITY of clustering for combined methods is 20% higher than content-only methods for users and 27% higher for source clustering.

(ii) When comparing ideology scores, IFD performs better than state of the art baselines KULSHRESTHA and BIASWATCH. Note that BIASWATCH almost has no correlation with ground truth and performs poorly. This could be due to reasons mentioned in Section 2. Though KULSHRESTHA works slightly better on the clustering task, it doesn't do so well with the correlation. This indicates that our method is able to identify the fine grained ideology scores better, where as KULSHRESTHA can do better at separating users into binary clusters.

(iii) NMF-SYMM, a network-only method, performs quite well for user clustering, whereas the results are not satisfactory for source clustering. Perhaps this can be attributed to the noise in the input matrix  $C^T C$ , caused by, e.g., sparsity and topical diversity of user interests, which is mitigated in IFD because of the joint factorization. Similarly, our method performs much better than RETWEET and FOLLOW, which do not use both content and network information.

**Audience of news sources.** Figure 3 shows a kernel density estimate (KDE) of ideology scores computed using the proposed method (solid line) and ground truth of ideology scores (dashed line) for all 6391 users for a selection of 5 representative news channels. The findings from our experiments are strikingly similar to the results computed using extensive user surveys [30]. We can clearly observe that there is a non-trivial association between news sources and polarization of users.

## 6 DIFFUSING THE INFORMATION BUBBLE

In the recent years, a few papers [17, 19] have aimed to solve the problem of echo chambers and filter bubbles on social media by connecting users with others outside of their bubble. The problem with most of these approaches is that they typically suggest technical solutions that do not take into account user biases, such as, cognitive dissonance [14] and biased assimilation [28]. In addition, in many cases, users themselves are not aware of being present in a bubble, due to the non-transparent nature of algorithmic filtering.<sup>2</sup>

To handle both these issues, we propose that a better solution to take users out of their bubble is to give them *information* and *choice*. We do this in two steps: (i) make a user aware of their information bubble, and (ii) provide content recommendations that can help users diffuse their bubble. In this section, we demonstrate how to use the latent space representations learned in this paper to perform both of these steps.

<sup>2</sup>E.g., the filter bubble is claimed to be a reason why many people did not predict correctly the results of Brexit or US elections  
<http://flamingogroup.com/trump-brexit-and-why-we-didnt-see-them-coming>

**Table 1: Comparison of the proposed method with baseline methods on (i) quality of ideological cluster separation (PURITY, ARI, AMI and NMI) and (ii) correlation (CORR) between computed ideology score and ground truth. CORR for ideology is represented as  $CORR_i$  and for popularity as  $CORR_\rho$ . Since, we do not have ground truth for popularity of users we do not compute  $CORR_\rho$  for users. Best results for each measure are marked in bold.**

**(a) Evaluation for users using BARBERA ground truth (upper value in each row) and AVG\_CONTENT ground truth (lower value in each row)**

Type	Method	PURITY	ARI	AMI	NMI	$CORR_i$	$CORR_\rho$
Network Only	NMF-SYMM	<b>0.928</b>	<b>0.733</b>	<b>0.628</b>	<b>0.629</b>	<b>0.912</b>	-
		0.861	0.522	0.418	0.418	0.744	-
	RETWEET	0.844	0.476	0.385	0.385	-	-
0.839		0.461	0.395	0.399	-	-	
Content Only	FOLLOW	0.867	0.538	0.454	0.456	-	-
		0.845	0.476	0.382	0.382	-	-
	ONMTF	0.743	0.234	0.233	0.266	0.756	-
0.749		0.247	0.223	0.246	0.715	-	
Network + Content	DMCC	0.74	0.229	0.23	0.263	0.755	-
		0.746	0.241	0.218	0.242	0.715	-
Network + Content	IFD	<b>0.925</b>	<b>0.722</b>	<b>0.62</b>	<b>0.621</b>	<b>0.904</b>	-
		<b>0.863</b>	<b>0.528</b>	<b>0.441</b>	<b>0.442</b>	<b>0.772</b>	-
	<b>0.925</b>	<b>0.722</b>	<b>0.62</b>	<b>0.621</b>	<b>0.904</b>	-	
	<b>0.863</b>	<b>0.528</b>	<b>0.441</b>	<b>0.442</b>	<b>0.772</b>	-	
	KULSHRESTHA	<b>0.931</b>	<b>0.744</b>	<b>0.637</b>	<b>0.638</b>	0.875	-
		<b>0.869</b>	<b>0.547</b>	<b>0.448</b>	<b>0.449</b>	0.744	-
BIASWATCH	0.541	0.000	0.000	0.000	0.002	-	
	0.543	0.000	0.000	0.000	0.018	-	


**(b) Evaluation for sources using CONTENT ground truth**

Network Only	NMF-SYMM	0.597	0.031	0.135	0.171	0.752	0.597
	FOLLOW	0.819	0.405	0.318	0.32	-	-
Content Only	ONMTF	0.606	0.039	0.145	0.181	0.746	0.593
	DMCC	0.606	0.039	0.145	0.181	0.746	0.592
Network + Content	IFD	<b>0.826</b>	<b>0.424</b>	<b>0.346</b>	<b>0.348</b>	<b>0.827</b>	<b>0.929</b>
	IFD-NGR	<b>0.822</b>	<b>0.415</b>	<b>0.339</b>	<b>0.341</b>	<b>0.813</b>	<b>0.93</b>

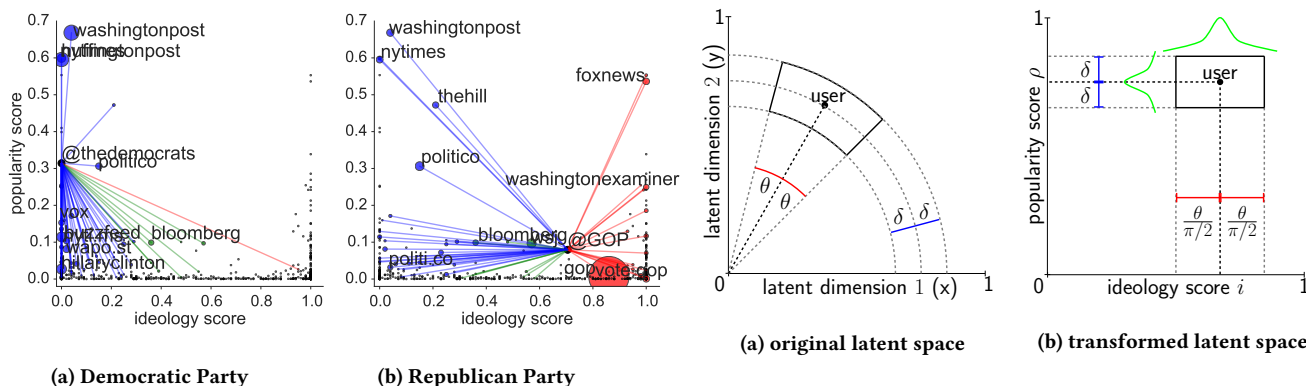
### 6.1 Visualizing the information bubble

Recent studies have shown that making users aware of their imbalance in media consumption can encourage them to make small yet significant improvements in increasing the diversity of their reading [19, 31]. Inspired by these studies, we suggest providing to the users a visual way to explore and explain their information filter bubble.

The way we approach this is by visualizing the user and content in the same space to allow users to explore their content consumption. To this end, we use the learned latent space to map the users and content in the same space. Using their estimated ideological positions (from Section 4), we project users as well as sources in a two dimensional ideology-popularity coordinate space with computed ideology score on  $x$ -axis and popularity score on  $y$ -axis. Since we do not have popularity score for users, we determine the position of a user on  $y$ -axis using the average popularity score of the content that the user engages with. Since all user scores are on the same scale and relative to each other, we would observe that users and sources with similar ideology lie close to each other. Finally, we connect users to the sources that they consume by drawing a link



**Figure 3: Polarization of the audience of news sources.** Values on the  $x$ -axis represent the ideology score of users and values on the  $y$ -axis represent the kernel density estimate of the number of users at each point.



**Figure 4: Ideological position of @thedemocrats and @gop (black dots) and their content engagement. Points in the grey are the sources that the user never interacted with.**

between them. The size of a source node is proportionate to the number of times a user has consumed content from the said source. In order to increase the ease of visual interpretation, we color the content according to the ideological leaning (blue: liberal, green: neutral and red: conservative). Content not consumed by the user is colored grey.

Figure 4 presents a prototype for two popular Twitter accounts from the two ends of the political spectrum: the Republican Party (@gop) and the Democratic party (@thedemocrats).<sup>3</sup> From this figure, one can visually observe their own ideological positioning as well as the ideology of the content that they engage with. For instance, @thedemocrats is heavily liberal in their ideology (ideology score 0.0). The content consumed by @thedemocrats is also heavily biased on the liberal side. As expected, a large fraction of the content they engage with is from the left (mainly liberal media like *nytimes.com* and *washingtonpost.com*), and negligible amount from the opposite point of view, whereas the opposite is true for @gop. It is interesting to observe that the Republican party account has a higher engagement with diverse view points than the Democrats.

## 6.2 Making ideologically diverse content recommendations

Garimella et al. [17] proposed an approach to diffuse a user’s filter bubble by connecting him to a user outside his bubble from the opposing viewpoint. Their approach is mainly based on identifying users from opposing sides and optimizing a global function.

<sup>3</sup>A link to an interactive web version of these plots generated for all other users will be added to in the full (non-anonymized) version of the paper.

**Figure 5: Logical diagram of user content recommendation by sampling from the Gaussian over “ideology” and “popularity” positioning.**

Here, we build on top of that idea and use our computed ideology to diffuse a user’s bubble by recommending him content from an opposing viewpoint, along with an *option* to choose how willing the user is to explore the other side. Recommending ideologically diverse content to a user can be controlled by the user using two parameters: ideology tolerance threshold  $\theta$  and popularity threshold  $\delta$ . Intuitively, a user is more likely to accept content within the region  $+\theta$  and  $-\theta$  on either side of the user’s ideological positioning, and  $+\delta$  and  $-\delta$  on either side of his popularity position. Figure 5a visualizes a hypothetical user in the original ideology latent space and the transformed ideology-popularity coordinate space (detailed in Section 4). Consider that we build two Gaussian distributions around the user *box* (see Figure 5b) with their means centered at user’s ideology and popularity score respectively, and variance as a function of the tolerance threshold given as input by the user. We can now sample content from these Gaussian distributions and use it for recommending content to the user. As desired, in such a sampling, the content close to the user’s own ideology and popularity score has a higher probability of being selected. As we move closer to the thresholds, the probability of an article being selected gradually decreases. This “*box*” gives the space of exploration for a user and depending on the user’s willingness to explore (based on parameters  $\theta, \delta$ ), they can see content outside their bubble.

## 7 CONCLUSIONS

We considered the problem of identifying ideological leaning of users and news sources (content) on Twitter. The paper tackles two main challenges: (i) learning the ideological latent factors of users and content in a joint model that explores simultaneously user-to-user and user-to-content relations; and (ii) embedding the



discovered factors in a common latent space so as to support visualization and exploration of the results. Our approach distinguishes itself from most existing work in the area in three major ways. First, our model aims to learn ideology on a continuous scale rather than a binary *liberal-conservative* opinion, which is a much simpler and well studied problem. Second, our model defines polarization as a multidimensional problem and allows for learning any number of dimensions (ideology, popularity, etc) rather than just one (ideology). Third, our model can identify the ideology of both the user and content simultaneously, thus making use of the interdependent structure between content and network. We also demonstrate how to use the learned latent space in an application of providing tools to visualize a users information bubble and for ideologically diverse content recommendation with the purpose of diffusing their information filter bubble.

**Future work.** One of the main focuses of this work was to create a strong technical foundation to understand the problem of online polarization. We believe that the methods presented in this paper provide several avenues for future work in multiple emerging interdisciplinary research areas, for instance, humanly interpretable and explanatory machine learning, transparent recommendations and a new era of social media platforms that encourage discussion and debates between users of diverse view points, thus helping to reduce the ideological segregation of users instead of reinforcing it.

## 8 REFERENCES

- [1] Eduardo Alemán, Ernesto Calvo, Mark P Jones, and Noah Kaplan. 2009. Comparing Cosponsorship and Roll-Call Ideal Points. *Legislative Studies Quarterly* 34, 1 (2009), 87–116.
- [2] Joseph Bafumi, Andrew Gelman, David K Park, and Noah Kaplan. 2005. Practical issues in implementing and understanding Bayesian ideal point estimation. *Political Analysis* 13, 2 (2005), 171–187.
- [3] Eytan Bakshy, Solomon Messing, and Lada A Adamic. 2015. Exposure to ideologically diverse news and opinion on Facebook. *Science* 348, 6239 (2015), 1130–1132.
- [4] Pablo Barberá. 2015. Birds of the Same Feather Tweet Together: Bayesian Ideal Point Estimation Using Twitter Data. *Political Analysis* 23, 1 (2015), 76–91.
- [5] Pablo Barberá, John T Jost, Jonathan Nagler, Joshua A Tucker, and Richard Bonneau. 2015. Tweeting from left to right: Is online political communication more than an echo chamber? *Psychological science* 26, 10 (2015), 1531–1542.
- [6] Robert Bond and Solomon Messing. 2015. Quantifying social media’s political space: Estimating ideology from publicly revealed preferences on Facebook. *American Political Science Review* 109, 1 (2015), 62–78.
- [7] Adam Bonica. 2013. Ideology and interests in the political marketplace. *American Journal of Political Science* 57 (2013).
- [8] Deng Cai, Xiaofei He, Xiaoyun Wu, and Jiawei Han. 2008. Non-negative matrix factorization on manifold. In *ICDM*. IEEE, 63–72.
- [9] Raviv Cohen and Derek Ruths. 2013. Classifying political orientation on Twitter: It’s not easy!. In *ICWSM*.
- [10] Michael D Conover, Bruno Gonçalves, Jacob Ratkiewicz, Alessandro Flammini, and Filippo Menczer. 2011. Predicting the political alignment of twitter users. In *SocialCom*. IEEE, 192–199.
- [11] Michael Dimock, Carroll Doherty, Jocelyn Kiley, and Russ Oates. 2014. Political polarization in the American public. *Pew Research Center* (2014).
- [12] Chris Ding, Xiaofeng He, and Horst D Simon. 2005. On the equivalence of nonnegative matrix factorization and spectral clustering. In *SDM*. SIAM, 606–610.
- [13] Chris Ding, Tao Li, Wei Peng, and Haesun Park. 2006. Orthogonal nonnegative matrix tri-factorizations for clustering. In *KDD*. ACM, 126–135.
- [14] Leon Festinger. 1962. *A theory of cognitive dissonance*. Vol. 2. Stanford university press.
- [15] Seth Flaxman, Sharad Goel, and Justin M Rao. 2016. Filter bubbles, echo chambers, and online news consumption. *Public Opinion Quarterly* 80, S1 (2016), 298–320.
- [16] Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. 2016. Quantifying controversy in social media. In *WSDM*. ACM, 33–42.
- [17] Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. 2017. Reducing controversy by connecting opposing views. In *WSDM*. ACM, 81–90.
- [18] Matthew Gentzkow and Jesse M Shapiro. 2011. Ideological segregation online and offline. *The Quarterly Journal of Economics* 126, 4 (2011), 1799–1839.
- [19] Eduardo Graells-Garrido, Mounia Lalmas, and Daniele Quercia. 2013. Data portraits: connecting people of opposing views. *arXiv preprint arXiv:1311.4658* (2013).
- [20] Eduardo Graells-Garrido, Mounia Lalmas, and Daniele Quercia. 2014. People of opposing views can share common interests. In *WWW Companion*. ACM, 281–282.
- [21] Tim Groeling. 2013. Media bias by the numbers: Challenges and opportunities in the empirical study of partisan news. *Annual Review of Political Science* 16 (2013).
- [22] Tim Groseclose and Jeffrey Milyo. 2005. A measure of media bias. *The Quarterly Journal of Economics* 120, 4 (2005), 1191–1237.
- [23] Quanquan Gu and Jie Zhou. 2009. Co-clustering on manifolds. In *KDD*. ACM, 359–368.
- [24] Juhi Kulshrestha, Motahhare Eslami, Johnnatan Messias, Muhammad Bilal Zafar, Saptarshi Ghosh, Krishna P Gummadi, and Karrie Karahalios. 2017. Quantifying Search Bias: Investigating Sources of Bias for Political Searches in Social Media. In *CSCW*. ACM, 417–432.
- [25] Tao Li and Chris Ding. 2006. The relationships among various nonnegative matrix factorization methods for clustering. In *ICDM*. IEEE, 362–371.
- [26] Q Vera Liao and Wai-Tat Fu. 2014. Can you hear me now?: mitigating the echo chamber effect by source position indicators. In *CSCW*. ACM, 184–196.
- [27] Q Vera Liao and Wai-Tat Fu. 2014. Expert voices in echo chambers: effects of source expertise indicators on exposure to diverse opinions. In *CHI*. ACM, 2745–2754.
- [28] Charles G Lord, Lee Ross, and Mark R Lepper. 1979. Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of personality and social psychology* 37, 11 (1979), 2098.
- [29] Haokai Lu, James Caverlee, and Wei Niu. 2015. Biaswatch: A lightweight system for discovering and tracking topic-sensitive opinion bias in social media. In *CIKM*. ACM, 213–222.
- [30] Amy Mitchell, Jeffrey Gottfried, Jocelyn Kiley, and Katerina Eva Matsa. 2014. Political polarization & media habits. *Pew Research Center* (2014).
- [31] Sean A Munson, Stephanie Y Lee, and Paul Resnick. 2013. Encouraging reading of diverse political viewpoints with a browser widget. In *ICWSM*.

- [32] Eli Pariser. 2011. *The filter bubble: What the Internet is hiding from you*. Penguin UK.
- [33] Marco Pennacchiotti and Ana-Maria Popescu. 2011. Democrats, republicans and starbucks aficionados: user classification in twitter. In *KDD*. ACM, 430–438.
- [34] Ferran Pla and Lluís F Hurtado. 2014. Political Tendency Identification in Twitter using Sentiment Analysis Techniques. In *COLING*. 183–192.
- [35] Keith T Poole and Howard L Rosenthal. 1997. *Ideology and congress*. Vol. 1. Transaction Publishers.
- [36] Markus Prior. 2007. *Post-broadcast democracy: How media choice increases inequality in political involvement and polarizes elections*. Cambridge University Press.
- [37] Cass R Sunstein. 2009. *Republic. com 2.0*. Princeton University Press.
- [38] Felix Ming Fai Wong, Chee Wei Tan, Soumya Sen, and Mung Chiang. 2016. Quantifying Political Leaning from Tweets, Retweets, and Retweeters. *TKDE* 28, 8 (2016), 2158–2172.
- [39] Jiho Yoo and Seungjin Choi. 2010. Orthogonal nonnegative matrix tri-factorization for co-clustering: Multiplicative updates on stiefel manifolds. *Information processing & management* 46, 5 (2010), 559–570.
- [40] Shenghuo Zhu, Kai Yu, Yun Chi, and Yihong Gong. 2007. Combining content and link for classification using matrix factorization. In *SIGIR*. ACM, 487–494.